

# THE INTELLIGENT OBSERVATORY

S. B. Potter<sup>†,1,2</sup>, S. Chandra<sup>1</sup>, N. Erasmus<sup>1</sup>, M. Hlakola<sup>1</sup>,  
R. Julie<sup>3</sup>, C. van Gend<sup>1</sup>, H. L. Worters<sup>1</sup>

<sup>1</sup>South African Astronomical Observatory, Cape Town, South Africa

<sup>2</sup>Department of Physics, University of Johannesburg, Auckland Park, South Africa

<sup>3</sup>South African Radio Astronomy Observatory, Cape Town, South Africa

## Abstract

The Intelligent Observatory (IO) is the vision of the South African Astronomical Observatory (SAAO) to enhance scientific support for the South African and global astronomical communities. By optimizing services, the IO initiative aims to create a seamless collaboration among diverse local and hosted astronomical facilities situated on the Sutherland Plateau. Achieving this vision necessitates strategic technological enhancements, such as upgrading telescopes for remote observations and automation, alongside a comprehensive redesign of the existing Sutherland operations model. The primary research driver is time-domain and transient science.

## INTRODUCTION

The South African Astronomical Observatory (SAAO) provides ground-based optical and IR observational facilities for astronomers across South Africa and internationally, conducts world-class astronomical research, and communicates the excitement of astronomy to the people of South Africa. The administrative headquarters, are based in Cape Town including the electronics and high precision mechanical and optical laboratories, whereas the telescopes are located on a barren plateau outside the Karoo town of Sutherland, approximately 350 kms from Cape Town (Fig. 1).

### *The Telescopes*

The SAAO is the premier facility for optical astronomy on the African continent, and owns and operates several optical and infrared telescopes which include the Southern African Large Telescope (SALT), Lesedi (1.0 m), the 1.0 m and 1.9 m telescopes. Other facilities, such as the Infra-Red Survey Facility are co-owned and operated with other international partners. The remaining ~20 facilities the SAAO hosts on behalf of other international institutes for which the SAAO is either permitted a percentage of time access or data access.

Each of these facilities have their own modes of operation primarily depending on the science driver and/or ownership. SALT issues 2 calls per year for applications from the partner institutes. The amount of time applied for and eventually awarded is tailored to meet the scientific goal of the application. The telescope uses a queue-based mechanism for executing the observing program, with dedicated SALT astronomers and operators performing the observations.

Applications to the SAAO 1.9 m, 1.0 m and the IRSF

\*† sbp@sao.ac.za

are made 3 times per year and the PI is expected to perform their own observations. Time is allocated in 1 week blocks as a result of the practicalities of travelling and accommodating in Sutherland. Applications are open to the entire national and international community and time is awarded primarily on scientific merit. The hosted facilities are all generally operated for specific science goals and consequently their mode of operation is mostly robotic and queue scheduled

## IO RATIONALE

The main science driver for the IO is time-domain and transient astronomy. The time over which astronomical objects can vary ranges from milliseconds to years depending of the characteristics of the objects. Synoptic monitoring of objects that are changing slowly are observed over time-scales of days and years, while the variability of close binary stars takes place on timescales of tens of minutes to hours, examples are cataclysmic variables and X-ray binaries. Then there are rapidly varying objects, such as pulsars and accretion instabilities in flows and discs that require high time resolution observations of sub-seconds. Many transients are found through e.g. all-sky surveys or high-energy events are discovered with X-ray, gamma-ray, and radio telescopes. More recently events are being discovered with gravitational wave and neutrino detectors. Such a range of objects and their corresponding phenomena result in a variety of emissions across the entire electromagnetic and multi-messenger spectrum.

It is through observing these timely events that we get the opportunity to improve our knowledge and understanding of these objects, their underlying physics and make new discoveries along the way. Observations will have to be tailored to the type of object and the nature of the optical/IR emissions. With ~25 different facilities operating on the Sutherland plateau with even more instruments ranging from wide-field imagers and spectrographs to specialised modes of high cadence observations or polarimetry, the SAAO is uniquely equipped to explore all transient and time domain phenomena - provided that all the facilities can be intelligently managed and operated with this vision in mind.

## THE IO VISION

The vision for the SAAO is to have all the SAAO facilities (telescopes and instruments) integrated into the IO. The operational model for the IO will add additional functionality to the operations of the observatory. For exam-

# ITER CONTROLS APPROACHING ONE MILLION INTEGRATED EPICS PROCESS VARIABLES

Anders Wallander\*, ITER Organization, St-Paul-lez-Durance, France  
Bertrand Bauvir\*, ITER Organization, St-Paul-lez-Durance, France

## Abstract

The ITER Tokamak is currently being assembled in southern France. In parallel, the supporting systems have completed installation and are under commissioning or operation. Over the last couple of years, the electrical distribution, building services, liquid & gas, cooling water, reactive power compensation and cryoplant have been integrated, adding close to one million process variables. Those systems are operated, or are under commissioning, from a temporary main control room or local control rooms close to the equipment using an integrated infrastructure. The ITER control system is therefore in production.

ITER procurement is 90% in-kind, so a major challenge has been the integration of the various systems provided by suppliers from the ITER members. Standardization, the CODAC Core System software distribution, training and coaching have all played a positive role. Nevertheless, integration has been more difficult than foreseen and the central team has been forced to rework much of the delivered software.

In this paper we report on the current status of the ITER integrated control system with emphasize on lessons learned from integration of in-kind contributions.

## INTRODUCTION

The goal of the ITER project is to demonstrate the technical feasibility of commercial fusion as a future energy source. It uses the technique of magnetic confinement to create and maintain a super-hot hydrogen plasma in a doughnut-shaped 1400 m<sup>3</sup> vacuum chamber, the Tokamak concept. The strong magnetic fields are created by superconducting magnets. There are extreme temperature differences between the magnets at 4 K and the plasma at hundreds of millions of degrees.

The Tokamak is currently under assembly at the ITER site in southern France. Many of the supporting systems, e.g. electrical supplies, building services, cooling water and cryoplant, are constructed and in commissioning or operation.

The ITER project is a collaboration between seven members (China, Europe, India, Japan, Korea, Russia and the USA) representing 35 countries. The members comprise half the world's population and 85 percent of the global gross domestic product. The project is based on a substantial proportion of in-kind procurement meaning the members provide components and systems, not money, to the project. This in-kind procurement poses a major challenge for the central ITER Organization responsible for the specification, integration and operation of the machine. This

challenge applies to the control system as well, as the software and hardware for each component or system delivered by a member is an in-kind contribution.

The main mitigation to address this challenge was established more than ten years ago with the publication of the standards (Plant Control Design Handbook [1]), a software distribution (CODAC Core System [2]) as well as active outreach and training as reported in [3]. Over the years these standards have evolved to address obsolescence and feedback of experience through regular updates. The software distribution has a release cycle period of less than one year. One fundamental decision, made ten years ago, was to base the control system on EPICS. This choice was driven by the prevalence of EPICS users in all ITER member states as well as its proven robustness record. In retrospect we can confirm this choice as sound, the EPICS core has proven its scalability and robustness.

Over the last four to five years, in-kind local control systems have been delivered to the site and been integrated with the central control system. Some of these have been successfully commissioned and put in operation, while others are still under commissioning. Commissioning is done incrementally as systems depend on each other. For example, the cooling water cannot run without an electrical supply, liquid & gas and building services; the cryoplant cannot run without cooling water and the superconducting magnets cannot be powered without the coil power supplies and cryoplant. In this paper we report on the experience of integrating in-kind local control systems, the global status of the integrated control system and plans for the near future.

## ARCHITECTURE AND INFRASTRUCTURE

The ITER integrated control system is hierarchical with 21 subsystems and 170 local control systems. The latter are delivered in-kind by the ITER members through 100 procurement arrangements. The architecture comprises three segregated vertical slices. In order of increasing integrity/criticality these are: (1) conventional control and operation, (2) machine protection and (3) occupational and nuclear safety. The local control systems interface sensors and actuators in the field and provide local controllers and data acquisition. Virtualized servers, at higher layer, interface with the local control system via networks. These servers provide orchestration, monitoring, configuration and data handling functions. Finally, the top layer is responsible for human machine interfaces (HMI) allowing operators to control and monitor the entire machine.

\* anders.wallander@iter.org, bertrand.bauvir@iter.org

# LCLS-II ACCELERATOR CONTROL SYSTEM STATUS\*

D. Rogind, S. Kwon

SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

The Linac Coherent Light Source complex at the SLAC National Accelerator Laboratory has been upgraded to add a new superconducting (SC) accelerator with beam rates up to 1 MHz. Though the majority of the more than twenty accelerator control systems are based on LCLS designs, to accommodate the increase in repetition rate from 120 Hz to 1 MHz, many of the diagnostics and global control systems were upgraded to high performance platforms with standalone CPUs running linuxRT to host the EPICS based controls. With installation and checkouts for control systems completing in 2022, the phased approach to integration and commissioning recently completed with demonstration of the threshold key performance parameters and first light occurring in the Summer of 2023. This paper provides an overview of the LCLS-II accelerator control system architecture, upgrades, the multi-year installation, checkout, integration, commissioning, and lessons learned.

## LCLS-II OVERVIEW

After close to a decade of work, the newly upgraded Linac Coherent Light Source (LCLS) X-ray free-electron laser (XFEL) at the Department of Energy's SLAC National Accelerator Laboratory (cartoon shown in Fig. 1) successfully produced its first X-rays this past summer of 2023. The brighter, more rapid bursts of X-rays will allow scientists to tackle challenges such as understanding how to adapt natural solutions for harvesting solar energy for a new generation of clean fuels, inventing sustainable manufacturing methods for industry, and designing a new generation of drugs based on the ability to create molecular movies of how our bodies respond to disease. LCLS-II delivers X-ray laser beams that are 10,000 times brighter than its predecessor [1]. Refer to Fig. 2, which graphs the




Figure 1: LCLS-II layout.

calculated spectral brightness from LCLS-II soft X-ray undulator (SXU) and hard X-ray undulator (HXU) at 4 GeV and high-repetition-rate operation.




Figure 2: Calculated spectral brightness.

Staff from four national laboratories – Argonne, Berkeley Lab, Fermilab, and Jefferson Lab – along with Cornell University, worked together to build the facility's next-generation components. LCLS-II increases beam rate from LCLS-I's 120 pulses per second to 1 million pulses per second. The newly constructed dual cryogenics cooling plant supplies helium to the 37 cryomodules at a temperature of two Kelvin. New variable gap Hard X-Ray undulators (HXU) and Soft X-Ray (SXU) undulators can each receive beam from either the normal conducting (NC) LCLS-I or the SC LCLS-II, in parallel. The series of undulator magnets force electrons to give off energy in the form of X-rays. These X-rays are then delivered to a suite of experimental instruments for users to conduct experiments.

## KEY FEATURES OF LCLS-II

The new 4 GeV SC linear accelerator (linac), shown in Figure 3: Linac Coherent Light Source – II, occupies the first kilometer of the 3 km SLAC linac tunnel. The existing legacy copper facility was removed from the first kilometer, thus eliminating need for excavation. The normal conducting (NC) copper LCLS-I linac, which began operation in 2009, occupies the last third of the linac tunnel. The middle kilometer is occupied by the Facility for Advanced Accelerator Experimental Tests (FACET)-II.

## ACCELERATOR CONTROLS

LCLS-II Accelerator Control System Scope is comprised of the following systems:

- **Global Controls:** Timing, Network, Computing Infrastructure, Machine Protection (MPS), Racks & Cables, ATCA Common Platform
- **Diagnostics:** Beam Position Monitor (BPM), Beam Current Monitor (BCM), Bunch Length Monitor (BLM), General Motion (Wire Scanner, Collimator, Bunch Compressor), Undulator Control, Profile Monitor

\*Work supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# EIC CONTROLS SYSTEM ARCHITECTURE STATUS AND PLANS\*

J. P. Jamilkowski†, S. Clark, M. Costanzo, T. D’Ottavio, M. Harvey, K. Mernick, S. Nemesure, F. Severino, K. Shroff, K. Kulmatycki, C. Montag, V. H. Ranjbar, K. S. Smith  
Brookhaven National Laboratory, Upton, USA  
L. B. Dalesio, Osprey DCS LLC, Ocean City, USA

## Abstract

Preparations are underway to build the Electron Ion Collider (EIC) once the Relativistic Heavy Ion Collider (RHIC) beam operations end in 2025, providing an enhanced probe into the building blocks of nuclear physics for decades into the future. With commissioning of the new facility in mind, Accelerator Controls will require modernization in order to keep up with recent improvements in the field as well as to match the fundamental requirements of the accelerators that will be constructed. We will describe the status of the Controls System architecture that has been developed and prototyped for EIC, as well as plans for future work. Major influences on the requirements will be discussed, including EIC Common Platform applications and our expectation that we'll need to support a hybrid environment covering both the proprietary RHIC Accelerator Device Object (ADO) environment as well as EPICS.

## NEARING THE END OF THE RHIC ERA

Beam operations of the Relativistic Heavy Ion Collider (RHIC) are expected to cease after the FY25 run, having successfully supported several key detectors that have been used for nuclear physics experiments over 23 years. Colliding beams in use at the facility have been comprised of many hadron species including polarized protons and gold with beam energies between 3.85 GeV [1] and 255 GeV [2] to date.

The RHIC Controls System includes Front End Computers (FECs) that utilize either VME-standard components or a proprietary XILINX FPGA-based platform. The latter was primarily developed and utilized to support the LLRF system at RHIC and elsewhere in the hadron injector chain. VME boards in use at RHIC include a significant component of proprietary boards for multiple systems as well as commercial equipment. Additionally, the software infrastructure of the Controls System is based on a proprietary standard called Accelerator Device Objects (ADOs) [3]. The current version of the Controls System was developed in preparation for the commissioning of RHIC, though it has continued to develop and expand in scope and complexity as the needs of RHIC and other independent machines have continued to evolve.

## WHAT IS THE EIC?

The US Department of Energy (DOE) approved the siting of the Electron Ion Collider (EIC) [4] at Brookhaven National Laboratory (BNL) in FY21, and a partnership has

\* Work supported by Contract Number DE-AC02-98CH10886 with the auspices of the US Department of Energy  
† jppjamilk@bnl.gov

General

Status Reports

been established with Thomas Jefferson National Accelerator Facility (TJNAF) to construct the new machine after FY25 pending project approval. The RHIC tunnel will be reused for the new project, along with a subset of the RHIC magnets that comprise the existing Blue and Yellow Rings for a new Hadron Storage Ring (HSR) (Fig. 1). A new turnkey LINAC will be added to accelerate polarized electrons to 400 MeV, and the resulting beam is accelerated in a new Rapid Cycling Synchrotron (RCS) Ring. Bunches are then transferred from the RCS to a new Electron Storage Ring (ESR), which supports collisions with beam in the HSR. Center of mass collision energies between the electron and hadron beams will range between 20 and 140 GeV [5] at one or two Interaction Regions, the first of which will be serviced by the electron-Proton Ion Collider (ePIC) Detector. Beam from an additional Strong Hadron Cooling (SHC) electron accelerator will be used to reduce the size of the hadron bunches in both HSR injection and store conditions. Completion of the commissioning phase of the project is currently expected to occur in FY34.

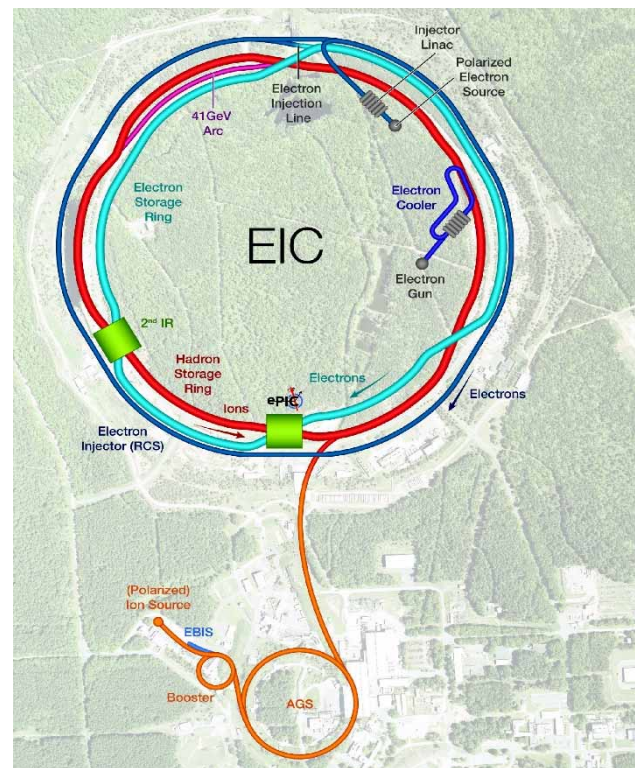


Figure 1: Cartoon Layout of the EIC.

The components required to achieve the functional and performance goals of the EIC Project will be typical of a large, modern particle accelerator. These include normal-conducting and super-conducting magnets, normal-

M01BC004

# DRIVING BEHAVIOURAL CHANGE OF SOFTWARE DEVELOPERS IN A GLOBAL ORGANISATION ASSISTED BY A PARANOID ANDROID

Ugur Yilmaz\*, SKA Observatory, Jodrell Bank, UK  
Adriaan De Beer†, SARAO, Cape Town, South Africa  
Marvin G.P.T. Android, SKA Observatory, the Universe‡

## Abstract

Ensuring code quality standards at the Square Kilometre Array (SKA) Observatory is of utmost importance, as the project spans multiple nations and encompasses a wide range of software products delivered by developers from around the world. To improve code quality and meet certain open-source software prerequisites for a wider collaboration, the SKA Observatory employs the use of a chatbot that provides witty, direct and qualified comments with detailed documentation that guide developers in improving their coding practices. The bot is modelled after a famous character albeit a depressed one, creating a relatable personality for developers. This has resulted in an increase in code quality and faster turnaround times. The bot has not only helped developers adhere to code standards but also fostered a culture of continuous improvement with an engaging and enjoyable process. Here we present the success story of the bot and how a chatbot can drive behavioural change within a global organisation, while helping DevOps teams to improve developer performance and agility through an innovative and engaging approach to code reviews.

GPT in Marvin G.P.T Android stands for Grumbly Pathos-filled Tinhead, at least, that's what we heard.

## INTRODUCTION

The Square Kilometre Array (SKA) Observatory is an international effort to build two radio telescopes in South Africa and Australia forming one Observatory monitored and controlled from global headquarters in the UK. When preparing releases for end-users, significant software projects often encounter the challenge of harmonising various components and deploying them in the production environment. This issue arises when multiple project segments have been developed independently for a period, leading to integration complexities and higher-than-anticipated developer resource allocation. In the dynamic realm of software development, merging independently developed components can lead to the notorious challenge known as "merge hell." This issue persists even with established practices and have been discussed extensively within the very large technology companies [1]. Within the SKA Observatory, involving over a hundred developers and repositories with varied technologies, such conflicts can impede timely releases. Thus, the imperative for standardised CI/CD practices is clear to start with a baseline quality. The SKA Observatory, employing

the SAFe Agile framework, leverages a dedicated Systems Team to bolster Continuous Integration, Continuous Deployment, test automation, and quality.

Enter the droll presence of Marvin the Paranoid Android, reluctantly drawn into yet another cosmic odyssey. Marvin's distinctive pessimism serves as a sobering reminder that even in the grand tapestry of the universe, neglecting the intricacies of code is a travesty that simply cannot be tolerated. After all, a neglected line of code, much like existential despair, is a matter of cosmic importance.<sup>1</sup>

In the subsequent sections of this paper, a brief overview of CI/CD practices at the SKA Observatory will be given focusing on how the SKA Observatory targets software quality by following best practices and with the help of automation materialised with the persona of Marvin the Paranoid Android. Then how Marvin helped foster a culture of collaboration amongst developers will be shared.

## IMPLEMENTATION

### *Continuous Integration and Merge Requests*

The advent of DevOps marks a pivotal cultural shift, transcending traditional silos between development and operations teams. In the SKA Observatory's context, this convergence is not merely a convergence of roles, but a convergence of responsibilities and ownership. Development teams actively engage in automating deployment operations, while operations teams gain insights into the applications they support. The shared accountability empowers development teams to deploy changes to production with confidence, underpinned by robust testing platforms and infrastructure management. This is achieved by following Continuous Integration and Deployment processes.

Continuous Integration (CI), a practice as integral as the flux capacitors of a hyperdrive, demands the seamless integration of code into a shared repository. It thrives on the rhythmic cadence of integration, allowing developers to align their contributions multiple times a day. It's a relentless pursuit of code integrity. Martin Fowler's foundational best practices resonate strongly, advocating for a unified source repository, automated builds, comprehensive testing, and a steadfast commitment to a stable main branch [2-4]. Complementing CI, Continuous Deployment (CD) extends the paradigm, ushering in the era of automated software releases. The codebase must consistently maintain a releasable state (called Release on Demand), subject to rigorous testing. Far from compromising stability, frequent deployments, when

\* ugur.yilmaz@skao.int

† adebeer@sarao.ac.za

‡ don't panic

<sup>1</sup> This section is written by Marvin the Paranoid Android himself

# CONCEPT AND DESIGN OF AN EXTENSIBLE MIDDLE-LAYER APPLICATION FRAMEWORK FOR ACCELERATOR OPERATIONS AND DEVELOPMENT\*

M. Schütte<sup>1†</sup>, A. Grünhagen<sup>2</sup>, J. Georg, H. Schlarb  
Deutsches Elektronen-Synchrotron DESY, Germany

<sup>1</sup>also at Hamburg University of Technology, Hamburg, Germany

<sup>2</sup>also at HAW Hamburg, Hamburg, Germany

## Abstract

Data collection and analysis are becoming increasingly vital not only for the experiments conducted with particle accelerators but also for their operation, maintenance, and development. Due to lack of feasible alternatives, experts regularly resort to writing task-specific scripts to perform actions such as (event triggered or temporary) data collection, system failure detection and recovery, and even simple high-level feedbacks. Often, these scripts are not shared and are deemed to have little reuse value, giving them a short lifetime and causing redundant work. We report on a modular Python framework for constructing middle-layer applications from a library of modules (parameterized functionality blocks) by writing a simple configuration file in a human-oriented format. This encourages the creation of maintainable and reusable modules while offering an increasingly powerful and flexible platform that has few requirements to the user. A core engine instantiates the modules according to the configuration file, collects the required data from the control system and distributes it to the individual module instances for processing. Additionally, a publisher-subscriber messaging system is provided for inter-module communication. We discuss architecture & design choices, current state and future goals of the framework as well as real use-case examples from the European XFEL.

## INTRODUCTION

Control systems for particle accelerators and other large experimental physics facilities provide a common infrastructure for a vast and heterogeneous collection of subsystems which expose a considerable amount of diagnostic, monitoring and configuration data. For example, over nine million data channels are exposed in the Distributed Object-Oriented Control System (DOOCS) control system at European XFEL [1]. This data is becoming increasingly vital for ensuring reliable operation, cutting-edge system performance, and failure detection and prevention. This is shown by many recent projects working on data-driven methods for accelerator setup [2, 3], operation [4], anomaly detection [5] and predictive maintenance [6]. Such project rely on a solid

data foundation which is as heterogeneous and complete as possible, as explained in [7], where a long term archive for the full state of an entire sub-system of the European XFEL is set up. Logging such vast amounts of data continuously over long time periods is however technically and financially challenging, and few subsystems today have such a data acquisition (DAQ) system readily available.

## Problem Description

To circumvent this limitation, we frequently observe operators and researchers utilizing control system library bindings in high-level scripting languages such as MATLAB or Python for loop-based polling of interesting data channels. For shot-based facilities, this is usually extended with data synchronization logic. Past the collection stage, the obtained data is (pre)processed with reoccurring operations, such as bit field extraction, sorting, outlier removal, averaging, Fourier transform computation etc. and eventually plotted. In a final stage, this offline analysis may result in the development of an online algorithm to automate routine tasks, generate reports, recover from system failure states or implement basic feedbacks e.g. for performance optimization.

This ad hoc scripting, in contrast to dedicated control system application development, has the advantages of rapid prototyping workflows and can be done directly by the scientists and operators without having to rely on scarce software developer resources. However, the do-it-yourself approach has severe downsides as well. Information about such development efforts is usually just spread by word-of-mouth, not everyone profits immediately. Scripts are seldom committed to a central code repository, making them hard to find and maintain. Multiple versions circulate and when in doubt, new users often rather start from scratch than understand and fix an older code. This is also because ad hoc scripts rarely follow good coding practices and speed is favoured over readability. This is fatal not only because the same task is solved redundantly over and over again, but also each time bugs may be introduced that cost additional time to fix or, even worse, stay unnoticed and produce incorrect results [8].

## Solution Proposal / Requirements

In order to keep and even extend on the rapid prototyping advantages of script based middle layer software, a guard rail system is required, providing immediate benefits to the operators and scientists using it and providing incentives

\* The authors acknowledge support from DESY Hamburg, Germany, a member of the Helmholtz Association HGF. © All figures and pictures by the authors under a CC BY 4.0 license.

† maximilian.schuette@desy.de

# STRATEGY AND TOOLS TO TEST SOFTWARE IN THE SKA PROJECT: THE CSP.LMC CASE

G. Marotta\*, C. Baffa, E. Giani, INAF-OA Arcetri, Firenze, Italy  
h I;^ < ' ~f/ < Gd; i < • | ð, Cosylab Switzerland, Brugg, Switzerland  
G. Brajnik, University of Udine and IDS, Udine, Italy

## Abstract

The Square Kilometre Array (SKA) Telescope will be one of the largest and most complex scientific instruments ever built. The development of a reliable software for monitoring and controlling its operations is critical to the success of the entire SKA project. The Local Monitoring and Control of the Central Signal Processor (CSP.LMC) is a software component responsible for controlling a key subsystem of the telescope, i.e. the Central Signal Processor (CSP). The software is implemented as a “device” within the TANGO framework, written in Python code.

In this paper we describe a testing strategy that addresses a series of problems typical of such a large and complex instrument. It is a multi-level strategy, based on a combination of automated tests (unit/component/integration), in the context of CI/CD practices. Software is also tested against possible errors and anomalous conditions that can occur while the CSP.LMC is interacting with external subsystems, which can be simulated.

The paper also discusses needs and solutions based on data mining test results. This allows us to obtain statistics of unexpected failures and to investigate their causes. Furthermore, a database containing test results over several weeks supports discovery of interesting and unexpected patterns of behaviors of the tests based on correlations about different test-related events and data. This helps us to develop a deeper understanding of the code’s functioning and to find suitable solutions to minimize unexpected behaviors. In addition it can be used also to support reliability testing.

## THE SKA PROJECT AND THE CSP.LMC

The Square Kilometre Array (SKA) telescope is an international effort to build the world’s largest radio telescope. Construction is underway at two primary sites in South Africa and Australia. These sites will house unprecedentedly large interferometers designed to observe the sky across two fundamental questions about the universe, including its evolution, the nature of gravity, and the search for extraterrestrial life. Consisting of thousands of antennas spread over long distances, the SKA will have a combined collecting area of approximately one square kilometer. Its unprecedented sensitivity and resolution will provide insights into cosmic phenomena, offering a deeper understanding of the early universe, black holes, and pulsars.

The Central Signal Processor (CSP) is a critical component of the Square Kilometre Array (SKA) telescope. It’s responsible for the real-time processing of the vast amount of data collected by the SKA antennas, in order to make it available for scientific data processing. Given the immense scale of SKA, the CSP is expected to handle unprecedented throughput of data. The estimated amount of data will be around 7.3 Tb/s for Low and 8.8 Tb/s for Mid [1].

CSP comprises three primary instruments (as shown in Fig. 1), each of which is a complex signal processing subsystem. These are:

- the Correlator and BeamFormer (CBF), that combines the signals from various antennas to create a unified and focused view of the sky;
- the Pulsar Search (PSS), that identifies potential candidates for pulsar discovery;
- the Pulsar Timing (PST), that measures the frequency of the radiation emitted from pulsar candidates.

On top of these subsystems, the Local Monitoring and Control (CSP.LMC) ensures the seamless operation of the CSP by providing real-time health checks, performance metrics, and adaptive controls [2]. In other words, CSP.LMC is the software component that represents the interface of the entire CSP instrument. It interacts with the Telescope Monitoring and Control (TMC) system, serving as the bridge between the TMC and the individual components of the CSP. It communicates to the TMC all the required information to monitor the CSP’s subsystems and provides the interface to send all the commands required to perform an observation. Even though the concept of CSP is the same for Mid and Low telescopes, some differences can occur in the data reduction hardware for the two telescopes. Therefore, two instances of CSP.LMC are developed to address the differences between the two telescopes.




Figure 1: Simplified schema for CSP.

\* gianluca.marotta@inaf.it

# APPLYING STANDARDISED SOFTWARE ARCHITECTURAL CONCEPTS TO DESIGN ROBUST AND ADAPTABLE PLC SOLUTIONS \*

S. T. Huynh<sup>†</sup>, B. Baranasic, M. Bueno, T. Freyermuth, P. Gessler, N. Jardón Bueno, N. Mashayekh, J. Tolkiehn, L. Zanellatto, European X-Ray Free-Electron Laser, Schenefeld, Germany

## Abstract

Between evolving requirements, additional feature requests and urgent maintenance tasks, the Programmable Logic Controllers (PLC) at the European X-Ray Free Electron Laser Facility (EuXFEL) have become subjected to an array of demands. As the maintainability effort towards the existing systems peak, the requirement for a sustainable solution become an ever pressing concern. Ultimately, in order to provide a PLC code base which can easily be supported and adapted to, a reworking was required from the ground up in the form of a new suite of libraries and tools. Through this, it was possible to bring standardised software principles into PLC design and development, conjunctively offering an interface into the existing code base for ongoing support of legacy code.

The set of libraries are developed by incorporating software engineering principles and design patterns in test driven development within a layered architecture. In defining clear interfaces across all the architectural layers - from hardware, to the software representation of hardware, and clusters of software devices, the complexity of PLC development decreases down into modular blocks of unit tested code. Regular tasks such as the addition of features, modifications or process control can easily be performed due to the adaptability, flexibility and modularity of the core PLC code base.

## INTRODUCTION

Over the decades, Programmable Logic Controllers (PLC) have been programmed and the world in which it exists, has also evolved into the standard that is known of today as the IEC 61131-3 [1] standard. With the increase in memory and the greater integration of Structured Text (ST) [2] PLCs have started to shift from an electrical hardwired concept to all that is offered within a software programming language. This cross discipline has opened up a powerful feature set which is often underutilised, and provides PLCs with the opportunity to integrate core software principles into their code base.

In the process of redeveloping the PLC code at the European X-Ray Free Electron Laser Facility (EuXFEL), the PLC developers are committed to developing a code base which is designed with thought and care, bringing into the design many of the software practices which are often embodied in large software projects, whilst taking into consideration the needs and functions of the PLC as the facility expands and

refines its needs. Taken into consideration is the continual support and maintenance of the current legacy code base, and its future integration.

## TECHNICAL DEBT

When juggling between time constraints and numerous feature requests, it would not be uncommon for code to be added to an existing code base in a haphazard manner, without thorough testing. Whilst this approach can seem to work, the consequences will inevitably catch up with the developers.

This was precisely the situation that the PLC developers at EuXFEL found themselves in: more time was being dedicated towards maintenance and resolving existing issues within the existing code base, than developing additional functionality. Within a scientific environment where experiments are in a constant state of flux, the demands on the PLC to be agile and adaptable, but also reliable was also further stressed.

It became customary to develop new patches to get around limitations, and the task of maintaining the array of existing tools became unmanageable. This is especially noted where core libraries and functions that were heavily relied upon within various programming languages, became deprecated or superseded.

Due to the amassed technical debt which accumulated over several years, the tipping point for redevelopment was reached. As such, a complete redevelopment was required for the PLC code base.

## KEY REQUIREMENTS

As with any major software project, the first step was the development of a Software Requirements Specification (SRS). As this paper will focus on the non-functional design aspects, the key outcomes of the SRS which ultimately defined the final PLC design, are detailed below.

### *Architectural Decision Records*

One of the challenges in maintaining legacy code is understanding what historical decisions were made and their reasoning. Unfortunately, this is often neither obvious nor something that can be garnered with going through the existing code. To be able to refactor, extend or at times, remove redundant code can be challenging without this information.

Knowing this will allow one to refine or edit existing work with an awareness of the original intention. Without this, code modification can result in unintended consequences, especially in a code base without a clean architecture. This

<sup>†</sup> sylvia.huynh@xfel.eu



# ENABLING TRANSFORMATIONAL SCIENCE THROUGH GLOBAL COLLABORATION AND INNOVATION USING THE SCALED AGILE FRAMEWORK

L. R. S. Brederode<sup>†</sup>, S. Ujjainkar, SKA Observatory, Macclesfield, United Kingdom  
F. Graser, Vivo SA, Somerset West, South Africa  
J. Coles, University of Cambridge, Cambridge, United Kingdom  
J. A. Kolatkar, PSL, Pune, India  
S. Valame, Sanikaizen Solutions, Pune, India

## Abstract

The Square Kilometre Array Observatory (SKAO) is one observatory, with two telescopes on three continents. It will be the world's largest radio telescope once constructed and will be able to observe the sky with unprecedented sensitivity and resolution. The SKAO software and computing systems will largely be responsible for orchestrating the observatory and associated telescopes, and processing the science data, before data products are distributed to regional science centres. The Scaled Agile Framework is being leveraged to coordinate forty lean agile development teams that are distributed throughout the world. In this paper, we report on our experience in using the Scaled Agile Framework, the successes we have enjoyed, as well as the impediments and challenges that have stood in our way.

## SKA OBSERVATORY

SKAO is an intergovernmental organisation (IGO) with 16 countries engaged in a partnership to design, build and operate the next-generation radio astronomy observatory. SKAO consists of a global headquarters located in the UK, a mid-frequency radio telescope (MID) and a low-frequency radio telescope (LOW). The MID telescope [1] is being constructed in the Karoo region of South Africa, and will consist of 197 dish antennas, while the LOW telescope [2] is being constructed in Western Australia and will consist of 131,072 log periodic antennas.

Both telescopes will use interferometry and a technique called aperture synthesis to combine signals from each antenna in a Central Signal Processor, located in Cape Town and Perth respectively. A Science Data Processor will leverage a supercomputer in each of these locations to ingest high-bandwidth data from the Central Signal Processor to produce data products, calibration solutions, and science alerts. Data products are delivered to SKA regional data centres scattered around the world for scientific analysis. Each telescope is orchestrated by a Telescope Monitoring and Control subsystem that provides configuration, observation execution, alarm handling, and monitoring services. An Observatory Science Operations subsystem provides proposal management and execution functions.

<sup>†</sup> ray.brederode@skao.int

As such, the telescopes are in many respects “software telescopes” requiring a range of developer skills, including:

- Platform Developers.
- Database Developers.
- Monitor and Control Developers.
- User Experience Developers.
- Radio Astronomy Data Scientists.
- High Performance Analysis Algorithm Developers.
- High Performance Computing Engineers.

Given the magnitude of the challenge, the range of skill requirements, and diverse membership of the observatory, it was anticipated that hundreds of developers spanning multiple geographic locations would need to collaborate to develop the software aspects of the observatory.

## FRAMEWORK DECISION

Waterfall approaches in similar projects had a bad track record, and so a large-scale lean-agile methodology and framework was favoured.

Key software and computing stakeholders had already bought into the basic agile principles [3]:

- Continuous delivery of valuable systems.
- Working systems as the primary measure of progress.
- Built-in quality and attention to technical excellence.
- Continuous improvements and plan-do-check-adjust cycles.
- Synchronised cadence with increments and iterations.
- Engaging with key stakeholders all the way.
- Leveraging motivated individuals that are enabled with decentralised decision making and autonomy.

SKAO considered five lean-agile frameworks: Disciplined Agile Delivery (DAD), Dynamic Systems Development Method (DSDM), Large Scale Scrum (LESS), Modular Framework for Scaling Scrum and the Scaled Agile Framework [4].

SKAO chose to implement the Scaled Agile Framework because of its large community of practitioners, the quality of its documentation and training material, and its support for developing cyber-physical systems.

# EMBEDDED CONTROLLER SOFTWARE DEVELOPMENT BEST PRACTICES AT THE NATIONAL IGNITION FACILITY

V. Gopalan, P. K. Singh, V. J. Hernandez, J. Fisher, C. M. Estes, P. Kale, A. Pao, A. Barnes  
Lawrence Livermore National Laboratory, Livermore, CA, USA

## Abstract

Software development practices such as continuous integration and continuous delivery (CI/CD) are widely adopted by the National Ignition Facility (NIF) which helps to automate the software development, build, test, and deployment processes. However, using CI/CD in an embedded controller project poses several challenges due to the limited computing resources such as processing power, memory capacity and storage availability in such systems. This paper will present how CI/CD best practices were tailored and used to develop and deploy software for one of the NIF Master Oscillator Room (MOR) embedded controllers, which is based on custom designed hardware consisting of a microcontroller and a variety of laser sensors and drivers. The approach included the use of automated testing frameworks, customized build scripts, simulation environments, and an optimized build and deployment pipeline, leading to quicker release cycles, improved quality assurance and quicker defect correction. The paper will also detail the challenges faced during the development and deployment phases and the strategies used to overcome them. The experience gained with this methodology on a pilot project demonstrated that using CI/CD in embedded controller projects can be challenging, yet feasible with the right tools and strategies, and has the potential to be scaled and applied to the vast number of embedded controllers in the NIF control system.

## INTRODUCTION

Software engineering best practices such as continuous integration and continuous delivery (CI/CD) [1] have transformed the development, testing and deployment of software applications in many domains. Once primarily used by web and cloud-based applications, these practices are now being increasingly extended to other areas, such as in the context of resource constrained embedded systems. Such systems typically have limited computing capabilities, real-time requirements, and may even lack an operating system. These practices will undoubtedly benefit such applications as well, given the increased complexity and faster development cycle demands of such systems, particularly in applications that involve control and monitoring systems that demand high reliability, and real-time responsiveness.

Embedded controllers used in the National Ignition Facility (NIF) [2] control systems are a case in point. An "embedded controller" in this context refers to a custom-built computing device that is designed to perform dedicated functions within the NIF control system. The controller is typically enclosed within the machine or equipment and directly controls all aspects of its operation autonomously

without needing to continuously communicate with other parts of the control system. However, they can have control or monitoring interfaces to integrate with the larger control system when required, while still maintaining autonomy.

These controllers are critical to the functioning of the overall NIF control system and typically require high reliability, resilience, and performance. This demands that their development, testing, and deployment procedures be rigorous and extensive. Incorporating software engineering best practices can address many of the challenges associated with these procedures. This will help in automating these procedures to achieve faster development cycles, enhanced quality, and improved system reliability.

However, applying best practices such as CI/CD to embedded controllers presents a unique set of challenges, due to the hardware-software interactions, real-time requirements, safety standards, and the limited computing capabilities. A careful approach is necessary when extending these practices to embedded controllers.

This paper explores the application of software engineering best practices to the development and deployment of embedded controllers used in the NIF control systems, specifically the controllers used in the Master Oscillator Room (MOR). It outlines the potential benefits of this approach, discusses the challenges that need to be addressed, proposes strategies for effectively implementing CI/CD best practices and presents a case study through a pilot project designed to validate this strategy.

## EMBEDDED CONTROLLERS IN NIF MASTER OSCILLATOR ROOM (MOR)

The Master Oscillator Room (MOR) [3] at NIF produces laser pulses from an optical fiber laser, originating with a few nanojoules and a beam diameter of a few micrometers. With four distinct oscillators (Fig. 1) corresponding to different beam cones on the target, each oscillator can be set up independently to generate the low-energy laser pulses.




Figure 1: The NIF master oscillator.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# CONTINUOUS INTEGRATION AND DEBIAN PACKAGING FOR RAPIDLY EVOLVING SOFTWARE\*

A. Barker<sup>†</sup>, J. Georg, M. Hierholzer, M. Killenberg, T. Kozak, D. Rothe, N. Shehzad, C. Willner  
Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, Hamburg, Germany

## Abstract

We describe our Jenkins-based continuous integration system and Debian packaging methods, and their application to the rapid development of the ChimeraTK framework. ChimeraTK is a C++ framework for control system applications and hardware access with a high level of abstraction and consists of more than 30 constantly changing interdependent libraries. Each component has its own release cycle for rapid development, yet API and ABI changes must be propagated to prevent problems in dependent libraries and over 60 applications. We present how we configured a Jenkins-based continuous integration system to detect problems quickly and systematically for the rapid development of ChimeraTK. The Debian packaging system is designed to ensure the compatibility of binary interfaces (ABI) and of development files (API). We present our approach using build scripts that allow the deployment of rapidly changing libraries and their dependent applications as Debian packages. These even permit applications to load runtime plugins that draw from the same core library, yet are compiled independently.

## INTRODUCTION

Minor release upgrades to software libraries typically exhibit binary compatibility. This ensures that projects reliant on these libraries can depend on a stable application binary interface (ABI), in addition to having compatible source code, known as the application programming interface (API).

Maintaining ABI compatibility requires significant effort and hampers the flexibility needed for a rapidly evolving software framework. On the ChimeraTK project [1], the needs of control system applications drive new feature development, and the release timelines should be short. Hence we have decided against preserving binary compatibility for minor software releases. However, this decision comes with the important drawback of always having to re-compile all dependent projects. This warrants special attention when creating software packages and setting up a continuous integration (CI) system.

## DEBIAN PACKAGING

Library packages are binary compatible for the lifetime of a distribution. A binary incompatible package requires a different package name. This is solved by incorporating major and minor version numbers into the package name, so that different versions are formally completely different

packages. Binary compatible patches remain viable since the patch level is not part of the package name.

A library depending on changing binary sources is not itself binary stable, even without source code changes. To enable new releases based on changes to the dependencies, we have extended the minor version to include a build version which changes as dependencies change. This build version also contains the distribution code name. Since this is part of the minor version, it also becomes part of the so-version, such that the C++ linking layer is taking it into account. For example: for major version 3, minor version 11, with Ubuntu20.04 code-named “focal”, and build number 1, the Debian package is named `libchimeratk-deviceaccess03-11-focal1` and the `.so`-file is `libChimeraTK-DeviceAccess.so.03.11focal1`.

Our packaging script [2] has a dependency database that stores all dependency versions for each build. Then the build version is increased if the dependency versions change. In addition, the script has an inverse dependency lookup mechanism to identify all libraries that depend on the library being rebuilt. Then their build versions are also increased and they are recompiled, so resulting in a new, binary consistent ensemble of libraries.

Debian packages for applications are usually not rebuilt, so untouched applications remain intact and unaffected by changes in the framework. Packages for applications also have the major, minor, and build version in their package names. Thus, there is a dedicated package for each release. In contrast to libraries, packages from the same application exclude each other’s installation. This is done by introducing a virtual package via the the package metadata, noting “provides” and “conflicts” with the package base name without version number. However, the use of package names with version number allows for easy roll-back to a previous version if needed.

## Binary Compatible Plugins

ChimeraTK DeviceAccess is a C++ library that provides access to hardware or device servers. Different communication protocols are accommodated by so-called device backends. These include PCI express and Linux UIO [3]. Other backends accommodate various control system middleware like DOOCS [4], EPICS [5] and OPC UA [6]. Generic applications, like the graphical user interface QtHardMon or the DeviceAccess Python bindings, are not linked against all the backends. This prevents the user from having to install all of those control system software stacks, even if they will not be used. DeviceAccess has a runtime loading mechanism to make these backends work with the generic applications, and this runtime loading requires binary compatible

\* We acknowledge support from DESY (Hamburg, Germany), a member of the Helmholtz Association HGF.

<sup>†</sup> anthony.barker@desy.de

# FACING THE CHALLENGES OF EXPERIMENT CONTROL AND DATA MANAGEMENT AT ESRF-EBS

J. Meyer, S. Debionne, W. De Nolf, S. Fisher, A. Götz,  
M. Guijarro, P. Guillou, A. Homs Puron, V. Valls  
ESRF, Grenoble, France

## Abstract

In 2020, the new European Synchrotron Radiation Facility Extremely Brilliant Source (ESRF-EBS) took-up operation. With the much higher photon flux, experiments are faster and produce more data. To meet these challenges, a complete revision of data acquisition, management and analysis tools was undertaken. The result is a suite of advanced software tools, deployed today on more than 30 beamlines. The main packages are Beamline Instrument Support Software (BLISS) for experiment control and data acquisition, Library for Image Acquisition 2 (LIMA2) for high-speed detector control, the ESRF Workflow System (EWOKS) for data reduction and analysis workflows, and Daiquiri the web GUI framework. These developments provide a solid foundation to tackle current and future challenges at the ESRF and to continue providing cutting-edge capabilities for synchrotron research.

BLISS is programmed in Python, to allow easy sequence programming for scientists and easy integration of scientific software. BLISS offers: Configuration of hardware and experimental set-ups, a generic scanning engine for step based and continuous data acquisition, live data display, frameworks to handle 1D and 2D detectors, spectrometers, monochromators, diffractometers and regulation loops.

For detectors producing very high data rates, data reduction at the source is important. LIMA2 allows parallel data processing to add the necessary compute power (CPU and GPU) for online data reduction in a flexible way.

The EWOKS workflow system can use online or off-line data to automate data reduction or analysis. Algorithms are wrapped as a workflow module and can be combined with other processing modules. Workflows can run locally or on a compute cluster, using CPUs or GPUs. Results are saved or fed back to the control system for display or to adapt the next data acquisition.

## ESRF-EBS

The ESRF-EBS is the first fourth generation high-energy synchrotron with a Hybrid Multi-Bend Achromat (HMBA) lattice [1]. The new X-ray beam is 100 times more brilliant and coherent than before, Fig. 1. For the construction the ESRF shut-down in December 2018. The new storage ring was built within the existing infrastructure and reduced the energy consumption by 30%. In September 2020 the new synchrotron took-up operation. With its much higher photon flux, experiments can be conducted faster and more in-situ experiments are carried-out.




Figure 1: View of the X-ray beam before (top) and with EBS (bottom).

## THE CHALLENGES

To cope with the new specifications of faster data acquisition, higher data rates and the increasing need for online data reduction, a large part of the ESRF software for beamlines had to be modernized after 30 years of continuous operation.

The main challenges to be addressed are:

- Data acquisition rates in the kHz and up to the MHz range.
- To limit the infrastructure cost and impact on the environment, the raw data to be stored must be reduced as early as possible during experiments.
- Standardization of all experimental raw and meta data for generic visualization, processing, automation and data curation.
- Automated online data analysis for immediate feedback of experiment results and avoid storing useless raw data.
- Experiment registration with all raw data, meta data and processed data produced. Remote users should have immediate access to experiment results.

The EBS shutdown was a unique opportunity to re-write the main data acquisition software, prepare for faster detectors with higher data rates, develop a workflow system for online data reduction and data analysis and to modernize the GUI framework

A clean management of the data flow is required on all levels. All the different software tools must work seamlessly together (see Fig. 2). All data visualization, data saving, data analysis and experiment registration depend on the data produced by the acquisition system.

## DATA ACQUISITION

### *BLISS: Scanning and Sequencing*

The replacement campaign of the old SPEC-based data acquisition system, was started during the EBS shutdown in 2018 to the situation today, where 35 beamlines are running with BLISS.

# A BEAMLINER AND EXPERIMENT CONTROL SYSTEM FOR SLS 2.0

K. Wakonig<sup>†</sup>, C. Appel, A. Ashton, S. Augustin, M. Holler, I. Usov, J. Wyzula, X. Yao  
Paul Scherrer Institute, Villigen PSI, Switzerland

## Abstract

The beamlines of the Swiss Light Source (SLS) predominantly rely on EPICS standards as their control interface but in contrast to many other facilities, there is up to now no standardized user interfacing component to orchestrate, monitor and provide feedback on the data acquisition. As a result, the beamlines have either adapted community solutions or developed their own high-level orchestration system.

For the upgrade project SLS 2.0, a sub-project was initiated to facilitate a unified beamline and experiment control system. During a pilot phase and a first development cycle, libraries of the Bluesky project were used, combined with additional in-house developed services, and embedded in a service-based approach with a message broker and in-memory database. Leveraging the community solutions paired with industry standards, enabled the development of a highly modular system which provides the flexibility needed for a constantly changing scientific environment. One year after the development started, the system was already tested during many weeks of user operation and recently received the official approval by the involved divisions to be rolled out as part of the SLS 2.0 upgrade.

## INTRODUCTION

The beamline and experiment control system is defined as the software layer tasked with the orchestration and the monitoring of the data acquisition routines and acts as the primary interface for the beamline scientist and users of a larger research facility. At the Swiss Light Source (SLS), Paul Scherrer Institute, Switzerland, this layer was initially often merged with the underlying control system, e.g., EPICS [1] resulting in EPICS as a control system and orchestration system. However, throughout the operation of the SLS, beamlines frequently developed or adopted high-level tools as dedicated orchestration systems. Yet in contrast to other facilities, the SLS currently does not provide a standardized interface for all beamlines. As a result, a large variety of solutions are currently in operation, *inter alia* SPEC[2], DA+ [3], FDA [4], PShell [5], GDA [6].

As these solutions were mostly driven and maintained by the beamline scientists themselves, it has led to an environment in which the beamlines were able to quickly adapt their acquisition routines to changes of the overall requirements. However, it also created an environment where sharing devices, feedback loops and data analysis routines between beamlines was frequently hindered by incompatible interfaces of the bespoke orchestration system. Additionally, already challenging topics such as FAIR

(Findable, Accessible, Interoperable, and Reusable) data, the integration of high-throughput devices or a more automated control of the beamline were hindered further by the increased complexity of supporting various high-level interfaces.

To remedy these shortcomings, a project was initiated to investigate the feasibility of leveraging the SLS 2.0 upgrade program and unifying these interfaces [7], [8]. Starting with the functional and non-functional requirements given by the beamline scientists, the users, and the support groups, the Beamline Instrument Support Software (BLISS) [9] and the Bluesky framework [10], two well-known experiment control systems were chosen to be investigated more thoroughly. Although these two solutions have a similar software architecture, their actual implementation and core concepts differ vastly. Fundamentally, both solutions rely on a Python-based, monolithic main component for steering and monitoring the acquisition. Events and metadata streams are forwarded to a dedicated message broker system (Redis [11] for BLISS, Kafka [12] for Bluesky). Another similarity is that the file writer on a production system should be listening to the message broker events instead of being attached to the core components directly to avoid performance penalties.

BLISS, marketed as an “all-in-one” solution [13], provides a plethora of features with a tight integration into the BLISS ecosystem. The Bluesky framework on the other hand provides fewer features out of the box but relies on a more modular environment, where components can be easily reused without committing into a fully-fledged ecosystem. BLISS is shipped with support for Tango [14] devices, yet often relies on direct device control interfaces. The Bluesky framework on the other hand uses Ophyd, a hardware abstraction layer on top of existing control systems. Although Ophyd was clearly designed with EPICS in mind and supports many different types of EPICS interfaces out of the box, Ophyd can also be used to integrate Tango devices and even custom communication protocols, bypassing the underlying control system.

As the SLS will keep using EPICS as a primary control system after the SLS 2.0 upgrade, a good support for EPICS devices is crucial for a successful adoption of a high-level interface. Although EPICS is predominantly used and is the only officially supported control system at PSI, various custom devices bypassing the EPICS layer have also been implemented. As a result, a hardware abstraction layer such as Ophyd provides a much-needed layer to standardize the diverse landscape at PSI.

## ARCHITECTURE

Following the initial evaluation phase, a prototype project was launched to investigate the potential integration of

<sup>†</sup> klaus.wakonig@psi.ch

# THE SOLID SAMPLE SCANNING WORKFLOW AT THE EUROPEAN XFEL

A. Garcia-Tabares\*, A. Kardoost, C. Deiter, L. Gelisio, S. Hauf, I. Karpics, J. Schulz, F. Shohn  
European X-ray Free Electron Laser Facility GmbH, Holzkoppel 4, Schenefeld, Germany

## Abstract

The fast solid sample scanner (FSSS) used at the European XFEL (EuXFEL) enables data collection from multiple samples while minimizing the need for a sample-holder exchange. To optimise scan duration, target positions can be identified in advance, thus preventing the (X-ray) exposure of empty locations. In this contribution, the automated sample delivery workflow for performing solid sample scanning using the FSSS is described. This workflow covers the entire process, from automatically identifying target positions within the sample, using machine learning algorithms, to setting the parameters needed to perform the scans. The integration of this solution into the EuXFEL control system, Karabo, not only allows to control and perform the scans with the existing scan tool but also provides tools for image annotation and data acquisition. The solution thus enables the storage of data and metadata for future correlation across a variety of beamline parameters set during the experiment.

## INTRODUCTION

Fixed targets are widely used among the instruments at the European XFEL, when using the 10 Hz train repetition mode [1], instead of the MHz rate burst rate. To scan such solid samples a common hardware component, the solid sample scanner, is used at 5 out of the 7 instrumental end stations of the European XFEL. The FSSS is a device equipped with two perpendicular stepper motors, enabling precise scanning along both the X and Y axes, with Z representing the beam direction. This scanner is situated above a hexapod that regulates the distance between the sample and the in-line microscope. Samples scanned using the FSSS include foil targets, wires, structured samples or powder samples encapsulated in Kapton and ultra-thin silicon nitride membranes in microfabricated silicon chips, among others. The sample so scanned thus cover a wide range of applications.

Despite the widespread use and recent hardware advancements [2,3], preparing the sample scanner for measurements remains a time-consuming and lengthy process. Typically, this procedure involves several steps: sample loading, individual sample characterization, data collection, and post-measurement analysis.

Samples are loaded into a designated holder outside the vacuum chamber and then individually placed into the interaction chamber without breaking the vacuum. Subsequent sample characterization occurs during the nominal beam time, usually during the night, and can last between 4 to 8 hours. The data collection process follows immediately after

alignment or is scheduled for the next day depending on the specific instrument.

While the sample loading process benefits from automation [2], the characterization of each target remains a manual task. It involves locating individual target locations before exposure to the X-ray source.

The solid sample scanner workflow of the HED instrument has undergone significant improvements, including pre-characterization of samples prior to beam time and the creation of a dedicated database for storing sample-related data. In this paper, we provide a comprehensive overview of the enhancements and automation implemented within the workflow, demonstrating the effectiveness of these improvements.

## WORKFLOW AUTOMATION

The fundamental workflow has remained largely unchanged; however, significant modifications have been introduced with the aim of automating, standardizing, and accelerating the process of illuminating multiple targets located in one sample. An overview is presented in Fig. 1.




Figure 1: Workflow schema follow during sample scanner after automation.

In the updated workflow, several modifications have been implemented to streamline and expedite the process:

- Target localization before the beam time: targets are located in the sample environment and characterization labs before the experiment. Due to the addition of reference marks, fiducials, to the sample, the targets coordinates can be measured in the lab and then used at the instrument.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

\* ana.garcia-tabares@xfel.eu

# EXPERIMENTAL DATA TAKING AND MANAGEMENT: THE UPGRADE PROCESS AT BESSY II AND HZB\*

R. Müller<sup>†</sup>, Heike Görzig, Gregor Hartmann, Klaus Kiefer, Ruslan Ovsyannikov,  
 William Smith<sup>‡</sup>, Simone Vadilonga, Jens Viefhhaus, HZB, Berlin, Germany  
 Daniel Allan, BNL, Upton, NY, USA

## Abstract

The endeavor of modernizing science data acquisition at BESSY II started 2019. Significant achievements have been made: the Bluesky software ecosystem is now the accepted framework for data acquisition, flow control and automation. It is operational at an increasing number of HZB beamlines, endstations and instruments. Participation in the global Bluesky collaboration is an extremely empowering experience. Promoting FAIR data principles at all levels developed a unifying momentum, providing guidance at less obvious design considerations. Now a joint demonstrator project of DESY, HZB, HZDR and KIT, named ROCK-IT (Remote Operando Controlled Knowledge-driven, IT-based), aims at portable solutions for fully automated measurements in the catalysis area of material science and is spearheading common developments. Foundation there is laid by Bluesky data acquisition, AI/ML support and analysis, modular sample environment, robotics and FAIR data handling. This paper puts present HZB controls projects as well as detailed HZB contributions to this conference into context. It outlines strategies providing appropriate digital tools at a successor 4th generation light source BESSY III.

## INTRODUCTION

From the origins of BESSY II there was no holistic control system environment, spanning accelerator, beamlines and instruments. Fragmentation of data acquisition (DAQ) tools led to specific, intertwined solutions, that blocked sustainable development plans. A concept for modernizing science data acquisition at BESSY II became an obvious need. A process has been started and plans have been documented in 2019 [1]. Due to the existing diversity of device access and software tools, EPICS channel access could be identified as necessary common denominator for integrating simple, complex and legacy systems (Fig. 1). A clear picture for flow control and the data handling software stack was far less obvious then.

Today significant progress has been made in software core areas, leveraging Bluesky [2] and ophyd enabled device integration and abstraction. Achieved status, and HZB relations to wider Bluesky community developments [3], are at the core of this update paper. Specific HZB contributions to Bluesky are presented in other submissions to this conference [4–8].

\* Work funded by BMBF and Land Berlin

<sup>†</sup> roland.mueller@helmholtz-berlin.de

<sup>‡</sup> william.smith@helmholtz-berlin.de

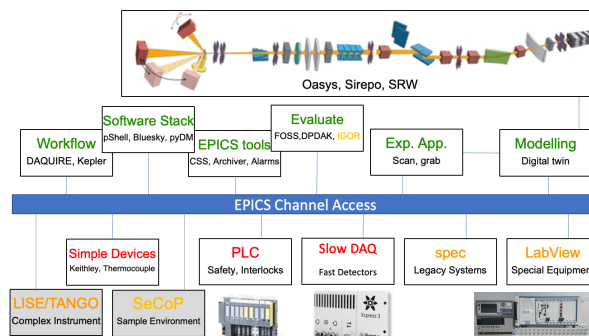


Figure 1: General device, legacy sub-system and digital twin integration guidelines of the 2019 modernisation proposal.

Today, other topics addressed in Ref. [1], like fly scan and machine learning are emphasized within a broader BESSY II facility upgrade plan BESSY II+, dubbed BII+. That plan aims at an elevated level of automation, lowering access barriers for non-expert users, higher instrument science data output and seeks to allow for more complex operando studies while being more sustainable (Fig. 2).




Figure 2: Intended goals of the BII+ facility re-adjustment plan, in terms of scientific focus, general resource saving, modernisation and climate protection.

In an unexpected turn of events, HZB was the victim of a ransomware attack in June 2023. While this caused a challenging set-back, the pause in operation and rebuilding process presents an unprecedented opportunity to speed up the implementation of the more unified controls infrastructure at BESSY II.

## STATUS OF BLUESKY AT BESSY II

Unusual compared to other facilities, progress in the experimental usability of the facility relies very much on grass root type initiatives from engineers and scientists, due to a lack of central steering and under-staffing of experimental control support.

## Deployment and Roll Out

Following initial success at the accelerator, EMIL beamlines and Sissy end stations in 2021, a team formed, en-

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

## DEPLOYMENT OF ADTimePix3 areaDetector DRIVER AT NEUTRON AND X-RAY FACILITIES

Kazimierz J. Gofron<sup>\*1</sup>, Bogdan Vacaliuc<sup>1</sup>, Jakub Wlodek<sup>2</sup>, Greg Guyotte<sup>1</sup>,  
Starra Lyons<sup>1</sup>, Seth Giles<sup>1</sup>, Fumiaki Funama<sup>1</sup>, Su-Ann Chong<sup>1</sup>

<sup>1</sup>Spallation Neutron Source and High Flux Isotope Reactor,  
Oak Ridge National Laboratory, Oak Ridge, TN, USA

<sup>2</sup>National Synchrotron Light Source II, Brookhaven National Laboratory, Upton, NY, USA

### Abstract

TimePix3 is a 65k hybrid pixel readout chip with simultaneous time-of-arrival (ToA) and time-over-threshold (ToT) recording in each pixel. The chip operates without a trigger signal, and employs a sparse readout (i.e., only pixels containing events are read out). The flexible architecture enables achieving up to 40 MHits/s/chip readout throughput @40 MHz clock and up to 80 MHits/s/chip @80 MHz clock, using simultaneous readout and acquisition by sharing readout logic with transport logic of a super pixel matrix formed using a  $2 \times 4$  pixel structure. The chip ToA records 1.5625 ns minimum timing resolution [1].

The X-ray and charged particle events are counted directly by the TimePix3 detector. However, indirect neutron counts use  ${}^6\text{Li}$  fission in a scintillator matrix. The fission energy is generated by the reaction  ${}^6\text{Li}(n, {}^4\text{He}){}^3\text{H} + 4.78 \text{ MeV}$ , where  ${}^6\text{Li}$  is embedded in the scintillator material, such as ZnS(Ag). The fission space-charge region is limited to 5–9  $\mu\text{m}$ . A photon from the scintillator material excites a photocathode electron, which is further multiplied in a dual-stack Microchannel plate (MCP) [2], with an integrated solution such as with an image intensifier tube [3].

Alternatively, a neutron event is detected in thermal neutron capture by  ${}^{10}\text{B}$  doped MCP glass [4] with the nuclear fission reaction  ${}^{10}\text{B}(n, {}^4\text{He}){}^7\text{Li} + 2.31 \text{ MeV}$ . The charged particle products enter microchannels of the first MCP. This design does not use a scintillator. For both thermal neutron detection designs, a neutron “event” is defined by a cluster of electron hits detected by the TimePix3 chip.

We report on the Experimental Physics Industrial Control System (EPICS) [5], areaDetector [6], and ADTimePix3 [7] driver that interfaces with the Amsterdam Scientific Instruments (ASI) [8] SERVAL software using JavaScript Object Notation (JSON) [9] objects. The Serval program controls and distributes data collected from the readout electronics. The ADTimePix3 driver directs data to storage and a real-time processing pipeline, and also configures the chip. The time-stamped data are either stored in a raw .tpx3 [8] file format or passed through a network socket for further processing of the hits. To identify individual neutrons, a clustering software module groups hits into clusters [10]. The conventional 2D images are available as image files for each exposure frame, and a preview is useful for sample alignment. The ADTimePix3 [7] areaDetector driver integrates the time-enhanced

capabilities of this detector into the neutron beamline controls at the Spallation Neutron Source (SNS) and High Flux Isotope Reactor (HFIR), resulting in unprecedented time resolution.

### AREA DETECTOR DRIVER

The ADTimePix3 areaDetector driver was developed using an emulator and SERVAL software provided by Amsterdam Scientific Instruments (ASI) [8]. The Java emulator executes concurrently with the Java SERVAL software to substitute for the real detector. Detector emulation allows for more efficient iterative testing with the EPICS [5] integration and provides for a simpler initial development environment.

After configuring the emulator, we identified the cpr [11] and json [12] software libraries as good candidates for handling our representational state transfer (REST) [13] application programming interface (API) communication—commands and responses are JSON [9] format—and began working on adapting these libraries to build with EPICS. This process required configuring a GNU Makefile to install the required header and library files to directories corresponding to the 64-bit x86 host architecture used for development.

Using a Python cookiecutter [14] template for an EPICS areaDetector [6] driver, ADDriverTemplate [15], we created the skeleton of a driver and determined the linker commands to link those dependencies into our driver executable so that it would compile via EPICS.

After an executable was generated and the external dependencies were linking correctly, we began to build the driver, focusing specifically on communication with the detector emulator. To begin, we wrote the input/output controller (IOC) shell function that initializes the driver to collect some basic diagnostic information upon start-up. Once communication with SERVAL from the EPICS IOC was established, additional readback functionality and control commands were incorporated. A REST API made it easy to break down the work into individual POST/GET requests. Each request was added to the driver, with various EPICS records tied to each field in the command and response.

Unlike many areaDetector drivers, the vendor API did not provide a callback function with direct access to a block of memory holding image data. The vendor software writes binary files directly to disk at a preconfigured location and only posts occasional preview images to a network location accessible by SERVAL via a GET request.

\* gofronkj@ornl.gov



# NEUTRON FROM A DISTANCE: REMOTE ACCESS TO EXPERIMENTS

P. Mutti\*, F. Cecillon, C. Cocho, A. Elaazzouzi, Y. Le Goc, J. Locatelli, H. Ortiz  
Institut Laue-Langevin, Grenoble, France

## Abstract

Large-scale experimental facilities such as the Institute Laue-Langevin (ILL) are designed to accommodate thousands of international visitors each year. Despite the annual influx of visitors, there has always been interest in options that don't require users to travel to ILL. Remote access to instruments and datasets would unlock scientific opportunities for those less able to travel and contribute to global challenges like pandemics and global warming. Remote access systems can also increase the efficiency of experiments. For measurements that last a long time scientists can check regularly on the progress of the data taking from a distance, adjusting the instrument remotely if needed. Based on the Virtual Infrastructure for Scientific Analysis (VISA) platform, the remote access becomes a cloud-based application which requires only a web browser and an internet connection. NOMAD Remote provides the same experience for users at home as though they were carrying out their experiment at the facility. VISA makes it easy for the experimental team to collaborate by allowing users and instrument scientists to share the same environment in real time. NOMAD Remote, an extension of the ILL instrument control software, enables researchers to take control of all instruments with continued hands-on support from local experts. Developed in-house, NOMAD Remote is a ground-breaking advance in remote access to neutron techniques. It allows full control of the extensive range of experimental environments with the highest security standards for data, and access to the instrument is carefully prioritised and authenticated.

## SETTING THE SCENE

Neutron experiments play a pivotal role in various scientific domains, ranging from fundamental physics to materials science and biology. The ability to control and manipulate the measurements remotely has become increasingly important, especially in the wake of global challenges such as the COVID-19 crisis which has dramatically transformed our work practices. The ILL, as many other research centres around the world, has played a crucial role during the pandemic. To allow researchers to dig into the secrets of the virus, the beam lines had to stay operational even during the confinement periods. We had to find quickly solutions to allow scientists, both internal as well as external users, to perform their experiment without being on-site. After the pandemic crisis was over, remote working and the increasing awareness on the climate impact of long air-plane travels, has strengthened those needs for remote access to experiments. Moreover, this possibility has gained prominence due to its potential to overcome geographical, logistical, and

safety constraints. NOMAD [1] is the ILL control software deployed on all 50 instruments in operation at the institute. Its client-server architecture was from the very beginning compliant with the model-view-controller software pattern principle, but it was necessary to confront our architecture with these new use-cases. The adopted solution is depicted in Fig. 1. The CAMEO [2] application manager is the core of the NOMAD ecosystem. It is handling the life-cycle of all the other applications, including NOMAD-Server and Nomad-GUI. It is also responsible for the execution of the real-time plot and, if any, of all third party applications as, for example, the automatic data reduction. Nomad-Server is handling the connection with the central data storage, the electronic logbook and the parameters survey writing into the database as well as the alarms management.

## A Multi-Client Solution

NOMAD has natively multi-client capabilities. Within the ILL we had so far three different clients: the one running on the instrument computer, those installed in the instrument cabins or the offices and a remote Android client running on tablets. All these GUI instances require a specific local software installation and configuration. To extend this concept to a remote client we had, first of all, to find a solution that did not require any intervention on the distant machine. The natural answer came with a web-based remote desktop of a virtual machine (VM). Using OpenStack [3] virtualization infrastructure coupled with Apache Guacamole [4] client-less remote desktop gateway, we could find an optimal compromise accounting for all our needs. First, the security aspects could be totally handled by the web connection, including double factor authentication. Second, the GUI application could be included directly in the image used to generate the VM and kept up-to-date by an automatic synchronization with the software repository. In such a way, we could preserve the data policy and prevent undesired connections thanks to a web-service which connect in a unique way the user demanding a connection with a specific instrument running a well defined experiment. The network integrity is guaranteed by the assignment of specific communication ports between server and client to a user for a specific experiment. The look-&-feel of the GUI is preserved since the user is running on the VM exactly the same application as on instrument.

In our use-case we can distinguish three different user's types:

- *superusers* who are members of the instrument control department and have permission to configure all devices,
- *instrument managers* who drive and supervise the experiment and who can change the instrument configu-

\* mutti@ill.fr

# DYNAMICAL MODELLING VALIDATION AND CONTROL DEVELOPMENT FOR THE NEW HIGH-DYNAMIC DOUBLE-CRYSTAL MONOCHROMATOR (HD-DCM-Lite) FOR SIRIUS/LNLS\*

T. R. S. Soares<sup>†</sup>, J. P. S. Furtado<sup>‡</sup>, G. S. de Albuquerque, M. Saveri Silva, R. R. Gerales  
Brazilian Synchrotron Light Laboratory, CNPEM, Campinas, Brazil

## Abstract

Two new High-Dynamic Double-Crystal Monochromators (HD-DCM-Lite) are under installation in Sirius/LNLS for the new beamlines QUATI (quick-EXAFS) and SAPUCAIA (SAXS), which requires high in-position stability (5 nrad RMS in terms of pitch) whereas QUATI's DCM demands the ability to perform quick sinusoidal scans in frequencies, for example 15 Hz at 4 mrad peak-to-peak amplitude. Therefore, this equipment aims to figure as an unparallel bridge between slow step-scan DCMs, and channel-cut quick-EXAFS monochromators. In the previous conference, the dynamical modelling of HD-DCM-Lite was presented, indicating the expected performance to achieve QUATI and SAPUCAIA requirements. In this work, the offline validation of the dynamical modelling is shown, comparing to the solutions achieved for the previous version of LNLS HD-DCMs. This work also presents the hardware-based control architecture development, discussing the loop shaping technique and upgrades in the system, such as the increase of the position resolution, synchronization of the rotary stages, and FPGA code optimization. Furthermore, it describes how the motion controller was developed, given the high-performance motion control, such as complex control algorithm in parallel with a minimal jitter and the expectations for the beamlines commissioning regarding detector and undulator synchronization.

## INTRODUCTION

Two new High-Dynamic Double-Crystal Monochromators (HD-DCM-Lite) are currently undergoing the offline commissioning at Sirius [1]. In this context, the dynamic behavior of all components that involve motion is being collected experimentally and validated together with the model provided by the mechanical design. The controller for each degree of freedom is also being designed. In addition to that, new mechatronic tools to control such unprecedented systems are being developed by the LNLS team members. This development is important because the HD-DCM-Lite is an engineering challenge, and the development of systems to solve problems related to precision and stability is essential for a laboratory which aims to be innovative and ambitious in technology.

The identification and control system present in this work are based on frequency domain analysis and the operation is

based on a NI CompactRIO (cRIO) 9049 [2] – so the mechatronic tools that will be shown require high performance programming in terms of resources and timing optimization for high-speed calculations in single precision, on a hardware-based platform. This work presents the results of the offline commissioning and perspectives for the online commissioning.

## ARCHITECTURE

### Cartesian Convention Guideline

In the next sections of this paper, the actuators and sensors that belong to HD-DCM-Lite will be explained. Each one of them controls a degree of freedom which is aligned with the standard axes of the laboratory. The goniometers (rotary stages) that control the 1st crystal – and consequently the Bragg angle ( $\theta$ ) – have a degree of freedom that is aligned with the  $x$  axis of the laboratory. The Short-Stroke, responsible for maintaining the beam at the same height  $H$  after meeting the 2nd crystal, controls three degrees of freedom: GAP (or gap, a translation aligned with  $y'$  axis), PTC (or pitch, a rotation aligned with the  $x$  axis) and RLL (or roll, a rotation aligned with the  $z'$  axis). These orientations are shown in Fig. 1.




Figure 1: Cartesian convention guideline for the monochromator.

### Goniometers

As seen in [3], the goniometers of HD-DCM-Lite are mechanically connected. To achieve scan results closer to the expectations demanded by the quick-EXAFS experiments, a gantry closed loop will be designed in the next few months. For the preliminary results and offline commissioning, the goniometers are controlled individually – each one with the respective controller. The goniometers – Aerotech APR260-DR180 for ultra-high vacuum [4] – are brushless actuators,

\* Work supported by the Ministry of Science, Technology and Innovation.

<sup>†</sup> telles.soares@lnls.br

<sup>‡</sup> joao.furtado@lnls.br

# OPTIMISATION OF THE TOUSCHEK LIFETIME IN SYNCHROTRON LIGHT SOURCES USING BADGER\*

S.M. Liuzzo<sup>†</sup>, N. Carmignani, L. Carver, L. Hoummi,  
D. Lacoste, A. Le Meillour, T. Perron, S. White, ESRF, Grenoble, France  
I. Agapov, M. Boese, L. Malina, E. Musa, J. Keil, B. Veglia  
Deutsches Elektronen-Synchrotron, 22607 Hamburg, Germany  
T. Hellert, Lawrence Berkeley National Laboratory, Berkeley, CA, USA  
A. Edelen, P. Raimondi, R. Roussel, Z. Zhang, SLAC, Menlo Park, CA, USA

## Abstract

Badger [1] is a software designed to easily access several optimizers (simplex, RCDS [2], Bayesian optimization, etc.) to solve a given multidimensional minimization/maximization task. The Badger software is very flexible and easy to adapt to different facilities. In the framework of the EURIZON European project, Badger was used for the EBS and PETRA III storage rings interfacing with the Tango and TINE control system. Among other tests, the optimisations of Touschek lifetime was performed and compared with the results obtained with existing tools during machine dedicated times.

## INTRODUCTION

In the framework of the EURIZON European project, a collaboration has been established between DESY and ESRF to work on the optimization of beam parameters for storage ring (SR) light sources such as lifetime and injection efficiency. In a first approach, the Extremum Seeking algorithm was tested in both SRs. It was soon found to be difficult to tune and did not bring major improvements for operation [3]. Contrary to this initial attempt, the software Badger [1] developed at SLAC has demonstrated to be extremely easy to set up, flexible to use and user friendly. It provides access to the Xopt algorithm library [4]. We present in these pages the user experience (UX) and the experimental results achieved thanks to the use of Badger and Xopt at the ESRF-EBS and DESY-PETRA III storage rings.

## XOPT

Xopt [4] is a high level python [5] package developed at SLAC National Accelerator Laboratory that provides a simple to use framework for connecting black box optimization algorithms with arbitrary optimization problems. Xopt decomposes optimization problems into three key components: defining parameter spaces and objectives, specifying how to evaluate these objectives, and implementing optimization algorithms. The VOCS class defines the optimization space, including the variables, objectives, constraints and constants (statics). The Evaluator class defines how to

evaluate objectives and constraints given points passed to it using serial or parallel (multithreading, MPI etc.) processes. Finally, the Generator class implements the optimization algorithm, and is used to generate points in variable space to be evaluated. The main Xopt object choreographs the execution and communication between these three modules in order to perform an optimization cycle with the step() command.

This modular, object-oriented approach enables easy modification and customization of optimization routines for specific use cases. For example, the same optimization algorithm can be applied to both simulation and experiment, or shared between different accelerator facilities by swapping out the Evaluator object. On the other hand, generators can also be swapped out to compare the performance of different algorithms on the same optimization problem. These objects can also be sub-classed to customize evaluation or generation of points to solve specific problems.

Optimization algorithms are defined by Generator objects, which are used to generate future points to be evaluated by calling their generate() method. While users are free to implement their own optimization algorithms, Xopt comes pre-packaged with a number of conventional and advanced optimization algorithms tuned by experts to be applicable to a wide variety of optimization problems “off-the-shelf”. Currently these algorithms include

- Autonomous Characterization
  - Bayesian Exploration [6]
- Single Objective Optimization
  - Nelder-Mead Simplex [7]
  - Robust Conjugate Direction Search [8]
  - Extremum Seeking [3]
  - Upper Confidence Bound BO [9]
  - Expected Improvement BO [10]
  - Trust Region BO (TuRBO) [11]
  - Multi-fidelity BO [12]
- Multi-Objective Optimization
  - Continuous NSGA-II [13]
  - Expected Hypervolume Improvement BO [14]

\* This project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No. 871072

<sup>†</sup> simone.liuzzo@esrf.fr

# IMPLEMENTATION OF MODEL PREDICTIVE CONTROL FOR SLOW ORBIT FEEDBACK CONTROL IN MAX IV ACCELERATORS USING PyTango FRAMEWORK

Carla Takahashi\*, Aureo Freitas, Magnus Sjöström, Jonas Breunlin  
MAX IV Laboratory, Lund, Sweden

Emory Jensen Gassheld, My Karlsson, Pontus Giselsson, LTH, Lund, Sweden

## Abstract

Achieving low emittance and high brightness in modern light sources requires stable beams, which are commonly achieved through feedback solutions. The MAX IV light source has two feedback systems, Fast Orbit Feedback (FOFB) and Slow Orbit Feedback (SOFB), operating in overlapping frequency regions. Currently in MAX IV, a general feedback device implemented in PyTango is used for slow orbit and trajectory correction, but an MPC controller for the beam orbit has been proposed to improve system robustness. The controller uses iterative optimisation of the system model, current measurements, dynamic states and system constraints to calculate changes in the controlled variables. The new device implements the MPC model according to the beam orbit response matrix, subscribes to change events on all beam position attributes and updates the control signal given to the slow magnets with a 10 Hz rate. This project aims to improve system robustness and reduce actuator saturation. The use of PyTango simplifies the implementation of the MPC controller by allowing access to high-level optimisation and control packages. This project will contribute to the development of a high-quality feedback control system for MAX IV accelerators.

## INTRODUCTION

MAX IV Laboratory is a fourth-generation synchrotron light-source facility comprising a 3 GeV storage ring, a 1.5 GeV storage ring, and a linear accelerator, which serves a dual role as a full-energy injector into the storage rings and as the driving source for the Short Pulse Facility. Annually, the laboratory accommodates approximately 1000 users spanning academia, research institutes, industry, and government agencies, all of whom gain access to the facility through dedicated user access programs. Remarkably, MAX IV consistently delivers beam currents of 400 mA and 500 mA on the 3 GeV and 1.5 GeV storage rings, respectively, meeting the diverse experimental requirements of its user community.

The control system at MAX IV is structured in a three-layer architecture. At its core, the middle layer employs the TANGO[1] distributed control framework, effectively serving as the interface between the multitude of equipment within the facility and the oversight of their operations. Crucial tasks are managed by dedicated hardware components, ensuring reliable and precise operation. Above everything,

from top client layer, users are able to employ Python and Matlab scripting languages to interface with the TANGO system.

## Orbit Feedback on MAX IV

Transverse stability of the beam is an important aspect of achieving the low emittance and high brightness goals of any modern light source. This is commonly achieved via feedback solutions but the implementations vary; at the MAX IV light source there are two separate feedback systems working together in two different but overlapping frequency regions as well as sets of sensors. The Fast Orbit Feedback, so named because of its 10 kHz repetition rate, is capable of attenuating noise up to 50 – 150 Hz in the most critical regions. This is hardware-based controller based on Libera Brilliance Plus beam position monitor systems. However, it also has the disadvantage of relying on relatively weak actuators that easily saturate. The SOFB controller, which on its own handles noise and orbit drifts from DC up to roughly 50 mHz, therefore also periodically off-loads the FOFB. It also covers the complete set of sensors. In contrast to the FOFB the SOFB, working at much lower rates, could be and was implemented in software as a feedback device in TANGO [2, 3].

The sensors of the both the FOFB and the SOFB systems are Beam Position Monitors (BPMs), for the horizontal and vertical planes, which are read out using commercial Libera Brilliance Plus electronics. Overall, there are 200 BPMs in the 3 GeV ring for each of the horizontal and vertical planes and 36 BPMs per plane on the 1.5 GeV ring. These devices are interfaced in TANGO and the beam positions in both planes are available as attributes. The attributes push events at the 10 Hz rate of the slow Libera data acquisition stream, which provides 2 Hz BW position data. The events are timestamped according to the local clock in the Libera crate, which is synced via NTP, so the position data from all BPMs are read and pushed every 0.1 s stamped with the same hardware time [4].

The actuators in the SOFB are the slow corrector magnets controlled by ITest BILT BE2811 power supplies. There are 380 power supplies in the 3 GeV ring, in which 200 drive horizontal plan magnets and 180 are used in the vertical plane, and 72 in the 1.5 GeV ring, 36 in each plane. Previous timing studies were performed to confirm that these can handle remote commands to change the current output at 10 Hz. However the power supplies of the corrector magnets often saturate, which sometimes require human intervention

\* carla.takahashi@maxiv.lu.se

# COMMISSIONING AND OPTIMIZATION OF THE SIRIUS FAST ORBIT FEEDBACK

D. O. Tavares\*, M. S. Aguiar, F. H. Cardoso, E. P. Coelho, G. R. Cruz, F. H. de Sá, A. F. Giachero,  
L. Lin, S. R. Marques, A. C. S. Oliveira, G. S. Ramirez, E. N. Rolim<sup>†</sup>, L. M. Russo  
Brazilian Synchrotron Light Laboratory (LNLS/CNPEM), Campinas, Brazil

## Abstract

The SIRIUS fast orbit feedback (FOFB) system started operation for users in November 2022. The system design targeted an output disturbance rejection crossover frequency of 1 kHz by pursuing the minimization of loop delay. This paper gives an overview of the system architecture and technology choices, and reports on the commissioning, system identification and feedback loop optimization performed along the system's first year of operation.

## INTRODUCTION

Figure 1 depicts the layout of the SIRIUS orbit feedback systems in one of the 20 sectors of the storage ring, showing the location of Beam Position Monitors (BPMs) and orbit corrector magnets.

A slow orbit feedback (SOFB) loop runs at an update rate of approximately 10 Hz to mitigate orbit drifts and errors caused by thermal effects, ground motion and residual misalignment. It comprises a set of 160 BPMs (320 beam position readings in total) as sensors and 281 actuators, more specifically, 120 horizontal and 160 vertical orbit corrector magnets plus the storage ring RF frequency.

The FOFB system, on the other hand, runs at an update rate of 48 kHz and is implemented upon a subset of 80 BPMs (160 beam position readings in total) and 80 dedicated fast orbit corrector magnets (160 corrector coils total)<sup>1</sup>, targeting the rejection of fast and small amplitude orbit disturbances in the range of 0.1 Hz to 1 kHz, typically given by magnets vibration, magnets' power supply electrical ripple, injection transients and insertion devices movements. Despite being ready to include in the loop all the existing electron BPMs and beamline front-end XBPMs, the current configuration only includes electron BPMs adjacent to the light sources points.

## DESIGN PRINCIPLES

The design of the SIRIUS FOFB system was driven by the minimization of the overall feedback loop delay. The main pursued design principles were:

1. Adopt the highest possible feedback loop update rate to decrease BPM decimation filters delay. Ideally, the loop

rate should be high enough to make the data distribution network latency the dominant component of the delay budget.

2. Deploy high bandwidth actuators, providing an ideally flat frequency response above the closed loop sensitivity function crossover frequency.
3. Implement the real-time processing (FOFB feedback controller, data network and power supply's feedback loop) entirely in hardware (e.g., FPGAs) and make system integration as tight as possible to prevent unnecessary delays.
4. Use the minimum amount of sensors and actuators in the loop, just enough to provide exact correction at the light sources and avoid extra latency for data distribution.

Principle 1 has driven the selection of the BPMs electrode-switching frequency at a submultiple of the FOFB update rate. This makes the switching spurs lie exactly at the notches of the BPM Cascaded Integrator-Comb (CIC) decimation filters and gives a wide band for the decay of the decimation filter's frequency response, hence largely simplifying the filter's specification. Currently, the update rate is limited to 48 kHz (12 turns in the storage ring) due to the particular implementation of the data network in the FPGA, which does not allow pipelining multiple loop iterations, however there is no fundamental hardware limitation that prevents reaching higher rates in the future, for instance 96 kHz (6 turns).

Principle 2 drove the design of a thin 0.3 mm stainless steel vacuum chamber and low power linear amplifiers for the fast orbit correctors' power supply. A target bandwidth of 10 kHz (1 decade above the target crossover frequency) has been established. In combination with principle 3, the low power requirements led to the design of a compact fast corrector power supply in MicroTCA.4 Rear Transition Module (RTM) form factor. Its digital current regulation loops implemented in the same FPGA of the main FOFB controller. Therefore, the interface between the orbit feedback controller and the correctors' feedback controller is nearly free of delay.

Principle 4 drove the separation of the slow and fast loops, in the same way other facilities have already adopted. It segments the requirements of high range low speed correctors and low range high speed correctors, greatly simplifying power supply design. The amount of data to be exchanged in the FOFB network and its associated latency is also minimized.

\* daniel.tavares@lnls.br

<sup>†</sup> erico.rolim@lnls.br

<sup>1</sup> The actual number of fast correctors included in the FOFB loop is currently 78 (156 coils). Two magnets of the injection straight section were temporarily repurposed to perform injection disturbance feed-forward with separate power supplies and control system.

# MODELLING AND CONTROL OF A MeerKAT ANTENNA

I. Dodia<sup>†</sup>, South African Radio Astronomy Observatory, Cape Town, South Africa  
also at University of Cape Town, Cape Town, South Africa

## Abstract

This paper presents a comprehensive approach to modelling for control system design for a MeerKAT antenna. It focuses on dynamic modelling using time and frequency domain techniques, and lays the foundation for the design of a control system to meet the telescope's stringent pointing and tracking requirements. The paper scope includes rigid body modelling of the antenna, system identification to obtain model parameters, and building a system model in Simulink. The Simulink model allows us to compare model performance with the measured antenna pointing, under various environmental conditions. The paper also integrates models for pointing disturbances, such as wind and friction. The integrated model is compared to the existing control setup. Wind disturbance plays a significant role in the pointing performance of the antenna, therefore the focus is placed on developing an appropriate wind model.

This research will conclude by providing a well-documented, systematic control system design that is owned by SARAO and can be implemented to improve the pointing performance of the telescope.

## INTRODUCTION

The MeerKAT Radio Telescope located in South Africa's Northern Cape province is a precursor instrument to the Square Kilometre Array (SKA) project. MeerKAT consists of 64 antennas each with 13.5 m dish structures. The antennas are arranged in a spiral array with a maximum baseline (point-to-point distance between antennas) of 8 km. MeerKAT observes in a range of frequencies ranging from about 500 MHz to upwards of 4 GHz with the use of multiple receiver systems. The telescope was commissioned in 2018 and is in active use.

Each antenna can move to point its dish in two axes: azimuth and elevation. Due to the large dish size and tight requirement specifications, pointing an individual antenna is a critical component in achieving optimal imaging fidelity. Antenna pointing and tracking is done by the Antenna Positioner (AP) subsystem - a commissioned "black box" subsystem that was integrated into MeerKAT during the telescope's development.

Now that MeerKAT is in its operational phase, issues related to the AP subsystem are being uncovered; without a comprehensive understanding of the design, troubleshooting and resolving these issues is challenging. This paper focusses on building a comprehensive model of the mechanical dynamics of the AP system, considering parameter uncertainty and accurately capturing disturbances. This will lay the foundation for future work towards a systematic feedback design for MeerKAT.

<sup>†</sup> idodia@sarao.ac.za

## BACKGROUND

### Drive Configuration

The individual antennas of the MeerKAT telescope use a 13.5 m Gregorian offset dish, configured in an elevation over azimuth arrangement. A pedestal provides a foundation on which a yoke is attached to achieve this configuration. Figure 1 identifies the major drive components of a MeerKAT antenna.

Rotation in the azimuth axis of the antenna is achieved by the use of two motor drives that are attached to the yoke structure. Each motor drive consists of a servo drive with a high stiffness, low friction planetary gearbox that drives output pinions. These pinions interface with the outer teeth of the static azimuth gear and drive the entire yoke (and elevation structure) in azimuth. The motor drives are electronically torque biased to minimise backlash. Figure 2 highlights the key components of the azimuth drive assembly. This drive configuration is typical in high precision radio astronomy applications [1]. An azimuth on-axis encoder is located inside the pedestal and reads out the position of the antenna in azimuth for the control loop.

Rotation in the elevation axis is achieved by means of a jackscrew drive (ball screw driven by a torque servo motor) that is mounted on the yoke and attaches to the back of the main reflector on the back up structure. Mechanical pre-loading of the jackscrew drive is done to minimise backlash. Linear actuation of the jackscrew causes rotational movement of the main reflector around a pivot point. Counterweights attached to the back up structure maintain the centre of mass of the dish to be about this pivot point, ensuring that there is no resulting unbalanced load torque on the elevation axis. The elevation encoder is also located on the axis of this pivot point, as shown in Fig. 2.




Figure 1: Azimuth and elevation drive assemblies, drawing taken from [1].

# PATH TO IGNITION AT THE NATIONAL IGNITION FACILITY(NIF): THE ROLE OF THE AUTOMATIC ALIGNMENT SYSTEM

Bela Praful Patel, Abdul A. S. Awwal, Mikhail Fedorov, Richard R. Leach, Jr.,  
Roger Lowe-Webb, Vicki Miller Kamm, Payal Kamlesh Singh

Integrated Computer Control System, National Ignition Facility, Lawrence Livermore National  
Laboratory, Livermore, CA 94550, United States of America

## Abstract

The historical breakthrough experiment at the National Ignition Facility (NIF) produced fusion ignition in a laboratory for the first time and made headlines around the world. This achievement was the result of decades of research, thousands of people, and hardware and software systems that rivaled the complexity of anything built before. The NIF laser Automatic Alignment (AA) system has played a major role in this accomplishment. Each high yield shot in the NIF laser system requires all 192 laser beams to arrive at the target within 30 picoseconds and be aligned within 50 microns-half the diameter of human hair-all with the correct wavelength and energy. AA makes it possible to align and fire the 192 NIF laser beams efficiently and reliably several times a day. AA is built on multiple layers of complex calculations and algorithms that implement data and image analysis to position optical devices in the beam path in a highly accurate and repeatable manner through the controlled movement of about 66,000 control points. The system was designed to have minimum or no human intervention. This paper will describe AA's evolution, its role in ignition, and future modernization.

## INTRODUCTION

On December 5<sup>th</sup>, 2022, a long sought-after and challenging milestone was achieved within the NIF facility at Lawrence Livermore National Laboratory [1, 2] by successfully producing fusion ignition for the first time. This important accomplishment enables a path towards solving countless national and world problems including the ability to provide abundant, clean energy for generations to come. The automated alignment system [3, 4] is a vital part of NIF and provided major contributions towards this achievement.



Figure 1: On December 5<sup>th</sup>, 2022, fusion ignition was achieved within the NIF facility at Lawrence Livermore National Laboratory. Imploding target within the hohlraum is depicted on the right.

Inside the NIF facility during what is called a shot, all 192 lasers are amplified, conditioned, aligned, and focused on their path to finally enter the hohlraum at the target chamber center, (Fig. 1) This generates X-rays that cause the target capsule (Fig. 2) within to implode that cause the deuterium and tritium atoms inside to fuse. As a result, ionized helium nuclei (alpha particles) are released into the surrounding fuel, and their deposited kinetic energy results in rapid heating of the surrounding fuel which causes a cascade of fusion events known as ignition when the deposited energy overcomes energy loss processes in the imploding fuel. The initial spark of fusion in the imploded hot spot needs to be sufficiently strong to cause ignition, and a higher temperature hot spot can increase the initial fusion spark [5].

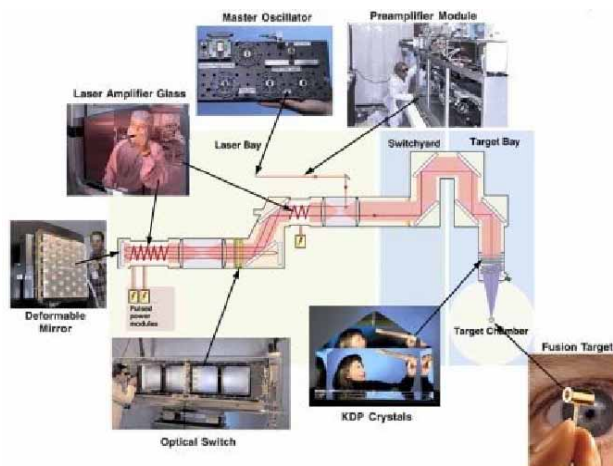


Figure 2: Schematic representation of a NIF laser beam line highlighting some of the key technologies.

The success of a fusion shot is contingent on precision alignment of 192 laser beams towards the target by moving thousands of optical and electro-mechanical components under the control of the AA system. The AA system is a data driven software framework with the ability to position optical and mechanical devices and align 192 high power laser beams accurately and consistently with minimal or no human interaction. To accomplish this, AA provides automated execution of multiple layers or loops of complex algorithms using data and image analysis which will be described in the following sections.

## AUTOMATIC ALIGNMENT STRUCTURE

NIF AA is a data driven software framework to set up the devices at correct positions and align 192 laser beams accurately and consistently with minimum or no human

System Modelling

Feedback Systems & Optimisation

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# ENERGY CONSUMPTION OPTIMISATION USING ADVANCED CONTROL ALGORITHMS

F. Ghawash<sup>1</sup>, B. Schofield\*, E. Blanco, CERN, Geneva, Switzerland

<sup>1</sup> also at Department of Engineering Cybernetics,  
Norwegian University of Science and Technology, Trondheim, Norway

## Abstract

Large industries operate energy-intensive equipment with energy efficiency now an important objective to optimize overall energy consumption. CERN, the European Laboratory for Particle Physics, uses a large amount of electrical energy to run its accelerator complex, with a total yearly consumption of 1.3 TWh and peak usage of up to about 200 MW. Energy consumption reduction can be achieved through technical solutions and advanced automation technologies. Optimization algorithms, in particular, have a crucial role not only in keeping the processes running within the required safety and operational conditions, but also in optimising the financial factors at play. Model-based Predictive Control (MPC) is a feedback control algorithm that naturally integrates the capability of achieving reduced energy consumption when including economic factors in the optimization formulation. This paper reports on the experience gathered when applying non-linear MPC on one of the contributors to the electricity bill at CERN: the cooling and ventilation plants (i.e. cooling towers, chillers, and air handling units). Simulation results on cooling tower control showed significant performance improvements, and energy savings close to 20 %, when compared to conventional heuristic solutions. The control problem formulation, the control strategy validation using a digital twin and the initial results in a real industrial plant are reported here, together with the experience gained implementing the algorithm in industrial controllers.

## INTRODUCTION

Environmental concerns together with rising energy prices are bringing energy consumption into focus across society as a whole. At CERN, there has been considerable effort recently to reduce energy consumption, which currently stands at roughly 1.3 TWh per year. Ideally, energy consumption reduction would be realized without impacting the operation of the accelerator complex and experiments. One attractive means to achieve this, which typically also requires minimal capital expenditure, is to improve control algorithms. In many cases, control algorithms are developed with focus on aspects such as performance or robustness. Optimization of energy usage is often not considered at the design stage. Furthermore, controllers in an industrial context are almost always designed using a setpoint-tracking paradigm, and generally implemented using Proportional-Integral-Derivative (PID) controllers. However, there are

numerous applications in which such a paradigm may be unnecessary, and may lead to unwanted excess energy consumption. Many such examples can be found in cooling and ventilation applications such as cooling towers and building HVAC systems. In these cases the control objective is almost always to keep process values (such as temperature and humidity) within a range of acceptable values, rather than at one specific setpoint. Using a classical setpoint-tracking approach here can lead to unwanted energy consumption by forcing the system to this (somewhat artificial) specific setpoint. At CERN, cooling and ventilation applications account for a considerable proportion of electrical consumption (66 GWh in 2019). More widely, studies indicate that in developed countries, energy consumption in buildings accounts for roughly a third of total energy consumption [1]. Improving control systems for these applications could therefore enable considerable energy savings. For these types of processes, it would be desirable to design a controller which ensures that process values are kept within their desired ranges, but which does not use excessive control effort to drive them to a specific setpoint. It would also be desirable if the control formulation could explicitly include the minimization of energy consumption as a control objective. MPC is a control design approach which allows this to be accomplished. By setting up a cost function which can contain both control objectives as well as energy or economic costs, and by including constraints, which in these applications is a more natural way of expressing the control objective, it is possible to formulate an optimization problem. By repeatedly solving this problem over a fixed, finite time horizon, controller output values can be obtained. In order to predict future process values over the time horizon, a model of the system is required. In this article, the principle of MPC will be introduced. A case study of an Air Handling Unit (AHU) will be considered. The derivation of a simple control-oriented process model will be presented, and a model predictive controller will be formulated. The validation of the controller using a process simulator (virtual commissioning) will be discussed. Finally, the results of applying the model predictive controller will be presented, with the focus on achievable energy savings.

## MODEL PREDICTIVE CONTROL

MPC uses a dynamic model of the process within a finite-horizon optimization problem to determine a set of control inputs. Only the first of these control inputs are applied to the plant, and the optimization problem is solved again in the next time step. This iterative approach provides the

\* brad.schofield@cern.ch



# CONTROL DESIGN OPTIMIZATIONS OF ROBOTS FOR THE MAINTENANCE AND INSPECTION OF PARTICLE ACCELERATORS

A. Díaz Rosales\*<sup>†1,2</sup>, H. Gamper\*<sup>1,3</sup>, M. Di Castro<sup>1</sup>

<sup>1</sup> European Organization for Nuclear Research (CERN), 1211 Meyrin, Switzerland

<sup>2</sup> Department of Cognitive Robotics, Delft University of Technology, 2628 Delft, The Netherlands

<sup>3</sup> Johannes Kepler University, Linz, 4040 Linz, Austria

## Abstract

Automated maintenance and inspection systems have become increasingly important over the last decade for the availability of the accelerators at CERN. This is mainly due to improvements in robotic perception, control, and cognition and especially because of the rapid advancement in artificial intelligence. The robotic service at CERN performed the first interventions in 2014 with robotic solutions from external companies. However, it soon became clear that a customized platform needed to be developed in order to satisfy the needs and in order to efficiently navigate through the cluttered, semi-structured environment. This led to the formation of a robotic fleet of about 20 different robotic systems that are currently active at CERN. In order to increase the efficiency and robustness of robotic platforms for future accelerators it is necessary to consider robotic interventions at the early design phase of such machines. Task-specific solutions tailored to the specific needs can then be designed, which in general show higher efficiency than multipurpose industrial robotic systems. This paper presents the latest developments and techniques for designing and developing robotic systems that are specific to certain tasks, such as maintenance and inspection in particle accelerators. We explore the necessary requirements for a robotic system, including the control strategies that are applied, and the optimization of the system's topology and geometry.

## INTRODUCTION

The fourth industrial revolution, the current trend of automation and data interconnection in industrial technologies, is leading to a boost for maintenance and availability for space applications, warehouse logistics, particle accelerators, and harsh environments [1]. The main pillars of Industry 4.0 are the Internet of Things (IoT), Wireless Sensors, Cloud Computing, Artificial Intelligence (AI) and Robotics. Core to success and future growth in this field is the use of robots to perform various tasks, particularly repetitive, unplanned, or dangerous tasks, which humans either prefer to avoid or are unable to carry out due to hazards, size constraints, or the extreme environments in which they take place.

During the last decade at the European Organization for Nuclear Research (CERN) [2], robotic technologies have been developed and integrated within the accelerators to support maintenance tasks reducing human exposure to hazards

and boosting machines availability [3]. The advancements in robotic perception, control, and cognition, particularly in artificial intelligence, have contributed to this development. The CERN robotic service initially used external company solutions for interventions but later had to create customized platforms to meet their specific requirements and navigate the cluttered and semi-structured environment efficiently. This led to a robotic fleet of about 20 different robotic systems [4,5]. In order to increase the efficiency and robustness of robotic platforms for future accelerators it is necessary to consider robotic interventions in the early design phase of new machines. Task-specific solutions tailored to particular needs can then be designed, which in general show higher efficiency than universal robotic systems [6, p. 284].

This approach is currently applied to the design of the new robotic manipulators at CERN. This paper presents the latest progress in producing a task-specific robotic system designed for maintenance and inspection. We will use the 100 km long main tunnel of the Future Circular Collider as an example, along with a system for examining Radio Frequency cavities. The requirements for such a robotic system are described in the Section on *Requirements and Restrictions*. Based on these findings a design optimization concerning the topology and geometry of the robotic system was performed and used as a starting point for the mechanical design of the different components (Section *Design Optimization*). The structure of the control strategies used on these robots is discussed in the Section on *Control Strategies*, while the Section on *Software Architecture* explains the workflow of the software developments and their deployment on the robots. The two different use cases are then explained in their own sections before concluding with opportunities for future work.

## RELATED WORK

We initially compared universal systems to task-specific solutions. Universal systems excel in unstructured environments demanding advanced locomotion, perception, and cognition, but they can become complex and less robust due to numerous components. In contrast, task-specific systems excel in structured environments, optimizing efficiency, availability, and robustness.

Considering the semi-structured nature of most of the environments found at CERN, we decided to focus on task-specific systems tailored to specific needs. We conducted a study looking at relevant examples from other facilities:

Joint European Torus (JET): JET's task-specific robotic system [7] excels in the challenging tokamak chamber, per-

\* Equal contribution of authors.

<sup>†</sup> alejandro.diaz.rosales@cern.ch

# CONTROL SYSTEM DEVELOPMENT AT THE SOUTH AFRICAN ISOTOPE FACILITY (SAIF)

J. K. Abraham, H. P. Anderson, L. S. Anthony, S. Baard, C. Baartman, A. H. Barnard, P. Beukes, J. I. Broodryk, P. J. Davids, M. E. Hogan, J. P. Mira, N. Pakade, A. Phillips, M. Robertson, J. J. Solomons, G. F. Steyn, N. P. Stodart, I. L. Strydom, iThemba LABS, Faure, South Africa  
P. J. Bester, N. de Wet, TF Design, Stellenbosch, South Africa  
W. D. Duckitt, Stellenbosch University, Stellenbosch, South Africa

## Abstract

The first phase of the South African Isotope Facility (SAIF) at iThemba LABS is well into commissioning. The intention of SAIF is to free up our existing separated sector cyclotron (SSC) facility to do more physics research and to increase our radioisotope production and research capacity. An EPICS based control system, primarily utilising EtherCAT hardware, has been developed that spans the control of beamline equipment, target handling and bombardment facilities, radiation protection and safety systems. Various building and peripheral services like cooling water and gases, HVAC and UPS have also been integrated into the control system via Modbus TCP/IP and OPC UA to allow for seamless control and monitoring. An overview of the SAIF facility and the EPICS based control system is presented. The control strategies, hardware and various EPICS and web-based software and tools utilised are presented.

## INTRODUCTION

iThemba Laboratory for Accelerator Based Sciences (iThemba LABS) is a multi-disciplinary research center offering accelerator facilities for various scientific research, training, radioisotope production for nuclear medicine, and accelerator mass spectrometry services. The K=200 separated sector cyclotron (SSC) facility has been in operation for more than 30 years. The facility operates two solid-pole injector cyclotrons, one for light ions and one for heavy ions and polarized protons, for the injection of beams into the K=200 SSC. The beamtime from the SSC facility is split between nuclear physics research and radioisotope production.

### Existing Radioisotope Production

The routine production of short-lived radioisotopes utilising the SSC facility started in 1988 at the horizontal beam target station (HBTS aka Elephant). Towards the late 1990s, production methods for the long-lived radioisotopes <sup>73</sup>As, <sup>68</sup>Ge, <sup>22</sup>Na and <sup>82</sup>Sr were developed. In 1996 a second target station (Babe) was designed and built for production of semi-permanent targets including <sup>18</sup>F. Production of <sup>18</sup>F was later transferred to a dedicated Siemens 11 MeV cyclotron. In 2006, the vertical beam target station (VBTS) was commissioned to exploit high-intensity proton beams delivered by the upgraded SSC. Development focussed on the production of long-lived and high-value radioisotopes such as <sup>68</sup>Ge, <sup>22</sup>Na and <sup>82</sup>Sr in a tandem configuration. These upgrades and developments were all undertaken to increase

radioisotope production in order to meet the high market demand.

### South African Isotope Facility

Increased radioisotope production using the SSC facility would require additional allocated beam time. However, this would come at the expense of research and other programmes. The South African Isotope Facility (SAIF) was established to increase the beam time for radioisotope production and research, whilst simultaneously freeing up the SSC facility to do more physics research. SAIF will be realised in two functional phases: (i) a centre for the production of exotic isotopes and (ii) a centre for the production of exotic beams.

The first phase of SAIF is well into commissioning. A Cyclone® 70 (C70) particle accelerator system, comprising of a high-current 70 MeV H<sup>-</sup> cyclotron and four attaching beamlines, has been procured from Ion Beam Applications (IBA) and installed [1]. Target bombardment and handling facilities and a number of beamline components have been designed and manufactured. Hardware and software for the control and monitoring of the manufactured systems, radiation protection, safety and various building and peripheral systems have been developed. The remainder of this paper addresses these SAIF control systems.

## ARCHITECTURE

The SAIF control system is largely built on the EPICS framework and mostly utilises off-the-shelf Beckhoff EtherCAT I/O terminals and third party PLCs to interface with equipment<sup>1</sup> [2]. As shown in Fig. 1, the SAIF control system consists of 3 hierarchical layers communicating through the *Control Network* - a flat Gigabit Ethernet network using the TCP/IP protocol.

**Network** Two Ethernet networks exist at iThemba LABS: the *Campus Network* and the *Control Network*. The first one, as its name states, is deployed across the iThemba LABS campus and accesses the Internet; the second one is restricted to equipment control and does not access the Internet.

<sup>1</sup> This discussion excludes the Cyclone® 70's control system. This is a turnkey proprietary control system developed and installed by IBA. A brief overview of the C70's control system, with a mention of integration with EPICS, is given in the next section. Contact the manufacturer for a detailed and up-to-date specification.

# ONLINE MODELS FOR X-RAY BEAMLINES USING SIREPO-BLUESKY

B. Nash\*, D. Abell, M. Keilman, P. Moeller, I. Pogorelov, RadiaSoft LLC, Boulder, Colorado, USA  
Y. Du, A. Giles, J. Lynch, T. Morris, M. Rakitin, A. Walter,  
NSLS-II, Brookhaven National Lab, Upton, New York, USA  
N. Goldring, State 33 Inc., Portland, Oregon, USA

## Abstract

Synchrotron radiation beamlines transport X-rays from the electron beam source to the experimental sample. Precise alignment of the beamline optics is required to achieve adequate beam properties at the sample. This process is often done manually and can be quite time consuming. Further, we would like to know the properties at the sample in order to provide metadata for X-ray experiments. Diagnostics may provide some of this information but important properties may remain unmeasured. In order to solve both of these problems, we are developing tools to create fast online models (also known as digital twins). For this purpose, we are creating reduced models that fit into a hierarchy of X-ray models of varying degrees of complexity and runtime. These are implemented within a software framework called Sirepo-Bluesky that allows for the computation of the model from within a Bluesky session which may control a real beamline. This work is done in collaboration with NSLS-II. We present the status of the software development and beamline measurements including results from the TES beamline. Finally, we present an outlook for continuing this work and applying it to more beamlines at NSLS-II and other synchrotron facilities around the world.

## INTRODUCTION

X-rays in synchrotron light sources or XFELs are produced by an electron beam radiating during passage through a magnetic field. The electron beam is stored in a storage ring or accelerated directly from a linac. In either case, there are typically substantial diagnostics to ensure that the electron beam has adequate properties to produce the required radiation. One of the important technologies enabling this well-controlled electron beam is the use of online models in which a physics model of the beam dynamics is coupled to beam diagnostics such that the physics model is continuously updated to match the measurements. A similar technology is typically not used on the photon beam. With increased X-ray brightness and coherence in both synchrotron light sources and XFELs, we propose to develop the tools to enable this technology.

There are several pieces required to implement an online model. One needs the reduced models to describe the beam propagation and one needs a software framework to integrate this into the beamline control system. Here we present further progress in the development of reduced models for use during real-time operation of X-ray beamlines. We also

describe how these models have been integrated into the Sirepo framework that includes interfaces to the full X-Ray beamline simulation codes such as Shadow [1] and SRW [2, 3].

A related technology for beamline control that we are developing is that of Bayesian Optimization. Here, the measured beam is directly related to beamline control parameters such as mirror positions or slit settings. A probabilistic model based on Gaussian processes is then used to directly optimize the measured beam distribution. This is a very effective approach to automate beamline alignment and reconfiguration, even if the resulting model is less interpretable. We describe substantial progress in this approach, as it fits into the same software ecosystem including Sirepo-Bluesky.

In [4], we introduced the concept of a matrix-aperture beamline composed of linear transport sections and physical apertures as shown in Fig. 1. This approach is an approximation with the hope of capturing important transport properties in a computationally efficient manner. Within this approach, there exists a hierarchy of methods<sup>1</sup> as shown in Fig. 2. The first row of the table involves second-moment propagation representing Gaussian Wigner functions [6]. The second row of the table involves propagating coherent electric fields via linear canonical transform (LCT). Progress in the creation of an LCT transport library is reported in [7]. The final row of the table represents generic partially coherent X-ray propagation via Wigner function passing through the matrix-aperture beamline. Some work towards developing this method was presented in [8]. The focus of this paper will be the top-level method of sigma matrix transport through the matrix-aperture beamline. We refer to this reduced model as the Gaussian Wigner function moment (GWFM) model. This model provides a computationally efficient calculation of the linear optics through the beamline, while also including effects of partial coherence. We apply the sigma matrix transport method to a KB mirror beamline with two apertures and compare results with SRW and Shadow. Finally, the realistic case of an NSLS-II beamline is treated with this method, and preliminary results are presented.

We remind the reader that the goal of such fast reduced models is to enable the creation of online models incorporating up-to-the-moment diagnostics data such that the model accurately reflects the true state of the beamline settings and X-ray transport from source to sample. Such an online model

<sup>1</sup> We do not intend to be comprehensive in our scope of all work pertaining to reduced models here. The hybrid method [5] may fit closely within our schema as an alternative to the LCT method, combining wavefront propagation with ray tracing.

\* bnash@radiasoft.net

# WEB TECHNOLOGY ENABLING FAST AND EASY LARGE EXPERIMENTAL FACILITY CONTROL SYSTEM IMPLEMENTATION

W. Zheng, L.Y. Wang, M. Zhang, X.H. Xie, P.L. Zhang, W.J. Ye, H.B. Ma, International Joint Research Laboratory of Magnetic Confinement Fusion and Plasma Physics, State Key Laboratory of Advanced Electromagnetic Engineering and Technology, School of Electrical and Electronic Engineering, Huazhong University of Science and Technology, Wuhan, China

## Abstract

Large experimental facilities are essential for pushing the frontier of fundamental research. Control system is the key for smooth operation for Large experimental facilities. Recently many new types of facilities have emerged, especially in fusion community, new machines with completely different designs are being built. They are not as mature as accelerators. They need flexible control systems to accommodate frequent changes in hardware and experiment workflow. The ability to quickly integrate new device and sub-systems into the control system as well as to easily adopt new operation modes important requirements for the control system. Here we present a control system framework that is built with standard web technology. The key is using HTTP RESTful web API as the fundamental protocol for maximum interoperability. This enables it to be integrated into the already well developed ecosystem of web technology. Many existing tools can be integrated with no or little development. Such as InfluxDB can be used as the archiver, NodeRed can be used as the scripter and docker can be used for quick deployment. It even made integration of in house developed embedded devices much easier. In this paper, we will present the capability of this control system framework, as well as a control system for field reversed configuration fusion experiment facility implemented with it.

## INTRODUCTION

Large experimental facilities play a crucial role in advancing fundamental research, requiring sophisticated control systems to ensure smooth operations. One key characteristic of as large experimental facility control system is its ability to adopt and integrate new systems. This necessity becomes even more evident in emerging fields such as fusion research, where innovative machines with distinct designs are being constructed. There are various control system framework existing in big physics community, each with different characteristics. For fusion community, ITER chose the Experimental physics and industrial control system (EPICS) and Channel Access protocol (CA) as the common language. EPICS has been the common language to the accelerator control system for decades [1]. Now chosen by ITER, it is used by many tokamaks as well [2-5]. It is mature and well supported by the community.

But the technologies used in tokamaks and other experiment facilities are different from those in accelerators. It is not a straight forward job to create EPICS support for equipment used in fusion experiment. Later emerged control system frameworks such as Tango uses object-oriented

technique to improve interoperability and flexibility [6-8]. But still, it is hard to have all the equipment in a control system supporting the control system framework that you chose.

Unlike traditional accelerators, many new facilities demand flexible control systems capable of adapting to frequent changes in hardware and experiment workflows. The ability to swiftly integrate new devices and subsystems, coupled with the ease of adopting novel operational modes, are critical requirements for such control systems.

This paper presents a control system framework developed using standard web technology, with HTTP RESTful web API serving as the foundational protocol to enhance interoperability. By leveraging web technology, the framework seamlessly integrates into the well-established web ecosystem, enabling easy integration with various tools and technologies. Existing tools like InfluxDB can be utilized as archivers, NodeRed as scripters, and Docker for rapid deployment. Additionally, the framework simplifies the integration of in-house developed embedded devices. This paper showcases the capabilities of this control system framework and its application in a field-reversed configuration fusion experiment facility.

The first section of this paper delves into the rationale behind the incorporation of web technology in control systems, highlighting its transformative potential [9]. The second part talks about a toolkit called CFET2 that is used to build a Web based control system. The 3rd part is about using CFET2 with existing web related technologies. Then in section 4 a few real world applications are presented.

## WEB TECHNOLOGIES FOR CONTROL SYSTEMS

Web technologies is an important part of our internet life. We keep using it every day. Web technologies not only power the web site. Today from mobile apps, online games, to smart sensor and IoT application, web plays a big role in them. Many devices have embedded web servers, and many client apps is running in browsers. They communicate using HTTP. The foundation of web technology is HTTP, HTML, and browser.

At its core, the web embodies the essence of interoperability. It serves as a unifying platform where diverse technologies seamlessly converge, facilitating the exchange of information across disparate systems and devices. One of the fundamental building blocks enabling this interoperability is the Hypertext Transfer Protocol (HTTP). HTTP operates as a text-based request and response protocol, cre-

# FAST BEAM DELIVERY FOR FLASH IRRADIATIONS AT THE HZB CYCLOTRON

J. Bundesmann<sup>†</sup>, G. Kourkafas, A. Denker<sup>1</sup>, Helmholtz-Zentrum Berlin, Berlin, Germany  
P. Mühldorfer, Berliner Hochschule für Technik, Berlin, Germany  
J. Heufelder, A. Weber, Charité – Universitätsmedizin Berlin, Berlin, Germany  
<sup>1</sup>also at Berliner Hochschule für Technik, Berlin, Germany

## Abstract

In the context of radiotherapy, Flash irradiations mean the delivery of high dose rates of more than 40 Gy/s, in a short time of less than one second. The expectation of the radio-oncologists are lesser side effects while maintaining the tumour control when using Flash. Clinically acceptable deviations of the applied dose to the described dose are less than 3%.

Our accelerator control system is well suited for the standard treatment of ocular melanomas with irradiation times of 30 s to 60 s. However, it is too slow for the short times required in Flash. Thus, a dedicated beam delivery control system has been developed, permitting irradiation times down to 7 ms with a maximal dose variation of less than 3%.

## INTRODUCTION

### Proton Therapy of Ocular Melanoma

Proton therapy of ocular melanoma is a well-established and successful treatment. At the Helmholtz-Zentrum Berlin (HZB), patients are treated in cooperation with Charité – Universitätsmedizin Berlin since 1998. Overall, more than 4500 patients have been treated. The local tumour control is 96% after five years [1, 2, 3].

Protons permit the confinement of the dose to the tumour. At HZB, we use a proton beam with 68 MeV having a range in water of 38 mm. With this energy, the dose drops at the end of the Bragg peak from 90% to 10% within less than 1 mm. This permits the sparing of organs at risk like the optical nerve or the macula.

For eye tumours the typically prescribed total dose is 60 Gy which is applied in four fractions over four days. The irradiation time is 30 s to 60 s.

### FLASH Irradiations

While the tumour control using protons for ocular melanoma is excellent, there is still a wish for further improvement: the reduction of side effects. A very vibrant research field for reducing side effects while maintaining tumour control is the so-called FLASH irradiation [4]: the dose rate is increased drastically, at least 40 Gy/s and the irradiation time is less than 1 s. At the moment, neither the ideal dose rate nor the ideal irradiation time is known. The normal dose rate for ocular melanomas is less than 0.5 Gy/s. The huge increase in dose rate provides challenges for the dosimetry like linearity of the dose monitors and saturation

effects. The medical physicists require that variations in dose delivery should be less than 3%. Hence, in order to start experiments with FLASH irradiations a new beam delivery control system had to be developed.

## LAYOUT OF THE ACCELERATORS

Figure 1 shows the layout of HZB's accelerator complex for proton therapy. The k130 cyclotron is served by two injectors: either a 6 MV Van-de-Graaff or a 2 MV Tandetron. Besides the treatment room is the experimental room. The control system for the accelerator is based on Vsystem [5].

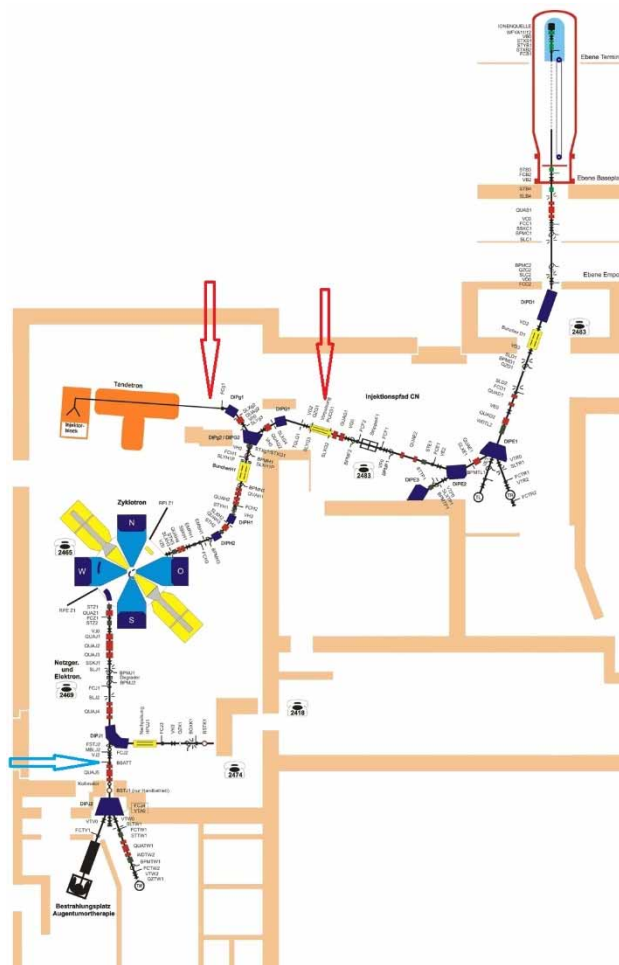


Figure 1: Layout of the proton therapy accelerator complex of HZB. The red arrows mark the position of electrostatic deflectors, the blue arrow the position of the fast beam shutter.

<sup>†</sup> bundesmann@helmholtz-berlin.de.

# USING BDD TESTING IN SKAO: CHALLENGES AND OPPORTUNITIES

V. L. Allan, University of Cambridge, Cambridge, UK\*  
G. Brajnik, University of Udine and IDS, Udine, Italy  
R. Brederode, SKAO, Macclesfield, UK  
on behalf of the The SKA Software Collaboration

## Abstract

The SKAO (Square Kilometre Array Observatory) is one observatory, with two telescopes on three continents. It will be the world's largest radio telescope, and will be able to observe the sky with unprecedented sensitivity and resolution. The SKAO software and computing systems will largely be responsible for orchestrating the observatory and associated telescopes, and processing the science data, before data products are distributed to regional science centres. The Scaled Agile Framework (SAFe™) is being leveraged to coordinate over thirty lean agile development teams that are distributed throughout the world. In this paper, we report on our experience in using the Scaled Agile Framework, the successes we have enjoyed, as well as the impediments and challenges that have stood in our way.

## INTRODUCTION

In this paper, we will provide an account of our attempts to adopt Behaviour Driven Development (BDD) and system testing, particularly for our control system based on Tango, with the goal of providing testers of control systems for other instruments with enough information to decide whether to use such an approach themselves. We will briefly provide the context of the SKAO (SKA Observatory), then look at our testing goals and challenges, focusing on automated testing. We will then explain what BDD testing is and what it has to offer us, and the progress we have made towards our goals while trying to use the approach. We look at the challenges imposed for creating testware to implement BDD tests for finite state automata. We will also discuss the issues we have experienced with roll-out, particularly in the context of the control system, and our current plans. We explore the organisational structures that hinder and help us, documenting our experience and conclusions for the benefit of future decision makers.

## THE SKA PROJECT

The Square Kilometre Array Observatory is one observatory, running two telescopes, over three continents. The headquarters are in the UK, and there are two telescopes: one is a low frequency (50-350MHz) telescope, consisting of hundred of thousands dipole antennas grouped together into stations (referred to as the Low telescope) in the desert in Western Australia, and the other is a mid frequency telescope (350MHz-15.4GHz) consisting of hundreds of dishes (the Mid telescope) in the Karoo desert in South Africa [1,2].

\* vla22@cam.ac.uk

To build these telescopes, a global collaboration spanning multiple countries and timezones has been established. Rees provides a more detailed account of the current state of the project [3].

In Fig. 1, we can see the two telescopes (the Low Array and Dish sub-systems on the diagram) are each connected to a Central Signal Processor (CSP) and a Science Data Processor (SDP), plus numerous supporting sub-systems (Synchronisation and Timing, Network Manager, Telescope Manager Control (TMC), a High Performance Compute Platform, plus eventually a Very Long Baseline Interferometry (VLBI) sub-system and SRCs (SKA Regional Centres)). These systems are used to collect astronomical signals from the sky (using the antennas and dishes in the desert), correlate those signals in the CSP for each telescope, then process those data on supercomputers, turning them into a product that can be delivered to SRCs for use by scientists. Both telescopes are controlled using the Tango control system, so most sub-systems will have one or more Tango devices which allow the transfer of commands and monitoring data, including system health data [4].

Each of these sub-systems is made up of multiple components. For example, the TMC contains a Central Node, which interfaces with the tools for defining telescope observations, a Subarray node, which controls a subset of the dishes or antennas for the relevant telescope, plus components to control and monitor the CSP, SDP, and dishes/antennas for each telescope.

To organize the software development work to create these components, we use the Scaled Agile Framework®(SAFe®), which allows us to co-ordinate multiple development teams on a common cadence [5]. Our teams are grouped into Agile Release Trains (ARTs), which deliver the telescope software. These ARTs span several countries; indeed, some teams are also highly distributed.

The work is cadenced using PIs (Planning Intervals) that last one quarter; each of these begins with a review of the previous PI, and then there is a Planning week where we converge on our plans for the subsequent PI, based on goals that are set by software architects and product managers (PMs). While we can pivot during a PI, that is generally not desirable; therefore major changes to or evolution of our testing process will usually occur not more than once per quarter.

## TESTING GOALS AND CHALLENGES

Brajnik *et al.* have noted that the SKAO has several goals and challenges when considering software testing [6]. Our

# LESSONS FROM USING PYTHON GraphQL LIBRARIES TO DEVELOP AN EPICS PV SERVER FOR WEB UIs

Rebecca Auger-Williams\*, Observatory Sciences Ltd, St Ives, UK

Abigail Alexander, Tom Cobb, Martin Gaughran, Austen Rose, Alexander Wells, Andrew Wilson  
Diamond Light Source, Harwell, UK

## Abstract

Diamond Light Source is currently developing a web-based EPICS control system User Interface (UI). This will replace the use of EDM and the Eclipse-based CS-Studio at Diamond, and it will integrate with future Acquisition and Analysis software. For interoperability, it will use the Phoebus BOB file format. The architecture consists of a back-end application using EPICS Python libraries to obtain PV data and the query language GraphQL to serve these data to a React-based front end. A prototype was made in 2021, and we are now doing further development from the prototype to meet the first use cases. Our current work focuses on the back-end application, Coniql, and for the query interface we have selected the Strawberry GraphQL implementation from the many GraphQL libraries available. We discuss the reasons for this decision, highlight the issues that arose with GraphQL, and outline our solutions. We also demonstrate how well these libraries perform within the context of the EPICS web UI requirements using a set of performance metrics. Finally, we provide a summary of our development plans.

## INTRODUCTION

Diamond Light Source is about to undergo a significant upgrade as part of its Diamond II project, including new beamlines and other accelerator upgrades. Improvements to existing technologies are also being considered as part of this initiative, which stimulated an assessment of the control system UIs currently being used at Diamond. These are predominantly EDM [1] and CS-Studio (Eclipse) [2], the latter of which is now deprecated and has been replaced by Phoebus [2]. Two alternatives are being considered, either Phoebus or a web-based UI, both of which would require a significant amount of effort to move to. A web browser UI has many advantages—there is no installation required, it is truly cross-platform, and it offers the best experience for remote usage—and Diamond therefore developed a prototype version of a web-based UI. This front-end application is built with React [3], one of the most popular JavaScript libraries for building web applications due to its use of components that make it fast, scalable, and simple to use. Redux [4] is used for the data management, which helps maintain a global state across the application. The application itself is written in TypeScript.

A back-end Python application has also been created, named Coniql [5], which uses EPICS [6] Python libraries to

\* rjw@observatorysciences.co.uk




Figure 1: A schematic showing how the front-end web UI receives PV data from EPICS. The web UI uses a WebSocket to connect to Coniql. It uses the GraphQL query language to send requests for data following a defined GraphQL schema. Coniql uses the Python aioca library to subscribe to updates from the requested EPICS PV and returns these to the web UI in a GraphQL query response.

access process variable (PV) data and a GraphQL [7] Python library to serve these data to the web UI via WebSockets [8]. A schematic of the application front-end and back-end is shown in Fig. 1.

## TECHNOLOGIES

This section outlines the technologies that Coniql uses to provide the back-end functionality. Figure 2 shows a diagram of how these fit together in order to supply the front-end web UI with EPICS PV data.

### EPICS Python Library

The Python library aioca [9] is used as the EPICS channel access (CA) client to provide access to EPICS PVs running on IOCs. This application is built on top of asyncio [10] to allow asynchronous querying. The API supports three main function calls: `caget`, `caput`, and `camonitor`. These resemble the EPICS CA command line tools.

### GraphQL

GraphQL is an open-source querying language and runtime engine that was originally developed by Facebook. It can be used to create fast and stable applications where the client can request only the data needed, with the results returned in a predictable format. This reduces the amount of data sent over the network and hence improves performance.

## PROTECTING YOUR CONTROLS INFRASTRUCTURE SUPPLY CHAIN

B. Copy <sup>\*</sup>, J.B. de Martel, F. Ehm, P. Elson, S. Page <sup>†</sup>, M. Pratoussy, L. Van Mol  
CERN Beams Department, 1211 Geneva, Switzerland

### Abstract

Supply chain attacks have been constantly increasing since being first documented in 2013. Profitable and relatively simple to put in place for a potential attacker, they compromise organizations at the core of their operation. The number of high profile supply chain attacks has more than quadrupled in the last four years and the trend is expected to continue unless countermeasures are widely adopted.

In the context of open science, the overwhelming reliance of scientific software development on open-source code, as well as the multiplicity of software technologies employed in large scale deployments make it increasingly difficult for asset owners to assess vulnerabilities threatening their activities.

Recently introduced regulations by both the US government (White House executive order EO14028) and the EU commission (E.U. Cyber Resilience Act) mandate that both Service and Equipment suppliers of government contracts provide Software Bills of Materials (SBOM) of their commercial products in a standard and open data format. Such SBOM documents can then be used to automate the discovery, and assess the impact of, known or future vulnerabilities and how to best mitigate them.

This paper will explain how CERN investigated the implementation of SBOM management in the context of its accelerator controls infrastructure, which solutions are available on the market today, and how they can be used to gradually enforce secure dependency lifecycle policies for the developer community.

### INTRODUCTION

Supply chain attacks involve a malicious third-party infiltrating an organization by exploiting vulnerabilities in third-party components or software used in critical systems. These attacks brought in the context of accelerator controls can compromise the integrity and reliability of operations, potentially leading to disruptions, data breaches, and unauthorized control over essential infrastructure.

The first documented software supply chain attack in June 2013 caused a distributed denial of service on the South Korean government and several news outlets [1] by leveraging the legitimate auto-update mechanism of the *SimDisk* file-sharing and storage service. When organizations rely on third-party suppliers for components, software, or services, they trust that these suppliers have adequate security measures in place. Suppliers with inadequate practices become vulnerable points of entry for attackers.

In the context of open science and its major reliance on

open-source software, open-source software maintainers replace commercial suppliers as essential pieces in securing the supply chain. While commercial suppliers are akin to regulated factories producing reliable components, open-source maintainers act as skilled artisans crafting components collaboratively. Their level of dedication and engagement is not dictated by commercial agreements and can vary over time, yet open source software constitutes the cornerstone of CERN accelerator operations.

Furthermore, this reliance on open-source software is globally following an accelerating trend, without signs of stabilising. The 2023 Open Source Security and Risk Analysis (OSSRA) on open source trends [2] reports that in 2022, the average software project depends on 595 third-party open-source libraries (a near 200% increase over the last four years), while 48% of the analyzed projects contained high-risk vulnerabilities such as documented proof-of-concept, active exploits or remote code execution opportunities (a minor decrease of 2% since 2021). The OSSRA 2023 report also shows that :

- Organizations are insufficiently fixing high-risk vulnerabilities, with a global 42% increase of their presence in codebases since 2018. Some industries more lucrative to attackers, such as e-commerce, are even reporting a 557% increase over the last five years.
- Open source maintenance is on the decline, with an increasing usage of open-source components that have been without activity over the past 24 months, from 85% in 2018 to 90% in 2022. Reasons most cited by open source maintainers are a lack of recognition and inadequate compensation.

In addition to the inexorable discovery and exploitation of software defects leading to vulnerabilities, open source software is subject to neglect with dreadful consequences. Lack of open source software maintenance opens the door to abuse such as the introduction of malware, *protestware* (software that contains some kind of political protest in addition to its primary function), *typo-squatting* (software that relies on common typographic or spelling errors to spread) and dependency confusion (when a malicious dependency library is downloaded from a public registry rather than the intended private/internal registry, for instance by exhibiting a higher version number).

Without appropriate countermeasures, such abuse is :

- Hard to detect and report upon, as both a suitable dependency taxonomy and communication channels need to be established.
- Difficult to track and to mitigate without the usage of an inventory and automation.

<sup>\*</sup> brice.copy@cern.ch  
<sup>†</sup> stephen.page@cern.ch



# IMPROVING CONTROL SYSTEM SOFTWARE DEPLOYMENT AT MAX IV

B. Bertrand\*, Á. Freitas, A. F. Joubert, MAX IV, Lund, Sweden  
J.T. Kowalczyk, S2Innovation, Kraków, Poland

## Abstract

The control systems of large research facilities like synchrotrons are composed of many different hardware and software parts. Deploying and maintaining such systems require proper workflows and tools. MAX IV has been using Ansible to manage and deploy its full control system, both software and infrastructure, for many years with great success. We detail further improvements: defining Tango devices as configuration, and automated deployment of specific packages when tagging Gitlab repos. We have now adopted conda as our primary packaging tool instead of the Red Hat Package Manager (RPM). This allows us to keep up with the rapidly changing Python ecosystem, while at the same time decoupling Operating System upgrades from the control system software. For better management, we have developed a Prometheus-based tool that reports on the installed versions of each package on each machine. This paper will describe our workflow and discuss the benefits and drawbacks of our approach.

## INTRODUCTION

The MAX IV synchrotron radiation facility in Lund, Sweden started user operations in 2016. It consists of a short-pulse facility and two storage rings with 16 beamlines. The control system has more than 500 virtual and physical machines to configure and maintain, including 24k Tango devices with 134k configurable properties. This requires automation, and the Software group has been using Ansible [1] to manage and deploy the full control system, both software and infrastructure, for 10 years [2]. We previously described [3] how we started moving away from RPM packages, tightly coupled to the Operating System version, to conda [4]. We now have deployments on many of our beamlines using solely conda. We discuss how the packages are created, and the pros and cons of our approach. Next, we describe how our Ansible deployment has changed to better support our common workflows. Finally, we report on a new monitoring tool we use to keep track of the actual state of the deployed code, and why it might differ from the Ansible configuration.

## PACKAGE MANAGEMENT

### Python

When investigating replacing RPM with conda, we started by creating a conda recipe in the source repository using a cookiecutter [5] template. Having to create a recipe in every repository made the adoption and transition quite slow. We developed a GitLab CI pipeline to automatically build conda

packages using Grayskull [6]. Grayskull is an automatic conda recipe generator. It could create recipes for Python packages available on PyPI and from GitHub repositories. We added support for local source distribution [7]. Figure 1 details how the GitLab job creates a sdist package to then automatically generate the conda recipe.

```
auto-build-conda-package:
  extends: .conda_build
  before_script:
    # Generate recipe with grayskull from local sdist
    - /grayskull/bin/python -m build -s
    - mkdir recipe
    - /grayskull/bin/grayskull pypi -m KITS -o recipe dist/*.tar.gz
    # Many entry points (like taurusgui) don't have a --help option... Skip entry point test...
    - sed -i "/ --help/d" recipe/*/meta.yaml
    - cat recipe/*/meta.yaml
    - >
    - grep -q "noarch: python" recipe/*/meta.yaml ||
      { echo "Recipe isn't noarch. Should script be replaced by entry_point? Aborting."; exit 1; }
```

Figure 1: auto-build-conda-package job.

If the recipe is not *noarch*, the job fails. For pure python package, this could be due to defining a script instead of an entry point, which should be fixed by the developer. For packages requiring compilation, a recipe has to be created manually. This is more for safety as grayskull is capable of generating such recipes. As we only have very few such repositories, this is not an issue. When an entry point is detected, grayskull automatically adds a test by running it with the *--help* flag. Unfortunately, applications based on taurus [8] do not support this flag. We remove this test but it would be better to keep it. The pipeline could be improved to keep it if taurus is not in the dependencies. This new pipeline allowed us to create conda packages very easily for all our internal repositories and made the transition from RPM to conda possible.

### C++

We work mostly with Python but also have some C++ Tango device servers. Those had a dependency at build time on Makefiles from Pogo [9], the Tango code generator. To ease the compilation with conda, we migrated the build system to CMake [10]. This was a manual process but was not a huge task as we do not have that many C++ repositories and the *CMakeLists.txt* to create is quite similar between projects. Once a project can be compiled with CMake, and without any Pogo dependency, creating the conda recipe is quite straightforward and similar to how upstream projects like TangoDatabase and TangoTest are built.

### Benefits

Moving from RPM to conda allowed us to separate the deployment from the operating system packaging and the system Python version. We could migrate from CentOS 7 to Rocky Linux 8 deploying exactly the same conda packages. Without this change, we would have had to rebuild all our

\* benjamin.bertrand@maxiv.lu.se

# APPLES TO ORANGES: A COMPARISON OF EPICS BUILD AND DEPLOYMENT SYSTEMS

S. Rose\*, D. Araujo<sup>†</sup>, A. Lindh Olsson<sup>‡</sup>, L. Magalhães<sup>§</sup>  
European Spallation Source ERIC, Lund, Sweden

## Abstract

ESS currently uses two different systems for managing the build and deployment of EPICS modules. Both of these use modules that are packaged and prepared to be dynamically loaded into soft IOCs, based on the `require` module developed at PSI. The difference is the deployment: For the accelerator, we use a custom python script to define and build an EPICS environment which is then distributed on a global NFS share to the production and lab networks. For the neutron instrumentation side, in contrast, we use Conda to build individual EPICS environments, where the individual packages are stored on a shared artifactory server.

In this paper we will provide an overview of some of the challenges, contrasts, and lessons learned from these two different but related approaches to EPICS module deployment.

## E3

### History

The ESS EPICS [1] Environment (e3) [2] is based on the approach developed at PSI which uses the standard EPICS base executable `softIocPVA` to run IOCs. Instead of statically or dynamically linking support modules and compiling them into a custom binary executable, we instead configure and dynamically load the provided shared libraries. This is all done using the module `require` [3] which acts as both a parallel build system for EPICS modules, as well as an EPICS module in its own right that dynamically loads other EPICS modules.

e3 was developed initially by Jeong Han Lee, who introduced the notion of a “wrapper”. After he left, the development and maintenance of e3 was taken over by a small team (the e3 team). This team initially consisted of representatives from all of the main groups that comprise ICS (Integrated Control Systems): Software, Hardware and Integration, and Infrastructure. While the team has evolved over time, it maintains its interdisciplinary nature and tight connections with all of its stakeholders.

### Structure

One challenge when working with community EPICS code is that one will often need to provide site-specific customisation to each support module; this presents, for example, a challenge when needing to update a module: how does one keep local changes in sync? How does one track that which comes from the community, and that which is specific to your site?

e3 resolves this by using a “wrapper”<sup>1</sup>: a separate repository which contains a reference to the community code together with any site-specific modifications that are necessary. This can include patches, site-specific database or template files, as well as any other site-specific build or run-time configuration. It also includes metadata such as a description of the dependencies required to build and load the module into an IOC; see Fig. 1.




Figure 1: Example e3 wrapper

This approach allows for development in line with community collaboration and release best-practices: we can update community releases if critical errors are found (such as a recent `caPutLog` memory leak), as well as develop and test patches before submitting them to the community for review.

In order to do this, e3 has opted to use a fork of PSI’s `require` module. In contrast to traditional EPICS IOCs, this module provides dynamic run-time loading and registering of support modules. It requires additional build configuration, but handles some rudimentary run-time dependency resolution.

One of our main divergences from the PSI `require` module is our use of wrappers. However, there are other significant changes such as our use of “virtual environments”, as well as a standard versioning schema which distinguishes site-specific changes from community releases. This is necessary to avoid (among other things) so-called ‘dependency hell’ (see Fig. 2).




Figure 2: Dependency hell

This is handled by appending a *revision number* to the version. Two builds that use the same source version but dif-

<sup>1</sup> Note that in Conda this is called a ‘recipe’; however, that name distinction is not particularly relevant.

\* simon.rose@ess.eu  
<sup>†</sup> douglas.araujo@ess.eu  
<sup>‡</sup> Anders.LindhOlsson@ess.eu  
<sup>§</sup> lucas.magalhaes@ess.eu

# XILINX ZYNQ UltraScale+ USED AS EMBEDDED IOC FOR BPMs

G. Marinkovic, D. Anicic, R. Ditter, B. Keil, J. Purtschert, M. Roggli  
Paul Scherrer Institut, Villigen, Switzerland

## Abstract

PSI is using the Xilinx/AMD Zynq UltraScale+ (ZynqU+) MultiProcessing System-on-Chip (MPSoC) on different hardware platforms, ranging from VME and CompactPCI-Serial (CPCI-S) cards to customized platforms for high-volume applications like beam position monitors (BPMs) and magnet power supplies (PS).

The first ZynqU+ application at PSI was “DBPM3”, a generic hardware platform for BPMs presently used for cavity BPMs in SwissFEL as well as button and stripline BPMs for the SLS 2.0 project. The DBPM3 platform has a digital backend with a ZynqU+, combined with an application-specific analog RF Front-End (RFFE) with integrated fast JESD204B ADCs.

This paper describes our experience with the application of ZynqU+ for EPICS IOCs of different applications at PSI, with focus on our first BPM applications. However, many of the topics and solutions discussed are also relevant and applicable for other applications.

## INTRODUCTION

PSI contributed to the European XFEL by developing the BPM electronics and Intra-Bunch Train Feedback in collaboration with DESY and CEA/Saclay. The BPM electronics design for the European XFEL is based on the Xilinx/AMD Virtex5 FXT field programmable gate arrays (FPGAs) for low-level digital processing and control system interfacing, forming a hierarchy of FPGAs and CPUs. Shortly after PSI built the SwissFEL accelerator using the Xilinx/AMD 7-Series FPGAs (Artix-7 and Kintex-7), where PSI hardware and FPGA firmware concepts from the European XFEL were reused and adapted to the 7 Series FPGAs and SwissFEL requirements. When the Zynq UltraScale+ became available, we checked if the chip was suitable for new BPM hardware required for a 2<sup>nd</sup> SwissFEL soft X-ray beamline (“Athos”) as well as for SLS 2.0, a major upgrade of the Swiss Light Source.

For the ZynqU+, we evaluated and verified that the ZynqU+ is able to integrate EPICS IOC, application specific real time processing, EVR triggering and other FPGA functionality all on a single chip. Moreover, we checked the performance of the IOC, compiled Petalinux, EPICS and implemented serial JESD204B ADC interfaces at 10 Gbit/s. After this successful evaluation, we chose the ZynqU+ for our future “DBPM3” hardware platform [1], combining functionalities of previously separated boards on one chip and printed circuit board (PCB), and thus reducing the number of connectors and potentially increasing the MTBF of our BPM systems.

## DBPM3 PLATFORM

The DBPM3 unit is designed with stability, reliability and redundancy in mind. The analog electronics is separated from the digital backend. Only digital signals are routed through the connectors between front-end and back-end PCBs, using coplanar multi-pin high speed connectors without a backplane.

### RF Front-End Boards

The RFFE boards of a DBPM3 system combine analog electronics and ADC on the same board, including an EEPROM for storage of calibration data. In order to minimize drift of the measurement data, the RFFE temperature is regulated by multiple heaters, combined with the fan speed regulation of DBPM3 19” housing (see Fig. 1). The fans are supervised in EPICS, enabling early automated detection of fan degradation long before a fan breaks, thus improving system MTBF by avoiding urgent repairs.



Figure 1: DBPM3 unit.

### Mechanics and Power Supply

The DBPM3 can incorporate various front-end board (“daughterboard”) sizes, ranging from six single-width daughterboards with 50 mm x 300 mm size each to two triple-width boards with 152 mm x 300 mm size. Six daughterboard connectors are available on the back-end (one per single-width daughterboard), providing power, clocking, digital IO and serial multi-gigabit transceiver (MGT) links between JESD204 ADCs or DACs on the daughterboards and the ZynqU+ on the back-end board.

The power supply of the DBPM3 unit is redundant but does not support live insertion, since it is only one compact space-saving unit with a common mu-metal shield to minimize stray fields to the RFFEs. The power supply status is monitored by the EPICS alarm handler. If one of the two power supplies breaks, the 2<sup>nd</sup> one is able to power the complete unit, where the supply can then be replaced at the next maintenance day of the accelerator.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# HydRA: A System-on-Chip TO RUN SOFTWARE IN RADIATION-EXPOSED AREAS

T. Gingold\*, A. Arias Vázquez, G. Daniluk, J. Serrano, T. Włostowski, CERN, Geneva, Switzerland  
M. Rizzi, PSI, Villigen, Switzerland

## Abstract

In the context of the High-Luminosity LHC project at CERN, a platform has been developed to support groups needing to host electronics in radiation-exposed areas. This platform, called DI/OT, is based on a modular kit consisting of a System Board, Peripheral Boards and a radiation-tolerant power converter, all housed in a standard 3U crate. Groups customise their systems by designing Peripheral Boards and developing custom gateway and software for the System Board, featuring an IGLOO2 flash-based FPGA. It is compulsory for gateway designs to be radiation-tested in dedicated facilities before deployment. This process can be cumbersome and affects iteration time because access to radiation testing facilities is a scarce commodity. To make customisation more agile, we have developed a radiation-tolerant System-on-Chip (SoC), so that a single gateway design, extensively validated, can serve as a basis for different applications by just changing the software running in the processing unit of the SoC. HydRA (Hydra-like Resilient Architecture) features a triplicated RISC-V processor for safely running software in a radiation environment. This paper describes the overall context for the project, and then moves on to provide detailed explanations of all the design decisions for making HydRA radiation-tolerant, including the protection of programme and data memories. Test harnesses are also described, along with a summary of the test results so far. It concludes with ideas for further development and plans for deployment in the LHC.

## BACKGROUND

The High-Luminosity Large Hadron Collider (HL-LHC) project [1] will increase the luminosity in the LHC in order to provide more frequent collisions in the experiments and maximise their discovery potential. This project brings a number of challenges to the control system of the accelerators. In particular, in the lowest tier, some areas will be exposed to increased levels of radiation, precluding the possibility of installing off-the-shelf electronics. A dedicated work package covers the development of a modular radiation-tolerant platform [2] which can serve as a basis for diverse systems. This allows equipment groups to capitalise on a set of basic building blocks and focus their efforts on their customisation, avoiding unnecessary duplication of developments and increasing overall quality.

\* tristan.gingold@cern.ch




Figure 1: Distributed I/O Tier hardware kit.

## THE DI/OT PLATFORM

The Distributed Input/Output Tier (DI/OT) project aims at providing a modular kit for building systems in the lowest tier of the control stack, directly interfacing to accelerator equipment. The basic kit is illustrated in Fig. 1. Modules are housed in a 3U Europa crate featuring a backplane compliant with the CompactPCI Serial (CPCI-S.0) standard. Having a fully passive backplane prevents problems related to radiation in this critical component.

The left-most slot in the crate is reserved for the system board. From that slot, a star topology in the backplane allows communication with the other cards, called peripheral boards. There are two variants of the kit:

- In the radiation-tolerant variant, the power supply is designed in-house and the system board features a flash-based IGLOO2 FPGA.
- In the non-radiation-tolerant variant, the power supply can be purchased off-the-shelf (a bonus of having chosen a standard format for the crate and backplane) and the system board is based on a Xilinx Zynq Ultrascale+ SoC.

The DI/OT platform does not aim to replace well-established solutions such as Programmable Logic Controllers (PLCs) and modular electronics platforms based on a bus (VME, uTCA, PXIe...). Instead, it focuses on two use cases not well covered by these platforms:

- Electronics exposed to radiation.
- Systems in which the connectivity among boards is fully custom for a given application. The fully-passive backplane of DI/OT and the configurable nature of the system board allow e.g. using some of the copper lanes in the backplane to directly stream ADC data from a peripheral board, connect interlocks, etc.

In the remainder of this article, we will focus on the

# THE DESY OPEN SOURCE FPGA FRAMEWORK

L. Butkowski\*, A. Bellandi, B. Dursun, C. Gümüs, K. Schulz,  
M. Büchler, N. Omidsajedi, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

## Abstract

Modern FPGA firmware development involves integrating various intellectual properties (IP), modules written in hardware description languages (HDL), high-level synthesis (HLS), and software/hardware CPUs with embedded Linux or bare-metal applications. This process may involve multiple tools from the same or different vendors, making it complex and challenging. Additionally, scientific institutions such as DESY require long-term maintenance and reproducibility for designs that may involve multiple developers, further complicating the process. To address these challenges, we have developed an open-source FPGA firmware framework (FWK) at DESY that streamlines development, facilitates collaboration, and reduces complexity. The FWK achieves this by providing an abstraction layer, a defined structure, and guidelines to create big FPGA designs with ease. FWK also generates documentation and address maps necessary for high-level software frameworks like ChimeraTK. This paper presents an overview and the idea of the FWK.

## INTRODUCTION

At the scientific institute DESY [1], Field-Programmable Gate Arrays (FPGAs) have become indispensable components within a wide array of control and diagnostic systems. The adoption of FPGAs in these applications is primarily driven by their exceptional multichannel computation power and unparalleled flexibility. However, the FPGA firmware development process at DESY presents a set of significant challenges. The institute's facilities, including EuXFEL and FLASH, demand long-term support and maintenance, spanning up to 20 years [2]. Over such extended periods, numerous features evolve or undergo modifications, toolchains are updated or replaced, and hardware experiences changes or upgrades. What further complicates the process is the involvement of multiple developers and collaborations, with projects transitioning between responsible persons.

These firmware development challenges coincide with the management of multiple projects, often handled by small teams, and run in parallel with a continuous stream of new developments and rapid prototyping efforts. To further complicate matters, the most recent developments encompass a convergence of multiple technologies, entailing the integration of code written in Hardware Description Language (HDL), High-Level Synthesis (HLS), Embedded C/C++, and Embedded Linux into a unified and coherent design.

To address these challenges of FPGA firmware development head-on, a dedicated firmware framework has been

developed. This framework is an abstraction layer that carries a set of rules, scripts, functions, and procedures that are universal and reusable. The initial framework idea is presented in Figure 1. This framework has been designed to tackle a wide range of issues:

- Universal support for multiple vendor tools through an abstraction layer.
- Consistent firmware construction for a unified approach across projects.
- Reproducible builds for enhanced stability and reliability.
- Code reusability for accelerated development and reduced redundancy.
- Collaboration support for multiple developers working on the same project.
- Faster adaptation to new hardware developments.
- Streamlined application deployment on ready hardware platforms.
- Efficient configuration management for easy customization and flexibility.
- Simplified address space management and integration with ChimeraTK framework [3].
- Automatically generating the technical and user documentation.
- Integrated design verification capabilities.
- Quality assurance enhancements through automation in versioning and builds, ready for continuous integration.
- Ability to combine multiple technologies such as HDL, HLS, Embedded C and Linux into one design.

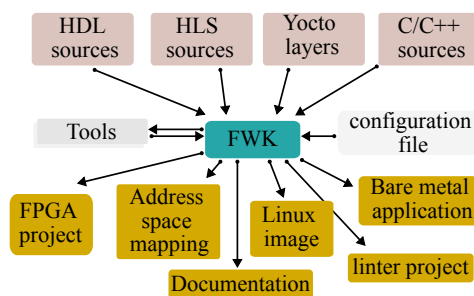


Figure 1: Framework concept.

The initial iteration of the firmware framework was created in 2013 for MTCA.4 systems at EuXFEL [4]. Initially, the framework was closely integrated with the code in a mono repository, leading to challenges related to sharing and scalability as project numbers grew. To address these issues, the decision was taken to decouple the framework from the code, restructure it, and release it as open source. This transition was motivated by the aim of supporting other institutions and facilitating collaboration through open sourcing [5].

\* lukasz.butkowski@desy.de

# DEVELOPMENT OF A TIMING AND DATA LINK FOR EIC COMMON HARDWARE PLATFORM\*

Paul Bachek<sup>†</sup>, Thomas Hayes, Kevin Mernick, Geetha Narayan, Freddy Severino  
C-AD, Brookhaven National Laboratory, Upton, NY, USA  
Joseph Mead, NSLS-II, Brookhaven National Laboratory, Upton, NY, USA

## Abstract

Modern timing distribution systems benefit from high configurability and the bidirectional transfer of timing data. The Electron Ion Collider (EIC) Common Hardware Platform (CHP) will integrate the functions of the existing RHIC Real Time Data Link (RTDL), Event Link, and Beam Sync Link, along with the Low-Level RF (LLRF) system Update Link (UL), into a common high speed serial link. One EIC CHP carrier board supports up to eight external 8 Gbps high speed links via SFP+ modules, as well as up to six 8 Gbps high speed links to each of two daughterboards. A daughterboard will be designed for the purpose of timing data link distribution for use with the CHP. This daughterboard will have two high speed digital crosspoint switches and a Xilinx Artix Ultrascale+ FPGA onboard with GTY transceivers. One of these will be dedicated for a high-speed control and data link directly between the onboard FPGA and the carrier FPGA. The remaining GTY transceivers will be routed through the crosspoint switches. The daughterboard will support sixteen external SFP+ ports for timing distribution infrastructure with some ports dedicated for transmit only link fan-out. The timing data link will support bidirectional data transfer including sending data or events from a downstream device back upstream. This flexibility will be achieved by routing the SFP+ ports through the crosspoint switches which allows the timing link datapaths to be forwarded directly through the daughterboard to the carrier and into the FPGA on the daughterboard in many different configurations.

## INTRODUCTION

To support the EIC project, an upgraded timing distribution network is being developed to take the place of several existing distributed timing links [1][2][3]. The upgraded timing data link (TDL) will be the primary mechanism for connecting devices in the accelerator complex that require real time information about the machine. The TDL provides high speed low latency deterministic timing data distribution to connected devices. It is used to send data and status information for distributed feedback and machine protection systems. Notable improvements over previous implementations are more efficient transfer of data bidirectionally, and more configurable data filters. The TDL defines the network protocol and structure which will be realized by interconnected CHP systems. The CHP systems will support various pluggable special function

daughterboards, one of which is dedicated to TDL distribution. This daughtercard in conjunction with the CHP carrier will constitute the backbone of the TDL network infrastructure.

## TIMING DATA LINK

The TDL is a networking protocol which defines layers one through four of the OSI model [4], consisting of the physical, data link, network, and transport layers. These layers represent a standard conceptual segmentation of functionality, the actual implementation doesn't necessarily have such a clear separation of components. The TDL network allows devices connected in a tree structure to reliably communicate prioritized timing critical information bidirectionally with deterministic timing when necessary.

### Physical Layer

The TDL physical layer is responsible for the transfer of raw data bytes between directly connected systems. Data bytes of 8 bits are encoded into symbols of 10 bits using the 8b/10b encoding scheme. Symbols are then transferred serially one bit at a time, LSB first, at a rate of 8 Gbps. Accounting for the encoding overhead, a data throughput of 6.4 Gbps is achieved. Data bits are electrically represented as differential pair CML signals. A link has two AC coupled differential pairs, one each for transmitted and received data. The 8b/10b encoding scheme keeps track of running disparity and ensures sufficient transitions for clock recovery and DC balance on the line. Data byte alignment is achieved by detecting valid 8b/10b K.28.1 comma symbols.

The CML electrical signal pairs are converted to optical signals using an SFP+ module optical transceiver. Each optical transceiver supports two optical fiber connections, one for each direction of data transfer. The physical medium for transmission of the optical signals is single-mode or multi-mode optical fiber. The optical fiber will make the long runs necessary to connect systems separated by a large physical distance.

A key feature of the physical layer is clock recovery from the received data. The 8 Gbps data line rate is transmitted synchronous to the accelerator master oscillator running at 100 MHz. A clock and data recovery (CDR) unit inside the PHY provides a divided down clock of the line frequency to the receiver. This recovered clock is then used as a synchronous copy of the 100 MHz master oscillator clock.

### Data Link Layer

The TDL data link layer is responsible for transferring data words between directly connected systems. No

Hardware

FPGA & DAQ Hardware

\* Work supported by Contract Number DE-AC02-98CH10886 with the auspices of the US Department of Energy  
<sup>†</sup> pbachek@bnl.gov

# OVERVIEW AND OUTLOOK OF FPGA BASED HARDWARE SOLUTIONS FOR DATA SYNCHRONIZATION, ACQUISITION AND PROCESSING AT THE EuXFEL

Bruno Fernandes\*, Hamed Sotoudi Namin, Frank Babies, Tobias Freyermuth, Patrick Gessler, Irem Soekmen, European X-Ray Free-Electron Laser, Schenefeld, Germany

## Abstract

The European X-Ray Free Electron Laser facility (EuXFEL) provides ultra short coherent X-Ray flashes, spaced by 220 nanoseconds and with a duration of less than 100 femtoseconds, in bursts of up to 2700 pulses every 100 ms to several instruments. The facility has been using standardized Field-Programmable Gate Array (FPGA) based hardware platforms since the beginning of user operation in 2017. These are used for timing distribution, data processing from large 2D detectors, high speed digitizers for acquisition and processing of pulse signals, monitoring beam characteristics, and low latency communication protocol for pulse data vetoing and Machine Protection System (MPS).

Our experience grows in tandem with user requests for more specific and challenging case studies, leading to tailor made hardware algorithms and setups. In some cases, these can be fulfilled with the integration of new hardware, where collaboration with companies for new and/or updated platforms is a key factor, or taking advantage of unused features in current setups.

In this overview, we present the FPGA hardware based solutions used to fulfill EuXFEL's requirements. We also present our efforts in integrating new solutions and possible development directions, including Machine Learning (ML) research, with the aim of bringing more accurate results and configurable setups to user experiments and facilitate communications with other platforms used in the facility, namely Programmable Logic Controllers (PLC).

## INTRODUCTION

Located in Hamburg, Germany, the European X-Ray Free-Electron Laser facility (EuXFEL) provides ultra short coherent X-Ray flashes for scientific applications. These bunches have an energy between 260 eV and 24 keV and can be generated at a repetition rate up to 4.5 MHz for 600  $\mu$ s [1]. A new batch of pulses, referred to as a train, is generated every 100 ms and distributed among the seven instruments currently operating at EuXFEL. Additionally, synchronized optical laser pulses are available to all instruments for time-resolved pump-probe studies and laser-controlled manipulation of electronic relaxation and excitation processes. These instruments use the properties of the generated pulses for a wide range of scientific research topics.

Bandwidths of 10 GBytes of data per second are generated from the unique 2D pixel detectors available in the facility, while detector types, like systems based on fast digitizing and

analog-to-digital converters, can reach up to 70 MBytes. The latter detectors are used not just for scientific experiments, but also for diagnostics purposes. Besides the data volume, synchronization with the main accelerator is a key aspect in the setup of the experiments.

Considering the variety of experiments that take place at the EuXFEL, together with the ever-changing user requirements and high-speed data throughput, the facility requires a hardware framework which can cope with the expected data bandwidth while offering flexibility, scalability and processing features to reduce the amount of data to be store.

Field Programmable Gate Array (FPGA) technology provides the simple hardware solutions together with the ability to implement functions during the firmware programming phase. The technology is relevant in data processing, allowing for reduction and filtering to be done at the hardware level together with implementation of low-latency feedback systems. Algorithms can be continuously upgraded without having to change the hardware, and the code reused and easily integrate into different projects.

## HARDWARE PLATFORMS

The main hardware platform is based on MicroTCA.4, a modular, open standard, created and maintained by the PCI Industrial Computer Manufacturers Group (PICMG) [2]. It provides its compatible boards, refer to as Advanced Mezzanine Cards (AMCs), high-speed interconnects via PCIe, Ethernet and point-to-point, with redundant cooling and power supply as well as full remote monitoring, control and automatic failure detection and reaction. These are particularly important for the EuXFEL, due to the many decentralized locations where sensitive measurement signals are detected and timing signals are required. The MicroTCA.4 standard extends the AMC card definition with a Rear Transition Module (RTM), a second board physically connected to the AMC to extend the amount of I/O. Other relevant and standard features of the platform include the possibility to distribute signals and clocks between the AMCs via the backplane, i.e. without the need for external equipment, and a MicroTCA Carrier Hub (MCH) which is the central managing device of a MicroTCA crate.

Based on this platform, we use FPGA equipped AMCs which provide digitization, synchronization, online processing and realtime data rejection (vetoing) capabilities. Commercial and custom design RTMs extend the functionality of these integrated AMCs to address the instruments' setup requirements. CPU AMCs allow for remote configuration, monitor and collection of the data acquire by the other AMCs

\* bruno.fernandes@xfel.eu

# STATUS OF THE MICROTCA BASED BEAM INSTRUMENTATION DAQ SYSTEMS AT GSI AND FAIR

T. Hoffmann, H. Bräuning, R. Geißler, T. Milosic  
GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany

## Abstract

While the first FAIR accelerator buildings are soon to be completed, MicroTCA-based data acquisition systems for FAIR beam instrumentation are ready for use. By using commercial off-the-shelf components as well as open hardware with in-house expertise in FPGA programming, there are now DAQ solutions for almost all major detector systems in MicroTCA in operation at the existing GSI accelerators. Applications span a wide range of detector systems and hardware, often taking advantage of the high channel density and data transmission bandwidth available with MicroTCA. All DAQ systems are synchronised and triggered using a comprehensive White-Rabbit-based timing system. This allows correlation of the data from the distributed acquisition systems on a nanosecond scale.

In this paper, we present some examples of our DAQ implemented in MicroTCA covering the range of beam current, tune, position and profile measurements. While the latter uses GigE cameras in combination with scintillating screens, the other applications are based on ADCs with different sampling frequencies between 125 MSa/s up to 2.5 GSa/s or latching scalers with up to 10 MHz latching frequency.

## INTRODUCTION

Already at the beginning of the planning of the FAIR project in 2008, there were efforts in the GSI Beam Instrumentation department to replace the proven but old standard form factor VME. A new hardware basis was required which should provide the following features:

- Higher data bandwidth on backplane
- State-of-the-art bus systems like PCIe (3rd Gen. or higher)
- Modularity (including power supply and fan trays)
- Hot swap or hot plug support
- Availability of relevant components, such as ADC, TDC, counter, I/O or carrier boards for other standard mezzanine technologies, e.g. FMC
- Scalability
- Extensive possibilities for remote maintenance
- High availability

In addition, the new platform should be based on an industry standard and be available as commercial off-the-shelf (COTS) products on the open market, at best with a second source.

During this period, DESY had evaluated the Micro-Telecommunications Computing Architecture ( $\mu$ TCA or MicroTCA) form factor for data acquisition at the XFEL and extended the PICMG standard by MTCA.4 for Physics, thus procuring a viable alternative to VME. DESY has

significantly advanced the positive development and establishment of the standard through numerous MicroTCA workshops [1], tutorials and the founding of the MicroTCA Technology Lab. The resulting growth of the community led to a stabilization and further development of the standard. With the momentum gained, MicroTCA has become a reliable and powerful option.

Therefore, GSI/FAIR Beam Instrumentation department now uses MicroTCA as its multi-channel data acquisition (DAQ) platform, complemented by Industrial-PC solutions for readout with a small number of channels. It is worth mentioning that our industry partners have also developed, partly at their own financial risk, MicroTCA modules, so-called advanced mezzanine cards (AMC), for specific FAIR measurement tasks. These are now available to everyone on the market.

The FAIR project schedule allowed various beam diagnostic systems at the existing GSI accelerator facility to be upgraded with MicroTCA systems, including integration into the FESA-based control system [2]. After completion of the FAIR facility, stable and mature systems can be quickly installed by simple duplication.

## MAIN SYSTEM SETUP

In most applications, a space-saving 2U - 6 slot Schroff (nVent) MTCA.4 chassis (RackPak/M5-1) with standard backplane and 600W or 1000W (AC) power supply is used. In few cases, where a higher number of AMC modules is required Schroff (nVent) 9U - 12 slot crates (RackPak/M12-41) with 1000 W (AC) power supply can be used. Each system requires a MicroTCA carrier hub (MCH, here NAT-MCH-PHYS80), which acts as an interface for remote maintenance and system configuration/control, as well as an internal PCIe and LAN switch. The systems are equipped with a Concurrent Technologies AM-G64 AMC CPU providing Dual PCIe x4 support on AMC ports 4-7 (FatPipe1) and 8-11 (FatPipe2) and a Quad-Core Xeon E3-1505M processor. Both crate types are equipped with a JTAG Switch Module (JSM) slot for FPGA programming via the backplane.

### FAIR Timing Receiver Node

The control system design for FAIR is based on a new machine timing system, which was developed by GSI, CERN and other partners on top of the White-Rabbit Protocol [3]. So-called FAIR Timing Receiver Nodes (FTRN) were developed for the targeted standards MicroTCA (AMC=Advanced Mezzanine Card) and PCIe. The development was carried out by Cosylab as an in-kind contribution of Slovenia to FAIR with support of the GSI timing group. The entire product is published under the CERN



# A WORKFLOW FOR TRAINING AND DEPLOYING MACHINE LEARNING MODELS IN EPICS

M. Leputa\*, K. R. L. Baker, M. Romanovschi  
ISIS Neutron and Muon Source, Didcot, United Kingdom

## Abstract

The transition to EPICS as the control system for the ISIS Neutron and Muon Source accelerators is an opportunity to more easily integrate machine learning into operations. However, developing high-quality machine learning (ML) models is insufficient. Integration into critical operations requires good development practices to ensure stability and reliability during deployment and to allow robust and easy maintenance. For these reasons we implemented a workflow for training and deploying models that utilise off-the-shelf, industry-standard tools such as MLflow. Our experience of how adoption of these tools can make developer's lives easier during the training phase of a project is discussed. We describe how these tools may be used in an automated deployment pipeline to allow the ML model to interact with our EPICS ecosystem through Python-based IOCs within a containerised environment. This reduces the developer effort required to produce GUIs to interact with the models within the ISIS Main Control Room as tools familiar to operators, such as Phoebus, may be used.

## INTRODUCTION

The ISIS Neutron and Muon Source accelerators [1] control system is presently undergoing a transition from a Vsystem (Vista Control system [2]) to an EPICS (Experimental Physics and Industrial Control system [3]) based control system [4].

Alongside this migration, an expansion of data monitoring and collection systems has been implemented, incorporating many new software stacks; including but not limited to EPICS Achiever Appliance [5], and Influx-DB [6] as well as the related metrics collection and analysis suites (Telegraf, InfluxDB, Grafana, etc.) [7].

Moreover, a significant milestone has been achieved in the form of digitisation of the Analog Waveform System (AWS) [8, 9], which has provided digital waveforms of key accelerator systems; to date only stored as hourly images of screen-captures from oscilloscopes.

All of these enhancements along with an increase in the quality and volume of data, as well as developments in machine learning frameworks (such as TensorFlow [10] and PyTorch [11]) and commonly available powerful hardware accelerators, have enabled the ISIS controls group to leverage these advances and begin development of advanced control, optimisation and monitoring systems.

In this paper, we will discuss the tooling, workflow and two sample deployments of machine learning projects using EPICS. We will demonstrate how the deployments conform

to a user interface that the operators in the ISIS main control room (MCR) are already familiar with and how the deployments require minimal setup from the operators side. The advantages of this workflow are discussed as well as plans for further development.

## MACHINE LEARNING OPERATIONS

### Motivation

The ISIS controls group is a relatively small team with many overlapping responsibilities throughout the accelerator; as such it is an implicit requirement that any software and infrastructure developed by the team is easy to maintain, develop and hand over to other members. To that end, the team follows a set of best practices and principles in their software and infrastructure development; the core of which include, but are not limited to:

- **Modular Architecture** - The team designs software in a modular fashion, ensuring minimal coupling between different objects both at the code level and service level (when said code is deployed).
- **Version Control Systems (VCS)** - As is standard in most software development teams the team uses a version control system (in our case git [12] and Git-Lab [13]) to manage the code-base. This allows for tracking changes, better collaboration, and facilitates continuous knowledge transfer within the team.
- **Continuous Integration and Continuous Deployment (CI/CD)** - CI/CD pipelines built on top of the controls group VCS allows for easy automation of repetitive tasks, including but not limited to testing, deployment, and evaluation of the code-base.

These practices have been singled out due to their significant impact on productivity relative to the time invested in their implementation. Consequently, they were chosen as a baseline requirement for the target machine learning workflow that the team would adopt.

Machine learning projects also face additional challenges, these are best exemplified in Burkov's chapter on why machine learning projects fail [14]. With the most relevant and actionable within the current scope of the project being:

- **Siloed Organizations and Lack of Collaboration** - Lack of standardised workflows and pipelines leads to every project being a "one-off" collection of scripts and code that only the original developer can realistically deploy and maintain.

\* mateusz.leputa@stfc.ac.uk

# INTEGRATING SYSTEM KNOWLEDGE IN UNSUPERVISED ANOMALY DETECTION ALGORITHMS FOR SIMULATION-BASED FAILURE PREDICTION OF ELECTRONIC CIRCUITS

Felix Waldhauser<sup>1,2,\*</sup>, Hamza Boukabache<sup>1</sup>, Daniel Perrin<sup>1</sup>, Stefan Roesler<sup>1</sup>, Martin Dazer<sup>2</sup>

<sup>1</sup>CERN, Geneva, Switzerland

<sup>2</sup>Institute of Machine Components, University of Stuttgart, Stuttgart, Germany

## Abstract

Machine learning algorithms enable failure prediction of large-scale, distributed systems using historical time-series datasets. Although unsupervised learning algorithms represent a possibility to detect an evolving variety of anomalies, they do not provide links between detected data events and system failures. Additional system knowledge is required for machine learning algorithms to determine the nature of detected anomalies, which may represent either healthy system behavior or failure precursors. However, knowledge on failure behavior is expensive to obtain and might only be available upon pre-selection of anomalous system states using unsupervised algorithms. Moreover, system knowledge obtained from evaluation of system states needs to be appropriately provided to the algorithms to enable performance improvements. In this paper, we will present an approach to efficiently configure the integration of system knowledge into unsupervised anomaly detection algorithms for failure prediction. The methodology is based on simulations of failure modes of electronic circuits. Triggering system failures based on synthetically generated failure behaviors enables analysis of the detectability of failures and generation of different types of datasets containing system knowledge. In this way, the requirements for type and extend of system knowledge from different sources can be determined, and suitable algorithms allowing the integration of additional data can be identified.

## INTRODUCTION

Electronic circuits used for safety-critical equipment must be fail-safe systems and meet high reliability requirements by design. Fail-safe radiation monitoring devices used in particle accelerator environments trigger interlocks upon failure detection which cut the particle beam. While this behavior is essential in terms of safety, it also affects the availability of the accelerators and experiments. Continuous condition monitoring and failure prediction allows to perform predictive maintenance and thus to reduce accelerator downtime caused by unexpected failures. This study pursues a data-driven approach allowing online system-level failure prediction of distributed electronic systems with numerous instances. However, characteristics for fault detection at a system level are often complex, diverse and depend on the circuit layout. A common approach to this problem is to use unsupervised anomaly detection algorithms to identify

outlying data samples. In this case, anomalies are detected empirically as deviating from the majority of samples but without considering their relation to the system condition or failure states. Hence, the presented approach aims at integrating system knowledge into unsupervised anomaly detection algorithms to establish the link between detected anomalies and true failure states. By quantifying the increase in performance of anomaly detection algorithms, the potential of various types of system knowledge for failure prediction can be assessed. The benchmarking of the algorithms is based on datasets obtained from SPICE simulations of an electronic radiation monitoring device. This novel approach combines failure simulations of electronic circuits and anomaly detection algorithms to select the most suitable type and extend of system knowledge for a failure prediction use case. As a result, the need for resources to generate system knowledge can be minimized, and predictive maintenance algorithms can be targeted at an early stage to reduce unexpected failures and thus increase the availability of large-scale systems.

## FRAMEWORK FOR SIMULATION-BASED FAILURE PREDICTION

Selecting the most promising approach for integrating system knowledge into anomaly detection algorithms requires quantifying the associated change in classification performance. Therefore, a fully labeled dataset to benchmark model predictions is essential. However, in most predictive maintenance use cases, labeled faults are scarce as the system failure conditions can be diverse and evolve over time, making it difficult to obtain labels for data events, especially at the transition between normal and abnormal states [1]. Hence, the presented study is based on SPICE simulations of a demonstrator circuit allowing full control of the circuit state on a component-level through netlist modifications. Using simulations as data source provides labeled datasets of both the healthy and faulty system state, and allows to generate various types of system knowledge.

The presented approach combines existing methods for failure simulation and anomaly detection to form a novel methodology allowing efficient configuration of the integration of system knowledge into failure prediction algorithms. The methodology can be subdivided into four main steps, as shown in Fig. 1. Step I requires definition of the failure behavior of the circuit components to generate a benchmarking dataset using simulations of the healthy and faulty state. The detectability study (Step II) aims at evaluating the ef-

\* felix.johannes.waldhauser@cern.ch

# SYSTEMS MODELLING, AI/ML ALGORITHMS APPLIED TO CONTROL SYSTEMS

S. A. Mnisi<sup>†</sup>, South African Radio Astronomy Observatory, Cape Town, South Africa

## Abstract

The 64 receptor (with 20 more being built) radio telescope in the Karoo, South Africa, comprises of a large number of devices and components connected to the Control And Monitoring (CAM) system via the Karoo Array Telescope Communication Protocol (KATCP). KATCP is used extensively for internal communications between CAM components and other subsystems. A KATCP interface exposes requests and sensors; sampling strategies are set on sensors, ranging from several updates per second to infrequent on-change updates. The sensor samples are of different types, from small integers to text fields. The samples and associated timestamps are permanently stored and made available for scientists, engineers and operators to query and analyse. In this paper, I present how to apply Machine Learning tools which utilise data-driven algorithms and statistical models to analyse data sets and then draw inferences from identified patterns or make predictions based on them. The algorithms learn from the data as they run against it, as opposed to traditional rules-based analytical systems that follow explicit instructions. Since this involves data preprocessing, I will touch on how the MeerKAT telescope data storage infrastructure (Katstore) manages the voluminous variety, velocity and volume of this data.

## OVERVIEW OF SENSOR DATA STORAGE

Before delving into how Machine Learning (ML) tools can be applied to the MeerKAT Control And Monitoring (CAM), it is important to first give an overview of the storage infrastructure, Katstore[1]. CAM Software components send data points(samples) to Katstore to save and make available for analysis. Katstore stores the values, status and other information about sensors in the CAM system[2]. These samples received by Katstore are keyed on time and sensor name. This makes Katstore a time series database (TSDB), purposely built to have a fixed index on time. The data in Katstore is immutable(no update on a sample is allowed) and only grows over time. It can be seen as an append-only database and samples do not need to arrive in chronological order.

The samples are packed as JavaScript Object Notation (JSON)[3] objects by the software components that collected the samples. Any valid JSON is accepted, with Katstore only requiring that each sample contains the keys name and time, where name is the sensor name and time is the time in Coordinated Universal Time (UTC).

Making each sample a document removes the need for application knowledge in Katstore and future-proofs the implementation. New fields can be added and removed without requiring changes to Katstore and there is no fixed schema for a sensor sample. The software components that collect the samples publish the samples at intervals that can be configured per sensor. Katstore subscribes to the per sensor archive subject on the message bus and stores the published samples.

A sensor is a fundamental concept in KATCP and a collection of sensor types are available. The following types are currently supported: integer, float, boolean, timestamp, discrete, address and string. Sensors always have a status and the following statuses are supported: unknown, nominal, warn, error, failure, unreachable and inactive. In KATCP sensor sampling is performed by the server based on a sampling strategy provided by the client, this allow every connection to set up a unique sampling strategy.

```
{  
  "name": "m000_rsc_rxl_cryostat_pressure",  
  "time": 1505982067.202219,  
  "value": 1013.25,  
  "status": "nominal",  
  "value_ts": 1505977839.44  
}
```

Example 1: Sensor data sample.

There are several sampling strategies available ranging from a fixed time interval (period) to on value change (event). The sensor data, Example 1, that is collected and stored in Katstore is enormous and continues to grow over time.

MeerKAT CAM has many software components, some components connect to hardware devices and others connect to software components. All inter-component communication is done with Karoo Array Telescope Communication Protocol (KATCP). Components can call requests on connected components for control purposes. A KATCP request is analogous to method or command calls of other platforms. For monitoring purposes, KATCP provides the concept of sensors. For the purpose of archiving, the components that make up the MeerKAT CAM system publish sensor samples to different subjects on the message bus. The publish rate is controlled by the system configuration. Katstore subscribes to the archive subjects and stores the samples to the buffer. For the samples to be stored, all the MeerKAT CAM software

\* smnisi@sarao.ac.za

# LASER FOCAL POSITION CORRECTION USING FPGA-BASED ML MODELS

J. Einstein-Curtis\*, S. J. Coleman, N. M. Cook, J. P. Edelen, RadiaSoft LLC, Boulder, CO, USA  
S. Barber, C. Berger, J. van Tilborg, Lawrence Berkeley National Lab, Berkeley, CA, USA

## Abstract

High repetition-rate, ultrafast laser systems play a critical role in a host of modern scientific and industrial applications. We present a diagnostic and correction scheme for controlling and determining laser focal position by utilizing fast wavefront sensor measurements from multiple positions to train a focal position predictor. This predictor and additional control algorithms have been integrated into a unified control interface and FPGA-based controller on beamlines at the Bella facility at LBNL. An optics section is adjusted online to provide the desired correction to the focal position on millisecond timescales by determining corrections for an actuator in a telescope section along the beamline. Our initial proof-of-principle demonstrations leveraged pre-compiled data and pre-trained networks operating ex-situ from the laser system. A framework for generating a low-level hardware description of ML-based correction algorithms on FPGA hardware was coupled directly to the beamline using the AMD Xilinx Vitis AI toolchain in conjunction with deployment scripts. Lastly, we consider the use of remote computing resources, such as the Sirepo scientific framework\*, to actively update these correction schemes and deploy models to a production environment.

## INTRODUCTION

Laser plasma accelerators (LPAs) rely upon accurate control of ultrafast lasers, typically Ti:Sapph and Nd:Yag amplifier systems [1]. The BELLA Center at Lawrence Berkeley National Laboratory (LBNL) features several ultra-short pulse, high-energy beamlines to develop LPAs. These accelerators require highly repeatable, stable interaction points to generate high-quality electron beams, which necessitates a collection of active and passive controls to mitigate environmental, mechanical, and component variations.

Recent work has primarily focused on enhancing transverse beam stability [2]. This paper describes a strategy to address focal position stability, leveraging a machine learning (ML) enhanced wavefront diagnostic in tandem with a Field Programmable Gate Array (FPGA) controller to correct focal position at a kHz-scale rate. By building a model of wavefront at the interaction point, it is possible to use a non-perturbative measurement to calculate the focal position.

\* joshhec@radiasoft.net

Table 1: Optimal lens movement vs focal shift and beam size change. Focus shift is per mm lens translation. Beam size change is change per mm lens translation.

	Shift	Size Change
Transmissive Amp3-in	2 mm	x1.348
Transmissive Amp4-in	2 mm	x1.046
Reflective Amp4-out	1 mm	x1.002

## FACILITY AND EQUIPMENT

The initial model was created for the BELLA HTU laser system, shown in Fig. 1. This beamline operates with 1 kHz seed pulses and a 1 Hz full-power pulse. A HASO FIRST Shack-Hartmann wavefront sensor was used as the ground-truth imaging device of the interaction and post-interaction region, with the pre-interaction region sensor a Thorlabs WFS20-7AR. A Xilinx Zynq ZCU104 FPGA evaluation kit was used for testing to provide flexibility during the prototype phase, including a variety of customizable I/O, well-supported manufacturer-provided software, and a variety of processing options in support of ML operations.

## FOCAL POSITION INVESTIGATIONS

To determine the optimal lenses to move for a focal shift, we looked at the magnitude of the shift at final focus and the (unwanted) increase in beam size throughout the optical chain. Table 1 summarizes these parameters for three different lenses in the telescope.

From these simulations we determined that the reflective Amp4-out is not ideal as a motorized correction optic for focal location because it is more weakly responsive, shifting the focus by only 1 mm per mm translation. Moreover, the off-axis reflective geometry introduces beam centroid kicks, even in response to relatively mild beam size variations. Ultimately, we determined the Amp4-in telescope is the best choice.

To verify our model, we measured the focal location vs lens separation at high power. Our measurement used a comparable method of capturing leakage from the final steering mirror thus measuring raw focal location without the need for further calibration or renormalization. The inset of Fig. 2 provides details of the measured focal position and radius of curvature taken from the wavefront sensor.

When comparing measurements to the simulation, we note that the focus shift per mm stage motion depends on the nominal Amp4-in lens separation. For a perfectly collimated beam entering the  $-f_1/+f_2$  telescope, and for a perfectly collimated beam leaving the telescope (lens separation is

# MODEL DRIVEN RECONFIGURATION OF LANSCE TUNING METHODS\*

C.E. Taylor, P. Anisimov, S.A. Baily, E.C. Huang, L. Rybarczyk, H.R. Leffler, A. Scheinker,  
H.A. Watkins, E.E. Westbrook, D.D. Zimmermann  
Los Alamos National Laboratory, Los Alamos, USA

## Abstract

This work presents a review of the shift in tuning methods employed at the Los Alamos Neutron Science Center (LANSCE). We explore the tuning categories and methods employed in four key sections of the accelerator, namely the Low-Energy Beam Transport (LEBT), the Drift Tube Linac (DTL), the side-Coupled Cavity Linac (CCL), and the High-Energy Beam Transport (HEBT). The study additionally presents the findings of employing novel software tools and algorithms to enhance each domain's beam quality and performance. This study showcases the efficacy of integrating model-driven and model-independent tuning techniques, along with acceptance and adaptive tuning strategies, to enhance the optimization of beam delivery to experimental facilities. The research additionally addresses the prospective strategies for augmenting the control system and diagnostics of LANSCE.

## INTRODUCTION

The Los Alamos Neutron Science Center is a renowned scientific establishment located at the Los Alamos National Laboratory (LANL) in New Mexico. This proton linear accelerator is globally recognized for its ability to accelerate protons up to 800 MeV at high power. LANSCE effectively sustains a dynamic program focused on fundamental scientific research by offering the scientific community with high-intensity sources of neutrons and protons. These sources are utilized for conducting experiments that contribute to both government and civilian research endeavors [1] and for the production of isotopes used in medical and research applications [2]. The scattering science research employs a high-powered proton and spallation neutron source with short-pulse characteristics, operating at a capacity of 100 kilowatts. These studies are conducted at multiple beamlines, enabling concurrent research across various topics. The experimental setup comprises the Coherent CAPTAIN-Mills (CCM) detector, a 10-ton liquid argon detector positioned at 20 meters from the high-intensity neutron/neutrino source. Its primary objective is to investigate the existence of sterile neutrinos and light-dark matter [3].

The control system of LANSCE uses the Experimental Physics and Industrial Control System (EPICS), developed at LANL [4]. This infrastructure is currently being used many accelerator facilities for performing data acquisition,

\* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). † cetaylor@lanl.gov, LA-UR-2023-##

supervisory control, closed-loop control, sequential control, and operational optimization. The EPICS architecture was developed through a collaborative effort between experts in both physics and industrial control. The LINAC encompasses five state-of-the-art research facilities that operate simultaneously: Lujan Center, Weapons Neutron Research, Proton Radiography, Isotope Production Facility, and Ultracold Neutrons. The intricate control system employed in LANSCE facilitates the concurrent execution of numerous experiments, making it a versatile instrument for enhancing scientific research. The control system of LANSCE is regarded as an impressive engineering achievement, as it operates its LINAC control system using technology that has been in use for almost three decades [5]. Over the course of time, there have been many modifications to peripheral components. However, it is anticipated that significant enhancements will be made to the control system in the future, with the aim of achieving optimal beam delivery to the experimental facilities.

## THE LANSCE ACCELERATOR CONFIGURATION

Figure 1 shows the four basic areas to tune in the LANSCE accelerator, each with their own set of diagnostics and control systems. Two species of beam ( $H^+$  and  $H^-$ ) are generated in the 1) *Low-Energy Beam Transport* (LEBT) at 750 keV and merged into the same beamline before entry into the 2) *Drift Tube Linac* (DTL). This first accelerator increases the beam energy of the beam up to 100 MeV. From there, the beam is transported through the transition region (TR) into the 3) *side-Coupled Cavity Linac* (CCL) for final acceleration up to 800 MeV. The beam is then transported to the experimental facilities through the 4) *High-Energy Beam Transport* (HEBT). The methods of control and tuning vary significantly within each region. To better understand the requirements of these areas, we will try to break the tuning processes down into two categories and two methods. The *categories* indicates if the tuning uses a model or is dependent solely on parameter minimization. The *method* describes if the control lies with the input parameters or the control elements (steering, focusing, etc.).

## TUNING CATEGORIES

In the conventional practice of operating accelerator facilities, the establishment of beam transport and acceleration is often initiated using a physics model. Diagnostics, such as emittance stations, beam position monitors, and wire scanners, are employed to quantify the beams'

# DISENTANGLING BEAM LOSSES IN THE FERMILAB MAIN INJECTOR ENCLOSURE USING REAL-TIME EDGE AI

K.J. Hazelwood, M.R. Austin, J.M. Arnold, J.R. Berlioz, P. Hanlet, M.A. Ibrahim, J. Mitrevski, V.P. Nagaslaev, D.J. Nicklaus, G. Pradhan, A.L. Saewert, B.A. Schupbach, K. Seiya, R.M. Thurman-Keup, N.V. Tran, A. Narayanan<sup>1</sup>  
Fermi National Accelerator Laboratory\*, Batavia, IL USA  
<sup>1</sup>also at Northern Illinois University, DeKalb, IL USA  
J.YC. Hu, C. Xu, J. Jiang, H. Liu, S. Memik, R. Shi, A.M. Shuping, M. Thieme  
Northwestern University<sup>†</sup>, Evanston, IL USA

## Abstract

The Fermilab Main Injector enclosure houses two accelerators, the Main Injector and Recycler Ring. During normal operation, high intensity proton beams exist simultaneously in both. The two accelerators share the same beam loss monitors (BLM) and monitoring system. Deciphering the origin of any of the 260 BLM readings is often difficult. The Accelerator Real-time Edge AI for Distributed Systems project, or READS, has developed an AI/ML model, and implemented it on fast FPGA hardware, that disentangles mixed beam losses and attributes probabilities to each BLM as to which machine(s) the loss originated from in real-time. The model inferences are then streamed to the Fermilab accelerator controls network (ACNET) where they are available for operators and experts alike to aid in tuning the machines.

## PROJECT OVERVIEW

The Accelerator Real-time Edge AI for Distributed Systems (READS) project is a collaboration between the Fermilab Accelerator Directorate and Northwestern University. It aims to implement ML models on edge hardware for use on the Fermilab accelerator complex. The project consists of two sub-projects; improving Delivery Ring resonant extraction regulation [2–5] for the future Mu2e experiment [6] and aiding in the machine attribution of beam loss in the Main Injector enclosure [7].

## Disentangling Beam Losses

The Fermilab Main Injector enclosure houses two accelerators; the Main Injector (MI) and the Recycler Ring (RR) (Fig. 1). The 8 GeV permanent magnet Recycler Ring acts as a proton stacker for the 120 GeV synchrotron Main Injector. To ensure the most protons are delivered to Fermilab's experiments, the Recycler Ring is loaded with Main Injector's next pulse of beam while the current MI pulse is accelerated and then extracted. During normal operations, there are high intensity proton beams in both Recycler Ring and Main Injector [8]. The two machines share the same beam loss monitors (BLM) and monitoring system. When beam losses

occur, it can be difficult to attribute the origin of the loss to either machine resulting in delays tuning the machines and unnecessary downtime. However, machine experts are often able to decipher loss origin from the time in the cycle of the loss, the current machine states, local and global loss patterns and tunnel residual dose rate surveys (Fig. 2) [9, 10]. This suggests that given enough information, a ML model can be created to replicate, automate and perhaps improve upon the machine experts ability to attribute beam loss to the correct machine.



Figure 1: The Main Injector enclosure. Main Injector (bottom), Recycler Ring (top), P1 extraction beamline (middle).

## DATASETS

For this project, the ML models were trained using Supervised Learning. The training data consists of readings from all 260 BLMs around the MI enclosure, machine readings such as Main Injector and Recycler Ring beam intensities, machine state, Main Injector dipole bus ramp current, and clock events.

## Beam Loss Monitor Location Recording

An assumption made at the beginning of the project was that any ML model created to attribute losses would be highly dependent on the placement and location of each BLM. While most BLMs are securely affixed to the machine, a good amount of BLMs have been attached to moveable fixtures and experts from time to time have moved these BLMs as they see fit to try and characterize problematic losses

\* Operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy. Additional funding provided by Grant Award No. LAB 20-2261 [1]

<sup>†</sup> Performed at Northwestern with support from the Departments of Computer Science and Electrical and Computer Engineering

## DATABASE'S DISASTER RECOVERY MEETS A RANSOMWARE ATTACK

M. Zambrano<sup>†</sup>, SKA Observatory  
Jodrell Bank, Macclesfield, UK

V. Gonzalez, ALMA Observatory, Joint ALMA Office, Santiago, Chile

### Abstract

Cyberattacks are a growing threat to organizations around the world, including observatories. These incidents can cause significant disruption to operations and can be costly to recover from. This paper provides an overview of the history of cyberattacks, the motivations of attackers, and the organization of cybercrime groups. It also discusses the steps that can be taken to quickly restore a key component of any organization, the database, and the lessons learned during the recovery process.

The paper concludes by identifying some areas for improvement in cybersecurity, such as the need for better training for employees, more secure networks, and more robust data backup and recovery procedures.

### INTRODUCTION

Hacking is now a multimillion-dollar industry where motivated people are looking for their next victims. Hospitals, schools, councils and observatories have also been targets of those groups. This threat will not disappear soon, and each organization should be prepared to handle a never-ending list of attacks. The average cost of a ransomware attack in the USA in 2023 was 4.45 million USD [1]. The median time to contain an attack is 73 days but this could be increased by 14 days when the disclosure is made by the attacker [1].

### A VERY SHORT STORY ON HACKING

Spring 1962 at MIT, Allan Scherr needed more than four hours to run his PhD simulation models of time-sharing systems [2]. He asked the system to print the password file on a punched card and he was able to steal credentials for the first time in history.

Sometime later in 1975 Kevin Mitnick started his hacking life by "traveling free" when he asked the driver of his bus where he bought the funny punch machine they use for a "school project". He continued his life by learning everything he could about the phone system. A friend gave him the phone number of a system used to develop an OS. Later he was searched by the FBI for intrusions at IBM, Sun and Motorola and four counts of wire fraud, two counts of computer fraud and one count of illegally intercepting a wire communication. He spent some 5 years in prison and returned as a "white hat" security professional consultant. He was the first to use "social engineering" to get what he wanted from people, and he continued to do so as a security consultant.

<sup>†</sup> mauricio.zambrano@skao.int

Hacking at this point started to be a for profit business with a lot of curious, talented young people with no guilt or remorse of doing what they were doing.

One of the main concepts that Mitnick exploited was the human factor. You can have the best security system in the world but if people are getting exploited, then you will be able to circumvent that system.

A new subculture emerged where breaking into other people's computers wasn't seen as bad. There was also a huge economic benefit of doing this and the perpetrators have no empathy for the real damage caused to real people. There is also an important cost of being caught by the authorities. If you "succeed" in this business, you might end up being prosecuted by the FBI. They now have a list of 119 individuals known as the Cyber's most wanted. Currently most of them are from Russia, North Korea, Iran and China and linked to their states and armies in some cases. A few of them have been captured when traveling abroad. Since then, they usually stay in their own countries. Some of the hacker activities also have had an important political impact for the targeted country.

The United Nations has been trying to define responsible cyber behavior for states with an Open-Ended Working Group negotiation but the talks, so far, have not reached an agreement. Russia has also submitted its vision for a Convention of the UN on ensuring International Information Security. Hence there is currently no change in the horizon for this kind of activity.

### BRAND NEW WORLD?

Hacking groups have existed since 1980. One of the first of them was named the 414s and was composed of "young, male, intelligent, highly motivated and energetic" from Milwaukee, Wisconsin. They met each other at an IBM youth program. By using information from the DEC manual to get the default password, they were able to breach systems at Los Alamos National Laboratories, a bank and deleted bill records of a hospital. Since then, groups have increased in complexity.

In each group there are specialized actors with different abilities with the final objective to get money from the victim.

There is also a trend to offer hacking as a service. In that case all the people involved share a portion of the ransom.

Most attacks start:

- With an innocent click on a phishing mail.
- Stolen credentials.
- Via an exploit on an exposed server. There is a lucrative market for zero-day exploits for all the different platforms.

# PROTECTION LAYER DESIGN FOR THE HIGH LUMINOSITY LHC FULL REMOTE ALIGNMENT SYSTEM

B. Fernández\*, A. Germinario, E. Blanco, M. Sosin, H. Mainaud  
CERN, Geneva, Switzerland

## Abstract

The Full Remote Alignment System (FRAS) is a complex measurement, alignment and control system designed to remotely align components of the Large Hadron Collider (LHC) following its High Luminosity upgrade. The purpose of FRAS is to guarantee optimal alignment of the strong focusing magnets and associated components near the experimental interaction points, while at the same time limiting the radiation dose to which surveyors in the LHC tunnel are subjected.

A failure in the FRAS control system, or an operator mistake, could provoke a non desired displacement of a component that could lead to damage of neighbouring equipment. Such an incident would incur a considerable repair cost both in terms of money and time.

To mitigate this possibility, an exhaustive risk analysis of FRAS has been performed, with the design of protection layers according to the IEC 61511 standard proposed.

This paper presents the different functional safety techniques applied to FRAS, reports on the current project status, and introduces the future activities to complete the safety life cycle.

## INTRODUCTION

The High-Luminosity Large Hadron Collider (HL-LHC) [2] is an ambitious project to upgrade the LHC and increase its discovery potential. During the next long shutdown, between 2026 and 2028, many LHC components will be upgraded and some completely new systems will be deployed. These upgrades aim to improve the overall LHC performance and increase the total number of particle collisions produced by a factor of 10.

One of the new systems that will be deployed during this long shutdown is the Full Remote Alignment System (FRAS) [1]. FRAS will allow to remotely align components in both sides of the Interaction Points (IP) 1 and 5 of the LHC. The reduction in the mechanical components misalignment will decrease the required orbit corrector strengths, improving the accelerator performance and will allow to reduce the radiation doses for surveyors working in the tunnel.

However, these benefits come with a risk. An excessive misalignment of more than  $\pm 2.5$  mm in the vertical and horizontal axes or 1 mrad in the rotational axis between two LHC components could damage the interconnecting bellows. This would provoke a downtime of the LHC between several months and one year for reparations. In addition, there are many potential failures that can cause this bellow damage

as FRAS is a complex control system with many hardware and software components and operator interactions.

To mitigate the risk to an acceptable level, two primary actions have been undertaken. Initially, a comprehensive risk analysis and assessment were conducted to identify the combinations of failures resulting in a possible bellow damage and ascertain the required risk reduction measures. Secondly, a number of protection layers were designed in alignment with functional safety standards, aiming to bring the risk down to tolerable level.

The paper is structured as follows: Section describes the FRAS controls architecture. Section shows the risk analysis and assessment methods that have been applied. Section presents the design and analysis of the protection layers. And finally some conclusions and future work are outlined.

## FRAS

Remote alignment of the LHC components (e.g. collimators, quadrupoles, dipoles, etc.) requires to equip them with high precision sensors that allow to determine the 3D position of each component and to compute their necessary displacement for an optimal alignment. The FRAS will enable remote positioning of 68 accelerator components, whose are installed across two Long Straight Sections (LSS), spanning a distance of 400 meters each. For such a big and complex installation, whose main role is to monitor and displace accelerator components, often weighing tenths of tons, the primary focus is ensuring its safe operation. This can be achieved thanks to a network of over 450 micrometric sensors, and a set of controllers and stepper motors that constitute the FRAS control system. Figure 1 depicts the controls architecture the FRAS.

This schematic shows 2 of the 17 LHC components controlled by the FRAS on each side of the IP. In this case, both components are equipped with Wire Positioning Sensors (WPS) based on a capacitive technology and 2 types of inclinometers, one based on Frequency Scanning Interferometry (FSI) and the other on capacitive technologies [1]. FRAS can also read the position of the 5 motorized actuators that are in charge of moving the jacks or Universal Adjustment Platform (UAP) that supports FRAS components. This is done by reading the resolvers that provide an absolute position of the motorized actuator assembly.

The control layer of FRAS is comprised of a combination of Commercial Off-The-Shelf (COTS) and in-house hardware devices that read each sensor, compute the 3D position of each component and provide the optimal movement commands to be transmitted to the stepper motors. Some of the COTS devices are the so-called FECs (Front End

\* borja.fernandez.adiego@cern.ch



# ACCELERATOR SYSTEMS CYBER SECURITY ACTIVITIES AT SLAC\*

G. White<sup>†</sup>, A. Edelen, SLAC National Accelerator Laboratory, California, U.S.A

## Abstract

We present the idea that the needs of accelerator beam tuning by the emerging methods of Machine Learning and multi-particle modeling, will disrupt the norms of computer networking and cyber-security architectures of large accelerator control systems. We review present SLAC activities in solving beam problems using techniques such as Bayesian optimization, surrogate and inverse Neural Network model inferencing. These are frequently trained on multi-particle simulations, the training itself is computationally expensive, and increasingly on very large stored datasets of beam synchronous observables' past values. High Performance Computing (HPC) and Big Data will therefore take a central role in the accelerator control system. SLAC is building a "digital twin" framework in which to run and manage these models in the HPC cluster. Increasingly, the results of the modeling in HPC, will be deployed into the running accelerator, implying that AIs, outside the classical secure control network, can and will deploy setpoints autonomously. Finally, we review new controls protocol architecture and technology being developed at SLAC, in advance of these realities to come.

## HISTORICAL CONTEXT AND CONTROLS ARCHITECTURE

Historically, low order models such as transfer matrix and Courant-Snyder parameters, have been central to "online" beam optimization, being the basis of beam orbit correction, bumps, and basic feedback. That is, beam tuning as carried out on the running accelerator controls, has been in opposition to offline lattice design and beam dynamics study, which have classically been done offline, using model methods able to investigate beam phase space, but that require High Performance Computing (HPC) and runtime periods of hours or days. Furthermore, our use of these methods for online tuning, has come with some assumptions that have become coded into norms. First, that we know *a priori* the basic the relation between actuators and sensors – and we approximate it largely linearly (plus some 2nd order). Second, that for global optimization, minimizing the orbit RMS, or timing, will optimize the true objective – minimize emittance or maximize luminosity. Third, that direct tuning 6D phase space, or beam time structure, was out of reach for the turn around time of accelerator operations.

Over the last years, model methods have evolved. Directly tuning the injector requires modeling space charge. To understand true linac optics, RF kicks, magnet errors, and dynamic initial conditions, must be included. These imply

multi-particle codes and Machine Learning brought to bear together online. ML is used to compute solutions under uncertainty, or to learn the dynamics for fast online execution, or simply to give empirical insight where the available physics or simulation have not accessed the parameter space with enough precision.

## MODELING ACTIVITIES AT SLAC

Many tuning problems at LCLS/LCLS-II and FACET-II at SLAC and elsewhere, now require detailed phase space customization for different experiments. The beam exists in 6-D position-momentum phase space. We measure 2-D projections and reconstruct based on perturbations of upstream controls (e.g. tomography, quad scans). However, we have incomplete empirical information and we have dozens-to-hundreds of controllable variables and hundreds-of-thousands (up to millions for LCLS-II). Beam optimization then, is a nonlinear, high-dimensional problem. We also have a wide variety of tuning needs, such as rapid FEL beam pulse energy spectra construction, beam parameters such as two-beam, or pulse-probe, and maintaining time-energy stability. To these problems we bring a collection of approaches; from model-free estimation like gradient Decent, to model guided optimization and Physics Informed Neural Networks, and inverse models for feed-forward corrections. Our strategy is to start with sample-efficient methods that do well on new systems, then build up to more data-intensive and heavily model-informed approaches.

The long term requirement then is for fast, accurate, system models of the beam and experiment dynamics. Accelerator simulations that include nonlinear and collective effects are powerful tools, but they can be computationally expensive. To some extent, ML models are able to provide fast approximations to simulations (the so-called "surrogate models") so in-situ optimization is orders of magnitude faster than problem error minimization by multi-particle code iteration. However, the surrogate models must themselves be trained on simulations, which must be computed, or on large, typically long baseline, data sets - which take a lot of storage space.

## DIGITAL TWIN FRAMEWORK

Tuning then depends on High Performance Computing, either for running the multi-particle codes or for training the ML. For cost reasons, HPC is typically a shared resource outside the production control system, so our codes and models often run "offline" in HPC facilities such as NERSC and, recently, the new Stanford Research Computing Facility (SRCF).

As the first, large step then, we're developing a "Digital Twin" framework for accelerator modelling, which specifically includes modeling and optimizing the extant accelera-

\* Work supported by US Dept of Energy DE-AC02-76SF00515, and grant from Executive Order 14028

<sup>†</sup> greg@slac.stanford.edu

# VERIFICATION AND VALIDATION OF THE ESS MACHINE PROTECTION SYSTEM OF SYSTEMS (MP-SoS)

A. Nordt, M. Carroll, J. Gustafsson, G. Ljungquist, S. Gabourin, S. Pavinato, S. Kövecses de Carvalho, A. Petrushenko, European Spallation Source, ERIC, Lund, Sweden

## Abstract

The European Spallation Source, ERIC (ESS) is a source of spallation neutrons used for neutron scattering experiments, complementary to synchrotron light sources. ESS has very ambitious goals and experimentation with neutrons at ESS should be one or two orders of magnitude more performing compared to other sources. Each proton beam pulse generated by the linear accelerator will have a peak power of 125 MW. The machine's equipment must be protected from damage due to beam losses, as such losses could lead to melting of e.g., the beam pipe within less than 5  $\mu$ s. System-of-Systems engineering has been applied to deploy systematic and robust protection of the ESS machine. The ESS Machine Protection System of Systems (MP-SoS) consists of large-scale distributed systems, of which components themselves are complex systems. Testing, verification and validation of the MP-SoS is rather challenging as each constituent system of the MP-SoS has its own management, functionality that is not necessarily designed for protection, and also the different system owners follow their own verification strategies. In this paper, we will present our experience gained through the first 3 beam commissioning phases, ESS has gone through so far. We will describe how we managed to declare MP-SoS to being ready for beam operation without complexifying the task, and we will present the challenges, issues, and lessons learned faced during the verification and validation campaigns.

## SYSTEM OF SYSTEMS ENGINEERING APPROACH FOR ESS MACHINE PROTECTION

Modern particle accelerator facilities, realised by complex constellations of interacting systems, serve a variety of users as research enablers. While the constituent systems exhibit a significant degree of technical and operational independence and distinct life cycles, the performance required to conduct research still needs to emerge from their integration into one overall system, the research facility. This renders a Systems of Systems oriented approach to engineering useful. Furthermore, accelerator-based research facilities face increasing availability expectations. Achieving those expectations can be supported through a tailored application of functional safety standards as engineering methodology guideline on an SoS level, as explained in [1]. An SoS-Engineering approach utilising functional safety standards (IEC 61511, IEC 61508) in this way is concretised in the Machine Protection Systems of Systems at ESS.

## ORGANISATION AND RESPONSIBILITIES

The ESS Machine Protection team is responsible to:

- Coordinate Machine Protection across ESS.
- Define global protection functions.
- Develop, operate, and maintain the Beam Interlock System (BIS).
- Ensure working interfaces with BIS.
- Foster awareness that things can break.
- Foster awareness that thorough testing leads to success.

The ESS System Owners are responsible to:

- Develop reliable systems.
- Implement local protection functions.
- Implement Machine Protection requirements in their system.
- Provide sensors needed for global protection.

## PROTECTION FUNCTIONS

The requirements for the different protection functions (PF) are derived from analysis of the different systems that contribute to Machine Protection at ESS. Once the tolerable risk has been set and the necessary risk reduction estimated, the protection integrity requirements for the PF can be allocated in terms of Probability of Failure on demand (PFD) or Probability of Failure per Hour (PFH). The PFD and PFH correspond to one of the Protection Integrity Levels (PIL) specified in Table 1. Table 2 and 3 include the required Safe Failure Fraction (SFF) and the Hardware Fault Tolerance (HFT) that is required for each PIL.

Table 1: PIL Specified PFD and PFH

PIL	PFH (h <sup>-1</sup> )	PFD	MTBO (kh)
0	$\geq 10^{-5}$ to $<10^{-4}$	$\geq 10^{-1}$ to $\leq 0,5$	10-100
1	$\geq 10^{-6}$ to $<10^{-5}$	$\geq 10^{-2}$ to $<10^{-1}$	100-1000
2	$\geq 10^{-7}$ to $<10^{-6}$	$\geq 10^{-3}$ to $<10^{-2}$	$10^3$ - $10^4$
3	$\geq 10^{-8}$ to $<10^{-7}$	$\geq 10^{-4}$ to $<10^{-3}$	$10^4$ - $10^5$
4	$\geq 10^{-9}$ to $<10^{-8}$	$\geq 10^{-5}$ to $<10^{-4}$	$10^5$ - $10^6$

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# THE HYBRID IDENTITY OF A CONTROL SYSTEM ORGANIZATION: BALANCING SUPPORT, PRODUCT, AND R&D EXPECTATIONS

S. Baymani<sup>†</sup>, Paul Scherrer Institute, Villigen, Switzerland

## Abstract

Controls organizations are often expected to fulfil a dual role as both a support organization and an R&D organization, providing advanced and innovative services. This creates a tension between the need to provide services and the desire and necessity to develop cutting-edge technology.

In addition, Controls organizations must balance the competing demands of product development, maintenance and operations, and innovation and R&D. These conflicting expectations can lead to neglect of long-term strategic issues and create imbalances within the organization, such as technical debt and lack of innovation.

This paper will explore the challenges of navigating these conflicting expectations and the common traps, risks, and consequences of imbalances. Drawing on our experience at PSI, we will discuss specific examples of conflicts and their consequences.

We will also propose solutions to overcome or improve these conflicts and identify a long-term, sustainable approach for a hybrid organization such as Controls. Our proposals will cover strategies for balancing support and product development, improving communication, and enabling a culture of innovation.

Our goal is to spark a broader discussion around the identity and role of control system organizations within large laboratory organizations, and to provide concrete proposals for organizations looking to balance competing demands and build a sustainable approach to control systems and services.

## INTRODUCTION

Accelerator facilities are complex in the sense that they are at the same time cutting edge technology (or were at the time of conception), while being production grade services, with the operational expectations that come from that. In many ways, particle accelerators at research labs are the ultimate “prototype gone production” system.

In such settings, any middleware becomes key: it makes the glue that keeps the parts together, and therefore directly affects the operation, fine-tuning, and usage of the system as well as the output of its users.

Controls organizations of labs and institutions across the world have somewhat different constitutions: some include PLCs and safety groups and some not. Some include central IT responsibilities and some not. Some include beam-line support and some not. Some include data analysis and scientific software, and some not. And so on.

Furthermore, the emphasis and focus on work changes over time: in the youth of each facility, the weight is put on development and slowly shifts towards maintenance.

Likewise, at the conception of a facility, the team is small but slowly grows over time, making room for more experiments, again changing the composure and dynamics of the organization.

Regardless of constitution and age, a common pattern exists: Controls organizations have a very wide range of responsibilities and expectations - both explicit and implicit.

It is relevant to approach this observation with an analysis: Which are the expectations? Which ones are explicit, and which are implicit? What actual roles are they connected to?

Using the Paul Scherrer Institute as a case study, this paper presents a novel analytical framework that frames some key roles and tensions experienced by controls organizations. From our observations, this analysis contributes to similar debates in other organizations.

## ABOUT PSI

The Paul Scherrer Institute [1], as it looks today, is the result of many changes over many years.

The institute, named after the Swiss physicist Paul Scherrer, was first created in 1988 when EIR (Swiss Federal Institute for Reactor Research, founded in 1960, “East side”) was merged with SIN (Swiss Institute for Nuclear Research, founded in 1968, “West side”).

The PSI accelerator complex comprises four facilities, HIPA, SLS, PROscan and SwissFEL. The oldest part of what is now High Intensity Proton Accelerator, HIPA - a cyclotron machine, was commissioned in 1974 with subsequent parts added later, building the chain that makes up the HIPA of today. Still in operation almost 40 years later, HIPA still delivers protons, muons and neutrons with more or less the same setup as from the start. The Swiss Light Source, SLS - a synchrotron, was commissioned in 2001, achieving a, for its time, ultra-thin beam. In 2007, a compact cyclotron, COMET, was built specifically for the proton therapy patients, who since 1984 had been served with a split off part of the beam from HIPA. The newest addition to the family is the SwissFEL, a free electron laser, commissioned in 2018. The next big projects on the accelerator side are the ongoing SLS upgrade project, SLS 2, and later on IMPACT, adding two targets in the HIPA accelerator. At SwissFEL, new end stations as well as a fourth transfer line, Porthos, are planned in the coming years. The accelerator complex at PSI is indeed complex, and under the responsibility of the division of Large Research Facilities, GFA.

## CONTROLS AT PSI

The PSI Controls section is part of the GFA division, with its stakeholders mainly split between the GFA machine side expert groups and the research groups of the

<sup>†</sup> simaolhoda.baymani@psi.ch

# TEXTUAL ANALYSIS OF ICALEPCS AND IPAC CONFERENCE PROCEEDINGS: REVEALING RESEARCH TRENDS, TOPICS, AND COLLABORATIONS FOR FUTURE INSIGHTS AND ADVANCED SEARCH

A. Sulc\*, A. Eichler, T. Wilksen, DESY, Hamburg, Germany

## Abstract

In this paper, we show a textual analysis of past ICALEPCS and IPAC conference proceedings to gain insights into the research trends and topics discussed in the field. We use natural language processing techniques to extract meaningful information from the abstracts and papers of past conference proceedings. We extract topics to visualize and identify trends, analyze their evolution to identify emerging research directions, and highlight interesting publications based solely on their content with an analysis of their network. Additionally, we will provide an advanced search tool to better search the existing papers to prevent duplication and easier reference findings. Our analysis provides a comprehensive overview of the research landscape in the field and helps researchers and practitioners to better understand the state-of-the-art and identify areas for future research.

## INTRODUCTION

The field of language processing has noticed remarkable advances and breakthroughs, shaping our understanding of the fundamental principles of working with written knowledge and our ability to automatically process it. As research facilities in the field of particle accelerators continue to push the boundaries in improving the particle accelerators and their controls, a large amount of text corpus has been created, capturing the collective knowledge and discoveries in the community. In this paper, we explore the rich archive of past papers on the particle accelerator from IPAC and ICALEPCS conference proceedings, employing recent language processing techniques. We aim to expose the evolution of ideas, explore the interconnectedness of research areas, and provide a tool for advanced search and discovery that adapts to the language used in the community.

Faced with many potentially relevant papers, researchers must either narrowly limit the scope of their review or rely on new methods to efficiently analyze large document collections. One of the key goals of this work is to help research orient in the overwhelmingly many papers being introduced every year in the community of particle accelerators. To overcome these limitations, researchers are increasingly adopting automated methods for topic modeling [1], semantic search, and knowledge extraction that can rapidly analyze word patterns across thousands of documents to reveal latent thematic structures.

For example, as researchers disseminate their findings, they reference and extend prior work, creating connections

within a knowledge network. A published study may later be cited by another, building links that indicate the flow of knowledge. Highly cited papers and their authors often represent influential roles within this network. There are also links between individual papers based solely on their text content, which, exposed to other papers, creates a hidden underlying complex network that can reveal papers that for instance cover important topics. Additionally, papers can be characterized by the topics they address, and their evolution. Each paper encapsulates a specific topic that may be relevant to a particular sub-community or provide insights into the development of a specific field. Moreover, semantic search plays a pivotal role in unveiling these concealed thematic structures. It empowers scientists to explore textual data beyond traditional full-text search, mitigating issues like typographical errors and synonyms.

This paper serves as an exploration into the possibilities of automatically exposing the collective wisdom stored within the particle accelerator community through contributions presented at the IPAC and ICALEPCS conferences. It provides a unique perspective on the historical evolution of this field, shedding light on previously unnoticed details and noteworthy contributions that might otherwise go unnoticed and help scientists find potential links between their work and potential future collaborations.

The structure of this paper is as follows: First, we provide background on existing approaches for topic modeling of complex corpora, summarizing the current landscape of automated paper processing. We then describe our methodology for enabling robust semantic search across the conference proceedings. This capability serves as the foundation for the analyses that follow, while also providing an invaluable tool for community members to efficiently find relevant papers. Building on this semantic search, we detail our techniques for uncovering latent topics in the corpus and analyzing their evolution over time for both the IPAC and ICALEPCS conferences. Finally, we introduce our approach to extracting knowledge from the abstracts using social network analysis of citation patterns. This reveals influential contributions and concepts based on both textual content and the networks formed by citations between papers.

## RELATED WORK

### *Topic Modeling*

Blei *et al.* [2] pioneered unsupervised topic modeling with Latent Dirichlet Allocation (LDA). LDA represents documents as mixtures of topics, where a topic is a distribution over words. Using Bayesian inference, LDA reverse engineers the latent per-topic word distributions and per-document

\* antonin.sulc@desy.de

# A SUCCESSFUL EMERGENCY RESPONSE PLAN: LESSONS IN THE CONTROLS SECTION OF THE ALBA SYNCHROTRON

G. Cuní\*, O. Matilla, J. Nicolás, M. Pont  
ALBA Synchrotron Light Source, Cerdanyola del Vallès, Spain

## Abstract

These are challenging times for research institutes in the field of software engineering. Our designs are becoming increasingly complex, and a software engineer needs years of experience to become productive. On the other hand, the software job market is very dynamic, and a computer engineer receives tens of offers from private companies with attractive salaries every year. Occasionally, the perfect storm can occur, and in a short period of time, several key people in a group with years of experience leave. The situation is even more critical when the institute is plunged into a high growth rate with several new instruments under way. Naturally, engaged teams will resist reducing operational service quality, but, on the other hand, the new installations milestones dates will approach quickly. This article outlines the decision-making process and the measures taken to cope with this situation in the ALBA Synchrotron's Controls Section. The plan included reorganizing teamwork, but more importantly, redefining the relationship with our clients and prioritization processes. As a result, the team was restructured and new roles were created. In addition, effective coordination was vital, and new communication channels were established to ensure smooth workflows. The emergency peak period is over in our case, but we have learned a lot of lessons and implemented many changes that will stay with us. They have made us more efficient and more resilient in case of future emergencies.

## TRIGGER

This story starts in October 2021, when four people of the Controls Section had recently left and there were three more people in different temporary leaves (i.e. sick leaves and paternity leave). In total, the group's manpower was reduced by 7 people. With an effective reduction of 7 Full-Time Equivalent (FTE) in a group of 16 people, there was a call to the Management Board for a critical ad hoc review and prioritization of the objectives and current developments for the Controls Section. By October 2021 there was a meeting with the Management Board presenting the current situation and developments, making visible the impossibility to do everything that was already started and going on, and start the next foreseen tasks. With a graphical visualization of people's main tasks and "orphaned tasks" as well as the big backlogs waiting to enter into the development cycle, it was clear that there was a need to do a thoughtful exercise at a facility level to plan and execute the most important tasks, hence, minimizing the risk of such situation. During December 2021 and January 2022 we held several meetings with

the three Division Heads affected (i.e. Accelerators, Experiments, and Engineering) to highlight the most important developments required.

## COORDINATION OFFICE CONTROLS CONTINGENCY PLAN FOR 2022

By February 2022 the Coordination Office Controls Contingency Plan was announced by our Director, and in March 2022 we held a kick-off meeting between Controls and the Coordination Office. One of the first tasks done was the creation of a Public Confluence Page where all the activities were highlighted so it was transparent on which developments there will be focus, and, as well, stating which activities were officially paused and stalled. The announcement of the Contingency Plan contained these clear messages:

- Work only on:
  - Incidents to assure operation
  - Activities to assure the progress of the New Beamlines Program (i.e. Beamlines in design, construction, and commissioning phases)
  - Few selected activities related to specific projects from Accelerators, Experiments, and Computing
- Assumptions:
  - Urgent tasks treated as exceptions
  - Internal activities self-managed (e.g. maintenance, critical bugs fixing, etc. )
  - Services reduced (i.e. no requests for changes or new features)
  - duration until December 2022
- Exclusions:
  - PLC Team activities (none of the staff that left was from the PLC Team, so productivity not affected)
  - Activities related to International Collaborations, Students, Newcomers, etc.

## Staff Evolution

Unfortunately, the situation worsened in the following months with four more leavers. In Fig. 1 you can see the evolution of the Controls Section Staff from three years prior the Contingency Plan where there are a standard rate of leavers and newcomers, until one year after the Contingency Plan. As a reference, there are some markers indicating the years of experience of the people leaving; the duration of the Contingency Plan (extended until the end of January 2023); the staff recovery progress from February 2022 to July 2023. Since ALBA is still building new beamlines, we were able to increase the Controls Section Staff size with two more positions, one person is expected to join us in November 2023 and the other position is expected to be opened during

\* guifre.cuni@cells.es

# ENSURING SMOOTH CONTROLS UPGRADES DURING OPERATION

M. Gourber-Pace<sup>†</sup>, F. Huguin, E. Matli, B. Urbaniec, W. Sliwinski, CERN, Geneva, Switzerland

## Abstract

The CERN Accelerator Controls systems have to remain as stable as possible for operations. However, there are inevitable needs to introduce changes to provide new functionalities and conduct important consolidation activities. To deal with this, a formal procedure and approval process, the Smooth Upgrades procedure, was introduced and refined over a number of years. This involves declaring foreseen Controls changes as a function of the accelerator schedules, validating them with stakeholders, and organising their deployment in the production environment. All of this with the aim of minimising the impact on accelerator operation. The scope of this activity is CERN-wide, covering changes developed by all CERN units involved in Controls and encompassing the whole CERN accelerator and facility complex. In 2022, the mandate was further extended with a more formal approach to coordinate changes of the software interfaces of the devices running on front-end computers, which form a critical part of the smooth deployment process. Today, Smooth Upgrades are considered a key contributor to the performance and stability of the CERN Control system.

This paper describes the Smooth Upgrades procedure and the underlying processes and tools such as schedule management, change management, and the monitoring of device usage. The paper also includes the major evolutions which allowed the current level of maturity and efficiency to be reached. Ideas for future improvements will also be covered.

## INTRODUCTION

Making controls changes during a beam run is desirable to deliver novel and enhanced functionality as requested by the CERN Operations team. It is a delicate procedure necessitating meticulous preparation and execution to ensure the preservation of the accelerator performance and stability. This paper describes the formal process applied at CERN, to prepare and execute the deployment of controls changes during beam operation. An emphasis is placed on the importance of documentation, approval, and communication along the process to mitigate adverse effects on operations.

## BACKGROUND

As part of the annual official CERN accelerator planning, a number of beam stops are scheduled to facilitate necessary upgrades and maintenance interventions aimed at enhancing the performance and reliability of the accelerator complex. These stops can be classified into two types:

- **A Technical Stop (TS)**, which occurs once or twice a year and typically lasts between 12 to 24 hours.

- **A Year-End Technical Stop (YETS)**, which spans several weeks starting in November or December.

Both the TS and YETS periods provide an opportunity for implementing and deploying controls upgrades. During these intervals, accelerators are stopped to facilitate various interventions on components such as radio frequency cavities, magnet power supplies, beam instrumentation, etc. Across the entire accelerator complex, an established protocol dictates that controls upgrades should only be executed during a TS or YETS. The only rare exceptions may be a bug fix or specific new feature deployment urgently requested by Operations teams.

This policy has grown more rigorous over time, stemming from lessons learned from past controls upgrades. In 2015 and 2016, a qualitative assessment was conducted, uncovering that controls upgrades had adverse effects on LHC performance. Following controls software deployments, numerous hours of beam operation were compromised. Several factors contributed to issues during the restart after TS and YETS events, including: inadequate pre-deployment testing, underestimation of the impact on interconnected systems, introduction of non-backward compatible changes, insufficient communication resulting in Operations' unfamiliarity with the deployed changes.

Drawing on this experience, the Smooth Upgrades procedure was formulated to facilitate a coordinated deployment of controls changes across the accelerator complex.

## SMOOTH UPGRADES PROCEDURE

### Mandate

The Smooth Upgrades (SU) procedure outlines a method to be applied during all TS and YETS periods, aiming to streamline deployments and minimize the risk of impact on accelerator operation. The SU procedure covers several interconnected needs:

- **Central Repository:** to document planned upgrades from controls teams in a centralized repository. The goal is to compile a comprehensive list of changes, allowing the Operations team to correlate issues during beam restarts with recently deployed modifications.
- **Approval and Compliance Workflow:** to ensure that upgrades and interface modifications have received approval from the Operations team before proceeding with the implementation.
- **Conflict Analysis and Mitigation:** to identify potential conflicts amongst the planned upgrades and establish priorities and / or deployment order.
- **Risk Evaluation:** to assess operational risks, ensure the existence of validation procedures, and establish contingency procedures for rolling back changes if necessary.
- **Process Review:** to evaluate the effectiveness of the procedure after each TS or YETS, gather user

# MAINTENANCE OF THE NATIONAL IGNITION FACILITY CONTROLS HARDWARE SYSTEM\*

J. Vaher†, G. Brunton, J. Dixon

Lawrence Livermore National Laboratory, Livermore, CA 94550 USA

## Abstract

At the National Ignition Facility (NIF), achieving fusion ignition for the first time ever in a laboratory required one of the most complex hardware control systems in the world. With approximately 1,200 control racks, 66,000 control points, and 100,000 cables, maintaining the NIF control system requires an exquisite choreography around experimental operations while adhering to NIF's safety, security, quality, and efficiency requirements. To ensure systems operate at peak performance and remain available at all times to avoid costly delays, preventative maintenance activities are performed two days per week as the foundation of our effective maintenance strategy. Reactive maintenance addresses critical path issues that impact experimental operations through a rapid response 24x7 on-call support team. Prioritized work requests are reviewed and approved daily by the facility operations scheduling team. NIF is now in the second decade of operations, and the aging of many control systems is threatening to affect performance and availability, potentially impacting planned progress of the fusion ignition program. The team is embarking on a large-scale refurbishment of systems to mitigate this threat. Our robust maintenance program will ensure NIF can capitalize on ignition and push the facility to even greater achievements. This paper will describe the processes, procedures, and metrics used to plan, coordinate, and perform controls hardware maintenance at NIF.

## INTRODUCTION

### *National Ignition Facility Background*

The National Ignition Facility (NIF), currently the world's largest and most energetic laser system, was built by the United States Department of Energy (DOE) National Nuclear Security Administration (NNSA) inside the one square mile boundary of the Lawrence Livermore

National Laboratory (LLNL) site in Livermore, California. The purpose of NIF was to execute high-energy-density (HED) laser experiments as a national and international user facility. These HED laser experiments would enable the study of physics for discovery science and inertial confinement fusion (ICF), while serving as a key experimental capability in the NNSA's establishment of a science-based stockpile stewardship program (SSP) in the absence of underground nuclear testing [1].

In 1997, construction started on the NIF building, which stands 10 stories tall and covers an area approximately the size of three football fields (Fig. 1). Construction involved the installation a target chamber weighing 130 metric tons and measuring 10 meters in diameter (Fig. 2), beampath for 192 laser beams, and the largest capacitor bank in the world capable of storing 300-400 megajoules of energy.




Figure 2: The 1999 installation of the NIF target chamber inside the half-constructed Target Bay. Credit: LLNL.



Figure 1: The National Ignition Facility at Lawrence Livermore National Laboratory in Livermore, California. Credit: Jason Laurea/LLNL.

\* This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Release # LLNL-CONF-855229.

† vaher1@llnl.gov

# ACCELERATOR CONTROL CLASS FOR GRADUATE STUDENTS IN SOKENDAI, KEK

N. Kamikubota<sup>1†</sup>, K. Furukawa<sup>1</sup>, M. Satoh<sup>1</sup>, S. Yamada<sup>1</sup>, and N. Yamamoto<sup>1</sup>, KEK, Ibaraki, Japan  
<sup>1</sup>also at the Graduate University for Advanced Studies, SOKENDAI, Hayama, Kanagawa, Japan

## Abstract

The Graduate University for Advanced Studies, known as SOKENDAI, provides educational opportunities for graduate students in collaboration with national research institutes in Japan. KEK is one of the institutes, and has a program *Accelerator Science*. Since 2019, we started two classes: *Introduction to accelerator control system* for one semester, and a two-day *Control of distributed devices for large systems*. The former consists of 12 lectures on various topics of accelerator controls by teachers, followed by a presentation day by students. The latter consists of lecture and hands-on, which enables students to practice EPICS with Raspberry-pi based devices. In the paper, status of accelerator control classes are reported.

## SOKENDAI AND KEK

### About SOKENDAI

SOKENDAI, the Graduate University for Advances Studies, was established in 1988, as a national university of Japan [1]. The headquarter is located in Hayama, Kanagawa, Japan. SOKENDAI does not have an undergraduate course. In close partnership and collaboration with research institutes, SOKENDAI operates Ph.D doctoral programs [2].

There are 20 research institutes in the scheme (see Fig. 1). They cover variety of fields: sciences of information, statistics, physics, accelerator, astronomy, fusion, space, molecular, material, environment, biology, physiology, polar, and cultural studies of anthropology and Japan. In 2023, 20 programs are available for education of graduate students, associated with the above research institutes.

### KEK and Accelerator Science Program

Since the foundation of SOKENDAI in 1988, KEK has been one of the research institutes of SOKENDAI. With the partnership, KEK provides three programs: a) Particle and Nuclear Science Program, b) Materials Structures Science Program, and c) Accelerator Science Program (see Table 1). The Accelerator Science Program is associated with two laboratories of KEK, Accelerator Laboratory and Applied Research Laboratory. The latter concerns radiation science, computing research, cryogenic research, and mechanical engineering. Table 2 shows the number of students of KEK's Programs in 2023. All of the programs contain certain amount of international students. They are mostly from Asian countries.

The Accelerator Science Program consists of various courses related to accelerator technologies as in Fig. 2. There are two types of courses, a half-year course and a

short course (less than a week). Most of them are the courses which will be held on a student's request. The complete list of the courses is given elsewhere [3].

Table 1: Programs and Associated Research Institutes or Laboratories of KEK

Program	Research Institutes
Particle and Nuclear Science	Institute of Particle and Nuclear Physics, KEK
Materials Structure Science	Institute of Material Structures Science, KEK
Accelerator Science	Accelerator Laboratory, Applied Research Laboratory, KEK




Figure 1: Partner research institutes of SOKENDAI [1] with highlights of the headquarter (HQ) and KEK.

Table 2: Numbers of Students of KEK's Programs

Program	Number of Students	Number of International Students
Particle and Nuclear Science	48	13 (27%)
Materials Structure Science	8	4 (50%)
Accelerator Science	14	5 (36%)

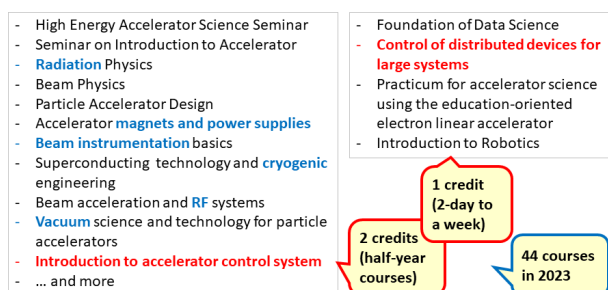


Figure 2: Part of courses of the *Accelerator Science Program*.

<sup>†</sup> norihiko.kamikubota@kek.jp



# EXTENDING THE COVERAGE OF AUTOMATED TESTING IN ITER'S CONTROL SYSTEM SOFTWARE DISTRIBUTION\*

R. Lange<sup>†</sup>, H. Kim, A. Zagar, ITER Organization, St. Paul lez Durance, France  
M. Ruiz, V. Costa, J. Nieto, Grupo de Investigación en Instrumentación y Acústica Aplicada,  
Universidad Politécnica de Madrid, Madrid, Spain

## Abstract

As part of the effort to standardize the control system environment of ITERs in-kind delivered >170 plant systems, the Controls Division publishes CODAC Core System (CCS), a complete Linux-based control system software distribution.

In the past, a large part of the integrated and end-to-end software testing for CCS was executed manually, using many long and complex test plan documents. As the project progress introduces increasing scope and higher quality requirements, that approach was not maintainable in the long term.

ITER CODAC and its partners have started a multi-year effort converting manual tests to automated tests, inside the so-called Framework for Integration Testing (FIT), which itself is being developed and gradually extended as part of the effort. This software framework is complemented by a dedicated hardware test stand setup, comprising specimens of the different controllers and I/O hardware supported by CCS. FIT and the test stand will allow to run fully scripted hardware-in-the-loop (HIL) tests and allow functional verification of specific software modules as well as different end-to-end use cases.

## INTRODUCTION

The ITER project is a collaboration between seven members (China, Europe, India, Japan, Korea, Russia and the USA) representing 35 countries. Construction of the ITER Tokamak facility in southern France is largely (>90 %) based on in-kind procurement, i.e. the members developing and delivering hardware, components and systems to the project. That poses major challenges to the central ITER Organization, which is responsible for the specification, integration and operation of the machine.

Within the scope of the control systems aspect of CODAC (Controls, Data Access and Communication), the most important mitigation strategy is standardization. Hardware is standardized by limiting choices through strictly applying hardware catalogues. As part of standardizing software, ITER publishes and distributes a complete software distribution, based on Red Hat Enterprise Linux and named CODAC Core System (CCS) [1, 2].

Running this CCS software distribution is mandatory for all ITER control system-related computers.

\* Work partially funded by PID2019-108377RB-C33/MCIN/AEI (Agencia Estatal de Investigación) /10.13039/501100011033 and PID2022-137680OB-C33/MCIN/AEI /10.13039/501100011033 / FEDER/ and the European Union.

<sup>†</sup> ralph.lange@iter.org

## Software Quality and Testing

Over the project's lifetime, the focus of CCS development activity has shifted. While the first versions were mainly adding features to reach the required functionality, recent versions were stabilizing the distribution, increasing its robustness and quality. With that shift, mandatory Quality Assurance (QA) tools have been introduced and software testing has gained importance.

Original concepts had foreseen the freezing of CCS software and plant system applications after site acceptance and integrated commissioning. As the project schedule evolved, it has become clear that continuing to update CCS is crucial to adapt to newer hardware and software, and to keep the Operating System parts under a support agreement.

## THE EXISTING TEST ENVIRONMENT

On the front-end controller level, ITER maintains around 20 proprietary software modules, most of them Linux drivers and EPICS Device Support modules for I/O boards in the ITER hardware catalogue.

ITER software development follows a workflow based on ISO 12207. The testing required as part of this workflow is defined as a Software Test Plan (STP) document for each of the software modules. The procedures in these test plans are executed manually and results are entered into a Software Test Report (STR). Some STPs contain well over a hundred pages and take days to execute.

During the first years of development, most of the manual testing was done by external contractors, using a considerable amount of resources.

## Changes Require a New Concept

With changes in the ITER project schedule and re-distribution of the budget, it became clear that this approach was not cost-effective and not sustainable.

To mitigate this, test coverage was cut down by only running a subset of tests and STPs with every release.

As a result of the limited testing, quality was decreasing. More bugs were found by users and had to be fixed and patched under time pressure.

## AN AUTOMATED APPROACH

### Automated Testing

Automated testing has always been used within CODAC Core System. The scope of automated testing, however, was limited to unit tests that are executed as part of the software module's build process.

## EPICS JAVA DEVELOPMENTS

K. Saintin<sup>†</sup>, L. Caouën<sup>‡</sup>, P. Lotrus  
CEA/DRF/IRFU, Saclay, France

### Abstract

The CEA IRFU/DIS [1] software control team is involved from feasibility studies to the deployment of equipment covering low level (hardware, PLC) to high level (GUI supervision). For their experiments, DIS software control solution is using two mains frameworks:

MUSCADE, a full Java in-house solution, an embedded SCADA dedicated to small and compact experiments controlled by PLC (Programmable Logic Controller), only compatible with Windows operating system (OS) for the server side.

EPICS [2], a distributed control systems to operate devices such as particles accelerators, large facilities and major telescopes, mostly deployed on Linux OS environments.

EPICS frameworks provides several languages for bindings and server interfaces such as C/C++, Python and Java. However, most of the servers also called IOC developed in the community are based on C/C++ and Linux OS System. EPICS also provides extensions developed in Java such as Appliance (the archiving tool), Phoebus Control-Studio [3] (GUI), and Display Web Runtime (Web Client). All these tools depend on CAJ a pure Java implementation Channel Access Library.

Today, MUSCADE users work under Windows operating system, and they need intuitive tools that provide the same features than MUSCADE in their EPICS projects. Thus, research and development activities mainly focus on EPICS solution adaptation. It aims to explore further CAJ library, especially on the server side aspect. In order to achieve this goal, several developments have been carried out since 2018:

PLCParserTool, PLC simulation to test EPICS synoptic without any hardware.

CAFEJava, Java IOC EPICS to provide process variables connected to any java application. We developed EPICS4MUSCADE bridge and Web Client for MUSCADE on this basis.

## INTRODUCTION

IRFU is working on different kind of experiments from particle accelerator installed in big facilities to gas station installed in the middle of the fields. IRFU participates regularly to the construction of accelerators around the world (ESS [4], SARAF [5], and SPIRAL2 [6] ...) and its control system team has developed EPICS skills for more than 20 years. EPICS and associated software are mainly designed for big facilities, which is not convenient for some of our experiment; especially small ones based on PLC control

To address this lack of feature, a large number of small control systems (more than 90 devices) are controlled using MUSCADE © Java Software. This framework is custom-built for our laboratory and offers a range of tools to control small-scale facilities and visualize experiments through synoptic (see Fig. 1).

The primary objective for our team is to migrate MUSCADE experiments to EPICS control system. EPICS knowledge is widely shared within a large community of developers and numerous facilities. However, because MUSCADE is a custom-built system, there are only a few persons capable of supporting these experiments.

To facilitate a gradual migration of all MUSCADE experiments, we have developed several tools that allow us to maintain the MUSCADE server-side functionality while leveraging the benefits of EPICS High-Level applications for display, archiving, and alarm notifications. In this article, we will focus on three of these developments and explain how they provide an adapted solution for a reliable migration.

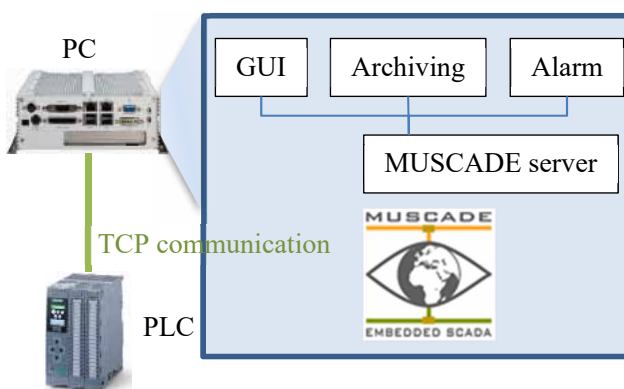


Figure 1: MUSCADE © SCADA an embedded solution.

## DXF2BOB CONVERTER

At IRFU, MUSCADE users build their supervision views with AutoCAD for 2D drawings software, especially for cryogenic sequential function chart [7] (SFC or GRAFCET), which can be tricky (see Fig. 2). AutoCAD is a vector drawing solution, which is convenient for precise drawing, and so for complex synoptic. It is also a popular software, especially in the industry, by architects, project managers, engineers, graphic designers, city planners and other professionals.

MUSCADE provides a tool developed in Java, which gives the possibility to import AutoCAD file for generating synoptic views.

Software

Control Frameworks for Accelerator & Experiment Control

<sup>†</sup> katy.saintin@cea.fr

<sup>‡</sup> loic.caouen@cea.fr

# REAL-TIME VISUALIZATION AND PEAK FITTING OF TIME-OF-FLIGHT NEUTRON DIFFRACTION AT VULCAN

B. A. Sobhani, Y. Chen, Oak Ridge National Laboratory, Oak Ridge, TN, USA

## Abstract

In situ neutron diffraction experiments on state-of-the-art diffractometers such as VULCAN at SNS, enable capturing key transients in the dynamic material evolution, for in-depth understanding of material science and engineering. Traditional data processing needs to transfer the time-of-flight neutron events to another computation cluster for post-experiment data reduction and analysis. This is not timely, and it thus cannot meet the demands of on-the-fly decision making and close-loop experiment control. This paper demonstrates an automatic and intuitive system at VULCAN that is developed in EPICS for a real-time visualization of live neutron diffraction patterns and Bragg peaks' fitting via generalized linear regression. The live plots and the results quantify the dynamic of the peak position, intensity, and width, which are indicators of materials straining, phase transformation, microstructure evolution and so on. It constructs a promising platform of future smart neutron experiment controls.

## REGRESSION

Linear regression is a versatile technique for not only linear function fitting but also the parameter estimation of complex functions or non-linear curves. The "linear" requirement only needs to apply to the fitting coefficients, so a class of functions that is non-linear in the independent variable can be fit by linear regression, as long as it is linear with respect to the fitting coefficients.

However, if one wants to fit a gaussian function, for example, to data using fitting parameters of amplitude, mean, and standard deviation, linearity does not hold.

$$f(A, \mu, \sigma) = A * e^{-(x-\mu)^2/(2*\sigma^2)}$$

Fortunately, this is of little concern because the Levenberg-Marquardt algorithm allows one to fit functions that are nonlinear even in the fitting coefficients. Implementations of the Levenberg-Marquardt are widely available on the internet. In particular, `scipy's optimize.curve_fit` function calls this algorithm so it can be used to fit functions that are nonlinear in the fitting coefficients. As a matter of curiosity, `scipy's optimize.curve_fit` function is largely not implemented in python but is instead only a python interface to the compiled Fortran fitting library called MINPACK,[1] which was developed at Argonne National Laboratory in the 1980s, and LAPACK for calculating the covariance matrix.

## PHYSICAL INTERPRETATIONS

The peak profile fitting exports some basic parameters such as peak position, area and width, which reflect the material status. In the following, the lattice spacing  $d$  as Bragg's peak position and the lattice strain are employed

for demonstration. When neutrons with the wavelength  $\lambda$  are diffracted by the lattice plane of a crystal with spacing of  $d$ , the Bragg's Law as the following must be satisfied:

$$\lambda = 2d \sin\theta$$

where  $2\theta$  is the scattering angle.

This can be related to time-of-flight data from SNS neutron detectors according to the following:

$$\frac{ht}{mL} = 2d \sin\theta$$

Given the instrument optics, the measured time-of-flight ( $t$ ) can be converted to the lattice  $d$  spacing [2]. This is presented as the Bragg's peak position. It is well known that the peak profile function at the TOF diffractometer can be as complex as a back-to-back exponential function convoluted with a pseudo-Voigt function. For the purpose on demonstrating the relative changes of the peaks, it is rational to employ a simpler Gaussian function in this paper. Therefore, the lattice spacing  $d$  of a particular crystal plane can be estimated as  $\mu$  by fitting the Gaussian function over the corresponding Bragg peak in the neutron diffraction pattern.

The lattice spacing  $d$  is responsive to the external stimulus such as temperature and stress. The relative change from a reference state  $d_0$  (such as room temperature and stress-free state) is calculated as lattice strain  $\epsilon = (d-d_0)/d_0$ . Over a dynamic process, the phase-specific and the lattice-plane-specific lattice strains evolve, reflecting materials properties such as thermal expansion and stiffness, and indicating possible material changes via different mechanisms.

The change of the lattice spacing shifts the Bragg's peak. In this way, the  $\mu$  fitting parameter can be used to quantify the strain.

## PEAK FITTING INTERFACE

User can select the peak to be fit by adjusting draggable bounds on a CSS Phoebus interface. The bounds then get written to PVs, and the peak fitting algorithm is run only on data between those two bounds.

## PEAK QUALITY

`Scipy's optimize.curve_fit` function returns a covariance matrix as its second return value.[3] The diagonals of this matrix are the variance of the fitting parameters – these are measures of certainty of the fit, not related to the variance of the gaussian peak itself.

Since these variances give us a way to quantify the quality of the peak fit, this provides an efficient way of automating the process of determining when enough data has been collected in a measurement, so that the next measurement can begin.

# PyDM DEVELOPMENT UPDATE

J. Bellister, Y. Yazar  
SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

PyDM is a PyQt-based framework for building user interfaces for control systems. It provides a no-code, drag-and-drop system to make simple screens, as well as a straightforward Python framework to build complex applications. Recent updates include expanded EPICS PVAccess support using the P4P module. A new widget has been added for displaying data received from NTTables. Performance improvements have been implemented to enhance the loading time of displays, particularly those that heavily utilize template repeaters. Additionally, improved documentation and tutorial materials, accompanied by a sample template application, make it easier for users to get started.

## NEW FEATURES

PyDM has had several feature releases during the past year. These releases include new widgets, better EPICS 7 support, and performance enhancements.

### EPICS 7

In order to support the EPICS pvAccess protocol, a new data plugin was created for widgets to communicate with. This plugin uses the PVAccess for Python (P4P) [1] wrapper under the hood to perform the standard set of EPICS requests for interacting with process variables (PVs).

To connect a widget to a device with EPICS 7 support is a very simple change to the address set in the widget. Instead of the channel access prefix, `ca://`, the prefix `pva://` is used in the address. This will route communication between the widget and the data source to the P4P data plugin.

The focus of the new plugin is on the normative types. It also includes an upgrade to the widget for displaying images represented by NTNDArrays. Multiple compression algorithms are supported to provide automatic decompression based on the algorithm specified within the structured data.

### New Widgets

One new widget of note is the PyDMNTTable. As can be seen from the name, this widget is for displaying and interacting with data formatted as a table according to the specification in the EPICS 7 normative types document [2]. Like other widgets this table can be used without writing any code, dragging and dropping it into a display from designer. The table can be run in read-only mode, or with writes enabled as well. When writes are allowed, individual cells in the table can be written to and all updates will be reflected in the source PV. It is also possible to pass a subfield of a table to other non-table widgets if only a certain part of the table would be useful to display rather than all of it at once. An image showing a full table is available in Fig. 1.

Another widget that has been added is the PyDMArchiverTimePlot. This widget enhances existing time plots with the ability to communicate with an instance of the EPICS archiver appliance. By setting the timespan on this plot, the widget will call the archiver appliance on startup and request that much historical data to backfill the plot. It is also possible to request archived data while the plot is running by scrolling the x-axis or zooming out on the plot. The signals generated by taking these actions on the plot will invoke calls to the archiver.

To support this new widget, the archiver data plugin was enhanced to make asynchronous calls to the archiver. This prevents calls for large amounts of data from blocking the main PyQt thread and freezing the user interface. The QNetworkAccessManager is utilized to make these non-blocking calls while integrating nicely with the standard signals and slots used by Qt.

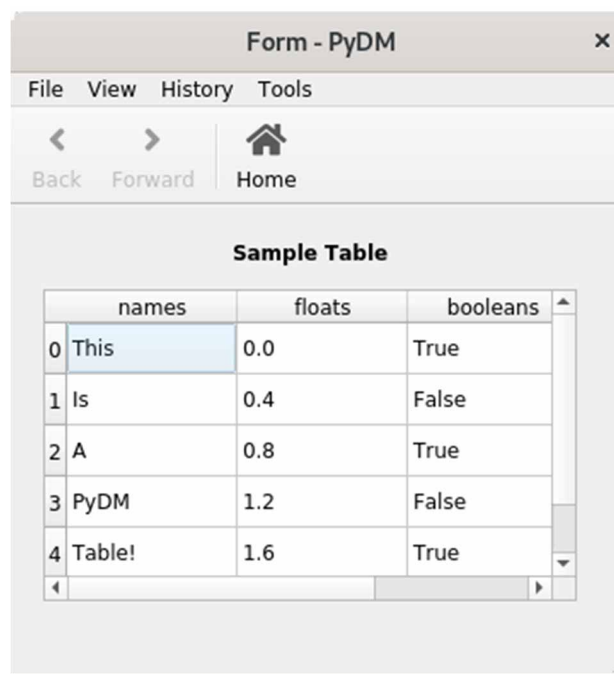


Figure 1: The new PyDMNTTable widget.

### Usability Enhancements

A number of enhancements have been made to improve the usability of PyDM for both display creators and end users. It is now easier to view EPICS field information on widgets thanks to an upgrade to tooltips. Specifying a field name preceded by a period will cause the tooltip to display the actual value of that field while the display is running.

To further customize their displays, creators can add custom menus to the toolbar to take actions relevant to their displays. An option has also been added to the menu for switching to a new stylesheet while a display is running.

# DYNAMIC CONTROL ROOM INTERFACES FOR COMPLEX PARTICLE ACCELERATOR SYSTEMS

B. E. Bolling\*, D. Nordt, G. Fedel, M. Munoz, European Spallation Source ERIC, Lund, Sweden

## Abstract

The European Spallation Source (ESS) is a research facility under construction aiming to be the world's most powerful pulsed neutron source. It is powered by a complex particle accelerator designed to provide a 2.86 ms long proton pulse at 2 GeV with a repetition rate of 14 Hz. Commissioning of the first part of the accelerator has begun and the requirements on the control system interfaces varies greatly as progress is made and new systems are added. In this paper, three such applications are discussed in separate sections.

A Navigator interface was developed for the control room interfaces aimed towards giving operators and users a clear and structured way towards quickly finding the needed interface(s) they need. The construction of this interface is made automatically via a Python-based application and is built on applications in any directory structure both with and without developer interference (fully and semi-automatic methods).

The second interface discussed in this paper is the Operations Accelerator Synoptic interface, which uses a set of input lattices and system interface templates to construct configurable synoptic view of the systems in various sections and a controller panel for any selected system.

Lastly for this paper there is a configurable Radio Frequency Orchestration interface for Operations, which allows in-situ modification of the interface depending on which systems and components are selected.

## INTRODUCTION

Complex machinery always poses multiple challenges when it comes to designing comprehensible graphical user interfaces (GUIs) such that the operators can work efficiently. Large-scale particle accelerators, with many unique subsystems are an extreme example of this. In this paper, the term Operator Interface (OPI) is used to describe any GUI that is used within the Phoebus framework [1] as it is intended to be used by an ESS Operator.

To address a few challenges of having to operate such complex machinery, two Python scripts were developed to dynamically build OPIs based on user inputs (accelerator lattice files). These OPIs are referred to as being dynamic as they are generated (updated) when needed. A third OPI is also described which is dynamic for the user during runtime.

## OPERATOR INTERFACES

At ESS the OPIs are structured in three different levels aimed for control room operators ("Operator level"), system expert functions ("System Expert level") and for lowest-level functions ("Engineering level"). With this methodology implemented already during commissioning and conditioning

phases, it became possible to find which settings the different levels of OPIs should be exposed to. Further on, each have a subset of directories for each part of the machine (e.g. Accelerator and Target), followed by another subset of directories for each system of that part of the machine (e.g. for the Ion Source section or Radio Frequency systems).

Within each machine-part directory lies the OPIs, which are designed in accordance with the internal OPI visual design rules document. The visual design document describes what colour codes are supposed to be used for what type of state and object, which type of object(s) to use for which type of setting/information, etc., in order to establish a site-wide continuity, with the outcome being a higher situational awareness for the control room operator and hence increasing the reliability of the facility [2].

## NAVIGATOR OPI

The openness of Phoebus introduces the issue that having a large number of user interfaces scattered across a large amount of directories, with many cases having support files that cannot be opened as standalone applications, the user interfaces needed become difficult/time-consuming to be found. The OPIs Navigator OPI is a prototype interface aimed at solving the issue by dynamically constructing an OPI to enable users to navigate amongst relevant user interfaces - saving users a lot of time and effort. A user flowchart is shown in Fig. 1.

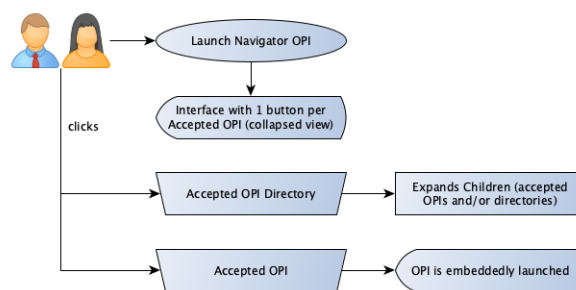


Figure 1: Navigator OPI user flowchart.

## Methodology

A Python script is used to render the Navigator OPI which loops through a given set of directories using a recursive strategy for any subdirectory and with a set of filters applied, such as methods for identifying support files such that they can be omitted. The script then launches a PyQt5-based user interface with a table of OPIs identified such that the developer may select which OPIs are to be included in the

\* benjamin.bolling@ess.eu

# EXTENDING PHOEBUS DATA BROWSER TO ALTERNATIVE DATA SOURCES\*

Mihnea Romanovschi<sup>†</sup>, Ivan Finch<sup>1</sup>, Gareth Howells<sup>1</sup>

<sup>1</sup>ISIS Neutron and Muon Source, Harwell Campus, Oxfordshire, OX11 0QX, UK

## Abstract

The Phoebus user interface to EPICS is an integral part of the upgraded control system for the ISIS Neutron and Muon Source accelerators and targets. Phoebus can use the EPICS Archiver Appliance, which has been deployed as part of the transition to EPICS, to display the history of PVs. However, ISIS data historically has and continues to be stored in the InfluxDB time series database. To enable access to this data, a Python application to interface between Phoebus and other databases has been developed. Our implementation utilises Quart, an asynchronous web framework, to allow multiple simultaneous data requests. Google Protocol Buffer (Protobuf), natively supported by Phoebus, is used for communication between Phoebus and the database. By employing subclassing, our system can in principle adapt to different databases, allowing flexibility and extensibility. Our open-source approach enhances Phoebus's capabilities, enabling the community to integrate it within a wider range of applications.

## INTRODUCTION

The Experimental Physics and Industrial Control System (EPICS) is a set of software tools and applications which provide a software infrastructure for use in building distributed control systems to operate devices such as particle accelerators, large experiments and major telescopes [1].

As part of the transition from the commercial VSystem [2] to the open-source EPICS, ISIS Accelerator Controls have deployed the EPICS Archiver Appliance [3] to record data across the control system. It has been used to archive the EPICS Process Variables (PVs) from Target Station 1 (TS1), following its recent upgrade [4].

While the EPICS Archiver Appliance offers several advantages, such as its seamless integration within the broader EPICS ecosystem and its lightweight deployment, it has limitations when applied to dynamic systems where the PV definitions are evolving. As ISIS has opted for a hybrid approach to its transition to EPICS [5], the PV definitions are likely to change over time. For example, we recently swapped from the `NTScalAr BOOLEAN` type to the `NTEnum` type for binary PVs. The current state of the EPICS Archiver Appliance does not easily provide the option to alter the data type of these PVs. The system lacks flexibility for retroactive data alterations. Its user API is challenging to integrate into novel systems, such as Machine Learning (ML) based systems. These ML systems might require asynchronous data for training, as mentioned in [6], or heavy use of statistics.

These statistics would need to be placed on the ML application. In contrast, InfluxDB offers a Python-like language called Flux, accessible through its API calls for more flexible data processing and benefits from a larger open-source community compared to the EPICS Archiver Appliance [7]. Our group's preference for InfluxDB stems from the fact that we've been archiving data with it since 2019, accumulating several more years' worth of data compared to the EPICS Archiver Appliance, which we only began using for archiving in 2022. One notable advantage of InfluxDB is its capability to easily back-fill data, whereas the EPICS Archiver Appliance lacks this feature.

To address the mentioned limitations while remaining within the EPICS ecosystem, our team has experimented with substituting the EPICS Archiver Appliance with InfluxDB as the data source for the Phoebus Data Browser. This involves introducing an additional application to serve as a mediator between Phoebus Data Browser [8] and InfluxDB.

## DATABASE WRAPPER

The Database Wrapper is an alternative endpoint for EPICS application requests, such as the Phoebus Data Browser, translating them to the chosen database's API, like InfluxDB, and sending the databases response in Protobuf [9] binary via HTTP to the requesting application.

The entire system is divided into four main components, adhering to a Model-View-Controller design approach (as depicted in Fig. 1):

- Phoebus Application: the view side of the system; it displays the control screens and the information from the archiving database. It also performs requests to the controller regarding what information is to be displayed.
- Quart Server [10]: an asynchronous server that acts as the controller. It unpacks the HTTP requests and delivers data from the model in an asynchronous manner. This approach minimizes the latency between processing information from a database and presenting it to Phoebus, ensuring optimal performance.
- Model: Performs the data conversion from the format that the databases API responds into the format that Phoebus expects, in this case a Protobuf binary.
- Database: for data collection, statistical analyses and metadata lookup. For example InfluxDB and CouchDB [11] respectively.

\* Work supported by Science and Technology Facilities Council (STFC)

<sup>†</sup> mihnea.romanovschi@stfc.ac.uk

# FRONT-END MONITOR AND CONTROL WEB APPLICATION FOR LARGE TELESCOPE INFRASTRUCTURES: A COMPARATIVE ANALYSIS

Stefano di Frischia\*, INAF-OAAb, Teramo, Italy  
Matteo Canzari, INAF-OAAb, Teramo, Italy  
Valentina Alberti, INAF-OAT, Trieste, Italy  
Athos Georgiou, CGI Scotland, Edinburgh, UK  
Hélder Ribeiro, Universidade do Porto, Porto, Portugal

## Abstract

A robust monitor and control front-end application is a crucial feature for large and scalable radio telescope infrastructures such LOFAR and SKA, whereas the control system is required to manage numerous attribute values at a high update rate, and thus the operators must rely on an affordable user-interface platform which covers the whole range of operations. In this paper two state-of-the-art web applications such Grafana and Taranta are taken into account, developing a comparative analysis between the two software suites. Such a choice is motivated mostly because of their widespread use together with the TANGO Controls Framework, and the necessity to offer a ground of comparison for large projects dealing with the development of a monitor and control GUI which interfaces to TANGO. We explain at first the general architecture of both systems, and then we create a typical use-case where an interactive dashboard is built to monitor and control a hardware device. Then, we set up some comparable metrics to evaluate the pros and cons of both platforms, regarding the technical and operational requirements, fault tolerances, developers and operators efforts, and so on. In conclusion, the comparative analysis and its results are summarized with the aim to offer the stakeholders a basis for future choices.

## INTRODUCTION

Nowadays, the complexity of the most important scientific projects related to radio-astronomy require the design of large infrastructures involving both hardware and software side. A prominent software architecture challenge for the correct operational behaviour of such extensive infrastructure is not only the acquisition and processing of a huge amount of data, but also the management of an affordable monitor and control system. This system must be capable of inspecting a large number of variegated attributes and values of the station, and therefore allowing an automated or user-controlled action if a certain event occur during the operational time of the station.

A crucial part of the monitor and control infrastructure is played by the front-end application which usually allows the operators to manage the whole station configuration and its range of operations, and let them be able to act in real time if any action is needed. Since the control system is required to manage numerous attribute values at a high update rate,

the front end application is required to possess significant ad strict requirements such affordability, consistency, security, fault-tolerance, user-friendly interface, and many other features which are covered in the present paper.

The two infrastructures being analysed in this paper are among the largest radio-astronomy facilities in the world, one fully operational and one under development: respectively LOFAR [1] and SKAO. LOFAR (Low-Frequency Array) is structured as an antenna network located mainly in the Netherlands, and spreading across 7 other European countries. Originally designed and built by ASTRON, it makes observations in the 10 MHz to 240 MHz frequency range with two types of antennas: Low Band Antenna (LBA) and High Band Antenna (HBA), optimized for 10-80 MHz and 120-240 MHz respectively [2]. The SKA (Square Kilometer Array) Observatory is an intergovernmental and international radio telescope project being built in Australia (low-frequency) and South Africa (mid-frequency) [3]. It is designed to reach a continuous frequency coverage from 50 MHz to 14 GHz. The frequency range from 50 MHz to 14 GHz requires more than one design of antenna, and so the SKA will comprise separate sub-arrays of different types of antenna elements that will make up the SKA-low, SKA-mid and survey arrays.

Regarding the front-end control interface, a comparative analysis between two different architectural and software choice of both the aforementioned observatories is covered in this paper. On LOFAR side, the software tool GRAFANA, a multi-platform open source analytics and interactive visualization web application, will be examined, since it is the LOFAR adopted choice at the present day. On the other hand, on SKAO side, the Taranta suite, a tool for creating dashboards and interacting with the devices within a TANGO Control System, will be examined as the SKAO preferred choice at the present day.

## FRONT-END CONTROL APPLICATIONS

In the present section a concise explanation of the chosen front-end frameworks from the point of view of the software architecture is carried out. In particular, the description will be focused on the context, the purpose, the features, the roles and several other characteristics related to the scope of this paper, emphasizing thus their use in radio-astronomy facilities.

The comparative analysis has been chosen to examine the front-end web interfaces, due to their high impact to the

\* stefano.difrischia@inaf.it

# UPGRADING AND ADAPTING TO CS-STUDIO PHOEBUS AT FACILITY FOR RARE ISOTOPE BEAMS

T. Ashwarya, J. LeTourneau, M. Ikegami, C. Morton

Facility for Rare Isotope Beams, Michigan State University, East Lansing, USA

## Abstract

The Facility for Rare Isotope Beams (FRIB) has been an early adopter of the CS-Studio ecosystem for its needs for a feature-rich and user-friendly interface with EPICS and the underlying accelerator controls infrastructure. For more than a decade, FRIB has developed thousands of operator displays spanning many areas of accelerator operations and engineering using the CS-Studio display tool “BOY”. CS-Studio has provided many useful control system tools like an alarm system (called “BEAST”), scan system, channels aggregator, display manager and more to support controls operations of the FRIB accelerator. In recent years, there has been a major redesign of the CS-Studio software architecture resulting in the new and upgraded CS-Studio Phoebus. Phoebus replaces the Eclipse RCP- and SWT-based Java framework with modern Java standards like JavaFX, SPI, Adapters and more. This paper details the efforts that have been made at FRIB to adapt and migrate to the upgraded CS-Studio Phoebus for FRIB’s operations and engineering needs.

## ALARM SYSTEM

FRIB deploys over 20 instances of the CS-Studio alarm system [1] to monitor thousands of EPICS [2] Process Variables (PVs) throughout the FRIB accelerator system [3]. Alarm systems are provided a configured list of PVs to monitor and report when the PV values are out of their normal value range or get disconnected. Operators in the FRIB control room use the alarm system to determine which PVs are not in an okay state and need attention. Some engineering groups have their own specialized alarm servers that notify them through emails and phone texts when their concerned PVs are not in the state that they should be.

The alarm system of CS-Studio changed its implementation from Apache ActiveMQ and a relational database to Apache Kafka for Phoebus [4]. At FRIB, we run a 3-node Kafka cluster to provide a fault-tolerant and load-balanced backend for the new Phoebus alarm system. Scripts provided with the Phoebus alarm server take care of correctly creating and configuring the Kafka topics for the alarm system. The alarm tree configuration, consisting of a list of PVs and their related alarm settings for legacy alarms, are automatically compatible with the new Phoebus alarm system, reducing setup efforts to a minimum.

## Adopting the New Alarm System

The new Phoebus alarm system’s client user interface has a very similar look and feel to the legacy alarm client.

We observed a much faster performance with the importing time of the alarm tree configuration to the new Phoebus alarm server in comparison to the legacy BEAST alarm server. There are a few differences to the email/text notification behaviour of the new alarm system which were added to address the issue of duplicate alarm notifications that existed in the legacy alarm system. There are two additional alarm features available with the new Phoebus alarm system for logging the history of alarm states for all PVs and the history of alarm configuration updates. Both of these features have been determined to be effective diagnostic tools for alarm users and system maintainers.

Additional features that were requested by FRIB users to be added to the Phoebus alarm system have also been implemented. These features included a mode to disable email notifications for alarms temporarily. It’s a feature used extensively by FRIB Operations to disable alarm email notifications when operators are present in the FRIB control room to live-monitor and address alarms in person. Alarm email notifications are later re-enabled so that alarms can be monitored remotely and catered to during offline operation hours. In Fig. 1, the yellow mail icon in the alarm table’s toolbar is used to disable and re-enable email notifications for alarm system.




Figure 1: Mode to disable/re-enable email notifications.

The Phoebus alarm system provides an authorization mechanism to allow only selected and authorized users to be able to interact with their alarm server, change its running alarm configuration, or edit the alarm PV tree. As per FRIB’s requirement for running many alarm servers within the same network, the alarm system’s authorization mechanism has been extended to have authorization rules set on a per-alarm-server instance basis. This allows authorized users of an alarm server to be able to interact only with their relevant alarm server and not with alarm servers belonging to other groups or parties. This feature helps to avoid accidental updates an unauthorized user might cause to an alarm server owned and maintained by somebody else. Figure 2 shows the various alarm views supported by the Phoebus alarm system.

\* Work supported by the U.S. Dept. of Energy Office of Science under cooperative Agreement DE-SC0023633



# MULTI-DIMENSIONAL SPECTROGRAM APPLICATION FOR LIVE VISUALIZATION AND MANIPULATION OF LARGE WAVEFORMS

B. E. Bolling\*, A. A. Gorzawski, J. Petersson, European Spallation Source ERIC, Lund, Sweden

## Abstract

The European Spallation Source (ESS) is a research facility under construction aiming to be the world's most powerful pulsed neutron source. It is powered by a complex particle accelerator designed to provide a 2.86 ms long proton pulse at 2 GeV with a repetition rate of 14 Hz. Protons are accelerated via cavity fields through various accelerating structures that are powered by Radio Frequency (RF) power. As the cavity fields may break down due to various reasons, usually post-mortem data of such events contain the information needed regarding the cause. In other events, the underlying cause may have been visible on previous beam pulses before the interlock triggering event.

The Multi-Dimensional Spectrogram Application is designed to be able to collect, manipulate and visualize large waveforms at high repetition rates, with the ESS goal being 14 Hz, for example cavity fields, showing otherwise unnoticed temporary breakdowns that may explain the sometimes-unknown reason for increased power (compensating for those invisible temporary breakdowns). The first physical event that was recorded with the tool was quenching of a superconducting RF cavity in real time in 3D. This paper describes the application developed using Python and the pure-python graphics and GUI library PyQtGraph and PyQt5 with Python-OpenGL bindings.

## INTRODUCTION

The ability to visualize data arrays as a spectrum on computers has developed as computer processing power increased, with many spectrograph software developed both by industry and research institutes. [1]

The first concept of this application was realized by developing a realtime spectrum visualization with a 2D heat map using Python [2] with the pyqtgraph library [3], combined with some other graphical elements via the PyQt5 Python library [4]. To integrate this with a proper handling of data streams (scalars and arrays with timestamps) as well as having the ability to import archived data, the pychiver library was used (see its subsection under Methodology). In-situ methods to manipulate were added via the standard Python NumPy library to e.g. apply a Discrete Fourier Transform on live data arrays, as well as custom-tailored functions within this package.

This paper will focus on the 3D functionality of the application, as the 3D functionalities seems in some cases that will be discussed in this paper to be more useful for visualizing live-streamed data-arrays in comparison with standard 2D waveform plots.

\* benjamin.bolling@ess.eu

## Quench of a SRF Cavity

The idea to begin using a 3D spectrum visualization came whilst doing conditioning of Superconducting Radio Frequency (SRF) cavities and observing the so-called quench phenomena. This occurs in SRF cavities when localized heating exceeds the critical temperature or when the critical superconducting magnetic field is exceeded, which causes the cavity material to lose its superconductivity. With the final day of an SRF cavity's cold conditioning approaching, the initial prototype to be able to visualize a quench in 3D was quickly and successfully implemented. With it, a provoked quench was captured as can be seen in Fig. 1.




Figure 1: SRF quench captured with the first prototype of the application.

Since then, the application has been further developed to increase its reliability and functionality.

## METHODOLOGY

Multiple Python libraries are used to span the Spectrogram application, which can be split into 3 main parts:

- Data retrieval (pychiver)
- GUI (PyQt5)
- Data visualization (pyqtgraph and OpenGL)

## EPICS

Experimental Physics and Industrial Control System (EPICS) is a set of software tools and applications that provide a software infrastructure that can be used to build distributed control systems to operate components in large-scale scientific experiments, such as particle accelerators or telescopes [5].

At ESS, physical components are controlled via EPICS Input-/Output-Controllers (IOCs) that communicate with physical hardware controllers (e.g. programmable logics controllers, PLCs). Each property of the physical component (e.g. applied current or detected pressure levels) has its own so-called process variable, which is processed within

# APPLICATIONS OF ARTIFICIAL INTELLIGENCE IN LASER ACCELERATOR CONTROL SYSTEM

F.N. Li<sup>\*</sup>, Z. Guo<sup>†</sup>, M.X. Zang, Y.D. Xia, Q.Y. He, K. Chen, Q. Wang<sup>1</sup>, C. Lin<sup>1</sup>  
Peking University, Beijing, China

<sup>1</sup>also at Beijing Laser Acceleration Innovation Center, Beijing, China

## Abstract

Ultra-intense laser-plasma interactions can produce TV/m acceleration gradients, making them promising for compact accelerators. Peking University is constructing a proton radiotherapy system prototype based on PW laser accelerators (CLAPA-II). This transient acceleration process becomes more challenging for stability control, which is critical for medical applications. This work demonstrates artificial intelligence's application in laser accelerator control systems.

Laser accelerator requires fast implementation of micro-precision alignment between the ultra-intense laser and the target. We proposed an automated positioning program using the YOLO algorithm. This real-time method employs the convolutional neural network, directly predicting object locations and class probabilities from input images. It enables precise, automatic solid target alignment in about a hundred milliseconds, reducing experimental preparation time. The YOLO algorithm is also integrated into the safety interlocking system for anti-tailing, allowing quick emergency response.

The intelligent control system also enables convenient, accurate beam tuning. We developed high-performance virtual accelerator software using OpenXAL and GPU-accelerated multi-particle beam transport simulations. The software allows real-time or custom parameter simulations and features control interfaces compatible with optimization algorithms. By designing tailored objective functions, the desired beam size and distribution can be achieved in a few iterations.

## INTRODUCTION

Ultra-intense laser interaction with solid targets can produce acceleration gradients up to TV/m [1], which is considered a promising candidate for future compact accelerators. Pulsed proton sources from laser-plasma accelerators (LPA) offer significant application advantages in high-dose-rate tumor radiotherapy [2,3], extreme environmental material irradiation [4,5], and proton radiography [6,7] due to their short time pulse and high peak current intensity. Among various acceleration mechanisms, Target Normal Sheath Acceleration (TNSA) has been extensively studied both theoretically and experimentally [8–10]. This mechanism stands out for simplicity, robustness, and smooth beam profile compared to other mechanisms. Theoretical and experimental studies have shown that the cut-off energy of protons accelerated under the TNSA mechanism is sensitive to the laser intensity  $I_0$  [11,12] following an  $E_{max} \sim I_0^{0.5}$  scaling law [13]. Laser

acceleration experiments have already demonstrated high-quality particle beams with energies approaching nearly 100 MeV [14].

To optimize laser energy deposition on the surface of the targets, the main laser with a center wavelength of  $\lambda$  is focused on a spot with the size of several micrometers, denoted as  $r_0$ , using an off-axis parabolic mirror (OAP). This configuration results in a Rayleigh length at the micrometer scale, as indicated by  $L_R = \pi r_0^2 / \lambda$ . Consequently, precise alignment at the micrometer level between the laser spot and the target surface becomes imperative and optical methods capable of measuring distances are employed.

Traditional automatic alignment depends on multi-step algorithms [15,16] that are time-consuming. The defocus distance is measured from the imaging system by designing a focus measure function. The focus measure function value will reach its extremum when the lens motor is approaching the optimal imaging position. Nowadays, artificial intelligence techniques for computer vision have been successfully used for rapidly processing imaging information [17,18]. This new method ensures highly efficient target positioning. In the first session, we explored a deep learning method to realize rapid automated positioning. We successfully demonstrated that the YOLO [19,20] (You Only Look Once) object detection network enables fast and high-precision automatic positioning. Subsequently, we integrated this deep learning model into the laser accelerator control system. Additionally, YOLO algorithm has also been implanted into the safety interlocking system of CLAPA-II. It can display information on the operator interface and issue a sound alarm to help operators quickly take necessary actions in emergency situations.

In the second session, we delineated the developmental process of virtual accelerator software within laser accelerators, encompassing the development of GPU-accelerated multi-particle beam transport simulation algorithm, the development of software interfaces, and the incorporation of applications utilizing genetic optimization algorithms. The virtual accelerator software has accomplished the bi-directional conversion between physical quantities and control quantities within the control system. This advancement facilitates the more convenient and efficient application of physical and artificial intelligence algorithms to accelerators, laying a foundational framework for the realization of intelligent control in future laser accelerators.

\* fnli@stu.pku.edu.cn

† zhen\_guo0327@stu.pku.edu.cn

# INITIAL TEST OF A SRF CAVITY ACTIVE RESONATE CONTROLLER BASED ON MACHINE LEARNING METHOD\*

Faya Wang<sup>†</sup>, Jorge Cruz, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

## Abstract

We will introduce an active motion controller that leverages machine learning technology and electric piezo actuators. This controller will be specifically designed for active resonance control (ARC) of superconducting radio frequency (SRF) cavities. Our approach involves the initial development of a data-driven model to capture the dynamic behavior of the system, followed by the construction of a model predictive controller (MPC). The accuracy of the model has been validated through real cavity testing, and we are currently in the process of implementing the MPC on the existing hardware of the LCLS-II LLRF system. In our paper, we will present simulation results along with preliminary test results using actual SRF cavities from the LCLS-II SRF linac.

## INTRODUCTION

Motion control plays an increasingly critical role in modern large accelerator facilities, such as 4th generation storage ring-based light sources, SRF accelerators, and high-performance photon beamlines. In the case of very high-Q SRF linacs like LCLS-II, precise cavity resonance control is essential to maintain stable operations. Failure to do so would necessitate a substantial increase in RF power, resulting in higher operational and capital costs due to the need for additional RF power sources.

The complexity of motion control in accelerator systems arises from the intricate interplay between the beam and electromagnetic fields with mechanical energy. For instance, in SRF cavities, electromagnetic modes are closely linked with their mechanical counterparts through Lorentz-force detuning and external microphonics. Due to the non-linear nature of these couplings, resonance control, especially for SRF cavities, poses significant challenges.

Recent developments in ARC have shown promise in mitigating microphonics-induced detuning by manipulating piezo tuners. Notable examples of model-based controllers have been demonstrated at facilities like CBETA[1], Fermilab[2], DESY[3], and more recently at SLAC[4].

In this paper, we will introduce a data-driven ARC approach that leverages machine learning, specifically based on the Dynamic Model Decomposition (DMD) method. Our discussion will begin with modelling of cavity dynamics in the presence of microphonics by DMD, and will include the simulations of MPC performance. We'll then

introduce the initial testing with real cavities. Finally, we will draw conclusions based on our study.

## SIMULATION OF CAVITY RESONANCE CONTROL BY DMD

To construct an accurate data-driven model for SRF cavity ARC, we employ an equivalent circuit model to simulate cavity dynamics in the presence of microphonics, as depicted in Fig. 1. where it has  $\omega_0 = \frac{1}{\sqrt{LC}}$ ,  $\frac{R}{Q} = \sqrt{\frac{L}{C}}$ , cavity half bandwidth  $\omega_{1/2} = \frac{\omega_0}{2Q_L}$ , and  $Q_L$  cavity loaded Q factor.




Figure 1: The equivalent circuit model of a cavity.

The transient behavior of the cavity can be described by the following equation:

$$\frac{dV}{dt} = \begin{bmatrix} -\omega_{1/2} & -\Delta\omega \\ \Delta\omega & -\omega_{1/2} \end{bmatrix} V + \begin{bmatrix} \omega_{1/2} & 0 \\ 0 & -\omega_{1/2} \end{bmatrix} V_g, \quad (1)$$

where  $V = [V_r V_i]$  is cavity voltage envelop and  $V_g = [V_{gr} V_{gi}]$  is the generator voltage envelop.

The variation in the resonance frequency of the SRF cavity arises from deformations in the cavity body, initiated by fluctuations in helium pressure or the Lorentz force pressure due to the electromagnetic field within the cavity. The coupling between cavity deformation and detuning can be described using a single mechanical model [5], as follows:

$$\frac{d}{dt} \begin{bmatrix} \Delta\omega_m \\ \Delta\dot{\omega}_m \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\omega_m^2 & -\frac{\omega_m}{Q_m} \end{bmatrix} \begin{bmatrix} \Delta\omega_m \\ \Delta\dot{\omega}_m \end{bmatrix} + \begin{bmatrix} 0 \\ -K_m \omega_m^2 \end{bmatrix} E_{cav}^2, \quad (2)$$

where  $\omega_m$ ,  $Q_m$  and  $K_m$  are respectively the mechanical mode's natural frequency, quality factor and coupling factor. The total detuning of a cavity is the sum of the contributions of all mechanical models, denoted as  $\Delta\omega = \Sigma \Delta\omega_m$ .

\* Work supported by the US Department of Energy, Laboratory Directed Research and Development Program at SLAC National Accelerator Laboratory, under Contract No. DE-AC02-76SF00515

<sup>†</sup> fywang@slac.stanford.edu

# ENHANCING ELECTRONIC LOGBOOKS USING MACHINE LEARNING

Jennefer Maldonado\*, Samuel Clark, Wenge Fu, Seth Nemesure  
Brookhaven National Laboratory, Upton, New York, United States

## Abstract

The electronic logbook (elog) system used at Brookhaven National Laboratory's Collider-Accelerator Department (CAD) allows users to customize logbook settings, including specification of favorite logbooks. Using machine learning techniques, customizations can be further personalized to provide users with a view of entries that match their specific interests. We will utilize natural language processing (NLP), optical character recognition (OCR), and topic models to augment the elog system. NLP techniques will be used to process and classify text entries. To analyze entries including images with text, such as screenshots of controls system applications, we will apply OCR. Topic models will generate entry recommendations that will be compared to previously tested language processing models. We will develop a command line interface tool to ease automation of NLP tasks in the controls system and create a web interface to test entry recommendations. This technique will create recommendations for each user, providing custom sets of entries and possibly eliminate the need for manual searching.

## INTRODUCTION

The electronic logbook (elog) system is used to record information related to machine and system operations as well as individual record keeping. Applications in the controls system send data, plots, and images to be uploaded. The system is shown in Fig. 1. Engineer and physicists often use the elog to document procedures and instructions. The more documentation we have available, the easier it is to diagnose new problems limiting down time and delays in science. The search feature in the system only provides entries with the exact search term the user enters. There are entries related to this search term that may not contain the term exactly. Natural language processing models can aid in this search process by analyzing similarity in entries and classifying them by topic group. If a user is interested in power supply failures and search the term "power supply", entries about magnet quenches without the search term will not be featured but are directly related. Analyzing and producing similarity metrics will help specific groups of people like operators, controls staff, and physicists narrow down entries of interest. The goal is to provide accurate results in order to improve productivity of users.

## DATA PROCESSING

As users enter entries in the system, each entry is stored in a MySQL database. This enables storage of information dating back to 2013 when the system was first implemented. To access this stored data we use a tool developed to connect

\* jmaldonad@bnl.gov

to our departmental databases. Along with each entry we have some meta data like the entry id, timestamp, author, and tag. This data is collected and stored in a Pandas dataframe. This dataframe is then processed to remove rows with empty

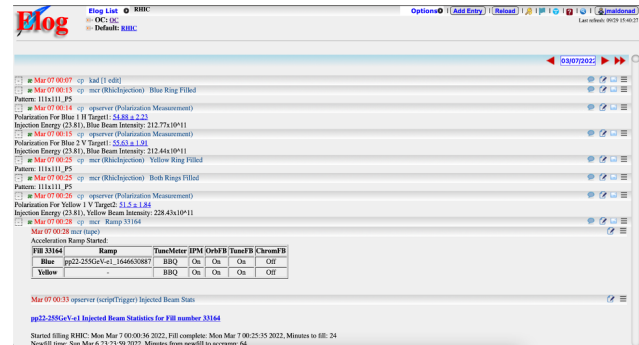


Figure 1: Screenshot of elog interface.

content. The empty rows correspond to entries with no text, only include images, or have been deleted. After processing, all rows in the dataframe contain content that need to be formatted in a language model interpretable format. First the entries are tokenized. The goal of tokenizing text is to separate the sentences into a list of words the model can utilize better for training purposes. After tokenizing, lemmatization is applied to convert all words in inflected forms to the dictionary form. For example, the word log has variations such as "logs", "logging", "logged". Each inflected form found in the dataframe is transformed into the dictionary form of the word. In our example the dictionary form is log. To ensure the model focuses on collider accelerator topics the most common words in the English language are removed from the data. Examples of these words are "the", "in", "go", and "had". If these words are not removed from text there is a risk the model will correlate sentences containing common words rather than ones related to the elog entry topics. Punctuation is also removed from the text vectors. The histograms in Fig. 2 display the total word counts of the most used words in elog entries before and after the common words have been removed. Once the contents of entries are processed, we can use Gensim's Doc2Vec (D2V) model to predict similar entries.

## NATURAL LANGUAGE MODELS

The Gensim package is a fast library for training large NLP models [1]. This package makes it easy to load a model and build a vocabulary from processed elog entry vectors. The model trained for approximately 1.5 hours on nearly a decades worth of entry data. This was done on a machine with a GPU and the total number of entries was about 1.5 million. Only 100 epochs were used for each prediction.

## RESEARCH AND DEVELOPMENT OF THE FAST ORBIT FEEDBACK SYSTEM FOR HEPS

P. Zhu<sup>†1,2,3</sup>, D. P. Jin<sup>1,2</sup>, Z. X. Xie<sup>1,2</sup>, Z. Lei<sup>1,2</sup>, Y. L. Zhang<sup>1,2</sup>, Y. C. He<sup>1,2</sup>, D. Y. Wang<sup>1,2</sup>

<sup>1</sup>Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China

<sup>2</sup>Spallation Neutron Source Science Center, Dongguan, China

<sup>3</sup>National Synchrotron Radiation Laboratory,  
 University of Science and Technology of China, Hefei, China

### Abstract

As a 4th-generation light source, High Energy Photon Source (HEPS) has much more stringent requirements to the beam orbit stability in both horizontal and vertical directions than the previous sources due to the much smaller beam sizes. A Fast Orbit Feedback (FOFB) system, with the closed-loop bandwidth around 500 Hz, is needed to meet the critical requirements. The latency of the FOFB system is the key to achieve these requirements.

This paper focuses on the design and implementation of the FOFB system. Based on the architecture of ATCA (Advanced Telecom Computing Architecture) standard, a total of 16 sub-stations are adopt to set up with bidirectional daisy-chained structure. The transmission between a sub-station with belong to BPMs or local fast correctors are connected in a "point-to-point" high-speed links, which is to minimize transmission delays and improve the system's closed-loop bandwidth. Meanwhile, it is optimized to calculate the large matrix based on the singular value decomposition (SVD) taking the digital signal processing (DSP) modules of V7 field programmable gate array (FPGA) with parallel pipeline. For performance tuning and additional flexibility, we implement an independent ethernet controller for remote operation using a full hardwired TCP / IP providing internet connectivity by using Serial Peripheral Interface (SPI), such as download a new matrix. In comparison to other standards, this architecture is advantageous for simplifying the hardware design of FOFB, improving system reliability, availability, maintainability, and scalability.

### INTRODUCTION

High Energy Photon Source (HEPS), a fourth generation of synchrotron radiation sources, is one of the major national scientific and technological infrastructures, as show in Fig. 1.




Figure 1: Layout of HEPS.

The size of modern light sources is getting smaller and smaller, and the brightness is getting higher and higher [1]. For HEPS, the brightest fourth-generation synchrotron radiation source in the world, it has higher requirements for the stability of the beam trajectory. The beam in the storage ring should have less than 10% jitter in both the horizontal and vertical directions (RMS).

To meet these stringent requirements, and also unlike most 3rd-generation light sources, such as SSRF, diamond and SLS, the fast orbit feedback system of HEPS is designed with a new architecture based on FPGA with Rocket IOs and built-in DSPs instead of real-time data transmission links and VME controllers with DSP boards. The response delay of the fast orbit feedback system should be minimized as much as possible, and the closed-loop bandwidth is designed to be 500 Hz. This new architecture is also adopted by latest designed and upgraded light sources, such as NSLS-II and APS-U [2].

### REQUIREMENT

There are many factors that affect the stability of the beam orbit, including the stability of the magnet power supply, ground vibration, temperature effects, etc. In order to suppress interference and keep the beam orbit stable, we must to adopt a high-intensity and high-speed orbit feedback system, a typical multiple-input and multiple-output (MIMO) system, to achieve long-term stable operation of the light source based on singular value decomposition (SVD) commonly, as shown in Fig. 2.

The correction algorithm is based on an SVD of the orbit response matrix:

$$\Delta \bar{X} = R \Delta \bar{\theta} \quad \text{and} \quad R = USV^T \quad (1)$$

$$\Delta \bar{\theta} = VS^{(-1)}U^T \Delta \bar{X}. \quad (2)$$

Where  $\Delta \bar{X}$  is error that the current orbit is compared to the golden orbit, U and V are matrices whose columns form an orthogonal basis in BPM(X) and corrector magnet  $\theta$  space, S is the diagonal matrix of singular values.

The proportional-integrator (PI) control algorithm operates in this diagonal space and the relevant parameters can be adjusted for each mode separately.

<sup>†</sup> zhup@ihep.ac.cn

# UPGRADE OF THE AGOR CYCLOTRON CONTROL SYSTEM AT UMCG-PARTREC

O. J. Kuiken<sup>†</sup>, J. Schwab<sup>‡</sup>, J. K. van Abbema, P. Schakel, A. Gerbershagen  
Particle Therapy Research Center (PARTREC),  
University Medical Center Groningen, University of Groningen, Groningen, Netherlands

## Abstract

The transfer of the AGOR cyclotron to the University Medical Center Groningen (UMCG) initiated investments for long-term reliability. Recent, current and upcoming upgrades and additions include replacing the Operational Technology (OT) network, control system software, PLCs and the RF resonator control. Several in-house developed electronic devices are also undergoing updates. Simultaneously, a new beamline is under construction. However, these investments also come with challenges such as keeping the cyclotron operational throughout the upgrade process, managing limited manpower, and mitigating supply chain disruptions exacerbated by COVID-19. Despite these hurdles, UMCG-PARTREC is committed to proactive investment in the cyclotron's future, securing its reliability for ongoing scientific research.

## INTRODUCTION

The superconducting AGOR cyclotron (see Fig. 1) began development in the late 1980s and was commissioned in 1997. In 2020, when the institute KVI was transferred from the University of Groningen (RUG) to the University Medical Center Groningen (UMCG) and became PARTREC [1], it marked the beginning of an upgrade process aimed at ensuring reliable operation for the foreseeable future. Through proactive investments, we are modernizing components that may be up to 30 years old to meet current standards, with the goal of averting potential reliability issues in the future. In this paper, we will describe some of the upgrades and challenges faced at PARTREC.



Figure 1: The AGOR cyclotron.

<sup>†</sup> o.j.kuiken@umcg.nl  
<sup>‡</sup> j.schwab@umcg.nl

## UPGRADES

The facility's recent, current, and forthcoming upgrades and additions encompass the following:

### Operational Technology (OT) Network

The current OT network is made up of controllers and I/O modules based on the Bitbus fieldbus, linked via Bitbus servers. Bitbus is considered to be a technology soon to be phased out, hence a pilot study was conducted to evaluate the feasibility of using a National Instruments CompactRIO (NI-cRIO) based subrack in place of Bitbus-based controllers for analog and digital I/O (see Fig. 2). Additionally, a similar PLC-based solution is currently under investigation. Both of these alternatives are able to accommodate our in-house developed Digital InterFace card (DIF) and InterFace for Analog card (IFA). The DIF cards provide current buffering and isolation and the IFA card provides analog output voltage drop compensation due to long cabling. They have the same type of connector receptacle as the equivalent Bitbus cards which enables the reuse of cables for existing equipment when desired, making it a drop-in replacement for our Bitbus based I/O cards.



Figure 2: A photograph of a pilot subrack for the replacement of Bitbus based controllers using an NI-cRIO controller and custom-made DIF and IFA cards.

For compatibility reasons, both pilot projects employ an Ethernet interface with Modbus TCP to connect to the servers hosting the control system software. Nevertheless, if we opt to update our control system software as described in the subsequent subsection, we are likely to transition to the OPC UA communication protocol. This shift is motivated by the observation that OPC UA appears to offer advantages in data throughput when compared to Modbus TCP.

### Control System Software

Our current control system software, Vsystem by Vista Control Systems, Inc. [2], though still manufacturer-supported, suffers from persistent instability issues,

# MAX IV LABORATORY'S CONTROL SYSTEM EVOLUTION AND FUTURE STRATEGIES

V. Hardion, P. Bell, T. Eriksson, M. Lindberg, P. Sjöblom, D. Spruce  
MAX IV Laboratory, Lund University, Lund, Sweden

## Abstract

The MAX IV Laboratory, a 4th generation synchrotron radiation facility located in southern Sweden, has been operational since 2016. With multiple beamlines and experimental stations completed and in steady use, the facility is now approaching the third phase of development, which includes the final two of the 16 planned beamlines in user operation. The focus is on achieving operational excellence by optimizing reliability and performance. Meanwhile, the strategy for the coming years is driven by the need to accommodate a growing user base, exploring the possibility of operating a Soft X-ray Laser (SXL), and achieving the diffraction limit for 10 keV of the 3 GeV storage ring.

The Technical Division is responsible for the control and computing systems of the entire laboratory. This new organization provides a coherent strategy and a clear vision, with the ultimate goal of enabling science. The increasing demand for more precise and efficient control systems has led to significant developments and maintenance efforts. Pushing the limits in remote access, data generation, time-resolved and fly-scan experiments, and beam stability requires the proper alignment of technology in IT infrastructure, electronics, software, data analysis, and management.

This article discusses the motivation behind the updates, emphasizing the expansion of the control system's capabilities and reliability. Lastly, the technological strategy will be presented to keep pace with the rapidly evolving technology landscape, ensuring that MAX IV is prepared for its next major upgrade.

## GENERAL STATUS

As of now, the MAX IV Laboratory boasts 15 operational beamlines, with a 16th set to be commissioned by this end of 2023. Discussions regarding the construction of additional beamlines are ongoing with stakeholders, but no firm decisions have been made. Remarkably, the facility's accelerators maintain an impressive availability rate of 98.5% for the 1.5 GeV ring, 98.2% for the 3 GeV and 96.8% for the linear accelerator, although the lab is committed to increase reliability to reach ideally the 99% range.

## Accelerators

In a particular effort toward the linear accelerator, the beam diagnostics have received an important upgrade. A Transverse Deflecting cavity (TDC) [1] conveys a perpendicular oscillation to the particle beam, acting as a diagnostic tool. By converting time differences within a beam pulse into spatial differences, it offers detailed insight into the pulse's temporal structure. Often likened to a "streak camera", the

TDC spreads the beam pulse across a spatial dimension for high-resolution examination. Beyond diagnostics, it also manipulates the beam for specific experiments, necessitating precise synchronization with the accelerator's components.

On the other hand, the development related to the Soft X-ray Free Electron Laser (SXFEL) is on-going. The Compact APPLE X undulator is an advanced device tailored for Soft X-ray beamlines. Not only is it more cost-efficient than its predecessor, the APPLE II, but it also boasts the ability to consistently control polarization, maintaining a uniform effective K-value across different polarization states. This feature positions it as the go-to source for the upcoming SXFEL at MAX IV. Furthermore, it can expand the energy spectrum below 2 keV for the FemtoMAX beamline at the Short Pulse Facility (SPF). When benchmarked against the APPLE II undulators at the 3 GeV MAX IV ring, it provides a broader energy range with complete polarization adjustability.

## Beamlines

On the Beamlines, a great effort has been made to reach the full capacity of the MAX IV Laboratory within the initial scope of 16 beamlines. Apart from the 2 last beamlines described below, all the beamlines are continuously improving their performance i.e an EXAFS measurement is done now in 3s at Balder, in contract to 30s in 2021.

ForMAX is a new beamline in operation tailored for in-depth research on tree materials. It enables multi-scale structural analysis ranging from nanometers to millimeters by seamlessly integrating tomographic imaging with small- and wide-angle X-ray scattering (SWAXS), as well as scanning SWAXS imaging. Operating between 8-25 keV, the beam size varies from approximately 1  $\mu\text{m}$  to 5mm based on the chosen mode.

MicroMAX, set to launch in the end of 2023, will revolutionize structural biology research. This beamline will facilitate 3D protein studies and their real-time analysis, particularly focusing on the challenging molecules that produce only microcrystals. Operating between 5-25 keV, its beam size varies from around 1  $\mu\text{m}$  to 5mm, boasting a photon flux of  $10^{15}$  photons/s based on the mode selected. An added capability of MicroMAX is its time-resolved mode, offering dynamic insights.

## TECHNICAL DIVISION

To ensure the sustainability of the technical support, MAX IV has adopted a new organization with the new Technical Division bringing together multiple technical disciplines, including mechanical engineering, electrical engineering,

## INTRODUCTION AND STATUS OF FERMILAB'S ACORN PROJECT\*

D. Finstrom<sup>†</sup>, E. Gottschalk, Fermilab, Batavia, USA

### Abstract

Modernizing the Fermilab accelerator control system is essential to future operations of the laboratory's accelerator complex. The existing control system has evolved over four decades and uses hardware that is no longer available and software that uses obsolete frameworks. The Accelerator Controls Operations Research Network (ACORN) Project will modernize the control system and replace end-of-life power supplies to enable future accelerator complex operations with megawatt particle beams. An overview of the ACORN Project and a summary of recent research and development activities will be presented.

### INTRODUCTION

The Fermilab Accelerator Complex is the largest accelerator complex in the United States and the second largest in the world. It currently operates with a single control system, ACNET [1], that contains hardware and software for controlling ten miles of accelerator components and beam transfer lines. The control system was originally developed for the start of colliding-beam operations in 1983. The control system initiates particle beam production, controls beam energy and intensity, transports particle beams to research facilities, measures beam parameters, and monitors beam transport through the accelerator complex to ensure safe, reliable, and effective operations. There are approximately 200,000 devices with 350,000 attributes and several million lines of software code in the existing system. Despite operating at peak performance, in 2018, an advisory committee formed to evaluate facility operational risks identified the need to invest in the accelerator control system and called out major issues, including a large amount of old hardware and software and an aging and declining in strength workforce with no software development related hires for 18 years. The Accelerator Controls Operations Research Network (ACORN) Project addresses these concerns and will replace control system hardware that is no longer available and software that is no longer maintainable. ACORN is a U.S. Department of Energy (DOE) Project for Acquisition of Capital Assets (O413.3B) to modernize the accelerator control system and replace end-of-life power supplies to enable future operations of the accelerator complex with megawatt particle beams. The total project cost range is \$100 - \$142 million dollars USD. The modernized control system will integrate with two projects, the Long Baseline Neutrino Facility/Deep Underground Neutrino Experiment (LBNF/DUNE) and the Proton Improvement Plan II

(PIP-II). ACORN received Critical Decision 0 (CD-0) milestone from DOE, which signifies approval of the project's mission need, on August 28, 2020. ACORN is working towards Critical Decision 1 (CD-1) milestone, which signifies approval of the alternative selection and cost range and is targeted to occur in Q3 FY24. As required by CD-1, the project is evaluating design alternatives and, consequently, has not settled on a technical design. Efforts are focused on building the project team, interviewing stakeholders, gathering requirements, and research and development efforts aimed at determining project cost and schedule estimates. Critical Decision 4 (CD-4) milestone, which signifies project completion, is projected to occur in the 2028 - 2030 time frame. The project Work Breakdown Structure consists of five main sections: Project Management, Accelerator Power Systems, Data Acquisition and Control, Control System Infrastructure, and Control System Applications.

### ACCELERATOR POWER SYSTEMS

The scope of work for the Accelerator Power Systems (APS) section includes the design, prototype, procurement, installation, and testing of power supplies and regulation systems for the Fermilab Accelerator Complex. Two types of power supplies are being considered: switch-mode power supplies and thyristor power supplies. The existing power supplies are almost all original and date back to the 1970s and 1980s. The power supplies were originally bid and built to specification and have a straightforward design and rugged construction. Fermilab has continued to buy power supplies using a combination of procurement strategies consisting of build-to-specification procurements for power components and build-to-print procurements for control components. Following the guidance of subject matter experts and our experience at Fermilab, we plan to separate the procurement of power conversion components from components needed for power regulation and control. The power conversion components are available from commercial vendors, and competition among vendors keeps costs reasonable. Power regulation components are highly specialized and can be applied to different vendors. This helps establish commonality among power supply systems and allows the replacement of power converters when requirements change.

Hiring expertise in accelerator power systems has been challenging, and consequently, progress in this section has been constrained. ACORN is working with Fermilab's Electrical Engineering Support department to understand the hazards associated with the installation, testing, and maintenance of power supplies and working with the ACORN Environment, Safety, and Health liaison to begin writing the preliminary hazard analysis report for CD-1. The project team is compiling an inventory of power

\* This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.  
<sup>†</sup> finstrom@fnal.gov



## SOLEIL II: TOWARDS A MAJOR TRANSFORMATION OF THE FACILITY

Y-M. Abiven<sup>†</sup>, B. Gagey<sup>‡</sup>, on behalf of the SOLEIL II technical details studies (TDR) team  
Synchrotron SOLEIL, Saint Aubin, France

### Abstract

Operational since 2008, SOLEIL provides users with access to a wide range of experimental techniques thanks to its 29 beamlines, covering a broad energy range from THz to hard X-rays. In response to new scientific and societal challenges, SOLEIL is undergoing a major transformation with the ongoing SOLEIL II project. This project includes the design of an ambitious Diffraction Limited Storage Ring (DLSR) to increase performances in terms of brilliance, coherence, and flux, upgrading the beamlines to provide advanced methods, and driving a digital transformation in data- and user-oriented approaches. This paper presents the project organization and technical details studies (TDR) for the ongoing upgrades, with a focus on the digital transformation required to address future scientific challenges. It will depict the computing and data management program with the presentation of the targeted IT architecture to improve automated and data-driven processes for optimizing instrumentation. The optimization program covers the facility reconstruction period as well as future operation, including the use of Artificial Intelligence (AI) techniques for data production management, decision-making, complex feedback systems, and data processing. Real-time processes will be applied in the acquisition scanning design, where detectors and robotic systems will be coupled to optimize beam time.

### SOLEIL II PROJECT ORGANISATION

The SOLEIL II project [1, 2] was initiated after the Conceptual Design Report (CDR) phase. Since 2021, the project has progressed to the Technical Design Report phase (TDR), which is structured into four distinct programs: one dedicated to beamlines and laboratories (referred to as BL<sup>2</sup>), another focused on accelerators, a third program dedicated to infrastructure, and an additional program encompassing Information Systems (IS) and data management. The latter program serves as a cross-functional support system for the other programs.

Currently, the project is targeting an 18-month shutdown to occur mid-2028.

The overall project is divided into two five years phases, each. The first one, known as the “Construction” phase encompasses the realization of the accelerators, the necessary adaptations of some beamlines, and the associated infrastructure upgrades. This phase includes the dark period and the commissioning of the new booster and the storage ring.

The second phase, “Towards Full Performance”, begins with the continuation of the storage ring commissioning and the initial commissioning of the beamlines. The

upgraded state-of-the-art beamlines will be significantly enhanced by the performance of the new accelerators. These improvements will allow experiments to take full advantage of the photon beams generated by the ultra-low emittance electron beam circulating in bending magnets and innovative insertion devices.

### THE SCIENTIFIC CHALLENGE

The scientific motivation behind this upgrade is to provide new tools and techniques to address the societal and technical challenges our society is facing. Based on the CDR findings [3], we have structured four main science use cases, which are listed here with their main benefits:

1. **Advanced Materials:** material engineering, quantum materials, information technologies
2. **Health and Well-being:** new pathogens, antibiotic resistance.
3. **Sustainable Energy Development:** batteries, catalysis, and green chemistry.
4. **Environment:** impact of pollutants and contributing to our understanding of global warming.

These use cases serve as a framework for shaping future instrumentation and methods. By conducting a comprehensive cross-cut review of scientific topics and drawing upon the common methods available across existing beamlines, our goal is to progressively tailor instruments to meet the needs of the scientific community. This entails achieving nanometric resolution scans/positioning, which will become the norm across most beamlines.

Moreover, we are working towards enabling operando analysis at millisecond or sub-millisecond timescales, made possible by enhancements in brilliance, coherence, and flux.

Imaging techniques will benefit highly and naturally from the lower electron emittance. To leverage on this gain, we are designing a multimodal and multi-technique approach to maximize the efficiency and flexibility of the SOLEIL II experimental stations. This strategy includes addressing the challenge of managing the growing data deluge from production to end user analysis, a concern shared and collaboratively addressed by both the computing and scientific teams.

### THE ACCELERATORS UPGRADE [4]

To establish this vital instrumental hub, a substantial upgrade of the accelerators is imperative. In this context, the new storage ring will completely replace the current 354-meter circumference ring. It will be based on a reference lattice featuring 20 straight sections alternating between 7BA and 4BA Higher-Order Achromat (HOA) cells, all housed in the same tunnel as the current setup. The primary

<sup>†</sup> yves-marie.abiven@synchrotron-soleil.fr

<sup>‡</sup> brigitte.gagey@synchrotron-soleil.fr

# INTEGRATION OF AN MPSoC-BASED ACQUISITION SYSTEM INTO THE CERN CONTROL SYSTEM

E. Balci\*, I. Degl'Innocenti, M. Gonzalez Berges, S. Jackson, M. Krupa  
CERN, Geneva, Switzerland

## Abstract

Future generations of Beam Instrumentation systems will be based on Multiprocessor System on Chip (MPSoC) technology. This new architecture will allow enhanced exploitation of instrumentation signals from CERN's accelerator complex, and has thus been chosen as the next platform for several emerging systems. One of these systems, for the HL-LHC BPM (High-Luminosity LHC Beam Position Monitors), is currently at a prototyping stage, and it is planned to test this prototype with signals from real monitors in CERN's accelerators during 2023. In order to facilitate the analysis of the prototype's performance, a strategy to integrate the setting, control and data acquisition within CERN's accelerator control system has been developed.

This paper describes the exploration of various options and eventual choices to achieve a functional system, covering all aspects from data acquisition from the gateway, through to eventual logging on the accelerator logging database. It also describes how the experiences of integrating this prototype will influence future common strategies within the accelerator sector, highlighting how specific problems were addressed, and quantifying the performance we can eventually expect in the final MPSoC-based systems.

## INTRODUCTION

To gain experience with System-on-Chip (SoC) technology and validate its suitability, a vertical slice of the new HL-LHC BPM system was developed. This vertical slice was designed to address the following requirements:

- configuration of the acquisition parameters;
- signal processing in alignment with the acquisition parameters;
- the ability to trigger acquisitions based on accelerator timing events, manual triggers, and timers;
- storage and display of both the acquisition data and system settings.

The chosen hardware is the development board ZCU208 [1], which includes as main component the Xilinx Zynq UltraScale+ RFSoc Gen3 [2]. The RFSoc is a System-on-Chip integrating a Processing System (PS), with Programmable Logic (PL) and high-resolution fast Analog to Digital Converters (ADC) and Digital to Analog Converters (DAC). The PS includes an Application Processing Unit (APU) with 4 Arm Cortex-A53 CPUs and a Real-Time Processing Unit (RPU) with 2 Arm Cortex-R5 CPUs. The PL includes UltraScale+ field programmable gate array (FPGA) fabric, which can be configured through hardware description language

to implement fast parallel signal processing, and it is fundamental in this chip to interface the fast RF converters data stream. The RFSoc was identified as potential candidate for the HL-LHC Beam Position Monitor (BPM) electronics [3]. With the requirements listed above in mind, various components available on the RFSoc are harnessed. The vertical slice consists of the signal processing and acquisition logic within the PL, communication software in the PS, and a Front End Server Architecture (FESA) device running on a Front End Computer (FEC) [4].

In addition, we implemented the necessary communication protocols and conducted lab tests and tests in the Super Proton Synchrotron (SPS) accelerator in order to verify the prototype and quantify its performance. The decisions regarding communication protocols have been made with the understanding that this is a prototype, and this selection may not be final. A new iteration of the prototype, incorporating future recommendations from the sector and group, is planned. The current state of the prototype will be further discussed in the following subsections.

## SYSTEM OVERVIEW

Figure 1 provides an overview of the system, highlighting its software and gateway components. The analog signals from BPMs are directed to the ADCs on the RFSoc board. The processing of these signals and the implementation of acquisition logic are carried out within the PL, as detailed in the gateway subsection.

A Linux based Operating System(OS) is prepared with the necessary modifications to the device tree in order to accommodate the interactions with the PS DDR memory, utilizing the Petalinux Tools. It is deployed to the APU of the PS. Decoupling of the operating system and the PL configuration is achieved through configuring the PL while the OS is running, through a start script added as a user application. Which is important in order to ensure that a time consuming task such as building an OS image is avoided for each change in the PL configuration.

Configuration of acquisition parameters and data transfer are managed within the APU. The integration of acquisition data and parameters into the CERN control system, as well as the logging of data, are achieved through a FESA device, which is further elaborated upon in its respective subsection.

### Gateway

The gateway consists of the RF ADC cores, the logic managing the data stream, the logic selecting the samples to be stored (depending on the acquisition settings), and the state machines controlling the acquisition operations

\* elif.balci@cern.ch

# DEVELOPMENT AND NEW PERSPECTIVES ON THE LMJ POWER CONDITIONING MODULES

Pascal Torrent, Jean Philippe Airiau, Irwin Issury  
CEA CESTA, Le Barp, France

## Abstract

The Laser MegaJoule (LMJ), the French 176-beam laser facility, is located at the CEA CESTA Laboratory near Bordeaux (France). It is designed to deliver about 1.4 MJ of energy on targets, for high energy density physics experiments, including fusion experiments. The first bundle of 8-beams was commissioned in October 2014. By the end of 2023, 15 bundles of 8-beams are expected to be fully operational.

In this paper, we will present:

- The LMJ Power Conditioning Modules
- The unexpected security shutdown of the PCM during experiments
- The software solution against the security shutdown.

## INTRODUCTION

The laser Megajoule (LMJ) facility, developed by the “Commissariat à l’Energie Atomique et aux Energies Alternatives” (CEA) [1], is designed to provide the experimental capabilities to study High Energy Density Physics (HEDP). The LMJ is a keystone of the Simulation Program, which combines improvement of physics models, high performance numerical simulation, and experimental validation, in order to guarantee the safety and the reliability of French deterrent nuclear weapons. When completed, the LMJ will deliver a total energy of 1.4 MJ of 0.35  $\mu\text{m}$  (3 $\omega$ ) light and a maximum power of 400 TW.

The LMJ is sized to accommodate 176 beams grouped into 22 bundles of 8 beams. These will be located in four laser bays arranged on both sides of the central target bay of 60-meter diameter and 40-meter height. The target chamber and the associated equipment are located in the center of the target bay.

The first bundle of eight beams has been commissioned at the end of 2014. The second bundle has been commissioned at the end of 2016 following the same commissioning process. Eleven additional bundles are now operational since the end of 2022, and the last fusion experiment using 80 operational beams took place in spring 2022.

The PETAL project consists in the addition of one short-pulse (0.5 to 10 ps) ultra-high-power (1 up to 7 PW), high-energy beam (1 up to 3.5 kJ) to the LMJ facility. PETAL offers a combination of a very high intensity petawatt beam, synchronized with the nanosecond beams of the LMJ [2].

The first phase of nuclear commissioning of LMJ has been achieved to take into account high-energy particles created by PETAL, and neutron production from DD fusion reaction [3]. A subsequent phase will take into account DT targets.

This paper describes the LMJ Power Conditioning Modules (PCM), the issue of security shutdown when there were more than 10 bundles during the LMJ-PETAL shots. And it will describe data tools and software patch to solve the problem.

## LMJ CONTROL SYSTEM

The LMJ facility has a Control System which is divided into 4 layers as shown in Fig. 1.

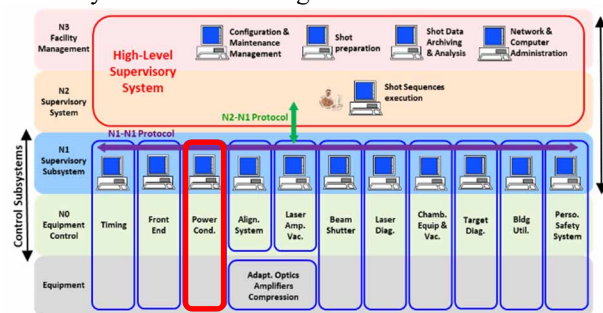


Figure 1: LMJ Control System architecture.

All control system software developed for the supervisory layers use a common framework based on the commercial SCADA Software PANORAMA E2.

In this framework the facility is represented as a hierarchy of objects called “Resources”. Resources represent devices (motors, instruments, diagnostics...) or high level functions (alignment, laser diagnostics). Resources are linked together through different kinds of relationships (composition, dependency, and incompatibility) and the resources life-cycle is described through states-charts. Control points, alarms, states and functions can be attached to any resource.

Dedicated mechanisms manage the resource reservation and propagate properties and states changes into the tree of resources through relationships. There are about 200 000 resources in order to describe the entire LMJ.

## PCM FOR LASER AMPLIFICATION

The laser pulse amplification is obtained by feeding it with photons from Neodym doped glass slab. Flash lamps powered by PCM stimulate the amplifying slab. The power pulse is provided by the PCM capacitors charged to 22 kV and this power is transferred to the flash lamps on the synchronisation triggering order. Because of material security specifications, at the end of the 22 kV charge, the PCM capacitors return to ground voltage after 15 seconds, by the security shutdown if the triggering order is not received.

# A NEW REAL-TIME PROCESSING PLATFORM FOR ELETTRA 2.0 STORAGE RING

G. Gaio<sup>†</sup>, A. Bogani, M. Cautero, L. Pivetta, G. Scalamera, I. Trovarelli  
Elettra-Sincrotrone Trieste, Trieste, Italy  
L. Anastasio, University of L'Aquila, L'Aquila, Italy

## Abstract

Processing synchronous data is essential to implement efficient control schemes. A new framework based on Linux and DPDK will be used to acquire and process sensors and actuators at very high repetition rate for Elettra 2.0. As part of the ongoing project, the current fast orbit feedback subsystem is going to be re-implemented with this new technology. Moreover, the communication performance with the new power converters for the new storage ring is presented.

## INTRODUCTION

At Elettra, a novel hardware/software platform [1] has been embraced to facilitate real-time interfacing and processing for devices equipped with Ethernet interfaces. This architectural paradigm, based on DPDK [2], an industrial-grade package primarily implementing network stack bypass techniques, was deployed last year in the FERMI free electron laser [3] as part of a comprehensive control system upgrade. Furthermore, it has recently been employed in the upgrade of the legacy Global Orbit Feedback (GOF) [4] for the Elettra synchrotron.

While this upgrade may not be deemed strictly necessary, it serves as an ideal testbed for the evaluation of hardware and software technologies over the next two years. These technologies are essential for the real-time control of the equipment slated for installation in the forthcoming particle accelerator, scheduled to replace the Elettra synchrotron in 2026.

## ELETTRA 2.0

The most critical devices installed in the new Elettra 2.0 accelerator will feature a dual Ethernet connection to the control system: the first one dedicated to configuration and supervision, the second for real-time control.

To justify the additional cost of a real-time connection, the communication performance in terms of latency and jitter with these devices must be at least a couple of orders of magnitude better than a mildly loaded conventional control system (10 ms), e.g. greater than 10 kHz. Although systems managing high-speed data do not improve the absolute performance of the accelerator (excluding bunch by bunch systems), they could be profitably used to enhance overall machine stability, speed up optimization and augment diagnostic capabilities.

For Elettra 2.0, equipment featuring dual interfaces will include 1344 magnet power converters (PS) and 171 beam

position monitors (BPM). Furthermore, high-speed dedicated network links will be extended to low-level RF systems for the radiofrequency cavities, insertion devices, and photon diagnostics in the beamlines, even though the hardware for these systems has not yet been completely defined.

Due to the heterogeneous nature of connected devices and the opportunity to develop applications capable of harnessing highly synchronized data (LOCO [5], beam based alignment, post-mortem analysis, instability detection, electron-photon correlations, electron/photon beam optimizations, power converter prognostics and beam position monitor fault detection), the decision was made to centralize all fast interfaces at a unique point, a dedicated Intel based server. This server will be underpinned by a system based on DPDK and coded in C, ensuring the flexibility to implement advanced control schemes surpassing mere orbit stabilization.

## GOF UPGRADE FOR ELETTRA

### Hardware

Currently, the Elettra storage ring hosts 96 BPM detectors (Libera-Electron) that transmit data via Ethernet at 10 kHz to twelve Motorola MVME6100 CPUs, with each receiving data from eight BPMs using four 48-port Extreme X440-G2 switches. Each CPU, equipped with digital to analog cards (DAC), can superimpose a maximum current equal to one-fortieth of the total strength onto the corresponding seven horizontal/vertical correctors.

To implement a global correction, the entire orbit data is shared among the CPUs via reflective memory with each one handling one-twelfth of the feedback processing at Libera transmission rate.




Figure 1: Network layout of the upgraded GOF.

<sup>†</sup> giulio.gaio@elettra.eu

# OPERATIONAL CONTROLS FOR ROBOTS INTEGRATED IN ACCELERATOR COMPLEXES

S. Fargier, M. Di Castro, M. Donze  
CERN, European Organization for Nuclear Research

## Abstract

The fourth industrial revolution, the current trend of automation and data interconnection in industrial technologies, is becoming an essential tool to boost maintenance and availability for space applications, warehouse logistics, particle accelerators and for harsh environments in general. The main pillars of Industry 4.0 are the Internet of Things (IoT), Wireless Sensors, Cloud Computing, Artificial Intelligence (AI) and Machine Learning. We are finding more and more ways to interconnect existing processes using technology as a connector between machines, operations, equipment and people. Facility maintenance and operation is becoming more streamlined with earlier notifications, simplifying the control and monitor of the operations. Core to success and future growth in this field is the use of robots to perform various tasks, particularly those that are repetitive, unplanned or dangerous, which humans either prefer to avoid or are unable to carry out due to hazards, size constraints, or the extreme environments in which they take place. To be operated in a reliable way within particle accelerator complexes, robot controls and interfaces need to be included in the accelerator control frameworks, which is not obvious when movable systems are operating within a harsh environment. In this paper, the operational controls for robots, integrated in accelerator complexes at the European Organization for Nuclear Research (CERN), is presented. Current robot controls at CERN will be detailed and the use case of the Train Inspection Monorail (TIM) robot control will be presented.

## INTRODUCTION

Maintenance and availability for space applications, warehouse logistics, particle accelerators, and harsh environments in general, increasingly rely on automation and data interconnection thanks to advancements in technology and the Industry 4.0 revolution [1]. Nuclear plants like Fukushima [2], fusion reactors like ITER [3], and particle accelerator facilities [4], such as the European Organization for Nuclear Research (CERN) [5], the European X-ray free-electron laser (XFEL) [6], or FERMILAB [7], present harsh environments, several kilometers of underground and semi-structured accelerator areas with thousands of different items of equipment which need to be inspected and maintained often through remote interventions to decrease human exposure to hazards and reduce intervention time.

Mechatronics and robotics have undergone several evolutionary steps over the past years [8], but particle accelerator environments present particular constraints, such as accessibility, long distances, communication possibilities with installed equipment, unknown objects and occlusions in



Figure 1: Overview of CERN's robots: (a) Train Inspection Monorail (CERN-made), (b) EXTRM robot (CERN-controlled), (c) CERNBot in different configurations (CERN-made), (d) drone for tele-op support, (e) quadrupeds for challenging terrain, (f) high-payload industrial arm for milling and repetitive tasks, (g) TEODOR robot, (h) Telemax robot.

cluttered areas. In addition, the equipment is delicate and expensive, thus in most cases, equipment owners and/or machine experts must operate the robots. This aspect requires a remote robotic system with a user-friendly Human-Robot Interface (HRI) that augments the proprioception [9] of the person assigned to its control and/or monitoring, sometimes also needed for autonomous systems. To improve their awareness a robot's environment, operators should have seamless information of the robot's environment, position, joint angles, velocities, torques and forces fundamental to teleoperation tasks.

Industrial robotic solutions are assigned to repetitive work without much modularity or intelligence and are not adapted to harsh or semi-structured environments. Operating robots for maintenance in dangerous environments on costly machines requires skilled and well trained, dedicated shift operators [10] and specific controls infrastructures. To be operated in a reliable way robot controls and interfaces need to be included in the accelerator control frameworks, which is not obvious when movable systems are operating within a harsh environment [11, 12]. Human robot communication is a fundamental aspect for the success of remote missions, and the communication channels used (e.g. WiFi, 4G, radio etc.) are key points to be addressed when designing robotic controls [13]. Table 1 shows possible connection types between a robot and the operator's computer, while Table 2 indicates the bandwidth and standard deviation, round-trip time, and jitter for all connection types.

# EPICS IOC INTEGRATION WITH Rexroth CONTROLLER FOR A T-Zero CHOPPER\*

Bhargavi Krishna, Mariano Ruiz Rodriguez  
Oak Ridge National Lab, TN, USA

## Abstract

A neutron chopper is not typically used as a filter, but rather as a way to modulate a beam of neutrons to select a certain energy range or to enable time-of-flight measurements. T-Zero neutron choppers have been incorporated into several beamlines at SNS and are operated via a Rexroth controller. However, the current OPC server is only compatible with Windows XP, which has led to the continued use of an XP machine to run both the Indradrive (Rexroth interface) and EPICS IOC. This setup has caused issues with integrating with our Data Acquisition server and requires separate maintenance. As a result, for a new beamline project, we opted to switch to the Rexroth XM22 controller with T-Zero chopper, which allows for the use of drivers provided by Rexroth in various programming languages. This paper will detail the XM22 controller drivers and explain how to utilize them to read PLC parameters from the controller into the EPICS application and its Phoebus/CSS interface.

## INTRODUCTION

The Oak Ridge National Laboratory's (ORNL) Spallation Neutron Source (SNS) is a cutting-edge pulsed neutron facility, supported by the U.S. Department of Energy's Office of Basic Energy Sciences. Its primary mission is to investigate the properties and behaviors of materials using neutron scattering techniques [1]. At ORNL, diagnostic instruments adhere to the Network Attached Device (NAD) concept. Each sensor possesses its dedicated resources, including networking, timing, data acquisition, and processing. NADs operate independently, mitigating the fragility commonly associated with tightly integrated systems. The fundamental concept behind a Network Attached Device is to design each instrument as a self-contained networked device with its set of resources [2]. The software suite is built upon EPICS, which facilitates communication through a shared memory interface and serves as the standard control system for the entire SNS facility. One such resource is a T-zero choppers that rotate a large mass through the beam to effectively place a beam stop in the path of the beam and eliminate high-energy neutrons that occur early in the neutron pulse. These choppers must be operated in phase with the production of neutrons so that the energy distribution in each neutron pulse in the instrument remains constant. The chopper controller is designed to couple the phase of the chopper's rotor to the phase command produced by the timing reference generator [3]. ORNL is planning to build a TOF neutron imaging facility at the SNS, called VENUS, located on a decoupled

poisoned parahydrogen moderator, which offers the sharp neutron time pulse widths needed for high wavelength discrimination. VENUS is optimized for the measurement of micro-scale structures in radiography (2D) and tomography (3D) modes. The optical components are comprised of a series of selected apertures, T0 and bandwidth choppers, beam scrapers [4]. Few other beamlines at SNS has T0 chopper for which the EPICS IOC was developed a decade ago. But in preparation for the VENUS neutron imaging facility, which demands sharp neutron time pulse widths, a transition to a more integrated control system was required. The legacy architecture utilized Windows XP and an OPC server, leading to compatibility issues and separate maintenance efforts. To address these challenges, the Rexroth XM22 controller with its drivers was adopted, providing a versatile solution for EPICS integration. In the following sections the current methodology will be discussed, followed by the newly implemented IOC for the new controller.

The EPICS control system environment has been a cornerstone of the SNS facility since its implementation in 2006, supporting accelerator operations, and later, beam-lines, and data acquisition systems.

## T0 EPICS ARCHITECTURE

### Legacy T0 Architecture

This EPICS architecture at SNS consisted of two primary Input/Output Controllers (IOCs): a Serial IOC responsible for managing PVs related to serial communications and other software PVs, and an OPC IOC designed to communicate with the Rexroth OPC server, facilitating communication with the Programmable Logic Controller (PLC) using the Open Platform Communications (OPC) standard. OPC is a widely used software interface standard that enables Windows programs to communicate with industrial hardware devices.

In the legacy system, the primary role of the OPC IOC was to establish communication with the PLC, making it a critical component of the control system.

The Serial IOC scanned various devices, including the Televac vacuum gauge controller, pressure readings for helium, and multiple interlocks such as water temperature, water flow, shaft temperature, bearing temperatures, and tachometers for overspeed interlocks and motor speed/vibration monitoring.

Despite being EPICS IOCs, these components resided on a Windows machine, operating outside the Linux EPICS ecosystem, leading to maintenance challenges.

\* This work was supported by the U.S. Department of Energy under contract DE-AC0500OR22725.

# EPICS BASED TOOL FOR LLRF OPERATION SUPPORT AND TESTING

K. Klys\*, W. Cichalewski

Lodz University of Technology Department  
 of Microelectronics and Computer Science, Lodz, Poland  
 P. Pierini, European Spallation Source ERIC, Lund, Sweden

## Abstract

Interruptions in linear superconductive accelerators LLRF (Low-Level Radio Frequency) systems can result in significant downtime. This can lead to lost productivity and revenue. Accelerators are foreseen to operate under various conditions and in different operating modes. As such, it is crucial to have flexibility in their operation to adapt to demands. Automation is a potential solution to address these challenges by reducing the need for human intervention and improving the control's quality over the accelerator. The paper describes EPICS-based tools for LLRF control system testing, optimization, and operations support. The proposed software implements procedures and applications that are usually extensions to the core LLRF systems functionalities and are performed by operators. This facilitates the maintenance of the accelerator increases its flexibility in adaptation to various work conditions and can increase its availability level. The paper focuses on the architecture of the solution. It also depicts its components related to superconducting cavities parameters identification and elements responsible for their tuning. Since the proposed solution is destined for the European Spallation Source control system, the application has a form of multiple IOCs (Input/Output Controllers) wrapped into E3 (ESS EPICS Environment) modules. Nevertheless, it can be adjusted to other control systems - its logic is universal and applicable (after adaptations) to other LLRF control systems with superconducting cavities.

## INTRODUCTION

The crucial aspect of the accelerator's functioning is its availability. It is vital to provide an operation period without any downtimes or interruptions that could limit users' time for experiments. The availability of an accelerator could be considered in terms of two aspects. The first is the accelerator's flexibility to adapt to various modes of functioning or to different operational conditions. Working conditions can be understood for example as operating in pulse mode or continuous wave mode. They can be also seen as maintaining proper parameters when conditions such as environmental temperature have changed. It is worth mentioning that the broader set of operational conditions offered may appeal to potential users of the facility and present possibilities for different types of experiments. The second aspect is to ensure that the accelerator operates properly without interruption caused by faulty components or misconfiguration of equipment or control system. The control system itself must detect dangerous situations and try to prevent them or

\* kkllys@mail.dmc.s.pl

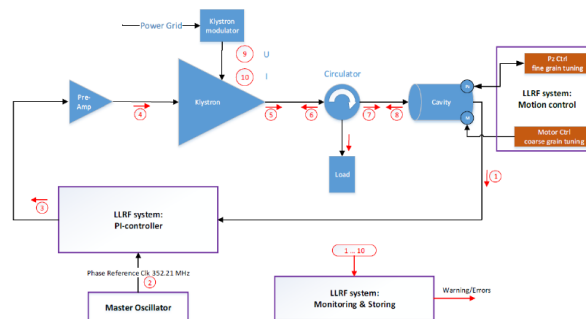


Figure 1: Schematic of superconducting (SC) RF system [3].

at least limit their effect by shutting down proper subsections of the accelerator.

One of the potential solutions to the described challenges is to introduce a certain level of automation to testing and operation support procedures. Automation helps to keep consistency in particle beam control. It allows for continuous and uninterrupted operation, optimizing efficiency and it limits downtimes needed for recovery after a failure. With minimal downtime for maintenance and streamlined processes, accelerators can operate at their optimum potential, maximizing productivity and contributing to the cost-effectiveness of research and industrial applications.

## LLRF CONTROL SYSTEM

### Role of LLRF Control System

The superconducting part of the accelerator is responsible for particle acceleration. It is made of niobium and resonant cavities. For each cavity, the electromagnetic wave with the appropriate gradient and frequency must be delivered. The LLRF (Low-Level Radio Frequency) control system manages the phase and amplitude within the cavity by gauging the field and comparing these measurements with a predefined set of target values [1]. In ESS (European Spallation Source), the control system uses a PI controller with the addition of adaptive feedforward compensation [2]. The schematic of the RF system is presented in Fig. 1. It is assumed that the LLRF should generate an input signal to the amplifier that drives the cavity to a field with an amplitude and phase precision that are within 0.5% and 0.5 degrees of a set value that is unique for each cavity [3, 4].

Another responsibility of the LLRF control system is compensation for cavities deformation caused either by Lorentz Force detuning (LFD) or microphonics effects. The cavities are tuned with two types of motion controllers: step

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# USE OF EPICS IN SMALL LABORATORIES

H. Junkes\*, W. Kirstaedter, A. Moshantaf, P. Oppermann<sup>1</sup>  
Fritz Haber Institute, Berlin, Germany

A. D. Fuchs<sup>2</sup>, J. A. F. Lehmeyer<sup>2</sup>, M. Schieber<sup>3</sup>, M. Krieger, H. B. Weber  
Friedrich-Alexander-Universität, Erlangen-Nürnberg, Germany

<sup>1</sup>also at SLAC, Menlo Park, CA, US

<sup>2</sup>also at Humboldt-Universität zu Berlin, Germany

<sup>3</sup>also at Universität der Bundeswehr München, Germany

## Abstract

Systems such as EPICS or TANGO are well-established solutions for instrument communication and data storage at large facilities (e. g. LIGO, ITER and Diamond Light Source). While these types of control systems have many inherent benefits such as time stamping, data archiving and timing solutions, they are also complex. The initial complexity and knowledge required to set up these systems is a high barrier to entry and often hinders adoption by new research facilities.

We demonstrate that such control systems can be used not only in highly sophisticated environments, but that it is possible to use and benefit from these systems in much smaller research labs and even in individual experimental setups. We present two separate use cases for which independent solutions have been developed.

The first use case demonstrates the use of an open-source configurable measurement software (NOMAD CAMELS) in a home-built Hall-effect measurement setup. Local instrument communication is combined with data acquisition from the laboratory infrastructure controlled by EPICS allowing flexible implementation of measurement protocols and automated recording of FAIR-compliant data and rich metadata.

The second use case presents the implementation of an all-EPICS environment in a catalysis laboratory allowing automated and standardized data acquisition and storage in a machine-readable and FAIR-compliant manner.

## MOTIVATION

Distributed, scalable and robust systems like EPICS (Experimental Physics and Industrial Control System) [1] and TANGO (TAco Next Generation Objects) [2] proved to be beneficial and reliable for the control of large-scale facilities, e. g. synchrotrons, and are well established [3]. However, such control interfaces offer opportunities also for smaller laboratories and even standalone equipment through standardized instrument communication and automated collection of metadata. The reason why EPICS or TANGO are rarely used in small labs is mainly the lack of training, documentation and successfully demonstrated use-cases.

\* junkes@fhi-berlin.mpg.de

The particular challenge in research areas like experimental condensed matter physics, physical chemistry and material sciences is the heterogeneity of experiments. In these disciplines, experiments are often performed in individual and specialized ad hoc experimental setups, which offer the required flexibility. However, at the same time, in contrast to standardized equipment, such setups require the development of specialized measurement and control software tools. This is typically done by the researchers themselves and requires a deep knowledge of software development including the communication protocols of the huge number of available interfaces and instruments from various vendors and ages. This often results in redundant software development as each lab or even researcher has to (re-)write the code to implement new measurement protocols. Moreover, frequently not much attention is paid to FAIR-compliant [4] data and metadata acquisition. In fact, metadata such as instrument settings are rarely stored along with the raw data.

Here, standardized control systems (e. g. EPICS or TANGO) can provide solutions when additionally equipped with user interfaces that provide a low-threshold entry. Documentation and training should enable even students at the bachelor level to implement new measurement protocols.

In this paper, we report on two successfully implemented use-cases for EPICS in small laboratories developed within the FAIRmat [5] NFDI [6] consortium which serves as a blueprint.

## NOMAD CAMELS AND THE DEMO LAB AT FAU ERLANGEN-NÜRNBERG

### NOMAD CAMELS

The difficulties of implementing measurement and process control protocols for specialized experimental setups that enables the acquisition of FAIR data motivates the development of a software which allows the user to quickly and dynamically control measurement setups. Within the FAIRmat project we are developing the configurable measurement and experiment control system NOMAD CAMELS (or just “CAMELS”) [7]. This software aims to simplify the implementation of virtually any measurement and process control protocol without the needs of programming skills.

A core concept is the automatic recording of data and metadata that accurately documents the experimental proce-



## DevPylon, DevVimba: GAME CHANGERS AT LULI

S. Marchand<sup>†</sup>, L. Ennelin, M. Sow, S. Minolli, J.-L. Bruneau, Laboratoire pour l'Utilisation des Lasers Intenses, Palaiseau, France

### Abstract

Apollon, LULI2000 and HERA are three Research Infrastructures of the Centre national de la recherche scientifique, École polytechnique, Commissariat à l'Énergie Atomique et aux Énergies Alternatives and Sorbonne University. Past-commissioning phase, Apollon is a four beam laser, multi-petawatt laser facility fitted with instrumentation technologies on the cutting edge with two experimental areas (short—up to 1m—and long focal—up to 20m, 32m in the future). To monitor the laser beam characteristics and control it from pilote through the interaction chambers, more than 500 devices are deployed in the facility and controlled through a Tango bus. This poster focuses on two linked software components: DevPylon and DevVimba. Each affected to a type of cameras: Basler via PyPylon wrapper interface of Pylon Software suite and Prosilica via Vimba's software development suite library, respectively. These two Tango devices are Python scripts constructed and generated via Code Generator POGO. They offer a specific way to monitor more than 100 CCD cameras in the facility at an image acquisition and display rate up to 10 Hz for a maximum of 300-shot at 1-minute rate per day and on an always-ON mode throughout the day.

### MULTI-PETAWATT APOLLON LASER FACILITY

In this chapter are presented some key figures of the facility, and some photos for a better comprehension of the areas of the facility. For the context, the Laboratoire pour l'Utilisation des Lasers Intenses (thereafter LULI) runs three facilities: historically LULI2000, Apollon and HERA. Apollo is the first of the three which was implemented with Tango control system [1].

- The Facility covers about 4,500m<sup>2</sup>
- LASER hall: ISO8 cleanroom
- Experimental rooms cover surfaces of 280m<sup>2</sup> and 490m<sup>2</sup> (focal lengths of 10+m)
- 5 m-thick concrete walls provide full radio protection
- 3PW on target in 2023, 8PW in 2025
- 500+ devices

With more than 500 devices running in the facility, a distributed control system was obviously needed. Tango was not so obvious ten years ago. Looking back, it was a good choice as the community is very helpful; larger and larger with new projects every year, and the system runs smoothly.

<sup>†</sup>stephane.marchand@polytechnique.edu

Some photos inside the facility, for the different rooms.



Figure 1: Supervision control room.



Figure 2: Long Focal Area (420m<sup>2</sup>).



Figure 3: Compression and switchyard laser subsystem.

# MOTION CONTROL ARCHITECTURE AND KINEMATICS FOR MULTI-DOF KIRK-PATRICK-BAEZ FOCUSING MIRRORS SYSTEM AT LNLS-SIRIUS \*

J. P. S. Furtado<sup>†</sup>, J. V. E. Matoso, T. R. S. Soares, G. B. Z. L. Moreno, C. S. C. Bueno, M. A. B. Montevechi Filho, Brazilian Synchrotron Light Laboratory, CNPEM, Campinas, Brazil

## Abstract

In modern 4th generation synchrotron facilities, piezo actuators are widely applied due to their nanometric precision in linear motion and stability. This work shows the implementation of a switching control architecture and a tripod kinematics for a set of 4 piezo actuators, responsible by positioning a short-stroke, the vertical and horizontal focusing mirrors of the Kirkpatrick-Baez mirror system at MOGNO Beamline (X-Ray Microtomography). The switching control architecture was chosen to balance timing to move through the working range (changing the beam incidence on stripes of low/high energy), resolution and infrastructure costs. This paper also shows the implementation and results of the layered kinematics, developed to uncouple short-stroke from long-stroke to fix any parasitic displacements that occur on the granite bench levellers due to motion uncertainty and mechanism non-linearities, and to match the required beam stability without losing alignment flexibility or adjustment repeatability. The architecture was built between a PIMikro-Move set of driver-actuators and an Omron Delta Tau Power Brick LV controller due to its standardization across the control systems solutions at Sirius, ease of control software scalability and its capability to perform calculations and signal switching for control in C language, with real-time performance to make adjustments to the angles responsible by focusing the beam in a speed that matches the required position stability, guaranteeing the necessary resolution for the experiments.

## INTRODUCTION

A new version of the Kirkpatrick-Baez focusing mirror system has been designed and installed in the MOGNO Beamline [1] (X-Ray Microtomography) at the Brazilian Synchrotron Light Laboratory [2] (LNLS), with extra features and implementations in comparison with other focusing systems currently working in other beamlines operation, in addition with smaller stability requirements — 8 nrad (RMS) for the horizontal focusing mirror (HFM) pitch angle, and 50 nrad for the pitch angle and 10 nm (RMS) for the translation of the vertical focusing mirror (VFM), in which both of them belong to the short-stroke subsystem, operating in high vacuum [3] (between  $10^{-9}$  mbar and  $10^{-9}$  mbar). The system implementation features a long-stroke, represented by a granite bench [4], which main function is to select the specific stripe for the experiment – low or high energy, present

in both mirrors –, but also to filter noises provided by floor vibrations and other scientific equipment. To accomplish with this functionality, the system was designed in order to implement a kinematics by layers. By doing this, there are two acting kinematic subsystems that, together, represent a numerical construction that allows a stripe selection in both horizontal and vertical focusing mirrors by moving both strokes in a determined cartesian direction – and that's what makes this implementation ubiquitous and innovative in the new Brazilian Synchrotron Light Source, as the reported results in terms of flexibility, stability and repeatability have become essential and unprecedented for the experiments in the beamline.

## KINEMATICS BY LAYERS

As mentioned before, the kinematic motion in a layered structure has been designed in order to provide a high resolution motion capacity for adjustment, and reasonable time for changing stripes – between high and low energies – in the same mechanism by moving the strokes in a single cartesian direction. Figure 1 shows the internal mechanism during its commissioning while the vacuum chamber was opened in a clean room (following the ISO 6 norm), responsible by focusing the synchrotron beam across the Z direction, following the laboratory convention guidelines [5]. The two stripes of the HFM are distributed in the Y direction, in which each center is separated in 7 mm – and that also happens with the vertical focusing mirror, but in the X direction.




Figure 1: Internal mechanism containing vertical and horizontal focusing mirrors during the commissioning of the short-stroke.

\* Work supported by the Ministry of Science, Technology and Innovation.  
<sup>†</sup> joao.furtado@lnls.br

# THE SILF ACCELERATOR CONTROLS PLAN

Z. Z. Zhou, G. M. Liu<sup>1</sup>, L. Hu, T. Liu, J. H. Zhu, T. Yu, M. T. Kang<sup>†</sup>

Institute of Advanced Science Facility, Shenzhen, Guangdong Province, 518107, China

<sup>1</sup>also at Guangdong Provincial Key Laboratory of Magnetolectric Physics and Devices and School of Physics, Sun Yat-Sen University, Guangzhou 510275, China

## Abstract

The Shenzhen Innovation Light Source Facility (SILF) is an accelerator-based multidiscipline user facility planned to be constructed in Shenzhen, Guangdong, China. This paper introduces controls design outline and progress. Some technical plans and schedules are also discussed.

## INTRODUCTION

The SILF is a fourth-generation medium-energy synchrotron radiation light source that envisions a future with over 50 beamlines. Its primary focus lies in supporting the development of domestic core industries, advancing frontiers in basic science research, and addressing strategic imperatives, including integrated circuits, biomedicine, advanced materials, and advanced manufacturing.

The SILF project received approval from the Shenzhen Government on September 10, 2020. The feasibility study was successfully completed in 2022, and the preliminary design phase is currently underway. Detailed plans for the construction phase of the project are currently being developed.

The accelerator complex is composed of a 200 MeV linac, a booster with ramping energy from 0.2 GeV to 3.0 GeV, and a 3.0 GeV storage ring as shown in Fig. 1. Two transport lines are designed to connect the linac, booster and storage ring. The circumference of the storage ring is 696 m, which includes 28 hybrid seven-bend achromat (H7BA) lattice periodic units to achieve the emittance below 100 pm-rad. The top-up operation mode (300 mA, 928 bunches) is considered, and a brightness of about  $10^{22} \text{ s}^{-1} \text{ mm}^{-2} \text{ m} \cdot \text{rad}^{-2} (0.1\% \text{ bandwidth})^{-1}$  is expected at the photon energy of 10 keV for SILF, as illustrated in Fig. 2. Considering machine errors and correction, the dynamic aperture of the storage ring is 7.2 mm, which meets the requirement for off-axis injection [1].

Furthermore, we have placed significant emphasis on the control system, which plays a pivotal role in providing and assuring high availability and reliability, laying a solid foundation for successful user experiments. In light of initial budget constraints, recent efforts have marked the commencement of research and development (R&D) activities. These R&D endeavours encompass the development of the EPICS7 development platform, the prototyping of vacuum control systems and PLCs, the creation of timing prototypes, machine protection prototypes, and the exploration of potential applications for orbit feedback.




Figure 1: Schematic layout of the SILF project.

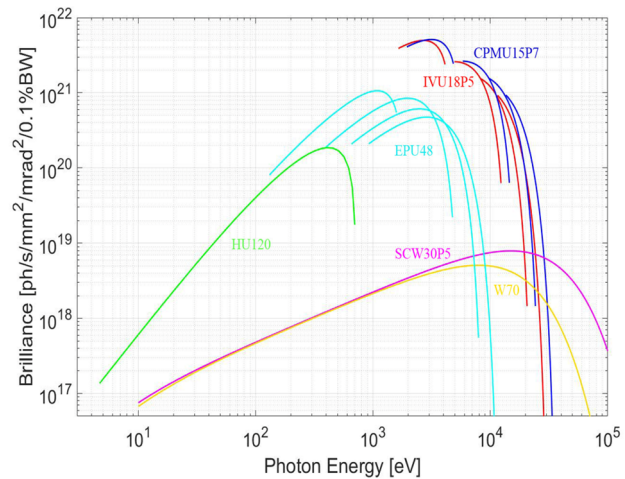


Figure 2: The available spectral brightness for SILF operated in the high brightness.

## CONTROL SYSTEM SCOPE

SILF uses the EPICS family of distributed control system software for the creation of a facility-wide data communication layer, which integrates all technical systems that participate in photon production and experiments [2]. The typical EPICS usage model of SILF is a three-tier system including the presentation layer, the middleware service layer and the frontend device layer as shown in Fig. 3.

<sup>†</sup>kangmingtao@mail.iasf.ac.cn

# PERSONNEL SAFETY SYSTEMS FOR ESS BEAM ON DUMP AND BEAM ON TARGET OPERATIONS

M. Mansouri†, A. Farshidfar, A. Abujame, A. Nordt, A. Andersson, A. Petrushenko, D. Paulic, D. Plotnikov, D. Daryadel, G. Ljungquist, J. Lastow, M. Carroll, M. Eriksson, N. Naicker, P. Holgersson, R. Foroozan, V. Harahap, Y. Takzare  
European Spallation Source ERIC, Lund, Sweden

## Abstract

The European Spallation Source (ESS) is a Pan-European project with 13 European nations as members, including the host nations Sweden and Denmark. ESS has been through staged installation and commissioning of the facility over the past few years. Along with the facility evolution, several Personnel Safety Systems, as key contributor to the overall personnel safety, have been developed and commissioned to support safe operation of e.g. test stand for cryomodules Site Acceptance Test, test stand for Ion Source and Low Energy Beam Transport, and trial operation of the Normal Conducting Linac. As ESS is preparing for Beam on Dump (BoD) and Beam on Target (BoT) operations in coming years, PSS development is ongoing to enable safe commissioning and operation of the Linear Accelerator, Target Station, Bunker and day-one Neutron Instruments. Personnel Safety Systems at ESS (ESS PSS) is an integrated system that is composed of several PSS systems across the facility. Following the experience gained from the earlier PSS built at ESS, modularized solutions have been adopted for ESS PSS that can adapt the evolving needs of the facility from BoD and BoT operations to installing new Neutron Instruments during facility steady-state operation. This paper provides an overview of the ESS PSS, and its commissioning plan to support BoD and BoT operations.

## INTRODUCTION

The European Spallation Source (ESS) represents a leap in the world of scientific research. Located in Lund, Sweden, ESS is a state-of-the-art facility under construction to address some of the most profound questions in the realm of materials science, chemistry, biology, and physics. Central to its mission is the utilization of neutrons produced through a process called spallation. This facility will boast the capability to accelerate protons to high energies, causing them to collide with a heavy metal target. This collision results in the release of neutrons, which are then harnessed to study the atomic and molecular structures of materials, providing insights into their fundamental properties and behaviours [1].

The operation of ESS is accompanied by the need to address a range of potential hazards that arise in a facility of this magnitude and complexity. These hazards can broadly be categorized as radiation, chemical, cryogenic, electrical, explosion, fire, laser, vacuum-related hazard, etc. Addressing these hazards effectively is paramount to ensuring the

continued success of ESS in advancing scientific knowledge while maintaining a safe and secure research environment. This paper delves into the *Personnel Safety Systems* at ESS, which plays a pivotal role in mitigating a designated number of hazards and safeguarding the individuals and infrastructure at ESS.

## PERSONNEL SAFETY SYSTEMS

At ESS, during the operation of proton accelerator and producing neutrons to be used at experimental stations, radiation is generated. This radiation includes neutron radiation, gamma radiation, fast neutrons, activation products, etc. The immediate radiation produced from these interactions is referred to as prompt radiation. ESS employs systems and administrative procedures to ensure the safety of personnel in the presence of prompt radiation. In this regard, one of the key systems is Personnel Safety Systems (PSS) that primarily encompasses a safety interlock system designed to protect personnel from prompt radiation. The PSS ensures that sources of prompt radiation, such as proton beam or Radio Frequency (RF)-powered cavities are switched off to protect personnel in beam enclosures (e.g. accelerator tunnel, experimental stations, etc.) from exposure to prompt radiation, prevent entry to beam enclosures while prompt radiation sources are energized, and switch them off when designated pre-defined access functions are violated.

The PSS at ESS is not confined solely to mitigating prompt ionizing radiation hazards but its function can be extended to mitigating a number of conventional hazards such as electrical and vacuum-related hazards. Moreover, where required, PSS contributes to access management to radiation areas in conjunction with ESS physical access control system.

## Methodology

The development of the PSS embodies a structured approach, guided by two foundational references: the ESS Handbook for Engineering Management and the Functional Safety Standard IEC 61511 [2]. The ESS Handbook for Engineering Management outlines the overarching principles that govern engineering practices at ESS. Within this framework, the development of the PSS is systematically planned and executed.

Complementing the ESS Handbook for Engineering Management, the Functional Safety Standard IEC 61511 provides a globally recognized framework for achieving functional safety in process industries. ESS draws upon the principles and best practices outlined in this standard to de-

† morteza.mansouri@ess.eu

# TOWARDS THE ZERO CODE WASTE TO INCREASE THE IMPACT OF SCIENCE

P. P. Goryl\*, L. Zytiniak, S2Innovation Sp. z o.o., Kraków, Poland  
A. Gotz, ESRF, Grenoble, France  
V. Hardion, MAX-IV, Lund, Sweden  
S. Hauf, European XFEL GmbH, Schenefeld, Germany  
K. S. White, ORNL, Oak Ridge, USA

## Abstract

Accelerators and other big science facilities rely heavily on internally developed technologies, including control system software. Much of it can and is shared between labs, like the Tango Controls and EPICS. Then, some of it finds broad application outside science, like the famous World Wide Web. However, there are still a lot of duplicating efforts in the labs, and a lot of software has the potential to be applied in other areas. Increasing collaboration and involving private companies can help avoid redundant work. It can decrease the overall costs of laboratory development and operation. Having private industry involved in technology development also increases the chances of new applications. This can positively impact society, which means effective spending of public funds. The talk will be based on the results of a survey looking at how much scientific institutes and companies focus on collaboration and dissemination in the field of software technologies. It will also include remarks based on the authors' experiences in building an innovative ecosystem.

## INTRODUCTION

An efficient economy is key for today's world challenges related to climate and limited resources. Zero Waste can be applied to all production cycles related to material resources and non-tangible assets like software.

The primary business of large scientific infrastructures, like particle accelerators or telescopes, is to conduct research in areas other than software technology. However, software is a core tool for them. No commercial software is often available due to specific functional and operational requirements and a limited market. This means the laboratories must develop or buy software development services to satisfy their needs. As the scientific labs compete on the scientific results, not the functionalities of the software tools, there is space for collaboration. So, these institutions already share an effort to provide software tools, minimising the costs of implementing functionalities. When they do this, sharing source code, they follow the Zero-Waste idea and build an efficient economy. The above concerns are the source of the Zero Code-Waste term.

\* piotr.goryl@s2innovation.com

## Zero Code-Waste

The Zero Code-Waste is not about keeping maintenance and running of the legacy or obsolete software.

The Zero Code-Waste is an idea of maximizing software reuse between projects. This shall decrease duplicated effort in providing the same functionalities within different software packages. In this context, it is directly related to **collaboration**. Ultimately, it should lower development and maintenance costs, increasing the so-called development capacity of the community, which means developing more functionalities with less human effort and other resources spent.

The Zero Code-Waste can also encompass methodologies and techniques to minimize source code (re-)writing within a project. However, it is not covered by this paper.

## Industry Involvement

Whereas current business models in the commercial software industry are often connected to closed-source IP rights, the scientific community, funded mainly with public money, relies heavily on open-source tools and collaborative models. It does not mean the software industry is not involved in the scientific project. Besides using standard commercial software, like operating systems or office applications, scientific institutes and collaborations often outsource or subcontract software development for their scientific needs. While the institutes do not regard this as increasing the chances of the software being reused, there are some reasons to claim it may help navigate towards Zero Code-Waste.

## Disclaimer

This paper covers only a few Zero Code-Waste and collaboration aspects. Some claims come from the author's experience but are not supported with references or discussed in-depth due to publication volume. The paper is more to get focus and start a discussion on the topic, which certainly needs more research.

## THE SURVEY

The survey [1] aimed to get wide feedback from software group leaders and managers about collaboration and its impact on the workload. The survey was sent to many organizations (70) all over the world. The 19 of them have answered.

# ENHANCED MAINTENANCE AND AVAILABILITY OF HANDLING EQUIPMENT USING IIoT TECHNOLOGIES

A. Garcia Fernandez, E. Blanco, D. Lafarge, G. Thomas, J. C. Tournier, CERN, Geneva, Switzerland

## Abstract

CERN currently houses 6000 handling equipment units categorized into 40 different families, such as electric overhead travelling cranes (EOT), hoists, trucks, and forklifts. These assets are spread throughout the CERN campus, on the surface (indoor and outdoor), as well as in underground tunnels and experimental caverns. Partial access to some areas, a large area to cover, thousands of units, radiation, and diverse needs among handling equipment makes maintenance a cumbersome task.

Without automatic monitoring solutions, the handling engineering team must conduct periodic on-site inspections to identify equipment in need of regulatory maintenance, leading to unnecessary inspections in hard-to-reach environments for underused equipment but also reliability risks for overused equipment between two technical visits. To overcome these challenges, a remote monitoring solution was introduced to extend the equipment lifetime and perform optimal maintenance.

This paper describes the implementation of a remote monitoring solution integrating IIoT (Industrial Internet of Things) technologies with the existing CERN control infrastructure and frameworks for control systems (UNICOS and WinCC OA). At the present time, over 600 handling equipment units are being monitored successfully and this number will grow thanks to the scalability this solution offers.

## INTRODUCTION

The Handling Engineering (HE) group plays a crucial role in ensuring the efficient handling and transport of equipment and materials at CERN, which are subject to regulatory constraints regarding maintenance, from equipment with unconventional shapes to extremely delicate detector parts. Their responsibilities include designing, and conducting feasibility studies related to transport and handling operations, as well as procuring, installing, and commissioning standard and custom-built equipment. The group also manages and maintains all industrial transport, handling, and lifting equipment to ensure optimal performance throughout its life cycle. This includes the purchase, installation, and maintenance of thousands of equipment such as lifts, cranes, forklifts, and custom-made transport machines. In this context, the HE group requires modern tools to streamline their tasks, reducing equipment downtime and maintenance expenses.

The handling equipment assets to be monitored and maintained are scattered across the CERN sites. These extends over a vast area, spanning across France and Switzerland with both surface and underground areas.

An optimal solution requires a system capable of efficiently acquiring geographically widespread data while minimizing installation costs and making the most of CERN's existing infrastructure, including communication networks, data acquisition and monitoring systems, and computerized asset management systems. Consequently, the desired characteristics for the final solution include interoperability and synergy with existing systems. A new system to manage handling equipment also needs to cope with a large variety of machinery that will be enhanced by a set of smart sensors to measure relevant data.

Hence, the adoption of a combination of IIoT (Industrial Internet of Things) technologies, industrial automation technologies, and computerized maintenance software suites appears to be the appropriate approach.

## HANDLING EQUIPMENT SCOPE

The HE group identified two principal areas where introducing the previously mentioned technologies would prove advantageous for their tasks.

The first area concerns overhead lifting equipment, including EOT (Electric Overhead Traveling) cranes and hoists. An analysis made in 2020 concluded that important savings on preventive maintenance could be done by changing the actual strategy based on a fixed periodicity to a strategy based on the usage of each crane. To be able to implement such change, the usage of each crane has to be known and remotely communicated on a daily basis. Thus, the HE group aimed to cut down on maintenance costs and prevent equipment maintenance periods from impacting the runtime of installations like the LHC accelerator. Their approach involved monitoring both the number of lifts (i.e., how often the equipment is in motion) and the total working hours. Given that cranes and hoists require maintenance every 30 hours of operation, they proposed implementing an automatic system to trigger maintenance orders accordingly. Depending on the level of operational priority of the cranes, EN-HE has been scheduling preventative maintenance every 6 months or 1 year, harmonizing the CERN Safety Rule GSI-M-1 prescriptions and the recommendations of a multitude of manufacturers. This is confirmed and benchmarked with the industry, where usage-based maintenance is starting to be the standard in terms of good practice. With this method, maintenance teams would carry out maintenance only on assets that genuinely required it, rather than adhering to a fixed periodic maintenance schedule regardless of the crane's actual usage. Additionally, administrators would be able to identify assets that are either overused or underused compared to the average, allowing for better anticipation of when a crane should be replaced.

# WORKING TOGETHER FOR SAFER SYSTEMS: A COLLABORATION MODEL FOR VERIFICATION OF PLC CODE

Ignacio D. Lopez-Miguel\*, TU Wien, Vienna, Austria,  
Borja Fernández Adiego<sup>†</sup>, Enrique Blanco Viñuela, CERN, Geneva, Switzerland,  
Matias Salinas, Christine Betz, GSI, Darmstadt, Germany

## Abstract

Formal verification techniques are widely used in critical industries to minimize software flaws. However, despite the benefits and recommendations of the functional safety standards, such as IEC 61508 and IEC 61511, formal verification is not yet a common practice in the process industry and large scientific installations. This is mainly due to its complexity and the need for formal methods experts. At CERN, the PLCverif tool was developed to verify PLC programs formally. Although PLCverif hides most of the complexity of using formal methods and removes barriers to formally verifying PLC programs, engineers trying to verify their developments still encounter different obstacles. These challenges include the formalization of program specifications or the creation of formal models. This paper discusses how to overcome these obstacles by proposing a collaboration model that effectively allows the verification of critical PLC programs and promotes knowledge transfer between organizations. By providing a simpler and more accessible way to carry out formal verification, tools like PLCverif can play a crucial role in achieving this goal. The collaboration model splits the specification, development, and verification tasks between organizations. This approach is illustrated through a case study between GSI and CERN.

## INTRODUCTION

Programmable Logic Controllers (PLCs) find extensive utilization in industrial automation, encompassing safety and standard control systems not only at establishments like CERN and GSI but also within diverse process industries and other scientific installations. The incorrect behavior of the PLC programs can yield considerable repercussions, such as property damage, environmental harm, or, in certain instances, even personal injuries. Consequently, the assurance of their accurate functionality holds paramount importance. Although testing has long been the conventional means of validating PLC programs, its effectiveness often falls short as the sole verification method. Even when automated, testing cannot achieve exhaustive coverage and thus lacks the ability to guarantee the absolute correctness of a given logic. Certain categories of requirements, such as safety specifications (which demand the prevention of unsafe states) or invariants (formulas that must hold throughout all conceivable system runs), can pose substantial challenges and might even be unfeasible to assess through testing alone.

Model checking is a formal verification technique that complements the testing activities to fully validate and verify PLC programs. It involves evaluating the fulfillment of formalized requirements upon a mathematical model of the system under scrutiny. This assessment encompasses every feasible combination of inputs and all potential execution paths. Additionally, if a violation is found, the technique provides the path leading to the breached requirement.

The broad adoption of model checking in the realm of PLCs encounters a two-fold challenge: Firstly, creating the mathematical model that represents the system being analyzed can require an in-depth understanding of the model-checking tools. Secondly, a multitude of real-life PLC logics are characterized by their complexity, leading to what is known as state-space explosion problem, that is, a large number of potential input combinations and execution paths that exceed the bounds of exhaustive exploration.

In September 2020, CERN released the PLCverif platform under an open-source license, aiming to facilitate the utilization of model-checking tools among PLC developers. This goal was achieved by automating the conversion of PLC programs into their corresponding mathematical models. Subsequently, a series of abstraction algorithms were integrated to address the state-space explosion challenge.

However, the incorporation of model-checking tools into the design and development process is far from simple, encountering various challenges. Large projects involve distinct teams that must engage in clear and unambiguous communication. Thus, achieving this integration is not trivial, particularly when the initial specifications are not formalized and lack precision. These specifications are usually written using natural language. Consequently, when the PLC developer implements the specifications, they must interpret these instructions, potentially resulting in a PLC program that deviates from the initially envisioned design conceived by the team responsible for creating the specifications.

Through collaborative efforts, by exchanging expertise and tools, constructing safe systems can become more viable. Such a collaboration took place between CERN and GSI, wherein CERN helped GSI to formalize requirements and to formally verify their PLC code.

Upon termination of the collaboration, a number of discrepancies between the specification and the PLC code were found. This discovery enhanced the understanding of safety engineers and PLC developers at GSI about the functioning of their PLC programs. It enabled them to refine the specification, rectify errors, and improve the system overall. Moreover, this collaboration yielded advantages for CERN

\* ignacio.lopez@tuwien.ac.at

<sup>†</sup> borja.fernandez.adiego@cern.ch

# REPLACING CORE COMPONENTS OF THE PROCESSING AND PRESENTATION TIERS OF THE MEDAUSTRON CONTROL SYSTEM

A. Höller<sup>†\*</sup>, S. Vörös<sup>‡\*</sup>, A. Kerschbaum-Gruber, C. Maderböck, D. Gostinski, L. Adler, M. Eichinger, M. Plöchl, EBG MedAustron GmbH, Wiener Neustadt, Austria

## Abstract

MedAustron is a synchrotron-based ion therapy and research facility in Austria, that has been successfully treating cancer patients since 2016. MedAustron acts as a manufacturer of its own accelerator with a strong commitment to continuous development and improvement for our customers, our users and our patients. The control system plays an integral role in this endeavour. The presented project focuses on replacing the well-established WinCC OA [1] SCADA system, enforcing separation of concerns mainly using .NET and web technologies, along with many upgrades of features and concepts where stakeholders had identified opportunities for improvement during our years of experience with the former control system setup for commissioning, operation and maintenance, as well as improving the user experience. Leveraging our newly developed control system API, we are currently working on an add-on called "Commissioning Worker". The concept foresees the functionality for users to create Python scripts, upload them to the Commissioning Worker, and execute them on demand or on a scheduled basis, making it easy and highly time-efficient to execute tasks and integrate with already established Python frameworks for analysis and optimization. This contribution outlines the key changes and provides examples of how the user experience has been improved.

## MEDAUSTRON

MedAustron is a synchrotron-based ion beam therapy and research centre located in Austria. The facility has been treating cancer patients since 2016, currently with the use of protons and carbon ions. In parallel to medical and research operation, MedAustron is also working on development and improvement projects, like commissioning the accelerator for use with a helium beam. MedAustron acts as a manufacturer of not only its own particle therapy accelerator, but also of further ion beam centres.

In consequence of being a manufacturer of multiple facilities, MedAustron has taken the decision to exchange some components of the MedAustron Control System (MACS) and enhance it with components more suitable to conform with the challenges of operating and further developing multiple facilities and the additional use case of initial commissioning (component commissioning, accelerator commissioning, beam commissioning) of new facilities.

\* These authors contributed equally to this work

<sup>†</sup> angelika.holler@medaustron.at

<sup>‡</sup> sandor.voros@medaustron.at

## MEDAUSTRON CONTROL SYSTEM

The MedAustron accelerator delivers proton and carbon ion beams for cancer treatment and research to four irradiation rooms including 3 horizontal, 1 vertical and a gantry beam line. In order to deliver beam from the source to the treatment room, the particle accelerator consists of:

- ~300 power converters controlling ~350 magnets
- ~270 vacuum devices
- ~80 beam diagnostics devices
- 2 RF systems and amplifiers
- 3 ion sources

The MedAustron Control System (MACS) creates a framework for all accelerator devices by providing a standardized set of essential "services" and interfaces in various tiers [2].

## Architecture

The architecture extends the industry best practice, 3-tier model [3, 4] in accordance with [5]: (1) presentation tier, (2) processing tier, (3) equipment tier and (4) frontend tier. Components in separate tiers that are distributed over a number of processing devices communicate with each other through a dedicated Ethernet network. Communication between equipment tier and frontend tier may also be achieved through dedicated field-bus and custom links, depending on the imposed constraints. [2] (see Fig. 1)

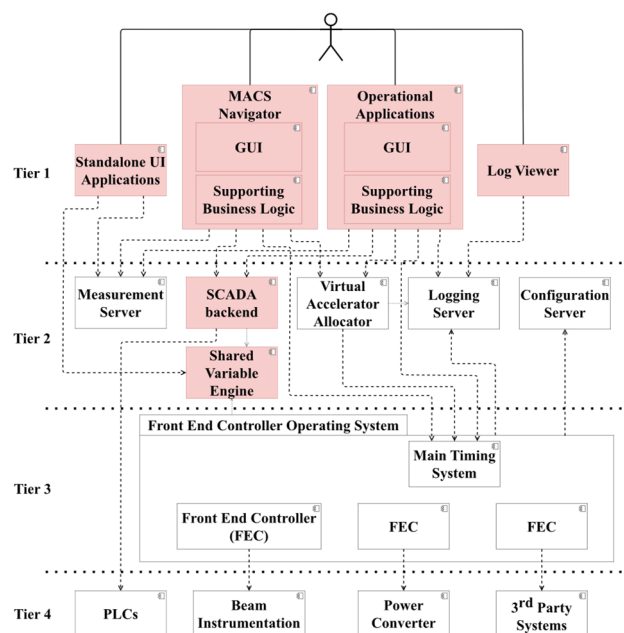


Figure 1: Architecture of the former MedAustron Control System (MACSv1).



# SYSTEM IDENTIFICATION VIA ARX MODEL AND CONTROL DESIGN FOR A GRANITE BENCH AT SIRIUS/LNLS \*

J. P. S. Furtado<sup>†</sup>, I. E. dos Santos, T. R. S. Soares  
Brazilian Synchrotron Light Laboratory, CNPEM, Campinas, Brazil

## Abstract

Modern 4th generation synchrotron facilities demand mechatronic systems capable of fine position control, improving the performance of experiments at the Beamlines. In this context, granite benches are widely used to position systems such as optical elements and magnetos, due to its capacity of isolating interferences from the ground. This work aims to identify the transfer function that describes the motion of the granite bench at the EMA Beamline (Extreme conditions Methods of Analysis) and then design the control gains to reach an acceptable motion performance in the simulation environment before embedding the configuration into the real system, followed by the validation at the beamline. This improvement avoids undesired behaviour in the hardware or in the mechanism when designing the controller. The bench, weighting 1.2 tons, is responsible by carrying a coil, weighting 1.8 tons, which objective is to apply a 3 T magnetic field to the sample that receives the beam provided by the electrons accelerator. The system identification method applied in this paper is based on the auto-regressive model with exogenous inputs (ARX). The standard servo control loop of the Omron Delta Tau Power Brick controller and the identified plant were simulated in Simulink in order to find the control parameters. This paper shows the results and comparison of the simulations and the final validation of the hardware performance over the real system.

## INTRODUCTION

The main theme of this work is the identification and controller design of a granite bench present in the Brazilian Synchrotron Light Laboratory (LNLS) [1], the 4th generation particles accelerator in the Brazilian Center for Research in Energy and Materials (CNPEM). The objective is to provide an effective method to identify the behavior of the granite bench responsible by positioning a coil that produces a magnetic field of extreme conditions for the scientific experiments that happen in this beamline [2]. In addition to that, the controller must be designed to move the system in a stable manner. The feedback transfer function model should be the same as the one present inside the Power Brick LV, commonly known as PBLV [3].

Figure 1 shows the coil responsible by providing the magnetic field under extreme conditions, essential for several kinds of experiments that happen inside the EMA Beamline. As mentioned in the last section, the coil weights 1.2 tons, while the bench [4], responsible by carrying the coil,

weights 1.8 tons – it is the biggest of the whole laboratory, between all Beamlines in operation. The design of a good stabilizing controller is essential to guarantee the success of the experiments, as the bench must be kept stopped during the process (acquisitions during incident beams), and also move in a safe manner in order to preserve the integrity of the whole mechanism.




Figure 1: Coil responsible by providing the magnetic field under extreme conditions.

The plant identification was done considering the black box approach [5], and this box encompasses both electrical and mechanical systems present in the granite bench: the operational amplifier, present in the PBLV power output, and the motor coupled to a belt and pulley subassy that provides the translation of the bench and controls the position of the magnetic coil. The identification was done using the ARX model [6], and the excitation signal was a smooth ramp in reference to the controller, in closed-loop mode.

The controllers tuning – in terms of gains, such as proportional, integrative, derivative, etc. [7] – are usually done in an empirical manner. So the main contribution of this work is to identify and design the controller of the system in order to have the performance evaluated even before embedding the gains inside the PBLV. This avoids the possibility to damage the system, by bringing rough movements to the granite bench or amplifiers burnout by over-current. Also, it is possible to test and validate different controllers setups and gains configurations – due to the simulation environment that applies the identified plant –, helping to find the best performance in terms of transient and stationary stability.

## EXPECTATIONS

The main expectation of this work is to provide a method to find control gains in order to stabilize critical systems,

\* Work supported by the Ministry of Science, Technology and Innovation.  
<sup>†</sup> joao.furtado@lnls.br

# IMPROVEMENTS ON KINEMATICS AND CONTROL OF GRANITE BENCHES AT LNL-SIRIUS\*

J. V. E. Matoso<sup>†</sup>, J. P. S. Furtado, J. P. B. Ishida, T. R. S. Soares  
 Brazilian Synchrotron Light Laboratory, Campinas, Brazil

## Abstract

At the Brazilian Synchrotron Light Laboratory, the radiation beam is conditioned by optical elements that must be positioned with high stability and precision. Many of the optical elements are positioned using granite benches that provide high coupling stiffness to the ground and position control in up to six degrees of freedom, using a set of stepper motors. The solution of the inverse kinematics was done numerically by the Newton Raphson method. By employing the property that these systems have small rotation angles, the Jacobian matrix used in this numerical method can be simplified to reduce computational execution time and allow high processing rates. This paper also shows the results of adding a notch filter to the position servo control loop of the granite benches to increase stability due to their mass-spring-damper characteristics. The kinematics and control of the granite benches are implemented in an Omron Power Brick LV controller, with the kinematics developed in MATLAB and the C-code generated by MATLAB C-Coder. Reducing the execution time of the kinematics improves the efficient use of the computational resources and allows the real-time clock rate to be increased.

## INTRODUCTION

Granite benches have many applications in the Sirius beamlines, such as positioning mirrors, monochromators, and experimental stations, providing high mechanical and thermal stability [1, 2].

The latest granite bench design has seven motors and six degrees of freedom (Fig. 1). The system consists of a stack of three granite pieces. The first layer is supported by three levelers that provide translation in Y and rotations along the X and Z axes. The next two layers of granite together provide rotation along the Y axis and translations in the X, Y and Z axes. These layers are positioned using a system of belts and pulleys driven by stepper motors. To reduce friction, the granites have air bearings to guide and lift the granite layers.

This article presents the latest advances in the calculation of the inverse kinematics of this mechanism since the last article [3] and introduces a new discussion on the position control of air-bearing granites using belt and pulley systems. All kinematic calculations and motor control are performed on an Omron Delta Tau Power Brick LV controller.



Figure 1: Mirror granite bench design used as a reference.

## KINEMATICS

### Model Overview

The kinematic calculations of this system are divided into two parts: the closed-chain tripod formed by the three levelers; and the remaining open-chain robot with 7 axes formed by the resulting effect of the levelers plus the movement of the granites (Fig. 2). The tripod kinematics for the levelers will not be covered in this article. For more information, see [3].




Figure 2: Mirror system kinematic scheme [3].

The control point is located at the  $S_8$  coordinate system with coordinates  $\vec{u}$  relative to the coordinate system  $S_0$  fixed on the beam, where  $\vec{f}$  is the forward kinematics and  $\vec{e}$  the axes positions. The axes  $R_{x_{lev}}$ ,  $R_{z_{lev}}$  and  $T_{y_{lev}}$  are

\* Work supported by the Ministry of Science, Technology & Innovation.  
<sup>†</sup> joao.matoso@lnls.br

# SYSTEM IDENTIFICATION EMBEDDED IN A HARDWARE-BASED CONTROL SYSTEM WITH CompactRIO\*

T. R. S. Soares<sup>†</sup>, J. P. S. Furtado, V. B. Falchetto, G. O. Brunheira, J. L. B. Neto, R. R. Gerales  
 Brazilian Synchrotron Light Laboratory, CNPEM, Campinas, Brazil

## Abstract

The development of innovative model-based design high bandwidth mechatronic systems with stringent performance specifications has become ubiquitous at LNLS-Sirius beamlines. To achieve such unprecedented specifications, closed loop control architecture must be implemented in a fast, flexible, and reliable platform such as NI's CompactRIO (cRIO) controller that combines FPGA and real-time capabilities. The design phase and life-cycle management of such mechatronics systems heavily depends on high quality experimental data either to enable rapid prototyping, or even to implement continuous improvement process during operation. This work aims to present and compare different techniques to stimulus signal generation approaching Schröder phasing and Tukey windowing for better crest factor, signal-to-noise ratio, minimum mechatronic stress, and plant identification. Also shows the LabVIEW implementation that requires specific signal synchronization and processing on the main application containing a hardware-based control architecture, increasing system diagnostic and maintenance ability. Finally, experimental results from the High-Dynamic Double-Crystal Monochromator (HD-DCM-Lite) of QUATI (quick absorption spectroscopy) and SAPUCAIA (small-angle scattering) beamlines and from the High-Dynamic Cryogenic Sample Stage (HD-CSS) from SAPOTI (multi-analytical X-ray technique) of CARNAÚBA beamline are also presented in this paper.

## INTRODUCTION

System identification is the initial process of Model-Based Control Design (MBD), that builds a mathematical model representing the behavior of a system based on measurements of stimulus and response signals. In LNLS-Sirius [1], for the High-Dynamic Double-Crystal Monochromators (HD-DCM-Lite) [2, 3] and High-Dynamic Cryogenic Sample (HD-CSS) Stage, system identification in the frequency domain is the first step followed by controller design based in the loop shaping technique [4], controller embedding and validation. A good system identification provides accurate models, allows reliable controller design and one of its critical aspects is the stimulus signal, that requires specific characteristics mentioned later in this article. Also, all signal paths should be the same as seen by the control system, ensuring every component is included in the identified system, with simultaneous sampling of stimulus and response signals.

\* Work supported by the Ministry of Science, Technology and Innovation.  
<sup>†</sup> telles.soares@lnls.br

In order to increase system diagnostic and validation capabilities, stimulus generation and acquisition of system response were embedded alongside the hardware-based control algorithm, which required few modifications given the integration flexibility of the cRIO [5], highly customized with LabVIEW Real-Time and LabVIEW FPGA. Acquired signals are post-processed in MATLAB to estimate frequency response of the system.

## ARCHITECTURE

HD-DCM-Lite and HD-CSS share a common high dynamic control architecture, presented in Fig. 1, where the controller, kinematics (homogeneous) transformations, and mechanical structure are specific for each application.




Figure 1: Common high dynamic control architecture.

The stimulus generation and response acquisition blocks were added in the Real-Time code to save FPGA resources, thus the FPGA works as a streaming bridge between the excitation signal generation (in Real-Time) and the input of the homogeneous transformations (also applied in FPGA), as shown in Fig. 2.

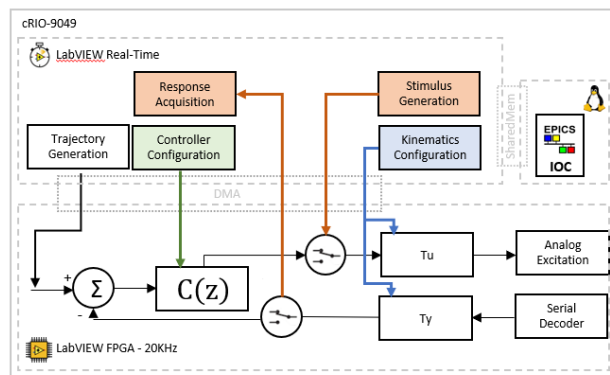


Figure 2: System identification architecture blocks.

When system identification mode is selected, the control algorithm is bypassed by two FIFO buffers where the stimulus signal will be applied in the actuation path and the

# MOBILE PUMPING UNITS FOR PARTICLE FREE BEAM VACUUM

T. Joannem<sup>†</sup>, D. Loiseau, C. Walter, CEA Paris-Saclay IRFU DIS, Saclay, France  
S. Berry<sup>\*</sup>, Q. Bertrand, C. Boulch, G. Monnereau, CEA Paris-Saclay IRFU DACM, Saclay, France

## Abstract

Commissariat à l’Energie Atomique et aux Energies Alternatives (CEA) is in charge of providing cryomodules for European Spallation Source (ESS). This collaboration lead CEA to design, build and provide cryomodules with industrial quality requirements and challenging milestones.

One critical element for cryomodules efficiency is vacuum quality and more specifically lack of particles. To reach this particle free beam vacuum, pumping units were specifically designed build and are still use at CEA Paris-Saclay.

This article will focus on these pumping unit for vacuum and control system point of views.

## INTRODUCTION

### Context

ESS cryomodules are built in several steps. Most of them are did in clean room to limit particles generation in order to avoid pollution and unwanted interactions with beam. One element to answer to this problematic is a mobile pump unit built by CEA Paris-Saclay vacuum and control teams.

### Constraints

The pumping system must answer to following constraints:

- It should be able to pump cavity up to -9mBar,
- It must be outside clean room to avoid particle generation, reliable for permanent use,
- Remotely control by operator in clean room and locally by other operators,
- Mobile from a pumping spot to another,
- Exchangeable and user friendly.

### Goals

Control system was designed to be the best assistant of clean room operator and to be forget, so when an operator starts an automatic pumping sequence, he can work without according attention to the system. Another goal of the system is to be safe and reliable. A lot of securities were developed, integrated and tested to prevent any operator error and to protect the pumping unit in case of hardware malfunction.

## REQUIREMENTS

### Presentation

The set-up needs to allow a leak check down to leak rates of  $10^{-10}$  mbar l/s, cavity assembly and pumping procedures will lead to reach this goal.

### Vacuum Process

The movement of particles in UHV systems has been studied systematically by DESY using an in-vacuum particle counter. Slow pumping process inherit of this study.

### Clean Room

Cryomodule assembly is a long and complex process that need specific tools to avoid particle generations. CEA Paris-Saclay clean rooms answer to this requirement with pumping slots providing a safe connection with regular atmospheric pressure environment.

**Pumping Slots** IRFU DACM main clean room is design and is wide enough to build cryomodules cavity train. Due to the cavity’s quantity and different state of each one they often are at different vacuum and clean status. Therefore, pumping slots must be accessible up to six at the same time. Valves and pumps movements are generating particles so pumping units can’t be located in a clean environment.

Thanks to sealed connections a cavity inside the clean room can be connected to a pumping unit outside clean area.

Therefore, independent pumping procedures could be start for each cavity.



Figure 1: Mobile pumping units.

**Mobile Flexibility** As mentioned in previous chapter pumping units are connected to pumping slots. CEA Paris-Saclay have height pumping units use for multiple purpose of clean pumping in different locations: three clean rooms and one cryomodule test stand. According to experiments needs they have to be move from a location to another. This is the reason why they are mobile (see Fig. 1).

<sup>†</sup> tom.joannem@cea.fr  
<sup>\*</sup> stephane.berry@cea.fr

# THE LASER MEGAJOULE FACILITY STATUS REPORT

I. Issury, J.P. Airiau, Y. Tranquille-Marques, CEA CESTA, Le Barp, France

## Abstract

The Laser MegaJoule (LMJ), a 176-beam laser facility developed by CEA, is located at the CESTA site near Bordeaux. The LMJ facility is part of the French Simulation Program, which combines improvement of theoretical models and data used in various domains of physics, high performance numerical simulations and experimental validation. It is designed to deliver about 1.4 MJ of energy on targets, for high energy density physics experiments, including fusion experiments.

In this paper, a review of the LMJ facility and the PETAL project is given with details on the status report and an update of the activities. Afterwards, a presentation of the Target Diagnostic is given. In addition, a brief description of the LMJ Control System is given with the major software developments during the last 2 years. Finally, the major recent experiments on LMJ are presented.

**Key words:** Laser facility, LMJ, PETAL, Control Systems.

## INTRODUCTION

Since it definitively abandoned nuclear testing, France relies on the Simulation Program to guarantee the operational performance and safety of its nuclear deterrent weapons throughout their lifetime.

Successful simulation requires both:

Qualified computer codes that integrate laboratory-validated physics models to simulate weapon functioning;

Teams of qualified physicists to use these codes.

In this respect, the Megajoule Laser (LMJ) [1] plays a vital role, as it is used to validate the numerical codes and certify the skills of French physicists.

In fact, the LMJ is designed to provide the experimental capabilities to study High Energy Density Physics (HEDP). The LMJ is a keystone of the Simulation Program, which combines improvement of physics models, high performance numerical simulation, and experimental validation, in order to guarantee the safety and the reliability of French deterrent weapons. When completed, the LMJ will deliver a total energy of 1.4 MJ of 0.35  $\mu\text{m}$  ( $3\omega$ ) light and a maximum power of 400 TW.

The LMJ is dimensioned to accommodate 176 beams grouped into 22 bundles of 8 beams. These beams are located in the four laser bays arranged on both sides of the central target bay of 60 meters length and 40 meters height. The target chamber and the associated equipment are located in the center of the target bay.

The LMJ technological choices were validated on the LIL, a scale-1 prototype composed of 1 bundle of 4 beams. The first bundle of 8 beams has been commissioned at the end of 2014. The second bundle has been commissioned at the end of 2016 following the same commissioning process. Fifteen bundles are now operational by the end of

2023, and the physics experiments using the 80 operational beams took place during the first semester of 2023.

Furthermore, there is the PETAL laser beam which consists in the addition of one short-pulse (0.5 to 10 ps) ultra-high-power (1 up to 7 PW) with a high-energy beam (1 up to 3.5 kJ) to the LMJ facility. PETAL offers a combination of a very high intensity petawatt beam, synchronized with the nanosecond beams of the LMJ.

The first phase of nuclear commissioning of LMJ has been achieved to take into account high-energy particles created by PETAL, and neutron production from  $\text{D}_2$  fusion reaction. A subsequent phase will take into account DT targets by 2030.

## THE LMJ PROJECT

### Presentation of the LMJ Facility

The LMJ facility is a flash-lamp-pumped neodymium-doped glass laser (1.053  $\mu\text{m}$  wavelength) configured in a multi-pass power amplifier system. The 1.053  $\mu\text{m}$  wavelength is converted to the third harmonic (0.351  $\mu\text{m}$ ) and focused, by means of gratings, on a target at the center of the target chamber. Once fully commissioned, with 176 beams (44 quads) operational, LMJ will deliver shaped pulses from 0.7 ns to 25 ns with a maximum energy of 1.4 MJ and a maximum power of 400 TW of UV light on the target (Figure 1).




Figure 1: Schematic view of the Laser Megajoule showing the main elements of the laser system.

At the center of the target bay, the target chamber consists of a 10 meter diameter aluminium sphere, equipped with two hundred ports for the injection of the laser beams, the location of diagnostics and target holders. It is a 10 cm thick aluminium sphere covered with a neutron shielding made of 40 cm thick borated concrete. The inside is covered by protection panels for X-ray and debris.

LMJ is configured to operate in the “indirect drive” scheme, which drives the laser beams into cones in the upper and lower hemispheres of the target chamber. Forty quads enter the target chamber through ports that are located on two cones at 33.2° and 49° polar angles. Four

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# THE LASER MEGAJOULE FULL AUTOMATED SEQUENCES

Yves Tranquille-Marques, Jean-Philippe Airiau, Pierre Baudon, Irwin Issury  
Alain Mugnier CEA, CESTA, Le Barp, France

## Abstract

The LMJ (Laser MegaJoule), a 176-beam laser facility developed by the French Nuclear Science directorate CEA (Commissariat à l’Energie Atomique et aux Energies Alternatives), is located at the \*CEA CESTA (Centre d’Etude Scientifique et Technique d’Aquitaine) site near Bordeaux. The LMJ facility is part of the French Simulation Program. It is designed to deliver about 1.4 MJ of energy on targets, for high energy density physics experiments, including fusion experiments. Since 2022, the LMJ facility aims at carrying out experiments with 13 bundles of 8 laser beams and 20 target diagnostics. In order to achieve daily shots including all the preparatory steps, the LMJ performs night activities from now on and the presence of technical operators is not required. These sequences perform vacuum windows inspection and beam alignment. They take into account all the prerequisites for their good performances and are scheduled automatically one after the other. They deal with material security and unexpected equipment alarms. They endeavour to required tasks success and give a detailed report of the night events to the shot director. This paper gives a presentation of the two sequences with solutions in order to answer the technical specifications and the last enhancements.

## LASER MEGAJOULE SHOT DAY

The LMJ [1] facility will count 176 laser beams. It is dimensioned to deliver 1.4 MJ shot of UV laser at 351 nm to a 10 mm target. The LMJ is designed to study high density plasma physics. It’s also a part of the French Simulation Program that forms the basis of the safety and reliability of French deterrent weapons.

The LMJ facility splits in 4 laser bays. In its center, takes place the target bay (Fig. 1). A shot is composed of laser pulses which are amplified in the laser bays and focused on a target in the target chamber through the vacuum windows. As prerequisites to shots, two actions are performed:




Figure 1: LMJ view.

- the alignment of laser beams and Target Diagnostics
- the vacuum windows damage inspection

Both these actions require positioning specific devices in the center of the target chamber. Two different positioners have to be inserted but not at the same moment [2].

At 8:30 PM, first of all, the vacuum windows inspection sequence is started. It is called SQMH (SeQuence Machine H). It lasts 4 hours. Right after, a second sequence is launched. It deals with beams alignment first part and it is called SQME (SeQuence Machine E). It lasts 9 hours, nevertheless it waits the end of SQMH, to start night alignment main part. From 9 PM to 6 AM, these sequences remains without any operator control.

Operators come back handling sequences at 6 AM. SQMH is finished. SQME is paused and displays night results. At that time, operator could play again part of the SQME if necessary.

Then SQME ends, in order operators to align Target diagnostics and beam transport section. They also carry out KDP configuration (Wavelength conversion) and CSO (Wavefront correction).

Next, target is aligned. And shot sequence starts with validation shot without power conditioning followed by power shot. Post shot steps handle power conditioning lamp test and results recovery (Fig. 2). And at 8:30 PM, a new shot day starts.

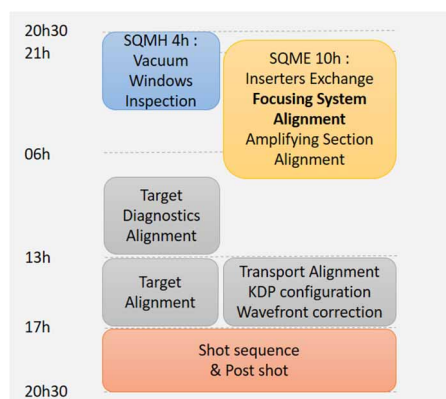


Figure 2: Shot day timeline.

Critical path is built of functions which require to position a device in the center of the target chamber. SQMH need to insert MDCC (Center Chamber Diagnostic Module) held by POCC (Center Chamber Object Positioner), whereas SQME inserts RC (Common Reference) fit at the end of PR (Reference Positioner). Changing inserters without operator increases collision risk, it is shown thereafter how it is reduced and controlled.

## CONTROL SYSTEM KEYS

### Control System Architecture

The LMJ control system architecture is built on a 4 layers architecture (Fig. 3):

## TANGO AT LULI

S. Marchand<sup>†</sup>, L. Ennelin, M. Sow, S. Minolli, J.-L. Bruneau, Laboratoire pour l'Utilisation des Lasers Intenses, Palaiseau, France

### Abstract

Apollon, LULI2000 and HERA are three Research Infrastructures of the Centre national de la recherche scientifique (CNRS), École polytechnique (X), Commissariat à l'Énergie Atomique et aux Energies Alternatives (CEA) and Sorbonne University (SU). Now in past-commissioning phase, Apollon is a four beam laser, multi-petawatt laser facility fitted with instrumentation technologies on the cutting edge with two experimental areas (short—up to 1m—and long focal—up to 20m, 32m in the future). To monitor the laser beam characteristics through the interaction chambers, more than 300 devices are distributed in the facility and controlled through a Tango bus. This poster presents primarily a synthetic view of the Apollon facility, from network to hardware and from virtual machines to software under Tango architecture. We can here have an overview of the different types of devices which are running on the facility and some GUIs developed with the exploitation team to insure the best possible way of running the lasers. While developments are still currently under work for this facility, upgrading the systems of LULI2000 from one side and HERA from the other side are underway by the Control-Command & Supervision team and would follow the same specifications to offer shared protocols and knowledge.

## MULTI-PETAWATT APOLLON

### LASER FACILITY

In this section are presented some key figures of the building, and some photos for a better comprehension of the areas of the facility. For the context, the Laboratoire pour l'Utilisation des Lasers Intenses (thereafter LULI) runs three facilities: historically LULI2000, Apollon and HERA. Apollo is the first of the three which was implemented with Tango control system.

- The Facility covers about 4,500m<sup>2</sup>

LASER hall: ISO8 cleanroom

- Experimental rooms cover surfaces of 280m<sup>2</sup> and 490m<sup>2</sup> (focal lengths of 10+m)
- 5 m-thick concrete walls provide full radio protection

With more than 500 devices running in the facility, a distributed control system was obviously needed. Tango [1] was not so obvious ten years ago. Looking back, it

was a good choice as the community is very helpful; larger and larger with new projects every year, and the system runs smoothly. Some photos inside the facility, for the different rooms.



Figure 1: Supervision control room.



Figure 2: Long Focal Area (420m<sup>2</sup>).




Figure 3: Compression and switchyard laser subsystem.

<sup>†</sup>stephane.marchand@polytechnique.edu

## STATUS ON CONTINUOUS SCANS AT BESSY II

N. Greve\*, G. Pfeiffer, M. Neu, D. Kraft, M. Brendike†

Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Berlin, Germany

### Abstract

Continuous energy scanning is an important feature for many beamlines at BESSY II. In 2015 this method was used at 11 undulator and 6 dipole beamlines. Since then the demand for this feature – especially among new build beamlines – increased, while the availability of the used hardware decreased. To tackle this problem, we investigate alternative hardware and software solutions. By introducing an independent high level controller between the device controllers we can compensate for communication incompatibilities and hence increase flexibility.

This paper shows the status of our research. The ideas leading to a first prototype, the prototype itself, and first results will be presented.

### INTRODUCTION

The synchrotron light source BESSY II supplies some 35 beamlines with brilliant soft x-rays from undulator and dipole sources with defined photon energy ranges. In this context, scanning across a provided photon energy range by synchronizing the movement of multiple actuators along a beamline without intermediate stops is referred to as continuous scanning (CS) or the continuous mode (CM) [1]. CS significantly decreases the time required for energy scanning experiments like x-ray absorption spectroscopy. Additionally it reduces sample exposure and hence increases sample lifetime, reduces optomechanical vibrations due to smoother acceleration and deceleration phases, and increases a beamline's experimental portfolio by including time-resolved experiments [2–4]. The above points combined with its wide use at BESSY II and other synchrotron facilities make CS an essential technique for today and future beamline operation.

The methods behind current CS at BESSY II beamlines are explained in detail by Balzer *et al.* [1]. The motion controllers (VME-based hardware) of the monochromator and the undulator communicate via a dedicated CAN Bus interface. The monochromator input-output-controller (MONO IOC) holds the gap-to-energy look-up-table (LUT) and can request the undulator to move from gap position A to gap position B. During this move the undulator will report its current gap position regularly via the CAN-bus to the MONO IOC. The monochromator then adjusts its optical elements according to the reported gap position and the stored LUT. The undulator moves with a symmetrical trapezoidal gap velocity profile with a constant maximum speed. This is equivalent to moving from energy X to energy Y with a non constant rate of change of photon energy. Sampling this energy at regular intervals – which is often wanted – is

thus more complex since the triggering needs to be done by additional hardware like Panda Box instead of a fixed time based sample rate [5].

Performing CS as described has limitations. First, scanning with a constant rate of energy change is not supported. This could be changed by implementing a trajectory generator on the MONO IOC, but this is challenging as the computational resources on the VME hardware are very limited. Second, current CS depends on the use of specific near-end-of-life hardware, which can lead to support and incompatibility issues in the future. To mitigate this risk other institutes also start searching for alternative hardware solutions [6, 7].

Off-the-shelf motion controllers are available to synchronize movement of multiple motor types. Integrating them into the existing BESSY II beamline control environment on the other hand is challenging. Their proprietary closed source nature often reduces their configurability, which makes it difficult and sometimes even impossible to integrate them with existing in-house applications. Hence a freely configurable solution, which can be integrated with existing applications is necessary.

In this paper we present such a solution in the form of a prototype. The following chapters discuss the prototype's design and implementation in more detail. We finish with first results from tests performed on real beamline equipment outside of BESSY II user operation.

### PROTOTYPE DESIGN

The main goal of this project is to move the undulator and monochromator axes synchronously using variable speed over time.

For the beginning we focused on controlling the movement of the undulator gap and the monochromator grating and mirror. Controlling the undulator shift is subject of future iterations.

Designing a modular software architecture is another requirement. Beamlines often vary in their configuration in regard to the type of monochromator and undulator present at the beamline, which makes a modular software design necessary. In a modular software design it is possible to include and exclude modules without breaking the whole software architecture. This makes it possible to combine several software modules, based on the beamline configuration at hand.

Additional requirements were:

- to further reduce implementation costs, this first prototype is implemented as an open-loop system, but extending the system to become a closed-loop system has to be possible;

\* nico.greve@proton.me

† maxim.brendike@helmholtz-berlin.de



## BLUESKY WEB CLIENT AT BESSY II

Huiling He<sup>1\*</sup>, William Smith<sup>1†</sup>, Sebastian Sachse<sup>1‡</sup>, Gabriel Preuß<sup>2§</sup>, Ruslan Ovsyannikov<sup>2¶</sup>  
Helmholtz Zentrum Berlin, Berlin, Germany

### Abstract

Considering the existing Bluesky control framework at BESSY II, a web client with React based on Bluesky HTTP Server is being developed. We hope to achieve a cross-platform and cross-device system to realize remote control and monitoring of experiments. The implemented functionalities for now are monitoring of the Bluesky Queue Server status, controlling over a Bluesky Run Engine environment, browsing of Queue Server history as well as editing and running of Bluesky plans. Challenges around the presentation of live data using Tiled are explored. This work builds on that of NSLS II who created a React based web interface and implements a tool for BESSY II.

### INTRODUCTION

In today's era of rapid digital innovation, web applications are playing an increasingly significant role in experiment control. This paper presents a comprehensive exploration of the Bluesky [1, 2] control system and its web user interface and offers insights into the integration, development, and core functionalities.

The motivation behind the innovative web-based interface is twofold. Firstly, it arises from the recognition of the potential of web applications. Secondly, it aligns with the adoption of the Bluesky project at Bessy II [3]. Although current Bluesky applications at Bessy II are underutilized, they hold great promise for the future of control systems. Understanding the scalability and versatility of Bluesky control is crucial. Web user interfaces are needed to meet the dynamic demands of experimental setups and simplify the experiment process. Moreover, a web client empowers remote control, enabling users to initiate, monitor and manage experiments from anywhere with internet access.

Our primary goal of the web client is to gain a comprehensive understanding of the Bluesky QueueServer [4] and its interaction with web operations. This encompasses real-time status monitoring, console tracking, experimental data display and experiment execution. These web-based capabilities aim to replace conventional reliance on manual instructions with command-line interfaces or Qt GUIs [5] at experimental stations.

This article delves into the architecture of web interactions with related back-end servers and databases and provides a detailed exploration of the functional components for the Bluesky web client. In the following sections, we will introduce the system architecture, back-end servers with Docker

usage, layout and functionalities of the web client, software implementation and conclusion.

### BLUESKY WEB CLIENT

#### System Architecture

The system architecture illustrated in Figure 1, follows a client-server model with two distinct facets: one for experiment control and the other for data display. Bluesky [1] components are responsible for experiment control, while Tiled [6] is dedicated to data presentation.

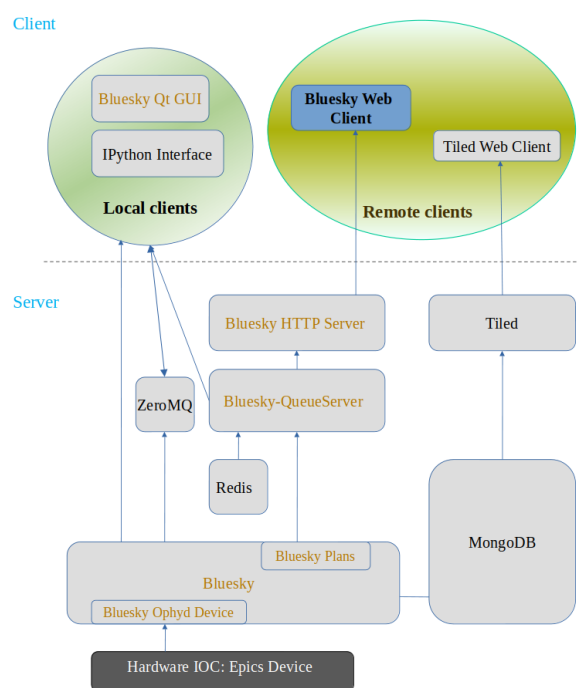


Figure 1: Server-Client architecture.

On the server side, Bluesky components are configured and managed within an IPython profile, which grants the Bluesky QueueServer [4] complete authority. Consequently, any authorized client can interface with the Bluesky environment via the Bluesky QueueServer [4]. This architecture has been successfully deployed at BESSY II [3], accommodate the demand for Bluesky control with IPython command-line tools, Qt GUIs [5] and web clients.

The Bluesky Queue Server [4] introduces a queuing system that facilitates experiment management, enabling researchers to execute experiments in a queue-based manner. This works as the cornerstone of the Bluesky web client, meanwhile, Bluesky HTTP Server [7] operates as a communication bridge to simplify and streamline the interactions between user operations on the web client and Bluesky QueueServer [4].

\* huiling.he@helmholtz-berlin.de

† william.smith@helmholtz-berlin.de

‡ sebastian.sachse@helmholtz-berlin.de

§ gabriel.preuss@helmholtz-berlin.de

¶ ovsyannikov@helmholtz-berlin.de

# TEST BENCH FOR MOTOR AND MOTION CONTROLLER CHARACTERIZATION

D. Kraft\*, M. Brendike†

Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Berlin, Germany

## Abstract

To maximize beamtime usage motorization of beamline equipment is crucial. Choosing the correct motor is complex, since performance depends largely on the combination of motor and motion controller. This challenge, alongside renewing the twenty years old infrastructure at BESSY II, led to the demand for a motor test bench.

The test bench was designed to be modular, so it fits different motors, loads and sensors. It allows independent performance verification and enables us to find a fitting combination of motor and motion controller. The test bench is operated via EPICS and BlueSky, allowing us usage of Python for automated data acquisition and testing. An overview of the mechanical and electrical setup, as well as some data from different performance tests will be presented.

## INTRODUCTION

BESSY II is nearly 25 years in service. Focusing on soft and tender x-ray radiation, over 2700 worldwide users per year conduct experiments on 48 beamlines [1]. Monochromators are needed to select the required photon energy from the spectrum. The monochromators depend on a precise positioning of their optical components. The positioning is usually done by stepper motors and controllers. Therefore, the combination of stepper motor and controller play a critical role in the success of an experiment. At BESSY II there is no uniform consensus about what stepper motor and controller combination should be used. Over the last 20 years many different philosophies and requirements led to a variety of different motion controllers and stepper motors in the experimental hall [2]. Various attributes of the controller and stepper motor combination add up to its performance for a dedicated task [3]. For example, the ability to perform  $\mu$ steps precisely and reliably plays a key role in positioning. Whereas moving large inertias focuses on high torque.

Previously, finding a well-suited combination for a specific task depended on manufacture's data sheets and experiences. Because BESSY II was built in the early 90s, many of the motion controllers are out of service by their manufacturers today. Aside from that the information provided by them is usually not sufficient. The construction of a motor test bench allows a decision based on results of self-defined tests to find the right use for a combination.

## IDEA

The test bench should help with characterization of the different motor-controller combinations used in the experi-

mental hall at BESSY II. The idea was to build a test bench, which allows custom defined test sequences. The requirements of the test bench are:

- fit a variety of components,
- simple prototype,
- mobile,
- similar to the beamline environment.

The results of the tests performed on the test bench allow us to determine whether a combination of stepper motor and motion controller fits the requirements for a task.

## IMPLEMENTATION

The test bench is divided into a mechanical setup (with experiment specific components mounted) and a control and data acquisition systems. Figure 1 displays the interconnections between the systems.

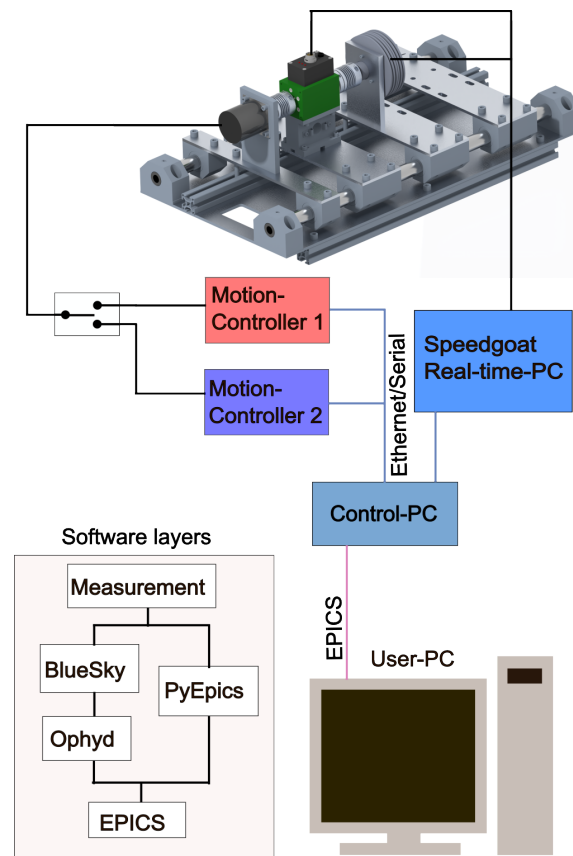


Figure 1: Schematic representation of the test bench.

## Mechanical Setup

Figure 2 shows the mechanical setup in two views. Modularity is a key requirement for the mechanical setup of the

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

\* david.kraft@helmholtz-berlin.de

† maxim.brendike@helmholtz-berlin.de

# MIGRATING FROM ALARM HANDLER TO PHOEBUS ALARM-SERVER AT BESSY II AND HZB

M. Gotz\*, T. Birke, Helmholtz-Zentrum Berlin, Berlin, Germany

## Abstract

The BESSY II lightsource has been in operation at Helmholtz-Zentrum Berlin (HZB) for 25 years and is expected to be operated for more than the next decade. The EPICS Alarm Handler (ALH) has served as the basis for a reliable alarm system for BESSY II as well as other facilities and laboratories operated by HZB. To preempt software obsolescence and enable a centralized architecture for the alarm systems running throughout HZB, it is being migrated to the alarm-service developed within the Control System Studio/Phoebus ecosystem. To facilitate simultaneous operation of the old alarm system while evaluating the new system, tools were developed to automate creation of the Phoebus alarm-service configuration files in the control systems' build process. Additionally, tools and configurations were devised to mirror the old system's key features in the new one. This contribution presents the tools developed and the infrastructure deployed to use the Phoebus alarm-service at HZB.

## MOTIVATION

The ALH [1] has served the HZB well as the basis for the current alarm system of the BESSY II light source for 25 years. However, with the prospect of operating BESSY II for another 10 years the desire arose to replace it with a more modern system. The key limitations of the ALH are that it is a single application requiring a graphical user interface (GUI) and that it is no longer actively maintained. Instances on different computers are mostly independent from one another, creating problems with synchronicity and requiring special care for automated actions which should only be performed once. Furthermore, no more active development means future obsolescence is always a risk and there will be no support for newer EPICS developments like pvAccess.

The Phoebus [2] alarm system uses an alarm server, of which only one instance is required, communicating with multiple GUIs via a Kafka message broker. This approach appealed greatly to us, eliminating most of the above mentioned shortcoming of the ALH. Furthermore, Phoebus seemed to us the only truly viable replacement candidate for ALH. There is, to our knowledge, no other alarm system in the EPICS ecosystem that is both actively developed and with contributions from more than one institute.

## OVERVIEW AND SPECIFIC NEEDS

The Phoebus alarm system is well documented [2]. At its heart is a Kafka message broker. An alarm server reads the configuration from Kafka and writes alarm state changes to

Kafka. Clients read alarm state changes from Kafka, display and allow changes to the configuration via Kafka. Additional services exist to log alarms to Elastic Search or log changes to the configuration.

The basic setup is simple enough. However, covering all our needs required additional developments. While manual addition of PVs to the alarm configuration of Phoebus is easy, at HZB we have a well established alarm configuration for BESSY II and our other labs. With roughly 10000 PVs in it, we needed an automated way to convert this configuration. Particularly, because parts of this configuration are generated from spreadsheets, which change fairly frequently.

Another aspect for us were access controls. The control system runs in a separate network, which should be able to operate in isolation. Also, changes to the alarm-system, like acknowledging alarms, should be possible from the entire network. Nevertheless, read-only remote access to the alarm-system should be possible from within the wider institute.

Finally, we required a method of notifying operators of alarms via an internal SMS-like service. For the ALH a script modified the configuration for one instance after deployment adding severity commands to each alarm. A more transparent system was desired for the new setup.

## CONFIGURATION CONVERSION AND GENERATION

ALH and Phoebus offer different features in the alarm configuration. For instance, the force mask in ALH allows changes to latching or logging behavior, in addition to the ability to disable an alarm based on the value of other PVs. On the other hand, the annunciation is only present in Phoebus. Therefore, a perfect conversion between these two formats is not possible. Our goal was a tool to convert the unambiguous parts and inform us about the problematic portions. We developed the python package *phoebusalarm* [3] to accomplish this. It parses alh-files into a python object structure. From that structure both alh and phoebus xml-files can be exported. In addition, the package can be used to programmatically create an alarm tree and export it into either format. Using this tool and manual intervention where necessary, we converted our alarm tree into the phoebus format. After the manual conversion was complete the process was also integrated into our build process. This ensures any changes to alh-files remain forward compatible with phoebus. Additionally, all the resulting alarm-configuration files are combined into a single file and checked against a schema with xmllint ensuring the presence of all included files and syntactic correctness.

\* malte.gotz@helmholtz-berlin.de

# ABOUT THE NEW LINEAR ACCELERATOR CONTROL SYSTEM AT GSI

P. Gerhard\*, GSI Helmholtzzentrum für Schwerionenforschung GmbH, 64291 Darmstadt, Germany

## Abstract

The first accelerator at GSI, UNILAC, went into operation in the early 1970s. Today, UNILAC is a small accelerator complex, consisting of several ion sources, injector and main linacs comprising 23 RF cavities, several strippers and other instrumentation, serving a number of experimental areas and the synchrotron SIS18. Three ion species can be provided at different energies simultaneously in a fast time multiplex scheme, two at a time. The UNILAC is going to be the heavy ion injector linac for FAIR currently under construction next to GSI, supported by a dedicated proton linac. The current linac control system dates back to the 1990s. It was initiated for SIS18 and ESR, which enlarged GSI at the time, and was retrofitted to the UNILAC. The linear decelerator HITRAP was added in the last decade, while an sc cw linac is under development. Today CRYRING, SIS18 and ESR are already operated by a new system based on the LHC Software Architecture LSA and other developments from CERN, as FAIR will be. In order to replace the outdated linac control system and simplify and unify future operation, a control system on the same basis is being developed for all GSI linacs. Recently the first data supply tests were performed during a dry run.

## INTRODUCTION

The GSI Helmholtzzentrum für Schwerionenforschung was founded in 1969 [1]. Its mission is to develop, build and operate heavy ion accelerators and conduct research with heavy ions. Starting with the UNiversal Linear ACcelerator UNILAC (as shown in Fig. 1) in the 1970s [2,3], the heavy ion synchrotron SIS18 [4] and the experimental storage ring ESR were added around 1990 [5]. In the context of this major extension of the facility, a highly sophisticated control system was developed in order to run the different accelerators in a well coordinated and highly efficient way. This system made extensive use of then state-of-the-art digital equipment and techniques, for instance data bases, computers with OpenVMS for the high level operation with programs written in Fortran, FPGAs for the complex online timing generation and device control by front end computers connected via MIL field buses.

Today, the construction of FAIR (Facility for Antiproton and Ion Research) [6, 7] is about to multiply the size and capabilities of the heavy ion accelerator facility on the Darmstadt site again, using the existing accelerator complex as its injector. A completely new control system for FAIR was developed in close cooperation with CERN. It is based on the LHC Software Architecture LSA [8] and the Front End Software Architecture FESA [9], while the new General Machine Timing is based on White Rabbit [10] to

enable accurate timing over the enlarged facility. The new control system was developed using the versatile CRYRING accelerator [11] as a testbed [12], and was put into regular operation for the ring machines SIS18 and ESR and the High Energy Beam Transport HEBT since 2016 [13, 14]. For the UNILAC, this has not been achieved so far.

The main feature of the facility, which is also reflected in the control system, is the efficient operation of the whole complex. The key is a virtual parallel operation of different ion beams, serving several users individually by a time multiplex operation. The increased number of storage rings for FAIR with their long cycle times adds a new dimension to this, while the UNILAC still has to be operated in the same manner on a cycle time of 20 ms.

End of 2019 a dedicated project was started to migrate the UNILAC to the new FAIR style control system, notwithstanding that some actions were already started earlier.

## LEGACY CONTROL SYSTEM

The legacy control system still in use today for UNILAC was introduced some 40 years ago. Some major upgrades and modernizations have taken place, for instance switching from OpenVMS to Linux as the underlying operating system, moving the production server cluster to virtual machines and recently replacing obsolete thin clients. But some outdated hard- and software components still remain and may be replaced or modernized only with large efforts or by replacing the control system completely. Among those are the front end computers and the corresponding software, software services not available for current operating systems, and operating software written in Fortran.

The current upgrade project aims at retrofitting the UNILAC to the modern FAIR control system standard. Because



Figure 1: The Alvarez type main linear accelerator of the UNILAC went into operation in 1975.

\* p.gerhard@gsi.de

# OPERATION OF THE ESR STORAGE RING WITH THE LSA CONTROL SYSTEM

S. Litvinov, R. Hess, B. Lorentz, M. Steck  
GSI, Darmstadt, Germany

## Abstract

The LHC Software Architecture (LSA) framework developed at CERN has been applied as a core for the new control system of the accelerator complex GSI and the future facility FAIR, Germany.

The Experimental Storage Ring (ESR) at GSI was recommissioned with the LSA and different accelerator and physics experiments were performed in the recent years. The overview of the ESR performance will be presented here. The features and challenges of the operation with the LSA system will be outlined as well.

## INTRODUCTION

The ESR [1] is the core instrument for unique physics experiments in the FAIR facility [2] (see Fig. 1) [3]. It is operated for accumulation, storage, cooling and deceleration of a wide range of heavy ion beams in the energy range from 4–400 MeV/u coming from the synchrotron SIS18 [4] via the FRagment Separator (FRS) [5] or a direct transport line. The ESR is a symmetric ring with two arcs and two straight sections and a circumference of 108.36 meters (see Fig. 2). It consists of 6 dipole magnets (deflection angle is 60°) and 10 quadrupole families (20 quadrupoles in total). For the second-order corrections 8 sextupole magnets are installed in the arcs. The ESR can be operated at a maximum magnetic rigidity of 10 Tm. For reducing transverse and longitudinal emittances of the stored ion beams, the ESR is equipped with the electron cooler which is installed in one of the straight sections of the ring. In another straight section the internal gas-jet target for in-ring reaction experiments is installed.

The ESR model for the control system is based on a generic circular accelerator hierarchical model, which was




Figure 2: Layout of the ESR.

implemented into LSA [6] and features of the ESR were added on top of common structures.

A generic model is a collection of all input parameters and algorithms based on the accelerator physics approach, which are needed to perform the calculation of hardware values according to the defined machine operation. Physics quantities and hardware set values are represented as parameters, which are grouped according to their parameter type, describing the physics category of the parameter. For each parameter there is a special make rule, which calculates its relations with other different parameters. Relations between parameters build a parameter hierarchy, which in turn form an accelerator hierarchy [7].

The ESR operational setting is called a pattern (see Fig. 3). It is a fully pre-planned chain of different subsequent blocks, named sub-chains, which consist of different beam processes arranged in a certain order. The beam process defines a special procedure of the machine, like the beam injection or the magnets' ramp. The order of beam processes in the certain sub-chain is built in such a way, to perform the corresponding beam procedure, pre-described in the sub-chain (e.g. beam deceleration). The sub-chains are also ordered in a certain way, required by the user to perform a full machine




Figure 1: Layout of present GSI and future FAIR facility.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# SUMMARY REPORT ON MACHINE LEARNING-BASED APPLICATIONS AT THE SYNCHROTRON LIGHT SOURCE DELTA

D. Schirmer\*, S. Khan, A. Radha Krishnan

Center for Synchrotron Radiation (DELTA), TU Dortmund University, Germany

## Abstract

In recent years, several control system applications using machine learning (ML) techniques have been developed and tested to automate the control and optimization of the 1.5 GeV synchrotron radiation source DELTA. These applications cover a wide range of tasks, including electron beam position correction, working point control, chromaticity adjustment, injection process optimization, as well as CHG spectra (coherent harmonic generation) analysis. Various machine learning techniques were utilized to implement these projects. This report provides an overview of the projects, outlines the basic concepts, and identifies ideas for future developments.

## INTRODUCTION

This article presents a general overview of recent advancements in the application of machine learning (ML) techniques for accelerator control and CHG (coherent harmonic generation) spectral analysis at the 1.5 GeV synchrotron light source DELTA [1–3].

The first study demonstrates the successful implementation of neural networks (NNs) for the correction of electron beam trajectories (orbits). These networks were trained using measured beam position data and corresponding corrector magnet strengths, showcasing competitive results but with fewer correction steps compared to conventional orbit correction methods.

In a parallel effort, ML-driven techniques were employed to control the working point (betatron tunes in both orbit planes). Therefore, a classical Proportional-Integral-Derivative (PID) system was replaced on a test basis by ML-trained NNs.

The paper further explores ML-based methods for adjusting storage ring chromaticity values. Gaussian Process Regressors (GPRs) and NNs were trained as surrogate models to predict optimal sextupole magnet settings using statistical and heuristic algorithms. Results showed significant improvements in chromaticity control within few iterations compared to manually performed settings.

Additionally, a novel approach was employed to optimize injection efficiency from the booster synchrotron BoDo to the storage ring. Through ML-trained predictive models in combination with heuristic optimization algorithms, injection settings were dynamically adjusted, resulting in improved electron transfer rates.

Lastly, the paper outlines the analysis of CHG-induced radiation spectra. CHG-generated ultrashort light pulses were examined through convolutional neural networks (CNNs) to

predict certain seeding laser settings from spectral measurements. This approach enabled the adjustment of the seeding laser parameters in order to optimize the spectra properties.

## ORBIT CORRECTION

At DELTA, the orbit is measured at 54 locations around the storage ring roughly every second using multiplexed beam position monitor (BPM) pick-up signals. Unwanted deviations of the beam position with respect to the ideal reference orbit are minimized by 56 corrector magnets, 30 for the horizontal plane ( $x$ ) and 26 for the vertical plane ( $y$ ). They are operated by a ‘slow’ (approx. 1 Hz) software feedback system applying numerical algorithms such as the SVD-method (singular value decomposition) [4] or the IPM-based (interior point-proximal method) software [5].




Figure 1: Orbit response at 108 BPM signals for the simulated  $x,y$ -coupled DELTA storage ring, calculated for 54 horizontal ( $x$ ) and 54 vertical ( $y$ ) deviations. In this example, 3000 randomly disturbed orbit vectors are shown. The data sets are used as inputs for supervised training of NNs.

As an alternative method, flat, non-deep NNs have been investigated for automatic orbit correction. ML-based studies on this topic, which initially were based on simulations (e.g., see Figs. 1 and 2), have then been successfully transferred to real accelerator operation. For this purpose, classical fully connected multi-layer feed-forward networks with more than 17.000 links were trained using supervised learning with measured beam position data and corresponding corrector magnet strengths. The trained networks subsequently served as inverse models for local and global beam position corrections. The supervised NN training was comparatively evaluated with various back-propagation learning algorithms (e.g., see Fig. 2). Finally, the performance of the ML-based beam control was benchmarked with conven-

\* detlev.schirmer@tu-dortmund.de

# MACHINE PROTECTION SYSTEM UPGRADE FOR A NEW TIMING SYSTEM AT ELBE

M. Justus, M. Kuntzsch, A. Schwarz, K. Zenker  
 Helmholtz-Zentrum Dresden-Rossendorf, Dresden, Germany  
 Z. Oven, L. Krmpotic, U. Legat, U. Rojec, Cosylab, Ljubljana, Slovenia

## Abstract

Running a C.W. electron accelerator as a user facility for more than two decades necessitates upgrades or even complete redesign of subsystems at some point. At ELBE, the outdated timing system needed a replacement due to obsolete components and functional limitations. Starting in 2019, with Cosylab as contractor and using hardware by Micro Research Finland, the new timing system has been developed and tested and is about to become operational. Besides the ability to generate a broader variety of beam patterns from single pulse mode to 26 MHz C.W. beams for the two electron sources, one of the benefits of the new system is improved machine safety. The ELBE control systems is mainly based on PLCs and industrial SCADA tools. This contribution depicts how the timing system implementation to the existing machine entailed extensions and modifications of the ELBE machine protection system, i.e. a new MPS PLC, and how they are being realized.

## ELBE TIMING SYSTEM UPGRADE

The electron accelerator at the ELBE Center for High Power Radiation Sources at HZDR [1] has provided beam-time as a user facility for more than two decades. Its unique feature is a 1 mA 35 MeV C.W. mode electron beam driven by a 235 kV thermal gun and SRF LINACs. It serves different sources of secondary radiation, i.e. infrared FELs (FELBE), THz sources (TELBE), a Bremsstrahlung facility (gELBE), as well as neutron (nELBE) and positron (pELBE) sources (Fig. 1). Having started as a one source one user facility, we could recently parallel beam options by a kicker, a scattering wire and reuse of the THz beam, but all of them only with limited range of beam power and pulse patterns. Over the past five years the ELBE SRF gun has become the standard electron source for THz and neutron beams [2]. It allows beam energies up to 40 MeV along with higher bunch charge and brightness, which implies further options of parallel beams.




Figure 1: ELBE facility layout.

The existing (hardware based) timing system for the thermal gun (injector 1) consists of a master oscillator at 13 MHz, a 26 MHz PLL, a gun clock pulse divider (ratio

1:2<sup>n</sup>, n = 0...8), and a single pulse generator to directly gate the gun from single bunches up to seconds long bunch trains at MHz down to mHz repetition rates (Fig. 2). A magnetic macro pulse generator in the 235 keV injector beamline gates the initial beam with repetition rates of 1 to 25 Hz. The distribution of pulses and triggers is done with hardware modules (PLLs, fanouts, fiber transceivers, ...). All components are managed by a beam control PLC that has the MPS master functionality. The SRF gun (IN2) timing patterns are currently generated by commercial off-the-shelf trigger sources in the range of 25 to 500 kHz.

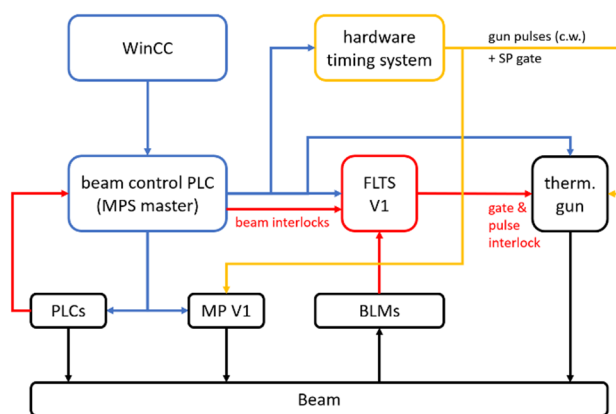


Figure 2: Existing timing system implementation with ELBE MPS for thermal injector.

Aged out electronics and insufficient flexibility of the existing timing system have led to the development of a new timing system for ELBE [3, 4]. It utilizes Micro Research Finland hardware based on the MTCA standard [5]. It is a modular, distributed system using event master modules (EVMs) to generate and send out events at a rate of 130 MHz and event receivers (EVRs) to build physical output signals from the received timing events. The gun pulse patterns become more variable, while SP and MP gating is preserved with enlarged parameter ranges. EVRs can be equipped with different universal output modules providing a variety of optical and electrical logic outputs for RF, beam diagnostics and user instrumentation. In future, the timing system shall support parallel user operation with very different beam patterns by use of RF kickers and parallel operation of two beam sources.

## MACHINE PROTECTION SYSTEM UPGRADE

We define the ELBE MPS as a set of functionalities, realized by different technical sub-systems (or parts of them), which are orchestrated by a central logic unit and system

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# DALI CONTROL SYSTEM CONSIDERATIONS

K. Zenker\*, M. Justus, R. Steinbrück,

Institute of Radiation Physics, Helmholtz-Zentrum Dresden-Rossendorf, Dresden, Germany

## Abstract

The Dresden Advanced Light Infrastructure is a future infrastructure under consideration at the Helmholtz-Zentrum Dresden-Rossendorf. In the current conceptional design phase, we are surveying different control system options. To benefit as much as possible from community experiences with different control systems, in 2023 a survey was conducted and participants from accelerator and light source facilities world-wide were invited. The results of that survey are presented and conclusions for our center are drawn.

## INTRODUCTION

The Dresden Advanced Light Infrastructure (DALI) [1] is part of the German national *Helmholtz Photon Science Roadmap* [2]. It will be a high-field source of intense terahertz radiation based on accelerated electrons and the successor of the Center for High-Power Radiation Sources (ELBE) operated at HZDR since 2002. In the current phase of DALI the conceptional design report is in preparation and there are ongoing considerations which control system to use best.

As a background note, due to historical reasons the ELBE control system is based on industrial automation products. Namely Siemens Simatic S7 PLCs [3, 4] on the field level and WinCC [5] as SCADA system for HMI and data archiving. Even when this system served the purposes of ELBE operation quite well during the last decades it lacks openness and interfaces for modern scientific tools like machine learning and artificial intelligence. Hence there is a strong request on the part of machine physicists and scientific users to search for control system alternatives for the new DALI facility.

To get an overview of the different control systems (CS), that are in operation at other accelerator based facilities we conducted a survey. With that we intended to benefit as much as possible from the community experience with different types of control systems.

## CONTROL SYSTEM SURVEY

In 2023 we conducted an online survey among invited participants. We invited 38 participants from 22 different particle accelerator facilities and light source facilities, respectively. We picked participants from facilities of similar size, similar machine types and field of research. Contacted participants are based in the USA, Brazil, Asia, and Europe. In total, we received 21 responses from 15 different facilities. Except for one American and one Asian facility all participating facilities are located in Europe. The results presented in the following are not representative for the whole community, but show a clear picture of the usage of different control

\* k.zenker@hzdr.de




Figure 1: CS categories and CS satisfaction.




Figure 2: Number of FTEs needed to maintain the CS.

systems in the community. It turned out that 46 % of the participants are using EPICS [6], 27 % are using Tango [7] and 27 % are using other control systems (e.g. DOOCS [8], TINE [9] or LabView [10]), as shown in Fig. 1a. Therefore, we decided to group the results into three control system categories – EPICS, Tango and Other. This allows to identify differences between those three categories.

In terms of control system operation time, we noticed that half of the control systems of participating facilities are in operation for more than 20 year and the other half is in operation less than 10 years. This observation is independent of the control system category.

The overall satisfaction with the control system in use, as shown in Fig. 1b, is quite high for EPICS (84 %) and Tango (94 %) and slightly lower for Other (70 %). The following control system features were mentioned as missing:

- High data rate support
- Synchronized data processing at a rate above 100 Hz
- Encryption of the network communication
- Support for multidimensional data type (> 2D)

## Development Efforts

As shown in Fig. 2 50 % of the EPICS participants and 67 % of the Other participants state that their CS is maintained by one to five full time equivalents (FTEs). 50 % of



# TECHNICAL DESIGN CONCEPT AND FIRST STEPS IN THE DEVELOPMENT OF THE NEW ACCELERATOR CONTROL SYSTEM FOR PETRAIV

R. Bacher, J. Behrens, T. Delfs, T. Tempel, J. Wilgen, T. Wilksen  
Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

## Abstract

At DESY, extensive technical planning and prototyping work is currently underway for the upgrade of the PETRAIII synchrotron light source to PETRAIV, a fourth-generation low-emittance machine. As part of this planned project, the accelerator's control system will also be modernized. This paper reports on the main decisions taken in this context and gives an overview of the scope of the development and implementation work.

## INTRODUCTION

With PETRAIII, DESY operates one of the best storage ring X-ray radiation sources in the world. PETRAIII is a 2300-metre-long storage ring feeding 24 user beamlines. It is operated either in brightness mode (480 equally distributed bunches, 120 mA stored beam) or in timing mode (40 equally distributed bunches, 100 mA stored beam), the latter a unique feature of PETRAIII. Research groups from all over the world use the particularly brilliant, intense X-ray light for a variety of experiments - from medical to materials research.

DESY plans to expand PETRAIII into an ultimate, high-resolution 3D X-ray microscope for chemical and physical processes. PETRAIV will extend the X-ray view to all length scales, from the atom size to millimetres. Researchers can thus analyse processes inside a catalyst, a battery or a microchip under realistic operating conditions and specifically tailor materials with nanostructures. PETRAIV offers outstanding possibilities and optimal experimental conditions for industry. Special emphasis is also placed on the sustainable operation of the facility.

PETRAIV will replace the PETRAIII facility and will be housed by the existing PETRAIII buildings. An additional experimental hall will provide space for another 18 user beamlines. A new synchrotron (DESYIV) will serve as booster between the existing electron source LINACII and PETRAIV. In addition, research has begun on the development and construction of a 6 GeV laser plasma injector.

In 2020, a preparatory phase for the future project PETRAIV was initiated with the aim of finishing a Technical Design Report by mid-2024. In the meantime, extensive planning work has been carried out with regard to buildings and facility infrastructure, specification of technical equipment and prototyping or beam-physical optimizations and simulations. In addition, the estimated project costs were determined and first steps were taken for the political decision-making process. The dismantling of PETRAIII and construction of PETRAIV will result in a two-year shutdown of operations at the PETRA complex.

## GENERAL CONCEPT

To consolidate and simplify the whole accelerator control system landscape at DESY and to take advantage of synergies between the accelerator facilities operated by DESY the control system of PETRAIV will closely follow the control system concept implemented at the European XFEL, which is a pulsed linear accelerator and free-electron laser. Therefore, the control system concept of PETRAIV will be adapted where necessary to the special needs of a storage ring X-ray radiation source. While many existing generic and basic software solutions from the European XFEL can be re-used for PETRAIV purposes, some have to be adapted or newly created to accommodate the different machine type, i.e. synchrotron vs. pulsed linear accelerator. These include, for example, the measurement of the beam positions or the fast orbit feedback system, but also the control of the magnetic power supplies. In addition, central control system services such as data acquisition and archiving, alarming or configuration management are to be extended and prepared for future requirements.

## CONTROL SYSTEM FRAMEWORK

The Distributed Object-Oriented Control System (DOOCS) [1] will form the basis of the accelerator control system. DOOCS is the leading accelerator control system at DESY. The initial development of DOOCS dates back to 1993. Since that time, it has steadily developed into a powerful, reliable and versatile control system. Recently, a road map was established to meet the increasing user demands over the next decade and to continue to keep pace with the rapid developments in IT and the controls community.

### Architecture

DOOCS follows an object-oriented design paradigm. Devices and data are objects. The basic entity is a device server representing some control system hardware or logic. The DOOCS naming scheme is hierarchical. The layout implements a three-tier approach (Fig 1). The front-end tier (resource layer) contains the device interface applications (device servers) that are connected to the accelerator devices through various field buses and hardware interfaces. The service tier (middle layer) provides common control system services (e.g. data archives) and cross-component or cross-system control functions (e.g. beam orbit). At the client tier (user interface layer) graphical user applications and tools for monitoring and operating the accelerator are located. The three tiers are interconnected through the network-based control and data bus of the control system.

# BOARD BRING-UP WITH FPGA FRAMEWORK AND ChimeraTK ON Yocto\*

J. Georg<sup>1†</sup>, A. Barker<sup>1</sup>, Ł. Butkowski<sup>1</sup>, M. Hierholzer<sup>1</sup>, M. Killenberg<sup>1</sup>,  
T. Kozak<sup>1</sup>, N. Omidajedi<sup>1</sup>, M. Randall<sup>1</sup>, D. Rothe<sup>1</sup>, N. Shehzad<sup>1</sup>, C. Willner<sup>1</sup>, K. Zenker<sup>2</sup>

<sup>1</sup>Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany

<sup>2</sup>Helmholtz-Zentrum Dresden-Rossendorf, Bautzner Landstr. 400, 01328 Dresden, Germany

## Abstract

This paper will showcase our experience in board bring-up using our Field Programmable Gate Arrays (FPGAs) Framework (FWK) and ChimeraTK, our C++ hardware abstraction libraries. The challenges involved in working with different FPGA vendors will be discussed, as well as how the framework and libraries help to abstract vendor-specific details to provide a consistent interface for applications. Our approach to integrating this framework and libraries with Yocto, a popular open-source project for building custom Linux distributions, will be discussed. We will show how we use Yocto's flexibility and extensibility to create a customized Linux image that includes our FPGA drivers and tools, and discuss the benefits of this approach for embedded development. Finally, we will share some of our best practices for board bring-up using our framework and libraries, including tips for debugging and testing. Our experience with FPGA-based board bring-up using ChimeraTK and Yocto should be valuable to anyone interested in developing embedded systems with FPGA technology.

## INTRODUCTION

New challenges in accelerator operations are frequently met with the introduction of custom-built hardware that can perform time-critical computing tasks on its own. While the bulk of these tasks are usually implemented in Field Programmable Gate Arrays (FPGA), they are not free from the need of user interaction – for example for run-time modification of algorithm parameters or status monitoring. Finally, the acquired data needs to be exposed to the surrounding control environment.

### Problem Description

Often enough, the user-facing part of such software is written from scratch for each new piece of hardware. This can lead to a considerable delay in the availability of the hardware to the introduction of the hardware into the control environment, slowing down feed-back loops between users of the hardware and the developers of the on-hardware algorithms. This can delay the uncovering of problems in the overall design happening well too late in the development cycles.

\* The authors acknowledge support from Deutsches Elektronen-Synchrotron DESY Hamburg, Germany, a member of the Helmholtz Association HGF.

† jens.georg@desy.de

On top of that, the resulting software is often tailored to the facility that developed the hardware, limiting the re-usability of the hard- and software elsewhere without extensive adaptation.

### Our Solution

The solution we are presenting below is tying together several building blocks that have been developed at the Accelerator Beam Controls (MSK) group in the past few years, as well as efforts from the global open source community.

We provide a generic software solution by taking advantage of the interoperability of the DESY FWK [1] with the ChimeraTK [2]. It exposes data to the control system with minor to no configuration effort, using industry standard protocols such OPC UA [3] and EPICS [4]. The data can then easily be consumed by all known scientific control systems or even industry-standard Supervisory Control and Data Acquisition (SCADA) systems, and hardware can be controlled likewise where necessary.

Furthermore, by taking advantage of the on-chip computing power of the new generation FPGAs that integrate a full System on Chip (SoC), it is possible to create a piece of hardware that can be used in quasi-standalone mode and directly plugged into the control system's network environment without having to write any code that runs on an external computer. This will cut down the time-to-machine considerably.

## OVERVIEW OF THE BUILDING BLOCKS

The final outcome of this work is not one stand-alone piece of software or configuration. Instead, it consists of several pieces that have already been in development for the past years. The major components of the setup shall be shown below.

### Firmware Generation & Hardware Description

The DESY FWK provides the developer with board support packages (BSP), reusable blocks for recurring firmware programming tasks and a unified build environment for the synthesis tools of different FPGA vendors. A complete description of the DESY FWK is beyond the scope of this publication. For further detail on the DESY FWK the reader shall be referred to [5].

One important artifact of the firmware generation process is the register map file. It provides a machine consumable description of hardware addresses to user-readable names and

# CHALLENGES OF THE COSY SYNCHROTRON CONTROL SYSTEM UPGRADE TO EPICS

C. Böhme, C. Deliege, M. Simon, M. Thelen, Forschungszentrum Jülich, Germany  
V. Kamerzhiev, GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany  
R. Modic, Ž. Oven, Cosylab d. d., Control System Laboratory, Ljubljana, Slovenia

## Abstract

The COSY (COoler SYnchrotron) at the Forschungszentrum Jülich is a hadron accelerator build in the early 90s, with work started in the late 80s. At this time the whole control system was based on a self-developed real-time operating system for Motorola m68k boards, utilizing, unusual for this time, IP-networks as transport layer. The GUI was completely based on Tcl/Tk. After 25 years of operation, in 20016, it was decided to upgrade the control system to EPICS and the GUI to CS-Studio, in order to e.g. allow a better automatization or automatized archiving of operational parameters. This was done together with Cosylab d. d. bit by bit while the synchrotron was in operation, and because of the complexity is still ongoing. The experiences of the stepwise upgrade process will be presented and a lessons learned will be emphasized.

## INTRODUCTION

### COoler SYnchrotron COSY

The COoler SYnchrotron (COSY) of the Forschungszentrum Jülich is a 184 m long racetrack-shaped synchrotron and storage ring for protons and deuterons from 300 MeV/c (protons) or 600 MeV/c (deuterons) up to 3.7 GeV/c. Built in are devices for stochastic as well as electron cooling. The stored ions can be polarized or unpolarized. COSY was commissioned in 1993.

### COSY Control System

Planning of the COSY control system date back to the mid 80s [1], which the original concept visualized in Fig. 1. The control system was self developed and based on VMEbus and G64 Hardware [2]. The communication was completely based on Ethernet 10BASE5 and 10BASE2. The communication protocol is a self developed system called "Single Command - Single Response" (SCSR), having a limited amount of commands and replies.

**Timing** The timing system is based on Ethernet 10BASE2 using only network hubs, as they deliver a deterministic delay, unlike network switches. The timing sender sends a broadcast every ms with an ID, timing receivers configured with that ID then send out a TTL pulse. With the used hardware the delay for each timing receiver is between 100 ns to 3.2 ms upon reception of the ID. The timing system only features timing relative to the cycle start. Each timing receiver is equipped with 2 timing outputs and 12 status bits output which can used as digital output with less accuracy.

### General

### Control System Upgrades




Figure 1: Original layout of the COSY control system, as described in [1].

**Central Clock** In order to synchronize all necessary devices, a central 8 MHz clock is distributed using as well Ethernet 10BASE2 hardware.

**Function Generators** The magnets at the synchrotron are controlled by self developed function generators. These controllers were designed as feed-forward controllers, because of the typical limited computing power of that time. Therefore the whole settings for one machine cycle are computed beforehand, and, after a trigger, the program is run with only the possibility of an emergency abort. As an option, the manual control of the controllers at a pre-defined time in the machine cycle, was implemented from the beginning.

## UPGRADE

### Motivation

With the JEDI experiment new requirements concerning the overall RMS beam orbit deviation were introduced [3]. Therefore an automated beam orbit control system had to be developed. Furthermore, other components were identified being in need to be upgraded or to be added. One example is the analog BPM readout electronics, whose signal offsets prevented an accurate position determination especially around the 0 position. The decision was made, instead of implementing the new sub-systems and features into the old control system, to upgrade the latter in a step-wise manner, in order to avoid long down-times of the machine. In addition the following considerations were taken into account:

- Add a logging and archiving mechanism
- Use software developed and supported by a larger community
- Software and compatible hardware available without protocol adaption

# ARCHITECTURE OF THE CONTROL SYSTEM FOR THE JÜLICH HIGH BRILLIANCE NEUTRON SOURCE

H. Kleines†, F. Suxdorf, A. Möller, S. Janaschke, A. Steffens, M. Glum, P. Kämmerling,  
J. Voigt, J. Baggemann, U. Rücker, P. Zakalek, T. Gutberlet, T. Brückel  
JCNS-2, Forschungszentrum Jülich, Jülich, Germany

O. Felden, A. Lehrach, R. Gebel, IKP-4, Forschungszentrum Jülich, Germany

Y. Beßler, R. Hanslik, D. Marshall, F. Palm, ZEA-1, Forschungszentrum Jülich, Jülich, Germany

H. Podlech, O. Meusel, IAP, Goethe-Universität Frankfurt, Germany

## Abstract

In the Jülich High Brilliance Neutron Source (HBS) project Forschungszentrum Jülich is developing a novel High Current Accelerator-driven Neutron Source (HiCANS) that is competitive to medium-flux fission-based research reactors or spallation neutron sources. The HBS will include a 70 MeV linear accelerator which delivers a pulsed proton beam with an average current of 100 mA to three target stations. At each target station the average power will be 100 kW generating neutrons for at least six neutron instruments. The concept for the controls system has been developed and published in the HBS technical design report. Main building blocks of the control system will be Control System Studio, EPCIS and Siemens PLC technology (for vacuum, motion, personnel protection...). The timing system will be based on commercially available components from Micro-Research Finland. The accelerator LLRF will rely on MTCA.4 developments of DESY that are commercially available, too. A small fraction of the control system has already been implemented for the new JULIC neutron platform, which is an HBS target station demonstrator that has been developed at the existing JULIC cyclotron at Forschungszentrum Jülich.

that fills the gap between neutron lab sources and high flux spallation sources. As shown in the planned layout in Fig. 1 it consists of an ion source, a LINAC and a multiplexer that distributes the pulsed proton beam via a high energy beam transport (HEBT) structure to three targets stations with more than 20 neutron instruments. The LINAC will provide 70 MeV protons at 100 mA peak current with different optimized pulse structures to each target station. Neutron production at the tantalum targets is based on a nuclear reaction, which yields a neutron flux at the neutron instruments similar to a medium flux research reactor, due to the optimized reflector/moderator setups very close to the targets.

## HBS CONTROL SYSTEM ARCHITECTURE

The HBS control system consists of components and tools, which connect all HBS equipment and present a homogenous and ergonomic interface to operators, engineers and physicists enabling safe and reliable operation of the HBS.

From a control system point of view, the HBS can be subdivided into the HBS machine consisting of ion source, accelerator system (LINAC, HEBT...), target stations and conventional facilities (cooling water, pressed air,...) which is responsible for neutron production, and the instruments using these neutrons for research. For the HBS machine a central control room is foreseen that is permanently manned with operators, whereas instruments are locally controlled by dedicated measurement scripts and programs. Ideally there will be one integrated control system for the HBS machine in order to reduce the development efforts and support a homogeneous user interface. Since the ion source and the conventional facilities will be provided by external partners, it is quite likely that they will come with their own control systems. In order to reduce the overall costs, these control systems will not be replaced but extended by gateway functions to the integrated HBS control system.

Operation of neutron instruments is typically fully automatic, requiring presence of instrument users or scientists only for measurement definition and start, for development and test of dedicated scripts or for sample change. From the perspective of the control system, instruments and HBS machine are only loosely coupled via the timing system,

## THE JÜLICH HIGH BRILLIANCE NEUTRON SOURCE PROJECT

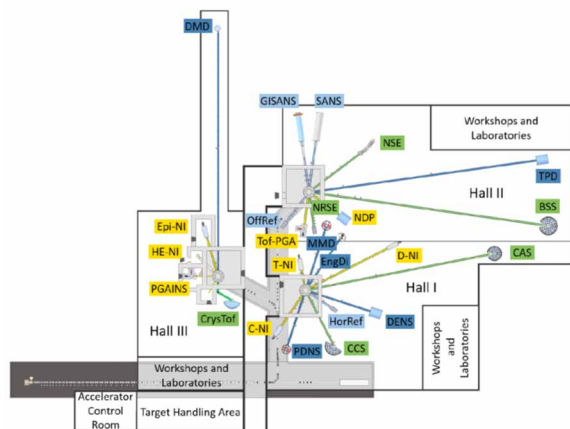


Figure 1: Layout of the planned HBS facility.

The Jülich HBS project [1] aims at the development of a novel High Current Accelerator-driven Neutron Source

† h.kleines@fz-juelich.de

# INTEGRATION OF AN OPTIMIZER FRAMEWORK INTO THE CONTROL SYSTEM AT KARA

C. Xu\*, E. Blomley, A. Santamaria Garcia, and A.-S. Müller  
Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany  
M. Zhang  
Princeton University, Princeton, United States

## Abstract

Tuning particle accelerators is not straightforward as they depend on a large number of non-linearly correlated parameters that drift over time. In recent years advanced numerical optimization tools have been developed to assist human operators in tuning tasks. A proper interface between the optimizers and the control system will encourage their daily use by the accelerator operators. In this contribution, we present our latest progress in integrating an optimizer framework into the control system of the KARA storage ring at KIT, which will allow the automatic tuning methods to be applied for routine tasks.

## INTRODUCTION

The Institute of Beam Physics and Technology (IBPT) at Karlsruhe Institute of Technology (KIT) operates several accelerators, including the Karlsruhe research accelerator (KARA) and the far-infrared linac and test experiment (FLUTE) [1], along with a compact storage ring cSTART and a plasma accelerator that are in the planning phase. Various machine learning methods have been implemented and studied at the accelerators of IBPT. For example, Bayesian optimization was used to optimize the injection efficiency into the storage ring [2], and reinforcement learning (RL) methods were employed to control the transverse motion and microbunching instability at KARA [3, 4]. In order to apply those methods at the accelerators, one often needs to implement specialized wrappers, such as the interface to the control system and results logging. They often don't have a user interface, which makes it more difficult for operators to use. In addition, the lack of common interfaces makes the software harder to maintain and requires the developer to be up-to-date with control system or operating system updates. Therefore, it is desired to have a generalizable framework for accelerator tuning tasks that can make common features and functionalities available across different particle accelerators.

One of the early examples of such a framework is OCELOT [5, 6], a multi-purpose software suite developed for the operation of European X-ray free-electron laser (EuXFEL). It contains a generic optimization framework that implemented several ready-to-use tuning algorithms like Nelder-Mead simplex, Bayesian optimization (BO) [7], and extremum seeking (ES) [8]. Nevertheless, many design choices of OCELOT have been made around FEL operation

and are more difficult to adapt to other tasks.

Recently, Badger [9] was developed as an alternative generic optimizer framework that focuses on efficiency and user-friendliness in the accelerator control room. After being tested at several facilities it has shown promising results in accelerator tuning tasks. At IBPT, we are continuously working on increasing the level of automation in daily operations and incorporating more standardized packages [10]. We decided to use Badger as the optimization framework for its simplicity and adaptability to new tasks. In this paper, we share our experience adapting Badger to the control system at KARA and demonstrate the first applications of automatic tuning during the accelerator commissioning process.

## INTEGRATION OF BADGER

### Badger Components

Here we briefly describe the the architecture of Badger and the steps needed to apply Badger to a new optimization task. Badger is written in Python with a modular approach, where the core package merely provides the basic functionality like a command line interface (CLI) and a graphical user interface (GUI) to control, monitor, and log the optimization process. It manages interfaces, environments, and algorithms through a plugin system, which allows easy modification and development of new plugins.

**Interface** In Badger, the communication with the control system is defined in an interface. Standard interfaces like *pyepics* and *pydoocs* are already implemented in Badger, providing basic read and write functionalities for the process variables (PVs). For the applications at KARA, we use an in-house developed Python package [11] for the channel access (CA) to our EPICS control system, which takes site-specific structures into account and uses the *caproto* [12] package as a backbone.

**Environment** The environment defines all the available variables and observables of interest for a subsystem of the accelerator. In the case of the EPICS control system, it contains mostly a list of the PVs that will be used in an optimization process.

**Algorithm** The algorithm is used to suggest the next point to evaluate in the optimization process. Since the interfaces with the machine are sufficiently abstracted, numerical optimization methods can be separately developed and easily incorporated into Badger. For example, the Xopt [13]

\* chenran.xu@kit.edu

# REFERENCE MEASUREMENT METHODS FOR PLANAR AND HELICAL UNDULATORS

S. Karabekyan<sup>†</sup>, European XFEL GmbH, Schenefeld, Germany

## Abstract

The modern permanent magnet undulators are usually equipped with motors that have integrated feedback electronics. These are essentially rotary encoders that indicate the position of the motor axis. In addition, undulators are also equipped with linear encoders that provide the absolute value of the gap between the magnetic structures or the position of the magnetic girders relative to the undulator frame. The operating conditions of undulators should take into account the risks of failure of electronic equipment under the influence of radiation. In case of encoder failure, the motor or encoder must be replaced. To avoid the need to return the undulator to the magnetic measurement laboratory, reference measurements are required to restore the position of the magnetic structure after replacement. In this article, reference measurement procedures for planar and helical APPLE-X undulators used at the European XFEL are presented.

## INTRODUCTION

Radiation damage is often a cause of failure of electronics installed in close proximity to a charged particle beam. This is the case for both synchrotron radiation sources and free-electron lasers. The causes of primary and secondary radiation can be errors in beam trajectory control as well as synchrotron radiation generated when the beam passes through deflecting magnets, undulators, wigglers, or other insertion devices (ID).

Such a situation occurred in particular with the helical undulators installed downstream of the SASE3 planar undulator system at the European XFEL [1]. A few months after the installation of these undulators in the tunnel, a considerable number of linear and rotary encoders were damaged by non-linear high-energy Compton scattering, as it turned out later. This radiation was apparently the result of scattering of spontaneous radiation from the planar undulators on the walls of the aluminum vacuum chamber of the helical APPLE-X undulator. As a result of this radiation damage to the feedback system, it was not possible to control the magnetic structures of the undulator [2].

To restore the functionality of the APPLE-X undulators, they had to be removed from the tunnel and the damaged encoders replaced. Since the position of the magnetic structures was lost after the encoders were replaced, new magnetic measurements had to be performed and the undulator had to be re-calibrated. Particularly, the magnetic measurements showed that the radiation that disabled the encoders did not affect the magnetic properties of the undulators. Thus, if reference measurements of the position of the magnetic structures had been made before the problems with

the encoders occurred, it would not have been necessary to remove the undulators from the tunnel. Indeed, after replacing the encoder or the motor, the magnetic structure could be returned to the position determined by the reference measurements.

The methods proposed below can be used to perform reference measurements on planar insertion devices such as an undulator or phase shifter, as well as on helical undulators. In the case of the helical undulator, two methods are proposed for measuring the undulator gap and the longitudinal position of the magnetic structures.

## REQUIREMENTS FOR REFERENCE GAP MEASUREMENT FOR PLANAR INSERTION DEVICES

If the construction of the ID, i.e., undulator or phase shifter, is not done in vacuum, then the vacuum chamber of the electron beam is located between the magnetic structures. In this case, the direct gap measurement must consider this, since the goal is to determine the position of the magnetic structures without retrieving the ID from its position. The proposed setup consists of two flat gap measurement sensors. Each gap measurement sensor consists of two capacitive sensors located on opposite sides of a printed circuit board (PCB). One of the arguments in favor of the measurement principle is the fact that the undulator or phase shifter is located in a tunnel where the temperature changes over time do not exceed a few hundred millikelvin.

Under constant temperature conditions, changes in the thickness of the vacuum chamber are negligible. In the case of the European XFEL, the thickness of the vacuum chamber is 9.6 mm. This means that a 1 °C change in the temperature of the vacuum chamber from its nominal value of 21 °C results in a 0.22 μm change in its thickness. The same can be said about the change in the thickness of the sensors. The thickness of the printed circuit board on which the capacitive sensors are mounted is 0.8 mm. This means that the thickness of the sensors changes by 0.015 μm for a temperature change of 1 °C, i.e. 0.03 μm for two sensors. Therefore, if the temperature of both sensors and the vacuum chamber changes by 1 °C, the total thickness change corresponds to 0.25 μm.

For free-electron lasers, a necessary condition for achieving lasing is that the spread of the K parameter of the undulator does not exceed the Pierce or FEL parameter. Thus, for high-energy free-electron lasers, the  $\Delta K/K$  must be less than  $3 \times 10^{-4}$ . The calculation of the boundary condition for the change of the gap is presented in chapter 2.3 of the article [3]. From this it can be seen that the change in the undulator gap should not exceed 3 μm. Therefore, the reproducibility of the relative reference measurements should not exceed this value.

<sup>†</sup> suren.karabekyan@xfel.eu.

# APPLYING MODEL PREDICTIVE CONTROL TO REGULATE THERMAL STABILITY OF A HARD X-RAY MONOCHROMATOR USING THE KARABO SCADA FRAMEWORK

M. A. Smith\*, G. Giovanetti, S. Hauf, I. Karpics, A. Parenti, A. Samadli,  
L. Samoylova, A. Silenzi, F. Sohn, P. Zalden, European XFEL, Schenefeld, Germany

## Abstract

Model Predictive Control (MPC) is an advanced method of process control whereby a model is developed for a real-life system and an optimal control solution is then calculated and applied to control the system. At each time step, the MPC controller uses the system model and system state to minimize a cost function for optimal control. The Karabo SCADA Framework is a distributed control system developed specifically for European XFEL facility, consisting of tens of thousands of hardware and software devices and over two million attributes to track system state.

This contribution describes the application of the Python MPC Toolbox within the Karabo SCADA Framework to solve a monochromator temperature control problem. Additionally, the experiences gained in this solution have led to a generic method to apply MPC to any group of Karabo SCADA devices.

## MONOCHROMATORS AT XFEL

European XFEL [1] operates three beamlines capable of delivering hard and soft X-rays. For the hard X-ray beamlines, silicon monochromators are used to select the pass band of X-ray energies that continue through to the instrument. The monochromator works by using Bragg's law, which gives the relationship between the incident angle of X-rays on a crystal lattice and the reflected angle. The position of the crystals is adjustable via motors, in order to control the incident angle of the X-rays to the silicon crystal's lattice structure. Python MPC Toolbox [2] was used to solve a monochromator temperature control problem.

Thermal drift in a monochromator causes a drift of the transmitted photon energy. Silicon is commonly used for Hard X-ray Monochromators due to its good commercial availability and point of zero thermal expansion around 125 K [3]. To mitigate the impact of temperature jumps caused by the X-ray pulse trains at European XFEL, the monochromator temperature has to stay just below this temperature. A CAD drawing of the two crystals mounted inside a monochromator is shown in Fig. 1. The orange arrow in the drawing shows the path of the X-ray beam through the two crystals of the monochromator in a 2-bounce configuration. To cool the overall monochromator, the silicon crystals are attached to copper plates which are in turn connected to a single cryogenic cold head to cool the plates down to around 100 K. Each crystal also has a local heating element for fine control of each crystal's individual temperature, along with




Figure 1: CAD drawing of the X-ray monochromator, showing the path of the X-ray beam through the two silicon crystals of the monochromator. Tx and HTx identify the temperature sensors and heater elements respectively [4].

a temperature sensor for feedback. Because each crystal is connected to a common cold head, heat applied to one crystal will have a time-delayed effect on the temperature of the other crystal via conduction.

In addition to this, the first crystal in the X-ray path is heated by the X-ray beam. This causes changes in the photon energy after the first Bragg reflection, which in turn causes a change in angle for the Bragg reflection from the second crystal. Tight thermal regulation is very important for ensuring stability of photon energies and the X-ray beam position.

## Modeling the Balance of Energies

The transfer of heat from one crystal to the cryo head can be modeled according to the thermal conduction Eq. (1), where  $U$  is the overall heat transfer coefficient of the cryo head,  $A$  is the cross sectional area, and  $T_{\text{crystal}} - T_{\text{cryo}}$  is the temperature difference between the crystal and the cryo head.

$$P_{\text{cryo}} = \frac{Q_{\text{cryo}}}{dt} = -U * A * (T_{\text{crystal}} - T_{\text{cryo}}) \quad (1)$$

Additionally, the X-ray beam itself imparts heat energy onto crystals it encounters and needs to be taken into account when setting the heater outputs. This power reading is readily available in the control system from an X-ray gas monitor.

\* michael.smith@xfel.eu

## GeCo: THE ELETTRA 2.0 BEAMLINE CONTROL SYSTEM

V. Chenda, A. Abrami, R. Borghes, A. Contillo, L. Cristaldi, M. Lucian, M. Prica, R. Pugliese, L. Rumiz, L. Sancin, M. Turcinovich, Elettra-Sincrotrone Trieste S.C.p.A., Trieste, Italy

### Abstract

The Elettra Synchrotron, located in Italy near Trieste, has been operating for users since 1994 being the first third generation light source for soft X-rays in Europe. To stay competitive for world-class photon science, a massive upgrade of the storage ring has been planned in 2025. The goal is to build an ultra-low emittance light source with ultra-high brilliance in the same building as the present storage ring. The downtime for installation and commissioning of Elettra 2.0 will last 18 months. In this plan, 20 of the present beamlines should be upgraded and 12 new beamlines are scheduled to be built. In this scenario, also the original beamline interlock and personnel safety systems are going to be upgraded using state of the art technologies. Siemens PLCs will be used for low level control, while higher level applications will be developed using the Tango framework. This work presents and describes the architecture of the future Elettra 2.0 beamline control system named GeCo, Gestione e Controllo in Italian.

### INTRODUCTION

Elettra-Sincrotrone Trieste is an Italian non-profit company of national interest that manages a multidisciplinary research centre with two light sources: the 3rd-generation synchrotron source Elettra, in operation since 1993 [1], and the seeded free-electron laser source FERMI, in operation since 2010 [2]. The light from Elettra feeds 28 beamlines with experimental stations that offer access to state of the art instrumentation for the most advanced spectroscopy, scattering and imaging techniques.

After nearly 30 years of operation, the Elettra synchrotron radiation source will be replaced by the new Elettra 2.0 4th generation light source. The Elettra 2.0 storage ring will employ a symmetric six-bend enhanced achromat lattice and will operate predominantly at 2.4 GeV. The brightness will increase ~35-fold at 1 keV and ~160-fold at 10 keV. The coherent fraction is expected ~30% and ~3% at 1 and 10 keV, respectively. These parameters of Elettra 2.0 will boost the spatial, energy and temporal resolution of all the experimental end-stations. The project is ongoing and the new storage ring and the first beamlines should start operation in 2026.

Elettra 2.0 will have up to 32 beamlines (Fig. 1): 20 of the present ones should be upgraded, and 12 new are scheduled to be built. For all of them a new control system infrastructure has been designed using state of the art technologies. The project involves the Interlock System, the Personnel Safety System and the Instrumentation Control System. The next sections provide the reader with a technical overview of the new beamline interlock system GeCo, a highly optimized and smart solution that allows to

operate on the beamline components and the vacuum elements in safe conditions.




Figure 1: Elettra 2.0 beamline layout.

### GECO INTERLOCK SYSTEM

#### Introduction

The present beamline control system (BCS) is based on VME industrial PCs running a LynxOS real time operating system. Beamline instrumentation is interfaced through digital and analog I/O boards or via serial or Ethernet communication protocols. The BCS is written in C language and is composed by a set of software drivers and a configuration ASCII database. It implements software interlocks and automatic actions that allow to operate the beamline components in safe conditions.

The BCS system has proven to be robust and reliable but, after almost 30 years, it suffers from obsolescence problems and a lack of spare parts. The new beamline interlock system GeCo, based on PLC (Programmable Logic Controller) technology, is the natural evolution of the present system. It has been designed and built to directly control the existing beamline components, getting rid of obsolete interface tools. Its name derives from two Italian words: Gestione (management) and Controllo (control).

#### Hardware Description

Each GeCo interlock system is based on a master PLC (Siemens S7 1500 – CPU 1513 - 1PN) and one or more peripheral slave units (Siemens ET200MP - IM155-5 PN ST3-1). A specialized 6U crate has been engineered in-house for hosting the PLC instrumentation (Figs. 2 and 3).



# NEW DEVELOPMENTS FOR eGiga2m HISTORIC DATABASE WEB VISUALIZER

L. Zambon\*, R. Passuello, Elettra Sincrotrone Trieste, Trieste, Italy

## Abstract

eGiga was an historic database web visualizer from 2002. At the beginning it was connected to a proprietary database schema, support for other schemas was added later, for example HDB and HDB++. eGiga was deeply refactored in 2015 becoming eGiga2m. Between 2022 and 2023 a few improvements have been made, among them, optimization of large data extraction, improvement of images and pdf exports, substitution of 3D chart library with a touch screen enabled one; the addition of: logger status info, a new canvas responsive chart library, adjustable splitter, support for TimescaleDB and HDF5 data format, correlations and time series analysis, and ARIMA (AutoRegressive Integrated Moving Average) forecast.

## INTRODUCTION

eGiga2m is a web application for displaying time series as charts. It was designed to be easily embedded in other applications and also to include other web pages. In order to get this result, it is possible to display only the chart without any menu, bar or configuration dialog and all configuration parameters are reported in the address bar (they are HTTP GET parameters).

There are also a few configuration parameters which are configurable only by the address bar, for example multiple X axes as shown in Fig. 1.



Figure 1: Multiple X axes.

Time series are provided by several data sources such as databases (for example HDB++), CSV files or numerical elaboration micro services.

The core of eGiga2m is written in JavaScript and doesn't receive data directly from a database, but it uses a web service which serves JSON encoded data.

JSON data schema is defined in JSON according to standards defined by json-schema.org [1].

Source code can be downloaded from <https://gitlab.elettra.eu/puma/client/egiga2m> eGiga2m architecture is reported in Fig. 2.

\* lucio.zambon@elettra.eu

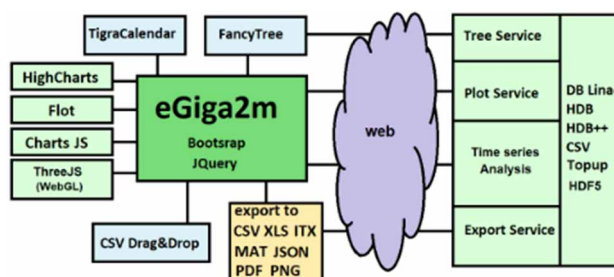


Figure 2: eGiga2m architecture.

Recent improvements are described in the following paragraphs.

## NEW FEATURES

### Canvas Chart Library

A new chart library, Chart.js [2], was added recently. This library isn't based on SVG but on HTML5 Canvas. SVG was defined by the W3C consortium and integrates perfectly in HTML, CSS and JavaScript, but if the chart requires a certain complexity, the browser must manage a very complex DOM and it can cause a sensible slowness. Using a canvas directly is significantly faster and makes it possible to be truly responsive; i.e. when the size of the charting area is modified the chart adapts instantly to the new dimensions.

### Adjustable Separator

In the main page on the left side there is a configurator dialog composed by the start and stop time selector and the time series tree selector; on the right side there is the chart. Between the configurator dialog and the chart there is a separator bar which was made adjustable recently. It is more useful if used with Chart.js because the chart adjusts instantly moving the bar. On mobile the ideal setup is obtained by rotating the device horizontally and adjusting the separation bar as shown in Fig. 3.



Figure 3: eGiga2m on a mobile device.

### Optimization of Large Data Extraction

Some PHP servers extract data from Databases and the whole data structure is returned in a unique JSON data stream; a quite fast way to produce it is to load all data in

# TOUCH-SCREEN WEB INTERFACES

L. Zambon\*, A. Apollonio, R. Passuello, Elettra Sincrotrone Trieste, Trieste, Italy

## Abstract

A touch screen (mobile or not mobile) has a significant impact on the kind of interaction between humans and control systems. This paper describes the development of some widgets and applications based on touch screens. The technologies used (for example PUMA, JavaScript and SVG) will be discussed in detail. Also a few tests and use-cases will be described compared with normal screens, mouse and keyboard interaction.

## INTRODUCTION

A prototype of touch screens was patented in 1946 but have become familiar to users only since the diffusion of smartphones and tablets.

Our development is based mainly on 2 JavaScript libraries: Hammer.js [1] for 2D applications, Three.js [2] for 3D models.

Data are taken from the Control System thorough PUMA [3]

## 2D APPLICATIONS

Hammer.js is a JavaScript library which efficiently captures gestures. Gestures like tap pinch etc are captured very quickly and are assimilated to the events produced by mouse. Any gesture also produces the event start, continue and stop. All aspects can be effectively customized. By default all gestures have a correspondence in mouse events (for example mouse wheel or right click) so the same web interface can be used on both touch and not touch screens. Hammer.js was used in our project in conjunction with SVG (Scalable Vector Graphics) [4]. SVG has been developed by W3C since 1999. It is an XML based 2D vectorial graphic format, it recognizes paradigms very similar to CSS and elements can be addressed by JavaScript like HTML elements.

SVG can be considered as "the" graphic extension of HTML. Graphic elements can be grouped in a symbol. A symbol can be repeated many times and can be customized (for example the filling color) as a unity. Every object can be translated by a transformation matrix or by a simple transformation command like translate or rotate. A significant trick in building circular design was to generate each element in the upper central position and then rotate it by a variable angle.

### The Knob

The "Knob" (Fig. 1) is a component which allows setting a value with a user experience similar to a physical knob. It is composed of two wheels with different colors.

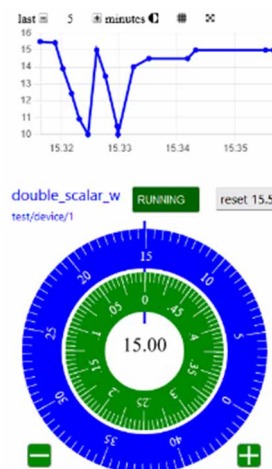


Figure 1: The Knob.

The knob is capable of more than 100 settings per second, but most control system servers cannot accept such a setting rate, so the knob implements a configurable throttling period by default 500 ms. It can be configured to set values only touching the central button; but the main task is a continuous flux of settings. The user can change the scale of the inner wheel (as well as almost all colors and sizes are configurable). There is an optional chart, a reset to initial value button and a status indicator. The inner wheel scale can be adjusted by pressing the plus or minus green buttons on the bottom. In this case the external wheel disappears until the original scale is restored.

### The Dodecagon

The "Dodecagon" (Fig. 2) is an application in which each of the 12 sections can be selected. Users can switch ON, OFF and standby High Voltage power supply and/or Feedbacks and set a value using the knob component.

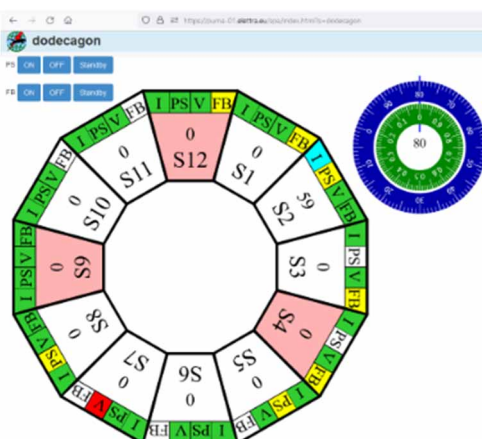


Figure 2: The Dodecagon.

\* lucio.zambon@elettra.eu

# STATUS OF VACUUM CONTROL SYSTEM UPGRADE OF ALPI ACCELERATOR

L. Antoniazzi, G. Savarese, C. Roncolato, A. Conte, INFN/LNL, Legnaro, Italy

## Abstract

The vacuum system of ALPI (Acceleratore Lineare Per Ioni) accelerator at LNL (Laboratori Nazionali di Legnaro), including around 40 pumping groups, was installed in the '90s. The control and supervision systems, composed by about 14 control Racks, were developed in the same period by an external company, which produced custom solutions for the HW and SW parts. Control devices are based on custom PLCs, while the supervision system is developed in C and C#. The communication network is composed of multiple levels from serial standard to Ethernet passing true different devices to collect the data. The obsolescence of the hardware, the rigid system infrastructure, the deficit of spares parts and the lack of external support, impose a complete renovation of the vacuum system and relative controls. In 2022 the legacy high level control system part was substituted with a new one developed in EPICS (Experimental Physics and Industrial Control System) and CSS (Control System Studio). After that, we started the renovation of the HW part with the installation and integration of two new flexible and configurable low level control system racks running on a Siemens PLC and exploiting serial server to control the renewed pumping groups and pressure gauges. The plan for the next years is to replace the legacy hardware with new one retrieving spare parts, provide service continuity, improve PLC software and extend the EPICS control system with new features. This paper describes the adopted strategy and the upgrade status.

## ARCHITECTURE OF THE LEGACY SYSTEM

The ALPI accelerator is the superconducting linac of the LNL, it is composed of 22 cryostats which were installed in the '90s together with their ancillary systems [1-4]. Most of the original hardware of the vacuum system, which is based on magnetic bearing turbomolecular pumps, have been maintained up to now and is still running. The control and supervision systems was deployed in the same years by a contractor company (TELEMA Computers), which fully developed dedicated HW and SW solutions from the field level up to the communication network and the HMI (Human Machine Interface).

The control infrastructure consists of three main parts: the vacuum controller, the collector and the supervisor. The controller represents the hardware mounted in the 14 racks placed along the accelerator. Each rack, also named VCS (Vacuum Control System) rack, can control the three pumping groups of a couple of cryostats and the diagnostic box between them (see Fig. 1). It is operable locally or remotely from the console. VCS rack is made up of: vacuum instru-



Figure 1: One of the TELEMA racks mounted in ALPI.

mentation (vacuometers and pump controllers); a custom serial multiplexer (MUX) and two crates containing ad-hoc microprocessor control boards called PLC (Programmable Logic Controller). Both MUX and PLC are based on Zilog Z80 microprocessor. The MUX collects the information from vacuometers and pump controllers, and forwards a subset of those data over a RS485 4-wires line. The two PLCs manage the valves and pumps actuation, implements the Machine Protection System (MPS) procedures, and manages high level procedures. The first named "Power Additional" controls the pumps, the valves connected to the pumping groups and the beam line; while the second "Additional" takes care of the valves connecting the cryostats to the centralized lines used for forevacuum, He and N2 inlet procedures. Both type of PLCs receive and send information over an RS-485 4-wires line, which is separate by the one for MUXs.

Each RS-485 serial networks for PLCs and MUXs extends along a section of the accelerator (low and high energy), so a total of 4 RS-485 networks are present. Two Windows PC (one for PLCs and one for MUXs) act as collector servers, receiving the data transmitted over the serial lines and forwarding them on the ethernet network up to the main console in the control room. The serial to ethernet conversion and

# INTEGRATING EPICS CONTROL SYSTEM IN VR ENVIRONMENT: PROOF OF CONCEPT

L. Pranovi<sup>†</sup>, M. Montis, INFN-LNL, Legnaro, Italy

## Abstract

Preliminary activities were performed to verify the feasibility of Virtual Reality (VR) and Augmented Reality (AR) technologies applied to nuclear physics laboratories, using them for different purposes: scientific dissemination events, data collection, training, and machine maintenance. In particular, this last field has been fascinating since it lets developers discover the possibility of redesigning the concept of the Human-Machine Interface. Based on the experience, it has been natural to try to provide to the final user (such as system operators and maintainers) with all the set of information describing the machine and control system parameters. For this reason, we tried to integrate the accelerator's control system environment and VR/AR application. In this contribution, the integration of an EPICS-based control system and VR environment will be described.

## PRELIMINARY ACTIVITIES

In our initial investigation [1] for the introduction of Virtual and Augmented Reality (VR and AR) technologies, we used the Microsoft<sup>®</sup> HoloLens 2 display [2] as an augmented reality (AR) device. This device combines waveguides and light projectors within its enclosure above the brow, utilizing laser light to illuminate the display.

In specialized environments like nuclear plants, the integration of such a device can enhance performance and assist operators in various ways, such as facilitating easy access to information during equipment maintenance or introducing an innovative Human-Machine Interface (HMI) that extends traditional control panels with virtual interfaces (as visible in Fig. 1).



Figure 1: Human-Machine Interface through AR technology for cyclotron supervision.

The primary use case involved the development of a virtual console for operators working with the cyclotron apparatus during its commissioning at LNL. This ongoing renewal process, conducted in parallel with regular operations, aimed to enhance the high-level control interface of the apparatus. In this context, augmented reality (AR) screens are generated alongside the standard control monitors. This setup allowed users to view the critical screens for monitoring and supervising the machine in augmented reality, while the standard computer-based HMI remains available for interaction, including issuing commands and executing procedures.

Conflicting news about the continuity of support provided by Microsoft during 2021-2022 moved our focus on HMI with the employment of VR devices such as Oculus and Valve products.

## CONTROLS AT INFN-LNL

The SPES project [3] is currently under construction at INFN-LNL (Laboratori Nazionali di Legnaro) and involves the integration of existing accelerator systems with a new setup comprising the primary beam and the ISOL target. To efficiently control this project, EPICS [4] was selected as the primary Control System framework. Consequently, a transition from the previous control system to the new framework is imperative to adapt the existing system for use in the upcoming facility.

In this moment several systems are under EPICS environment, in particular most of the original accelerator lines through a migration campaign and part of the new upcoming lines which are directly developed in EPICS. However not all the apparatus are EPICS compliant (i.e., cyclotron apparatus – visible in Fig. 2).

As consequence, to reach the goal of extending VR technology to controls supervision using Unreal Engine tool, data exchange between controls and VR environment requires a certain grade of flexibility and the possibility to interface using different protocols.



Figure 2: The 70 MeV cyclotron apparatus used as source for the proton line in the ISOL facility.

<sup>†</sup> lorenzo.pranovi@lnl.infn.it

# SAFETY SYSTEM FINAL DESIGN FOR THE ITER HEATING NEUTRAL BEAM INJECTOR TEST BED

A. Luchetta<sup>†1</sup>, M. Battistella, S. Dal Bello, L. Grando<sup>1</sup>, M. Moressa<sup>1</sup>, Consorzio RFX, Padova, Italy  
C. Labate, F. Paolucci, Fusion for Energy, Barcelona, Spain  
J. M. Arias, ITER Organization, St. Paul-Lez-Durance, France  
<sup>1</sup>also at CNR – Institute for Plasma Science and Technology, Padova, Italy

## Abstract

MITICA, the test bed for the ITER heating neutral beam injector, will use an extensive computer-based safety system (MS) to provide occupational safety. The MS will integrate all personnel safety aspects. After a detailed risk analysis to identify the possible hazards and associated risks, we determined the safety instrumented functions (SIFs), needed to mitigate safety risks, and the associated Safety Integrity Levels (SIL), as prescribed in the IEC 61508 technical standard on functional safety of electrical/electronic/programmable electronic safety-related systems. Finally, we verified the SIFs versus the required SIL. We identified about 50, allocated to SIL2 and SIL1.

Based on the system analysis, we defined the MS architecture, also considering the following design criteria: Using IEC 61508 and IEC 61511 (Safety instrumented systems for the process industry) as guidelines; Using system hardware to allow up to SIL3 SIFs; Using certified software tools to allow programming up to SIL3 SIFs. The SIL3 requirement for hardware/software derives from the need to minimize the share of the failure probability, thus allowing maximum share to sensors and actuators.

The paper presents the requirements for the MITICA safety systems and the system design to meet them. Due to the required system reliability and availability, the hardware architecture is fully redundant for all components involved in safety functions. Given the requirement to choose proven solutions, the system implementation adopts industrial components.

## INTRODUCTION

ITER requires powerful neutral beam injectors (HNB), for plasma additional heating and current drive [1]. The heating beams are produced through electrostatic acceleration of H<sup>+</sup> or D- (Deuterium) ions, up to 1 MeV, followed by ion neutralization. Atoms need to be neutral to penetrate the high magnetic field surrounding the plasma. Negative ions are used since their neutralization efficiency at ion energy exceeding 100 keV is much greater than that of positive ions.

HNBs with the ITER requirements in terms of beam power (16.5 MW), ion energy (1 MeV), accelerated beam current (40 A), divergence (7 mrad), and pulse length (3600 s) do not exist and, therefore, the HNB development is carried out through a dedicated facility, called the Neutral Beam Test Facility (NBTF) [2], aimed at developing

the ITER full-size HNB test bed [3], called MITICA, and testing it up to nominal performance.

This paper reports the design of the MITICA safety system. After a brief discussion on safety in MITICA, we summarize in the paper the system analysis executed and briefly discuss the apparent antinomy of using proven and also innovative solutions. We then present the system requirements, the applied design concepts, the proposed hardware architecture, and the approach followed for software development. Finally, we discuss the planned system commissioning.

## OVERVIEW OF MITICA SAFETY

While ITER HNBs require nuclear and occupational safety, MITICA is only required to ensure personnel safety, as the nuclear risk does not demand for a specific nuclear-class safety system. The hazards in MITICA are mainly related to high voltage, explosive and asphyxiating gasses, radiation, fire, and high pressure coolants.

We decided to develop a dedicated system, called MITICA safety system (MS) and based on programmable electronics, to manage and coordinate all safety issues at MITICA experiment level. This choice was also pushed by the Regulatory Authority, which requires a centralized safety system to issue the license to operate the plant.

### Functional Safety

The purpose of MS is to reduce the risk of serious injury to personnel to an acceptable level. Functional safety of programmable electronics is the subject of the technical standard IEC 61508, which is a general standard accompanied with specific standards dedicated to specific applications, such as oil, automotive and process industry (IEC 61511). As we can figure out MITICA as a process, we decided to base the MS development on the technical standards IEC 61508 and IEC 61511. Adhering to these standards during the whole life-cycle of the safety systems provides a methodology for the risk analysis and mitigation and for system design, verification, implementation, testing, and maintenance. It also supports the developers in defining the correct level of risk mitigation and improves the quality of the final product.

## SAFETY ANALYSIS

Safety analysis must be carried out with the support of safety experts as this is a very sensitive activity. The purpose of the analysis is to identify all possible hazards, to quantify the risks, to define the required safety instrumented functions (SIF) to mitigate the unacceptable risks,

\* Work supported by Fusion for Energy, ITER Organization, and EUROfusion

<sup>†</sup> adriano.luchetta@igi.cnr.it

# CONTROL AND DATA ACQUISITION SYSTEM UPGRADE IN RFX-MOD2

G. Martini\*, G. Manduchi<sup>1</sup>, A. Luchetta<sup>1</sup>, C. Taliercio<sup>1</sup>, A. Rigoni, N. Ferron, P. Barbato<sup>1</sup>  
Consorzio RFX, Padova, Italy

<sup>1</sup>also at CNR Institute for Plasma Science and Technology, Padova, Italy

## Abstract

RFX-mod2, currently under construction at Consorzio RFX, is an evolution of the former RFX-mod experiment, with an improved shell and a larger set of electromagnetic sensors. The extended set of sensors allows for exploring a wide range of plasma control schemes but also poses a challenge to its Control and Data Acquisition System (CODAS). RFX-mod2 CODAS is required to provide the high-speed acquisition of a large set of signals and their processing in the Plasma Control System (PCS). The PCS must provide a sub-millisecond response to plasma instabilities. While brand-new solutions provide for the acquisition of electromagnetic signals, involving Zynq-based ADC devices, other parts of the CODAS system have been retained from the former RFX-mod CODAS.

## INTRODUCTION

RFX-mod2 is an upgrade of RFX-mod [1, 2] with a modified shell and mechanical structure to enhance plasma-shell proximity and improve plasma control. The main goal of RFX-mod2 is an improved plasma that will be achieved thanks to a much larger number of electromagnetic (EM) probes (1500) than that (800) used in RFX-mod. The higher number of signals to be acquired and possibly used for real-time plasma control significantly impacts the requirements of the Control and Data Acquisition System (CODAS), consequently requiring new extensions and improvements.

Moreover, even if the new CODAS inherits the overall architecture of the previous one, it must cope with the obsolescence of several hardware and software components. Luckily, several architectural choices in software and hardware made almost 20 years ago for RFX-mod, are still valid. Hence, the new system will maintain them. Indeed, the use of Linux is now even more widespread than at the time of initial CODAS development, and the software frameworks MDSplus [3] and MARTE [4], used for data acquisition and real-time control, respectively, are currently in use in many fusion experiments and actively maintained.

As for the hardware, the new system will maintain the CompactPCI (CPCI) technology, adopted for a large part of the data acquisition system, since it is still widely used. However, some legacy systems still using CAMAC will be replaced by more modern hardware. Similar considerations hold for plant control, where Siemens S7 400 PLCs used in some plant systems will be retained, while legacy S5 PLCs used in other subsystems will be replaced.

\* giulio.martini@igi.cnr.it

## ELECTROMAGNETIC SIGNALS

Acquisition of signals from EM probes represents an important CODAS function in fusion because it allows the reconstruction of the plasma equilibrium profiles that represent the first step in the analysis of plasma performance and in the plasma control itself. While the former is usually performed offline, relying on the signal database built by CODAS (called pulse file in the Fusion research jargon), the latter requires real-time data acquisition and control computation to meet the strict temporal requirements of the Plasma Control System (PCS), which require a sub-millisecond system response.

Two independent systems performed the corresponding functions in RFX-mod: CPCI-based solution for non-real-time data acquisition and VME for real-time data acquisition. It involved costly hardware duplication, complicating data acquisition management and hardware maintenance.

Moreover, most signals derived from EM probes must be integrated over time to derive the measured magnetic field. In RFX-mod, signal integration was performed before acquisition via hardware.

However, Plasma Control often requires the time derivative of the input signals (e.g. to perform PID control), which can be achieved either by (1) digitally differentiating the integrated original signal or (2) duplicating acquisition by additionally acquiring the original EM signal. The second solution is preferable for control as it avoids any error derived from a digital differentiation, but requires duplication of the acquisition channels.

## FPGA-Based Architecture

The new FPGA-based architecture currently under development for RFX-mod2 provides the definitive solution:

1. The ADC board hosting the FPGA will be used for (1) high-speed data acquisition (up to 1 MHz) with DMA transfer to local memory and (2) low-speed (10 kHz) data sampling and streaming for plasma control.
2. The same FPGA will also perform digital integration to provide real-time integrated signals, avoiding the need for an additional front end for analog integration.

Long-lasting experiments require sustained streaming for data acquisition; conversely, due to the foreseen pulse duration of RFX-mod (up to 1 s), it is possible to use local ADC board memory in the Transient Recorder configuration, i.e. storing signals in local memory during the pulse and reading the memory content afterwards.

Moreover, thanks to the short plasma discharge duration, it is possible to achieve the desired precision in integration

# FINAL DESIGN OF CONTROL AND DATA ACQUISITION SYSTEM FOR THE ITER HEATING NEUTRAL BEAM INJECTOR TEST BED

L. Trevisan \*, N. Cruz<sup>1</sup>, A. Luchetta<sup>2</sup>, G. Manduchi<sup>2</sup>, G. Martini, A. Rigoni,  
 C. Taliercio<sup>2</sup>, N. Ferron, R. Delogu, Consorzio RFX, Padova, Italy  
 C. Labate, F. Paolucci, Fusion for Energy, Barcelona, Spain  
 J.M. Arias, ITER Organization, St. Paul-lez-Durance, France  
<sup>1</sup>also at Instituto de Plasmas e Fusão Nuclear, Instituto Superior Técnico,  
 Universidade de Lisboa, Portugal  
<sup>2</sup>also at Institute for Plasma Science and Technology , Padova, Italy

## Abstract

Tokamaks use heating neutral beam (HNB) injectors to reach fusion conditions and drive plasma current. ITER, the large international tokamak, will have three high-energy, high-power (1MeV, 16.5MW) HNBs. MITICA, the ITER HNB test bed, is being built at the ITER Neutral Beam Test Facility, Italy, to develop and test the ITER HNB, whose requirements are far beyond the current HNB technology. MITICA operates in a pulsed way with pulse duration up to 3600s and 25% duty cycle. It requires a complex control and data acquisition system (CODAS) to provide supervisory and plant control, monitoring, fast real-time control, data acquisition and archiving, data access, and operator interface. The control infrastructure consists of two parts: central and plant system CODAS. The former provides high-level resources such as servers and a central archive for experimental data. The latter manages the MITICA plant units, i.e., components that generally execute a specific function, such as power supply, vacuum pumping, or scientific parameter measurements. CODAS integrates various technologies to implement the required functions and meet the associated requirements. Our paper presents the CODAS requirements and architecture based on the experience gained with SPIDER, the ITER full-size beam synchronization, fast real-time control, software development for long-lasting experiments, system commissioning and integration.

## INTRODUCTION

MITICA (Megavolt ITER Injector and Concept Advancement) is an experimental device located at the Neutral Beam Test Facilities (NBTF) in Padova, Italy [1]. Its primary purpose is to develop and test high-energy heating neutral beam injectors (HNB) for ITER, the International Thermonuclear Experimental Reactor [2]. Table 1 reports the MITICA requirements.

The MITICA experiment [3] consists of various components, including the beam source to generate the ionized gas, the accelerator to accelerate the ions through an electrostatic voltage gap of up to 1MV, the neutralizer to make the beam neutral, the residual ion dump to deflect the residual ionized particles from the neutral beam, and the calorimeter, which serves as a target for the neutralized beam.

\* luca.trevisan@igi.cnr.it

Table 1: MITICA Key Requirements

	Unit	H	D
Beam Energy	keV	870	1000
Acceleration ion current	A	46	40
Max beam source filling pressure	Pa	0.3	0.3
Max deviation from uniformity	%	±10	±10
Beamlet divergence	mrad	≤ 7	≤ 7
Beam on time	s	3600	3600
Duty cycle on/off		25%	25%
Co-extracted electron fraction		< 5%	< 1%

Various power supply systems provide the necessary electrical power. A series of power supply units known as the Ion Source and Extraction Power Supply (ISEPS) power the beam source. The Acceleration Grids Power Supply (AGPS) and the Ground Related Power Supply (GRPS) energize the acceleration grids and the residual ion dump, respectively.

In MITICA, there are other components: the Cooling System to remove heat from all heated components, the Gas and Vacuum System to create the high-vacuum conditions in the MITICA vessels and provide the gas feed for beam creation and neutralization, and the Cryogenic System to feed the cryogenic pumps with cryogenic coolant. These components are referred to as auxiliary plant systems because they are needed for the HNB operation in the NBTF but are not part of the HNB system itself.

This paper reports the design of the MITICA Control and Data Acquisition System (CODAS) that monitors and controls the other plant systems and acquires and manages the experimental data. After a brief overview of the CODAS architecture, we will present its integration with the other plant systems. Then, we will introduce the strategies adopted to develop CODAS in a long-lasting experiment framework. Finally, we will discuss the advantages of the synchronization network already implemented in MITICA.

## MITICA CODAS ARCHITECTURE

### Functions

MITICA CODAS consists of two main parts: Central and Plant System CODAS. The former provides central functions, such as long-term data storage, supervisory control,

# IMPROVING THE PERFORMANCE OF TARANTA: ANALYSIS OF MEMORY REQUESTS AND IMPLEMENTATION OF THE SOLUTION\*

Matteo Canzari<sup>†</sup>, INAF - Osservatorio Astronomico d'Abruzzo, Teramo, Italy

Hélder Ribeiro, Atlas Innovation, Rua Rangel de Lima, Portugal

Ajaykumar Dubey, PSL, Pune, India

Athos Georgiou, CGI Scotland, Scotland

Valentina Alberti, INAF - Osservatorio Astronomico di Trieste, Trieste, Italy

Yimeng Li, Vincent Hardion, Mikel Eguiraun, Johan Forsberg, Max-IV Institute, Lund, Sweden

## Abstract

Taranta is a software suite for generating graphical interfaces for Tango Controls software, currently adopted by MaxIV for scientific experiment usage, SKA during the current construction phase for the development of engineering interfaces for device debugging, and other institutions. A key feature of Taranta is the ability to create customizable dashboards without writing code, making it easy to create and share views among users by linking the dashboards to their own tango devices. However, due to the simplicity and capabilities of Taranta's widgets, more and more users are creating complex dashboards, which can cause client-side resource problems. Through an analysis of dashboards, we have found that excessive memory requests are generated by a large amount of data. In this article, we report on the process we believe will help us solve this performance issue. Starting with an analysis of the existing architecture, the issues encountered, and performance tests, we identify the causes of these problems. We then study a new architecture exploiting all the potential of the Javascript framework React on which Taranta is built, before moving on to implementation of the solution.

## INTRODUCTION

Taranta [1] is a tool for creating user interfaces for Tango devices without the need for coding [2]. In the Problem Impact section, we will analyze how the use of Taranta for creating complex dashboards has led to performance and system stability issues, resulting in a loss of confidence within the community. We then explored an architectural solution to the problem described in the Problem Analysis section.

In the Implementation section, we describe the implementation of this solution, which was released in the latest version of Taranta [3], and its results are detailed in the Benchmark Test section. Finally, in the conclusions, we highlight the architectural improvements introduced into the system through the implemented changes and discuss potential enhancements.

\* Work supported by the financial support by the Italian Government (MEF - Ministero dell'Economia e delle Finanze, MIUR - Ministero dell'Istruzione, dell'Università e della Ricerca)

<sup>†</sup> matteo.canzari@inaf.it

## PROBLEM IMPACT

The issue reported by Taranta users involved a gradual slowdown of dashboards containing numerous widgets, eventually leading to a complete halt in updates. Additionally, there was a progressive increase in browser RAM usage. This problem was highly inconvenient and, at times, detrimental for several reasons. Primarily, users were unaware that the values from Tango were not being updated at the correct frequency due to the lack of a runtime notification mechanism. Worse still, when the slowdown was severe enough to freeze the dashboard, it provided no feedback to the user. In this state, the user viewing the dashboard assumed that updates from Tango were not coming through when in reality, they had a way to ascertain the current state of the control system. Moreover, the system might not respond to commands, meaning a user could issue a command to change an attribute or state and not observe the result. This situation must be avoided, as the user may need to discern whether the issue stems from malfunctioning devices or an unresponsive user interface. Lastly, there was a concern regarding the escalating browser RAM usage. This was particularly troublesome, given that Taranta is a JavaScript application, hence client-side, and could significantly slow down the user's computer or even cause it to crash.

Addressing this problem promptly was imperative since these often severe issues were eroding user confidence in using Taranta. It is important to note that this type of issue arose only in dashboards with a high number of running widgets over an extended period. For simpler dashboards executed for a limited time, the problem was negligible.

The developers of Taranta were already aware of the performance limitations stemming from an architecture inadequate for supporting a high number of widgets and updates. However, priority was given to developing new features rather than resolving this issue. Taranta users had been using the software for several years without encountering significant problems and appreciated the newly added features. The community's reporting of this issue is a clear indication that users are increasingly adopting the software for more complex scenarios, not just as a simple prototyping and testing tool. This underscores the software's good quality, the utility of the introduced features, and a growing confidence in using Taranta.



# MONITORING THE SKA INFRASTRUCTURE FOR CICD

M. Di Carlo\*, M. Dolci, INAF Osservatorio Astronomico d’Abruzzo, Teramo, Italy  
U. Yilmaz, P. Harding, SKA Observatory, Macclesfield, UK  
P. Osório, Atlar Innovation, Portugal  
J. B. Morgado, CICGE, Faculdade de Ciências da Universidade do Porto, Portugal

## Abstract

The Square Kilometre Array (SKA) is an international effort to build two radio interferometers in South Africa and Australia, forming one Observatory monitored and controlled from global headquarters (GHQ) based in the United Kingdom at Jodrell Bank. The selected solution for monitoring the SKA CICD (continuous integration and continuous deployment) Infrastructure is Prometheus with the help of Thanos. Thanos is used for high availability, resilience, and long term storage retention for monitoring data. For data visualisation, the Grafana project emerged as an important tool for displaying data in order to make specific reasoning and debugging of particular aspect of the infrastructure in place. In this paper, the monitoring platform is presented while considering quality aspect such as performance, scalability, and data preservation. Keywords: prometheus, monitoring, thanos, grafana

## INTRODUCTION

The Square Kilometre Array (SKA) project has selected SAFe (Scaled Agile Framework) as an incremental and iterative development process. A specialized team – named System Team – is devoted to support the Continuous Integration, Continuous Deployment (CI/CD) [1], test automation, and the project components’ quality. Building an infrastructure to support the CI/CD together with a monitoring solution, was one of the first goals of the team.

The SKA infrastructure consists of a standard footprint of VPN/SSH gateway, monitoring, logging, storage and Kubernetes [2] services tailored to support the GitLab [3] runner architecture that is shown in Fig. 1. Furthermore, this infrastructure is used to support Development and Integration testing facilities for the project’s many subsystems. The selected logging solution is Elasticsearch [4], storage is handled through Ceph [5], while Prometheus [6] handles monitoring (see the Prometheus section below), and the (central) artefact repository (CAR) is Nexus [7]. It is important to realize that only artefacts produced by GitLab pipelines that have been marked for a release (i.e. triggered by a git tag), are allowed to be stored on the CAR. On all other cases, GitLab’s own artefact repository is used.

The infrastructure shown in Fig. 1 is replicated in multiple locations spread over different continents, with specific hardware on top of OpenStack [8] or bare metal/virtual machine instances. Meaning, that for every infrastructure, the same components will be present – at a different scale depending on the available resources. Some components, such as

Elasticsearch and Thanos, exist on a single location, centralizing the access to information of several data centres to allow aggregated analysis across different datacentres while reducing the maintenance overhead.

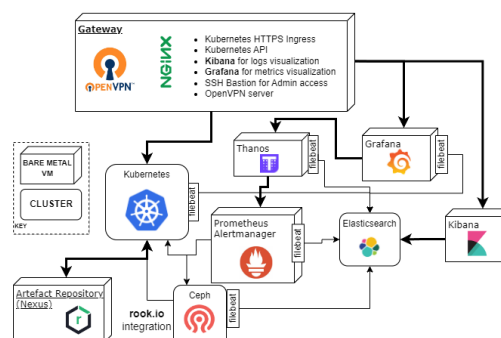


Figure 1: Simplified infrastructure.

## PROMETHEUS

The selected monitoring solution, for every infrastructure, is Prometheus. It works in client-server architecture, where Prometheus acts as a client that reads (\*scrapes\* in its domain-specific language) timestamped information from multiple servers – called targets or exporters. The data is stored on a disk in TSDB format [9] and pushed to an object store. Figure 2 shows a detailed diagram that illustrates the monitoring architecture on SKA. The main components of the diagram are:

- **Prometheus server**, which is composed by the scraper, the TSDB storage, an HTTP API, and Web interface for data querying;
- **Thanos**, which is composed by many different components to help with high availability, data retention, retrieval and long term storage;
- **Jobs/exporters**, where Prometheus scrapes information as a time series;
- **Grafana**, as a data visualization and export tool that integrates with Prometheus/Thanos using PromQL – a specific query language;
- **Alertmanager**, that delivers alarm notifications to end-user systems;

Each exporter provides a time series uniquely identified by its metric name and some optional key-value pairs – called labels. It is important to note that the exporter must give an

\* matteo.dicarlo@inaf.it

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# BEAM OPERATION FOR PARTICLE PHYSICS AND PHOTON SCIENCE WITH PULSE-TO-PULSE MODULATION AT KEK INJECTOR LINAC

K. Furukawa\*, M. Satoh, Injector LINAC group,  
 KEK, Tsukuba, Japan, SOKENDAI, Tsukuba, Japan

## Abstract

The electron and positron accelerator complex at KEK offers unique experimental opportunities in the fields of elementary particle physics with SuperKEKB collider and photon science with two light sources. In order to maximize the experimental performances at those facilities the injector LINAC employs pulse-to-pulse modulation at 50 Hz, injecting beams with diverse properties. The event-based control system effectively manages different beam configurations. This injection scheme was initially designed 15 years ago and has been in full operation since 2019. Over the years, quite a few enhancements have been implemented. As the event-based controls are tightly coupled with microwave systems, machine protection systems and so on, their modifications require meticulous planning. However, the diverse requirements from particle physics and photon science, stemming from the distinct nature of those experiments, often necessitate patient negotiation to meet the demands of both fields. This presentation discusses those operational aspects of the multidisciplinary facility.

## INTRODUCTION

Particle accelerators require large resources with advanced technologies and experienced personnel in order to construct and to operate them. Therefore, beams from certain accelerators have been shared between several different purposes. A long beam switching interval may be accepted between some accelerators. On the other hand, in order to switch beams frequently the mechanism called pulse-to-pulse beam modulation (PPM) have been developed and employed for advanced accelerators [1,2]. However, different beam users may request variety of beam properties as well as conflicting operational concepts. Thus, meticulous arbitration for multidisciplinary operation would be necessary.

The injector LINAC at KEK has been operated for more than 40 years and has served for several storage ring accelerators in the past [3]. Presently, it makes multidisciplinary injection operation to support two light sources and a B factory collider storage rings with pulse-to-pulse beam modulation as in Fig. 1. It is often referred to as simultaneous top-up injections [4].

## INJECTOR LINAC AT KEK

In 1978 a dedicated (the second generation) light source project was approved, and then a high energy electron positron collider project was approved as well. A 400-m

electron linear accelerator was constructed to support both synchrotron radiation research and high energy physics collider experiment [5]. Since then, it has been serving for various projects which required diverse beam properties as in Fig. 2.

Additionally, each project required an upgrade of the injector LINAC which might take up to six years in between projects, however, even during the upgrade the other project needed the beam injection. Thus, the equipment tests, fabrications and installations had to be carefully arranged in order to support a sustainable injector operation.

For example, it took 5 years to upgrade the accelerators to realize a KEKB B-factory machine [6]. While the collider ring was shutdown for 5 years, the injector LINAC had to deliver the beam every year to the light source. The injector performed 3-month injection operation and 3-month upgrade installation repeatedly. The SuperKEKB upgrade required 6 years from 2010 and another year in 2017 [7]. The injector utilized the last downstream part of LINAC to inject the beam for the light sources. And the rest of the LINAC was reconstructed to support high intensity and low emittance injection for SuperKEKB as well as to recover from the Great East Japan Earthquake in 2011. The longest complete shutdown of the injector LINAC was 5 months in 2017. During the upgrade period a particular effort was made for the alignment in order to suppress the wakefield effect in the accelerating structure.

In such way, the injector LINAC has negotiated with both communities of photon science and particle physics, and they understood each other to maximize the overall experimental performance.




Figure 1: Beam injection from electron positron injector LINAC to SuperKEKB dual ring collider as well as PF and PF-AR light source storage rings.

\* kazuro.furukawa@kek.jp

# DEVELOPMENT OF OPERATOR INTERFACE USING ANGULAR AT THE KEK e-/e+ INJECTOR LINAC

M. Satoh <sup>†,1</sup>, I. Satake, High Energy Accelerator Organization (KEK),  
Accelerator Laboratory, Tsukuba, Japan

T. Kudou, S. Kusano, Mitsubishi Electric System & Service Co., Ltd (MSC), Tsukuba, Japan

<sup>1</sup>also at The Graduate University for Advanced Studies (SOKENDAI),  
Department of Accelerator Science, Tsukuba, Japan

## Abstract

At the KEK e-/e+ injector LINAC, the first electronic operation logbook system was developed in 1995 using a relational database. This logbook system can automatically record the detailed operational status. In addition, operators can manually input detailed information about operational problems, which is highly helpful for future troubleshooting. In 2010, the logbook system was improved with the implementation of a redundant database, an Adobe Flash-based frontend, and an image file handling feature. In 2011, we started using the CSS archiver system using PostgreSQL as a backend database and a new web-based archiver viewer developed by Adobe Flash. However, with the discontinuation of Adobe Flash support at the end of 2020, it became necessary to develop a new frontend without Adobe Flash for both the operation logbook and archiver viewer systems. For this purpose, the authors adopted the Angular framework, which is widely used for the development of web applications using JavaScript. In this paper, we report the development of operator interfaces using Angular for the injector LINAC.

## INTRODUCTION

The KEK e-/e+ injector LINAC [1, 2] established the simultaneous top-up injection to five independent rings in 2019 to support both the SuperKEKB particle collider experiment with DR, LER and HER rings [3], and the photon science experiment at the PF ring and PF-AR, as shown in Fig. 1. It succeeded in improving the efficiency of the SuperKEKB collision experiment more than 200% before and after the introduction of the simultaneous top-up injection [4]. This noble injection scheme became indispensable because the beam lifetime of the SuperKEKB ring is short, especially at the LER positron ring, which was less than 10 min in 2021. On the basis of this operation arrangement, the injector LINAC has gradually improves the injection performance and contributed to the achievement of the world-record collision luminosity of SuperKEKB [5].

At the injector LINAC, the beam operation status and its history have been automatically recorded with an electronic operation logbook system. In 1995, the first operation logbook system was developed with Microsoft Access (MS-ACCESS) and Microsoft SQL (MS-SQL) on the Windows 2000 server because of the easy database management and development of the user interface. MS-ACCESS is a frontend component, and it decreases the CPU load of

MS-SQL. The console PC and MS-SQL are connected through Open Database Connectivity (ODBC), and an operator can enter the current operation status of LINAC in the Japanese language. After 1999, the beam injection into the KEKB ring started and the data size markedly increased. After that, the electrical operation logbook system has been gradually improved. Finally, the backend database was replaced with PostgreSQL, and the frontend component, which operates within web browsers, was developed using Angular [6], a well-known web application framework, in addition to React. Additionally, the viewer web application for the data archiver system was also developed and has been operated using Angular.




Figure 1: Layout of the KEK e-/e+ accelerator complex with beam properties from the injector LINAC to five independent rings.

## PREVIOUS OPERATOR INTERFACE

In injector operations, following the transformation of the electronic operation logbook system into a web application in 2010 [7], Adobe Flash was extensively employed in the frontend of web applications. Flash applications offered enhanced usability and expressive capabilities compared with simple HTML-based pages. They could run on computers with Flash Player, regardless of the operating system, and being web applications, they facilitated easy redistribution, significantly reducing operational complexity compared with native applications that required installation. These applications had been stably operated since their development. However, with the discontinuation of Adobe Flash Player support by major web browsers in 2020, a transition to a new framework became imperative.

At the injector LINAC, Angular emerged as a promising framework to replace Adobe Flash. Angular is a web application framework developed by a community led by

<sup>†</sup> masanori.satoh@kek.jp

# THE UPGRADE OF PULSED MAGNET CONTROL SYSTEM USING PXIe DEVICES AT KEK LINAC

D. Wang\*, M. Satoh  
KEK, Ibaraki 305-0801, Japan

## Abstract

The pulsed magnet control system (PMCS) at KEK electron positron injector LINAC operates at every 20 ms to achieve simultaneous injection for four rings, a 2.5 GeV Photon Factory (PF), a 6.5 GeV PF-AR, a 4 GeV SuperKEKB lower energy ring (LER) and a 7 GeV SuperKEKB high energy ring (HER). The system consists of a control server that operates on the Windows 8.1 platform, in conjunction with a PXIe chassis equipped with a DAC, an ADC, and an event timing module. The PXIe DAC board responds to the trigger signal which is generated from the event receiver and sets the current of the pulsed magnet. The readback value of the magnet current is retrieved from the ADC module. Despite its long success, the PMCS has several problems in use. One is the discontinued support of Windows 8.1. Another key concern is the unsatisfactory long-term stability. To solve the problems, an upgraded system using real-time Linux to communicate with PXIe modules is adopted. The EPICS driver for PXIe devices is developed to integrate with the LINAC control system. The development of the new Linux-based PMCS is introduced in this work.

## INTRODUCTION

The injector LINAC at KEK is responsible for the injection of 4 target rings which consist of a 7 GeV electron high energy ring (HER), a 4 GeV positron low energy ring (LER), a 2.5 GeV Photon Factory (PF) and a 6.5 GeV PF-AR ring, as shown in Fig. 1. LINAC is responsible for performing a simultaneous top-up injections into four target storage rings and a DR using the method called pulse-to-pulse modulation (PPM) [1].




Figure 1: A schematic view of LINAC, SuperKEKB, and PF/PF-AR.

In order to meet the PPM requirements, a total of 16 pulsed magnet control units have been deployed across the 600-meter LINAC since 2017. Upon receiving a specific event

code, indicative of a particular beam mode, these control units are activated to adjust the magnet current. By using this configuration, the magnetic field can undergo pulse-to-pulse modifications within 20 ms. This rapid response ensures that the beam profile is optimized to the requirements of each destination ring.

## PULSED MAGNET CONTROL SYSTEM

### Hardware

Figure 2 shows the rack of one pulsed magnet control unit. Each unit consists of a homemade server and a National Instruments (NI) PXIe-1082 chassis fitted with four modules, a controller control module (NI PXIe-8381), an event receiver (EVR) board (MRF PXI-EVR-230), a DAC board (NI PXI-6733), and an ADC board (NI PXIe-6356) [2, 3].




Figure 2: Rack of a pulsed magnet control unit.

The established units facilitates the independent control and monitoring of the output currents for as many as 8 power supplies. It offers a resolution of 16 bits and operates with a sampling rate of 1 MSA/s. To counteract potential instabilities that may arise from power disruptions or external

\* di.wang@kek.jp

# 15 YEARS OF THE J-PARC MAIN RING CONTROL SYSTEM OPERATION AND ITS FUTURE PLAN

Shuei Yamada\*, KEK/J-PARC, Ibaraki, Japan

## Abstract

The accelerator control system of the J-PARC MR started operation in 2008. Most of the components of the control computers, such as servers, disks, operation terminals, front-end computers and software, which were introduced during the construction phase, have went through one or two generational changes in the last 15 years. Alongside, the policies for the operation of control computers has changed. This paper reviews the renewal of those components and discusses the philosophy behind the configuration and operational policy. It is also discussed the approach to matters that did not exist at the beginning of the project, such as virtualization or Haber security.

## INTRODUCTION

J-PARC (Japan Proton Accelerator Research Complex) is a high-intensity proton accelerator facility jointly planned, developed and operated by KEK (High Energy Accelerator Research Organisation) and JAEA (Japan Atomic Energy Agency). Construction began in 2001 and the operation started in 2007 [1]. The control system of the J-PARC accelerator consists of separate accelerator control systems for two different machine cycles. The accelerator control system for Linac (LI) and Rapid Cycle Synchrotron (RCS), whose repetition is 25 Hz, is managed by JAEA. While the control system for the Main Ring (MR) with 1.36 s or 4.24 s cycles, depending on the operation mode, is managed by KEK. These two control systems are built on a common infrastructure: timing system, network, storage system and EPICS [2], and work closely together to control the entire accelerator as a whole.

This paper discusses the 15-year operational history and future prospects of MR control system. The accelerator control system in the broadest sense also includes timing systems [3,4], personnel protection system [5], and machine protection system [6,7], but these are only mentioned in the references. The components of a networked distributed control system using EPICS, such as network, storage system, computing and software environment, will be focused on.

## ACCELERATOR CONTROL NETWORK

### Structure of the Network

The logical configuration of the J-PARC accelerator control LAN is shown in Fig. 1. The core switch is located in the Central Control Building (CCR) and is wired by optical fibre to the edge switches in the LI, RCS, Materials and Life Science Experimental Facility (MLF), and MR 3rd Power Supply Building (D3). Optical fibres are further wired from

D3 to the edge switches at D1, D2, Neutrino Experimental Facility (NU), and Hadron Experimental Facility (HD).




Figure 1: Logical Topology of the J-PARC Accelerator Network.

The core and edge switches are redundant and fails over from primary to the secondary system upon the failure. On-site maintenance services will promptly quickly identify the cause of the failure, take countermeasures and replace faulty components. VLAN divides the network into seven segments, namely, LI, RCS, MR, MLF, HD, NU, and CCR, minimizing the impact of network failures in one facility on other facilities.

The entire accelerator control network uses approximately 250 edge switches. LI and RCS are configured with edge switches right down to the end switch. MR uses a total of 12 edge switches for intermediate switches and SOHO switching hubs at the ends to reduce deployment and maintenance costs.

### History of the Network

Equipment for the J-PARC accelerator control network has been supplied by Extreme, and all core and edge switches are replaced every 7–8 years. The control network started operation at LI in 2005 and was extended to MR in 2007. The bandwidth was 10 Gbps on the backbone and 1 Gbps between buildings. Network equipment was upgraded between 2011 and 2015, with bandwidth of 40 Gbps backbone and 10 Gbps between buildings. A second renewal of equipment has been underway since 2019. It will be completed in 2024, with a 100 Gbps backbone and 10 Gbps between buildings.

### Network Security

The relationship between the J-PARC accelerator control network, the office network and the Internet is shown in Fig. 2.

\* shuei@post.kek.jp

# DEVELOPMENT AND TEST OPERATION OF THE PROTOTYPE OF THE NEW BEAM INTERLOCK SYSTEM FOR MACHINE PROTECTION OF THE RIKEN RI BEAM FACTORY

M. Komiyama<sup>†</sup>, M. Fujimaki, A. Kamoshida<sup>1</sup>, A. Uchiyama, N. Fukunishi,  
RIKEN Nishina Center for Accelerator-Based Science, Wako, Japan  
M. Hamanaka, M. Nishimura, R. Koyama, K. Kaneko, H. Yamauchi,  
SHI Accelerator Service Ltd., Tokyo, Japan  
<sup>1</sup>also at National Instruments Japan Corporation, Tokyo, Japan

## Abstract

Since 2006, we have been operating a beam interlock system (BIS) for machine protection at the RIKEN RI Beam Factory (RIBF). The system has a reaction rate of approximately 15 ms after receiving an alert signal from the accelerator and beamline components. Considering that the beam intensity of the RIBF will increase, we are currently developing a successor system to stop the beam within 1 ms. After comparing multiple systems, CompactRIO, a National Instruments product, was selected as the successor system. Interlock logic for signal input/output is implemented on a field-programmable gate array (FPGA) to ensure fast processing speed. However, signal condition setting and monitoring do not require the same speed as the interlock logic. Therefore, they are implemented on a real-time operating system (RT-OS) and controlled using the Experimental Physics and Industrial Control System (EPICS) by setting up an EPICS server on the RT-OS. Furthermore, in the successor system, a system for fast alert signals, such as digital signals, immediately outputs from the equipment. In addition, a system for slow alert signals, such as analog values of the beam current sampled at a certain period, is developed as a linked system. Development of the successor system began in 2021. As of the summer of 2023, a prototype consisting of three stations for processing digital input signals and one station for processing analog input signals was installed in parallel with part of the BIS. The time required for the prototype to output a signal required to stop a beam after receiving an alert signal averaged 130  $\mu$ s for a digital input signal and 470  $\mu$ s for an analog input signal. These results show that the prototype can achieve much better performance than the target system response time, and the target performance can be expected to be achieved when the prototype is extended to the entire system in the future.

## INTRODUCTION

The RIKEN Radioactive Isotope Beam Factory (RIBF) at the RIKEN Nishina Center for Accelerator-Based Science is a complex accelerator facility comprising multiple linear accelerators, including a superconducting linac, multiple cyclotrons, and the world's largest superconducting cyclotron. Furthermore, RIBF can provide the world's most intense RI beams over the entire atomic mass range

using the fragmentation or fission of high-energy heavy ions [1].

The components of the RIBF accelerator complex (such as the magnet power supplies, beam diagnostic devices, and vacuum systems) are controlled by Experimental Physics and Industrial Control System (EPICS) [2]. However, exceptions exist, such as the control system dedicated to the RIBF's radio frequency system [3]. All essential operation datasets of EPICS and other control systems were integrated into an EPICS-based control system [4]. Two independent safety systems operate in RIBF facilities: a radiation safety interlock system for human protection [5] and a machine protection interlock system [6].

## OVERVIEW OF THE THREE TYPES OF BEAM INTERLOCK SYSTEMS IN OPERATION

RIBF currently has three machine protection interlock systems: the beam interlock system (BIS), AVF-BIS, and superconducting RIKEN Linear Accelerator (SRILAC)-BIS. First, the BIS is the largest of the three abovementioned beam interlock types and began operation in 2006. It was developed based on Melsec-Q series programmable logic controllers (PLCs), a product of the Mitsubishi Electric Corporation [7]. Figure 1 shows the hardware configuration and process flow of the BIS. Furthermore, it was designed to stop the beams within 10 ms of receiving an alert signal from the accelerator and beamline components. Upon receiving an alert signal, the BIS outputs a signal to one of the beam choppers, which immediately deflects the beam downstream of the ion source (the time required for this is called the system response time). It also inserts a beam stopper (Faraday cup) installed upstream of the problematic component. The BIS ignores the problems that occur downstream of the beam stopper insertion point. After inserting the relevant beam stopper, the beam chopper can be switched off, and beam delivery can resume up to the inserted beam stopper. This feature is particularly important because the problem recovery time can be effectively used to readjust the beam to the inserted beam stopper when the problematic component cannot be recovered within a short time. Moreover, the inserted beam stopper can be extracted from the beamline after the problem is resolved.

The BIS consists of five PLC stations with 384 digital inputs (DI), 80 digital outputs (DO), and 112 analog inputs

<sup>†</sup> misaki@riken.jp

# THE PROGRESS AND STATUS OF HEPS BEAMLINE CONTROL SYSTEM\*

G. Li<sup>†</sup>, AI. Y. Zhou, Y. Liu, Z. Y. Yue, X. W. Dong, CH. X. Yin, X. B. Deng, Q. Zhang, D. SH. Zhang, ZH. H. Gao, ZH. Zhao, N. Xie, G. Lei<sup>†</sup>, Institute of High Energy Physics, CAS, Beijing, China

## Abstract

HEPS (High Energy Photon Source) will be the first high-energy (6 GeV) synchrotron radiation light source in China, which is mainly composed of accelerator, beamlines and end-stations. In the phase I, “14+1” beamlines and corresponding experimental stations will be constructed. The beamline control system has completed the design of the control system scheme based on the EPICS framework. And it will soon enter the stage of engineering construction and united commissioning. Here, the progress and status of beamline control system are presented.

## INTRODUCTION

The High Energy Photon Source (HEPS) [1], which is a high energy kilometre-scale ring-based light source with a double-bend achromat (DBA) design, is the first high energy diffraction-limited storage ring (DLSR) light source under construction in China (Fig. 1).



Figure 1: A bird's eye view of HEPS.

To satisfy some major demands of the nation and to further cutting edge scientific development, “14+1” beamlines and corresponding experimental stations will be constructed [2] (Fig. 2).

In order to simultaneously complete the construction of 15 beamlines control system on schedule, a set of standards is developed to design and build control systems for beamlines, minimizing heterogeneousness, to ensure that the control system has good stability, reliability, flexibility, availability, reliability etc.

Beamline control system of HEPS is designed and developed, based on the EPICS framework. This paper presents the progress and status of HEPS beamline control system.




Figure 2: Beamlines layout in HEPS Phase I.

## CURRENT STATUS OF BEAMLINE CONTROL SYSTEMS

Beamline control system of HEPS is designed and developed, based on the EPICS [3] framework (Fig. 3).




Figure 3: Overall structure of beamline control system.

This paper presents the progress and status of HEPS beamline control system.

### Automated Deployment of the Software

Each beamline has many front-end controllers that need to deploy OS (operating system) and EPICS applications, so it is necessary to design and build a software deployment system to complete the above tasks.

In the process of HEPS beamline control system construction, there are multiple controllers (terminal PC, IPC, workstation and server, etc.) need to configure the OS, which only rely on access to the traditional installation media on such as CD/DVD-ROM or USB driver, etc., the workload is large and inefficient, and the distribution of the OS is not easy to manage, therefore, a set of parallelized batch deployment of OS schemes is required.

PXE (Pre-boot Execution Environment) is technology created by Intel Corporation, which allows automated provisioning OS of servers or workstations over a network [4]. It works in Client/Server mode, where the client controller downloads the OS image from the server over the network to implement the deployment of the client's OS (Fig. 4).

\* Work supported by HEPS project  
<sup>†</sup>li75gang@ihep.ac.cn; leige@ihep.ac.cn

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# INTRODUCTION TO THE CONTROL SYSTEM OF THE PAL-XFEL BEAMLINES

Gisu Park<sup>†</sup>, Sunmin Hwang, Chaeyong Lim, Minzy Jeong, Wonup Kang  
 Pohang Accelerator Laboratory, Pohang University of Science and Technology,  
 Pohang, Kyungbuk Republic of Korea

## Abstract

The PAL-XFEL beamlines are composed of two different types of beamlines: a hard X-ray beamline and a soft X-ray beamline. The hard X-ray beamline generates free electron lasers with pulse energies ranging from 2-15 keV, pulse lengths of 10-35 fs, and arrival time errors of less than 20 fs from 4-11 GeV electron beams for X-ray Scattering & Spectroscopy (XSS) and Nano Crystallography & Coherent Imaging (NCI) experiments. On the other hand, the soft X-ray beamline generates free electron lasers with photon energies ranging from 0.25-1.25 keV, and with more than  $10^{12}$  photons, along with 3 GeV electron beams for soft X-ray Scattering & Spectroscopy (SSS) experiments. To conduct experiments using the XFEL, precise beam alignment, diagnostics, and control of experimental devices are necessary. The devices of the three beamlines are composed of control systems based on the Experimental Physics and Industrial Control System (EPICS), which is a widely-used open-source software framework for distributed control systems. The beam diagnostic devices include QBPM (Quad Beam Position Monitor), photodiode, Pop-in monitor, and inline spectrometer, among others. Additionally, there are other systems such as CRL (Compound Refractive Lenses), KB mirror (Kirkpatrick-Baez mirror), attenuator, and vacuum that are used in the PAL-XFEL beamlines. We would like to introduce the control system, event timing, and network configuration for PAL-XFEL experiments.

## INTRODUCTION

Introduces the configuration of the PAL-XFEL beamlines and the experimental techniques used in each experiment hutch, introduces the event timing system for devices requiring synchronization, introduces network configuration for device control and data transmission, and introduces EPICS IOC for controlling experimental devices. finally, we will introduce the data monitoring system for EPICS IOC management.

### PAL-XFEL Beamlines

The PAL-XFEL beamlines is composed of HX (Hard X-ray) and SX (Soft X-ray) FEL beamlines. The HX beamline consists of a 780-meter-long accelerator line, a 250 meter-long undulator line, and 80-meter-long experimental halls. The SX beamline branches off at the 260-meter point from the beginning and includes a 170-meter-long accelerator line, a 130-meter-long undulator line, and a 30-meter-long experimental hall [1].

Figure 1 below is a schematic diagram of the PAL-XFEL beamlines. HX has two experiment hutches and supports X-ray scattering and spectroscopy (XSS) and nano-crystallography and coherent imaging (NCI) experiments. XSS supports femtosecond X-ray scattering (FXS)/femtosecond X-ray liquidography (FXL), and NCI supports coherent X-ray imaging/scattering/spectroscopy (CXI)/serial femtosecond crystallography (SFX) experiment techniques. SX has two endstations and supports soft X-ray scattering and spectroscopy (SSS) experiments. X-ray emission/absorption spectroscopy (XES/XAS) and resonant soft X-ray scattering (RSXS) experiment techniques are supported.




Figure 1: Schematic diagram of PAL-XFEL beamlines.

Table 1 below shows the beam parameter values of HX and SX, the photon energy of hard X-ray is 2 ~ 15 keV, and the repetition rate supports up to 60Hz. Band width of pink beam is ~0.4% and Photon flux (pink beam) is  $>10^{11}$  phs/pulse. the photon energy of Soft X-ray is 250 ~ 1250 eV, and the repetition rate supports up to 60 Hz. Band width of pink beam is ~ 0.5% and Photon flux (pink beam) is  $>10^{12}$  phs/pulse @ 800 eV [2-3].

Table 1: Conditions of Hard X-ray and Soft X-ray

	Hard X-ray	Soft X-ray
Photon energy	2.0 ~ 15 keV (0.6 ~0.8 nm)	250 ~ 1250 eV (5~1 nm)
Repetition rate	10 Hz, 30 Hz, 60 Hz	10 Hz, 30 Hz, 60 Hz
Band width of pink beam ( $\Delta E/E$ )	~ 0.4 %	~ 0.5 %
Photon flux (pink beam)	$>1.0 \times 10^{11}$ phs/pulse	$>1.0 \times 10^{12}$ phs/pulse @800 eV

### Event Timing System

The PAL-XFEL uses an event timing system developed by Micro-Research Finland (MRF). Essentially, this system is VME-based and comprises components such as the MVME6100 CPU, VME-EVG-230, VME-EVR-230, all running on the RTEMS operating system. However, owing to the recent discontinuation of the VME-EVR card, we are in the process of transitioning to a MircoTCA (MTCA)-based system.

<sup>†</sup> gspark86@poatech.ac.kr



# IMPLEMENTATION OF AN EXTERNAL DELAY CALCULATOR FOR MeerKAT

B. Ngcebetsha\*, South African Radio Astronomy Observatory, Cape Town, South Africa

## Abstract

The MeerKAT is an array of dishes designed to study the mysteries of the distant universe, it is made up of 64 dishes that operate as one telescope. The signal from the distant celestial objects encounters several distortions and is corrupted when it arrives at the receiver. Most of the distortions are due to the medium between the observatory and the object of interest. Each corruption can be quantified and corrected, the corruptions are termed "propagation effects". The process of correcting the propagation effects constitutes calibration and takes on various stages during and after the observation is done. The context of this paper takes a close look at signal path delay correction. This is the adjustment of the time of arrival of the signal at the correlator from all 64 antennas. This is required as the signal arrives at different times at every antenna, and the cable from each antenna is of differing length. The MeerKAT CAM system implements a delay update manager. The delay update manager calculates the delay terms, and submits the corrections to the correlator. In this paper, we describe how this solution was evolved when katpoint (the underlying library on which the delay corrections depend) had a change in its own dependencies. There were two major changes to katpoint 1) utilising astropy instead of ephem, and this meant 2) migrating telescope code from version 2 to 3. In this paper we explore the lessons learned when katpoint started to implement astropy which is implemented in Python3 whilst the rest of the code-base was still in Python2. The technical benefit of this update was an improvement in the astrometry for delay calculations which will enhance the MeerKAT science and images.

## INTRODUCTION

One of the objectives of radio observations by telescopes such as the MeerKAT is to make a map of the patch of sky observed and create a catalogue with a list of astronomical objects and their physical properties. The telescope is composed of pair-wise groupings of the antennas known as baselines. An illustration of a baseline is Fig. 1, where a signal arriving at both antennas is recorded ideally at the same time for further processing at the correlator. Prior to arriving at the correlator, the signal is a voltage reading, and is converted to a digital signal by the analog-to-digital converter (ADC). The main task of the correlator (dashed-box), is to multiply and time-average (correlation). An extra step in collecting the data is the tweaking of the time of arrival. This is achieved by the correlator introducing a signal

\* bngcebetsha@sarao.ac.za




Figure 1: An two-antenna interferometer, antenna1(left) detects the signal  $\tau_g$  seconds later than antenna2(right). Where  $\tau_g$  is called the geometric delay.

path delay. The delay is usually in the range of nanoseconds (billionth of a second), defined as the time it takes light to travel the distance  $\vec{B} \cdot \vec{s}$ . This is the distance between the tip of antenna2 on the right and that of antenna1 on the left in Fig. 1.

The antennas sample the signal in an arbitrary plane known as the uvw plane, and this is the coordinate system in which the baselines have been positioned. There exists, theoretically - a Fourier relationship between the uvw plane and the actual image plane, see Fig. 2. The output of the correlator is a table of complex numbers known as visibility data. The data represents snapshots at every feasible uvw coordinate at the location of every baseline for the duration of the observation. In order to make the data scientifically useful, astronomers make an image of the sky through the use of deconvolution algorithms built into imaging pipelines. This process produces better results with improved and accurate measurement of sky positions. The sky position measurements also rely on accurate timing of the arrival of the signal, which has travelled vast distances from outer space. The incident signal is electromagnetic in nature, the wave nature of light implies the phase at which it arrives also needs to be taken into consideration and corrected. MeerKAT generally uses an in-house Python library known as katpoint for all target-related motions in the control software. Earlier

# AVN RADIO TELESCOPE CONVERSION SOFTWARE SYSTEMS

R.L. Schwartz\*, P.J. Pretorius†, R.M. Ebrahim§,  
South African Radio Astronomy Observatory (SARAO),  
Cape Town, South Africa

## Abstract

The African VLBI Network (AVN) is a proposed network of Radio Telescopes involving 8 partner countries across the African continent. The AVN project aims to convert redundant satellite data communications ground stations, where viable, to Radio Telescopes. One of the main objectives of AVN is human capital development in Science, Engineering, Technology and Mathematics (STEM) with regards to radio astronomy in SKA (Square Kilometer Array) African Partner countries.

This paper will outline the software systems used for control and monitoring of a single radio telescope. The control and monitoring software consists of the User Interface, Antenna Control System, Receiver Control System and monitoring of all proprietary and Off-The-Shelf (OTS) components. All proprietary and OTS interfaces are converted to the open protocol (KATCP).

## INTRODUCTION

The African Very Long Baseline Interferometry Network (AVN) is a pan-African project that will develop Very Long Baseline Interferometry (VLBI) observing capability in several countries across the African continent, either by conversion of existing telecommunications antennas into radio telescopes or by building new ones.

This paper focuses on the conversion of the Kuntunse satellite communication station (near Accra, Ghana), specifically the software systems developed by the software groups of SARAO.

## BACKGROUND

### *African VLBI Network*

The AVN is an ambitious project that will establish a network of radio telescopes across the African continent to support the existing international VLBI network. Stations located across Africa would greatly improve the image fidelity of VLBI observations, since the geographic location of these stations improves sensitivity to angular scales on the sky that are not well sampled using the currently available global configuration [1].

### *Objective of the AVN*

The AVN will help to develop the skills, regulations and institutional capacity required in SKA partner countries to optimise African participation in the SKA and enable participation in SKA pathfinder technology development and

science. The SKA AVN partners of South Africa are: Botswana, Ghana, Kenya, Madagascar, Mauritius, Mozambique, Namibia, and Zambia.

The AVN programme will transfer skills and knowledge to African partner countries to build, maintain, operate and use radio telescopes.

It will bring new science opportunities to Africa on a relatively short time scale and develop radio astronomy science communities in the SKA partner countries.

### *VLBI*

VLBI is astronomical interferometry used in radio astronomy, where the same astronomical source signal is recorded by multiple radio telescopes on Earth or in space. The data is timestamped with very accurate time sources (e.g., GPS and hydrogen maser). These signals are then correlated at a later stage to produce the resulting image, using the timestamps. This emulates a telescope with a diameter equal to the maximum distance between the telescopes and enables distances between antennas to be much greater than possible with conventional interferometry.

## STATION CONTROL AND MONITORING SOFTWARE OVERVIEW

### *The Software Group*

The AVN Software Group currently consists of three members who are responsible for the main Control and Monitoring software systems design, implementation and commissioning, as well as the interconnection and networking of various subsystems on the telescope.

### *KATCP*

KATCP (KAT Communications Protocol) is a simple ASCII communications protocol layered on top of TCP/IP. It was developed as part of the Karoo Array Telescope (KAT), and is used for the monitoring and control of hardware devices.

The protocol consists of newline separated text messages sent asynchronously over a TCP/IP link. There are three categories of messages: requests, replies and informs. The AVN control and monitoring system uses KATCP as the primary station communications protocol.

### *Station CAM Software*

AVN Station Control and Monitoring (CAM) Software consists of the following software subsystems:

- Station Controller software
- Protocol Translation software
- Antenna Steering Control System software
- Environmental Monitoring System software
- VLBI Backend software

\*rschwartz@sarao.ac.za

†ppretorius@sarao.ac.za

§rebrahim@sarao.ac.za

# OPEN TIME PROPOSAL SUBMISSION SYSTEM FOR THE MeerKAT RADIO TELESCOPE

R.L. Schwartz\*, T. Baloyi†, S.S. Sithole§  
South African Radio Astronomy Observatory (SARAO),  
Cape Town, South Africa

## Abstract

Through periodic Calls for Proposals, the South African Radio Astronomy Observatory (SARAO) allocates time on the MeerKAT Radio Telescope to the international community to maximise the scientific impact through radio astronomy; while contributing to South African scientific leadership and human capital development.

Proposals are submitted through the proposal submission system, followed by a stringent review process where they are graded based on specific criteria. Time on the telescope is then allocated based on the grade and rank achieved.

This paper outlines the details of the Open Time proposal submission and review process and the design and implementation of the software used to grade the proposals and allocate the time on the MeerKAT Radio Telescope.

## INTRODUCTION

The MeerKAT Radio Telescope is a project of the South African Radio Astronomy Observatory (SARAO) [1] under the National Research Foundation (NRF) [2]. Inaugurated in 2017, it comprises interlinked receptors in the Meerkat National Park, located in the Northern Cape, South Africa. It was initially named the Karoo Array Telescope (KAT) and would have only consisted of 20 receptors, however when the South African government increased the budget, a total of 64 receptors were commissioned. For this reason, it was renamed “MeerKAT”, using the Afrikaans word “meer”, meaning more.

The MeerKAT Radio telescope was born as a precursor to the intergovernmental Square Kilometer Array (SKA) telescope [3] that will combine thousands of receptors and observatory infrastructure in South Africa and Australia. The SKA will be the largest and most powerful radio telescope in the Southern Hemisphere spanning an effective area of 1 million square meters.

As with any other observatory, SARAO offers telescope time through Open Time Calls (OTC). An OTC is the workflow that SARAO utilises in gathering radio observation proposals from astronomers all over the world, which is then followed by a rigorous review process that ultimately produces proposal observations that the MeerKAT telescope will observe for a specific period.

Following the inauguration of the MeerKAT radio telescope in 2017, the first Open Time Call (OTC1) for proposals was launched in 2018/2019 to a select community.

OTC1 was a manual process that employed document collaboration tools for numerous tasks, such as recording user information, observation parameters, and tracking the proposal statuses. Consequently, bulk correspondence with persons who submitted proposals was inherently laborious.

Despite the success of OTC1, the complexity and volume of information necessitated a change for the next OTC. Thus, the arduous workflow of OTC1 motivated the design and development of an information system that would cater for subsequent OTCs, and so the Proposal Workflow System (PAWS) was conceived.

This paper aims to describe and explain, on a high level, the inner workings of the PAWS system that is used at SARAO to acquire potential radio observation proposals for the MeerKAT Radio Telescope.

## PROPOSAL WORKFLOW SYSTEM

PAWS is a front-facing web application that was developed to aid the process of acquiring observation proposals for the second Open Time Call (OTC2) scheduled for 2020.

The OTC process follows ordered events that are initiated and managed by PAWS. The workflow is as follows:

- Proposals are submitted during the submission period through a collection form.
- A rigorous review process is then initiated to evaluate the science and feasibility aspects of the proposal and provide a numerical score.
- The proposals are then ranked, reviewed and graded by a panel of experts. Based on the reviews, the top proposals are selected to be given time on the telescope within the next year.
- A report is sent out at the end of the review process for each proposal, and those that have been given time are informed to submit their observation plans.
- The Science Operations team records each project's progress and telescope utilisation time to ensure that observation plans match the time allocated during the review process.

## DESIGN AND DEVELOPMENT

### *Requirements and Process*

Requirements for PAWS were provided by the Chief Scientist. They were captured and tracked through the Agile project management software called Jira.

A Kanban Agile approach was taken as the development process throughout the project, with the progress tracked on a Jira Kanban board.

\* rschwartz@sarao.ac.za

† tbaloi@sarao.ac.za

§ ssithole@sarao.ac.za

# OVERVIEW OF OBSERVATION PREPARATION AND SCHEDULING ON THE MeerKAT RADIO TELESCOPE

L.P. Williams, R. Schwartz, South African Radio Astronomy Observatory, Cape Town, South Africa

## Abstract

The MeerKAT radio telescope performs a wide variety of scientific observations. Observation durations range from a few minutes, to many hours, and may form part of observing campaigns that span many weeks.

Static observation requirements, such as resources or array configuration, may be determined and verified months in advance. Other requirements however, such as atmospheric conditions, can only be verified hours before the planned observation event.

This wide variety of configuration, scheduling and control parameters are managed with features provided by the MeerKAT software.

The short term scheduling functionality has expanded from simple queues to support for automatic scheduling (queuing). To support long term schedule planning, the MeerKAT telescope includes an Observation Planning Tool which provides configuration checking as well as dry-run environments that can interact with the production system. Observations are atomized to support simpler specification, facilitating machine learning projects and more flexibility in scheduling around engineering and maintenance events.

This paper will provide an overview of observation specification, configuration, and scheduling on the MeerKAT telescope. The support for integration with engineering subsystems is also described. Engineering subsystems include User Supplied Equipment (USE) which are hardware and computing resources integrated to expand the MeerKAT telescope's capabilities.

## OBSERVATION SPECIFICATION

An observation is a general term for any activity that is performed using the combined telescope systems. Observations are loosely divided in *science* and *engineering*. Science observations begin with a *proposal*. The proposal describes the science case/motivation, in addition to certain key requirements that may be fulfilled by the telescope. This initial information is described in a *proposal* submission. A proposal is a formal document that describes the full observation campaign including: the scientific merit, the required observing time and the telescope features/functionality required.

The technical feasibility of the proposal is assessed by the resident astronomers who are familiar with the telescope systems and capabilities. These individuals are able to translate the required telescope functionalities into the real telescope subsystems and resources that enable the required functionalities.

If the proposal is accepted, then the parameters are captured into a set that will enable configuration. The MeerKAT telescope [1] has a particular set of proposals

that will consume a majority of the observing time with only a smaller fraction being reserved for other proposals awarded at the director's discretion.

Engineering observations are generally shorter activities. The planning and scheduling of engineering observations is done directly with the Telescope Operations team, but uses most of the same mechanisms and data structures used for science observations.

## OBSERVATION CONFIGURATION

The observation configuration (also referred to as *experiment configuration*) on the MeerKAT telescope, inherits from the operational learnings of the Karoo Array Telescope (KAT-7) [2]. The KAT-7 telescope was composed of 7 antennas operating as a single array. It was operated 24 hours a day, 7 days a week with ad hoc maintenance events in daylight hours of a work week. This predecessor instrument experienced a relatively brief time as an operational telescope servicing a variety of engineering and scientific observations. For MeerKAT the two types of observations, engineering and science, are described further.

### Engineering Observations

Engineering observations are requested by engineers, technicians and science commissioners to characterise, test and monitor the telescope and its subsystems. The periods of engineering observations are defined in local time and they will require specific components to be exercised. This may or may not entail movement of the antennas in the array.

Real time sensor data is of prime interest to system engineers, rather than data collected from the receivers with some exceptions. Sensor data is used to assess hardware performance, predict failures and inform development of new features. The data set of interest is encapsulated by the period of the observation and contains sufficient information for the requesting party to continue their work.

### Science Observations

Science observations are requested by science commissioners and external scientists. These observations require specific movement patterns on or around sky coordinates. The observing periods are defined in sidereal time and will utilise specific telescope resources. Scientific data captured from the receivers is of prime interest but sensor data is also required for calibration purposes.

The astronomical sources which need to be observed are not always above the horizon. To gather sufficient data, the same observation would need to be repeated on different

# CAN MONITORING SOFTWARE FOR AN ANTENNA POSITIONER EMULATOR

V. van Tonder\*, G. Adams, M. Welz

South African Radio Astronomy Observatory, Cape Town, South Africa

## Abstract

The original Controller Area Network (CAN) protocol, was developed for control and monitoring within vehicular systems. It has since been expanded and today, the Open CAN bus protocol is a leading protocol used within servo-control systems for telescope positioning systems. Development of a CAN bus monitoring component is currently underway. This component forms part of a greater software package, designed for an Antenna Positioner Emulator (APE), which is under construction.

The APE will mimic movement of a MeerKAT antenna, in both the azimuth and elevation axes, as well as the positioning of the receiver indexer. It will be fitted with the same servo-drives and controller hardware as MeerKAT, however there will be no main dish, sub-reflector, or receiver. The APE monitoring software will receive data from a variety of communication protocols used by different devices within the MeerKAT control system, these include: CAN, Profibus, EnDAT, Resolver and Hiperface data.

The monitoring software will run on a BeagleBone Black (BBB) fitted with an ARM processor. Local and remote logging capabilities are provided along with a user interface to initiate the reception of data. The CAN component makes use of the standard SocketCAN driver which is shipped as part of the Linux kernel. Initial laboratory tests have been conducted using a CAN system bus adapter that transmits previously captured telescope data. The bespoke CAN receiver hardware connects in-line on the CAN bus and produces the data to a BBB, where the monitoring software logs the data.

## INTRODUCTION

The MeerKAT radio telescope consists of 64 Gregorian offset antennas, located in the Karoo desert of South Africa. Any individual telescope can move in both the azimuth and elevation axes to direct the telescope towards a specific astronomical source. In the azimuth axis, the telescope can move from  $-185^\circ$  to  $+275^\circ$ , where  $0^\circ$  points towards the North. In the elevation axis, the telescope can move from  $15^\circ$  to  $9^\circ$ . The azimuth slew rate is  $2^\circ/\text{s}$  with  $1^\circ/\text{s}^2$  acceleration whereas the elevation slew rate is  $1^\circ/\text{s}$  with  $0.5^\circ/\text{s}^2$  acceleration. A cable wrap exists for rotation in the azimuth axis. Each telescope is fitted with a rotating turret with four slots for different receivers. Currently three receivers are fitted onto the turret making the telescope receptive to UHF, L, and S-band frequencies.

The telescope makes use of a servo-control system that has been developed by CPI VERTEX ANTENNENTECHNIK

\* vereese@sarao.ac.za

GmbH (VA) [1]. The control system consist of an Antenna Conditioning Unit (ACU) and Lenze 9400 driver units. The driver units provide current for the motors. The azimuth drive assembly consists of two servo motors that are electrically torqued to avoid backlash. The controller uses three loops to achieve tracking: current, velocity, and position. There are various communication protocols involved in the control system including Controller Area Network (CAN), Profibus, EnDAT, Resolver and Hiperface data.

The Antenna Positioner Emulator (APE) project will consist of the same ACU, Lenze driver units, encoders and motors that a MeerKAT antenna is fitted with. For this project a monitoring system is currently being built and the CAN communication bus is the first protocol to be supported.

## CONTROLLER AREA NETWORK

We make use of the CANopen standard. The standard CANopen frame consists of an 11-bit frame ID, a remote transmission request (RTR) bit, followed by zero to eight bytes of data. The 11-bit frame ID consists of a 4-bit code to describe the functionality followed by a Node ID. Each device in the CAN network must have a unique node ID [2].

Table 1 summarises the CAN bus node ID's assigned to the different drivers.

Table 1: CAN Bus Nodes

Drive	Node ID
Azimuth 1	1
Azimuth 2	2
Elevation	5
Receiver indexer	9

In the control system, a function code of  $0x180$  has been assigned to describe sensors. Each of the  $0x180+\text{CAN}$  node ID frame consists of the four sensors listed in Table 2.

Table 2: CAN Sensors

Sensor	Data Length	Data Offset
Status	1	0
Motor temperature	1	1
Torque	2	2
Rotor position	4	4

## SOFTWARE ARCHITECTURE

The monitoring software will be implemented on a Beagle Bone Black (BBB) which is an ARM based system. The different software components will run as daemonized processes. Each communication protocol exists in its own

# MANAGING ROBOTICS AND DIGITIZATION RISK

D. Marais\*, J. Mostert, R. Prinsloo, Necsas SOC Ltd, Pretoria, South Africa

## Abstract

Robotic and digitization risks refer to the potential negative consequences that can arise from the use of robots and digital technologies in various industries, which include experimental physics control systems. Risks include the compromising or malfunctioning of these systems, resulting in injury, equipment damage, loss of data or disruptions to critical infrastructure and services. Notwithstanding the negative consequences, the benefits, including enhanced efficiency, productivity, accuracy, safety, cost savings, reduced human error, and real-time data access, typically outweigh the associated risks.

This paper provides a summary of how to moderate these risks by taking proactive steps to reduce the likelihood of negative consequences and minimize their impact if they do occur. A comprehensive risk management approach is proposed, that incorporates a combination of technical, organizational, and cultural strategies which can help mitigate the potential risks.

## INTRODUCTION

Robotic risks primarily concern physical machinery and autonomous systems, such as manufacturing robots causing workplace injuries due to malfunctions, or self-driving cars involved in accidents due to unforeseen road conditions. Digitization risks are associated with the use of digital technologies and data; for example, data breaches due to insufficient cybersecurity measures, or business disruptions from software failures. Though there is an overlap, like a hacked robotic system causing harm, robotic and digitization risks are distinct categories, each presenting unique challenges and consequences.

In the field of experimental physics, where researchers not only delve into the fundamental mysteries of the universe, but also engage in more routine pursuits such as calibrating sensors or analysing large data sets, the integration of robotic and digital technologies brings both unprecedented potential and a distinctive set of risks.

## Anecdotes

The advent of cutting-edge technologies like Large Language Models (LLMs) and text-to-image generators, used for tasks ranging from data analysis to experiment design, further amplify this potential and risk, setting the stage for a new era in experimental physics. These tools underscore the evolving landscape of risks in this field.

The following examples anecdotally highlights potential pitfalls and unforeseen consequences that can arise from the interaction between humans, machines, and digital transformation. These cases shed light on the delicate balance between innovation and vulnerability, emphasizing

the importance of proactive measures to mitigate the adverse outcomes that can result from the rapid evolution of technology.

**Bing Image Creator** Bing Image Creator generates images from natural language descriptions using the DALL-E Artificial Intelligence (AI) backend [1]. This system offers the potential to significantly reduce both time and costs when generating improved visualizations across a wide array of scientific fields. These applications span from conceiving and portraying novel materials to simulating and visualizing robot behaviours and interactions in diverse scenarios. Care in using such a system must however be taken as natural language can easily be misinterpreted as is demonstrated in Fig. 1, where Bing Image Creator was asked to “generate an image of a neutron powder diffractometer with a banana detector”. In the field of neutron scattering, a ‘banana detector’ is a neutron detector with a specific configuration and does not refer to a physical banana [2].



Figure 1: AI generated image of a neutron powder diffractometer with a banana detector.

**ChatGPT** ChatGPT is a language model developed by OpenAI, based on the GPT (Generative Pre-trained Transformer) architecture [3]. It is designed to generate human-like text based on the input it receives. ChatGPT is particularly well-suited for natural language processing tasks, including generating conversational responses, answering questions, providing explanations, and engaging in interactive text-based communication.

Whilst performing a literature review, ChatGPT was posed the question “What is the best software for neutron guide design?”. It responded with descriptions of “Monte Carlo Simulation of Time-of-Flight Spectroscopy (McStas)”, “Virtual Instrument for the Simulation of Spectrometer Experiments (Vitess)” and “Neutron Beam Line Analysis (NEBULA)”. Upon posing a follow-up question as to where NEBULA can be obtained, the reply was as

\* deon.marais@necsas.co.za

# OPC UA EPICS BRIDGE

W. D. Duckitt. Stellenbosch University, Stellenbosch, South Africa  
J. K. Abraham. iThemba LABS, Cape Town, South Africa

## Abstract

OPC UA is a service-orientated communication architecture that supports platform-independent, data exchange between embedded microcontrollers, PLCs or PCs and cloud-based infrastructure. This makes OPC UA ideal for developing manufacturer independent communication to vendor specific PLCs, for example. With this in mind, we present an OPC UA to EPICS bridge that has been containerized with Docker to provide a microservice for communicating between EPICS and OPC UA variables.

## INTRODUCTION

Open Platform Communications United Architecture (OPC UA) [1] supports platform-independent, data exchange between embedded microcontrollers, programmable logic controllers (PLCs) or personal computers (PCs).

Of particular interest to us, is the use of OPC UA for systems integration between the Experimental Physics and Industrial Control System (EPICS) [2] and infrastructure PLCs for heating, ventilation, and air conditioning (HVAC), water supply and electrical reticulation systems.

Our experience has shown that infrastructure suppliers typically have very little knowledge of EPICS, as EPICS is mainly used for the control of large scientific experiments. These suppliers, also typically implement their control systems in proprietary technology from PLC manufacturers.

In our opinion, it is left up to the EPICS developer to specify which variables within the PLCs are of importance and should be integrated with EPICS.

With all the major PLC suppliers [3-5] it is possible to install an OPC UA server within the PLC and export and serve these PLC variables over OPC UA. Once available via OPC UA, the systems can then be integrated into EPICS.

Within the EPICS community, OPC UA integration is possible using the EPICS device support module for OPC UA [6, 7]. The device support module, in the initial release, required the integration of a commercial C++ software development kit. This limitation prompted us to investigate open-source alternatives.

With our investigation, we found that it was possible to develop open-source OPC UA clients using Python [8, 9].

With this in mind, we have drawn from our experience in developing containerized EPICS systems in the React-Automation-Studio project [10], and present a system in the next sections which is containerized with Docker [11], deployable as microservices and implemented in Python using the Python SoftIOc project [12] and OPCUA-asyncio [9] modules.

Software

Control Frameworks for Accelerator & Experiment Control

## SYSTEM REQUIREMENTS

The goal was to develop a containerized software bridge capable of synchronizing and relaying process variable information between the OPC UA and EPICS domains.

It should be designed purely with open-source libraries and should be containerized with Docker [11] and version controlled as a mono-repository using Git [13].

Furthermore, the configuration mode of entry should be as docker compose files and EPICS record files.

Finally, the system should be sufficiently documented with use cases, examples and an implementation guide.

Each of the goals has been achieved, and the system overview is given below.

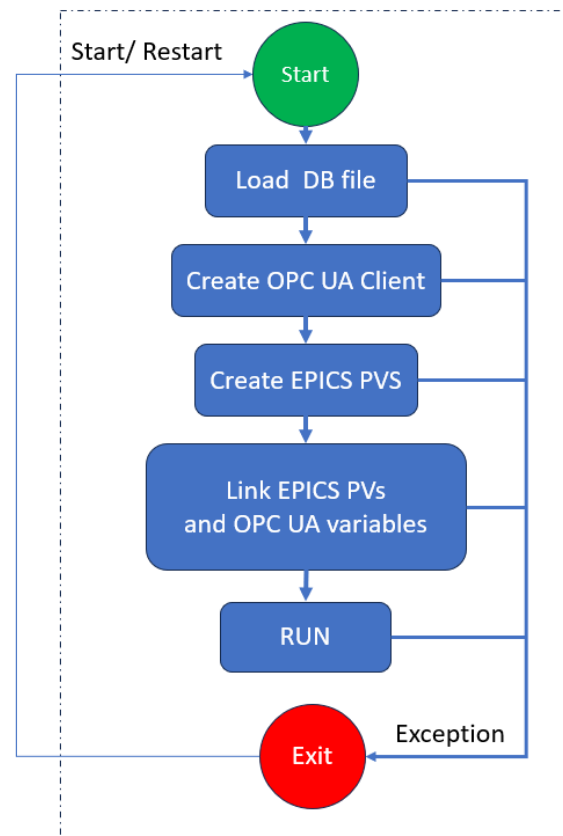


Figure 1: A state-machine diagram of the OPC UA EPICS bridge implementation.

## SYSTEM OVERVIEW

A high-level state machine diagram of implementation of the OPC UA EPICS bridge microservice is shown in Fig. 1. The microservice is written in Python [8] and

TUPDP075

681

# PRELIMINARY DESIGN FOR THE ALBA II CONTROL SYSTEM STACK

Sergi Rubio-Manrique, Fulvio Becheri, Guifre Cuni, Roberto Homs-Puron, Zbigniew Reszela  
ALBA-CELLS, Cerdanyola del Vallès

## Abstract

One of the main pillars of the ALBA Synchrotron Light Source (Barcelona, Spain) Strategy Plan is the preparation of ALBA to be upgraded to a fourth-generation light source. To accomplish this, a preliminary design of the accelerator has already been initiated in 2021. On the Computing side, the upgrade of the accelerator will require a comprehensive overhaul of most parts of the Control System, DAQ, Timing, and many other systems as well as DevOps strategies. This need for a major redesign will bring new architectural challenges, and opportunities to benefit from new technologies that were not present at the time ALBA was designed and build. This paper presents the preliminary design studies, pilot projects, new approaches to development coordination and management, and the preparation plan to acquire the knowledge and experience needed to excel in the ALBA II Control System Stack design.

## INTRODUCTION

### *ALBA II Preliminary Studies*

In 2019 the ALBA Synchrotron Accelerators Division published its first studies [1] on upgrading the existing storage ring to a 4th generation synchrotron lightsource with a low emittance and high brilliance machine; while reusing current injectors, tunnel and building. ALBA II White Paper [2] was finally released in May 2023, being the construction and commissioning of the new machine expected between 2028 and 2031.

ALBA's Computing and Controls Division have developed its own ALBA II preliminary study [3], structured in 11 different areas: Input/Output controller architecture, power supplies, timing system, equipment protection system, personnel safety system, IT architecture, motion control, control system stack, configuration management and stock management, machine learning, realtime processing needs.

This study focuses on the common software of the accelerators and beamlines control. Experiments control needs are explored in a separate article [4].

### *Control System Challenges of ALBA II*

ALBA II poses unique challenges as a 4th Generation Synchrotron Light Source. It demands smaller, more stable photon beams and a greater number of independently controlled magnets, using faster orbit feedback loops. Faster RF control loops and precise diagnostic elements are essential, especially for nano-focus beamlines.

To ensure control system alignment with final specifications, we started preliminary studies to foresee and evaluate the technologies required for ALBA II.

## *Current Status of ALBA Control System*

The ALBA Control System (CS) Stack consists of thousands of Tango [5, 6] Device Servers controlling hardware and software processes, user interfaces like Taurus GUI [7, 8] and distributed Tango Control System services (Archiving [9], Alarms [10], Sardana [11], etc.)

Recognizing component obsolescence, an upgrade began in 2019 [12], migrating to modern technology (64-bit CPUs, Tango 9.3, Debian Linux 10, Python 3.x) and applying methodologies that guarantee a continuous evolution of the stack: continuous integration and delivery pipelines, guided incremental change, continuous delivery and testing, fitness functions measurement and focus on performance. Paving the way to the ALBA II transition.

## *Control System Upgrade Work Packages*

The ALBA II control system upgrade requirements are being categorized into four distinct areas:

- Distributed System & Event. Related to IOCs, timing, IT infrastructure and real time loops.
- Tango Control System. New features to be introduced in our control framework.
- Graphical User Interfaces. Addressing new security and performance requirements, along with web technologies.
- DevOps: software development, packaging, integration, deployment and operation. Related with IOCs, IT Architecture and Configuration Management.

## DISTRIBUTED CONTROL SYSTEM

### *Preliminary Study*

A study of all communication processes have been started in order to detect control system issues (event overload, communication bottlenecks, RAM/CPU issues), and to improve the current interaction between accelerators and beamlines via gateway machines.

Our investigation identified bottlenecks mostly in the magnet power supplies subsystem, the one that will become the most critical in the upgrade to ALBA II. The problems that have been identified are:

- Dependency in old hardware and legacy event system (notifd) that limits the event rate.
- Bursts of events saturating applications.
- Slowness in event subscription due to the high number of attributes accessed by applications.
- CPU peaks causing python memory leaks.

Several strategies have been implemented to minimize these problems: upgrade of most control hosts, monitoring and restart of legacy event daemons, reconnection mecha-



## TOWARDS THE ALBA II : THE COMPUTING DIVISION PRELIMINARY STUDY

Oscar Matilla, Jose Avila-Abellan, Fulvio Becheri, Sergi Blanch-Torné†,  
 Albert Miret Burillo, Antonio Camps Gimenez, Isidre Costa, Guifre Cuni, Toni Fernández Maltas,  
 Roberto Homs, Jairo Moldes, Emilio Morales, Carlos Pascual-Izarra†, Antoni Pérez Font,  
 Sergi Pusó Gallart, Zbigniew Reszela, Borja Revuelta†, Alberto Rubio,  
 Sergi Rubio-Manrique, Jordi Salabert, Nil Serra, Xavier Serra Gallifa, Nicolas Soler,  
 Sergio Vicente Molina, Jorge Villanueva, ALBA Synchrotron, Cerdanyola del Vallés, Spain

### Abstract

The ALBA Synchrotron has started the work for upgrading the accelerator and beamlines towards a 4th generation source, the future ALBA II, in 2030. A complete redesign of the magnets lattice and an upgrade of the beamlines will be required. But in addition, the success of the ALBA II project will depend on multiple factors. First, after thirteen years in operation, all the subsystems of the current accelerator must be revised. To guarantee their lifetime until 2060, all the possible ageing and obsolescence factors must be considered. Besides, many technical enhancements have improved performance and reliability in recent years. Using the latest technologies will also avoid obsolescence in the medium term, both in the hardware and the software. Considering this, the project ALBA II Computing Preliminary Study (ALBA II CPS) was launched in mid-2021, identifying 11 work packages. In each one, a group of experts were selected to analyze the different challenges and needs in the computing and electronics fields for future accelerator design: from power supplies technologies, IOC architectures, or PLC-based automation systems to synchronization needs, controls software stack, IT Systems infrastructure or machine learning opportunities. Now, we have a clearer picture of what is required. Hence, we can build a realistic project plan to ensure the success of the ALBA II. It is time to get ALBA II off the ground.

### THE ALBA II

The ALBA is a 3rd generation Synchrotron Light facility located in Spain. ALBA received its first users in 2012 and serves over 2,000 scientists annually. Currently, in 2023, ALBA has ten beamlines in operation, which will increase up to fourteen, one per year, when XAIRA, FAXTOR, MINERVA, and 3SBAR begin providing services to the scientific community (see Fig. 1).

In parallel, the facility is starting to leap from the 3rd to the 4th generation by upgrading the accelerator, the beamlines and building new instruments to complete the scientific portfolio [1], giving birth to the ALBA II.

Currently, most of the light sources are facing a complete overhaul of their sources, upgrading the magnet lattice to the Multi-Bend Achromat (MBA) already successfully used by SIRIUS, MAX IV and the ESRF. The upgrade brings an enormous increase in brilliance at the sample point (see Fig. 2).

† on leave




Figure 1: Current Layout of ALBA.




Figure 2: Landscape of current and future photon sources after MBA lattice upgrade [2].

Taking advantage of this increased photon beam quality also requires upgrades on the most critical parts of the beamlines. In addition, the increased luminosity allows for tackling sub-nanometer focusing beamlines. This latter case requires designing very long beamlines, with a length higher than 150m, which exceeds the boundaries of the building that houses the ALBA Synchrotron. Therefore, the endstation of these beamlines will require the construction of new technical buildings (see Fig. 3).

The Spanish and Catalanian governments have determined that the ALBA II project is a top science priority project for the country. The plots needed for placing the long beamlines have already been donated, and additional funding has been arrived for recruiting new staff for the studies of the new MBA lattice in the past two years. The next step, a multiannual funding agreement that will allow start executing the largest required investments, is expected for next year. The ALBA II project will completely upgrade the storage ring and maintain the linac and booster accelerators. In addition, the current operating beamlines will also

# MANAGEMENT OF CONFIGURATION FOR PROTECTION SYSTEMS AT ESS

M. Carroll, G. Ljungquist, M. Mansouri, D. Paulic, A. Nordt  
European Spallation Source, Lund, Sweden

## Abstract

The European Spallation Source (ESS) in Sweden is one of the largest science and technology infrastructure projects being built today. The facility design and construction includes the most powerful linear proton accelerator ever built, a five-tonne, helium-cooled tungsten target wheel and 22 state-of-the-art neutron instruments. The Protection Systems Group (PSG), as part of the Integrated Control Systems (ICS) Division at ESS, are responsible for the delivery and management of all the Personnel Safety Systems (PSS) and Machine Protection Systems (MPS), consisting of up to 30 PSS control systems and 6 machine protection systems. Due to the bespoke and evolving nature of the facility, managing the configuration of all these systems poses a significant challenge for the team. This paper will describe the methodology followed to ensure that the correct configuration is correctly implemented and maintained throughout the full engineering lifecycle for these systems.

## INTRODUCTION

In any facility, good configuration management is essential for maintaining the reliability and integrity of complex systems by preventing system inconsistencies, uncontrolled changes and conformity issues when related to regulatory bodies. Further when it comes to safety and protection systems, the failure to implement robust configuration management can be catastrophic as was demonstrated in high profile accidents such as the piper alpha disaster in 1988 [1] and the deep water horizon accident in 2010 [2]. In both these cases, poor configuration management was credited as one of the key contributing factors where uncontrolled or improperly made changes ultimately led to multiple fatalities, irreversible environmental damage and ultimately severe regulatory penalties for the companies involved. Accidents like these and personal experience from working at multiple facilities with various levels of configuration management, has demonstrated the cost of overlooking this process, especially when designing and implementing safety and protection systems. Further the requirement for configuration management is mandated from the functional safety standards IEC 61511-1:2016 [3] used for the development of personnel safety systems, and EN 61508 [4] used for the development of machine protection systems at ESS.

## CONFIGURATION MANAGEMENT

The key elements of the configuration management strategy used by the PSG team are shown below and in Figure 1 with further details in the following chapters.

1. Configuration Identification
2. Version Control
3. Change Management
4. Configuration Auditing




Figure 1: Configuration management elements.

## CONFIGURATION IDENTIFICATION

The correct management of the configuration for PSG systems at ESS first requires that a system configuration is accurately identified in the initial design phase. This establishes an approved baseline for the system through a detailed documentation strategy. To achieve this each system has the following key documentation package developed through the design process.

- Concept of Operation (Conops)
- System Requirement Specification (SRS)
- Interface Control Documents (ICDs)
- Detailed Design Specification (DDS) / Electrical Schematics
- Test Specifications (see configuration auditing)

### Concept of Operation

The ‘Conops’ document identifies at a high level the key conceptual requirements of the system, such as who the main interfacing systems and stakeholders are, how the stakeholders are expected to interact with the system and how the system is expected to interact with the rest of the facility.

### System Requirement Specification

The *SRS* is a document used to collect all the requirements for a system, with a unique ‘SRS-ID’ for each requirement and a link back to the source of the requirement. Sources typically come from operational requirements originated in the Conops or as Safety Implemented Functions (SIFs) described in Safety Integrity Level (SIL) assessments or machine protection (MP) functions (PFs) from an MP analysis. An example can be seen in Table 1.

# AUTOMATED PROCEDURE FOR CONDITIONING OF NORMAL CONDUCTING ACCELERATOR CAVITIES

E. Trachanas\*, G.Fedel, S. Haghtalab, B. Jones, R. Zeng  
European Spallation Source, ESS-ERIC Lund, Sweden

A. Gaget, O. Piquet- CEA-IRFU, University Paris-Saclay, Gif-sur-Yvette, France  
Luca Bellan, Carlo Baltador, Francesco Grespan- INFN/LNL, Legnaro

## Abstract

Radio frequency (RF) conditioning is an essential stage during the preparation of particle accelerator cavities for operation. During this process the cavity field is gradually increased to the nominal parameters enabling the out-gassing of the cavity and the elimination of surface defects through electrical arcing. However, this process can be time-consuming and labor-intensive, requiring skilled operators to carefully adjust the RF parameters. This contribution presents the software tools for the development of an automated EPICS (Experimental Physics and Industrial Control System) control application with the aim to accelerate and introduce flexibility to the conditioning process. The results from the conditioning process of the ESS Radio-Frequency Quadrupole (RFQ) and the parallel conditioning of Drift-Tube Linac (DTL) tanks will be presented demonstrating the potential to save considerable time and resources in future RF conditioning campaigns.

## INTRODUCTION

The European Spallation Source (ESS) linear accelerator is designed to accelerate a 62.5 mA, 2.86 ms, 14 Hz proton beam up to 2 GeV to the entrance of a rotating tungsten target and is divided in a normal-conducting (NCL) and a superconducting part (SCL). The normal conducting part is comprised of electromagnetic cavities at room temperature operating at 352.21 MHz containing the ion source, a Low Energy Transport (LEBT), a Radiofrequency Quadrupole (RFQ), a Medium Energy Beam Transport (MEBT) and five (5) Drift-Tube Linac (DTL) tanks (as shown in Fig. 1) [1].

The ESS RFQ was delivered by CEA-Saclay equipped with four (4) modulated vanes and a total length of 4.55 m accelerating the proton beam up to 3.6 MeV with an inter-vane voltage of 120 keV and a duty cycle of 4% [1]. The ESS RFQ was installed at ESS accelerator tunnel in 2019 and following the local and integrated testing phase, power conditioning was successfully completed in summer 2021 (see Fig. 2) and the first proton beam injected in the RFQ in October of the same year [2–4].

The Drift Tube Linac (DTL) is an in-Kind collaboration from INFN composed of five (5) tanks of 8 m each accelerating the beam up to 90 MeV. DTL 1 has been conditioned and commissioned with beam in summer 2022 [5, 6] followed by the installation of DTL tanks 2,3,4 in late 2022. The RF conditioning process for DTL 2,3,4 along with beam

commissioning process started at the beginning of 2023 and completed in late summer of the same year [7,8]. DTL 5 was installed at the ESS accelerator tunnel in September 2023 and the preparatory testing phase before RF conditioning is ongoing.




Figure 1: The ESS Radio Frequency Quadrupole and an ESS Drift Tube Linac tank.

RF conditioning is a time consuming and challenging process that lasts for several weeks and requires multiple parameters constant monitoring especially when performed for multiple cavities simultaneously. RF conditioning is a standard process before operation where the cavity field is progressively increased to the nominal level.




Figure 2: Peak and Average power in the RFQ cavity between 25-06-2021 – 17.07.2021 showing the gradual increase of cavity field to the nominal values.

During this process, different mechanisms such as out-gassing, multipacting or field emission can lead to electrical arcing which to some degree is a desirable effect for surface defect elimination. Extensive arcing and vacuum breakdowns can potentially lead to cavity and equipment damage, making the need for a versatile interlock system essential. In

\* Emmanouil.Trachanas@ess.eu

# THE ESS FAST BEAM INTERLOCK SYSTEM - DESIGN, DEPLOYMENT AND COMMISSIONING OF THE NORMAL CONDUCTING LINAC

S. Pavinato\*, M. Carroll, S. Gabourin, A. Gorzawski, A. Nordt  
European Spallation Source, Lund, Sweden

## Abstract

The European Spallation Source (ESS) is a research facility based in Lund, Sweden. Its linac will have an high peak current of 62.5 mA and long pulse length of 2.86 ms with a repetition rate of 14 Hz. The Fast Beam Interlock System (FBIS), as core system of the Beam Interlock System at ESS, is a critical system for ensuring the safe and reliable operation of the ESS machine. It is a modular and distributed system. FBIS will collect data from all relevant accelerator and target systems through 300 direct inputs and decides whether beam operation can start or must stop. The FBIS operates at high data speed and requires low-latency decision-making capability to avoid introducing delays and to ensure the protection of the accelerator. This is achieved through two main hardware blocks equipped with FPGA based boards: a mTCA 'Decision Logic Node' (DLN), executing the protection logic and realizing interfaces to Higher-Level Safety, Timing and EPICS Control Systems. The second block, a cPCI form-factor 'Signal Conversion Unit' (SCU), implements the interface between FBIS inputs/outputs and DLNs. In this paper we present the implementation of the FBIS control system, the integration of different hardware and software components and a summary on its performance during the latest beam commissioning phase to DTL4 Faraday Cup in 2023. Keywords: EPICS, xTCA, FPGA, commissioning, operation.

## INTRODUCTION

ESS [1], located in Lund, Sweden, is on the forefront of neutron science research. ESS has designed its Machine Protection system [2,3] to strike a delicate balance between safeguarding equipment from potential damage and ensuring high beam availability. This equilibrium is of paramount importance due to the unprecedented proton beam power of 125 MW per pulse (with an average of 5 MW). Uncontrolled release of this energy could result in catastrophic damage to equipment within microseconds. To address this critical concern, ESS has developed the FBIS, which is engineered for minimal latency. The FBIS plays a pivotal role in ESS operations, tasked with collecting a diverse range of data from various Sensor Systems, including both slow systems such as Vacuum, Magnets, Personal Safety System (PSS) as well as fast systems like Radio Frequency Control Local Protection System and Beam Instrumentation. These data inputs serve as vital signals for processing units (PU) within the FBIS, which are responsible for making real-time decisions on whether to continue or halt beam production,

thereby ensuring the safe and reliable operation of the ESS accelerator.

## ARCHITECTURE OF FBIS

The FBIS architecture is fully redundant to ensure it can reach the Protection Integrity Level requested by the Protection Functions. It was developed by a team constituted of the ESS in-kind partner ZHAW [4], the ESS electronic standard manufacturer IOxOS Technologies [5], and ESS FBIS team. Hundreds of Sensor Systems connection are foreseen along the 600 m Linac, leading to a specific architecture with two component types: the "Signal Conversion Unit" (SCU) and the "Decision Logic Node" (DLN).

### Signal Conversion Unit (SCU)

The SCU, Fig. 1, is a concentrator for Sensor Systems connections. It is based on a cPCI standard chassis with custom electronic cards. An SCU hosts up to 12 Mezzanine Cards (MC) on which Sensor Systems connect. Two more cards, called Serializers, host a MPSoC Zynq Ultra-scale+ to manage the connection to the MC by the backplane, and the communication with the DLN through serial links (Slink), redundant optical connections using the Xilinx Aurora 64b/66b core IP.




Figure 1: Signal Conversion Unit - SCU.

Several kind of MC were designed to interface with various types of Sensor Systems. One of them manages PLC-based Sensor Systems [6], and two others fast electronic Sensor Systems, mainly FPGA-based. Twenty five SCUs are foreseen along the Linac and in the Target building. The two first SCUs, installed close to the Ion Source, host also specific MC to interface with five hardware Actuators, acting to stop the Beam Production upon DLNs request.

### Decision Logic Node (DLN)

The DLN, as depicted in Fig. 2, is responsible for implementing the protection logic utilized by the FBIS. Furthermore, it facilitates connections to the Higher-Level Safety System, Control System, and the ESS Timing System. The

\* stefano.pavinato@ess.eu

# ESS TARGET SAFETY SYSTEM MAINTENANCE

A. Sadeghzadeh<sup>†</sup>, O. Ingemansson, O. Janson, M. Olsson, L. Coney  
 European Spallation Source, Lund, Sweden

## Abstract

The Target Safety System (TSS) is part of the overall radiation safety plan for the Target Station in the European Spallation Source (ESS). The Target Controls and Safety group (ESS – Target Division) is responsible for the design and construction of the TSS.

The TSS stops proton production if vital process conditions measured at the Target Station are outside the set boundaries with the potential to cause an unacceptable radiation dose to third parties (public outside ESS).

The TSS is a 3-channel fail-safe safety system consisting of independent sensors, a two redundant train system based on relay and safety PLC techniques and independent ways of stopping the proton beam accelerator.

The TSS will continuously monitor safety parameters in the target He cooling, wheel and monolith atmosphere systems, evaluate their conditions, and turn off the proton beam if necessary.

After passing several stages of off-site testing, the TSS cabinets are now installed on site and successfully passed internal integration.

In this paper, we explain the features built into the system to ease emergency repairs, system modification and system safety verification and general maintainability of the system.

## TSS SYSTEM OVERVIEW

The TSS is a distributed system, including electrical, I&C (Instrumentation and Control) and mechanical equipment with components placed in locations across the ESS facility according to Fig 1. The TSS consists of sensors, signal collecting units, signal intermediaries, logic equipment, software and actuators.




Figure 1: TSS distribution in ESS facility.

The basic design for the TSS consists of a three-channel (A, B and C) system divided into two main trains (relay trains and PLC-train), each of them with separated 2003 logic for tripping the proton beam. The TSS follows the

fail-safe, redundancy, diversity and separation principles. The fundamental architecture of the TSS is shown in Fig 2.

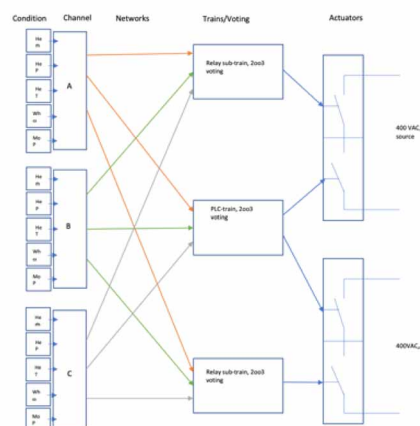


Figure 2: Basic architecture of the TSS.

## Physical Separation

Physical separation for the TSS is obtained through a design containing mechanical barriers and/or distance between three channels to ensure that one event, internal or external, does not impact more than one channel. To achieve this, electrical cabinets of the three TSS channels are located in different floors with different fire zones and their cables are routed through different metal conduits.

## Functional Separation

Functional separation for the TSS is obtained via the following principles:

- The channels are galvanically separated at the 2003 voting relay setup; each channel has a dedicated relay. There is no communication between channels.
- The trains are galvanically separated by having dedicated actuators, i.e., one train cannot affect any other train. There is no communication between the trains.

## Diversity

In order to have protection against common cause failure, diversity for the TSS is obtained through:

- A two-train solution with different technologies: one train with electromechanical components (relay, divided into two sub-trains) and one train with programmable electronics (PLC)
- Diverse components to disconnect the electrical supply to the Ion Source and RFQ (radio-frequency quadrupole), using contactors or switch disconnectors;
- Diverse mechanisms to stop proton production via the RFQ or Ion Source;

<sup>†</sup> atefeh.sadeghzadeh@ess.eu

# DAQ SYSTEM BASED ON TANGO, SARDANA AND PANDABOX FOR MILLISECOND TIME RESOLVED EXPERIMENT AT THE COSAXS BEAMLINE OF MAX IV LABORATORY

Vanessa Da Silva\*, Roberto Appio, Aureo Freitas, Byungnam Ahn, Marcelo Alcocer, Mirjam Lindberg, Ann Elizabeth Terry, Tomas Sigfrido Plivelic  
MAX IV Laboratory, Lund, Sweden

## Abstract

CoSAXS is the Coherent and Small Angle X-ray Scattering (SAXS) beamline placed at the diffraction-limited 3 GeV storage ring at MAX IV Laboratory. The beamline can deliver a very high photon flux of  $10^{13}$  ph/s and is equipped with state-of-the-art pixel detectors, suitable for experiments with a high time-resolution to be performed. In this work we present the upgraded beamline data acquisition strategy for a millisecond time-resolved SAXS/WAXS experiment, using laser light to induce temperature jumps or UV-excitation with the consequent structural changes on the system. In general terms, the beamline control system is based on TANGO and built on top of it, Sardana provides an advanced scan framework. In order to synchronize the laser light pulse on the sample, the X-ray fast shutter opening time and the X-ray detectors readout, hardware triggers are used. The implementation is done using PandABox, which generates the pulse train for the laser and for all active experimental channels, such as counters and detectors, in synchronization with the fast shutter opening time. PandABox integration is done with a Sardana Trigger Gate Controller, used to configure the pulses parameters as well to orchestrate the hardware triggers during a scan. This paper describes the experiment orchestration, laser light synchronization with multiple X-ray detectors.

## INTRODUCTION

CoSAXS (Coherent and Small Angle X-ray Scattering) beamline [1] is located at the 3 GeV ring of the Swedish 4th generation synchrotron, MAX IV, featuring high photon flux  $1 \times 10^{12} - 1 \times 10^{13}$  ph/s and small X-ray spot size. As the state-of-the-art Small Angle X-ray Scattering (SAXS) beamline can deliver high intensity and enable experiments with a high time resolution to be performed.

Time-resolved (TR) SAXS/WAXS (where WAXS stands for Wide Angle X-ray Scattering) experiments require a source of perturbation which might cause a structural change in the sample, while fast data acquisition is performed along the time to monitor the dynamics. Currently, CoSAXS beamline offers lasers for temperature jumps time-resolved studies as one of the sample environments available for the users [2].

Following the MAX IV standard, the beamline control system is based on TANGO [3]. Built on top of TANGO, Sardana [4] it is used for both scan configuration and execution.

\* vanessa.silva@maxiv.lu.se

In the present configuration, CoSAXS time-resolved experiment uses either an Infrared (IR) or an Ultraviolet (UV) laser pulse to induce a perturbation in the sample system. The laser pulse duration and the time synchronization for data acquisition of both SAXS and WAXS detectors are controlled by PandABox [5]. PandABox is a FPGA based system and it is integrated in the control system as a TANGO device. On top of it, a Sardana trigger/gate controller is the main interface used for configuration during the scan.

Besides reaching small time-resolutions, avoiding radiation damage is also another challenge faced on this experiment, due to very high beam intensity on the sample. A fast shutter is used to control the sample exposure to radiation in each acquisition window.

## EXPERIMENTAL SETUP

The time-resolved sample environment at CoSAXS is primarily developed to study proteins in aqueous solutions. The experimental setup is illustrated in Fig. 1. The sample holder is a temperature-controlled flow cell, consisting of a 1.5 mm inner-diameter quartz capillary and 10  $\mu$ m wall thickness. The sample delivery is done from a reservoir using a Reglo ICC peristaltic pump.

Two lasers are available for the TR setup: a continuous wave IR laser (1470 nm) and a Q-Switch nanosecond laser with a 532 nm pump laser and an OPO which can tune the wavelength over ca. 700-2000 nm. The laser pulse is delivered to the sample through a fiber physically attached to the flow cell positioned a few millimeters from the capillary wall. The other end of the light guide is coupled to the laser, either IR or UV, and it consists of a 200 mm and 0.22 NA fiber.

Data acquisition is performed by a dual hybrid pixel X-ray detectors system (see Fig. 1). The SAXS detector, positioned inside a vacuum vessel, is a Eiger2 4M and the readout can be performed at 500 Hz. Two WAXS detectors are available at the beamline and can be used, not simultaneously, in TR experiments. The first one is a Mythen2 1K that can be operated at 1 kHz and it is positioned in air, right after the flow cell. The second one is a Pilatus3 2M in an L-shape configuration, positioned in front of Eiger2 in the vacuum vessel, and it can be operated at 250 Hz.

Premature sample radiation damage is avoided with CEDRAT fast shutter. The fast shutter takes up to 8 ms to fully open/close and it can be easily controlled with TTL external signals. It also provides a feedback signal reflecting the hardware current state. In the TR experiment setup, the fast

# EPICS AT FREIA LABORATORY

K. Gajewski<sup>†</sup>, K. Fransson, FREIA Laboratory, Uppsala University, Uppsala, Sweden

## Abstract

The FREIA laboratory is a Facility for REsearch Instrumentation and Accelerator development at Uppsala University, Sweden. It was officially opened in 2013 to test and develop superconducting accelerating cavities and their high-power RF sources. The laboratory focuses on superconducting technology and accelerator development and conducts research on beam physics and light generation with charged particles, accelerator technology and instrumentation. From the very beginning EPICS has been chosen as a control system for all the infrastructure and equipment in the lab. Use of EPICS allowed us to build a robust, expandable and maintainable control system with a very limited man power. The paper will present the choices we made and the problems we've solved to achieve this goal. We will show the current status of the control system and the strategy for the future.

## INTRODUCTION

FREIA laboratory was established as part of the Department of Physics and Astronomy at Uppsala University. A 1000 square meter experimental hall has been built at the Ångström Laboratory campus [1]. The construction started in May 2012 and was ready for installation of the equipment in June 2023. The laboratory was jointly funded by Uppsala University and a grant from the Swedish Government. The helium liquefaction plant was sponsored by the Kurt and Alice Wallenberg Foundation.

In the initial phase of operation FREIA was functioning as a cryogenic centre, delivering liquid helium and liquid nitrogen for the users at Uppsala University and other institutes in Uppsala and as a test facility for superconducting RF cavities followed by the testing of the cryomodules for the European Spallation Source (ESS) in Lund, Sweden. After installation and commissioning of a, so called, vertical cryostat, tests of bare superconducting cavities and superconducting magnets have been conducted for CERN and other European laboratories. Figure 1 shows an overview picture of the experimental hall. Behind the bunker there is a liquefier and the vertical cryostat (below the floor level) as shown in Fig. 2.

The decision to use EPICS [2] as a control system was made from the very beginning. The idea was to connect all the equipment and instrumentation to a common platform that would support a uniform access and services like alarms, archiving, save/restore etc. The choice of EPICS was highly influenced by the fact that ESS supported us with equipment running EPICS: LLRF controls, timing system as well as stepper and piezo motor controllers for the cavity tuning systems.

<sup>†</sup> konrad.gajewski@physics.uu.se.

The specification for all ordered subsystems included the requirement for easy integration with EPICS.



Figure 1: FREIA hall.



Figure 2: FREIA hall – liquefier and vertical cryostat.

## INFRASTRUCTURE AND EQUIPMENT

### Helium Liquefaction Plant

The first equipment installed in FREIA laboratory was a helium liquefier system. It is a system delivered by Linde capable to produce up to 140 l/h of liquid helium at 4 K. It came with its local controller based on the Siemens S7-300 series PLC and an operator's interface built on WinCC. The PLC gives us read-only access in EPICS to almost all process variables and operator level write possibilities. The engineer level controls are possible only through WinCC.

### Sub-Atmospheric Pump Station

The station, consisting of 7 pumps working together, is used for cooling the 4 K liquid helium down to 1.9 K. The total capacity of the pumps is 4.3 g/s at 15 mbar. The local controller delivered by the vendor is based on Siemens S7-

# OPERATIONAL TOOL FOR AUTOMATIC SETUP OF CONTROLLED LONGITUDINAL EMITTANCE BLOW-UP IN THE CERN SPS

N. Bruchon\*, I. Karpov, N. Madysa<sup>1</sup>, G. Papotti, D. Quartullo, CERN, 1211 Geneva, Switzerland  
<sup>1</sup>currently at GSI, 64291 Darmstadt, Germany

## Abstract

The controlled longitudinal emittance blow-up is necessary to ensure the stability of high-intensity LHC-type beams in the CERN SPS. It consists of diffusing the particles in the bunch core by injecting a bandwidth-limited noise into the beam phase loop of the main 200 MHz RF system. Obtaining the correct amplitude and bandwidth of this noise signal is non-trivial, and it may be tedious and time-demanding if done manually. An automatic approach was developed to speed up the determination of optimal settings. The problem complexity is reduced by splitting the blow-up into multiple sub-intervals for which the noise parameters are optimized by observing the longitudinal profiles at the end of each sub-interval. The derived bunch lengths are used to determine the objective function which measures the error with respect to the requirements. The sub-intervals are tackled sequentially. The optimization moves to the next one only when the previous sub-interval is completed. The proposed tool is integrated into the CERN generic optimization framework that features pre-implemented optimization algorithms. Both single- and multi-bunch high-intensity beams are quickly and efficiently stabilized by the optimizer, used so far in high-intensity studies. A possible extension to Bayesian optimization is being investigated.

## INTRODUCTION

Given the strict beam parameter constraints to be met before injecting into the LHC, maintaining longitudinal stability in the SPS is an essential task. Stability relies on the increased synchrotron frequency spread thanks to a double-harmonic RF system, as well as on the controlled longitudinal emittance blow-up. Both techniques increase the synchrotron frequency spread within the bunch, enhancing Landau damping [1, 2].

The controlled longitudinal emittance blow-up is based on the injection of bandwidth-limited phase noise into the beam phase loop which locks the main RF system operating at 200 MHz to the bunch phases. A noise signal with a bandwidth-limited excitation spectrum is needed. The spectrum, as well as the bandwidth, is defined by the cutoff frequencies that follow the variation of the small-amplitude synchrotron frequency,  $f_{s0}$ , during the acceleration ramp. The low and high cutoff frequencies are named  $f_{low}$ , and  $f_{high}$  respectively. By normalizing those values with respect to  $f_{s0}$ , the ratios called “margin low”  $m_{low} = f_{low}/f_{s0}$  and “margin

high”  $m_{high} = f_{high}/f_{s0}$  are defined. Figure 1 illustrates the relation between  $f_{s0}$ , the frequency band, and the emittance.




Figure 1: Example of a normalized synchrotron frequency distribution (black) in a double-harmonic RF system. The horizontal dashed lines indicate  $m_{low}$  and  $m_{high}$  respectively. The vertical blue line marks the longitudinal emittance.

These normalized settings are easier to manage for tuning the noise bandwidth since their value remains constant despite the changes of  $f_{s0}$  during acceleration. The aim of the blow-up is to impact the bunch core exclusively, without increasing the tail population of the particle distribution, which would risk generating losses. In addition, the effects on the longitudinal beam profiles are also dependent on the blow-up amplitude,  $a$ , defined in rms degrees of the 200 MHz RF system. A low-amplitude noise signal may be ineffective, while too high amplitude can negatively affect the bunch distribution. An additional parameter is the time interval during the cycle the RF manipulation is applied.

Fine-tuning the blow-up settings is challenging since multiple time-dependent settings are involved. The manual procedures are time-intensive and cannot guarantee optimal noise settings. Moreover, even optimal settings necessitate reevaluation when parameter changes occur, e.g., higher bunch intensity or a different voltage program. To simplify the optimization, efforts have been recently dedicated to automating the setup of the noise for the controlled blow-up in the SPS. Initial studies in this direction were outlined in [3], where the automatic control of the blow-up was demonstrated for single-bunch beams. The software is integrated as part of the CERN Machine Learning (ML) platform [4], which provides pre-implemented generic optimization algorithms. The advantage of relying on this framework lies in its ambition to integrate numerical optimization, machine learning, and reinforcement learning into routine accelerator operation. The extension of the software for single- to multi-bunch beams has been straightforward; the management of the settings did not change, as well as obtaining the observations. However, a more careful reassessment was necessary for the cost function, which must take into consideration

\* niky.bruchon@cern.ch



# ENHANCING MEASUREMENT QUALITY IN HL-LHC MAGNET TESTING USING SOFTWARE TECHNIQUES ON DIGITAL MULTIMETER CARDS-BASED SYSTEM

H. Reymond, O.O. Andreassen, C. Charrondiere, M. Charrondiere,  
P. Jankowski, CERN, Geneva, Switzerland

## Abstract

The HL-LHC magnets play a critical role in the ongoing High-Luminosity Large Hadron Collider project, which aims to increase the luminosity of the LHC and enable more precise studies of fundamental physics. Ensuring the performance and reliability of these magnets requires high-precision measurements of their electrical properties during testing. To meet the R&D program needs of the new superconducting magnet technology, an accurate and generic voltage measurement system was developed after the testing and validation campaign of the LHC magnets. The system was based on a set of digital multimeter (DMM) cards installed in a PXI modular chassis and controlled using CERN's in-house software development. It allowed for the measurement of the electrical properties of the magnet prototypes during their study phase. However, during the renovation of the magnet test benches and in preparation for the HL-LHC magnet series measurement, some limitations and instabilities were discovered during long recording measurements. As a result, it was decided to redesign the measurement system. The emergence and promises of the new PXIe platform, along with the requirement to build eight new systems to be operated similarly to the existing four, led to a complete redesign of the software. This article describes the various software techniques employed to address platform compatibility issues and significantly improve measurement accuracy, thus ensuring the reliability and quality of the data obtained from the HL-LHC magnet tests.

## INTRODUCTION

During the LHC magnet series tests in early 2000's, several mobile measurement systems were built to allow accurate analysis of their properties, with the same temperature and pressure conditions expected in the LHC emulated in the SM18 magnet test hall. A generic voltage measurement device [1], based on six KEITHLEY® 2001 digital multimeters (DMM) was used to record data for the specific acceptance tests: Residual Resistivity Ratio (RRR), energy losses in superconducting cable strands (Loss) and coil splice resistance measurement (Splice). The multimeters were remotely controlled via a GPIB bus from a SUN microsystem workstation, running a dedicated LabVIEW® software. Three mobile racks were used for RRR measurements and two for the Loss/Splice tests. They were extensively used for the 1232 dipoles and 480 quadrupoles magnet tests, prior to their installation in the LHC tunnel.

## USING THE PXI PLATFORM AS A NEW STANDARD

After the intensive qualification campaign of the LHC magnets, a thorough review of the existing measurement and control systems took place, with the aim of preparing for the future HL-LHC [2] project. The availability of the PXI platform [3], which has become a new industry standard for Commercial Off-The-Shelf (COTS) modular systems, coupled with the approaching end-of-life for VME components, influenced the choice for the next generation of equipment.

Numerous measurement systems have been redesigned using PXI components. This same approach was applied to the legacy RRR and Loss/Splice mobile racks. After conducting a market survey, the PXI-4071 card from National Instruments (NI) was selected. This 7½-Digit Digital Multimeter (DMM) was deemed a strong candidate to handle all three test types within the same system.

## THE NEW DMM MOBILE RACKS

The redesign plan aimed not only to introduce a new, more accurate, and versatile voltage measurement system but also to incorporate new features specifically designed to streamline the operation and testing of magnets.

### *The Hardware Considerations*

The DMM system consists of an 18-slot PXI-1045 chassis that contains the following components: a PXI-8108 CPU running Windows 7 and the LabVIEW® application, up to 16 DMM cards, and a PXI-6221 DAQ card. This entire setup is installed in a 29 U mobile rack, complete with a screen, keyboard, and mouse for user interaction.

The DAQ card was used to read digital signals from a dedicated connector attached to the local magnet test bench. This simple connection reduced many sources of error, and allowed error tracking in the result files by automatically embedding the bench name as a 3-bit code.

The GUI incorporated a control panel for power converters. It featured a current profile editor and allowed for the selection and control of the converter, enabling comprehensive management of magnet current during DMM measurements.

The choice of the chassis and CPU was made in conjunction with hardware selections made two years prior for the renovation of the HF/LF DAQ systems. Ten of these systems were deployed across SM18 test benches, and the decision to use the same chassis and CPU was driven by maintenance considerations.

# LABVIEW-BASED TEMPLATE FOR ENHANCED ACCELERATOR SYSTEMS CONTROL: SOFTWARE SOLUTIONS FOR THE CERN-ISOLDE FACILITIES

C. Charrondiere, L. Le, R. Heinke, R. E. Rossel, S. Rothe, E. Galetti,  
O. Ø. Andreassen, B. A. Marsh, S. Sudak, CERN, Geneva  
A. Benoit, G. Boorman, ANGARA Technology, Geneva

## Abstract

ISOLDE is part of the experimental infrastructure within the CERN accelerator complex that provides radioactive ion beams for studies of fundamental nuclear physics, astrophysics, condensed matter physics and medical applications. Complementing the available controls infrastructure, an easy-to-use set of applications was developed to allow operators to record and display signals from multiple sources, as well as to provide drivers for non-standard, custom-made instruments and specialized off-the-shelf components.

Aimed not only at software engineers but in general targeting developers without any specific background in software engineering, a generic and modular software template was developed in LabVIEW following a collaboration between CERN and ANGARA Technology. This unified template can be extended to support interaction with any instrument and incorporate any newly developed application to the existing control system and integrated into the CERN control and monitoring infrastructure. New modules and instrument drivers are easy to maintain as the structure and communication layers are all derived from the same template and based on the same components.

In this paper, we will explain the implementation, architecture, and structure of the template, as well as a wide variety of use cases - from motor control to image acquisition and laser-specific equipment control. We will also show use cases of applications developed and deployed within a few days in the ISOLDE facility. **Keywords:** reconfigurable, control system, rapid deployment, maintenance, LabVIEW.

## MOTIVATION AND PURPOSE

The Isotope Separator On-Line (ISOLDE) [1] at CERN is a facility dedicated to the production of radioactive beams using a Resonance Ionization Laser Ion Source (RILIS) [2] to efficiently ionize specific isotopes of interest. In addition to ISOLDE, there is the MEDICIS [3] facility, including both RILIS and its counterpart MELISSA [4], which focuses on the production of medical radioisotopes. There are also several auxiliary offline test benches employed for quality control and research and development purposes [5], including two offline isotope mass separators known as OFFLINE1 and OFFLINE2 [6].

In this paper, we discuss the significance of avoiding large monolithic and complex applications, as such systems are typically challenging to maintain and sustain over time. We introduce the Module Template, which empowers

developers to construct independent hardware drivers and software projects. These projects are designed to assist both machine operators and code maintainers when changes or updates are required.

The Module Template serves as a project template that already incorporates default functionality, offering a modular inter-process communication architecture. This framework enables developers to swiftly deploy an interface, promotes the upkeep of clean and organized code, and establishes a common communication approach for both software and hardware communication layers. Additionally, it streamlines the integration of existing software, such as that for the laser system [7], ensuring greater consistency and seamless integration into CERN's general network and data handling infrastructure.

## CHALLENGES

An important variable in making design choices for the template is the consideration of the users and the code maintainers perspective: a balance must be struck between the amount of abstraction and hidden communication layers versus ease of use. For short-term affiliates such as students, for example, the aim is to reduce the initial effort required to make contributions. The template should also be comprehensible for the physicists and academics working at ISOLDE, allowing them to develop their own domain-specific applications.

The Module Template must be able to integrate a wide range of communication interfaces, such as Control Middleware (CMW) [8], a software communication protocol widely used at CERN, various hardware communication protocols, as well as a command line interface and the capability to handle a variety of configuration files.

## TEMPLATE DESIGN

### General Layout

Figure 1 shows the Module Template application architecture consisting of software modules that represent a distributed monitoring and control system. Each module can be comprised of one or more components (e.g. Manager and Worker). A module can be responsible for a specific task, such as interacting with a hardware device, performing stabilization tasks, or presenting a user interface. The Module Template is dependent on the LabVIEW RADE framework [9] and provides a coherent integration with the CERN control infrastructure [10].

# IMPROVING CERN'S WEB-BASED RAPID APPLICATION PLATFORM

E. Galatas\*, S. Deghaye, J. Raban, C. Roderick, D. Saxena, A. Solomou  
CERN, Geneva, Switzerland

## Abstract

The Web-based Rapid Application Platform (WRAP) aims to provide a centralised, zero-code, drag-n-drop means of GUI creation. It was created at CERN to address the high maintenance cost of supporting multiple, often deprecated, solutions and potential duplication of effort. WRAP leverages web technologies and existing controls system infrastructure to provide a drop in solution for a range of use cases. However, providing a centralized platform to cater for diverse needs and to interact with a multitude of data sources presented performance, design, and deployment challenges. This paper describes how the architecture evolved to address technological limitations and increase usability and adoption.

## INTRODUCTION

The development of CERN's Web-based Rapid Application Platform (WRAP) aims to replace the expensive task of specialised application development, and the time investment required by many people for maintenance [1]. Such a platform needs to encapsulate a great deal of complexity in its implementation, far greater than any individual user application previously developed at CERN. Focusing on WRAP's large component parts individually (including both back-end services and front-end modules), has helped navigate this challenge. This paper describes some of the more difficult challenges faced in each of these parts, together with the implemented solutions, and limitations that remain to be solved.

A summary breakdown of CERN's accelerator control system core concepts [2], in simple terms, will serve as a foundation to understanding the WRAP architecture and the aspects that it needs to address.

## CONTROL SYSTEM OVERVIEW

### *Devices and Properties*

Data produced by a device follows the so-called device-property model. Each device having one or more properties, and every property having one or more fields. For WRAP device interactions, two types of operations are possible and in most cases they are made on the property level:

1. Subscriptions to receive property updates.
2. Sets to load new values into a device property.

\* epameinondas.galatas@cern.ch,  
stephane.deghaye@cern.ch,  
jakub.raban@cern.ch,  
chris.roderick@cern.ch,  
dinika.saxena@cern.ch,  
andreas.solomou@cern.ch

In some cases, operations accept a filter called "selector" that refers to a specific particle accelerator state for which the operation should be performed. A Timing system coordinates the state of the accelerators and devices react accordingly. Properties that accept such filter are marked as "multiplexed". For example, an extraction kicker magnet's strength setting depends on the extraction energy, and the property to control the strength is multiplexed to allow different settings for different beams.

The Controls Configuration Service (CCS) [3] centralizes the configuration of the entire Control System. Its CCDA (Controls Configuration Data Access) API [4] provides, amongst other things, services to retrieve the available devices, their property models and their supporting metadata and available interactions.

### *Subscriptions*

A subscription is established directly to a device, using CMW [5], and references a property. A so-called "first update" provides an initial value when a client subscribes, which is highly valuable for device properties that update infrequently. For multiplexed properties, if a selector is provided, a first update is received for the given selector only, otherwise a value is received for each possible selector.

### *Data Archival*

Device property data can be configured (using the CCS) to have its data logged as times series in the CERN Accelerator Logging Service (NXCAL) [6], which provides numerous facilities to process and extract logged data.

### *Set Operations*

In general, applying settings to devices requires the user to provide values for each individual field in the property concerned, resulting in one atomic "Set" operation. CERN's setting management system "LSA" [7] supports more sophisticated interactions, including "partial sets", which allow to emit some fields from a property setting, by internally using the latest cached values. LSA also provides high-level virtual devices, with virtual properties, that can also be set and subscribed to. Behind the scenes, sets on these high-level device properties will typically modify multiple property fields of multiple underlying devices at the same time. For example, when directly controlling the high-level physics parameters of a particle accelerator, such as the tune, chromaticities, etc.

## PROJECT ARCHITECTURE

The WRAP architecture can be split into four major parts:

1. **Metadata Service:** a back-end service for providing metadata.

Software

User Interfaces & User Experience

# WEB APPLICATION PACKAGING – DEPLOYING WEB APPLICATIONS AS TRADITIONAL DESKTOP APPLICATIONS IN CERN'S CONTROL CENTRE

M. von Hohenbühel\*, S. Deghaye, E. Galatas, E. Matli, E. Roux  
CERN, Geneva, Switzerland

## Abstract

Web applications are becoming increasingly performant and are now capable, in many cases, of replacing traditional desktop applications. There is also a user demand for web-based applications, surely linked to their modern look & feel, their ease of access, and the overall familiarity of the users with web applications due to their pervasive nature. However, when it comes to a Controls environment, the limitations caused by the fact that web applications run inside a web browser are often seen as a major disadvantage when compared to native desktop applications. In addition, applications deployed in CERN's Control Centre (CCC) are tightly integrated with the control system and use a CERN-specific launcher and manager that does not easily integrate with web browsers. This paper presents an analysis of the approaches that have been considered for deploying web applications and integrating them with CERN's control system. The implications on the development process, the IT infrastructure, the deployment methods as well as the performance impact on the resources of the target computers are also discussed.

## INTRODUCTION

Over the last decade, the use of web-applications has become increasingly prominent and CERN's Controls applications have not escaped this trend. Even if nowadays, web applications are still a relative minority in CERN's Control Centre (CCC), it is envisaged that they could represent more than 50% of the applications used in the control rooms within the next 5-10 years. The workstations deployed in CERN's control rooms are typically computers connected to two or three displays as depicted in Fig. 1, and are responsible for running several controls applications concurrently. Most of these applications require access to special data like the dynamic beam configurations orchestrated by the Controls Timing system. As such, the applications are typically controlled through a CERN-made tool, the Common Console Manager (CCM) [1].

The CCM is used to launch and manage all the applications, keeping them in the correct context and providing them with all necessary information. In a nutshell, the CCM allows the operators to run their applications, control their screen positions, and minimise/maximise them with a feature comparable to virtual desktops, depending on which acceler-

ator beam they want to control or observe. The introduction of web-applications caused a few integration issues with the CCM, ranging from inconsistent window titles, dependency on browser support, impact on the computing resources, and more.

One possible solution to this problem is to *transform* the web-applications into desktop applications. This paper reports on the work done to compare the industry standard called 'Electron' with a very new, but very promising framework called 'Tauri', which focuses more on resource optimisation. In addition, studies on the performance of these solutions compared with a normal web browser running the current web-applications are also reported.



Figure 1: Workstations in the CCC.

## PROOF OF CONCEPT PROJECT

A Proof of Concept (PoC) project was launched to select the best tool for web application deployment in the CERN Control Centre (CCC), and to gain a practical understanding of the impact of these solutions on the CCC Technical Console (TC) computers. As such, the PoC had 2 main objectives:

1. Create a working prototype of a representative web application, running as a desktop application (i.e. no browser), using several packaging solutions.
2. Run various benchmarks to compare the performance, and the load on the TCs, with respect to that of a conventional web browser.

\* maximilian.freiherr.von.hohenbuehel@cern.ch,  
stephane.deghaye@cern.ch,  
epameinondas.galatas@cern.ch,  
emanuele.matli@cern.ch,  
eric.roux@cern.ch

# UPGRADE OF THE PROCESS CONTROL SYSTEM FOR THE CRYOGENIC INSTALLATION OF THE CERN LHC ATLAS LIQUID ARGON CALORIMETER

Czeslaw Fluder, Caroline Fabre, Lukasz Goralczyk, Marco Pezzetti,  
Andrzej Zmuda, CERN, Meyrin, Switzerland  
Krzysztof Mastyna, AGH University, Cracow, Poland

## Abstract

The ATLAS (LHC detector) Liquid Argon Calorimeter is classified as a critical cryogenic system due to its requirement for uninterrupted operation. The system has been in continuous nominal operation since the start-up of the LHC, operating with very high reliability and availability. Over this period, control system maintenance was focused on the most critical hardware and software interventions, without direct impact on the process control system. Consequently, after several years of steady state operation, the process control system became obsolete (reached End of Life), requiring complex support and without the possibility of further improvements. This led to a detailed review towards a complete upgrade of the PLC hardware and process control software. To ensure uninterrupted operation, longer equipment lifecycle, and further system maintainability, the latest technology was chosen.

This paper presents the methodology used for the process control system upgrade during development and testing phases, as well as the experience gained during deployment. It details the architecture of the new system based on a redundant (Hot Standby) PLC solution, the quality assurance protocol used during the hardware validation and software testing phases, and the deployment procedure.

## INTRODUCTION

The ATLAS Liquid Argon (LAr) Calorimeter measures the energy and timing of photons, electrons and hadrons produced by proton-proton collisions in the Large Hadron Collider (LHC) [1]. The particles from the collisions are absorbed on particular metal layers, resulting in a shower of new, low-energy particles. The newly formed particles ionize liquid argon between the layers, producing an electric current that is measured by high-voltage electrodes.

The LAr Calorimeter (550 t of cold mass) is housed in three individual cryostats (one barrel and two end-caps) which are filled with a total volume of 80 m<sup>3</sup> of liquid argon at a temperature of 88.4 K. It has been kept in continuous nominal conditions for the past 17 years thanks to a dedicated cryogenic system (the LAr cryogenic installation). Liquid nitrogen, used as cooling medium, is provided to the cryostats (the clients) by a set of back-up supply processes (the infrastructure). In addition, 114 feed-throughs, through which a total of 228,000 wired detector signals are routed to the outside of the cryostats, are kept under vacuum and at a controlled temperature to protect pin-carriers against




Figure 1: Typical DCS architecture.

humidity damage.

## CONTROL SYSTEM OVERVIEW

The LAr cryogenic installation is designed to ensure uninterrupted functioning and safe handling of large quantities of liquid argon in underground areas. Its dedicated LAr control system is classified as critical due to these requirements. In addition to controlling nominal cryogenic conditions and protecting the feed-throughs it also monitors critical insulation vacuum levels. In case of degradation of such vacuum, indicating a potential liquid argon leakage and oxygen deficiency hazard, the control system triggers alarms for personal safety (cavern evacuation, air ventilation system, fire brigade) and for material safety (Detector Safety System interlocks).

### Legacy Architecture

To automate the process control of the LAr cryogenic installation, a typical Distributed Control System (DCS) has been implemented (see Fig. 1).

The visualization and operation of the cryogenic process are ensured through Supervisory Control And Data Acquisition (SCADA) based on the WinCC OA<sup>®</sup> application. One SCADA data server is connected to six autonomous Programmable Logic Controllers (PLCs), which use Modicon M580<sup>®</sup> and Quantum<sup>®</sup> technology. Two recently upgraded M580<sup>®</sup> PLCs are utilized to control a nitrogen refrigerator (cryoplant). Additionally, four old Quantum<sup>®</sup> PLCs are

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# LIFE CYCLE MANAGEMENT AND RELIABILITY ANALYSIS OF CONTROLS HARDWARE USING OPERATIONAL DATA FROM EAM

E. Fortescue<sup>†</sup>, I. Kozsar, CERN, Geneva, Switzerland

V. Schramm, Institute of Machine Components, University of Stuttgart, Germany

## Abstract

The use of operational data from Enterprise Asset Management (EAM) systems has become an increasingly popular approach for conducting reliability analysis of industrial equipment. This paper presents a case study of how EAM data was used to analyse the reliability of CERN's standard controls hardware, deployed and maintained by the Controls Electronics and Mechatronics (CEM) group. The first part of the study involved the extraction, treatment and analysis of state-transition data to detect failures. The analysis was conducted using statistical methods, including failure-rate analysis and time to failure analysis to identify trends in equipment performance and plan for future obsolescence, upgrades and replacement strategies. The results of the analysis are available via a dynamic online dashboard. The second part of the study considers Front-End computers as repairable systems, composed of the previously studied non-repairable modules. The faults were recorded and analysed using the Accelerator Fault Tracking (AFT) system. The study brought to light the need for high quality data, which led to improvements in the data recording process and refinement of the infrastructure team's workflow. In the future, reliability analysis will become even more critical for ensuring the cost-effective and efficient operation of controls systems for accelerators. This study demonstrates the potential of EAM operational data to provide valuable insights into equipment reliability and inform decision-making for repairable and non-repairable systems.

## INTRODUCTION

CEM's Infrastructure (IN) section is responsible for the life-cycle management of a very large collection of hardware components for the Front-End Computers (FECs) distributed across the entire CERN accelerator complex. Planning for replacement and upgrades has become a real challenge. Therefore, having readily available reliability metrics such as the Mean Time To Failure (MTTF), or statistics to determine whether the failure rate is increasing for a given equipment type, known as a part in asset management terminology, can significantly aid strategic decision-making.

Since 2005, instigated by the need for industrial-scale asset management during the installation of the controls hardware in the LHC, the CEM-IN section has invested heavily in using CERN's Enterprise Asset Management (EAM) system [1], to follow operational equipment through its' life cycle. The EAM system provides many

functionalities including history tracking, user-defined state machines, spare part and store management, and user reporting.

Assets are registered in the EAM system for instances of any part that need to be individually tracked, such as electronic boards, chassis, power supplies, fans, etc. The corresponding barcodes, or QR codes, are attached to the physical equipment upon reception at CERN. All movements and events affecting the asset are recorded in the EAM system, such as installation in a new position, attachment to, or detachment from a parent asset, return or issue from a store or changes in asset status e.g. Installed, In Repair, Out of Service etc.

## USE OF STATE-TRANSITION DATA

The Infrastructure section registered and traced an increasing number of electronic assets in the EAM system over a period of more than 15 years. CEM-IN currently manage approximately 50,000 operational assets and 300 distinct parts, and in 2021, it was considered that sufficient historical data was present in EAM to be able to attempt a rudimentary reliability analysis.

The first initiative was to extract, process and analyse the state-transition data to detect failures in the electronic modules and generate reliability statistics per part.

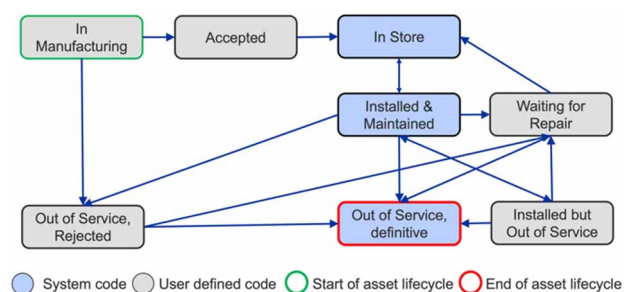


Figure 1: Custom state machine for CEM-IN assets.

Figure 1 shows the custom state machine applied by the EAM system to the CEM-IN assets between 2005 and 2021. Whenever an asset changes from one state to another, the state-transition is recorded and timestamped in the asset's history. A failure, as defined in this context, was deemed to have occurred on the date that an asset's state passed directly from "Installed and Maintained" to "Waiting for Repair". By querying this basic state-transition data, it is straightforward to determine how many failures occur per calendar year for a given type of electronic module.

<sup>†</sup> Eve.Fortescue@cern.ch

# CERN PROTON IRRADIATION FACILITY (IRRAD) DATA MANAGEMENT, CONTROL AND MONITORING SYSTEM INFRASTRUCTURE FOR POST-LS2 EXPERIMENTS\*

B. Gkotse<sup>†1</sup>, G. Pezzullo, F. Ravotti,  
Experimental Physics Department, CERN, Geneva, Switzerland  
P. Jouvelot, Mines Paris, PSL University, Paris, France  
<sup>1</sup>also at Mines Paris, PSL University, Paris, France

## Abstract

Since upgrades of the CERN Large Hadron Collider are planned and design studies for a post-LHC particle accelerator are ongoing, it is key to ensure that the detectors and electronic components used in the CERN experiments and accelerators can withstand the high amount of radiation produced during particle collisions. To comply with this requirement, scientists perform radiation testing experiments, which consist in exposing these components to high levels of particle radiation to simulate the real operational conditions. The CERN Proton Irradiation Facility (IRRAD) is a well-established reference facility for conducting such experiments. Over the years, the IRRAD facility has developed a dedicated software infrastructure to support the control and monitoring systems used to manage these experiments, as well as to handle other important aspects such as dosimetry, spectrometry, and material traceability. In this paper, new developments and upgrades to the IRRAD software infrastructure are presented. These advances are crucial to ensure that the facility remains up-to-date and able to cope with the increasing (and always more complex) user needs. These software upgrades (some of them carried out within the EU-funded project AIDAInnova and EURO-LABS) will help to improve the efficiency and accuracy of the experiments performed at IRRAD and enhance the capabilities of this facility.

## INTRODUCTION

The CERN Large Hadron Collider (LHC) upgrades and the post-LHC particle-accelerator design studies make necessary the qualification of various components against radiation. The CERN Proton Irradiation Facility, IRRAD (see Fig. 1), is a reference facility in High-Energy Physics (HEP) community, dedicated to radiation tests of components, materials and irradiation experiments on complex detector or accelerator systems. Detectors, materials and electronic components of different dimensions (from a few mm<sup>2</sup> to several cm<sup>2</sup>) are being exposed to the proton beam in order to assess their radiation hardness. The proton beam

is delivered from the Proton Synchrotron (PS) accelerator at CERN with a momentum of 24 GeV/c, in spills of 400 ms every 10 s on average, and it has a typical Gaussian shape of 12×12 mm<sup>2</sup> full width at half maximum (FWHM). It impinges on the devices under test positioned along the beam line. The Devices Under Test (DUT), are usually positioned on stages (IRRAD Tables) that can be remotely controlled and moved horizontally, vertically or rotated by a certain angle w.r.t. the beam axis, while providing also in-beam scanning capabilities. The samples can be installed or withdrawn from the irradiation zones only once a week. Therefore, for smaller samples (e.g., 5×5 mm<sup>2</sup>) and shorter irradiation experiments, a 9-m long shuttle system is used for moving the samples in and out of the beam line without stopping the beam operation. The IRRAD facility is also equipped with two cold boxes for irradiation experiments down to -25 °C, but also with a LHe cryostat for experiments performed down to 4.2 K. Beam Profile Monitors (BPM) are also installed in IRRAD for the real-time monitoring of the beam profile and quality [1] but also for aligning precisely the IRRAD Tables in the beam. Since 2022, IRRAD is also one of the research infrastructures that provide Transnational Access to users coming from other countries through the European Project EURO-LABS [2].




Figure 1: IRRAD Facility.

For the successful operation of the IRRAD facility, a robust hardware and software infrastructure had to be established. This includes data management systems, beam and environmental monitoring/logging tools, and control systems. In order to cope with the continuous demand of experiments, improve and facilitate the operation as well as be compliant with the CERN IT infrastructure updates and security rules, several upgrades have been performed and detailed in the following sections.

\* THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GA NO. 101004761.

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON EUROPE RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NO 101057511.

<sup>†</sup> Blerina.Gkotse@cern.ch

# EPICS NTTABLES FOR MACHINE TIMING CONFIGURATION

A. Gorzawski\*, J. P. S. Martins, N. Milas,  
 ESS, European Spallation Source, Lund, Sweden

## Abstract

The European Spallation Source (ESS), currently under construction and initial commissioning in Lund, Sweden, will be the brightest spallation neutron source in the world, when its driving proton linac achieves the design power of 5 MW at 2 GeV. Such a high power requires production, efficient acceleration, and almost no-loss transport of a high current beam, thus making the design and beam commissioning of this machine challenging. The recent commissioning runs (2021-2023) showed a need for a consistent and robust way of setting up the machine for beam production. One of the big challenges at ESS is joining the machine setup and the timing setup limiting the need for operator actions. In this paper, we show a concept of using EPICS 7 NTTables to enable this machine settings consistency. Related to that, we also highlight a few challenges related to other EPICS tools like *Save And Restore* and *EPICS Archiver*.

## INTRODUCTION

ESS is a collaboration of 17 European nations and its objective is to be the world's most powerful spallation neutron source [1]. The neutrons are produced by a pulsed proton beam hitting a solid, rotating tungsten target at a distance of 600 m from the ion source. The ESS linac, the driver of the protons onto the target, requires site-wide synchronization in order to accelerate the desired beam through its components. The production of the proton beam begins with the ion source providing the pulse of protons with an optimized current and of a given length [2]. Later, the two choppers (LEBT and MEBT) shape the pulse to the desired pulse length. The acceleration is given by the radio-frequency cavities pulsing at the correct moment. The overall repetition rate is a consequence of the definition of the arbitrary number of consecutive sets of timing events (cycles), that are pre-loaded before the ion source starts producing the beam. By design the actual maximum repetition rate is 14 Hz.

The ESS distributed control system is based on Experimental Physics and Industrial Control System (EPICS [3]). The full synchronization is provided by a distributed timing system with its own network infrastructure, and it is operated within a *Beam Production* environment. The entire timing system is configured with the pre-created timing tables, that consist of a definition of super-cycles, namely sequenced definitions of what the RF, actuators, and instrumentation do in the given cycle.

This paper summarises the efforts put in the beam commissioning of the Normal Conducting Linac throughout the multiple commissioning periods, i.e. from the ion source to the DTL4 FC with a special look at the machine tim-




Figure 1: The successful beam pulse signatures read by available beam instrumentation in the fully available NCL part of ESS linac.

ing setup along with the development and implementation of the new approach to handling timing tables using the `epics:NTTTable` available in EPICS 7.

## BEAM COMMISSIONING

Recently four designated periods of hardware and beam commissioning were successfully completed [4]. In each of the commissioning periods, we tested, verified, and established the ways the timing setup is used with an increasing number of commissioned equipment. The following top-level milestones were achieved:

1. *October 2021 - December 2021*, the MEBT FC commissioning with a small current beam, the first time with the timing system available,
2. *February 2022 - March 2022*, MEBT FC, commissioning with high current beam,
3. *May 2022 - July 2022*, DTL1 FC commissioning, low and high current beam,
4. *April 2023 - July 2023*, DTL4 FC commissioning, low and high current beam,

Figure 1 illustrates the Beam Current Monitor (BCM) signals from different locations at the ESS linac. One can see the shortening of the pulse in the more downstream locations (MEBT and DTL4). This is due to the acting MEBT chopper for the head of the beam. The responsibility for the success shown in that plot lies in a series of functioning timing-related events (see next section).

## TIMING SYSTEM FOR ACCELERATOR

While the hardware details of the ESS timing system are described in [5] and [6], the top-level timing parameters are

\* arek.gorzawski@ess.eu



# DESIGN OF THE CONTROL SYSTEM FOR THE CERN PSB RF SYSTEM

D. Landré, Y. Brischetto, M. Haase, M. Niccolini, CERN, Geneva, Switzerland

## Abstract

The RF system of the CERN PS Booster (PSB) has been renovated to allow the extraction energy increase and the higher beam intensities required by the LHC Injectors Upgrade (LIU) project. It relies on accelerating cells installed in three straight sections of each of the four accelerating rings of PSB. Each cell is powered by one solid-state RF amplifier. This modularity is also embedded in its controls architecture, which is based on PLCs (Programmable Logic Controllers), several FESA (Front-End Software Architecture) classes, and specialized graphical user interfaces for both operation and expert use. The control system was commissioned during the Long Shutdown 2 (LS2) and allows for the nominal operation of the machine. This paper describes the design and implementation of the control system, as well as the system performance and achieved results.

## INTRODUCTION

The PSB RF (Radio-Frequency) system was entirely replaced in the framework of the LIU project, to ensure the performance required for the High-Luminosity Large Hadron Collider (HL-LHC) project. In conjunction with the hardware renovation, a state-of-the-art control system was developed and commissioned to ensure the system's optimal operation.

## OVERVIEW OF THE SYSTEM

The PSB machine consists of four superimposed rings, and the acceleration of particles is ensured by the new wide-band Finemet RF system [1, 2], which is installed in three straight sections of the machine. In each sector two groups of Finemet cells are installed per ring, resulting in structures composed of forty-eight cells per sector. Figure 1 shows one of such structures before installation in one of the sectors. The new RF system thus consists of a total of one hundred and forty-four cells, each with independent power supplies, amplifiers and control electronics. Air and water cooling distribution are shared by the cells in each sector.

## HARDWARE CONTROL

In order to ensure the safety of the system, which consists of accelerating cells, amplifiers, power supplies and auxiliary services such as cooling, machine vacuum and fieldbus communication, an advanced industrial control and interlocks system was developed. Each structure installed in a sector is controlled independently through one PLC and measurement and control cards (AI – analog input, AO – analog output, DI – digital input, DO – digital output) tailored to the equipment installed. The control program is structured to follow a sequence of data acquisitions, each authorizing




Figure 1: One of the three structures of the new PSB RF system before installation (source: CERN).

the start of the next step and ensuring that all necessary conditions are met before the RF can be pulsed. The control sequence is composed of sub-sequences, which span from overseeing general services to granting authorization for RF pulsing on a cell.

## Control Sequence

**Level 1** The initial part of the first level of the sequence runs checks on the fieldbus communication, water temperature and pressure, water leaks and machine vacuum. In the event of any equipment failure within this segment, an immediate shutdown of the RF system is triggered in the affected sector to ensure safety and prevent potential hazards. In the latter part of the sequence, the water cooling of the amplifiers is started independently for each ring, such that any fault happening in this stage only affects the specific ring and does not influence the operation of the others. Moreover, the program runs ventilation system checks. Any malfunctions exclusively impact the cells connected to that specific ventilation, ensuring that the neighboring cells within the sector remain unaffected.

**Level 2** At this stage the designated power is distributed to the cells. Each cell is controlled independently, allowing for power redistribution among cells in the event of a malfunction as long as the minimum required power is attained. The power redistribution is made possible thanks to an automated procedure which is presented later in this document. The first step of this stage is the check of the position of the sector safety key. If this key is either removed or placed in the safety position, no power is sent to any of the cells within the relevant sector, thereby allowing safety conditions for

# VERY EARLY FIRE DETECTION FOR HIGH POWER PULSED SYSTEMS

S. Pavis, N. Magnin, E. Carlier, C. Lolliot, CERN, Geneva, Switzerland

## Abstract

Equipment racks containing high current power supplies and solid-state electronic switching circuits associated with accelerator kicker installations at CERN can fail in ways that risk fire outbreak, and building and tunnel fire detection systems may not be well placed to detect fire outbreaks in such racks until the fire is well established. The risks associated with late fire detection can be worse in normally unmanned surface buildings and normally non-accessible underground areas and accelerator service tunnels. Very early fire detection directly in the racks is, therefore, highly desirable to give local power interlock in the event of smouldering, with interlock levels being configurable as determined appropriate for each individual environment. This paper describes the specific risk situations and circumstances, and the detection technologies considered. The final choice of detection and interlocking strategy is demonstrated to be successful in detecting the very early incipient stages of a fire and drastically reducing the risks of covert fire development leading to major fire outbreak with all its associated consequences. Several of these early fire detection systems have already been installed in LHC and SPS accelerator kicker installations, with many more planned.

## INTRODUCTION

### Beam Injection and Extraction Kicker Systems

CERN's accelerator beam transfer kicker systems inject particles into a circular accelerator, or extract particles from an accelerator to a transfer line or beam dump. A Kicker magnet is a pulsed dipole magnet which produces a rectangular field pulse with very fast rise and/or fall time (typically 50 ns to 1 us). The magnet pulse may be several kA. To achieve this, for each kicker magnet a kicker system uses high current power supplies known as the DCPS which charge capacitor banks. This charge is then discharged via a thyristor to a step up transformer, the secondary of which charges the Pulse Forming Line (PFL) or Pulse Forming Network (PFN) to the required kick strength. The combination of these stages is called the Resonant Charging Power Supply (RCPS), and the resultant charge is then rapidly discharged into the magnet via fast, high power switches (thyratrons or solid-state electronic switches) called the Main Switch and the Dump Switch which act together to produce a clean rectangular pulse [1] (see Fig. 1).

### Kicker Racks and Fire Risk

Enclosed cabinets can be considered micro-climates, being somewhat isolated from their surrounding environment to protect personnel from hazardous voltages and charges, populated with power electronics and busbars, in some cases electrically noisy, containing a mixture of signal and high




Figure 1: Simplified schematic of a kicker system.

voltage and/or high current cables, some with locking/interlocking door mechanisms, sometimes featuring forced air circulation. The high energy power supplies and associated solid state switching electronics and interconnecting cables, under certain fault conditions such as a degrading power component, a loose connector, material fatigue of a connector or cable, damage to a connector or cable or disconnection of a cable resulting in a short circuit, arcs and abnormal heating effects can occur, potentially causing smoldering and even igniting of nearby flammable material. In such a contained environment, any breakout of fire, especially at the early smouldering stages, may remain covert for some time. Building/Tunnel smoke and fire detection, due to the lack of propagation of smoke from the cabinet, is often not capable of detecting such fires until they are well established, meaning significant damage can already occur before any alarm is raised or extinguishing action is carried out. If the zone in which kicker racks are installed is underground, this may also be inaccessible during accelerator operation, further complicating the detection and identification of fire situations. Such a situation occurred in the AD Horn installation (see Fig. 2). The probable cause of this fire was a loose trigger cable for the ignitron such that the power capacitor energy was shorted to ground and the resulting electric arc ignited nearby cables and oil hoses. The loose cable was caused by material fatigue or damage of the connector. While the occurrences of such fires have so far proven to be rare, the consequence can be significant damage to equipment resulting in long machine shutdown for repair, and therefore having impact on CERN's scientific programs. To mitigate the in-rack fire risks and non-optimal conditions for the main building/tunnel fire systems to be able to detect the early stages of such fires, a more localised and rapid detection strategy is needed which can detect the incipient stages of smouldering and cut power to the rack in order to prevent an actual fire. In order to interface to and cut power to the kicker system and considering the accessibility issues during accelerator operation, the chosen technology must feature good diagnostic and alarm reporting capability via standard industrial 4-20 mA and dry contact type

# AUTOMATIC CONDITIONING OF KICKER MAGNETS

C. Lolliot\*, M. J. Barnes, N. Magnin, S. Pavis, CERN, Geneva, Switzerland  
C. Monier, INSA, Lyon, France

## Abstract

Fast pulsed, in vacuum, kicker magnets are used across the accelerators of the CERN complex to inject and extract beam. The kicker magnets are powered by high voltage (HV) pulse generators. The kicker magnets may suffer from electrical breakdown during the pulse. Hence, to prepare them for reliable operation at full voltage, a pulsed HV conditioning is necessary. The magnet is conditioned by first gradually increasing pulse voltage, with a short pulse length, and subsequently increasing the pulse length up to values beyond operational conditions. Before LS2, this conditioning was typically carried out manually, which was a workforce intensive procedure because a magnet conditioning can last for up to several weeks. To overcome this drawback, a standardised industrial controller algorithm has been developed to control the pulse generators. The voltage and pulse length ramps are fully adjustable, and now different modes are available - in particular a voltage logarithmic mode. If a breakdown occurs during the conditioning, the controller will automatically reduce the voltage by a specified percentage and then continue the conditioning procedure. Furthermore, a simulation mode has been developed to allow a quick visualisation of the whole conditioning as well as the simulation of HV breakdowns. This functionality has been implemented in several kicker installations across the various accelerators, as well as in test cages in the lab, and it will be implemented in most installations in the future.

## INTRODUCTION

Kicker magnets are used throughout the CERN accelerator chain to inject or extract beam from a ring. These kicker magnets generally have fast rise and/or fall times and are driven by high voltage pulses of up to 40 kV. To achieve fast field rise and/or fall times, transmission line kicker magnets are typically used [1]: these magnets are installed in vacuum. The kicker magnets are terminated either in a short circuit or with a resistor whose value is equal to the characteristic impedance of the system. A pulse generator is used to provide the energy required for powering the kicker magnet: this is either a Pulse Forming Line (PFL) or a Pulse Forming Network (PFN) [1].

As part of the process of preparing the kicker magnet for reliable operation in the accelerators, the magnets must be conditioned under representative conditions, i.e. with High Voltage (HV) pulses.

## HV PULSE CONDITIONING

### *Considerations for HV Pulse Conditioning*

Kicker magnet conditioning is a process to prepare the magnet for reliable operation in the accelerators. Dust, contaminants or small features increase locally the electric field and thus can cause HV breakdown or corona; see Fig. 1. To condition the surfaces and hence reduce locally high electrical field, a voltage is applied: the value is incremented in small steps up to a specified value, slightly above the operating voltage. Failure to increment the voltage slowly can result in HV breakdown in the magnet, which can in turn damage the kicker magnet. The phase of increasing pulse voltage is normally carried out, where possible, with relatively short pulse lengths: this limits the energy that is dissipated in the site of the breakdown and hence minimizes the possibility of damaging the surfaces.

Failure to properly condition the magnets can result in strong HV breakdown and may damage the surface of components. The HV breakdown also degrades the vacuum. An important risk of magnetic breakdown during operation is a transient increase or decrease in the magnetic field, which can cause the injected/extracted beam to be mis-kicked.




Figure 1: HV corona inside a kicker magnet.

During the conditioning the kicker magnet is subject to gradually higher and higher pulse voltages (ramping mode), and subsequently greater pulse lengths (elongation mode), until the voltage holding capability is sufficiently beyond the nominal operating conditions. Historically, the “pulsed conditioning” was carried out manually: this is a workforce-intensive and also technically non-optimal procedure. During a manual conditioning process, the operator has to survey the vacuum pressure, and monitor and record the conditioning parameters. In addition, in the case of HV breakdown of the magnet, causing a significant pressure rise, the system may interlock, preventing pulse conditioning until a human operator is available to intervene and restart the process.

The HV conditioning by a human operator requires constant attention: voltage and pulse length remain unchanged for long periods, e.g. evening, night and weekends or if the operator has to deal with something else – by contrast a

\* christophe.lolliot@cern.ch

# SPARK ACTIVITY MONITORING FOR LHC BEAM DUMP SYSTEM

C. B. Durmus, E. Carlier, N. Magnin, T. D. Mottram, V. Senaj  
CERN, Geneva, Switzerland

## Abstract

LHC Beam Dump System is composed of 25 fast-pulsed magnets per beam to extract and dilute the beam onto an external absorber block. Each magnet is powered by a high-voltage generator that discharges the energy stored in capacitors into the magnet by using high-voltage switches. These switches are housed in air in cabinets which are not dust-protected. In the past years of LHC operation, we noticed electrical sparks on the high-voltage switches due to the release of accumulated charges on the surfaces of the insulators and the switches. These sparks can potentially cause a self-trigger of the generators yielding to asynchronous dumps which should be avoided as much as possible. In order to detect dangerous spark activity in the generators before a self-trigger occurs, a Spark Activity Monitoring (SAM) system was developed. SAM consists of 50 detection and acquisition systems deployed at the level of each high-voltage generator and one external global surveillance process. The detection and acquisition systems are based on digitisers to detect and capture spark waveforms coming from current pick-ups placed in various electrical paths inside each generator. The global surveillance process collects data from all the acquisition systems in order to assess the risk of self-trigger based on the detected sparks' amplitude and rate. This paper describes the architecture, implementation, optimisation, deployment and operational experience of the SAM system.

## INTRODUCTION

The LHC Beam Dump System (LBDS) plays a critical role in LHC machine protection, by ensuring the safe extraction of the beam from the LHC. The beam is directed to the extraction channel through the beam-synchronised activation of 15 extraction kicker magnets (MKD). Subsequently, the beam is diluted by 4 horizontal (MKBH) and 6 vertical dilution magnets (MKBV) before reaching the beam dump absorber (TDE). To accommodate the rising edge of the MKD magnetic field, a particle-free Beam Abort Gap (BAG) of approximately  $3 \mu\text{s}$  is meticulously maintained within the LHC [1].

### High-Voltage Pulse Generators

High-voltage pulse generators (HVPG) are used to generate the current pulse in the MKD and MKBs. The HVPGs consist of capacitors and fast high-voltage switches to discharge the energy stored in the capacitor into the magnet (see Fig. 1). The capacitor charge depends on the LHC beam energy. The conduction of high-voltage switches is started by power trigger modules (PTM) [2].

These HVPGs can self-trigger for different reasons which can result in the MKD rising edge occurring outside of the




Figure 1: High-Voltage Switch used in HVPG.

BAG, and the beam circulating in LHC might be spread over the machine aperture, which could yield damage to various machine elements. [3].

### Re-Trigger System

To limit the impact of a self-trigger of an MKD HVPG, a re-triggering system is implemented to trigger as fast as possible all other MKD HVPG. To this extent, re-trigger boxes (RTB) are installed on each generator, interconnected using re-trigger lines (RTL). The re-trigger boxes are connected on one side to sensors inside the HVPG to generate a pulse on the re-trigger lines in case of HV switch conduction, and on the other side to the PTM to trigger conduction of HV switch in case pulses are present on the re-trigger line. [4] These sensors are installed on different power circuits of the HVPG to detect small current flows and capacitor voltage changes.

### Spark Activity in HVPG

In the past years of LHC operation, HVPG self-trigger events occurred. The current waveform captured from the HVPG indicated a sign of high-voltage sparks a few  $\mu\text{s}$  before the start of the conduction of the high-voltage switch. Before starting of LHC operation, we manually checked for the high-voltage spark activity by using oscilloscopes on-site. This has been a time-consuming operation and it was impossible to monitor all generators continuously. To detect dangerous high-voltage spark activity before a self-trigger occurs, we decided to develop a continuous spark activity monitoring system.

## SAM ANALYSERS

### Spark Waveform Acquisition

SAM waveform acquisition and analysis is implemented using the Internal Post Operation Check (IPOC) framework

# A MODULAR APPROACH FOR ACCELERATOR CONTROLS COMPONENTS DEPLOYMENT FOR HIGH POWER PULSED SYSTEMS

R. A Barlow\*, C. Boucly, E. Carlier, C. Lolliot, C. Chanavat, N. Magnin, S. Pavis, P. Van Trappen, N. Voumard, CERN, Geneva, Switzerland

## Abstract

In the context of the LHC Injector Upgrade (LIU) project, significant advancements have been made in enhancing the control systems of the PSB and PS (injection) kickers at CERN during the Long Shutdown 2 (LS2). This upgrade has transitioned these systems from heterogeneous, custom-built electronic solutions to a more flexible and open architecture. Despite the distinct functionalities, power circuit topology and operational requirements of these kickers, a uniform approach was adopted by employing standardized hardware and software control blocks for both accelerator facilities.

The newly implemented control architecture is structured around a series of subsystems, each tailored to perform specific, essential functions necessary for managing high power fast pulsed systems. These functions encompass aspects such as equipment and personnel safety, slow control and protection, high-precision fast timing systems, fast interlocking and protection mechanisms, as well as pulsed signal acquisition and diagnostic analysis. Each of these subsystems involves the integrated utilization of both hardware components and associated software.

This paper aims to present a comprehensive overview of the functionalities inherent to the various subsystems. Additionally, it offers insights into how these subsystems have been effectively integrated to cater to the distinct requirements of the two different use cases.

## FINITE STATE MACHINE

In recent years, significant advancements have been achieved in the refinement and enhancement of kicker control systems, driven by a growing emphasis on modularity and the optimization of the creative process and its associated hardware components. Central to this progress is the adoption of a finite state machine approach, which definitively encapsulates each modular aspect, as depicted in Fig. 1. This approach adheres to a well-defined sequence, commencing with the activation of the Power Distributor Controller (PDC) system in the initial stage, responsible for supplying power to the entire installation. Subsequently, industrial distributed PLC systems combined with a generic finite state machine approach involves the careful management of kicker thyratrons, inclusive of their respective heater controllers and heater power supplies. These general PLC library components ensure also the efficient management of capacitor bank switches, a key element in the majority of kicker systems within Accelerator and Beam Transfer (ABT) kicker systems facility. Typically, these systems incorporate an initial energy storage element, such as a capacitor bank, to effectively

\* roger.andrew.barlow@cern.ch

store energy from the Direct Current Power Supply (DCPS) and subsequently deliver it by resonant charging to a Pulse Forming Network (PFN), a task tackled by the Capacitor Charger and Protection Unit (CDPU). Upon successful validation of these initial stages within the finite state machine, the DCPS can be activated. The ultimate functionality of the entire state machine hinges on the precise timing system, critical for coordinating the delivery of pulses to the magnet, a process typically governed by the Kicker Timing System (KiTs). Additionally, for post-operational verification, there exists a fast acquisition system known as Internal Post Operation Check (IPOC), enabling pulse-to-pulse acquisition of magnet current and other relevant signals.

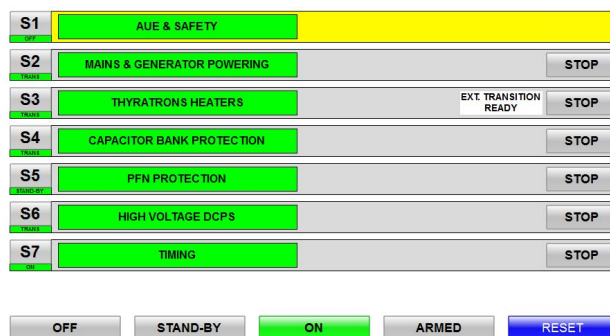


Figure 1: Typical finite state machine view for supervisory control of kicker systems.

Typically the following stages exist:

*Stage 1 (AUE and Safety):* This initial stage primarily encompasses various low-level security ancillaries and peripheral devices, including emergency stops, thermal switches, and other safety-related components.

*Stage 2 (Mains and Generator Powering):* This stage is responsible for delivering power to the entire system, both from the mains and generators. It's an automated process that controls most of the low-voltage distribution to critical subsystems.

*Stage 3 (Thyratrons Heaters):* This stage handles the power supply for thyatron heaters. It is specifically designed to manage the heating process and reservoir pressure settings. See subsection *Thyratron Heater Power Supplies*.

*Stage 4 (Capacitor Bank Protection):* In this stage of the state machine, the focus is on managing the earthing scheme and ensuring the security of activating and deactivating mains earthing points within the system.

*Stage 5 (PFN Protection):* Similar to the previous stage, the PFN system requires careful control to ensure human safety and equipment protection. This stage also manages the relevant components of the PFN.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# LEVERAGING LOCAL INTELLIGENCE TO CERN INDUSTRIAL CONTROL SYSTEMS THROUGH EDGE TECHNOLOGIES

A. Patil, F. Varela, F. Ghawash, B. Schofield, CERN, Geneva, Switzerland,  
T. Kaufmann, A. Sundermann, D. Schall, Siemens AT - T DAI DAS, Austria,  
C. Kern, Siemens DE - T CED SES, Germany

## Abstract

Industrial processes often use advanced control algorithms such as Model Predictive Control (MPC) and Machine Learning (ML) to improve performance and efficiency. However, deploying these algorithms can be challenging, particularly when they require significant computational resources and involve complex communication protocols between different control system components. To address these challenges, we showcase an approach leveraging industrial edge technologies to deploy such algorithms. An edge device is a compact and powerful computing device placed at the network's edge, close to the process control. It executes the algorithms without extensive communication with other control system components, thus reducing latency and load on the central control system. We also employ an analytics function platform to manage the life cycle of the algorithms, including modifications and replacements, without disrupting the industrial process.

Furthermore, we demonstrate a use case where an MPC algorithm is run on an edge device to control a Heating, Ventilation, and Air Conditioning (HVAC) system. An edge device running the algorithm can analyze data from temperature sensors, perform complex calculations, and adjust the operation of the HVAC system accordingly. In summary, our approach of utilizing edge technologies enables us to overcome the limitations of traditional approaches to deploying advanced control algorithms in industrial settings, providing more intelligent and efficient control of industrial processes.

## INTRODUCTION

The latest advances in AI and ML, along with time-tested methods like MPC, offer new ways to enhance the functionality of industrial control systems [1]. For instance, these techniques can improve system reliability through anomaly detection, enable energy-efficient operation of complex industrial processes, and extend equipment life through predictive maintenance. Nevertheless, enhancing industrial control systems through such techniques poses several challenges.

One significant challenge is ensuring that the core processes of the system and the algorithms operate independently and do not interfere with one another. This process independence ensures that the demands of complex algorithms do not jeopardize the safe operation of the core process and overburden its resources. Also, deploying complex algorithms on the existing control infrastructure may only be possible if specialized hardware components like GPUs or AI processors are available. These components were relatively uncommon in industrial control setups until recently.

However, new control hardware, such as multi-processor PLCs and AI expansion cards, have emerged, making this deployment possible.

Another challenge is the notable disparities between the focus areas of control engineers and data scientists when devising control systems. Control engineers primarily concentrate on industrial communication protocols, control devices, PLC programming, and SCADA development. In contrast, data scientists and software engineers focus on creating new control strategies using Python or C++ and utilize software development tools like package managers and containers. New computing paradigms tailored to industrial control systems have been developed that bridge this divide and integrate information technology (IT) tools into operational technology (OT). Examples include integrating control systems with Cloud computing, High-Performance Computing (HPC), and Edge computing.

This article mainly focuses on solutions that address these challenges and provide local intelligence to a control system, i.e., intelligence close to the process, allowing faster analysis of streamed data and lightening the load on the different layers of the control system by reducing network latency and traffic. We start by comparing various techniques for leveraging local intelligence. We emphasize Industrial Edge Computing as an emerging solution that provides benefits such as separation of concern, simplification of algorithm development, and easy application lifecycle management. Finally, we will share insights from implementing an advanced optimization algorithm on state-of-the-art edge technologies and validating its use in a real-world setting at CERN.

## LEVERAGING INTELLIGENCE TO INDUSTRIAL CONTROL SYSTEMS

In recent years, the capability to deliver intelligence close to industrial processes has progressed from conventional setups such as bare-metal Industrial PCs to more advanced systems like multi-process controllers and edge technologies. Some of these setups are outlined below.

- **Multi-process controllers:** PLC vendors have acknowledged developers' needs to program control components in languages beyond the IEC 61131-3 standard [2], adopting higher-level languages like C++ and Python. For instance, the Siemens S7-1518 Multi-Functional Platform (MFP) has a Linux OS alongside its standard PLC OS that primarily supports C++. Communication between the OSs is via an Ethernet virtual switch, eliminating additional hardware and separating

# INTERLOCK SUPER AGENT: ENHANCING MACHINE EFFICIENCY AND PERFORMANCE AT CERN'S SUPER PROTON SYNCHROTRON

E. Veyrunes, J. Wenninger, G. Trad, A. Asko  
CERN, Geneva, Switzerland

## Abstract

In the CERN Super Proton Synchrotron (SPS), finding the source of an interlock signal has become increasingly unmanageable due to the complex interdependencies between the agents in both the beam interlock system (BIS) and the software interlock system (SIS). This often leads to delays, with the inefficiency in diagnosing beam stops impacting the overall performance of the accelerator. The Interlock Super Agent (ISA) was introduced to address this challenge. It traces the interlocks responsible for beam stops, regardless of whether they originated in BIS or SIS. By providing a better understanding of interdependencies, ISA significantly improves machine efficiency by reducing time for diagnosis and by documenting such events through platforms such as the Accelerator Fault Tracking system. The paper will discuss the practical implementation of ISA and its potential application throughout the CERN accelerator complex.

## INTRODUCTION

The CERN Super Proton Synchrotron (SPS) has recently implemented a new Interlock Super Agent (ISA). ISA presents the complete chain of logic that has led to these beam dumps, providing the operator with a powerful diagnostic tool and systematic guidance during the setup phase for the different beams. Using so-called exporters, a variety of follow-up actions can be initiated automatically via the agent, such as automated logbook entries or fault registration in the CERN Accelerator Fault Tracking (AFT).

## BEAM INTERLOCK SYSTEM



Figure 1: Inputs Beam Interlock System.

The Beam Interlock System [1] (BIS) is a global beam interlocking system at CERN. It is responsible for monitoring the parameters of the particle beam in the accelerators and experiments in real time. If abnormal or dangerous conditions are detected, the BIS activates safety measures to stop the beam and protect equipment. The BIS

system is based on several Beam Interlock Controllers (BIC) receiving signals from various equipment (Fig. 1), each client must be decoded via a database.

## SOFTWARE INTERLOCK SYSTEM

The Software Interlock System [2] (SIS) is a key component of operation with over 1000 parameters monitored (at the SPS) to ensure optimum protection in all the different systems that are linked by equations (Fig. 2). Over time it has become an assistant that helps the operation team to achieve its mission, providing high quality beams and thus in a safe way.

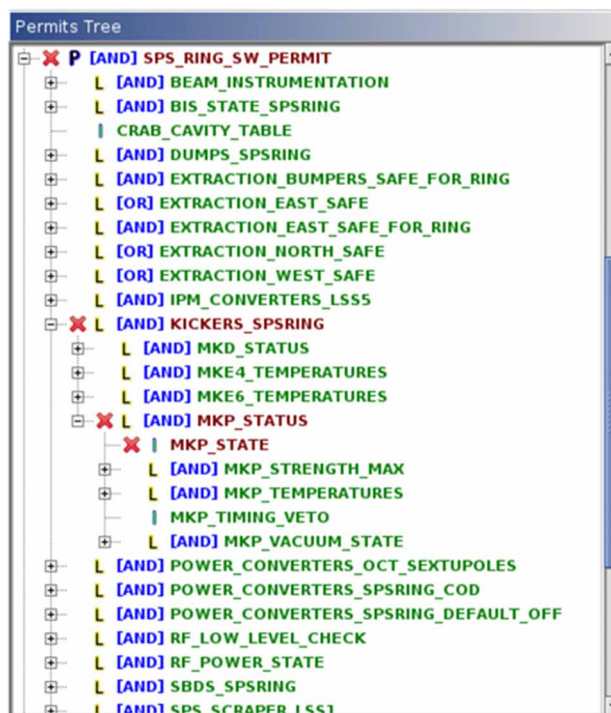


Figure 2: Tree structure with logic function AND and OR with as final result a permit.

## LOGBOOK

The logbook is used by the operator to record all the information needed to operate the machine, and to distribute the information to the machine coordinators and other team members. Information such as beam stops (Fig. 3) and masks applied to interlock systems are also recorded in the logbook.

# PROGRESS TOWARDS THE COMMISSIONING AND INSTALLATION OF THE 2PACL CO<sub>2</sub> COOLING CONTROL SYSTEM FOR PHASE II UPGRADE OF THE ATLAS AND CMS EXPERIMENTS

L. Zwalinski<sup>†</sup>, V. Bhanot, M. Ciupinski, J. Daguin, L. Davoine, M. Doubek, S. Galuszka, Y. Herpin, W. Hulek, T. Pakulski, P. Petagna, K. Sliwa, D. Teixeira, B. Verlaat  
CERN European Organization for Nuclear Research, 1211 Geneva 23, Switzerland

## Abstract

In the scope of the High Luminosity program of the Large Hadron Collider at CERN, the ATLAS and CMS experiments are advancing the preparation for the production, commissioning and installation of their new environment-friendly low-temperature detector cooling systems for their new trackers, calorimeters and timing layers. The selected secondary “on-detector” CO<sub>2</sub> pumped loop concept is the evolution of the successful 2PACL technique allowing for oil-free, stable, low-temperature control. The new systems are of unprecedented scale and largely more complex for both mechanics and controls than installations of today. This paper will present a general system overview and the technical progress achieved by the EP-DT group at CERN over the last few years in the development and construction of the future CO<sub>2</sub> cooling systems for silicon detectors at ATLAS and CMS. We will describe in detail a homogenised infrastructure and control system architecture which spreads between surface and underground and has been applied to both experiments. Systems will be equipped with multi-level redundancy (electrical, mechanical and control) described in detail herein. We will discuss numerous controls-related challenges faced during the prototyping program and solutions deployed that spread from electrical design organization to instrumentation selection and PLC programming. We will finally present how we plan to organise commissioning and system performance check out.

## INTRODUCTION

New Environment-friendly, non-flammable and low-temperature detector cooling systems, shown on Fig. 1, based on the evolution of the 2-Phase Accumulator Control Loop (2PACL) [1] concept are being constructed for ATLAS and CMS experiments to cope with challenging cooling needs of the new trackers, timing detectors and silicon based calorimeters. For many years, smaller systems with cooling powers up to 15 kW have been in continuous operation at CERN for LHCb Velo, ATLAS IBL and CMS Pixel Phase I detectors. The next generation systems will have much higher cooling power needs and complexity with 310 kW for ATLAS and 550 kW for CMS and about 1000 – 1800 evaporators each. Detector temperature requirements are also shifting lower to -40 °C. The cooling plants will be located in underground service caverns, displaced more than 120 m away from the detector proximity distribution manifolds and interconnected via concentric insulated transfer lines. The detector dissipated power removed by the oil-free 2PACL will be rejected to an R744 (CO<sub>2</sub>)-based primary refrigeration system.

The R744 refrigeration plants are located in a newly constructed surface building interconnecting the underground plants via two warm long transfer lines placed into ~90 m vertical shaft. The R744 system can pre-cool the 2PACL loop to about -53 °C (very close to the freezing point of




Figure 1: ATLAS and CMS CO<sub>2</sub> cooling system overview.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI



# THE SLS 2.0 BEAMLINE CONTROL SYSTEM UPGRADE STRATEGY

T. Celcer<sup>†</sup>, E. Zimoch, X. Yao, Paul Scherrer Institut, Villigen PSI, Switzerland

## Abstract

With over two decades of successful operation, the SLS facility is now undergoing a major upgrade that includes the complete replacement of the storage ring, yielding substantial improvements in beam emittance and brilliance and setting the stage for a new era of scientific exploration. As a critical component of the SLS 2.0 project, beamline upgrades are integral to harnessing the full potential of these enhanced beam characteristics.

To ensure that our users enjoy an optimal beamtime experience and maximize the scientific output, it is imperative to elevate the capabilities of our beamline control and data acquisition tools. Therefore, a thorough modernization and upgrade of our current control system stack is not just desirable but essential.

This article provides a comprehensive overview of the planned beamline control system upgrade, examining it from both technical and project management perspectives. We investigate the key sub-areas encompassed within a beamline control system upgrade and explore the strategies for efficiently integrating them together.

## INTRODUCTION

The Swiss Light Source (SLS) is a third-generation light source which has seen the first light in 2001 and has been one of the leading accelerator facilities in the last two decades. Nevertheless, with the latest innovations in storage ring designs, it is time for the SLS facility to undergo a major upgrade yielding the improved beam performance of the fourth-generation light sources.

The SLS dark time started on October 2<sup>nd</sup> this year, and is planned to last until the end of 2024, giving a very tight time window for such a major upgrade. The SLS 2.0 project is focused on the upgrade of the storage ring, and it does not cover the linac and the booster part [1]. Only a smaller part of the total project budget is allocated for the beamline upgrades, therefore only selected beamlines are included in the upgrade under the scope of the SLS 2.0 project. Nevertheless, all beamlines will go through larger or smaller upgrades and for this purpose, a parallel project, ESup (End Station upgrade), was started, in which the upgrades of beamlines outside of the SLS 2.0 project scope are planned. From control system point of view, we do not differentiate between these two projects.

Currently, the SLS hosts 18 user operation beamlines. During the SLS 2.0 project, two additional beamlines will be built (Debye and I-TOMCAT), two beamlines will be merged (PEARL and SIS), and one beamline will move to a different sector (microXAS), hence 19 beamlines will be available for user operation after the SLS 2.0 upgrade. While the SLS 2.0 project officially finishes end of 2024,

the beamline upgrades will continue to take place until the last planned milestone in mid 2026 and further.

To guarantee the availability of limited resources, the beamlines will be upgraded in three phases. During Phase 0, which concluded with the start of the dark time, one new beamline (Debye) was (partially) commissioned while one crystallography beamline (PX III) and few systems on other selected beamlines were upgraded. This phase was important not only for the time distribution of allocated resources, but also for testing and confirming the feasibility of the chosen SLS 2.0 control system solutions. During Phase 1 another seven beamlines in user program will be upgraded and 1 new beamline will be commissioned (I-Tomcat) and they are planned to restart the user operation in mid 2025, while the remaining beamlines will be upgraded during Phase 2, mainly during the additional 6-months shutdown in 2026.

## SLS Legacy and Challenges Ahead

The upgrade of the SLS and the project's timeline will pose many challenges, but also offer important opportunities. During the two decades of SLS operation the technical debt was growing, along with user demands. For instance, the chosen HW portfolio, mainly based on the VME platform, as well as existing SW capabilities could over the time not keep up with rising demands of the end user community. Meanwhile, e.g., fly scanning, which brings together complex motion control, device synchronisation, detector integration and data acquisition and handling, became a highly complex and critical task.

Furthermore, there was no centralised approach in providing a higher layer tool above the standardized EPICS layer, which would be responsible for beamline experiment control (BEC) and orchestration. Consequently, different solutions were adopted at different beamlines, but none of those solution were properly supported by a central organisation at PSI. Several years ago, controls section attempted to bridge this gap and developed pShell [2], however it was not adopted by all beamlines and could not resolve the lack of standardisation and lack of proper centralised strategy when it comes to the higher-level beamline experiment orchestration tools.

## BEAMLINE CONTROL SYSTEM UPGRADE STRATEGY

Learning from these experiences, creating a clear strategy for the SLS 2.0 control system upgrade, together with higher level SW solutions is critical for the success of the SLS 2.0 project, as its impact on the end user experience will be comparable to the impact of the improved beam characteristics obtained with the storage ring upgrade. This fact was also recognised by the SLS 2.0 project management, making Controls systems and Scientific IT area one

<sup>†</sup> tine.celcer@psi.ch

# SwissFEL RESONANT KICKER CONTROL SYSTEM

R. Krempaska, A. Alarcon, S. Dordevic, Ch. Gough, W. Portmann, M. Paraliiev  
Paul Scherrer Institute, Villigen PSI, Switzerland

## Abstract

SwissFEL X-ray Free Electron Laser at the Paul Scherrer Institute is a user facility designed to run in two electron bunch mode in order to serve simultaneously two experimental beamline stations. Two closely spaced (28 ns) electron bunches are accelerated in one RF macro pulse up to 3 GeV. A high stability resonant kicker system and a Lambertson septum magnet are used to separate bunches and to send them to the respective beamlines. The resonant kickers control system consists of various hardware and software components whose tasks are the synchronization of the kickers with the electron beam, pulse-to-pulse amplitude and phase monitoring, generating pulsed RF power to excite a resonating deflection current, as well as movement of the mechanical tuning vanes of the resonant kickers. The feedback SW monitors and controls all the important parameters. We present the integration solutions of these components into EPICS.

## INTRODUCTION

The Swiss X-ray Free Electron Laser [1] (SwissFEL) is a linear electron accelerator-based 4<sup>th</sup> generation light source capable of producing short, high brilliance x-ray photon pulses with duration of tens of femtoseconds down to hundreds of attoseconds for experimental research in material science, biochemistry, biophysics, and other fields [2]. SwissFEL has two undulator lines: a hard x-ray (Aramis) and soft x-ray (Athos) undulator line. To increase efficiency two closely spaced electron bunches are separated and sent to different undulator beamlines as schematically indicated in Fig. 1. This is done by a system of two high stability resonant kicker magnets followed by a septum magnet [3]. Both lines can operate simultaneously and independently up to the maximum machine repetition rate (100 Hz). In order to separate the two electron bunches both of them are deflected by a fast resonant kicker system: one up and the other down. Compensating dipoles counteract the deflection of the down-kicked bunch and send it straight through the zero field region of the Lambertson septum to the Aramis beamline. The up-kicked bunch is deflected by the Lambertson septum field sideways and is sent to the Athos beamline [3].

## RESONANT KICKER CONTROL SYSTEM COMPONENTS

The Resonant Kicker control system was developed to meet high stability pulse-to-pulse beam position requirements necessary for the stable operation of the two FEL

lines. Two identical kicker systems are used to reach the necessary beam deflection. The main components of each kicker magnet are: the timing and synchronisation system (TCS), the driver (DRV), the full range measurement (FRM), the precision amplitude measurement system (PMS) and finally the motion system used to remotely tune the resonant kicker (MOT).

### *Synchronisation and Timing System (TCS)*

The goal of the TCS is to synchronize the kicker with the electron beam and to produce the so-called ‘pulse train’ in order to control the DRV. The TCS receives a start signal from an event receiver and resynchronizes it to a stable 142.8 MHz SwissFEL machine clock signal. The event receiver signal is provided by an EVR-300 card controlled by a VME IOC running on the IFC1210 module.

### *Driver (DRV)*

A high amplitude stability and low phase-noise solid-state RF-driver developed at PSI provides up to ~7 kW RF power during the excitation period [3]. It is powered by a TDK-Lambda programmable DC power supply.

### *Full Range Measurement System (FRM)*

The full range measurement system provides pulse-to-pulse amplitude and phase information about the oscillating current of the kickers. It is based on fast (250 MHz), high resolution (16-bit) ADCs interfaced by a field-programmable gate array logic (Virtex 6, Xilinx) [3]. Two ADCs are running in parallel to further reduce the measurement system jitter. The amplitude and phase information provided by the FRM is used for general resonant kicker monitoring and for a phase feedback [3].

### *Precision Measurement System (PMS)*

The precision measurement system is a complementary offset-based system for higher performance amplitude measurements [3]. It provides higher (part-per-million) resolution measurement of the kicker’s sin-wave current amplitude using an offset or utilizing the so called balanced measurement method. However, this method has limited dynamic range (about 0.1% of the full range).

### *Motion System (MOT)*

This module controls the stepper motors which move the mechanical tuning vanes of the resonant kickers and reads their position via incremental encoder attached to the motor. The motor control is used by the feedback SW for the resonant kickers’ tuning.

# PROGRESS OF THE EPICS TRANSITION AT THE ISIS ACCELERATORS

I. D. Finch\*, B. R. Aljamal<sup>1</sup>, K. R. L. Baker, R. Brodie, J. L. Fernandez-Hernando,  
G. Howells, M. Leputa, S. A. Medley, M. Romanovschi  
STFC/RAL/ISIS, Chilton, Didcot, Oxon, United Kingdom  
A. Kurup, Imperial College of Science and Technology, London, United Kingdom  
<sup>1</sup>now at SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

## Abstract

The ISIS Neutron and Muon Source accelerators have been controlled using Vsystem running on OpenVMS / Itaniums, while beamlines and instruments are controlled using EPICS. We outline the work in migrating accelerator controls to EPICS using the pvAccess protocol with a mixture of conventional EPICS IOCs and custom Python-based IOCs primarily deployed in containers on Linux servers. The challenges in maintaining operations with two control systems running in parallel are discussed, including work in migrating data archives and maintaining their continuity. Semi-automated conversion of the existing Vsystem HMIs to EPICS and the creation of new EPICS control screens required by the Target Station 1 upgrade are reported. The existing organisation of our controls network and the constraints this imposes on remote access via EPICS and the solution implemented are described. The successful deployment of an end-to-end EPICS system to control the post-upgrade Target Station 1 PLCs at ISIS is discussed as a highlight of the migration.

## INTRODUCTION

The accelerators at the ISIS Neutron and Muon Source [1, 2] at Rutherford Appleton Laboratory have operated using a combination of Vsystem [3] commercial and EPICS [4] open-source control system software since Nov 2022. A full transition to EPICS is underway, but hybrid operation is expected to last several more years. (Note that ISIS Experiment Controls has already transitioned to the use of EPICS [5, 6].) During the transition, the two control systems must be run in parallel without interrupting operations. A software package called PVEcho [7, 8] has been developed to reliably bridge between the two systems.

Here we highlight the progress of the EPICS transition at the ISIS accelerators, with emphasis on lessons learned and differences with regard to the deployment at other facilities.

## ARCHITECTURE OVERVIEW

EPICS controls systems may be deployed in a variety of different configurations. We have chosen to prefer the pvAccess protocol [9], prefer to run our Input/Output Controllers (IOCs) in containers [10], and use Phoebus [11] as our primary user interface.

\* ivan.finch@stfc.ac.uk

## Servers

The Vsystem control system software suite runs on a cluster of four HP Itanium servers running the OpenVMS operating system. Previously a transition to Vsystem on Linux was considered [12], but loss of expertise means that this system is now expected to be maintained as is until obsolescence.

The operational EPICS control system is deployed on a trio of Linux servers running Docker in swarm mode [13] for failover capability. Docker in swarm mode was chosen over the more complex Kubernetes (K8) [14] for orchestration of our containers as we have little need for the more advanced K8 features such as the ability to adapt to dynamic workloads. A pair of load-balancing Linux servers running haproxy [15] are deployed to manage traffic to the operational servers for further failover capability.

An additional pair of Linux servers, also running Docker in swarm mode, are used as a development and CI/CD system. Both the operational and development swarms are managed via Portainer [16]. The development system is also used for Machine Learning Operations (MLOps) [17], including training of ML models.

## IOCs

All EPICS IOCs are deployed in containers running on the operational Linux servers. An unusual feature of our current deployment is that we have no conventional C/C++-based IOCs. The majority of PVs are bridged from Vsystem and are produced by PVEcho using the pvapy [18] Python library. Our first fully EPICS PVs are read from and written to ISIS Target Station 1 Omron PLCs using the CIP protocol, implemented in Python using the pvapy and CPPPO [19] libraries. For performance and maintainability reasons it is intended to migrate communications with these Omron PLCs to the MQTT protocol [20] and migrate PV management to the pvAccess for Python (p4p) [21] library.

The majority of the existing PVs managed by PVEcho are read from or written to our internally-developed CPS crates [22] via Vsystem. Prototype Python software using the p4p library has been developed to interface with the existing XML over HTTP protocol used to initialise and communicate with these crates. In the near-term the CPS-derived channels will be moved from Vsystem to EPICS, with the direction of the PVEcho bridging for these channels reversing.

The first conventional C/C++-based IOCs are being developed under contract by Mobiiis [23] to interface with Omron PLCs using the FINS protocol. The code is based on an

# TICKIT: AN EVENT-BASED MULTI-DEVICE SIMULATION FRAMEWORK

A. Emery, G. O'Donnell, C. Forrester, T. Cobb  
Diamond Light Source, Harwell, United Kingdom

## Abstract

Tickit is an event-based multi-device simulation framework providing configuration and orchestration of complex simulations. It was developed at Diamond Light Source in order to overcome limitations presented to us by some of our existing hardware simulations. With the Tickit framework, simulations can be addressed using a compositional approach. It allows devices to be simulated individually while still maintaining the interconnected behaviour exhibited by their hardware counterparts by modelling the interactions between devices. Devices can be collated into larger simulated systems providing a layer of simulated hardware against which to test the full stack of Data Acquisition and Controls tools.

We aim to use this framework to extend the scope and improve the interoperability of our simulations; enabling us to further improve the testing of current systems and providing a preferential platform to assist in development of the new Acquisition and Controls tools.

## MOTIVATION

Tickit's development was driven by the desire to simulate hardware triggered scans. To simulate such scans, multiple devices need to be able to communicate with one another and have linked behaviour triggered by events.

There were initial efforts to use Lewis [1], a cycle-driven hardware simulation framework for isolated devices from the European Spallation Source (ESS) and the ISIS Neutron and Muon Source. However, our requirements were not well matched to the function of this framework. This resulted in us writing bespoke solutions for our devices that grew increasingly complicated and verbose. Eventually a decision was made to develop a framework more appropriately suited to our needs.

The scans we wish to simulate require use of a Zebra [2], an event handling system utilising an FPGA from Quantum Detectors, an Eiger x-ray detector from Dectris, and a PMAC motion controller from Omron. To support scans of this nature, we would need to simulate the above devices as well as numerous motors and the communication between them. To provide this the framework needed to be:

- Multi device. To simulate a full scan we need to have many linked devices.
- Event driven. Each device in the system needs to update only when relevant, either when it changes state or when a device downstream changes an input to it.

A significant limiting factor in our hardware triggered scan simulations was the high FPGA frequency. The Zebra FPGA operates at 50 Mhz [3], a rate unachievable in a time

driven system. Even with the ability to match this rate, using a time driven approach would drive the system intensely with the majority of these updates being redundant. As the simulation progresses, simulation time and real time would slowly diverge, the rate of which increasing with added complexity. By using an event driven approach instead we only need to update each device when there is a relevant change. This enables the simulation to be run at a slower overall rate, lagging behind when operations are made, but then synchronising back to real time when there are periods of no change.

## DESIGN

The resulting framework consists predominantly of two parts: a scheduler, and components. Components encapsulate the simulated devices and their network interfaces, and the scheduler contains the logic to run the simulation. Components possess the operational logic for running and updating the devices they contain, and provide the interface by which the scheduler orchestrates the updating each device.

All devices possess optional inputs and outputs, which can be wired together to produce a directed acyclic graph of dependent devices. This device graphing is determined with a configuration yaml file which is used to build and run the simulation. Simple device graphing is presented in Fig. 1.




Figure 1: Device graphing. Wiring the device's inputs and outputs forms a directed acyclic graph. The scheduler ensures devices are updated in order, such that B is only updated when A has finished, and D is only updated once B and C have finished.

A summary of the framework's design and its constituent parts can be seen in Fig. 2.

### The Scheduler

The scheduler orchestrates the running of the simulation. It contains references to all the components in the simulation and the wiring of all of their inputs and outputs. It is responsible for routing all the changes through the system and ensuring time is maintained.

# CONTROL SYSTEM DESIGN OF THE CHIMERA FUSION TEST FACILITY

P. T. Smith, A. J. Greer, B. A. Roberts, P. B. Taylor, Observatory Sciences Ltd, St Ives, UK  
M. Roberts, D. Mccubbin, Jacobs Clean Energy Ltd, Warrington, UK

## Abstract

CHIMERA is an experimental nuclear fusion test facility which aims to simulate the intense magnetic fields and temperature gradients found within a tokamak fusion reactor. The control system at CHIMERA is based on EPICS and will have approximately 30 input/output controllers (IOCs) when it comes online in 2024. It will make heavy use of CSS Phoebus for its user interface, sequencer and alarm system. CHIMERA will use EPICS Archiver Appliance for data archiving and EPICS areaDetector to acquire high speed data which is stored in the HDF5 format. The control philosophy at CHIMERA emphasises PLC based control logic using mostly Siemens S7-1500 PLCs and using OPCUA to communicate with EPICS. EPICS AUTOSAVE is used both for manually setting lists of process variables (PVs) and for automatic restoration of PVs if an IOC must be restarted.

## INTRODUCTION

CHIMERA is being developed for the United Kingdom Atomic Energy Authority (UKAEA) [1] with Jacobs Clean Energy Ltd as principle designer and constructor and Observatory Sciences Ltd developing the SCADA system. The CHIMERA nuclear fusion test facility (see Figures 1 and 2) aims to test various in-vessel components of a nuclear fusion reactor in conditions that replicate the harsh environment in which they will operate [1]. It will recreate the magnetic field gradient, magnetic field transients and heat fluxes that exist within a tokamak reactor. CHIMERA is designed to study a subject under test (SUT) up to a volume of  $1.67 \times 0.96 \times 0.46 \text{ m}^3$ . This is the size of the Test Blanket Module (TBM) being developed for the International Thermonuclear Experimental Reactor (ITER) and tested at CHIMERA. To properly test the TBM, a cooling system which runs at up to pressurised water reactor conditions (328 °C, 155 bars) is being built at CHIMERA and will be connected to the TBM during testing.

At the end of phase 1, CHIMERA will be capable of producing a peak static magnetic field of 5 tesla, a  $\pm 0.25$  tesla pulsed magnetic field, surface heating of  $0.5 \text{ MW/m}^2$ , and volumetric heating of 100 kW. CHIMERA phase 2 will introduce a high heat flux continuous-wave laser producing heat fluxes of  $200 \text{ MW/m}^2$  over  $100 \text{ mm}^2$  or  $20 \text{ MW/m}^2$  over  $1500 \text{ mm}^2$ . It will also introduce a liquid metal loop which can be used to circulate liquid PbLi around a water-cooled lithium lead (WCLL) blanket [2].

A separate, but equally important aspect of project CHIMERA, is the facility's digital twin. This term describes a suite of software which will make use of CHIMERA test

data to develop a sophisticated simulation of the experiments being carried out at CHIMERA. The intention is that this simulation will run in real-time and use live test data from the data acquisition (DAQ) system to predict the outcome of experiments and further improve the simulation. It will also be used to predict upcoming problems with live experiments before they occur [1].



Figure 1: Top down view of the CHIMERA facility. (Courtesy of UKAEA).




Figure 2: Cross-sectional view into the CHIMERA vacuum test chamber. The superconducting magnets are not shown (Courtesy of UKAEA).

## CONTROL SYSTEM OVERVIEW

The SCADA system is being developed using the Experimental Physics and Industrial Control System (EPICS) framework. EPICS was chosen in part, due to its open source and flexible nature which will allow this experimental facility to develop over time. The key software can be divided into “high-level” components and “low-level” components. The high-level components are the DAQ system, data archiver, server monitoring system, experiment sequencer, alarm handling system and web server for external access of data. The “low-level” components are 7 IOCs which directly control and monitor various hardware components. All of these SCADA systems communicate over a controls network via TCP/IP, most communications are done with the EPICS channel access (CA) protocol. This paper will focus on giving an overview of the design and current development of the SCADA system prior to the commissioning of the CHIMERA facility.

# SOFTWARE AND FIRMWARE-LOGIC DESIGN FOR THE PIP-II MACHINE PROTECTION SYSTEM MODE AND CONFIGURATION CONTROL AT FERMILAB\*

L. Carmichael, M. Austin, J. Eisch, E. Harms, R. Neswold, A. Prosser, A. Warner, J. Wu  
 Fermilab, Batavia, IL 60510, U.S. A

## Abstract

The PIP-II Machine Protection System (MPS) requires a dedicated set of tools for configuration control and management of the machine modes and beam modes of the accelerator. The protection system reacts to signals from various elements of the machine according to rules established in a setup database in the form of a Look-Up-Table filtered by the program Mode Controller. This is achieved in accordance with commands from the operator and governed by the firmware logic of the MPS. This paper describes the architecture, firmware logic, and implementation of the program mode controller.

## INTRODUCTION

The Proton Improvement Plan-II (PIP-II) is an enhancement to the Fermilab accelerator complex [1] that will provide intense high energy neutrino beam to the Deep Underground Neutrino Experiment (DUNE) [2]. PIP-II (Fig.1) will consist of an 800 MeV H<sup>-</sup> Superconducting linac which includes a Warm Front-end (WFE), and a 300-meter-long beam transfer line to the Fermilab Booster. The WFE of the linac plays a critical role in the accelerator. It generates a 30 keV H<sup>-</sup> beam, defines the beam parameters, accelerates the beam to an energy of 2.1 MeV with its RFQ for compatibility with downstream accelerating structures, and generates a required bunch pattern. One of the high-level goals of the machine is to deliver a proton beam power to target in excess of 1 MW with sustained high reliability along with multiuser operations of the Fermilab complex. A PIP-II Injector Test Facility [3] (PIP2IT) was assembled in multiple stages in 2014-2021 as a testbed for PIP-II technologies and protection schemes.




Figure 1: Schematic of PIP2 Facility.

A Machine Protection System (MPS) capable of managing high power is needed to protect the accelerator and its components from beam induced damage. This system also needs to allow for seamless transition from the defined operational modes while maintaining MPS integrity. Four Beam Modes specify the beam pulse length and repetition rate, and seven Machine Modes specify beam path configurations. In each Beam Mode the machine can be in one of three states – Standby, Ready or Beam on. The required fast response time to shut off the linac beam in response to critical detected faults and beam losses above measured thresholds is 10 microseconds, accomplished by inhibiting the Low Energy Beam Transport (LEBT) line Chopper which is one of four (4) Beam Inhibit Devices (BIDs). Commissioning of the WFE is scheduled for 2026 with phased commissioning of the full linac thereafter. The final design of the MPS have been approved and is commensurate with this timescale.

## MPS SYSTEM ARCHITECTURE

The system architecture is shown in Fig. 2. The MPS is FPGA based and consists of a Main MPS (MPSM) which issues the system permits and interfaces with the BIDs, an Analog MPS (MPSA) for post-processing of digitized signals derived from certain beam current measuring devices and a Digital MPS (MPSD) which processes serialized inputs from machine subsystems coming from the field via serializer hardware.

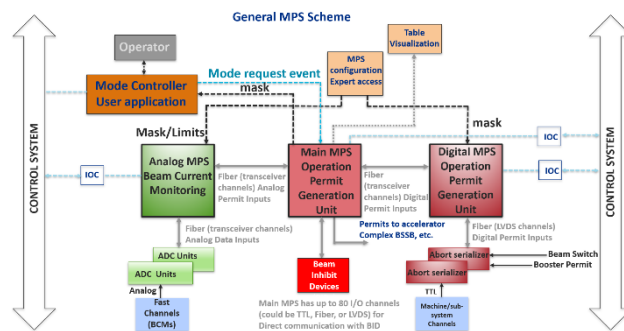


Figure 2: MPS Integrated Overview Diagram.

The MPSM receives all permit input status bits from the MPSA and MPSD Units. For expansion purposes it can accommodate up to 20 fiber transceiver channels from multiple MPSA and MPSD units. It can send data to other MPS units, and it can accept up to 80 I/O channels for direct connection with the BIDs. These can be TTL, Fiber or LVDS. Based on the Operation Mode or Masked channel configuration the system will read all the inputs from the MPSD and MPSA units and decide whether to

\* This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# A FLEXIBLE EPICS FRAMEWORK FOR SAMPLE ALIGNMENT AT NEUTRON BEAMLINES

M. Henderson, J. Edelen\*, M. Kilpatrick, RadiaSoft LLC, Boulder, USA  
S. Calder, B. Vacaliuc, ORNL RAD, Oak Ridge, USA

R. D. Gregory, G. Guyotte, C. Hoffmann, B. Krishna, ORNL, Oak Ridge, USA

## Abstract

RadiaSoft has developed a flexible front-end framework, written in Python, for rapidly developing and testing automated sample alignment IOCs at Oak Ridge National Laboratory. We utilize YAML-formatted configuration files to construct a thin abstraction layer of custom classes which provide an internal representation of the external hardware within a controls system. The abstraction layer employs the PCASPy and PyEpics libraries in order to serve EPICS process variables and respond to read/write requests via Channel Access, with future developments planned for PV Access through the P4P library. Our framework allows users to build a new IOC that has access to information about the sample environment in addition to user-defined machine learning models and data processing methods. The IOC monitors for user inputs, performs user-defined operations on the beamline, and reports its status back to the control system. Our IOCs can be booted from the command line, and we have developed command line tools for rapidly running and testing alignment processes. These tools can also be accessed through EPICS GUIs or separate Python scripts. These proceedings provide an overview of our software structure and showcases its use at two beamlines at ORNL.

## INTRODUCTION

Robust, accessible controls software for beamline sample environments are a critical need for operators at neutron science user facilities like the Spallation Neutron Source (SNS) and High-Flux Isotope Reactor (HFIR) at Oak Ridge National Laboratory (ORNL). Existing input-output controller (IOC) software and control workflows allow beamline operators to meet the needs of users, but require operators to expend significant amounts of time and resources on trivial tasks like sample alignment which make poor use of their expertise and are good targets for automation. Additionally, the creation of new workflows or the extension of existing ones typically require appreciable effort from controls experts who already experience heavy workloads unrelated to day-to-day beamline operations at user facilities. To that end, RadiaSoft has developed *rscontrols*: a flexible front-end controls framework, written in Python, for rapidly developing IOCs with embedded machine learning (ML) models and other modern automation tools. *rscontrols* leverages common configuration tools and a Pythonic API to enhance user accessibility and is built on existing Python packages for implementing controls via the EPICS framework.

\* jedelen@radiasoft.net

## SUPPORTING PACKAGES

In addition to supporting broad user accessibility, our choice to employ Python for the *rscontrols* framework has allowed us to take advantage of pre-existing tools for executing controls operations through EPICS. For client-side access to existing networks of EPICS process variables (PVs) we use the PyEpics package [1], a Python API for the channel access (CA) protocol within EPICS. For serving new PVs associated with the *rscontrols* framework to the network we use the PCASPy [2] package, a Python API for the Portable CA Server (PCAS) support module for EPICS.

## USER INPUTS

One of the strengths of the *rscontrols* framework is the reduced workload placed on users in comparison to pure EPICS or the Python APIs employed by *rscontrols*. To create new IOCs using those tools typically requires significant development time and the efforts of expert programmers. Even in the case of Python APIs like PyEpics and PCASPy, these efforts generally produce IOCs dedicated to specific processes that therefore feature low levels of reusability. In contrast, user inputs for *rscontrols* IOCs have been reduced to a single configuration file with a human-readable layout and hooks for control and server processes which take the form of simple Python functions. This allows new IOCs to be developed quickly by operators familiar with local beamline equipment and processes and a basic understanding of Python scripting.

### Configuration Files

The primary user input for an IOC created with *rscontrols* is a YAML-formatted configuration file. This file provides a list of all hardware elements to be abstracted by the software, including the PVs associated with each element, as well the available ML models, control processes, and any new PVs to be served by the IOC (see Appendix).

### Controls Process Scripts

One of the most important aspects of *rscontrols* is its method of representing user-defined controls processes as Python functions. A single Python module containing these functions is specified by the user via the *path* entry of the *Processes* section of a config file. Only processes listed in the config file are imported for use into the framework, though functions in the processes module can freely depend on one another (including functions which do not correspond to any active processes).

# MACHINE LEARNING BASED NOISE REDUCTION OF NEUTRON CAMERA IMAGES AT ORNL

I.V. Pogorelov\*, J. Edelen<sup>†</sup>, M. Henderson, M. Kilpatrick, RadiaSoft LLC, Boulder, USA  
S. Calder, B. Vacaliuc, ORNL RAD, Oak Ridge, USA  
R.D. Gregory, G. Guyotte, C. Hoffmann, B. Krishna,  
ORNL, Oak Ridge, USA

## Abstract

Neutron cameras are utilized at the HB2A powder diffractometer to image the sample for alignment in the beam. Typically, neutron cameras are quite noisy as they are constantly being irradiated. Removal of this noise is challenging due to the irregular nature of the pixel intensity fluctuations and the tendency for it to change over time. RadiaSoft has developed a novel noise reduction method for neutron cameras that inscribes a lower envelope of the image signal. This process is then sped up using machine learning. Here we report on the results of our noise reduction method and describe our machine learning approach for speeding up the algorithm for use during operations.

## INTRODUCTION

Neutron Cameras are a fairly ubiquitous piece of instrumentation specifically in neutron scattering experiments. Here instrument scientists utilize neutron cameras to visualize the sample location relative to the incident beam to ensure proper alignment in the beam. While these cameras are highly useful, due to the fact that they experience direct radiation the pixels degrade over time and noise develops that degrades the image quality. For many cases this degradation doesn't impact the ability to visualize the sample but for low absorbing samples it can be nearly indistinguishable from the beam. Over the years numerous efforts have been made to remedy noise in neutron camera images. The advent of machine learning has also led to a resurgence of these efforts as the problem is notoriously difficult to solve. Median filters are common approach to removing this noise [1] due to the nonuniform nature of the noise. Recent work on adaptive median filters [2] has shown some promise but median filters can be slow and are not always robust to changes in the noise characteristics over time. Another recent effort utilizes Generative Adversarial Networks to remove signal noise [3]. There is a strong machine learning upside to this technique however it relies on machine learning as the primary mechanism for removal of the noise as opposed to understanding the noise and using machine learning as a speed up tool. Other methods using principle component analysis have also shown promise in recent years [4]. Here we develop a denoising algorithm based on the fundamental characteristics of the noise in the neutron camera data. We then use convolutional neural networks to speed up this calculation for use in real time during operations. Our algorithm

has been fully tested and deployed at the HB2A beamline at Oak Ridge National Laboratory.

## DENOISING ALGORITHM

Speckle/"salt-and-pepper" noise is conspicuous in the raw images collected from the neutron camera at HB2A, as can be seen in the left panels of Fig. 1. Cleaning up the images improves our ability to generalize sample identification algorithms in images from neutron cameras and improves the ability for instrument scientists to visualize the sample in the beam. Our machine learning algorithm is trained using data generated using a first principles algorithm developed to remove noise by identifying its key characteristics.

We studied the noise in the images by examining the data on a slice-by-slice basis. We chose a horizontal slice as the aspect ratio of the beam is more favorable for understanding the noise characteristics in this direction. An inspection of the signal contained in a typical row of 640 pixels, such as that pictured in Fig. 2, shows that the statistical properties of the speckle noise are quite different from those of the Gaussian white noise that is predominantly encountered in practice and for which numerous denoising algorithms exist. Two salient features are, the large noise amplitude that is comparable in magnitude to the signal itself, and the "one-sidedness" of the noise, in the sense that the noise only assumes positive values. This special character of the noise makes it possible to attempt its removal by approximating the targeted denoised signal via inscribing an envelope from below, as shown in Fig. 2.

The inscribed envelope can be computed for each row of pixels, whereupon the denoised rows replace the original pixel rows in the image. Interestingly, not every way of inscribing an envelope is equally robust; the best-performing of the algorithms that we explored is based on a kind of min-pooling within a window moving along a line of pixels. It proved important to properly match the size of the min-pooling window to the size of the footprint of the spikes that make up the noise in the images. Figure 1 illustrates the results of application of this procedure to two of the more feature-rich sample images.

## MACHINE LEARNING ADAPTATION

We developed a machine learning adaptation of our denoising algorithm for the purpose of speeding up execution during operations. While our baseline algorithm is not prohibitively slow, higher data rates demand faster image

\* ilya@radiasoft.net

<sup>†</sup> jedelen@radiasoft.net



# MACHINE LEARNING FOR COMPACT INDUSTRIAL ACCELERATORS

J. P. Edelen, M. J. Henderson, J. Einstein-Curtis and C. C. Hall, RadiaSoft LLC, Boulder CO, USA,  
J. A. Diaz Cruz and A. L. Edelen, SLAC, Menlo Park CA, USA

## Abstract

The industrial and medical accelerator industry is an ever-growing field with advancements in accelerator technology enabling its adoption for new applications. As the complexity of industrial accelerators grows so does the need for more sophisticated control systems to regulate their operation. Moreover, the environment for industrial and medical accelerators is often harsh and noisy as opposed to the more controlled environment of a laboratory-based machine. This environment makes control more challenging. Additionally, instrumentation for industrial accelerators is limited making it difficult at times to identify and diagnose problems when they occur. RadiaSoft has partnered with SLAC to develop new machine learning methods for control and anomaly detection for industrial accelerators. Our approach is to develop our methods using simulation models followed by testing on experimental systems. Here we present initial results using simulations of a room temperature s-band system.

## INTRODUCTION

In recent years machine learning (ML) has been identified as having the potential for significant impact on the modeling, operation, and control of particle accelerators (e.g. see [1, 2]). Specifically, in the diagnostics space, there have been many efforts focused on improving measurement capabilities and detecting faulty instruments. When it comes to diagnostics, developments for beam position monitors have been quite ubiquitous over the years. Relatively recently, ML methods have been utilized to improve optics measurements from beam position monitor data [3]. Additionally, machine learning has been used to identify and remove malfunctioning beam position monitors in the Large Hadron Collider (LHC), prior to application of standard optics correction algorithms [4]. Furthermore, we have developed techniques for automation of noise removal in BPM data using machine learning [5]. On the contrary, there is a real dearth of knowledge when it comes to the application of machine learning for industrial accelerators. Moreover, the developments for using machine learning to improve RF signal processing are considerably further behind than other diagnostics in use at accelerators. The ability to remove noise from RF measurements would greatly improve our ability to extract meaningful information from RF systems especially in an industrial setting.

Machine learning methods such as autoencoders and variational autoencoders (VAEs) are well established for the removal of noise from various signals. For VAEs specifically noise reduction due to the enforcement of a smoothness condition in the latent-space representation. This feature of VAEs has been applied to gravitational wave research [6, 7] and geophysical data [8], for example. Recurrent autoen-

coders have the added advantage of being well suited to work with data sequences. In this paper we explore the use of Variational Recurrent Autoencoders (VRAEs) to remove different power law spectra (colors) of noise from simulated BPM data in a ring.

Our work utilizes a combination of approaches to understand which is best when considering RF waveform data. Our work has explored the use of model based approaches such as Kalman filters and machine learning approaches such as convolutional neural networks and variational autoencoders. Here we begin with a review of our data generation model followed by an analysis of Kalman filters, convolutional autoencoders, and variational autoencoders for the removal of noise from RF signals.

## DATA GENERATION

Our data was generated using a RF simulator that reproduces waveforms as they would be seen in industrial systems. Over the past year, RadiaSoft has been developing a full RF simulation tool that is integrated with EPICS for the development of new control algorithms, developing IOC software, and testing user interfaces. The simulator can be run through various APIs including a command line interface, via a Jupyter notebook, or directly through an EPICS connection. The simulator is based on a linear circuit model that takes into account coupling factor, quality factor, frequency, drive amplitude and phase, pulse duration, detuning, etc. The dynamics of our model are based off of equations derived here [9–11].

The data were generated by varying the RF pulse characteristics and the cavity characteristics. For the pulse the length of the pulse was varied from 3  $\mu\text{s}$  to 7.5  $\mu\text{s}$  which is a reasonable range for industrial accelerator applications operating at S-Band. Additionally we varied the start time of the RF pulse in the data window. While we typically don't expect the RF pulse to vary in position along the DAQ window adding in this flexibility will ensure better generalization when transferring from simulations to measurement.

The RF cavity parameters of interest for this study are  $Q_0$  and  $\beta$  which were varied over a range of 10,000 to 225,000 for the  $Q_0$  and 1 to 3 for  $\beta$ . The detuning was also varied within a range of plus minus one half bandwidth, a fairly typical range seen on industrial systems. In all the parameter range chosen represents a reasonable range of industrial RF systems and will allow us to develop simulation based algorithms that should be readily transferable to measurement when the time arises.

## KALMAN FILTERS

First we consider the Kalman filter for noise reduction. Kalman filters, also referred to as linear quadratic estimators,

# MACHINE LEARNING BASED SAMPLE ALIGNMENT AT TOPAZ

M. Henderson, J. Edelen\*, M. Kilpatrick, I.V. Pogorelov, RadiaSoft LLC, Boulder, USA  
S. Calder, B. Vacaliuc, ORNL RAD, Oak Ridge, USA  
R.D. Gregory, G. Guyotte, C. Hoffmann, B. Krishna, ORNL, Oak Ridge, USA

## Abstract

Neutron scattering experiments are a critical tool for the exploration of molecular structure in compounds. The TOPAZ single crystal diffractometer at the Spallation Neutron Source studies these samples by illuminating samples with different energy neutron beams and recording the scattered neutrons. During the experiments the user will change temperature and sample position in order to illuminate different crystal faces and to study the sample in different environments. Maintaining alignment of the sample during these processes is key to ensuring high quality data are collected. At present this process is performed manually by beamline scientists. RadiaSoft in collaboration with the beamline scientists and engineers at ORNL has developed a new machine learning (ML) based alignment software automating this process. We utilize a fully-connected convolutional neural network with dropout, configured in a U-net architecture, to produce sample segmentation masks from which we compute the sample center of mass. We then move the sample using a custom python-based EPICS IOC interfaced with the motors. In these proceedings we provide an overview of our ML tools and initial results aligning samples at ORNL.

## INTRODUCTION

The TOPAZ instrument is a high-resolution neutron time-of-flight (TOF) Laue diffractometer for single-crystal diffraction located on a beamline at the Spallation Neutron Source (SNS) user facility at Oak Ridge National Lab (ORNL) [1]. Along with the diffractometer itself (which already features a variable number of individual neutron detectors), the TOPAZ beamline is host to a large array of instruments and hardware for executing experimental controls and maintaining desirable sample environments. Implementing these controls and collecting the data needed for users to achieve their scientific goals are the responsibility of a dedicated group of beamline operators with a high degree of expertise in experimental neutron science. Simplifying the workflows of these operators, and automating controls processes wherever possible, is a critical effort for maximizing the scientific output of user facilities like the SNS.

## CONTROLS AT TOPAZ

In addition to the overall mechanical complexity of beamline hardware and instrumentation, operational logistics at TOPAZ include the navigation of two distinct operational modes and a rich but complex network of controls software.

## Operational Modes

The operational modes in use at TOPAZ are an ambient-temperature (ambient) mode and a cryogenically-controlled (cryo) mode. These modes differ not only in the range of temperatures they represent within the sample chamber, but also in the hardware and controls available during their operation. Most notably, because the cryostream instrument used in cryo-mode enters the chamber from a port which typically hosts a diagnostic camera in ambient-mode, an alternative camera mounted to a side port (from which the sample arm is not clearly visible) must be used. Figure 1 demonstrates the difference in sample views provided in ambient- (top) and cryo-mode (bottom). Additionally, sample shields sometimes used in cryo-mode are visible in and can partially obscure the sample images.

## Controls Network

Experimental controls at TOPAZ are implemented using Channel Access (CA) protocols within the EPICS framework. Like most instruments at the SNS, TOPAZ encompasses an extensive network of EPICS controls and process variables (PVs). This network covers sample arm motor positions, diagnostic camera controls, thermal conditions, neutron guide environments, detector settings, and much more [1]. In addition to low-level controls, the network also features several layers of abstracted controls mechanisms. These include mechanisms for security purposes, such as virtual motors used to validate proposed motor controls, and input-output controllers (IOCs) for simplifying workflows and automating controls processes.

## Sample Alignment

One feature of the EPICS network at TOPAZ is an IOC for executing semi-automated sample alignment. This IOC automates the determination of motor positions, but requires a human operator to identify approximate sample centroids in diagnostic camera images. Operators use a graphical user interface (GUI) to initiate an alignment state, and then must click a sample image in the GUI to identify the centroid. This process is repeated over several adjustments to motor positions, after which the center of the sample will have been aligned with the current beam position. Although this process is much faster than manually and incrementally updating motor positions until the sample centroid and beam are aligned, it still requires operators to expend time and attention on alignment which could be better spent on more complex tasks. In cases where samples must be re-aligned frequently (e.g., due to sample shifting during thermal variations), this expense becomes significant.

\* jedelen@radiasoft.net

# CLASSIFICATION AND PREDICTION OF SUPERCONDUCTING MAGNET QUENCHES

J. Einstein-Curtis \*, J. P. Edelen, M. Kilpatrick, R. O'Rourke, RadiaSoft LLC, Boulder, Colorado  
K. A. Drees, J. S. Laster, M. Valette, BNL, Upton, New York

## Abstract

Robust and reliable quench detection for superconducting magnets is increasingly important as facilities push the boundaries of intensity and operational runtime. RadiaSoft has been working with Brookhaven National Lab on quench detection and prediction for superconducting magnets installed in the RHIC storage rings. This project has analyzed several years of power supply and beam position monitor data to train automated classification tools and automated quench precursor determination based on input sequences. Classification was performed using supervised multilayer perceptron and boosted decision tree architectures, while models of the expected operation of the ring were developed using a variety of autoencoder architectures. We have continued efforts to maximize area under the receiver operating characteristic curve for the multiple classification problem of real quench, fake quench, and no-quench events. We have also begun work on long short-term memory (LSTM) and other recurrent architectures for quench prediction. Examinations of future work utilizing more robust architectures, such as variational autoencoders and Siamese models, as well as methods necessary for uncertainty quantification will be discussed.

## INTRODUCTION

Quench protection systems have been in service since the advent of superconducting magnet technology [1, 2]. In spite of the long history of quench detection technology, there has been a continuous effort to improve quench protection through the 90's with the construction of the LHC [3, 4] and in the 2000's with the construction of large detector magnets such as CMS [5] and the MICE experiment [6]. For a single magnet, once a quench is detected the power supply is switched off and then the magnet energy dumped into a load resistor through a cold diode. When concerned with multiple magnets either bridge circuits or isolation amplifiers are usually considered [7]. For the MICE experiment for example, the magnet is divided into subdivisions in order to reduce the impact of a quench in a single subdivision. At RHIC there are multiple magnets chained together which introduces additional complexity [8]. In addition to the hardware requirements for protecting the magnet from quench events, there have been a number of efforts to ensure robust timing and triggering systems for coordinating the different quench protection systems and to ensure there is not a dirty beam dump. At RHIC, specifically, a redundant fiber optic communications system has been developed and installed to ensure effective permitting [9].

\* joshec@radiasoft.net

In addition to refining quench protection systems, quench detection also remains an area of interest. Traditional quench detection relies on measurement of the magnet resistance through the voltage and current delivered by the power supply. More recently efforts to identify quenches in advance to take preventative action has led to some interesting developments. For example, computation of the parasitic capacitance has demonstrated early detection of quench events [10]. Additional efforts have focused on using acoustic sensors to detect quenches [11]. This method monitors the change in the magnet's acoustic transfer function induced by a local temperature rise or an epoxy crack. In fact, a recent machine learning effort to detect precursors to magnet quenches using similar acoustic data has been quite successful [12]. While these methods show significant promise, much work can be done to refine the machine learning methods and to integrate them with operational accelerators. Here we develop tools for the classification of quench events in the hopes of being able to detect quench precursors in power supply or BPM data. We begin with a quench overview of our dataset and then provide details for our quench classification results.

## QUENCH DATA AND INITIAL ANALYSIS

The data included in our studies are from two separate datasets provided by BNL: beam position monitor (BPM) and power supply (PS). The original data was received as text files for each device, inside a hierarchy of directories representing the run, fill number, event type, and ring. The BPM text files contain timestamped beam position, difference, and coherence data at 10 kHz for 100 milliseconds around the event. The PS text files contain timestamped reference current, current, voltage, and voltage error data at 720 Hz for the 3 seconds leading up to the event and 1 second after the event. Additionally, Excel files were provided by BNL that contain information about the events, including the names of the specific power supply device names that were involved in reporting a quench event.

Before development of our ML models, it was first necessary to determine the characteristics of the data to determine the necessity of any additional data pre-processing. This included detailed investigations of the data itself around quench events, including gaining an understanding of the proper data sampling periods, expected signal behavior, and if there are any sort of label generation methods readily available. We also performed metadata analyses of the devices involved in the quenches; histograms of these studies can be seen in Figures 1 and 2.

Examining individual waveforms, we determined that quench vs non-quench datasets can be identified with high

# HOW EMBRACING A COMMON TECH STACK CAN IMPROVE THE LEGACY SOFTWARE MIGRATION EXPERIENCE

C. D. Burgoyne, C. R. Albiston, R. G. Beeler, M. Fedorov, J. J. Mello, E. R. Pernice, M. Shor  
Lawrence Livermore National Laboratory, Livermore, USA

## Abstract

Over the last several years, the National Ignition Facility (NIF), the world's largest and most energetic laser, has regularly conducted approximately 400 shots per year. Each experiment is defined by up to 48 unique pulse shapes, with each pulse shape potentially having thousands of configurable data points. The importance of accurately representing small changes in pulse shape, illustrated by the historic ignition experiment in December 2022, highlights the necessity for pulse designers at NIF to have access to robust, easy to use, and accurate design software that can integrate with the existing and future ecosystem of software at NIF. To develop and maintain this type of complex software, the Shot Data Systems (SDS) group has recently embraced leveraging a common set of recommended technologies and frameworks for software development across their suite of applications. This paper will detail SDS's experience migrating an existing legacy Java Swing-based pulse shape editor into a modern web application leveraging technologies recommended by the common tech stack, including Spring Boot, TypeScript, React and Docker with Kubernetes, as well as discuss how embracing a common set of technologies influenced the migration path, improved the developer experience, and how it will benefit the extensibility and maintainability of the application for years to come.

## INTRODUCTION

On December 5, 2022, the National Ignition Facility (NIF), the world's largest and most energetic laser, made history by achieving ignition in a laboratory setting for the first time [1]. This milestone was made possible in part by various software applications which are instrumental in the NIF's ability to regularly and efficiently conduct up to 400 experiments per year. Given the speed of technological innovations in software, and with NIF now in its second decade of operations, several of these software systems require various upgrades or rewrites.

Recently, a team within the Shot Data Systems (SDS) group worked on rewriting one of these legacy tools, a Java Swing based desktop application called Pulse Shape Editor (PSE), to a new single-page, modern, web-application known as Pulse Shape Tool (PST). This application, whose purpose is to create and define the laser pulses used on NIF, is instrumental in the shot process, empowering pulse designers to configure potentially thousands of individual pulse points per experiment. The importance of PST functioning efficiently and accurately was illustrated by the successful December 2022 shot, where small adjustments to the energy and timing of laser pulses, achieved in part

through manipulation of pulse point and spline point data in PST, played a role in reaching ignition [2].

With limited resources and many disparate applications to develop and maintain, when choosing a migration path for PST it was critical for SDS developers to choose a set of technologies, also known as a technology stack or tech stack, that was modern yet leveraged the knowledge already available within the team. With so many existing technologies available, and more becoming available nearly every day, choosing an ideal tech stack for a given application and development team can be challenging and time-consuming. Ultimately, it was chosen to develop PST using a tech stack that was already being utilized among a subset of other SDS applications, resulting in benefits such as the ability to leverage existing developer knowledge of technologies that were also well-supported by the developer community at large.

## REQUIREMENTS

The original PSE was designed as a stand-alone Java Swing application (Fig. 1). Users would download PSE onto their desktops prior to use, yet running the application still required an Internet connection to use all the features due to reliance on a remote database connection for storing and retrieving pulse data. Besides using outdated technology containing vulnerabilities, many of PSE's limitations were due to architectural decisions made by the original developers. For example, since PSE was a desktop application, it was difficult to extend it to dynamically integrate with other applications in real-time, a feature increasingly expected in modern applications. Furthermore, these same architectural decisions made it difficult to maintain in general. After thorough examination of the PSE code, it was decided performing an in-place upgrade would not provide enough added value and would be needlessly difficult and time consuming to complete. Instead, the decision was made to migrate the functionality to a new application.

Key functionalities in PSE included the ability to edit, save and plot pulse and spline points, simple data validation, and the ability to import and export pulse data to and from external files and the database. Users required these key functionalities preserved in the migration to PST. Users also requested additional modern features be integrated into the new application such as the ability for users to modify and plot multiple pulses simultaneously, support for multiple users editing the same pulses concurrently, complex data integrity checks, per pulse permission management, and complex integrations with other applications, all while ensuring the software remained easy for developers to maintain and extend as required. These requirements

# CONCEPTUAL DESIGN OF THE MATTER IN EXTREME CONDITIONS UPGRADE (MEC-U) REP-RATED CONTROL SYSTEM

B. Fishler<sup>†</sup>, F. Batysta, J. Galbraith, V. Gopalan, J. Jimenez, L. Kiani, E. Koh, E. Sistrunk,  
J. McCarrick, A. Patel, R. Plummer, B. Reagan, T. Spinka, K. Terzi, K. Velas  
Lawrence Livermore National Laboratory, Livermore, USA  
M. Cabral, A. Wallace, J. Yin, SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

The Lawrence Livermore National Laboratory (LLNL) is delivering the Dual-mode Energetic Laser for Plasma and High Intensity Science (DELPHI) system to SLAC as part of the MEC-U project to create an unprecedented platform for high energy density experiments. The DELPHI control system is required to deliver short and/or long pulses at a 10 Hz firing rate with femto/pico-second accuracy sustained over fourteen 12-hour operator shifts to a common shared target chamber. The MEC-U system requires the integration of the control system with SLAC provided controls related to personnel safety, machine safety, precision timing, data analysis and visualization, amongst others. To meet these needs along with the system's reliability, availability, and maintainability requirements, LLNL is delivering an EPICS based control system leveraging proven SLAC technology. This paper presents the conceptual design of the DELPHI control system and the methods planned to ensure its successful commissioning and delivery to SLAC.

## INTRODUCTION

As part of the MEC-U project at SLAC, LLNL is developing DELPHI, a dual-mode rep-rated laser that provides short-pulse lasers to a target chamber (TCX) shared with the Linac Coherent Light Source (LCLS) X-ray Free Electron Laser (XFEL). The combination of these two systems, plus a third system being provided by the Laboratory for Laser Energetics (LLE) provides a leading experimental platform for high energy density (HED) plasmas under extreme environments [1]. The current project scope for DELPHI is a single Petawatt beamline. We have deferred or do not preclude requirements for a 2<sup>nd</sup> beamline, plus the addition of a long-pulse rep-rate laser (Fig. 1).

There's a total of 628 laser control points in the conceptual design, with an additional 213 control points in the utility systems. The system is expected to fire at a fixed enumerated rates up to 10Hz delivering optical pulses to the beam transport treaty point with a timing jitter of less than 200fs rms. Experiments are expected to run across multiple 12-hour shifts, which emphasizes the need for high reliability, availability, and maintainability (RAM). To help meet these requirements, LLNL is prioritizing the re-use of LCLS' EPICS based control system software, controllers, devices, and instrumentation in the DELPHI control system design.

<sup>†</sup> fishler2@llnl.gov




Figure 1: Conceptual Layout of DELPHI.

## ENGINEERING PROCESS

The original proposal for the DELPHI control system included the use of the National Ignition Facility (NIF) & Photon Science (PS) Photon Control System (PSC) framework. This technology is based on NI LabVIEW and relies heavily on NI cRIO hardware and other related technologies. It allows for the rapid development of scalable distributed control systems, including the addition of new device level controls including common control system features such as configuration user sets, configuration locking, etc. The introduction of this technology into SLAC's disparate control system ecosystem presented significant RAM challenges. SLAC, LLNL, and LLE agreed to take a unified approach to the MEC-U control system and adopted LCLS' EPICS baseline. To address the lack of in-house expertise with EPICS and LCLS hardware, we placed the engineering agency Cosylab, on contract. Cosylab is an active member of the SLAC control system engineering team and has been responsible for the development and deployment of several of their systems. Given their level of knowledge of EPICS based control systems and SLAC technology, they have been given responsibility for the hardware detailed design and controller/PLC level software engineering through LLNL's final design review (FDR). Part of their mandate is to work with SLAC to identify hardware already in production use to avoid introducing new or unique hardware as part of the DELPHI system. This reduces the overall software and hardware engineering effort, eases maintenance by deploying familiar hardware, and helps with an overall hardware sparing strategy at MEC-U.

The DELPHI control system recently completed its conceptual design review (CDR) that included

# FAST WIRE SCANNER MOTION CONTROL SOFTWARE UPGRADE FOR LCLS-II \*

Z. Huang<sup>†</sup>, T. Thayer, N. Balakrishnan, J. Bong, M. Campell  
SLAC, Menlo Park, CA 94025, USA

## Abstract

LCLS-II is the first XFEL to be based on continuous-wave superconducting accelerator technology (CW-SCRF), with the X-ray pulses at repetition rates of up to 1 MHz. LCLS-II's wire scanner motion control is based on Aerotech Ensemble controller. The position feedback and the beam loss monitor readings during a wire scan aim to measure the beam profile. To meet the measurement requirements under both low and high beam repetition rates, we redesign the software program for EPICS IOC, Aerotech controller, and develop a new User Interface (UI) based on PyDM. This paper will describe the software development details and the software commissioning result under LCLS-II's production environment.

## INTRODUCTION

The Linac Coherent Light Source II (LCLS-II) is a new X-ray Free-electron laser (XFEL) facility located at SLAC in California. It is the first XFEL to be based on continuous-wave superconducting (SC) accelerator technology (CW-SCRF), enabling it to deliver X-ray laser pulses that are 10,000 times brighter than those produced by the normal conducting (NC) linac [1]. Wire scanners are one of the primary diagnostic tools to measure the transverse profile of the electron beam at LCLS-II. The fast wire scanner has 3 wires with x, y, and u axes, assembled on an interchangeable card mounted on a linear motor stage [2, 3]. Figure 1 shows the fast wire scanner in the LCLS-II.




Figure 1: LCLS-II Fast Wire Scanner.

Controlled by an Aerotech Ensemble controller, the fast wire scanner can rapidly intercept the beam while providing high-resolution beam profile on-the-fly. To reconstruct the transverse beam profile at different electron bunch repetition rates, which ranges from 1Hz to 1 MHz, the wire scanner must execute a varied range of scan profiles. Meanwhile, both NC linac and SC linac utilize the emittance diagnostic section in Beam Transport Hall(BTH) leading up to the Hard X-ray and Soft X-ray beamlines in the undulator hall. Therefore, it is essential for the fast wire scanner controls software in this region to be compatible with both the facilities.

To meet these new requirements, the motion control software program in EPICS and Aerotech controllers have been redesigned. In this paper, at first, we will define the design requirements in detail. The software development details, including the software architecture and design specifics, will be described in this paper. Additionally, the commissioning and field performance test results will be presented. Finally, ongoing work and future plans will be introduced.

## MOTION CONTROL SOFTWARE DESIGN

### Design Requirements:

In the existing SLAC linac tunnel, the LCLS-II system is equipped with 20 fast wire scanners that span from the RF Gun to the Beam-Transfer Hall (BTH). The schematic of LCLS NC and SC beamlines are depicted in Fig. 2, which showcases 6 out of the 20 newly installed fast wire scanners. NC and SC linac have the capability to deliver beam to the Hard X-ray beamline and the Soft X-ray beamline. Based on the desired beam path of NC and SC linac, the shared beamline sections of the emittance diagnostic region in the BTH receive different bunch patterns. Consequently, it is essential for the wire scanners in this shared beamline to be capable of automatically changing their scan modes based on the facility mode. To facilitate testing and debugging of software, a “local mode” is introduced that can read local simulated beam parameters in addition to the software having the ability to read global operating beam parameters.

Wire scanners are crucial for scanning beams with desired velocity to gather enough measurement points to obtain a beam profile. The wire scanner's desired moving velocity is calculated as [4, 5]:

$$v_{desired} = \frac{f_{rep} \times r_{wire}}{N_{desired}} mm/s \quad (1)$$

where  $r_{wire}$  is the scan range and  $N_{desired}$  is the desired number of measurement points for each wire. Since the SC beam repetition rate can be continuously adjusted from 1 Hz

\* Work supported by U.S. Department of Energy under contract number DE- AC02-76SF00515

<sup>†</sup> zyhuang@slac.stanford.edu

# SLAC ATCA SCOPE - UPGRADING THE EPICS SUPPORT PACKAGE\*

D. Alnajjar<sup>†</sup>, M. P. Donadio, K. Kim, L. Ruckman  
 SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

The SLAC ATCA Scope, a 4-channel dual oscilloscope, has an EPICS support package that runs on top of SLAC’s Common Platform software and firmware and communicates with several high-performance systems in LCLS running on the 7-slot Advanced Telecommunications Computing Architecture (ATCA) crate. The software was completely refactored to improve the usability for EPICS IOC engineers. Once linked with an EPICS IOC, it initializes the oscilloscope hardware and instantiates the upper software stack providing a set of PVs to control the hardware and to operate the oscilloscope. The exported PVs provide a seamless means to configure triggers and obtain data acquisitions similar to a real oscilloscope. The ATCA scope probes are configured dynamically by the user to probe up to four inputs of the ATCA ADC daughter cards. The EPICS support package automatically manages available ATCA carrier board DRAM resources based on the number of samples requested by the user, allowing acquisitions of up to 8 GBytes per trigger. The user can also specify a desired sampling rate, and the ATCA Scope will estimate the nearest possible sampling rate using the current sampling frequency, and perform downsampling to try to match that rate. Adding the EPICS module to an IOC is simple and straightforward. The ATCA Scope support package works for all high-performance systems that have the oscilloscope common hardware implemented in its FPGAs. Generic interfaces developed in PyDM are also provided to the user to control the oscilloscope and enrich the user’s seamless overall experience.

## INTRODUCTION

The 7-slot Advanced Telecommunications Computing Architecture (ATCA) crate is used for numerous high-performance systems (HPS) at SLAC, such as the bunch charge monitor [1], bunch length monitor [1], beam position monitor [2], low-level radio frequency [3], machine protection system [4], timing system, and a few others. Each slot takes an advanced mezzanine card (AMC) generic carrier board equipped with an FPGA and two double-wide, full-height AMB bays. A variety of AMC cards (commercial and custom) can be plugged into the carrier. Some are application-specific and others are generic [5]. The most common is called the General Purpose ADC/DAC AMC Card, which contains DACs and ADCs. Other examples of AMCs designed at SLAC can be found in [4].

Having FPGA-level visibility in most of the stream-based applications mentioned earlier is always practical for obtaining raw data, debugging, and analysis. For this purpose, a

standard FPGA firmware component to receive streams, encapsulate them, and send them upstream was developed and integrated into all high-performance systems. This standard firmware component requires accompanying easy-to-use software resembling an oscilloscope as much as possible. With that in mind, the SLAC ATCA Scope EPICS [6] support package was completely refactored to improve the usability for SLAC’s EPICS IOC engineers. In this paper, we will discuss the upgraded EPICS support package.




Figure 1: High-performance system SW/HW overview.

## HPS OVERVIEW

A general overview of the high-performance systems is shown in Fig. 1. Each AMC card is linked to a Scope component in the firmware, where most streams passing in and out of the AMC are fed into the Scope module, and data is stored in the DRAM. Communication between the server and the ATCA is established through the common platform firmware and software [7]. The common platform firmware and software also provide means to read data from the DRAM and send it upstream. The Scopes can be triggered using multiple trigger sources and rates. The Scope software is linked as

\* Work supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515

<sup>†</sup> dnajjar@slac.stanford.edu

# DESIGN AND IMPLEMENTATION OF THE LCLS-II MACHINE PROTECTION SYSTEM\*

J. A. Mock<sup>†</sup>, Z. Domke, R. T. Herbst, P. Krejcik, L. Ruckman, L. Sapozhnikov  
SLAC National Accelerator Laboratory, Menlo Park, United States

## Abstract

The linear accelerator complex at the SLAC National Accelerator Laboratory has been upgraded to include LCLS-II, a new linac capable of producing beam power as high as several hundred kW with CW beam rates up to 1 MHz while maintaining existing capabilities from the copper machine. Because of these high-power beams, a new Machine Protection System with a latency of less than 100  $\mu$ s was designed and installed to prevent damage to the machine when a fault or beam loss is detected. The new LCLS-II MPS must work in parallel with the existing MPS from the respective sources all the way through the user hutches to provide a mechanism to reduce the beam rate or shut down operation in a beam line without impacting the neighboring beam line when a fault condition is detected. Because either beam line can use either accelerator as its source and each accelerator has different operating requirements, great care was taken in the overall system design to ensure the necessary operation can be achieved with a seamless experience for the accelerator operators. The overall system design of the LCLS-II MPS software including the ability to interact with the existing systems and the tools developed for the control room to provide the user operation experience is described.

## INTRODUCTION

The Linear Accelerator Facility at the SLAC National Accelerator Laboratory (SLAC) has seen a significant upgrade to the Linac Coherent Light Source (LCLS) facility with the installation of LCLS-II [1]. This new accelerator and undulator complex was designed to increase capacity at SLAC for photon science using the free electron lasers. A high level diagram of the new facility is shown in Fig. 1. Because it uses cryogenic RF cavities, the LCLS-II accelerator is capable of delivering electrons with a minimum bunch spacing of 1  $\mu$ s and maximum electron energy of 4 GeV, leading to a maximum overall beam power of 120 kW. The high rate and high power potential of these beams requires safety systems capable of reacting to beam events such that operations remain safe. The personnel protection and beam containment systems (PPS and BCS) are safety systems designed to protect people from radiation-based hazards, and the Machine Protection System (MPS) is designed to protect the accelerator from damaging itself. The MPS is not a credited safety system, so it is designed to be more agile to allow a variety of operating conditions while still protecting the accelerator.

\* SLAC is supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-76SF00515

<sup>†</sup> jmock@slac.stanford.edu




Figure 1: A high level map of the new LCLS-II facility. The MPS has a mitigation device to inhibit beam or reduce beam rate to each of the destinations shown. Additionally, both the new LCLS-II superconducting accelerator and legacy normal conducting accelerator can deliver to the undulator complex, though not yet at the same time.

The scope of the MPS is confined exclusively to shutting off the electron beam when a fault condition occurs that can potentially damage beam line hardware. Other systems that protect high power devices such as power supplies, RF power sources, vacuum systems, or cryogenic systems are handled separately as equipment protection. A key driving parameter of the MPS is the maximum allowable time interval in which the beam must be shut off before damage can occur. The MPS requirement for the original LCLS dictated that the electron beam be shut off within one beam pulse at the full repetition rate of 120 Hz. This is not possible in LCLS-II where the minimum bunch spacing is only 1  $\mu$ s and propagation delay for a signal in a cable from one end of the accelerator to the other can be as long as 20  $\mu$ s, not including additional processing delays incurred from electronics. The MPS baseline beam shutoff time, defined as the time between detection of fault and loss of photo-current, is required to not exceed 100  $\mu$ s to avoid catastrophic damage to the beam line, though in principle the MPS physics requirement is as low as reasonable achievable. Not every fault condition requires the fast shutoff time of 100  $\mu$ s. For example, a slow change in some state, such as a temperature rising, allows ample time for the control system to warn of the impending change. Therefore, MPS responses to prompt events such as beam loss mitigate within the fast response window, and other, slower events and more complicated logic process within a 360 Hz window, equal to the processing time of the LCLS-I MPS.

As shown in Fig. 1, the LCLS-II accelerator can deliver to some combination of an injector diagnostic line, one of two undulator beam lines, or through the linac to a high power dump in the SLAC Beam Switch Yard (BSY). The default destination for the electron beam is this high power dump. Pulsed kicker magnets are used to kick the beam into one of the other beam lines, which means if the pulser does not kick, the beam will need to travel to the BSY beam dump. Therefore, the MPS uses this high power dump as its pri-



# SLAC LINAC MODE MANAGER INTERFACE\*

T. Summers<sup>†</sup>, C. Bianchini Mattison, M. Gibbs, T. J. Kabana, P. Krejcik, J. A Mock  
 SLAC National Accelerator Laboratory, Menlo Park, CA, USA

## Abstract

With the successful commissioning of the new superconducting (SC) Linac, the Linac Coherent Light Source (LCLS) now has the capability of interleaving beams from either the normal conducting (NC) Linac or the SC Linac to two different destinations, the soft (SXR) and hard (HXR) X-ray undulator beamlines. A mode manager user interface has been created to manage the beamline configuration to transport beam pulses to multiple destinations, which include the numerous intermediate tune-up dumps and safety dumps between the injectors and the final beam dumps. The mode manager interfaces with the timing system which controls the bunch patterns to the various locations, and the machine protection system which prevents excess beam power from being sent to the wrong destination. This paper describes the implementation method for handling the mode switching, as well as the operator user interface which allows users to graphically select the desired beam paths.

## INTRODUCTION

The Linear Accelerator Facility (LAF) at SLAC National Accelerator Laboratory has recently met the LCLS-II project [1] key performance parameter milestones, replacing the furthest upstream 1 km section of the original SLAC linac with a superconducting electron linac capable of beam rates up to 1 MHz. This new SC Linac will complement the 120 Hz normal conducting electron linac accelerator and undulator beamlines which have been serving the LCLS user community since 2009.




Figure 1: Schematic view of the SLAC LCLS SC and NC linacs, transport lines and switchyard, undulator lines and experiment halls.

In addition to the superconducting linac itself, the LCLS-II project included a bypass line to transport the accelerator beam past the FACET-II and LCLS-I NC Linac, as well as a reconfiguration of the beam switchyard (BSY) area where beam from either SC or NC linac can be directed to either SXR or HXR undulator lines. A schematic of the facility can be seen in Fig. 1.

\* LCLS IS SUPPORTED BY THE U.S. DEPARTMENT OF ENERGY, OFFICE OF SCIENCE, UNDER CONTRACT DE-AC02-76SF00515  
<sup>†</sup> tsummers@slac.stanford.edu

The LCLS facility is designed for beam from each linac to be sent to either or both undulator lines, but not beam from both linacs to the same undulator line. We define the operating ‘mode’ as the combination of beamline components that define the path of possible beam from electron gun to final destination (E.g., dump). The modes are named starting with the source (E.g., NC or SC) and a unique number. Table 1 lists a sample of the most commonly used modes.

Table 1: Sample SC and NC Beam Path Definitions

Mode	Name	Description
SC10	SC Laser	Full power, 1 MHz
SC13	SC DIAG0 Line	120 Hz max
SC14	SC BSY Dump	120 kW, 100 kHz
SC17	SC HXR Dump	120 kW, 100 kHz
SC18	SC SXR Dump	120 kW, 100 kHz
NC0	NC Laser	Full charge, 120 Hz
NC7	NC HXR Tuning	Full charge, 10 Hz
NC8	NC HXR Dump	Full charge, 120 Hz
NC11	NC SXR Tuning	Full charge, 10 Hz
NC12	NC SXR Dump	Full charge, 120 Hz

The various beam paths from the NC and SC electron guns to the final HXR and SXR beam dumps, and all the various intermediate dumps can be seen in Fig. 2. This image also identifies the kickers, magnets, and stoppers that are needed to define each path. For scale, a geographical layout of the facility can be seen in Fig. 3.

Several of these stoppers are Personnel Protection System (PPS) devices and can only be inserted or retracted with a physical button in the accelerator control room (ACR). Most of the devices on this map are also inputs into the LCLS SC and NC Machine Protection Systems (MPS).




Figure 2: Overview of beam paths from SC and NC electron guns to their intermediate or final destination.

The NC Linac MPS system is designed to disable the beam permit if magnets and beamline components are not in the correct state, or reduce the allowed beam rate if specific equipment is inserted, for example maximum 10 Hz if the NC beam destination is one of the tune up dumps. The

# THE LCLS-II EXPERIMENT CONTROL SYSTEM PREEMPTIVE MACHINE PROTECTION SYSTEM \*

A. Wallace<sup>†</sup>, SLAC National Accelerator Laboratory, Menlo Park CA, USA

## Abstract

The LCLS-II Preemptive Machine Protection System (PMPS) safeguards diagnostics, optics, beam-shaping components and experiment apparatuses from damage due to excess XFEL average power or single shots. The dynamic nature of these systems requires a somewhat novel approach to a machine protection system design, relying more heavily on preemptive interlocks and automation to avoid mismatches between device states and beam parameters. This is in contrast to standard reactive machine protection systems deployed at other laboratories. Safe beam parameter sets are determined from the combination of all integrated devices using a hierarchical arrangement and all state changes are held until beam conditions are assured to be safe. This machine protection system design utilizes the Beckhoff industrial controls platform and EtherCAT. It is woven into the LCLS subsystem controllers as a code library and standardized hardware interface.

## INTRODUCTION

Machine Protection Systems (MPS) are ubiquitous in high-energy physics machines around the world. Typically this term is reserved for the subsystem which protects the physical machine hardware from the very phenomenon it produces. Machine protection in the context of the LCLS Experiment Control Systems (ECS) is protection from damage by the x-ray photon beam (XFEL). It is distinguished from other terms such as Personnel Protection Systems (PPS), and Equipment Protection Systems (EPS). EPS typically refers to interlock logic for actuators, or vacuum equipment to protect those devices from damaging themselves.

LCLS consists of two control system domains, characterized by different operating requirements. The electron beam production aspect consists of accelerator and undulator technology. Physical access in this domain is more rare and the configuration is generally more stable. Once produced, the XFEL photon x-ray beam proceeds through a front-end where it is measured, attenuated, steered, shaped and focused by x-ray photon devices. After the front-end the beam is delivered to experiment interaction points where it is typically used as a probe. This part of LCLS is more accessible, and requires more in the way of experimental reconfiguration. This dynamic area of the LCLS thus lends itself to a different control system architecture - that of the ECS.

The ECS of LCLS-II is a SCADA type system consisting of Beckhoff PLCs which integrate mechatronics and vacuum system components using a common framework

of hardware templates and IEC61131-3 Structured Text libraries [1]. These libraries provide standardized EPS and MPS functionality as well as a uniform EPICS interface. The ECS machine protection functionality is completely implemented within this PLC framework.

## REQUIREMENTS

MPS are typically designed to rate-limit or zero-rate beam as fast as possible in reaction to anomalous or fault conditions. Some MPS for pulsed machines are designed to react to an error by turning beam off within the time between pulse bunches (macro-pulses) or even individual pulses [2]. This approach is not technically feasible with the LCLS-II machine parameters.

LCLS-I uses the SLAC Normal Conducting (NC) linear accelerator (linac) with a maximum repetition rate of 120 Hz, leaving ample time between pulses for fault detection and mitigation. LCLS-II, in contrast, has a nominal repetition rate of just less than 1 MHz, while the electronic signal propagation time along the length of the linac along is at least 20  $\mu$ s.




Figure 1: The increase in brightness and average power from LCLS-I to LCLS-II.

Inter-pulse mitigation is not necessarily required for the ECS MPS. Instead, the primary challenge is preventing even a single pulse of the wrong wavelength and energy from entering the experiment system domain, see Fig. 1. Mismatches between optic coatings or diagnostic elements and XFEL pulse wavelength can cause damage. Additionally, with the increased repetition rate, average power delivery can cause harm on a longer timescale. Finally, an additional

\* Work supported by U.S. D.O.E. Contract DE-AC02-76SF00515

<sup>†</sup> awallace@slac.stanford.edu

# THE DESIGN OF A PYTHON-BASED LIVE & ARCHIVED DATA VIEWER

Y. Yazar\*, T. Summers, J. Bellister, Z. Domke, F. Osman  
SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

A new open-source PyQT-based archive viewer application is under development at SLAC National Accelerator Laboratory (see Fig. 1). The viewer's main purpose is to visualize both live values and historical Process Variable (PV) data retrieved from the EPICS Archive Appliances. It is designed as both a stand-alone application and to be easily launched from widgets on PyDM operator interfaces. In addition to providing standard configurability for things like traces, formulas, style and data exporting, it provides post-processing capabilities for filtering and curve fitting. The development teaming is current working on a release which will support standard enumerated and analog data types as well as waveforms. Extension of this support to EPICS7 normative data types such as NTTTable and NTNDArray is to follow.

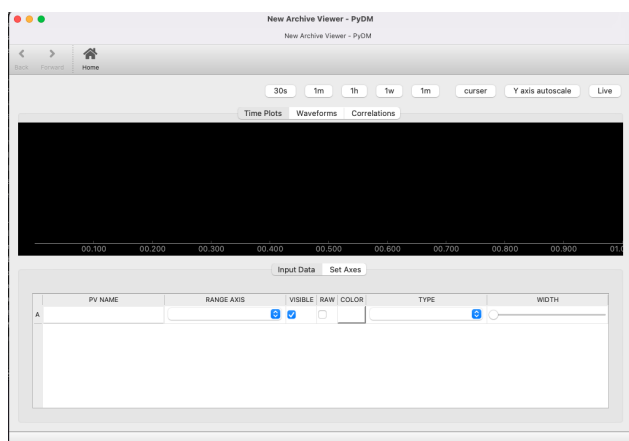


Figure 1: New Archive Viewer Under Development.

## INTRODUCTION

SLAC has a need to replace its current Java-based archive-viewer. This need was brought about through a combination of a need for new functionality to view archived and live data along with a degradation of the current archive viewer over the last decade or so.

### Problems

One of the main impetus to replace the current java-based archive viewer (see Fig. 2), which was developed around 2005 at SNS, is that it is no longer maintained at SLAC, this has led to a loss in functionality. In addition to the degradation, there is a controls department desire to coalesce around a smaller set of technologies, in this case part of its ongoing effort to move many of the labs older Java-based

\* email: yazar@slac.stanford.edu

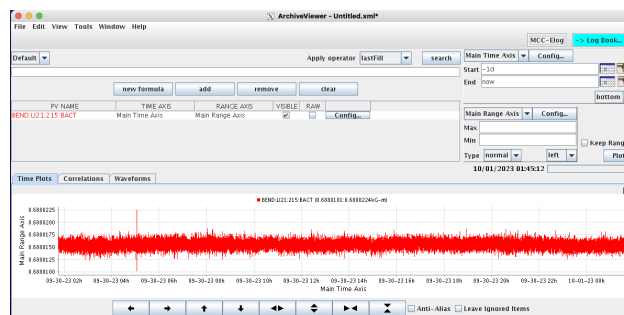


Figure 2: Current Java-Based Archive Viewer.

application to python-based ones. In addition to these two main points, there are a slew of additional challenges facing the lab with its current archive viewer:

- New functionality needed of the archive viewer going forward
- A slew of applications have been developed/used at SLAC to access archived data through the Archiver Appliance API, and although the objective is not to replace them all, there is a desire to include enough functionality to reduce the total number of applications being maintained by the lab
- The current viewer is less than intuitive with new users needing ample instructions (see Fig. 3 for some examples of this)
- The code is not well documented, making modification to the existing code challenging

These collective points are behind the decision to build a new archive viewer.




Figure 3: Current Viewers less than intuitive design.

## REQUIREMENTS

The requirements for the new archive viewer can be succinctly summarized as follows:

- Replicate the functionality of the current Archive Viewer
- Provide enhanced post-processing capabilities
- Include EPICS 7 data type integration

# LONGITUDINAL FEEDBACK FOR THE LCLS-II SUPERCONDUCTING LINEAR ACCELERATOR AT SLAC\*

C. Zimmer†, D. Chabot, W. Colucho, Y. Ding, J. Nelson, SLAC National Accelerator Laboratory, Menlo Park, CA, US

## Abstract

SLAC recently commissioned a new continuous-wave, MHz repetition-rate Superconducting (SC) Linear Accelerator (Linac). This accelerator can produce a 4 GeV electron beam that drives two dedicated Hard and Soft X-ray Undulator lines as part of the Linac Coherent Light Source (LCLS) Free Electron Laser. A new Python-based longitudinal feedback is used to control the electron beam energy and bunch length along the accelerator. This feedback was written to be simple, easily maintainable and easily portable for use on other accelerators or systems as a general-purpose feedback with minimal dependencies. Design and operational results of the feedback will be discussed, along with the Graphical User Interfaces built using Python Display Manager (PyDM).

## SUPERCONDUCTING LINAC LAYOUT

The Superconducting Linear Accelerator is divided into four main accelerating regions defined in terms of the included cryomodules (CMs): L0B (CM01), L1B (CM02-3), L2B (CM04-15) and L3B (CM16-35). Each cryomodule contains eight individual 9-cell radiofrequency (RF) cavities.

Figure 1 illustrates these regions by showing an expanded view of the original 3 km SLAC Linac tunnel. At the end of each of these regions is a bend which can be used to measure the beam energy. After L1B and L2B, there are dedicated bunch compressor chicanes (BC1B and BC2B respectively) which also have bunch length monitors. This makes for a total of six parameters we want to precisely control – electron beam energy at four locations and electron bunch length at two locations.

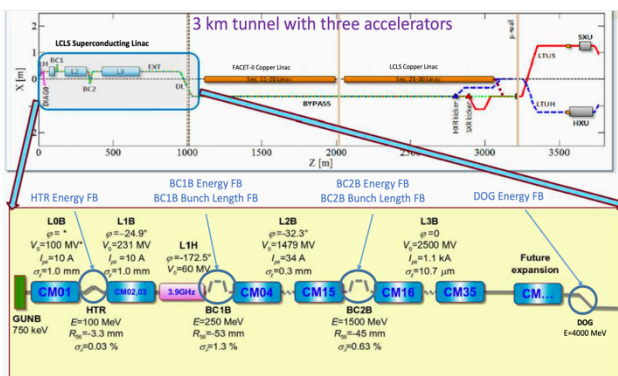


Figure 1: Layout of SC linac.

\* Work supported by US DOE under grant No. DE-AC02-76SF00515  
 † zimmerc@slac.stanford.edu

## RF ABSTRACTION LAYER

The SC Linac currently has 296 superconducting RF cavities (thirty-five 1.3 GHz accelerating CMs and two 3.9 GHz harmonic linearizer CMs). Each of these cavities is powered by an individual solid-state amplifier with its own amplitude and phase controls. Trying to control the overall electron beam energy and longitudinal profile by adjusting each of these individual cavities would be next to impossible.

For this reason, the RF Abstraction Layer (RFAL) was developed as a tool to orchestrate the management of all cavities in order to achieve the desired energy and energy spread (chirp) in each region. It provides a way of abstracting away the mathematical details of achieving these parameters using large collections of RF cavities, making it easy and efficient for an end-user (or feedback) to control these parameters. The RFAL is a dedicated EPICS soft IOC that seamlessly integrates Python control code (using pyDevSup [1]) to perform real-time calculation, vector plotting and distribution of settings to individual RF cavities at up to ~20 Hz.

### RFAL Working Principles

Continuous-wave SC RF cavities experience Lorentz Force Detuning that changes with RF field amplitude. If the amplitude is kept constant, the corresponding Lorentz Force Detuning effect is static [2] and can more easily be compensated for. This motivates a strong desire to keep the cavity amplitudes fixed. The RFAL therefore only changes cavity phases to control the electron beam energy and chirp.

A prerequisite for calculating and distributing phases is that the RFAL needs to know which cavities should be used for pure acceleration, imparting energy spread, or maintaining average energy. Four cavity roles are specified in Table 1.

Table 1: RF Abstraction Layer Cavity Roles

Cavity Role	Purpose	RFAL Sets Phase
NOTA	Acceleration only	No
Chirp Only	Add energy spread and increase energy	Yes
FB -	Maintain total energy, preserve chirp	Yes
FB +	Maintain total energy, preserve chirp	Yes

# TEMPERATURE CONTROL OF CRYSTAL OPTICS FOR ULTRAHIGH-RESOLUTION APPLICATIONS\*

Kazimierz J. Gofron†

Spallation Neutron Source, Oak Ridge National Laboratory, Oak Ridge, USA

David Scott Coburn, Alexey Suvorov, Yong Q. Cai‡

National Synchrotron Light Source II, Brookhaven National Laboratory, Upton, USA

## Abstract

The temperature control of crystal optics is critical for ultrahigh-resolution applications such as those used in meV-resolved Inelastic X-ray Scattering (IXS). Due to the low count rate and long acquisition time of these experiments, for 1-meV energy resolution at ~10 keV, the absolute temperature stability of the crystal optics must be maintained below 4 mK for days to ensure the required stability of the lattice constant, thereby ensuring the energy stability of the optics. Furthermore, the temperature control with sub-mK resolution enables setting the absolute temperature of the individual crystal, making it possible to align the reflection energy of each crystal's rocking curve in sub-meV precision thereby maximizing the combined efficiency of the crystal optics.

## INTRODUCTION

The crystal optics employed in X-ray beamlines at synchrotron facilities utilize perfect single crystals such as silicon and diamond. Their performance in resolving power is governed by the perfection of their crystal lattice and dynamic X-ray diffraction:  $\Delta E/E = -\Delta d/d$ , where  $\Delta d$  is the lattice (d) variation across the crystal for the chosen reflection. When subjected to varying temperature environments, the changing temperature introduces changes in the lattice constant given by  $\Delta d/d = \alpha \Delta T$ , where  $\alpha$  is the thermal expansion coefficient of the crystal. For Si ( $\alpha = 2.6 \times 10^{-6} \text{ K}^{-1}$ ), therefore, a 40 mK temperature variation will shift the diffraction energy by ~1 meV for 10 keV X-rays. To achieve 1 meV resolution in a stable condition, a 4 mK temperature stability or better, corresponding to 1/10 of the energy resolution, is required.

In this article, we report the details of an EPICS temperature control system using PT1000 sensors, Keithley 3706A 7.5 digits sensor scanner, and Wiener MPOD LV power supply for the analyzer optics of the IXS beamline 10-ID at NSLS-II [1]. The crystal optics part of the analyzer is depicted in Fig. 1 schematically, where the dispersing (D) crystal has a total length of 1.2 m to achieve the required angular acceptance and is made up of 6 highly asymmetrically cut silicon crystals of 200 mm in length each. The lattice homogeneity and temperature stability of these D-crystals are the most critical for the performance

of the analyzer. The temperature control system was applied to them. We were able to achieve absolute temperature stability below 1 mK and sub-meV energy alignment for these D-crystals. The EPICS ePID [2] record was used for the control of the power supplies based on the PT1000 sensor input that was read with a 7.5-digit resolution from the Keithley 3706A scanner. The system enhances the performance of the meV-resolved IXS spectrometer achieving a 1.4 meV total energy resolution with unprecedented spectral sharpness and stability for the studies of atomic dynamics in a broad range of materials.

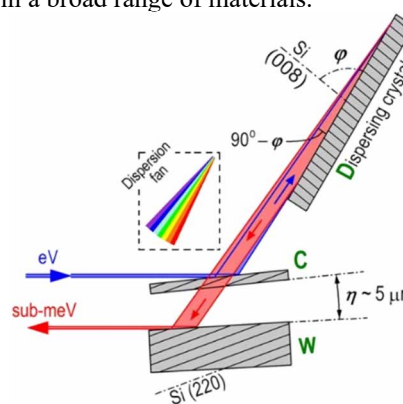


Figure 1: Schematic layout of the analyzer crystal optics employed at the meV-IXS spectrometer of the IXS 10-ID beamline at NSLS-II.

## CONTROLLERS

Crystal temperature controls consist of a Keithley 3706A [Fig. 2 (top)] 7.5-digit resolution scanner. The Keithley Instruments Digital Multi-Meter (DMM) 3706A 6 slot system switch utilizes a 3724 dual 1/30 FET Card (Auto CJC with 3724-ST). Such a system allows us to measure up to 30 PT1000 sensors per 3724-ST card by scanning selected inputs. The 3706A takes input from the in-vacuum PT1000 temperature sensors attached to the copper base of each D-crystal of the analyzer optics. The copper base of each D-crystal is fitted with three PT1000 sensors, one at each end, and one near the heating element which is used for control. The control driver uses EPICS State Notation Language sequencer [3] which reads individual PT1000 inputs. The values read by the Keithley 3706A DMM are passed to EPICS PVs [2]. The Temperature PVs are processed by ePID EPICS PVs and provided as feedback to adjust the power output of the MPOD [Fig. 2 (bottom)] power supply channel providing heat to each D-crystal.

\* Work supported by U.S. Department of Energy, Office of Science, Scientific User Facilities Division under Contract No. DE-AC05-00OR22725, and Office of Basic Energy Sciences under Contract No. DE-SC0012704.

† gofronkj@ornl.gov

‡ cai@bnl.gov

# CONTROL SYSTEMS DESIGN FOR STS ACCELERATOR\*

J. Yan<sup>†</sup>, S. M. Hartman, K. Kasemir  
Oak Ridge National Laboratory, Oak Ridge, USA

## Abstract

STS accelerator systems will build a Ring-to-Second-Target (RTST) transport beamline from the present Ring to Target Beam Transport (RTBT) to the second target. The integrated Control Systems (ICS) will provide remote control, monitoring, OPI, alarms, and archivers for the accelerator systems, such as magnets power supply, vacuum devices, and beam instrumentation. The ICS will upgrade the existing Linac LLRF controls to allow independent operation of the FTS and STS and support different power levels of the FTS and STS proton beam. The ICS accelerator controls are in the phase of preliminary design for the control systems of magnet power supply, vacuum, LLRF, Timing, Machine Protection System (MPS), and computing and machine network. The accelerator control systems build upon the existing SNS Machine Control systems, use the SNS standard hardware and software, and take full advantage of the performance gains delivered by the PPU Project and SNS.

## INTRODUCTION

The Second Target Station (STS) Project includes the comprehensive design, construction, installation, and commissioning of the essential facilities and equipment aimed at establishing a cutting-edge source of cold neutrons with unprecedented peak brightness at the Spallation Neutron Source (SNS). The project leverages the capacity of the existing SNS accelerator, accumulator ring, and infrastructure and takes full advantage of the performance gains delivered by the Proton Power Upgrade (PPU) Project [1]. The PPU will double the SNS proton beam power from 1.4 MW to 2.8 MW. This increase is achieved by integrating seven additional superconducting cryomodules into the existing linear accelerator (Linac), resulting in a 30% rise in beam energy combined with a 50% rise in beam current. Operating at a frequency of 60 Hz, the SNS accelerator, when combined with the STS, will supply 45 pulses per second to the First Target Station (FTS) and 15 pulses per second at a frequency of 15 Hz to the STS. STS Accelerator Systems will build a new 231.9-meter-long Ring to Second Target (RTST) beamline to transport protons from an extraction point in the existing Ring to Target Beam Transport (RTBT) line to the Second Target. The RTST system integrates diverse components, including vacuum systems, magnets, power supplies, beam instrumentation and diagnostics, and utilities, all working in tandem to support its operation. This paper describes the integrated controls of the STS accelerator systems.

## STS ACCELERATOR CONTROLS

STS Accelerator Control systems build upon the existing Experimental Physics and Industrial Control System (EPICS) SNS Control systems, effectively utilizing SNS standard hardware and software components where applicable to govern the operation of the RTST [2]. The purview of the Accelerator Controls encompasses an array of vital aspects, including vacuum controls, magnet power supply controls, the Machine Protection System (MPS), the Run Permit system (RPS), the Timing system, core software, machine networking and computing, and upgrades to the Low-level Radio Frequency (LLRF) control for the pre-existing Linac.

### *Vacuum Controls*

The RTST vacuum system will extend the vacuum configuration from the existing RTBT to the STS. This comprehensive system comprises various crucial elements, such as vacuum assemblies, sensors, pumps, and associated instrumentation, which include valves, gauges, gauge controllers, and pump controllers. To effectively regulate and monitor the vacuum equipment within the RTST, a dedicated vacuum control system has been devised. The architecture of this control system is modelled on the well-established RTBT vacuum control system, which is grounded in the utilization of Allen-Bradley (AB) Programmable Logic Controllers (PLCs) and EPICS Soft Input/Output Controllers (IOCs) [3]. Specifically, an Allen-Bradley ControlLogix PLC will be deployed to oversee gauge and pump operations, control vacuum valves, and provide interlocking to the Machine Protection System (MPS). Communication between the PLC and the EPICS IOCs will be facilitated via EtherNet/IP [4], while Linux workstations will serve as the operator interface.

### *Magnet Power Supply Controls*

The 15 Hz beam will be extracted from the RTBT, directed to the RTST using pulsed dipole magnets, and then transported to the second target with quadrupole, dipole, and corrector magnets. Both DC magnet power supplies and pulsed dipole power supplies are utilized to provide the necessary current for all these magnets. The control interface for the DC magnet power supplies is established through Ethernet or serial communication protocols. The EPICS control software will be developed to control all power supplies. Furthermore, an interlock system will be implemented between the power supply and the Machine Protection System (MPS).

<sup>†</sup> yanjl@ornl.gov

# EXPLORATORY DATA ANALYSIS ON THE RHIC CRYOGENICS SYSTEM COMPRESSOR DATASET\*

Y. Gao<sup>†</sup>, K. A. Brown, R. Michnoff, L. Nguyen, B. van Kuik, A. Zarcone  
Brookhaven National Laboratory, Upton, NY, USA  
A. Tran, Facility for Rare Isotope Beams, East Lansing, MI, USA

## Abstract

The Relativistic Heavy Ion Collider (RHIC) Cryogenic Refrigerator System is the cryogenic heart that allows RHIC superconducting magnets to operate. Parts of the refrigerator are two stages of compression composed of ten first and five second-stage compressors. Compressors are critical for operations. When a compressor faults, it can impact RHIC beam operations if a spare compressor is not brought on-line as soon as possible. The potential of applying machine learning to detect compressor problems before a fault occurs would greatly enhance Cryo operations, allowing an operator to switch to a spare compressor before a running compressor fails, minimizing impacts on RHIC operations. In this work, various data analysis results on historical compressor data are presented. It demonstrates an autoencoder-based method, which can catch early signs of compressor trips so that advance notices can be sent for the operators to take action.

## INTRODUCTION

The Relativistic Heavy Ion Collider (RHIC) Cryogenic Compressor System at Brookhaven National Laboratory (BNL) is comprised of ten first-stage compressors, four second-stage compressors, and a redundant compressor that can function as a first, second, or full-stage compressor, as shown in Fig. 1. Initially, the compressors were controlled through 120VAC relay logic with minimal data available for Operations and only a local enunciator to indicate faults during unscheduled shutdowns of a compressor. Since 2014, the compressor controls have been upgraded to a more modern 24VDC PLC-controlled system. To date, six first-stage and all second-stage compressors have been upgraded. A part of the modernization is the increased availability of data for operators to monitor and track the health of each running compressor. The total data acquired is 163 variables for a first-stage compressor and 100 variables for a second-stage compressor, the result of one less motor-compressor set. All data are logged at a one-point-per-second rate. The data focus of this study is on a first-stage compressor, which comprises 27 analog variables, i.e., 19 temperature sensors (names starting with “TT”), 5 pressure transducers (names starting with “PT”), 2 horsepower monitors (M77, M79), and an oil level probe. The oil level probe parameter is omitted in this study since it is not as informative as the other parameters.

\* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.

<sup>†</sup> ygao@bnl.gov




Figure 1: RHIC cryogenic compressor system overview. It comprises ten first-stage compressors, four second-stage compressors, and a redundant compressor.

The 16 to 27 variables per compressor are just a small fraction of the 10,000+ data points an operator must monitor to understand the health of the Cryogenic system. Manually monitoring the system takes valuable time from operators, and it takes much more resources to recover from a system failure than to detect and prevent it beforehand. In this work, we present the initial results of analyzing historical compressor data to determine if developing faults with a compressor can be detected early enough to minimize the impact on operations and to narrow the cause of faults to facilitate quicker repairs, increasing the run-time availability of each compressor.

## DATASET AND METHODS

The datasets are collected from the upgraded first and second-stage compressors. Compressor First Stage 1 (FS1) is chosen for analysis because it is the only upgraded compressor with a documented fault during the 493 days of recorded data. The documented trip happened on Apr. 7th, 2022. So the training data were selected from Jan. 15th to Mar. 5th, 2022, and testing data were from Mar. 6th to Apr. 5th, 2022, to test if the algorithm can detect any early precursors. The data were acquired at 1 Hz.

In this work, we focus on analyzing 26 float-type variables, as discussed above. Those time series data are shown in Fig. 2. The “TT” sensors monitor different system parts’ temperatures, “PT” sensors monitor different parts’ pressures, “M77” and “M79” are the horsepower monitors for the two motors. A pictogram of the FS1 compressor with corresponding parameters is shown in Fig. 3.

In the first step, we applied time series K-means to cluster the datasets to better understand the data patterns. Next, we

# THE POINTING STABILIZATION ALGORITHM FOR THE COHERENT ELECTRON COOLING LASER TRANSPORT AT RHIC\*

L. K. Nguyen<sup>†</sup>, Collider-Accelerator Dept., Brookhaven National Laboratory, Upton, NY, USA

## Abstract

Coherent electron cooling (CeC) is a novel cooling technique being studied in the Relativistic Heavy Ion Collider (RHIC) as a candidate for strong hadron cooling in the Electron-Ion Collider (EIC). The electron beam used for cooling is generated by laser light illuminating a photocathode after that light has traveled approximately 40 m from the laser output. This propagation is facilitated by three independent optical tables that move relative to one another in response to changes in time of day, weather, and season. The alignment drifts induced by these environmental changes, if left uncorrected, eventually render the electron beam useless for cooling. They are therefore mitigated by an active "slow" pointing stabilization system found along the length of the transport, copied from the system that transversely stabilized the Low Energy RHIC electron Cooling (LEReC) laser beam during the 2020 and 2021 RHIC runs. However, the system-specific optical configuration and laser operating conditions of the CeC experiment required an adapted algorithm to address inadequate beam position data and achieve greater dynamic range. The resulting algorithm was successfully demonstrated during the 2022 run of the CeC experiment and will continue to stabilize the laser transport for the upcoming run. A summary of the algorithm is provided.

## INTRODUCTION

The luminosity demands of future colliders and the prospect of more productive runs for existing ones have made investigations into novel cooling techniques increasingly important. At Brookhaven National Laboratory (BNL), site of the future Electron-Ion Collider (EIC), coherent electron cooling (CeC) is being experimentally studied in the Relativistic Heavy Ion Collider (RHIC) for strong hadron cooling in the EIC. Although CeC employs many proven technologies and processes, its successful implementation still involves the overcoming of many challenges [1]. As a result, the CeC experiment is accompanied by several novel systems.

Among these is the CeC laser beam trajectory stabilization system, also known as the pointing stabilization system or the (transverse) position stabilization system, found along the laser transport. This system was developed and installed to help meet the demanding requirements regarding alignment between the electron beam and the ion beam, as the electron beam is generated by laser light striking a photocathode [1]. Like the system monitoring and stabilizing the laser beam along the laser transport of the Low Energy RHIC electron Cooler (LEReC) [2, 3], the CeC

version of the system is a "slow" feedback system controlled by two MATLAB scripts. However, several key differences with the LEReC scripts exist to address CeC's unique setup and laser operating conditions. Chief among these are the need for even greater dynamic range than the LEReC system, and the lack of adequate laser beam position data from the controls system for the purposes of the stabilization system.

## ENVIRONMENT AND SETUP

In the description that follows, an "Operations camera" is a camera that is part of the Controls System used by operators of CeC; these cameras were installed during the construction/commissioning of the cooler itself and are on a common timing system with CeC instrumentation. "Stabilization camera", on the other hand, refers to a camera (there are two in all) that is on a local network with the computer running the pointing stabilization scripts. In terms of reliability, this means that Controls System downtime does not affect the monitoring capability of the slow position stabilization system. However, a network connection is needed to make changes to the voltages of the piezo steering mirrors.




Figure 1: Aerial view of Interaction Region 2 (IR2) at RHIC, showing the approximate layout of the CeC laser transport.

Figure 1 shows an aerial view of the CeC laser transport. The laser beam is generated in the laser trailer, outside the RHIC tunnel. The first piezo steering mirror, controllable remotely in the Controls System, is located on an optical table here in the trailer. From the trailer, the laser beam travels through an evacuated pipe to reach the so-called relay table, just inside the tunnel. The first Operations camera and, slightly past it, the first stabilization camera is located here. Downstream of the cameras but still on the relay table is the second piezo steering mirror, also remotely controllable in the Controls System. After reflecting off this relay piezo mirror, the laser beam travels down another

\* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.

<sup>†</sup> l.nguyen@bnl.gov



# POSITION-BASED CONTINUOUS ENERGY SCAN STATUS AT MAX IV

Á. Freitas\*, N. Al-Habib, B. Bertrand, M. Eguiraun, I. Gorgisyan,  
A. Joubert, J. Lidón-Simón, M. Lindberg, C. Takahashi  
MAX IV Laboratory, Lund, Sweden

## Abstract

The traditional approach of step scanning in X-ray experiments is often inefficient and may increase the risk of sample radiation damage. In order to overcome these challenges, a new position-based continuous energy scanning system has been developed at MAX IV Laboratory. This system enables stable and repeatable measurements by continuously moving the motors during the scan. Triggers are accurately generated by hardware based on the motor encoder positions to ensure precise data acquisition. Prior to the scan, a list of positions is generated, and triggers are produced as each position is reached. The system uses TANGO and Sardana for control and a TriggerGate controller to calculate motor positions and configure the PandABox, which is the equipment in charge of trigger signal generation. The system is capable of scanning a single motor, such as a sample positioner, or a combined motion like a monochromator and undulator. In addition, the system can use the parametric trajectory mode of the IcePAP driver, which enables continuous scans of coupled axes with non-linear paths.

This paper presents the current status of the position-based continuous energy scanning system for BioMAX, FlexPES, and FinEstBeAMS beamlines at MAX IV and discusses its potential to enhance the efficiency and accuracy of data acquisition at beamline endstations.

Keywords: controls, software, hardware, experiment

## INTRODUCTION

In the domain of experimental science, the efficiency and speed of data acquisition have always been of extreme importance. A prior MAX IV work [1] tackled the challenges associated with conventional step scans, where in each acquisition process a series of time-consuming steps, including motor movements, detector activation, and data retrieval. Although widely employed, this method has inherent dead time, restricting scan speed to just a few points per second, especially when dealing with millisecond-range measurements.

Building upon the foundation laid in previous research, this subsequent paper delves deeper into our quest for a more streamlined and expedited data collection methodology. The primary focus continues to be on continuous scans, a concept previously explored using a combination of software and hardware solutions. In this work, we further refined the approach, placing a heightened emphasis on hardware-based solutions to optimize motion control and precise trigger pulse generation. Our objective is to significantly enhance

the speed and efficiency of data acquisition, ultimately pushing the boundaries of what is attainable in the realm of scientific experimentation.

This paper outlines the ongoing efforts to develop and implement a robust hardware-centric system, shedding light on its novel capabilities, performance enhancements, and applications across various experimental setups.

## SYSTEM

The system is based on the TANGO Controls toolkit [2] and Sardana [3] for scan orchestration – Fig. 1 illustrates the implemented system. If necessary, non-linear motions and multi-motor synchronization are effected using parametric trajectories [4] with trapezoidal profile movements (ensuring constant speed while respecting acceleration and deceleration constraints) in IcePAP [5]. Such trajectories can also be applied using any other motor driver with parametric trajectory support. IcePAP features a position closed-loop system with an integral controller for each individual axis. It is imperative to fine-tune this controller to ensure that the motor accurately follows its trajectory over its full range. All interlocks, limit switches, and error stops actuate on all associated axes, ensuring coordinated movements. This setup permits the use of a single motor encoder as a position reference since the driver guarantees synchronization of the other axes.




Figure 1: Schematic of the implemented system.

The timing system is based on PandABox [6] and allows the user to choose between position-based or time-based triggering. For position-based, each trigger is determined using the incremental or absolute encoder readings, sampled at 1 MHz by PandABox, and a pre-loaded look-up table. The values in the table are equally spaced in eV, based on the user inputs. For time-based, only the first scan point is determined by the encoder position, afterwards, all remaining pulses are generated with a constant time interval. The system accepts combinations of multiple motors, e.g., undulator, grating and mirror for PGM monochromator, or a single motor used as the master. PandABox supports up to eight encoder inputs. The combination will vary based on the resolution needed for each application. For incremental encoders, the scan routine synchronizes the IcePAP readings with PandABox

\* aureo.freitas@maxiv.lu.se

# INTEGRATING ONLINE ANALYSIS WITH EXPERIMENTS TO IMPROVE X-RAY LIGHT SOURCE OPERATIONS \*

N.M. Cook<sup>†</sup>, E. Carlin, J. Einstein-Curtis, P. Moeller, R. Nagler, R. O'Rourke,  
RadiaSoft LLC Boulder, CO 80301, USA

A. Barbour, M. Rakitin, L. Wiegart, H. Wijesinghe,  
National Synchrotron Light Source II, Brookhaven National Laboratory, NY, 11973, USA

## Abstract

The design, execution, and analysis of light source experiments requires the use of sophisticated simulation, controls and data management tools. Existing workflows require significant specialization to accommodate specific beamline operations and data pre-processing steps necessary for more intensive analysis. Recent efforts to address these needs at the National Synchrotron Light Source II (NSLS-II) have resulted in the creation of the Bluesky data collection framework, an open-source library for coordinating experimental control and data collection. Bluesky provides high level abstraction of experimental procedures and instrument readouts to encapsulate generic workflows. We present a prototype data analysis platform for integrating data collection with real time analysis at the beamline. Our application leverages Bluesky in combination with a flexible run engine to execute user configurable Python-based analyses with customizable queuing and resource management. We discuss initial demonstrations to support X-ray photon correlation spectroscopy experiments and future efforts to expand the platform's features.

## INTRODUCTION

X-ray light sources are prominent drivers of scientific discovery across a range of disciplines. These facilities serve a diverse user community, often providing concurrent beam time and user support to tens of domain scientists with unique backgrounds. Increasing demand for beam time, coupled with the increasing sophistication of experiments, places constraints on the software infrastructure required to successfully carry out experiments within available time and resource allocations. Recently, significant development efforts have been made towards improving experimental planning and execution; however, significant challenges remain to integrating real-time analysis tools within the experimental workflow. In this proceedings, we discuss a strategy for incorporating analysis pipelines within common experimental workflows, focusing on applications at the NSLS-II light source. We present a schematic workflow for orchestrating analysis in concert with experimental execution. We then demonstrate this workflow via an open source, browser-based interface furnishing beamline agnostic analysis pipelines.

\* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Award Number DE-SC00215553.

<sup>†</sup> ncook@radiasoft.net

## AN INTEGRATED FRAMEWORK FOR EXPERIMENT AND ANALYSIS

Our platform supports customizable analysis coupled with pre-existing experimental and data management frameworks. The Bluesky Data Collection Framework is used to coordinate data and metadata access and processing [1]. Bluesky aims to provide end-to-end experimental planning, execution, and data acquisition tools through a set of interoperable Python libraries; it is currently used across all active beamlines at NSLS-II. Bluesky is comprised of a set of distinctive libraries, and here we highlight a few critical ones. First, the eponymous bluesky library implements a run engine and event model to permit experimental control and data collection through the execution of high level plans. The ophyd library provides hardware abstraction to communicate plans to devices along the beamline. The databroker library implements an API for structured access to experimental data and metadata generated during an experiment executed by Bluesky. These libraries are under active development, with new capabilities added frequently.

To orchestrate analyses, we have developed a custom scan monitor to provide asynchronous execution of Python-based analyses, operating in tandem with a Sirepo front end user interface. Sirepo is an open-source scientific computing gateway that provides access to community codes through custom, browser-based interfaces and an embedded JupyterHub instance. Sirepo is designed to be hardware agnostic; simulation environments are deployed via Docker containers, and can be executed across a range of computing systems, ranging from a laptop to a GPU cluster at a high performance computing facility. Sirepo provides support for numerous accelerator modeling and related tracking codes; existing applications have been employed to provide customized simulations of X-ray beamlines at the NSLS-II using the Synchrotron Radiation Workshop code [2]. Sirepo has also been integrated with Bluesky to enable the asynchronous execution of long-running SRW simulations to support multi-parametric optimizations of beamlines [3].

## ANALYSIS SERVICE AT NSLS-II

We have deployed a prototype analysis web-service at NSLS-II. The application is hosted on the NSLS-II science network, accessible from beamline workstations or via remote desktop applications such as Apache Guacamole or VMware Horizon by all credentialed users. The application is deployed via a Podman container providing the environ-

# INTEGRATING TOOLS TO AID THE AUTOMATION OF PLC DEVELOPMENT WITHIN THE TwinCAT ENVIRONMENT

N. Mashayekh\*, B. Baranasic, M. Bueno, T. Freyermuth, P. Gessler, S. T. Huynh, N. Jardón Bueno, J. Tolkiehn, L. Zanellatto, European X-Ray Free-Electron Laser, Schenefeld, Germany

## Abstract

Within the myriad of day to day activities, a consistent and standardised code base can be hard to achieve, especially when a diverse array of developers across different fields are involved. By creating tools and wizards, it becomes possible to guide the developer and/or user through many of the development and generic tasks associated with a Programmable Logic Controller (PLC).

At the European X-Ray Free Electron Laser Facility (EuXFEL), we have striven to achieve structure and consistency within the PLC framework through the use of C# tools which are embedded into the TwinCAT environment (Visual Studio) as Extensions. These tools aid PLC development and deployment, and provide a clean and consistent way to develop, configure and integrate code from the hardware level, across to the Supervisory Control And Data Acquisition (SCADA) system.

Keywords: PLC, TwinCAT, C#, software tools, extension development

## INTRODUCTION

Whilst there were tools [1] previously developed in order to aid Programmable Logic Controller (PLC) project generation at the European X-Ray Free-Electron Laser (EuXFEL), overtime, these tools became harder to manage. Many of the tools were developed in an array of programming languages, and require the PLC developers to become adept in multiple Integrated Development Environment (IDE) and languages. In turn, in order to enhance or add to an existing tool or function, edits would have had to be made across the multiple applications to ensure consistency. This approach can work within a diverse and well integrated team, however, also caused bottle necks and had a high dependency on all of the tools being kept up-to-date. This constriction was highlighted as an area which could definitely be improved upon.

A new approach was envisioned where all of the function previously being performed either manually, or via some means of automation, was collated together into a Single Point of Contact (SPoC). Provided with the backdrop of the TwinCAT environment, it was a logical step to build upon this platform by integrating this new functionality and interface into TwinCAT itself via the means of Visual Studio extensions.

\* navid.mashayekh@xfel.eu

## THE NEED FOR TOOLS AND WIZARDS

TwinCAT's integration with Visual Studio makes it convenient to combine Visual Studio extensions with the TwinCAT Automation Interface library. This combination is highly beneficial for automated PLC project generation, code injection, and hardware linking. Visual Studio extensions enable custom tools and workflows that seamlessly integrate into the TwinCAT environment, while the TwinCAT Automation Interface library provides programmatic access to TwinCAT's features; allowing for the automation of tasks and the adaptation of useful features associated with modern programming languages. Together, these tools enhance productivity and reduce the need for switching between different applications whilst reducing the potential for errors within automation workflows at EuXFEL.

## VISUAL STUDIO EXTENSIONS

A solid foundation provides the ideal canvas upon which to develop auxiliary tools. By utilising Visual Studio extensions within the realm of C#, it becomes feasible to import an in-house configuration seamlessly and to also dynamically modify both the project node, and the hardware node of the project in an adhoc manner.

This capability enables developers to make on-the-fly adjustments and closely monitor the success or failure of each step throughout the process. Consequently, this real-time feedback loop facilitates efficient and flexible project management, allowing for swift adaptation to changing requirements in addition to a more streamlined development process. Also in response to TwinCAT's lack of support for generic data types, developers can create Visual Studio extensions as a workaround to generate duplicated function blocks (Fig. 1) for different data types (Fig. 2).




Figure 1: List implementation for multiple interfaces.

# STATE MACHINE OPERATION OF COMPLEX SYSTEMS\*

P. Hanlet<sup>†</sup>, Fermi National Accelerator Laboratory, Batavia, Illinois, USA

## Abstract

Operation of complex systems which depend on one, or more, sub-systems with many process variables often operate in more than one state. For each state there may be a variety of parameters of interest, and for each of these, one may require different alarm limits, different archiving needs, and have different critical parameters. Relying on operators to reliably change  $10^2$ - $10^5$  of parameters for each system for each state is unreasonable. Not changing these parameters results in alarms being ignored or disabled, critical changes missed, and/or data archiving inefficiencies.

To reliably manage the operation of complex systems, such as cryomodules (CMs), Fermilab is implementing state machines for each CM and an over-arching state machine for the PIP-II superconducting linac (SCL). The state machine transitions and operating parameters are stored/restored to/from a configuration database. Proper implementation of the state machines will not only ensure safe and reliable operation of the CMs, but will help ensure reliable data quality. A description of PIP-II SCL, details of the state machines, and lessons learned from limited use of the state machines in recent CM testing will be discussed.

## INTRODUCTION

Fermi National Accelerator Laboratory, or Fermilab, in Batavia, Illinois, USA is constructing a new superconducting linear accelerator (LINAC) with twice the energy of the existing LINAC and significantly higher power. The LINAC, PIP-II, will power the rest of the Fermilab accelerator complex, generating the world's most intense high-energy neutrino beam, as well as providing beam to other experiments and test beams, see Fig. 1.




Figure 1: Fermilab's accelerator chain and experiments.

The capabilities of PIP-II [1] are to provide an 800 MeV proton beam of 1.2 MW using a superconducting RF LINAC, see Fig. 2. The beam is upgradeable to multi-MW and CW-compatible as well as customizable for a variety of user

requirements. The scope of PIP-II includes a beam transfer line to the existing Booster ring and accelerator complex upgrades to the Booster, Recycler, and Main Injector,




Figure 2: PIP-II scope: new LINAC and beam transfer line to Booster.

## MOTIVATION

In this context, a “complex system” refers to a hardware device in a control system which has multiple subsystems and multiple operational states. This frequently means that it has a large number of process variables (PVs) and for each state: there are different PVs of interest, different alarm limits and severities, different archiving needs, and different critical PVs (those used to identify/notify subsystem experts).

The result of ignoring these differences in states can potentially be severe: e.g. incorrect alarm limits may be too loose and thus not notify operators of problems; or worse, if the alarm limits are too tight, the alarms are continuous and likely ignored and/or disabled. For archiving controls data, collection may be inefficient, or worse, changes may not be recorded if deadbands are too loose. Incorrect identification of states may result in delays in operation if subsystem experts are not contacted promptly.

## Example

An example of a complex system is an accelerator superconducting RF cryomodule (CM). At PIP-II, each CM will have subsystems: vacuum, cryogenics, safety, machine protection (MPS), RF permits (RFPI), low level RF (LLRF), and high power RF (HPRF). A partial example of these states for this system are shown in Table 1. Note that as one transitions through the states, the numbers of PVs change.

Note also the distinction in the “Types” of states. “Static” states refer to those in which all of the PVs are expected to remain constant, within deadbands. Here, the alarm limits are tight and one archives data in a monitor mode. In contrast, for the “Dynamic” states, some of the PVs will remain static, but others are expected to change. Those which change

\* This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

<sup>†</sup> hanlet@fnal.gov

## COMPONENTS OF A SCALE TRAINING TELESCOPE FOR RADIO ASTRONOMY TRAINING

A. C. Linde, X. P. Baloyi, P. Dube, L. Lekganyane, A. M. Lethole,  
V. Mlipha, P. Pretorius, U. C. Silere, S. Sithole, SARAQ, Cape Town, South Africa

### *Abstract*

To establish the engineering and science background of Radio Astronomy in SKA African partner countries, a need was identified to develop a training telescope which would serve as a vehicle for demonstrating the principles. The Scale Training Telescope (STT) will be used as an interactive teaching tool for experiencing the basics of a radio telescope structure and control system development as well as its operation in practical applications, tracking sources and finding stationary satellites. The telescope aims to work as closely to a real radio telescope as possible. The STT allows students at various academic levels in different educational institutions the ability to access a scale version of a radio telescope design that can be assembled and operated by the students.

The paper describes the mechanical, electrical, electronics and software elements of the STT. The mechanical elements range from the structural base to the rotating dish of the radio telescope. The electrical elements incorporate the electromechanical components used to move the telescope as well as the wiring and powering of the telescope. The software is used to control the telescope system as well as collect, process and visualize the resulting data. A software-based user interface will allow the students to control and monitor the telescope system. The PLC-based (Programmable Logic Controller) control system facilitates the motion control of the telescope, in both the azimuth and elevation axes.

### BACKGROUND OF THE STT PROJECT

The establishment of the Square Kilometer Array (SKA) Telescope in South Africa has provided an opportunity to develop Radio Astronomy not only in South Africa but also in the SKA African partner countries. As a result, an interactive training tool is required to assist in the introduction of radio astronomy.

By demonstrating the engineering and science of a full-scale radio telescope through a scaled model, radio astronomy on the African continent is given the opportunity to grow. Both the final product and the development processes are replicated [1].

The project is currently run at the Cape Town South African Radio Astronomy Observatory (SARAQ) offices as part of a graduate-in-training group assignment designed to allow the graduates to explore the workings of an engineering project, collaboration between the different disciplines and, most importantly, the workings of a radio telescope. The departments that the graduates are in include project management, mechanical engineering, civil engineering, software development and electronic engineering.

The platform allows the graduates to network in the workplace as well as gain valuable experience in working in multi-disciplinary teams, all while in the learning phase of their employment. The current STT project is under construction.

While an overview of the mechanical and electronic systems will be given, it should be noted that the focus of this paper is the control system of the STT.

### OVERVIEW OF STT PROJECT

A radio telescope is broken down into several subsystems that work together to achieve functionality, namely the telescope structure, the signal chain, the software system and the control system. In each of these subsystems there are hardware, specifically mechanical and electronic components, and software elements.

These subsystems are adapted for the scaled model to functionally resemble a full-scale telescope as closely as possible. The components are scaled down in size but retain functionality, allowing the possibility of the telescope software to be integrated to an extent that provides a further training platform for telescope operators.

The functions of the telescope include rotation of the dish in both azimuth and elevation, location of true north and 0° elevation, stowing, and displaying position as well as speed in both the azimuth and elevation positions. The telescope design also accounts for protection from damage during use in both elevation and azimuth directions, including an emergency stop feature.

### MECHANICAL SYSTEM

By considering the mechanical functionality of a full-scale telescope, the sub-systems for the scale telescope are derived and designed. The main mechanical subsystems to be considered are the yoke and pedestal assemblies, the elevation drive assembly and the reflector assembly. In order to allow for maximum pointing and tracking accuracy, anti-backlash mechanisms are implemented in both the azimuth and elevation directions. The overall STT design can be seen in Figs. 1 and 2 [2].

# WEB DASHBOARDS FOR CERN RADIATION AND ENVIRONMENTAL PROTECTION MONITORING

Adrien Ledoul\*, Alexandru Savulescu, Gustavo Segura  
CERN, Geneva, Switzerland

## Abstract

CERN has developed and operates a SCADA system for radiation and environmental monitoring, which is used by many users with different needs and profiles. To provide tailored access to this control system's data, the CERN's Occupational Health & Safety and Environmental Protection (HSE) Unit has developed a web-based dashboard editor that allows users to create custom dashboards for data analysis.

In this paper, we present a technology stack comprising Spring Boot, React, Apache Kafka, WebSockets, and WebGL that provides a powerful tool for a web-based presentation layer for the SCADA system. This stack leverages WebSocket for near-real-time communication between the web browser and the server. Additionally, it provides high-performant, reliable, and scalable data delivery using low-latency data streaming with Apache Kafka. Furthermore, it takes advantage of the GPU's power with WebGL for data visualization.

This web-based dashboard editor and the technology stack provide a faster, more integrated, and accessible solution for building custom dashboards and analyzing data.

Keywords: WebSocket, Kafka, React, HighCharts, WebGL, Dashboard.

## INTRODUCTION

### *CERN Radiation and Environmental Monitoring*

The highest priority of CERN is to ensure that its research infrastructure, whose core assets are particle accelerators and experimental areas, operates safely for the workers and for the public, while minimizing the organization's activities environmental impact. To that regard, CERN has established and implemented a Safety Policy where radiation safety and environment protection are central.

The implementation of the safety policy includes the setup and operation of a radiation and environment monitoring infrastructure that provides a real-time view of the effects of the operation of CERN. This monitoring infrastructure helps the organization to early detect deviations and gives the possibility to steer the processes (accelerators or operational parameters) to avoid approaching safety boundaries. The infrastructure also allows recording radiation and environment measurements, which are reported to the authorities under the form of comprehensive regulatory reports and are also published in public environment reports [1].

The monitoring infrastructure is composed of the extensive network of instruments that continuously measure

radiation and conventional environment parameters in accelerators, experimental areas and in public CERN places across the campus, but also within closed perimeters outside CERN fences.

### *Radiation and Environment Monitoring Unified Supervision*

The Radiation and Environment Monitoring Unified Supervision (REMUS) [2] is the Supervisory Control and Data Acquisition (SCADA) system that orchestrates the instrumentation. REMUS is built on WinCC Open Architecture (WinCC OA) [3]. It collects measurements from the instrumentation and immediately makes them available to the CERN control rooms operators and to experts in charge of radiation protection and environment monitoring. REMUS also supervises the correct operation of the monitoring infrastructure by detecting faulty instruments, network communication issues, power failure, and faults in its internal software components. This is supported by the fully redundant architecture of the REMUS system.

REMUS supervises an infrastructure of 2 800 instruments. These instruments are very heterogeneous with up to 87 different types of devices. REMUS contains 900 000 SCADA tags, manages 125 000 distinct alarm tags and handles a throughput of 25 000 input/output operations per second. REMUS archives 80 billion of measured values per year.

The big amount of data generated and recorded by REMUS are made available to its more than 250 active users in real-time SCADA client software and in long-term data stores namely Oracle and the Next CERN Accelerator Logging Service (NXCALS) [4].

REMUS accommodates a wide range of professional profiles (accelerator or experiment operators, radiation protection specialist, environmental protection specialist, physicist, instrumentation engineers and technicians, SCADA software engineers, etc.), with different needs and functional expectations from REMUS. For these users, REMUS provides software tools adapted or adaptable to their specific needs, which furthermore can be tailored and adaptable down to each individual user. However, the need for adaptive user interfaces with wider customization capabilities and openness grows constantly. In order to respond to these increasing user needs, REMUS has implemented and deployed a modern web dashboard editor.

## RATIONALE

This section introduces the rationale behind the creation of web dashboards, and solutions that were considered to

Software

User Interfaces & User Experience

\* adrien.ledoul@cern.ch

## PHOEBUS TOOLS AND SERVICES

K. Shroff, Brookhaven National Laboratory, Upton, NY, USA  
R. Lange, ITER Organization, St. Paul lez Durance, France  
T. Ford, Lawrence Berkeley National Laboratory, Berkeley, CA, USA  
G. Weiss, European Spallation Source ERIC, Lund, Sweden  
K. Kasemir, Oak Ridge National Laboratory, Oak Ridge, TN, USA  
T. Ashwarya, FRIB, East Lansing, MI, USA

### Abstract

The Phoebus toolkit consists of a variety of control system applications providing user interfaces to control systems and middle-layer services. Phoebus is the latest incarnation of Control System Studio (CS-Studio), which has been redesigned replacing the underlying Eclipse RCP framework with standard Java alternatives like SPI, preferences, etc. Additionally, the GUI toolkit was switched from SWT to JavaFX. This new architecture has not only simplified the development process while preserving the extensible and pluggable aspects of RCP, but also improved the performance and reliability of the entire toolkit.

The Phoebus technology stack includes a set of middle-layer services that provide functionality like archiving, creating and restoring system snapshots, consolidating and organizing alarms, user logging, name lookup, etc. Designed around modern and widely used web and storage technologies like Spring Boot, Elasticsearch, MongoDB, Kafka, the Phoebus middle-layer services are thin, scalable, and can be easily incorporated in CI/CD pipelines. The clients in Phoebus leverage the toolkit's integration features, including common interfaces and utility services like adapter and selection, to provide users with a seamless experience when interacting with multiple services and control systems.

### MOTIVATION

The workflows of operators, engineers, and scientists interacting with control systems like EPICS [1] often involve interacting with multiple applications, each designed for specific use cases. Some applications are used to display real-time data from various signals, some visualize historical data or consolidated alarms, while others provide a set of operations to write directly to the control system or trigger actions in middle-layer services, such as setting predefined values from a snapshot.

Phoebus/Control System Studio (CSS) [2] was developed to streamline these workflows by offering a suite of integrated applications and a framework that simplifies the development of such applications. For end-users, this translates to a collection of applications that are easy to navigate between, where data can be seamlessly transferred from one application to another, ensuring a consistent user experience. For developers, CSS provides a framework for creating applications that seamlessly

integrate into this ecosystem without requiring tight dependencies on other applications. It also supports modules that offer access to shared and optimized resources, such as connections to EPICS Process Variables (PVs), REST clients, and more.

### Stepping into the Sun

The initial incarnation of Control System Studio (CSS) was built on top of the Eclipse Rich Client Platform (RCP) which provided a great foundation, offering a plethora of essential features required for building extensible, pluggable applications. Over time, most of the core functionalities and capabilities originally provided by the Eclipse RCP framework found their way into the standard Java Development Kit (JDK) library itself. The incorporation of these Eclipse RCP features into the JDK introduced an opportunity to simplify the application development by allowing developers to leverage the standardized Java library, and a more cohesive and streamlined development process that comes with it. This transition was further accelerated by the increasing complexity of the Eclipse RCP framework, which prompted us to embrace the native Java ecosystem for building Phoebus as a modern, efficient, and maintainable application.

## PHOEBUS ARCHITECTURE

### Java Service Provider Interface

Phoebus/CSS is employed across diverse international scientific and industrial setups, each with its unique requirements. Hence, enabling each site to curate customized products comprising relevant applications and tools within their control environment is crucial. Equally important is Phoebus's seamless integration with existing software stacks.

Phoebus facilitates this adaptability through a list of Service Provider Interfaces (SPIs) for contributing applications. Java SPI is a mechanism to modularize a software framework. In the case of Phoebus, SPI allows applications to register for file extensions, contribute menu or toolbar entries, offer data sources or access to a site-specific logbook. To add a new application, one simply needs to provide an implementation of the app SPI and include it in the classpath. This approach encourages the contribution of new applications and simplifies the creation of site-specific products tailored to each site's specific needs.

# VME2E: VME TO ETHERNET – COMMON HARDWARE PLATFORM FOR LEGACY VME MODULES

Y. Tian<sup>†</sup>, NSLS-II, Brookhaven National Laboratory\*, Upton, New York, USA  
J. Jamilkowski, Electron-Ion Collider, Brookhaven National Laboratory, Upton, New York, USA

## Abstract

The VME architecture, developed in the late 1970s, has been a reliable control system platform for decades. Today, it faces challenges including throughput, computing power, the operating system, and obsolescence of legacy modules. Next-gen platforms like ATCA and microTCA require significant changes but may not always be suitable for legacy upgrades.

In this paper, we introduce VME2E (VME to Ethernet), an open-source hardware platform for replacing legacy VME modules without disassembling the system. VME2E retains the VME form factor and utilizes stable power and cooling. It features a Xilinx FPGA system-on-module (SOM) with GigE for high-speed computing and communication, along with a high pin count (HPC) FPGA mezzanine connector (FMC) for IO support. VME2E is designed as a low-cost, open-source solution for VME legacy upgrades.

## INTRODUCTION

The VME architecture [1, 2, 3], originally introduced in the late 1970s, has established itself as a robust hardware platform for controlling systems over the past four decades. VME, short for 'VERSAbus Module Eurocard,' was developed to serve various industrial, commercial, and military applications. The standard specifies both electrical and mechanical specifications.

The VME bus follows a master-slave computer architecture. Notably, its signalling scheme is asynchronous, meaning that data transfer is not tied to the timing of a bus clock. Additionally, a single VME crate can accommodate up to 21 modules.

However, today, the VME hardware platform faces the following four significant challenges.

### Backplane Throughput

In a standard VME chassis, modules communicate with each other using shared bandwidth. Over the past four decades, the throughput of the VME backplane has been improved multiple times, including:

- VMEbus, introduced in 1981 with a throughput of 40MB/s.
- VME64, introduced in 1994 with a throughput of 80MB/s.

- VME64x, introduced in 1997 with a throughput of 160MB/s.
- VME320/2eSST, introduced in 1997 with a throughput of 320MB/s.

However, as a shared bus, the VME backplane's throughput is inherently limited. The VME protocol is based on master-slave communication, where the backplane is shared by all VME boards, essentially operating sequentially.

Today, each module can have a Gigabit Ethernet port to send data or communicate with external systems using high-speed fiber links. The shared nature of the VME backplane restricts the overall chassis throughput.

### FEC Computing Power

In a typical VME chassis, the Front-End Controller (FEC) handles most of the computing. As computing power has improved, so has the performance of VME FECs. Here are some snapshots of the history of VME FECs:

- 1979: Motorola 68000 (m68k): 68k transistors, 3.5 $\mu$ m technology, 4-16.67MHz.
- 1988: MVME147 MC68030: 16-33MHz, 4K SRAM, 32MB DRAM.
- 1992: MVME162, 1.6 million transistors, 90-120MHz, 0.5 $\mu$ m technology.
- 2010: MVME3100, PowerPC MPC8540, 1GHz, 500MB DDR, 2eSST.

However, today, each module can host a powerful computing unit, such as a microprocessor or FPGA. Concentrating all the computing power solely on the FEC has its limits.

### Operating System

Due to its centralized architecture, where all VME modules share computing power on the FEC and communication bandwidth on the backplane, it is often necessary to run a real-time operating system (OS) on the FEC. Over the years, several real-time operating systems have been developed to run on VME FECs, including VxWorks/Tornado (from WindRiver Systems), RTEMS (from OAR Corp), LynxOS (from Lynx Real Time Systems), OS-9 (from MicroWare), and SPECTRA/VRTX (from Microtech Research).

However, supporting the real-time OSs is becoming increasingly challenging, and many have ceased development.

\* Work supported by NSLS-II, BNL, US Department of Energy

<sup>†</sup> Email address: ytian@bnl.gov



# DATA MANAGEMENT INFRASTRUCTURE FOR EUROPEAN XFEL

J. Malka\*, S. Aplin, D. Boukhelef, K. Filippakopoulos, L. Maia,  
 T. Piszczek, G. Previtali, J. Szuba, K. Wrona  
 European XFEL, Schenefeld, Germany

S. Dietrich, M. Gasthuber, J. Hannappel, M. Karimi, Y. Kemp, R. Lueken,  
 T. Mkrtchyan, K. Ohrenberg, F. Schluenzen, P. Suchowski, Ch. Voss  
 Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany

## Abstract

Effective data management is crucial to ensure research data is easily accessible and usable. We will present the design and implementation of the European XFEL data management infrastructure supporting high-level data management services. The system architecture comprises four layers of storage systems, each designed to address specific challenges. The first layer, referred to as Online, is designed as a fast cache to accommodate extremely high rates where up to 15 GB/s of data is generated during experiments with a single scientific instrument. The second layer, called high-performance storage, provides the necessary capabilities for data processing both during and after experiments. The layers are incorporated into a single InfiniBand fabric and connected through a 4.4 km long 1 Tb/s link. This allows fast data transfer from the European XFEL experiment hall to the DESY computing centre. The third layer, mass storage, extends the capacity of the data storage system to allow mid-term data access for detailed analysis. Finally, the tape archive provides data safety and a long-term archive (10+ years). The high-performance and mass storage systems are connected to computing clusters. This allows users to perform near-online and offline data analysis, or alternatively export data outside the European XFEL facility. The data management infrastructure at the European XFEL has the capacity to accept and process up to 2 PB of data per day, which demonstrates the remarkable capabilities of all the sub-services involved in this process.

## EUROPEAN XFEL FACILITY

The European XFEL Facility is one of the most advanced sources of pulsed, extremely intense and coherent radiation in the hard and soft X-ray range. These properties of X-ray flashes attract various research communities and scientists, who use them in a diversity of scientific investigations. The facility spanning a length of 3.4 km, extends from the DESY campus in Hamburg to the town of Schenefeld situated in Schleswig-Holstein, where the European XFEL experiment stations are located. Currently, a scientific experiment can be performed using one of the seven available instruments and utilising one of the three self-amplified spontaneous emission light sources called SASE. The layout of scientific instruments and SASEs is illustrated in Fig. 1.

\* janusz.malka@xfel.eu




Figure 1: Beamlines and instruments of the European XFEL

## Scientific Data Policy

The European XFEL Scientific Data Policy defines the rights and responsibilities for data of all parties involved in experiments and sets rules for data management. It was approved and published in 2017, shortly before the first user experiments were performed. Based on the policy, a set of data services has been created together with underlying storage and computing infrastructure. Through the implemented services, scientists are able to catalogue and later find the data, assess its quality and trigger some actions like, for example, data archiving, and specialised data processing pipelines. The data policy is currently under revision in order to be aligned with the FAIR data principles. It will also address the challenges related to the vast data volumes by introducing data management plans, data reduction concepts and a new data retention policy. Along with the updated policy, the data services are being reviewed with the goal of implementing new data handling concepts.

## Meta Data Catalogue

Data generated in the context of scientific experiments at the European XFEL facility is curated with the help of metadata catalogue service (myMdC) [1]. The metadata catalogue service besides storing information about experiments data, acts as a hub for integration with other services (e.g. data acquisition system, electronic logbook, calibration service, storage infrastructure) and is able to execute various data management workflows. The myMdC is built on top of the relational database, which links experiment meta-data and data together. It allows scientists to define experiment techniques, measurements and sample types for each experiment dataset. Through extensive use of microservices, it allows the automation of data infrastructure tasks such as creating an underlying file system structure for each experiment and defining data access roles within the experiment team. It is also used to trigger data migration processes, ini-

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

## DESIGN OF THE HALF CONTROL SYSTEM\*

G. Liu<sup>†</sup>, L. Chen, C. Li, X. Sun, K. Xuan, D. Zhang, NSRL, USTC, Hefei, Anhui 230029, China

### Abstract

The Hefei Advanced Light Facility (HALF) is a 2.2 GeV 4th synchrotron radiation light source, which has been under construction in Hefei, China since June 2023. The HALF contains an injector and a 480-m diffraction limited storage ring, and 10 beamlines for Phase-I. The HALF control system is built on EPICS with integrated application and data platforms for the entire facility, which includes accelerator, beamlines, and utilities. The unified infrastructure and network architecture are designed to build the control system. The infrastructure provides resources for the EPICS development and operation through virtualization technology, and provides resources for the storage and process of experimental data through distributed storage and computing clusters. The network is divided into the control network and the dedicated high-speed data network by physical separation, the control network is subdivided into multiple subnets by VLAN technology. Through estimating the scale of the control system, the 10Gbps control backbone network and the data network that can be expanded to 100Gbps can fully meet the communication requirements of the control system. This paper reports the control system architecture design and the development work of some key technologies in details.

### INTRODUCTION

The Hefei Advanced Light Facility (HALF) is a 2.2 GeV 4th synchrotron radiation light source, which has been under construction in Hefei, China since June 2023. The main parameters of the storage ring are listed in Table 1 [1].

Table 1: Main Parameters of the HALF Storage Ring

Parameter	Value
Beam Energy	2.2 GeV
Circumference	480 m
Emittance	86.3 pm·rad
Beam current	350 mA
Brightness	$1.15 \times 10^{21}$ phs/mm <sup>2</sup> /mrad <sup>2</sup> /0.1%BW/s
Injection	Full energy, Top-off
Lattice	6BA
Straight sections	20×5.3 m + 20×2.2 m
RF	500 MHz

As shown in Fig. 1, HALF is composed of a 2.2 GeV full energy linac, a transport line, a 480 m diffraction limited storage ring and 10 beamlines at phase-I.

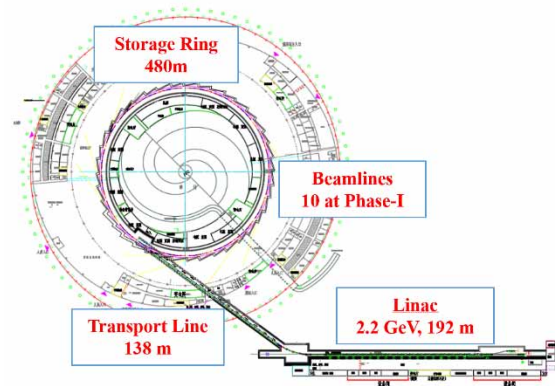


Figure 1: Layout of HALF.

The HALF control system serves as the platform for the commissioning and operation of the entire facility. It adopts a unified architecture to enable device control, process control, data acquisition and analysis for the accelerator, beamlines, and utilities, thereby addressing the issue of insufficient manpower. HALF is a user facility, and its availability is a primary requirement in design to ensure high-quality operations. With approximately 250,000 process variables (PVs), efficient development, deployment, and operation of the control system are required. Additionally, scalability, flexibility, and security must be considered. The users of HALF control system can be divided into accelerator physicists, operators, engineers, beamline staff, experiment station users, and facility managers. To meet their respective requirements, targeted design should be conducted, and corresponding user interfaces should be provided.

Considering the design requirements and development team's experience and foundation, EPICS7 [2] will be used as the software development platform. Server virtualization, VLAN, and commercial off-the-shelf (COTS) products will be utilized to complete the control system development within the constraints of human resources, budget, and time.

### ARCHITECTURE

The HALF control system is a distributed control system based on EPICS7, which is divided into 3 layers: device control layer, middle layer, and presentation layer, as illustrated in Fig. 2. The device control layer is the basis of the HALF control system, and it controls the facility-wide devices from the accelerator, beamlines, and utilities, while providing equipment protection and timing. The middle layer provides basic IT services such as DNS, NTP, and data collection and processing. In the presentation layer,

\*Work supported by the National Development and Reform Commission, Anhui Province and Hefei City

<sup>†</sup> gfliu@ustc.edu.cn

# LCLS-II EXPERIMENT SYSTEMS VACUUM CONTROLS ARCHITECTURE\*

M. Ghaly<sup>†</sup>, A. Wallace, SLAC National Accelerator Laboratory, Menlo Park CA, U.S.A

## Abstract

The LCLS-II Experiment System Vacuum Controls Architecture is a collection of vacuum system design templates, interlock logic, supported components (eg. gauges, pumps, valves), interface I/O, and associated software libraries, which implement a baseline functionality and simulation. The architecture also includes a complement of engineering and deployment tools including cable test boxes or hardware simulators, as well as some automatic configuration tools. Vacuum controls at LCLS span from managing rough vacuum in complex pumping manifolds, protection of highly-sensitive x-ray optics using fast shutters, maintaining ultra-high vacuum in experimental sample delivery setups, and beyond. Often, the vacuum standards for LCLS systems exceed what most vendors are experienced with. The system must maintain high-availability while remaining flexible and adaptable to accommodate ongoing modifications. This paper provides an overview of the comprehensive architecture and the specific requirements of the LCLS systems. Additionally, it introduces how to utilize this architecture for new vacuum system designs. The architecture is intended to influence all phases of a vacuum system life-cycle, and ideally with the goal of becoming a shared project for installations beyond LCLS-II.

## INTRODUCTION

The majority of LCLS experiments and measurements necessitate a vacuum environment as they rely on the absence of any background interference. Furthermore, various optic devices are highly sensitive to vacuum levels, with any level above Ultra High Vacuum (UHV) potentially accelerating the rate of contamination on the mirror surface.

The overarching Experiment Controls System (ECS) vacuum controls system is responsible for controlling vacuum devices and protecting both vacuum components and vacuum itself. Operating at a high availability, the system supports a diverse array of vacuum devices, including gauges, pumps, valves, and other related devices and controllers. Utilizing a single vacuum Programmable Logic Controller (PLC) for numerous devices, the system efficiently handles complex logic necessary to enforce protection interlock requirements. Vacuum interlocks established to prevent the system from entering an unsafe or undesired state. For instance, the system should inhibit the operator from opening a valve when the differential pressure across this valve exceeds a certain limit. Additionally, reactive interlocks are implemented for the closure and isolation of specific vacuum

volumes to prevent the propagation of pressure events from adjacent sections.

Furthermore, the vacuum control system interfaces with other systems such as the Machine Protection System (MPS), enabling it to shut off the beam while a beamline valve is closing. This functionality protects the valve from potential damage caused by the beam.

## ARCHITECTURE

The ECS vacuum controls architecture is a collection of software libraries, scripts, widgets and a complement of software and hardware tools that are fully integrated cross all layers of the controls stack.

The ECS vacuum controls system design is based on Beckhoff embedded controllers and associated industrial hardware and EtherCAT. The system's software stack includes the Experimental Physics and Industrial Control System (EPICS) layer, the Python layer, and the User Interface (UI) Layer, as illustrated in Fig. 1.

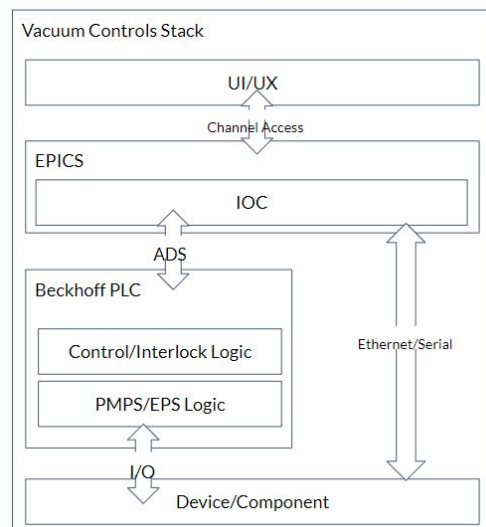


Figure 1: ECS Vacuum Controls Stack.

## Device Hardware

The design process of every new vacuum device starts at the connector. Each vacuum device has a corresponding custom cable designed for it. These comprehensive cable design drawings include all the necessary information for cable fabrication, including specifications for the cable type, connector type, and fabrication instructions. Additionally, the cable designs contain pin-outs for the signals needed to control and monitor each device, as illustrated in Fig. 2.

For every device cable a number of I/O terminals are designated based on the device type enabling signals readouts

\* Work supported by U.S. D.O.E Contract DE-AC02-76F00515.

<sup>†</sup> mghaly@slac.stanford.edu

# THE LCLS-II PRECISION TIMING CONTROL SYSTEM\*

T. Johnson<sup>†</sup>, M. Browne, C. Pino  
SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

The LCLS-II precision timing system is responsible for the synchronization of experiment optical lasers with the LCLS-II XFEL. The system uses both RF and optical references for synchronization. In contrast to previous systems used at LCLS the optical lasers are shared resources, and must be managed during operations.

The timing system consists of three primary functionalities: RF reference distribution, optical reference distribution, and a phase-locked loop (PLL). This PLL may use either the RF or the optical reference as a feedback source. The RF allows for phase comparisons over a relatively wide range, albeit with limited resolution, while the optical reference enables very fine phase comparison (down to attoseconds), but with limited operational range.

These systems must be managed using high levels of automation. Much of this automation is done via high-level applications developed in EPICS. The beamline users are presented with relatively simple interfaces that streamline operation and abstract much of the system complexity away. The system provides both PyDM GUIs as well as python interfaces to enable time delay scanning in the LCLS-II DAQ.

## BACKGROUND

### What is "Timing"?

LCLS is a pulsed machine, capable of operating at a variety of fixed frequencies, depending on machine state and experiment need. The X-ray pulses emitted by LCLS are on the order of 100's to 10's of femtoseconds (fs) in pulse duration. The optical lasers used to perform experiments at LCLS often have pulse durations on the same order, leading to extremely tight synchronization requirements for optical pump-probe experiments [1]. The precision timing system provides a level of synchronization beyond that which can be achieved with the event timing system. This system is capable of not only synchronizing the fs pulses, but also controlling relative arrival time of the pump laser with respect to the X-rays, enabling optical pump-probe experiments that investigate transient states of matter [1].

### Changes for LCLS-II

The LCLS-I laser locker system has been described previously [2]. This system was constructed using a combination of commercial and SLAC-designed hardware, and uses a combination of EPICS [3] and Python [4] code to create a high level application for facility users. This system has

been used quite successfully since its inception, with every hutch at LCLS utilizing these systems for Ti:Sapphire laser experiments. However, to meet the precision timing needs of LCLS-II, a new system had to be developed in order to accommodate changes in both the timing requirements as well as changes in the physical layout and deployment of the laser systems. The major timing changes from LCLS-I to LCLS-II are outlined below.

- RF reference. The LCLS-II precision timing system uses the Phase Reference Line (PRL) RF reference used by the LCLS-II accelerator for synchronization of accelerator hardware.
- Laser technology. The LCLS-II laser systems are based on OPCPA technology, rather than the Ti:Sapphire lasers used in LCLS-I, allowing for higher laser repetition rate.
- Laser delivery. Rather than installing a laser system in every hutch for every interaction point, shared laser systems are used in LCLS-II. An evacuated, remotely controlled laser beam transport system is used to deliver beam from any of 4 possible laser bays in the NEH to any experiment interaction point.
- Optical locking. The LCLS-II precision timing system offers two modes of phase locking: RF locking, and optical locking.

## SYSTEM FUNCTIONALITY

The LCLS-I precision timing system had a high level of automation, providing automated system operation and a single "target time" control point, controlling the arrival time of the optical laser pulses with respect to the X-ray pulses. The changes from LCLS-I to LCLS-II described above have necessitated commensurate changes in the LCLS-II precision timing system. The LCLS-II precision timing system implements many of the features present in LCLS-I, as well as others that are completely new for LCLS-II. The major features of the LCLS-II precision timing system are listed below.

- RF reference distribution.
- Frequency locking.
- RF phase locking.
- RF bucket jump detection and correction.
- Target time control and read-back.
- Automated lock transition.

\* WORK SUPPORTED BY U.S. D.O.E. CONTRACT DE-AC02-76SF00515

<sup>†</sup> tjohnson@slac.stanford.edu

# IN THE MIDST OF FUSION IGNITION: A LOOK AT THE STATE OF THE NATIONAL IGNITION FACILITY CONTROL AND INFORMATION SYSTEMS

M. Fedorov, L. Beaulac, A. Casey, J. Castro Morales, J. Dixon, C. M. Estes, M. S. Flegel, S. Heerey, V. Miller Kamm, B. Patel, M. Paul, N. Spafford, A. Barnes, V. Gopalan, R. Lacuata, J. L. Vaher, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550

## Abstract

The National Ignition Facility (NIF) is the world's largest and most energetic 192-laser-beam system which conducts experiments in High Energy Density (HED) physics and Inertial Confinement Fusion (ICF). In December 2022, the NIF achieved a scientific breakthrough when, for the first time ever, the ICF ignition occurred under laboratory conditions. The key to the NIF's experimental prowess and versatility is not only its power but also its precise control. The NIF controls and data systems place the experimenter in full command of the laser and target diagnostics capabilities. The recently upgraded Master Oscillator Room (MOR) system precisely shapes NIF laser pulses in the temporal, spatial, and spectral domains. While the increasing neutron yields mark the NIF's steady progress towards exciting experimental regimes, they also require new mitigations for radiation damage in control and diagnostic electronics. With many NIF components approaching 20 years of age, a Sustainment Plan is now underway to recapitalize NIF subsystems, including controls hardware and software, to assure operations through 2040. Keywords: fusion, ignition, status, upgrade, long-term support.

## INTRODUCTION

The National Ignition Facility is a large (3 football fields) and complex (192 laser beams) experimental physics system (Fig. 1) [1]. Experiments at NIF support several programmatic missions: Stockpile Stewardship, Discovery Science, National Security Applications, and Inertial Confinement Fusion (ICF). For ICF, thermonuclear ignition has been the long-term goal of the facility, defined by the U.S. National Academy of Sciences as producing more energy from the DT target fusion than the laser energy on the target. The NIF was pursuing the ignition goal for almost 10 years, and it proved to be a scientific and engineering challenge.



Figure 1: NIF building layout.

A breakthrough was achieved in December 2022, when an experiment at NIF produced 3.15 MJ of fusion energy, more than the laser energy on the target of 2.05 MJ [2], thus achieving the fusion ignition. The achievement was broadly covered by scientific and news media, Fig. 2.

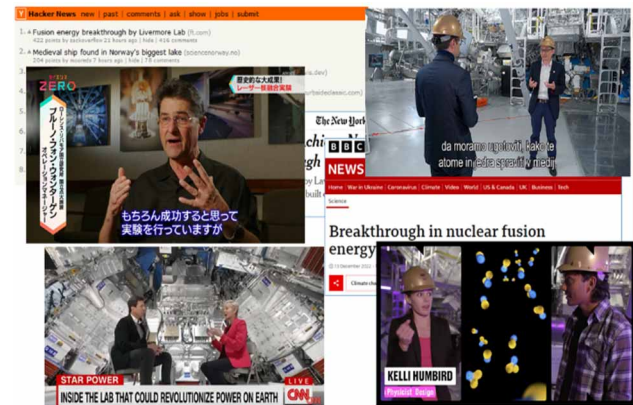


Figure 2: NIF Ignition covered by the news media.

NIF is operated 24x7 by a shift of the 12-14 Control Room operators with the help of the Integrated Computer Control System (ICCS). Over 66,000 devices with rich APIs are distributed over 2,300 front-end-processors (FEPs) and embedded controllers (ECs). NIF experiments are structured around laser shots. Each shot takes 4-8 hours and involves the control system executing over 2 million device operations.

## ROLE OF CONTROL SYSTEMS

The NIF has achieved its remarkable fusion ignition results due to its precise control over energy, optics, and target parameters. This report highlights the pivotal role of control systems in ensuring the success of NIF's experiments. Figure 3 illustrates the multitude of parameters which are controlled and monitored by NIF systems.

# ONGOING IMPROVEMENTS TO THE INSTRUMENTATION AND CONTROLS SYSTEM AT LANSCE\*

M. Pieck<sup>†</sup>, C. D. Hatch, H. A. Watkins, E. E. Westbrook  
Los Alamos National Laboratory, Los Alamos, USA

## Abstract

Recent upgrades to the Instrumentation and Controls System at Los Alamos Neutron Science Center (LANSCE) have significantly improved its maintainability and performance. These changes were the first strategic steps towards a larger vision to standardize the hardware form factors and software methodologies. Upgrade efforts are being prioritized through a risk-based approach and funded at various levels. With a major recapitalization project finished in 2022 and modernization project scheduled to start possibly in 2025, current efforts focus on the continuation of upgrade efforts that started in the former time frame and will be finished in the latter. Planning and executing these upgrades are challenging, considering that some of the changes are architectural in nature; however, the functionality needs to be preserved while taking advantage of technology progressions. This is compounded by the fact that those upgrades can only be implemented during the annual 4-month outage. This paper will provide an overview of our vision, strategy, challenges, recent accomplishments, as well as future planned activities to transform our 50-year-old control system into a modern state-of-the-art design.

## INTRODUCTION

Los Alamos Neutron Science Center (LANSCE) has been in operations for over 50 years. The multifunctional facility has grown over the years and has now five distinct state-of-the-art experimental facilities which provide the scientific community with intense sources of protons and neutrons, with the capability of performing experiments supporting civilian and national security research [1].

LANSCE was one of the first linear accelerators that had computerized control. Over the years some of the instrumentation & controls equipment (ICE) interfacing with a variety of beam line devices have been added and upgraded. As a result, the facility accumulated a wide variety of ICE hardware form factors and software methodologies that challenge the maintainability and longevity of the LANSCE Control System (LCS).

To address these concerns, a methodical and relatively well-funded upgrade effort between 2011 and 2022 focused on the replacement of the original 50-year-old RICE (Remote Instrumentation and Control Equipment) system. It was replaced with a modern customized control system

in stages during each annual 4-month accelerator outage period. Despite the recent modernization success, the LCS still contains equipment up to 40+ years old and requires significant long-term investments beyond current base funding levels to upgrade it to current controls hardware and software technology maturation levels [2, 3].

Once that has been achieved, the focus should shift from upgrading obsolete and past-end of life equipment to active forward-looking life-cycle management. Accelerator instrumentation and control systems (from here on called simply “control system”) are modern information technology systems, and as such require constant investments. Risks associated with neglecting active lifecycle management increases the risk of disruption to operations, extended downtime, increase of system complexity which will increase maintenance cost, and loss of upgrade paths resulting in potentially costly greenfield recapitalization projects.

## CONTROL SYSTEM DESCRIPTION

To understand the LCS upgrade and maintenance challenges, it helps to present some of the boundary conditions. First, we will describe the LCS size/complexity, and in the next section (Control System Group) we will discuss the number of people supporting it, their background, and funding levels.

The Accelerator Operations and Technology division – Instrument and Controls group (AOT-IC) is responsible for maintaining LCS which utilizes EPICS (Experimental Physics and Industrial Control System). Geographically, AOT-IC’s responsibilities include all ICE starting from the two Injectors through the nearly half-mile long accelerator, a proton storage ring, and the distribution lines to the experimental facilities including some target stations but not the experimental end stations.

Technically, AOT-IC is responsible for the computer-based system that gathers and analyzes industrial process and real-time data to monitor and control accelerator relevant equipment that deals with critical and time-sensitive information or events. This includes a distributed event-based timing system to trigger actions at different locations at the same time and in sequence with predefined time intervals, while synchronizing the local time at different locations with high precision and timestamping events at different locations to analyze what happened first.

In addition, the group is responsible for the LCS network star-like infrastructure including firewalls, core switches, distribution, and leaf switches as well as the CAT6 and fiber optics cable plant. Moreover, AOT-IC maintains its own server infrastructure. Finally, the group develops, deploys, and maintains all software application needs and custom services.

\* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001).

<sup>†</sup> pieck@lanl.gov, LA-UR-23-31088

# MAINTAINING A HYBRID CONTROL SYSTEM AT ISIS WITH A VSYSTEM/EPICS BRIDGE\*

K. R. L. Baker<sup>†</sup>, I. D. Finch, M. Romanovschi  
ISIS Neutron and Muon Source, Didcot, United Kingdom

## Abstract

The migration of the controls system for the ISIS accelerator from VSystem to EPICS presents a significant challenge and risk to day-to-day operations. To minimise this impact throughout the transition, a software bridge between the two control systems has been developed that allows the phased porting of HMIs and hardware. The hybrid VSystem and EPICS system also allows the continued use of existing feedback control applications that now require interaction between both control systems, for example the halo steering operation in Target Station 1. This work describes the implementation of this bridge, referred to as PVEcho, for the mapping of VSystem channels to EPICS PVs and vice versa. The position within the wider ISIS controls software stack is outlined as well as how it utilises Python libraries for EPICS. Finally, we will discuss the software development practices applied that have allowed the bridge to run reliably for months at a time.

## INTRODUCTION

VSystem is a commercial product distributed by Vista Control Systems [1] which saw first use for the accelerators at the ISIS Neutron and Muon Source [2, 3] in 1998. It is run using the OpenVMS [4] operating system on Itanium servers. Like most other control systems, components in the accelerator are distributed into different channels and databases. Metadata associated with them is summarised in database (.adb) files which determine the data type of the channel as well as alarm configurations, display limits and labels, among other things.

The imminent obsolescence of the Itanium servers on which VSystem is run prompted a recent evaluation of alternative control systems that could be implemented at the ISIS accelerators. The evaluation concluded that the ISIS accelerator should be migrated over to EPICS to reap the benefits of being part of a wider community of developers and users of the control system, as well as access expertise of other facilities on the Rutherford Appleton Laboratory site.

The nature of ISIS as an operational facility mandates that the transition from VSystem to EPICS needs to be done in a way that minimizes impact on user runs. It was therefore decided that a phased porting of control of hardware is the desired option for the transition [5–7]. This way, conversion of Graphical User Interfaces (GUIs), referred to as control screens at ISIS, could be decoupled from the porting of control hardware and updates to systems could progress in

the background. Similarly, it would mean that new systems and services could be slowly introduced part way through the transition to allow a period of adjustment and training for machine operators [8].

In order to complete the transition in this phased manner, the current state of VSystem channels needs to be replicated within EPICS to mirror the behaviour of the machine in real-time as EPICS process variables (PVs). This work outlines the overall purpose and structure of this bridging software, as well as explaining where it sits within the larger software stack of the ISIS accelerator controls. The software is named PVEcho [9] and is divided into two halves. The first is VSystem to EPICS, referred to as V2E, which replicates channels where VSystem is still the source of truth as EPICS PVs. The second is the reverse. EPICS to VSystem (E2V) provides updates from PVs where the hardware has already been ported to EPICS to Vsystem channels which do not link to hardware but are still available for operators and software to interact with. We refer to these channels as 'soft' channels. More detail about each of these services will be provided in the following sections.

## SYSTEM OVERVIEW

The current ISIS software stack consists of a mix of what we have operating currently (the OpenVMS servers and VSystem) and what we are moving towards, as well as a variety of applications that bridge the gap between the two. Recent work has involved development of systems that allow us to interface more easily with Vsystem. These services are deployed onto a cluster of Linux servers as Docker containers in Swarm mode [10] (hereafter referred to as Docker Swarm), as can be seen in Fig. 1.

The aforementioned services include an MQTT [11] broker that allows us to stream updates to VSystem channels via topics to which clients subscribe. An application running natively in the OpenVMS server transmits messages to readback topics, as well as listening for messages published on an equivalent set topic. This allows a client to set values in VSystem from outside of the OpenVMS system.

Likewise, we host a NoSQL CouchDB database to store metadata associated with the Vsystem channels. Usually stored in the proprietary .adb files, a service has been built that scans the live Vsystem channels at regular intervals (approximately every 15 minutes) and updates the CouchDB database with any new contents such as modified alarm limits. This application makes the configuration of the channels available to developers via read only access to the CouchDB instance.

PVEcho leverages both of these services extensively. Both

\* Work supported by Science and Technology Facilities Council

<sup>†</sup> k.baker@stfc.ac.uk

# CONTINUOUS MODERNIZATION OF CONTROL SYSTEMS FOR RESEARCH FACILITIES

K. Vodopivec, K. S. White, Oak Ridge National Laboratory, Oak Ridge, USA

## Abstract

The Spallation Neutron Source at Oak Ridge National Laboratory has been in operation since 2006. In order to achieve high operating reliability and availability as mandated by the sponsor, all systems participating in the production of neutrons need to be maintained to the highest achievable standard. This includes the SNS integrated control system, comprising of specialized hardware and software, as well as computing and networking infrastructure. While machine upgrades are extending the control system with new and modern components, the established part of control system requires continuous modernization efforts due to hardware obsolescence, limited lifetime of electronic components, and software updates that can break backwards compatibility. This article discusses challenges of sustaining control system operations through decades of facility lifecycle, and presents a methodology used at SNS for continuous control system improvements that was developed by analyzing operational data and experience.

## INTRODUCTION

The Spallation Neutron Source (SNS) is a neutron scattering research facility with the world's most powerful pulsed neutron source. The facility started operation in 2006 and is currently undergoing a major Proton Power Upgrade (PPU) that will increase beam power from 1.4MW to 2.8MW when fully commissioned. The SNS Controls Integration group is responsible for global systems and control system infrastructure to form an Integrated Control System (ICS) for the SNS accelerator and target complex. The ICS has been largely successful at supporting SNS operations since initial commissioning and has played a key role of achieving overall SNS beam availability of over 90 %.

To achieve a sustained high availability, continuous maintenance and modernization efforts are required for all components of the ICS. The SNS Controls Integration group is performing regular activities during the scheduled maintenance periods. These include ensuring controlled temperature environment by cleaning or replacing air filters and fans for increased air flow through the electrical circuits, tracking inventory operational time and replacing components at about 80 % expected lifetime as provided by vendor or based on data collected at SNS, maintaining a healthy amount of fully functional spares with at least 10 % based on the number of installed components and in some cases much higher when components are becoming obsolete.

The preventive maintenance activities help reduce unexpected failures and allow for overall high availability. But the expected SNS facility lifetime is several more decades, and these activities are not addressing issues like

electronics and software obsolescence, technological advances, increased demand for data, or functionality upgrades. The rest of the document will describe these challenges in detail and present modernization efforts by the SNS Controls Integration group to address them.

## HIGH AVAILABILITY OF COMPUTING AND NETWORKING INFRASTRUCTURE

The computing and networking infrastructure are an essential part of every control system. Networking allows devices to be connected into one integrated control system and is the carrier of all the controls information including delivering alarms in a timely manner, reliably transporting archive data from source to data storage, and providing real-time monitoring and control of any aspect of the system. Computing is needed to host various applications for automation, data acquisition, diagnostics, operator tools and many others. Performance and throughput requirements are usually low for the control system environment and are certainly within reach of modern technology. But recent advances in technologies allow improvements towards reliability of these systems.

Originally the SNS ICS network infrastructure was divided in 3 logical layers as depicted in Fig. 1. The top layer consisted of a single network node called the network core switch. Its job is to define the network topology and route the traffic between segments as well as to connect to an external network through a firewall. Next down is a distribution layer whose role is to provide enough links to connect all end-point network switches on the next layer. End-point network switches' responsibility is to connect network devices to the network. With about 3,000 network devices and about 600,000 sensors and actuators on the SNS ICS network, neither throughput nor latency are a challenge for our 1 Gbps network. However, availability of every network switch and general network health are of the




Figure 1: Complete redundancy and isolation of the SNS network for Integrated Control System allows high availability even when complete separation might be needed.



# EPICS DEPLOYMENT AT FERMILAB\*

P. Hanlet<sup>†</sup>, M. Gonzalez, J. Diamond, f I; n . Martin  
 Fermi National Accelerator Laboratory, Batavia, Illinois, USA

## Abstract

Fermilab has traditionally not been an EPICS house; as such expertise in EPICS is limited and scattered. PIP-II will be using EPICS for its control system. When in operation, it will need to interface with the existing, modernized (see ACORN) legacy control system. Treating EPICS controls at Fermilab as a green field, we have developed and deployed a software pipeline which addresses these needs and presents to developers a tested and robust software framework, including template IOCs from which new developers can quickly deploy new front ends, aka IOCs. In this presentation, motivation for this work, implementation of a continuous integration/continuous deployment pipeline, testing, template IOCs, and the deployment of user services/applications will be discussed. This new infrastructure of IOCs and services is being developed and used in the PIP-II cryomodule test-stand; our experiences and lessons learned will be also be discussed.

and conventional facilities. Space is also reserved in the LINAC for 2 additional SRF cryomodules (CMs). For more on PIP-II, see [3].




Figure 2: PIP-II scope: new LINAC and beam transfer line to Booster.

## INTRODUCTION

Fermi National Accelerator Laboratory, or Fermilab, in Batavia, Illinois, USA is constructing a new superconducting RF (SRF) linear accelerator (LINAC) with twice the energy of the existing LINAC and significantly higher power. This LINAC, known as PIP-II, will power the rest of the Fermilab accelerator complex, generating the world’s most intense high-energy neutrino beam, as well as providing beam to other experiments and test beams, see Fig. 1.




Figure 1: Fermilab’s accelerator chain and experiments.

The capabilities of PIP-II [1, 2] are to provide an 800 MeV proton beam of 1.2 MW using a superconducting RF LINAC, see Fig. 2. The beam is CW-compatible and customizable for a variety of user requirements, and is upgradeable to multi-MW power. The scope of PIP-II includes a beam transfer line to the existing Booster ring and accelerator complex upgrades to the Booster, Recycler, Main Injector,

## GOALS FOR DEPLOYMENT

Though EPICS [4] has been long been deployed in small pockets at Fermilab, one could say that Fermilab has never been an EPICS shop. As such, the expertise is limited. Furthermore the controls engineers who support the remainder of the accelerator complex are stretched thinly. Therefore, for a small controls team, we want to deploy EPICS in a sustainable way by requiring 1) a robust build of infrastructure, 2) automated build procedures, 3) extensive testing, and 4) minimal functionality to automate testing, deployment, and monitoring of IOCs.

We are treating EPICS deployment as a green field to simplify development for non-experts. This follows from the fact PIP-II, and likely new components from ACORN [5], will not rely on old hardware, such as VME and CAMAC. With this in mind, we use current versions of EPICS software, both for base and Support modules, with an expectation that new software versions will continue to operate for the foreseeable future. This is an ideal environment to adopt modern computing procedures, in particular employing a Continuous Integration/Continuous Deployment (CI/CD) pipeline. This allows us to build all of EPICS base and EPICS Support/Modules as “standard” Fermilab versions and deploy it on an NSF host for all to use. Furthermore, we have developed “template IOCs” which have the boiler plate functionality which we require for all IOCs on the Fermilab controls network.

In addition to standardizing code, we are exploring standardizing hardware platforms for hosting IOCs. We are using Buildroot to build for raspberry pi, and two System on Module (SoM) platforms: the arm/Cyclone-V and arm/Arria-10 platforms. As of the writing of this document, we are adding

\* This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

<sup>†</sup> hanlet@fnal.gov

## 15 YEARS OF ALICE DCS

P. Chochula<sup>†,1</sup>, A. Augustinus<sup>1</sup>, P. M. Bond<sup>1</sup>, A. Kurepin<sup>1,4</sup>, M. Lechman<sup>1,3</sup>, O. Pinazza<sup>1,2</sup>, D. Voscek<sup>1</sup>  
<sup>†</sup>corresponding author; <sup>1</sup>CERN, Geneva (CH); <sup>2</sup>INFN, Sezione di Bologna, Italy;  
<sup>3</sup>University of the Witwatersrand, Johannesburg, South Africa; <sup>4</sup>Affiliated with an Institute Covered  
by a Cooperation Agreement with CERN, Geneva, Switzerland

### Abstract

The ALICE experiment studies ultrarelativistic heavy-ion collisions at the Large Hadron Collider at CERN. Its Detector Control System (DCS) has been ensuring the experiment safety and stability of data collection since 2008. A small central team at CERN coordinated the developments with collaborating institutes and defined the operational principles and tools. Although the basic architecture of the system remains valid, it has had to adapt to the changes and evolution of its components.

The introduction of new detectors into ALICE has required the redesign of several parts of the system, especially the front-end electronics control, which triggered new developments.

Now, the DCS enters the domain of data acquisition, and the controls data is interleaved with the physics data stream, sharing the same optical links. The processing of conditions data has moved from batch collection at the end of data-taking to constant streaming. The growing complexity of the system has led to a big focus on the operator environment, with efforts to minimize the risk of human errors.

This presentation describes the evolution of the ALICE control system over the past 15 years and highlights the significant improvements made to its architecture. We discuss how the challenges of integrating components developed in tens of institutes worldwide have been mastered in ALICE.

## THE ALICE DCS ARCHITECTURE

### *The ALICE Experiment*

The ALICE experiment [1] involves the deployment of a massive detector system at the Large Hadron Collider (LHC), positioned approximately 65 meters beneath the Earth's surface. With dimensions measuring 15x15x20 m<sup>3</sup> and a total weight of 11,000 tons, it stands as a remarkable feat of engineering, posing formidable operational challenges.

ALICE has been designed as a modular device, comprising 15 subdetectors employing diverse technologies to cover the entire spectrum from solid-state silicon modules to gaseous detectors. The operation of these subdetectors is notably demanding, given the distinct operational conditions within the rigorous ALICE environment. For instance, the TPC field cage, housing 90 cubic meters of gas, must maintain a stable temperature of 20 degrees Celsius with a precision of 0.1 degrees Celsius, even in proximity to electronics modules emitting several

kilowatts of thermal energy and adjacent to lead crystals held at -22 degrees Celsius.

The majority of ALICE's detectors are installed within a 9,000-ton solenoid magnet, necessitating their operation within a strong magnetic field. During LHC operations, the cavern housing ALICE remains sealed, with access for maintenance being granted only under exceptional circumstances.

### *ALICE Distributed Control System Architecture*

The initiation of ALICE DCS traces back to 2008, coinciding with the first delivery of particle beams by the LHC to experimental setups. This milestone marked the conclusion of nearly seven years of preparatory work and testing [2]. Since then, an uninterrupted DCS service has been responsible for upholding the stability and safety of the detectors.

The current DCS architecture is a product of years of experience refined during its operational history. While it has adapted to evolving requirements, the fundamental architectural principles remain steadfast. The design places significant emphasis on the autonomous operation of system components, employs a rigorous hierarchical structure, and incorporates an intuitive user interface equipped with mechanisms to preempt human errors. These fundamental elements are complemented by extensive standardization across all control domains.

At the core of the DCS lies the WINCC OA SCADA system by Siemens [3], a configuration in line with other CERN experiments. It is extended through the JCOP framework, which offers customizations tailored to the CERN environment. Additional framework layers are developed by ALICE teams to introduce ALICE-specific extensions, including detector safety protocols, responses to beam conditions, and user interfaces, among others.

### *Control System*

Each ALICE subdetector is designed as an autonomous unit, capable of independent operation within the DCS. To minimize interdependencies among detectors and to ensure the integrity of hardware modules, such as power supplies and crates, sharing is prohibited. An initial architectural decision was the subdivision of control systems into subsystems, wherein devices with similar functions were grouped and isolated from other subsystems. The specific requirements of various detectors led to physical distinctions among subsystems. For instance, the powering system necessitates finer granularity to differentiate between electronics power and detector power. This separation allows high-voltage power supplies to be located outside the experimental cavern, while low-voltage

# MODULAR AND SCALABLE ARCHIVING FOR EPICS CHANNEL ACCESS AND GENERAL TIME SERIES USING SCYLLA AND RUST

D. Werder<sup>§</sup>, T. Humar<sup>‡</sup>, Paul Scherrer Institut, 5232 Villigen PSI, Switzerland

## Abstract

To unify, simplify and improve our archiving setup we have developed a more modular archiver architecture based on an existing database with industry adoption and enterprise support, combined with additional services which ingest data from EPICS sources into the database and serve user requests for channel data from the database. The prototype is currently tested at the Swiss Free Electron Laser (SwissFEL) and planned to run in production mode from the beginning of 2024. Additional support for more input sources beside EPICS sources is also under development.

## CURRENT ARCHIVING SITUATION

At the Paul Scherrer Institut we operate at the moment several different software products to archive our EPICS [1] process variables (PV) as well as data from non-EPICS sources. At the Swiss Light Source (SLS) we use the EPICS Channel Archiver [2], while at the SwissFEL [3] we use the EPICS Archiver Appliance [4] for the EPICS channels as well as the SwissFEL-Databuffer to buffer beam-synchronous data at the machine pulse rate, where the difference between the archiver and buffer being only that the buffer retains the data only for a limited time. Also HIPA, Proscan and several other smaller systems operate archivers in differing versions.

The operation of these heterogeneous setups binds resources, requires extensive expertise and makes the operation more difficult than it needs to be. Also, the current products are not easy to scale and hard to enhance for replication and high availability. Furthermore, it is often difficult, unspecified or impossible to access data from an external process in a defined, synchronized way, and also the file formats themselves are often used in rather small communities so that the available tooling can be limited.

At this time, the upgrade program for SLS 2.0 is now ongoing, which will among other changes bring a substantial increase in the number of archived channels. To meet the evolving requirements and to simplify and unify our archiving and buffering architecture, we have conducted a design study with the purpose to find a more modular and hopefully more sustainable approach.

## DESIGN GOALS AND ARCHITECTURE

Instead of the current rather monolithic archiver solutions we would like to de-couple and modularize the setups. Our solution should be more easy to scale, should

offer tunable replication and availability. We would like to be able to access the data concurrently in a well-defined way with more clear interfaces. The core storage engine should therefore be a dedicated database product which runs in its own process(s) and offers access via a network protocol. At the same time, we want to avoid having to maintain a custom solution and prefer an existing product which sees a wider usage also in industry.

The other components around the database to handle the different source types, including now Channel Access and in the future Beam-Synchronous-Readout and PV-Access, should be able to interact independently and concurrently with the common database. We give a brief overview of the architecture, after which we describe the individual components and their interaction.

## Overview

The first component of this architecture, named Ingest service, handles the communication with the EPICS Channel Access Input Output Controllers (IOCs), maintains and monitors the channels (“virtual circuits”), and inserts the received updates from the PVs into the database.

A second component, named Retrieval service, is responsible to serve user requests for channel data and can deliver full events as well as aggregated and time-binned data. It communicates with the database as well. Other direct interaction between Ingest and Retrieval is not required, the database is the only connecting interface between the services.

As our database we have chosen Scylla [5]. Channel information and some other data which lends itself more to transactional databases is kept in Postgres [6].

## Ingest service

The Ingest service [7] takes a list of channel names as input. It finds the IP addresses of the corresponding IOCs, maintains communication with the IOCs via the EPICS Channel Access Protocol [8], opens the channels on the IOCs and monitors for updates.

Updates to the PV values are inserted into the database, at which point we can also already maintain a reduced, aggregated and/or binned time series which is meant for a typically longer retention period.

The status of the channel and of the TCP connection is monitored and written as a separate time series to the database. To distinguish communication silence due to a lack of PV updates from a broken IOC connection, the Ingest service can issue a Channel Access Echo message.

<sup>§</sup> dominik.werder@psi.ch

<sup>‡</sup> tadej.humar@psi.ch

# DATA MANAGEMENT FOR TRACKING OPTIC LIFETIMES AT THE NATIONAL IGNITION FACILITY

R. D. Clark\*, L. M. Kegelmeyer  
Lawrence Livermore National Laboratory, Livermore, USA

## Abstract

The National Ignition Facility (NIF), the most energetic laser in the world, employs over 9000 optics to reshape, amplify, redirect, smooth, focus, and convert the wavelength of laser light as it travels along 192 beamlines. Underlying the management of these optics is an extensive Oracle database storing details of the entire life of each optic from the time it leaves the vendor to the time it is retired. This journey includes testing and verification, preparing, installing, monitoring, removing, and in some cases repairing and re-using the optics. This talk will address data structures and processes that enable storing information about each step like identifying where an optic is in its lifecycle and tracking damage through time. We will describe tools for reporting status and enabling key decisions like which damage sites should be blocked or repaired and which optics exchanged. Managing relational information and ensuring its integrity is key to managing the status and inventory of optics for NIF.

## INTRODUCTION

The National Ignition Facility (NIF), the most energetic laser in the world, facilitates experiments that create temperatures and pressures that exist at the center of stars, giant planets and nuclear weapons. Roughly the size of three American football fields placed next to each other, NIF has two identical laser bays each containing two clusters of 48 beamlines. Laser light is propagated along these beamlines where mirrors, spatial filters and other devices amplify it from 1 billionth of a joule to 4 million joules and ensure the beams are correctly shaped. The beams are then grouped into 48 quads of 2x2 arrays and pass through the final optics assembly where they are converted from infrared to ultraviolet energy before entering the target chamber.

The laser energy travels nearly a kilometer before focusing on a 2mm fuel capsule (roughly the size of two grains of rice placed side by side). Along that path are thousands of serialized parts including over 6200 exchangeable frames (called LRUs) that house over 9000 optics that reshape, amplify, redirect, smooth, focus, and convert the wavelength of laser light as it travels along the 192 beamlines. The total surface area of the NIF optics is three-quarters of an acre - 40 times the surface area of the giant Keck telescope in Hawaii [1]. The optics are vital to NIF operations and are monitored and maintained throughout their lifetime.

## DATABASE

Thousands of pieces of data are captured and stored on each of the NIF optics in an extensive Oracle Relational Database Management System (RDBMS) with Real Application Clusters (RAC). The RAC system allows multiple database instances to run in parallel accessing the same database storage [2]. The optics database has hundreds of tables, managing almost 300 GB of data, that details the entire life of each optic. Oracle Application Express (APEX) is used to provide a User Interface (UI) application for managers and operations personnel to access required data from across several different applications. The Production Optics Reporting and Tracking (PORT) tool consolidates this data, provides processing functionality, and generates customized reports to address the health and status of optics.

## OPTIC PROCESSING

### Optic States

An optic's journey includes being tested and verified, then prepared, installed, monitored, removed, and in some cases repaired and reused (Figure 1). Here we address data structures and processes that enable storing information about each step like where an optic is in its lifecycle, the calibration data that verifies the optic meets required specifications, and tracking damage through time. We will describe tools for reporting status and enabling key decisions like which damage sites should be blocked or repaired and which optics exchanged. Managing relational information and ensuring its integrity is key to managing the status and inventory of optics for NIF.

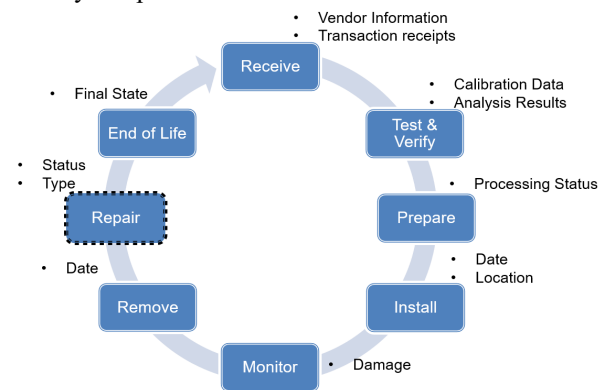


Figure 1: Optic Lifetime.

\*clark94@llnl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# IMPROVING OBSERVABILITY OF THE SCADA SYSTEMS USING ELASTIC APM, REACTIVE STREAMS AND ASYNCHRONOUS COMMUNICATION\*

I. Khokhriakov<sup>†</sup>, V. Mazalova, O. Merkulova, DESY, Hamburg, Germany

## Abstract

As modern control systems grow in complexity, ensuring observability and traceability becomes more challenging. To meet this challenge, we present a novel solution that seamlessly integrates with multiple SCADA frameworks to provide end-to-end visibility into complex system interactions. Our solution utilizes Elastic APM to monitor and trace the performance of system components, allowing for real-time analysis and diagnosis of issues. In addition, our solution is built using reactive design principles and asynchronous communication, enabling it to scale to meet the demands of large, distributed systems. This paper will describe our approach and discuss how it can be applied to various use cases, including particle accelerators and other scientific facilities. We will also discuss the benefits of our solution, such as improved system observability and traceability, reduced downtime, and better resource allocation. We believe that our approach represents a significant step forward in the development of modern control systems, and we look forward to sharing our work with the community at ICALEPCS 2023.

## INTRODUCTION

With the advancement of technology, modern scientific instruments have evolved to exhibit remarkable levels of complexity. This evolution is characterized by complicated architectures, multifaceted functionalities, and the integration of diverse subsystems and technologies. Alongside this development, there is a notable increase in data acquisition rates and volumes which are boosted by the pursuit of higher precision and the exploration of broader scientific phenomena [1, 2, 3, 4].

This increase in data acquisition creates a challenge, as instruments generate vast amounts of data at an accelerated pace, making it necessary to create advanced solutions for efficient data handling, processing, and storage. The complex nature of these instruments, coupled with big and high-velocity data they produce, underscores the imperative for enhanced observability and traceability within Supervisory Control and Data Acquisition (SCADA) systems. Addressing these challenges is crucial for ensuring the reliability and efficacy of scientific experiments and for unlocking new approaches in science.

Navigating the complexity of modern scientific instruments and managing data they produce demands a focus on the observability of SCADA systems. Observability [5, 6], in this context, is not just a luxury but a critical requirement. It enables operators to gain deep insights into system performance, identify anomalies in real-time, and troubleshoot issues promptly, thereby ensuring reliability of the entire control system.

Given the vital role of SCADA systems in managing and monitoring complex processes, a lack of observability can lead to inefficiencies, operational disruptions, and compromised data quality. Furthermore, in an era when scientific experiments are rapidly progressing, the ability to observe and understand the inner workings of SCADA systems is crucial for optimizing system performance and adapting to evolving requirements.

Thus, enhancing the observability of SCADA systems emerges as an essential undertaking, driving the development of innovative solutions and methodologies to meet this imperative need. This paper introduces a novel approach, leveraging Elastic APM, reactive design principles, and asynchronous communication, to significantly improve the observability and traceability of SCADA systems across various scientific facilities [7, 8].

In the subsequent sections of this paper, we are going to look into two distinctive approaches to incorporating observability into SCADA systems, addressing the unique challenges posed by the increasing complexity and high data acquisition rates of modern scientific instruments.

- 1. Dedicated Microservice Approach:** The first approach we explore involves the development and implementation of a dedicated microservice designed to listen to messages as they pass through the system. By forming transactions from these messages, this method aims to provide a granular view of system interactions, thereby enhancing the ability to monitor, trace, and diagnose potential issues within the SCADA environment.
- 2. Integrating Agents into Software Components:** The second approach takes a more integrated route, embedding agents directly into corresponding software components of the SCADA system. This method seeks to ensure a seamless and comprehensive monitoring capability, offering real-time insights and quick responsiveness to any anomalies or system performance deviations.

\* work supported by the consortium DAPHNE4NFDI

<sup>†</sup> igor.khokhriakov@desy.de

# THE CMS DETECTOR CONTROL SYSTEMS ARCHIVING

W. Karimeh<sup>1</sup>, CERN, Geneva, Switzerland

C. Vazquez, CERN, Geneva, Switzerland

F. Glege, CERN, Geneva, Switzerland

M. Chamoun, Université Saint-Joseph de Beyrouth, Beirut, Lebanon

<sup>1</sup>also at Université Saint-Joseph de Beyrouth, Beirut, Lebanon

## Abstract

The CMS experiment relies on its Detector Control System (DCS) to monitor and control over 10 million channels, ensuring a safe and operable detector that is ready to take physics data. The data is archived in the CMS Oracle conditions database, which is accessed by operators, trigger, data acquisition, and offline data reconstruction systems. In the upcoming extended year-end technical stop of 2023/2024, the CMS DCS software will be upgraded to the latest WinCC-OA release, which will utilise the SQLite database and the Next Generation Archiver (NGA), replacing the current Raima database and RDB manager. Taking advantage of this opportunity, CMS has developed its own version of the NGA backend to improve its DCS database interface. This paper presents the CMS DCS NGA backend design and mechanism to improve the efficiency of the read-and-write data flow. This is achieved by simplifying the current Oracle conditions schema and introducing a new caching mechanism. The proposed backend will enable faster data access and retrieval, ultimately improving the overall performance of the CMS DCS.

## INTRODUCTION

Aimed at probing the deepest questions of fundamental physics, the Compact Muon Solenoid (CMS) is one of the flagship experiments at CERN's Large Hadron Collider (LHC). The Detector Control System (DCS) is integral to its functionality, enabling safe operations and efficient control of the LHC experiments. From the early stages of design, all LHC experiments have adopted WinCC Open Architecture (OA) [1] from ETM as their default Supervisory Control and Data Acquisition (SCADA) software.

At CMS, a myriad of sensors and control units require unwavering precision, addressed by the distributed and redundant DCS projects [2]. With its real-time data assimilation, the DCS projects play an indispensable role in ensuring operational consistency while also offering invaluable archival records. The vast, real-time data collected by the DCS not only helps in ongoing operations but also serves as an essential historical record, enabling researchers to understand anomalies, enhance the system, and plan future experiments.

Data collected by WinCC OA is archived in an Oracle database: The Conditions Database. The paper describes the evolution of the CMS DCS conditions database over the past 15 years of operations and unveils the latest development: the CMS Next Generation Archiver backend.

## CMS CONDITIONS DATABASE

WinCC OA uses an RDB manager that serves as a bridge between its internal database, RAIMA, and external databases, notably the Oracle DB used at CERN. The conditions database schema, provided by ETM, categorizes data into metadata tables—representing WinCC OA projects internal datapoints—and real data tables, which hold events and alerts records.

At CMS, the vast volume of data in the event and alert tables presented a challenge, making data retrieval particularly time-consuming from the expansive tables that housed these records.

To address these challenges, the CMS conditions database schema was developed on top of the official one, incorporating PL/SQL scripts. Over the years, it has seen significant enhancements, introducing several key features:

### *DPT Tables*

In WinCC OA, data structuring revolves around Data Points (DP), each being categorized under a specific Data Point Type (DPT). Each DP can encompass one or more Data Point Elements (DPEs), where every DPE represents a unique value or state.

To optimize the distribution and manage the extensive load of the events table, the CMS schema employs PL/SQL scripts. These scripts facilitate the creation of new DPT tables, where columns are designated for DPE names, and they systematically channel DPE values into their respective tables.

### *Last Value Tables*

Due to the large size of DPT tables and the need for certain applications to promptly access the most recent values to understand the current state of the detector, a 'last value' (LV) mechanism has been incorporated into the CMS schema using PL/SQL functions.

This approach led to the creation of an LV table corresponding to each DPT table. Each LV table is consistently updated with the latest values for every datapoint element. When a new value is received, it's compared to its predecessor in the LV table. If the two values match, the new one isn't archived in the DPT table. Furthermore, a dead band configuration is introduced for every DPE within the LV table, ensuring the historical DPT tables only store significant data variations and filter out minor deviations.

# ASSONANT: A BEAMLINe-AGNOSTIC EVENT PROCESSING ENGINE FOR DATA COLLECTION AND STANDARDIZATION

Paulo Baraldi Mausbach\*, Eduardo X. Miqueles, and Allan Pinto  
Brazilian Synchrotron Light Laboratory (LNLS),  
Brazilian Center for Research in Energy and Materials (CNPEM),  
Zip Code 13083-100, Campinas, Sao Paulo, Brazil.

## Abstract

Synchrotron radiation facilities comprise beamlines designed to perform a wide range of X-ray experimental techniques which require complex instruments to monitor thermodynamic variables, sample-related variables, among others. Thus, synchrotron beamlines can produce heterogeneous sets of data and metadata, hereafter referred to as data, which impose several challenges to standardizing them. For open science and FAIR principles, such standardization is paramount for research reproducibility, besides accelerating the development of scalable and reusable data-driven solutions. To address this issue, the Assonant was devised to collect and standardize the data produced at beamlines of Sirius, the Brazilian fourth-generation synchrotron light source. This solution enables a NeXus-compliant technique-centric data standard at Sirius transparently for beamline teams by removing the burden of standardization tasks from them and providing a unified standardization solution for several techniques at Sirius. The Assonant implements a software interface to abstract data format-related specificities and to send the produced data to an event-driven infrastructure composed of streaming processing and microservices, able to transform the data flow according to NeXus. This paper presents the development process of Assonant, the strategy adopted to standardize beamlines with different operating stages, and challenges faced during the standardization process for macromolecular crystallography and imaging data at Sirius.

## INTRODUCTION

For decades, data standardization has been one of the main concerns of industry and research institutes toward providing a better solution to find, transport, and interchange information across different computational platforms. Looking at the main progresses in computer science and engineering, retrospectively, we notice that data standardization brings significant progresses in disciplines such as telecommunication networks, processing, database technology, and web technologies. In telecommunication, data standardization plays a crucial role in communication protocols, which are fundamental for exchanging data between computing devices [1]. Data standardization also contributed to advances in database technology, which provides tools for storing, retrieving, updating, and deleting data from a collection of objects. For instance, the DICOM standard has

been established to facilitate data exchange between medical database systems in digital format [2]. Data standardization was also paramount for the emergence of large X-ray-related datasets [3, 4]. Thus, standardized data formats and structures are critical for efficient and effective data management. Data standardization ensures that data stored within a database follows a consistent format and structure. This consistency simplifies data management, reduces errors, and makes it easier to enforce data integrity constraints [5]. Finally, data management systems often interact with various applications, and standardized data format enables seamless integration and data exchange among systems.

Nowadays, with the increasing use of machine learning and data science pipelines in the industry, data standardization is still crucial to enable novel processing technologies by helping the automation of data analysis pipelines and business needs. A common practice of data scientists who works on data lakes repositories, i.e., the pull of structure and unstructured data that came from multiple sources with different structure and format, consists of extracting, transforming, and loading (ETL) such data towards combining them and prepare them in a more suitable data format to facilitate further analysis. While ETL is an effective data integration method, the main drawback of this approach is the generation of specialized scripts for extracting meaningful data for further analysis, which is time-consuming due to the learning curve needed to understand the data and algorithms for efficient mining data. Furthermore, it is challenging monitoring the quality of data in a data lake repository, which is an essential aspect of the Findability, Accessibility, Interoperability, and Reusability (FAIR) principles [6], since the main goal of FAIR is to optimize the reuse of data.

In synchrotron radiation facilities, data management is challenging mainly due to the heterogeneity of data sources, which are devices that compound the beamlines which in turn are designed to perform a specific X-ray experimental technique (e.g., X-ray powder diffraction, nanotomography). Besides, there are some proprietary devices in the beamlines that do not allow the modification of generated data structure, which also imposes several challenges for managing the data. Sirius is a fourth-generation synchrotron source of radiation which are still under development [7]. The Sirius beamlines were designed towards offering to the users a wide range of X-ray experimental techniques by providing a modern scientific instrument. Nowadays, Sirius is in its first phase of operation with 14 experimental stations, of which six are already fully operational and work with different experiment

\* paulo.mausbach@lnls.br

# EXTENDING THE ICAT METADATA CATALOGUE TO NEW SCIENTIFIC USE CASES

A. Götz\*, A. de Maria, M. Bodin, M. Gaonach, ESRF, Grenoble, France  
A. Gonzalez Beltran, P. Austin, K. Phipps, V. Bozhinov, L. Davies, STFC, Harwell, United Kingdom  
R. Krahl, HZB, Berlin, Germany  
M. Al Mohammed, S. Ali Matalgah, SESAME, Allan, Jordan  
R. Cabezas Quirós, ALBA, Barcelona, Spain  
K. Syder, Diamond Light Source, Harwell, United Kingdom  
A. Pinto, Brazilian Synchrotron Light Laboratory, Campinas, Sao Paulo, Brazil

## Abstract

The ICAT metadata catalogue is a flexible solution for managing scientific metadata and data from a wide variety of domains following the FAIR data principles. This paper will present an update of recent developments of the ICAT metadata catalogue and the latest status of the ICAT collaboration. ICAT was originally developed by UK Science and Technology Facilities Council (STFC) to manage the scientific data of ISIS Neutron and Muon Source and Diamond Light Source. They have since been joined by a number of other institutes including ESRF, HZB, SESAME, ALBA and SIRIUS who together now form the ICAT Collaboration<sup>1</sup>. ICAT has been used to manage petabytes (12.4 PBs for ESRF and over 50 PBs for DLS) of scientific data for ISIS, DLS, ESRF, HZB, and in the future SESAME and ALBA and make these data FAIR. The latest version of the ICAT core as well as the new user interfaces, DataGateway and DataHub, and extensions to ICAT for implementing free text searching, a common search interface across Photon and Neutron catalogues, a protocol-based interface that allows making the metadata findable, electronic logbooks, sample tracking, and web-based data and domain specific viewers developed by the community will be presented.

## INTRODUCTION

Science needs data in order to make new discoveries and verify existing theories. Data are therefore essential to science [1] and need to be managed for the short and long term i.e. during and after an experiment. Managing data during the experiment helps provide feedback to the experimentalists, while in the long-term, a record of the data needs to be kept so that scientists can get back to their own data after the experiment, and others can verify and eventually reproduce the results. Preserving data requires having a metadata catalogue for storing all metadata and references to data. The metadata catalogue supports browsing, searching and displaying of relevant metadata, as well as providing access to the data themselves. Over a thousand scientific data repositories exist today built on diverse solutions ranging from bespoke closed source data software to common open source metadata catalogue software. The choice of which solution

to use will depend on the needs of the scientific domains being considered, as well as the technical preferences and available resources of the software engineers. This paper describes the ICAT solution, its architecture, implementation and which sites have adopted ICAT and how they use it.

## ICAT ARCHITECTURE

### Data Model

The metadata stored in ICAT follows a schema [2] based on the Core Scientific Meta-Data Model (CSMD), originally conceived in 2003 [3] and most recently updated in 2013 [4]. CSMD captures high level information about scientific studies and the data that they produce, and contains 27 entities, which have remained fairly stable since their inception. At the core of the ICAT schema (see Fig. 1), includes:

- **Investigations**, which set high level and scientific context. They may correspond to a proposal or visit to a facility.
- **Datasets**, which define context for creating data. They may correspond to a single experiment, measurement or simulation. A single **Investigation** can contain many **Datasets**.
- **Datafiles**, which represent the actual files that make up the data. A single **Dataset** can contain many **Datafiles**.

**Users** are associated with data at the **Investigation** level. Other aspects of the metadata are represented with their own entities and related to these core entities, such as at which **Instrument** an **Investigation** was carried out. It is also worth noting that these entities and their fields can be used to suit the needs of the **Facility** using ICAT. Many are optional which allows the schema to be applied flexibly.

In particular, **Types**, **Parameters** and **Technique** allow a highly configurable method for categorising data and capturing the experimental conditions. Each of the three main entities and the **Sample** is required to have a named **Type** (with the exception of the **Datafile** which has the optional **Datafile-Format**, intended to explicitly capture the file format). All four can possess any number of **Parameters**, which in turn must have a defined **ParameterType** (as shown in Fig 2). The former records a value, the latter what it is measuring

\* andy.gotz@esrf.fr

<sup>1</sup> <https://icatproject.org>



# EFFICIENT AND AUTOMATED METADATA RECORDING AND VIEWING FOR SCIENTIFIC EXPERIMENTS AT MAX IV

D. van Dijken\*, M. Leorato, M. Eguiraun, V. Hardion, M. Klingberg,  
M. Lindberg, V. Da Silva, MAX IV Laboratory, Lund, Sweden

## Abstract

With the advancements in beamline instrumentation, synchrotron research facilities have seen a significant improvement. The detectors used today can generate thousands of frames within seconds. Consequently, an organized and adaptable framework is essential to facilitate the efficient access and assessment of the enormous volumes of data produced. Our communication presents a metadata management solution recently implemented at MAX IV, which automatically retrieves and records metadata from Tango devices relevant to the current experiment. The solution includes user-selected scientific metadata and predefined defaults related to the beamline setup, which are integrated into the Sardana control system and automatically recorded during each scan via the SciFish library. The metadata recorded is stored in the SciCat database, which can be accessed through a web-based interface called Scanlog. The interface, built on ReactJS, allows users to easily sort, filter, and extract important information from the recorded metadata. The tool also provides real-time access to metadata, enabling users to monitor experiments and export data for post-processing. These new software tools ensure that recorded data is findable, accessible, interoperable and reusable (FAIR) for many years to come. Collaborations are on-going to develop these tools at other particle accelerator research facilities.

## INTRODUCTION

Metadata, often referred as data about data, has become more and more important in the scientific field to complete and accompany the results of an experiment. Metadata is often very important to contextualise scientific result and help with the reproducibility of an experiment.

The increasing importance highlights the need to treat metadata with the respect. To collect the metadata in a complete and efficient way, to store it safely and most importantly, in a transparent way for the scientific user, to not increase the burden of work for that already follow the execution of an experiment.

It is also essential to guarantee ways to explore and use such metadata otherwise the whole collection process become meaningless.

Solving this issue has been an important priority to MAX IV with the aim of creating a system that would allow the collection of metadata related to an experiment and provide additional functionalities for the coming scientist during their experiments.

Lastly it is also essential to be able to navigate and use the metadata collected in a efficient way to improve the experience and enable as much as possible scientific operations.

## SCIZOO

SciZoo is the collection name of SciCat and the associated services running at MAX IV. SciCat is a collaboration project between several research facilities to support a uniform way of storing metadata. SciCat serves as a web-based graphical user interface designed to showcase scans alongside their accompanying metadata. Its primary aim is to provide a straightforward overview of conducted scans, linking them to specific proposals, detailing relevant parameters, and indicating the data storage location.

The central SciCat webpage presents a comprehensive scan overview, with each scan occupying a dedicated page. Users can log in using their MAX IV user credentials to access scans within their purview. For beamline staff, this encompasses all beamline scans, whereas members of proposal groups can access scans linked to their respective proposals. Publicly shared scans are visible to anyone without the need to log in.

Upon selecting a specific scan, users gain access to all of metadata associated with that scan. This includes details such as scan name, description, a unique PID (also available in the URL for easy reference), data type (raw or processed), and creation timestamp. Additionally, information pertaining to the associated proposal, such as the proposal owner and principal investigator's email, is provided. The storage path of the data is also disclosed. Finally, a list of scientific metadata recorded during the scan is presented.

A schematic overview of SciCat and its associated services is depicted in Fig. 1. At MAX IV, all sub-systems combined is often referred as SciZoo. A detailed description of each of the services included in the SciZoo suite and other relevant services at MAX IV follows below.

## DUO

DUO, the Digital User Office at MAX IV, plays a pivotal role in managing user access and data. Before arrival at the facility, each user is assigned a proposal within DUO, and only individuals affiliated with this specific proposal possess the authorization to access the associated data stored on the facility's data storage system. The Tango control system [1] employs a dedicated Tango Device known as "PathFixer" to efficiently generate the requisite file path on the disk with the appropriate permissions for each proposal. Furthermore, PathFixer is responsible for disseminating this path information to the scanning orchestration systems such as Sardana,

\* daphne.van\_dijken@maxiv.lu.se

## IR OF FAIR: PRINCIPLES AT THE INSTRUMENT LEVEL

Gerrit Günther<sup>1\*</sup>, Sebastian Baunack<sup>2</sup>, Luigi Capozza<sup>2</sup>, Oliver Freyermuth<sup>3</sup>, Pau Gonzalez-Caminal<sup>4</sup>, Boxing Gou<sup>5</sup>, Johann Isaak<sup>6</sup>, Sven Karstensen<sup>7</sup>, Axel Lindner<sup>7</sup>, Frank Maas<sup>2</sup>, Oonagh Mannix<sup>1</sup>, Andrew Mistry<sup>8</sup>, Isabella Oceano<sup>7</sup>, Christiane Schneide<sup>7</sup>, Thomas Schörner-Sadenius<sup>7</sup>, Kilian Schwarz<sup>7</sup>, Vivien Serve<sup>1</sup>, Lisa-Marie Stein<sup>7</sup>, Stefan Typel<sup>6</sup>, Malte C. Wilfert<sup>2</sup>.

<sup>1</sup> Helmholtz-Zentrum Berlin für Materialien und Energie GmbH (HZB),

Hahn-Meitner-Platz 1, 14109 Berlin, Germany

<sup>2</sup> Helmholtz-Institut Mainz, Johannes Gutenberg-Universität Mainz,

Staudingerweg 18, D-55099 Mainz, Germany

<sup>3</sup> Rheinische Friedrich-Wilhelms-Universität Bonn,

Nußallee 12, 53115 Bonn, Germany

<sup>4</sup> Fusion for Energy (ATG Science & Technology, S.L.),

c/ Josep Pla 2, Blg. B3, 08019 Barcelona, Spain

<sup>5</sup> Institute of Modern Physics, Chinese Academy of Sciences,

Lanzhou 730000, China

<sup>6</sup> Technische Universität Darmstadt, Institut für Kernphysik,

Schlossgartenstraße 9, 64289 Darmstadt, Germany

<sup>7</sup> Deutsches Elektronen-Synchrotron DESY,

Notkestraße 85, 22607 Hamburg, Germany

<sup>8</sup> GSI Helmholtzzentrum für Schwerionenforschung GmbH,

Planckstraße 1, 64291 Darmstadt, Germany

### Abstract

Awareness of the need for FAIR data management has increased in recent years but examples of how to achieve this are often missing. Focusing on the large-scale instrument A4 at the MAMI accelerator, we transfer findings of the EMIL project at the BESSY synchrotron to improve raw data, i.e. the primary output stored on long-term basis, according to the FAIR principles. Here, the instrument control software plays a key role as the central authority to start measurements and orchestrate connected (meta)data-taking processes. In regular discussions we incorporate the experiences of a wider community and engage to optimize instrument output through various measures from conversion to machine-readable formats over metadata enrichment to additional files creating scientific context. The improvements were already applied to currently built next generation instruments and could serve as a general guideline for publishing data sets.

### INTRODUCTION

In recent years the concept of FAIR data, addressing their Findability, Accessibility, Interoperability, and Reusability [1], evolved from a theoretical framework to realization where processes and infrastructure enable a data life cycle from creation over long-term storage to its re-use. Large-scale research instrumentation represents a starting point of this life cycle as they generate data from sophisticated measurement processes. As a consequence, the quality and

precision of these raw data are crucial for the quality of future data that may originate from this basis. Here, raw data refers to the first version of experimental output that is stored on long-term basis and serves for the instrument scientists as original data source for subsequent analysis.

In previous work we have found that responsibility for the findability and accessibility aspects of FAIR lies with repositories and higher level services [2, 3], while the interoperability and re-usability according to the FAIR principles is inherent in the raw data produced by the instrument. To implement a FAIR data life cycle, the instrument's raw data must be improved as a starting point for subsequent automatic processing. Unfortunately, there is only a vague understanding about the FAIRness of raw data and concrete examples are rare. However, there are frameworks to assess the FAIRness of data, such as the FAIR Data Maturity Model, which can serve as a guideline [4, 5].

In this paper, we report on our findings and measures to improve the FAIRness of raw data produced by two particle physics instruments at the beginning and the end of the instrument life cycle. The A4 instrument at the MAMI is already dismantled and is particularly interesting since the output has been designed before the FAIR principles were published. As the A4 data still generate results of scientific interest [6], we create a post-processing workflow to convert and enrich the existing raw data sets improving their FAIRness. As a second use case, the Any-Light-Particle-Search experiment ALPS II at DESY [7] represents an ideal test bed to implement workflows applied at the instrument level since it is currently in the commissioning phase. We concentrate

\* gerrit.guenther@helmholtz-berlin.de.

# RADIATION-TOLERANT MULTI-APPLICATION WIRELESS IoT PLATFORM FOR HARSH ENVIRONMENTS

S. Danzeca\*<sup>1</sup>, J.C. Luna Duran<sup>1</sup>, A. Masi<sup>1</sup>, R. Sierra<sup>1</sup>, A. Zimmaro<sup>1,2</sup>

<sup>1</sup>CERN, Geneva, Switzerland

<sup>2</sup>IES, Université de Montpellier, CNRS, Montpellier, France

## Abstract

The rise of the Internet of Things (IoT) has been a driving force behind the digital transformation of industries and innovation in consumer electronics, enabling real-time data collection, automation, and the creation of interconnected devices tailored to personalized user experiences. However, within the unique environment of particle accelerators, the adoption of wireless approaches is still in its early stages. Given this context, the primary objective of this study is to detail the efforts made in the development and implementation of a flexible IoT-based platform designed for remote monitoring and control within particle accelerators. Founded on the principles established by a previous radiation monitoring device at CERN, this platform stands out as a versatile hardware solution, having been pre-validated for radiation tolerance and energy efficiency. While it integrates seamlessly with CERN's existing LoRa Wireless network, it can also be configured for dedicated gateway setups or even operate in a stand-alone mode. The platform is designed to interface with a variety of sensors and conditioning circuits, representing a significant advancement in the evolution of monitoring and control systems.

## INTRODUCTION

A particle accelerator is surrounded by a myriad of sensors and actuators, all functioning in perfect harmony. In this context, the concept of the Internet of Things (IoT) isn't new. At CERN, the devices that control and monitor these machines are usually called the "control system" [1]. Historically, this system has been deeply associated with cabled infrastructure, predominantly composed of "copper cables" and optical fibers [1]. The wireless IoT concept can be applied to a particle accelerator context but introduces new challenges, some of which were discussed and asserted in [2]. This showed the challenges of qualifying low-power components typical of IoT design, while in [3] the LPWAN solution that best suits the LHC environment was discussed and identified. This focused on the advantages and features that a stand-alone IoT wireless application could bring to the LHC.

As the operational dynamics of a particle accelerator evolve, so do its requirements. The instrumentation, therefore, must exhibit a similar adaptability. Given the expansive nature of these accelerators and their distributed instrumentation, each new requirement typically mandates the deployment of new systems, complete with their own power lines

and communication cables. The methodology followed to develop such systems has ensured their high reliability and operational efficiency for years. However, when it comes to addressing simpler operational needs, conducting specific investigations, or undertaking cross-validations, this traditional approach becomes impractical. The labor intensive nature of such development and the cost implications of the associated cabling and integration often render such projects prohibitive.

The wireless IoT platform discussed in this study offers a new solution. It provides teams with a tool to address challenges that were once too difficult or expensive to tackle. This includes installation of monitoring devices for specific purposes in just a few hours, new installations in temporary experiments, and installations in remote locations or locations with high radioactivity. In this paper we talk about the hardware platform and its performance and characteristics. We also look at CERN's wireless network, focusing mainly on the LoRaWAN network and its data publishing and collection architecture, and present how LoRaWAN devices work in the CERN accelerators along with deployment strategies. The last part is dedicated to some examples of how this technology can be used.

## CERN IOT HARDWARE PLATFORM

Developing a radiation-tolerant wireless hardware platform suitable in the accelerator context presents significant challenges. First, a deep understanding of radiation effects on electronics is necessary to properly select components and qualify them [4]. Achieving a balance between radiation resistance, power efficiency, and wireless capability requires a meticulous approach and rigorous validation. The design flow follows two paths: one purely electrical-performance driven and the other focused on radiation tolerance, driven by the radiation environment. The radiation qualification path has been discussed in other works [2] and will not be the focus of this paper.

### Hardware Choices

To fulfill the requirements, several key criteria must be addressed in the design choices. First, to avoid the use of cables and exploit wireless powering, a design leading to extended battery life is essential. The choice of components is thus driven by these requirements and only components with low power features are selected. As depicted in Fig. 1, the platform is designed to be modular and consists of three units, the power supply board, the main board and the application sensor board. To meet the typical characteristics

\* salvatore.danzeca@cern.ch

# HIGH FIDELITY PULSE SHAPING FOR THE NATIONAL IGNITION FACILITY

A. S. Gowda<sup>†</sup>, A. Barnes, B. Buckley, A. Calonico-Soto, E. Carr, J. Chou, P. Devore, V. Gopalan,  
JM. Di Nicola, J. Heebner, V. Hernandez, R. Muir, A. Pao, L. Pelz, L. Wang, A. Wargo  
Lawrence Livermore National Lab, Livermore, U.S.A

## Abstract

The National Ignition Facility (NIF) is the world's most energetic laser capable of delivering 2.05 MJ of energy with peak powers up to 500 terawatts on targets a few mms in diameter. This enables extreme conditions in temperature and pressure allowing a wide variety of exploratory experiments from triggering fusion ignition to emulating temperatures and pressures at the center of stars or giant planets. This capability enabled the groundbreaking results of December 5<sup>th</sup>, 2022 when scientific breakeven in fusion was demonstrated with a target gain of 1.5. A key aspect of supporting various experiments at NIF is the ability to custom shape the pulses of the 48 quads independently with high fidelity as needed by the experimentalists. For more than 20 years, the Master Oscillator Room's (MOR) pulse shaping system has served NIF well. However, a pulse shaping system that would provide better shot-to-shot stability, power balance and accuracy across the 192 beams as well as mitigate obsolescence issues is required for future NIF experiments including ignition. The pulse shapes requests vary drastically at NIF which lead to challenging requirements for the hardware, timing and closed loop shaping systems. In the past two years, a High-Fidelity Pulse Shaping System was designed, and a proof-of-concept prototype was shown to meet all requirements. This talk will discuss design challenges, solutions and how modernization of the pulse shaping hardware helped simple control algorithms meet the stringent requirements set by the experimentalists.

## INTRODUCTION

The capability of imploding small capsules to study the physics and interaction of materials at high pressures, matter and radiation temperatures and densities is a valuable tool across fields, from understanding the mechanisms driving stars and the interiors of giant planets to creating a self-sustaining thermonuclear fusion reaction. While multiple technologies exist to create focused energy and power, with the demonstration of the laser in 1960, the option of focusing the energy of laser beams into a small volume was pursued by Lawrence Livermore National Labs to produce mini-fusion explosions. Decades of research and development on a practical laser system capable of delivering the power and energy required for such implosions culminated in the National Ignition Facility (NIF), the world's most energetic laser that can achieve a high-density, symmetric implosion enabling the study of extreme conditions for

Stockpile Stewardship [1]. NIF focusses the energy of 192 beams into a cylindrical capsule a few mms in size to create matter temperatures  $> 10^8 K$ , radiation temperatures  $> 3.5 \times 10^6 K$ , densities  $> 10^2 g/cm^3$  and pressures  $> 10^{11} atm$  [2]. The 192 beams combined deliver about 2.05 MJ energy [2]. With this capability, on December 5<sup>th</sup>, 2022, a self-sustaining fusion reaction that generated 3.15 MJ from 2.05 MJ of laser energy, a target gain of 1.5, was demonstrated for the first time [2, 3]. Recently, this achievement was repeated on July 30<sup>th</sup>, 2023 with even higher energy out 3.88 MJ with 2.04 MJ laser energy, a target gain of 1.9.

The facility is a powerful tool for experiments in fields spanning high energy density to discovery science. An important feature of the facility that supports the vast breadth of experiments is flexibility of shaping 48 independent shapes with custom delays. Figure 1 illustrates the breadth of pulses that users request at Target Chamber Center (TCC), from high contrast ratio (defined as the ratio of the power at the peak divided by the lowest power in the foot) pulses for ignition experiments to long, slow but precise power ramps for high energy density science experiments.

This paper describes the legacy pulse shaping system used at NIF and a recent upgrade that was deployed to improve performance of the shaping system as well as address obsolescence issues. The paper is broken up into three main sections, the first section describes how pulse shaping is accomplished at NIF and the more stringent performance requirements needed for future experiments, the second section describes the upgrade that was designed for higher fidelity and NIF sustainment and the third section outlines the performance of the upgrade. The conclusions are outlined in the final section.

## PULSE SHAPING AT NIF

Figure 2 illustrates the physical and control path of an individual beamline [1]. The pulse starts in the Master Oscillator Room (MOR) as a nominally Flat-in-Time (FIT) pulse near  $\sim 1053$  nm wavelength. The pulse is shaped in the MOR and passed through the rest of the optics which amplifies  $10^{15}$  times and frequency converts twice from  $1\omega \rightarrow 2\omega \rightarrow 3\omega$ . The MOR is the only point in the beamline that has active feedback control of the pulse and can correct for the distortions to the pulse shape from the amplification and frequency conversion process. A typical shot starts with the user inputting the required pulse shape at Target Chamber Center (TCC) into the Laser Performance and Operations Model (LPOM). LPOM with the Virtual Beamline engine back propagates the pulse from

\* Prepared by LLNL under Contract DE-AC52-07NA27344. LLNL-PROC-855426

<sup>†</sup>gowda1@llnl.gov

# NOISE MITIGATION FOR NEUTRON DETECTOR DATA TRANSPORT\*

Kazimierz J. Gofron<sup>†</sup>, Bogdan Vacaliuc, Rob Knudson IV, Steve Hicks  
Spallation Neutron Source and High Flux Isotope Reactor,  
Oak Ridge National Laboratory, Oak Ridge, TN, USA

## Abstract

Detector events at user facilities require real-time fast transport of large data sets. Since its construction, the Spallation Neutron Source (SNS) user facility at the US Department of Energy's Oak Ridge National Laboratory successfully transported data using an in-house solution based on Channel Link low-voltage differential signalling (LVDS) point-to-point data protocol. Data transport solutions developed more recently have higher speed and more robustness; however, the significant hardware infrastructure investment limits migration to them. Compared with newer solutions, the existing SNS LVDS data transport uses only parity error detection and LVDS frame error detection. The used channel link is direct current (DC) coupled and thus is sensitive to electrical environment noise. It is difficult to maintain the same LVDS common reference potential over an extensive system of electronic boards in detector array networks.

## INTRODUCTION

The SNS existing Channel Link uses LVDS [1] for data transport with a clock of about 40 MHz and a mixture of parallel and serial data transport. The 7 data bits per twisted pair in each clock cycle are transported over three pairs of typically Cat5e or Cat7 shielded cable. The maximum data rate is about 840 Mbps per cable. The DS90CR217 [2] and DS90CR218 [3] chipset pairs transport LVDS data, and SN65LVDS32 [4] and SN65LVDS31 [5] chipsets send timing-trigger, clock, and resend to lower-level boards.

Discussed herein are noise mitigation methods to improve data transport within the existing as-built infrastructure. We consider the role of shielding, ground loops, and specifically the use of toric ferrite isolation transformer [6] for radio frequency (RF) common mode noise filtering [7] on power input.

The paper investigates LVDS noise at two levels. At one level we look directly at LVDS eye diagram signals using an oscilloscope. For this work, we used Cat8 cables from FiberStore (New Castle, Delaware) [8] to identify propagation limitations. At another level, we look at the system with data loss caused by system environment noise in an

extended network of detectors within the neutron beam-line.

## EYE DIAGRAM

The SNS data collected by read-out card (ROC) are sent through the concentrator (FEM) and data system packetizer (DSP) over Cat5e cables. Reported herein are studies using higher-specification Cat8 cables from FiberStore [8]. The Cat8 shielded cables support transport frequencies up to 2 GHz.

A Tektronix MDO4104B-6 Mixed Domain Oscilloscope, 6 GHz (1 GHz, 5 GS/s) and a Tektronix TDP3500 Differential Probe 3.5 GHz, shown in Fig. 1 and Fig. 2, measured the LVDS signals for the eye diagram. The 40 MHz clock, was measured using a Tektronix TAP1500 1.5 GHz active probe. The measurement points are head-board pins soldered directly to the printed circuit board as close to the LVDS receiver as physically possible.




Figure 1: The eye diagram measurement setup with loop-back hex AAAA5555 test pattern signals sent from the transmit to receive connection of the ROC board.




Figure 2: A Tektronix oscilloscope, with two differential probes and an active probe, is used for eye diagram measurement.

\* Work supported by US Department of Energy (DOE), Office of Science, Scientific User Facilities Division under Contract No. DE-AC05-00OR22725. This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with DOE. The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/doe-public-access-plan>).

<sup>†</sup> gofronkj@ornl.gov

# HELIUM MASS FLOW SYSTEM INTEGRATED INTO EPICS FOR ONLINE SRF CAVITY $Q_0$ MEASUREMENTS\*

K. Jordan<sup>†</sup>, G. Croke, D. P. Christian, J. P. Jayne, M. G. Tiefenback, and C. Wilson  
Jefferson Lab, Newport News, Virginia, USA  
G. Biallas, Hyperboloid LLC, Yorktown, Virginia, USA

## Abstract

The SBIR funded Helium Mass Flow Monitor System, developed by Jefferson Lab and Hyperboloid LLC, is designed to measure the health of cavities in a Cryomodule in real-time. It addresses the problem of cavities with low  $Q_0$ , which generate excess heat and evaporation from the 2 K super-fluid helium bath used to cool the cavities. The system utilizes a unique meter that is based on a superconducting component. This device enables high-resolution measurements of the power dissipated in the cryomodule while the accelerator is operating. It can also measure individual Cavity  $Q_0$ s when the beam is turned off. The Linux-based control system is an integral part of this device, providing the necessary control and data processing capabilities. The initial implementation of the Helium Mass Flow Monitor System at Jefferson Lab was done using LabView, a couple of current sources & a nanovoltmeter. The device is a superconducting element tightly coupled to a heater installed in the helium return transfer line. The amount of power dissipated in the cryomodule is directly proportionate to the amount of heat required to quench the superconducting element. Once the device was proven to work at 2K the controls transitioned from LabView to a hand wired PCB & Raspberry Pi & finally to a PCB interface to a LabJack T7. This is interfaced to the open-source Experimental Physics and Industrial Control System (EPICS) control system. The EE support group preferred to support a LabJack T7 over the Raspberry Pi. 12 chassis were built and the system is being deployed as the cryogenic U-Tubes become available.

## INTRODUCTION

The 2 K Refrigerator used to supply the cold helium to the CEBAF Superconducting Radio Frequency (SRF) Accelerator needs a STEADY Heat Load. Sudden variations “trip” the refrigerator causing hours long “Beam-off” times. The static heat load is the ~20-W heat leak into each of the 52 Cryomodules with 8 SRF Cavities each. More important is the 200 to 300-Watt dynamic heat load from each cryomodule at full gradient in the cavities. The sources of the variable heat are contamination on the inner surface of the cavities, field emission and others. They are lumped into the term  $Q_0$ , which varies with the cavity’s accelerating gradient. When changing the accelerating gradients or turning on or off RF in a Cryomodule, an automatic control system in EPICS changes power in resistive heaters built into all the cryomodules so the Cryogenic Load remains steady. But that automatic system needs accurate  $Q_0$ s for each cavity at operating gradients to work well. Presently there is a ~500 W discrepancy when all the

Cryomodules are turned off and the resistive heaters are set to the compensating values. This is near the difference that could cause a Central Helium Liquefier (CHL) Trip.

The  $Q_0$ s in the automatic system’s internal table are obtained by a combination of the original  $Q_0$ s determined for each cavity in a new cryomodule by tedious Bomb Calorimetry methods or a less tedious “Heat Monster” method using the position settings of the not-well-calibrated Joule Thompson Valves that regulate the level of the 2 K helium bath in each cryomodule to determine heat dissipation. There was no good way to non-invasively determine the heat load of a CM. For over 30 years, Operators wanted an accurate gas flow meter placed in the helium return pipe of Cryomodules. The 2 to 3 K and 1/30 atm conditions for such a meter are an exceptional challenge. If successful, the flow signal could be calibrated as a Watt Meter signal. Thermodynamics reveals that one gram per second of helium evaporation from the 2 K bath of superfluid helium that cools the cavities is caused by 23 Watts of dissipation. Such a flow meter does not require accelerator access & can be used parasitically. Placed in the separable return pipe called a U-Tube, the meter could run while beam is on, showing total cryomodule dissipation – no tedious procedure. To assess individual cavity  $Q_0$ s at various gradients, the beam has to be off and a simple closure and re-opening of the JT valve would be required.

## HISTORY

A hot wire meter was attempted at JLab/SRF in 2004 [1] getting no good results. Several other failed attempts were made using commercial flow meters including hot wire and Coriolis designs. A paper by Japanese researchers [2] using a tin-plated quartz fiber and a superposed resistive heater coating of gold, utilizing the superconducting to non-superconducting transition to get a strong signal was successful in a laboratory setting.

A JLab attempt that improved the robustness of the instrument head was unsuccessful when placed in the return u-tube from a cryomodule because of signal did not rise above system noise. The sensor in Fig. 1 was installed in the LCLS-II test stand return transfer line. This superconductor was only a centimeter long and the heater was insufficiently coupled to the niobium wire. If it did quench there was insufficient voltage drop to measure.

## The Solution

An SBIR topic to make “devices and methods for accurate in-situ measurement of SRF cavity  $Q_0$ s” was suggested by JLab to DOE and posted. A grant was awarded by the Office of Nuclear Physics, DOE Office of Science, to Hyperboloid LLC in February of 2022. The robust

# DEPLOYMENT AND OPERATION OF THE REMOTELY OPERATED ACCELERATOR MONITOR (ROAM) ROBOT \*

Thomas C. Thayer, Namrata Balakrishnan, Maria Alessandra Montironi, Alessandro Ratti  
SLAC National Accelerator Laboratory, Menlo Park, California, USA †

## Abstract

Monitoring the harsh environment within an operating accelerator is a notoriously challenging problem. High radiation, lack of space, poor network connectivity, or extreme temperatures are just some of the challenges that often make ad-hoc, fixed sensor networks the only viable option. In an attempt to increase the flexibility of deploying different types of sensors on an as-needed basis, we have built upon the existing body of work in the field and developed a robotic platform to be used as a mobile sensor platform. The robot is constructed with the objective of minimizing costs and development time, strongly leveraging the use of Commercial-Off-The-Shelf (COTS) hardware and open-source software (ROS). Although designed to be remotely operated by a user, the robot control system incorporates sensors and algorithms for autonomous obstacle detection and avoidance. We have deployed the robot to a number of missions within the SLAC LCLS accelerator complex with the double objective of collecting data to assist accelerator operations and of gaining experience on how to improve the robustness and reliability of the platform. In this work we describe our deployment scenarios, challenges encountered, solutions implemented and future improvement plans. Keywords: Robotics, Infrastructure Monitoring, Remote Control

## INTRODUCTION

The Linac Coherent Light Source (LCLS) accelerator at SLAC National Accelerator Laboratory is a very extensive machine, spanning multiple miles in length and employing hundreds of interdependent devices, all operating together create hard X-ray free-electron laser pulses. Originally designed to produce pulses at up to 120 Hz, the upgraded LCLS-II accelerator is capable of producing pulses on the order of 1 MHz. In order for the machine to function optimally at such a high repetition rate, the numerous individual devices need to be monitored while the machine is operating. This is ordinarily accomplished using an array of sensors placed in carefully chosen fixed locations around the accelerator housing supervised through remote network connections. In most cases, these sensors are adequate to remotely monitor the accelerator function, however there are many instances where ad-hoc sensor placement is necessary

for diagnostics or troubleshooting. Fulfilling the requirements for as-needed sensor placements have, in the past, been executed with movable one-off sensor carts in lieu of permanent installations, but this still requires accelerator downtime for human access to move the carts into their positions. Thus, there is a need for remotely configurable sensor arrangements that do not require beam off conditions.

First presented in [1], the Remotely Operated Accelerator Monitor (ROAM) robot was created as a sensor platform that can be mobilized to extend the capabilities of the fixed location sensor network. The design of ROAM relies on the use of Commercial Off-the-Shelf (COTS) components and open-source software, which permits streamlined development and allows the platform to be easily configurable for different deployment scenarios with minimal modifications. Because of ROAM's intended environment, prudent attention was paid when developing the remote control software to minimize the chances of collision with accelerator components.

In this work, we discuss the deployments and operational challenges faced by ROAM in the LCLS and LCLS-II accelerators. After providing an overview of some similar robots and applications within accelerator environments in Sec. *Related Work*, a review of the hardware and software features of the ROAM robot is given in Sec. *ROAM Overview*. Then in Sec. *Deployment and Challenges*, we describe the challenges and lessons learned from the deployment of ROAM in the accelerator complex. Finally, we provide closing remarks in Sec. *Conclusions* and consider future plans for the application, development, and deployment of the ROAM platform.

## RELATED WORK

The deployment of robots for the purpose of accelerator monitoring is a relatively new field with a limited body of research associated with it. Mainly, larger institutions such as the European Organization for Nuclear Research (CERN) have more developed robotics programs. As an example, [2] discussed an omnidirectional wheeled robot with to perform environmental monitoring for radiation, temperature, and oxygen concentration within the Super Proton Synchrotron at CERN. This robot works autonomously, and has its radiation sensor mounted on a movable arm for flexible positioning. Its purpose built, and is designed to work with CERN's already developed robotics infrastructure. This infrastructure is known as CERNTAURO [3], and was created as a framework for real-time controls of mobile robots. It is similar to Robot Operating System (ROS), however it was conceived

\* This work was supported in part by the U.S. Department of Energy under contract number DE-AC02-76SF00515. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of their affiliated agencies.

† {thayer, namrata, alexmon, ratti}@slac.stanford.edu

# MEASUREMENT OF MAGNETIC FIELD USING SYSTEM-ON-CHIP SENSORS\*

A. Sukhanov<sup>†</sup>, Brookhaven National Laboratory, Upton, NY, USA

## Abstract

Magnetic sensors have been developed utilizing various physical phenomena such as Electromagnetic Induction, Hall Effect, Tunnel Magnetoresistance (TMR), Giant Magnetoresistance (GMR), Anisotropic Magnetoresistance (AMR) and Giant Magnetoimpedance (GMI). The compatibility of solid-state magnetic sensors with complementary metal-oxide-semiconductor (CMOS) fabrication processes makes it feasible to achieve integration of sensor with sensing and computing circuitry at the same time, resulting in systems on chip. We discuss cost-effective multi-channel applications of AMR, TMR and Hall effect integrated sensors for precise measurements of 3D static magnetic fields in wide range of magnitudes from  $10^{-6}$  T to 0.2 T and for pulsed magnetic fields up to 0.2 T.

## INTRODUCTION

Rapid grow of development and production of system-on-chip magnetic field sensors started in around 2010, driven by demand in consumer electronics and automotive industry. The commercially-available system-on-chip 3-axis magnetic sensors (smart sensors) are very compact multi-chip modules designed for low to medium field magnetic sensing with a digital interface for applications such as low-cost compassing, magnetometry and mechanical measurements. The Smart Sensor consist of a 3-axis magneto-sensitive sensor and an ASIC containing analog signal processing, calibration control and a digital interface. The typical area of application of magnetic field sensors is shown in Fig.1




Figure 1: Typical area of applications of magnetic field sensors [1].

For very small magnetic field SQUID sensors are used, for small magnetic fields flux-gate sensors, for medium values are used MR (magnetoresistive) sensors and for high magnetic field Hall sensors.

\* This work was supported by Brookhaven Science Associates, LLC under contract No.DE-SC0012704 with the U.S. Department of Energy.

<sup>†</sup> sukhanov@bnl.gov

The majority of commercial Smart Sensors are build using AMR effect for magnetic fields lower than 1 mT. The TMR sensors are available for low and high magnetic fields. The Hall-effect Smart Sensors cover magnetic field range up to  $\pm 0.26$  T.




Figure 2: Schematic diagram of a smart sensor.

The schematic diagram of a typical Smart Sensor is shown in Fig. 2.

## MAGNETIC FIELD SENSORS

**AMR sensors [1].** Very simple in preparation, produced in large volumes, mainly for cell phones. They have better sensitivity than GMR and Hall effect sensors. AMR technology is robust against ionizing radiation up to 200 krad [2]. The typical sensor size 3x3 mm. The disadvantages of AMR sensors are as follows:

- relative small change of resistance, not exceeding 2%, magnetic remanence,
- sensitivity to orthogonal component.

**Hall effect sensors [1].** They have many advantages but one quite important drawback – low sensitivity, not exceeding 5mV/mT. The Hall sensors are sensitive to ionizing radiation. The advantages are as follows: simple design and technology of manufacturing,

- possibility to design very small sensors, with dimensions of several microns,
- non-invasive measurements – lack of ferromagnetic elements.

**TMR sensors [1].** The TMR sensor element is using the Tunnel Magneto-Resistance effect, discovered in 1995. They have higher sensitivity than AMR and GMR sensors but suffer from higher noise. The advantages are as follows:

- high sensitivity, ten times higher than that of AMR,
- high tolerance to ionizing radiation (100 krad) [3],
- high sampling rate.
- small size and low power consumption.

### Analog Sensors

Analog sensors cover wider range of magnetic field and have wider bandwidth than the digital sensors. They mainly measure one axis of the field. Output of the analog



## FIVE YEARS OF EPICS 7 – STATUS UPDATE AND ROADMAP\*

R. Lange<sup>†</sup>, ITER Organization, St. Paul lez Durance, France  
A. N. Johnson, S. Veseli, Argonne National Laboratory, Lemont, IL, USA  
K. Shroff, Brookhaven National Laboratory, Upton, NY, USA  
T. Korhonen, S. Rose, European Spallation Source ERIC, Lund, Sweden  
H. Junkes, Fritz-Haber-Institute, Berlin, Germany  
S. M. Hartman, K. Kasemir, Oak Ridge National Laboratory, Oak Ridge, TN, USA  
L. B. Dalesio, M. Davidsaver, G. McIntyre, Osprey DCS LLC, Ocean City, MD, USA  
G. White, SLAC National Accelerator Laboratory, Menlo Park, CA, USA  
M. R. Kraimer<sup>‡</sup>, Self Employment, USA

### Abstract

After its first release in 2017, EPICS version 7 has been introduced into production at several sites.

The central feature of EPICS 7, the support of structured data through the new pvAccess network protocol, has been proven to work in large production systems. EPICS 7 facilitates the implementation of new functionality, including developing AI/ML applications in controls, managing large data volumes, interfacing to middle-layer services, and more. Other features like support for the IPv6 protocol and enhancements to access control have been implemented.

Future work includes integrating a refactored API into the core distribution, adding modern network security features, as well as developing new and enhancing existing services that take advantage of these new capabilities.

The talk will give an overview of the status of deployments, new additions to the EPICS Core, and an overview of its planned future development.

## INTRODUCTION

### EPICS 7

The EPICS (Experimental Physics and Industrial Control System) software toolkit has been publishing regular updates of its major version 3 release series for more than 30 years [1].

In December 2017, the 10-year initial development phase of the next-generation EPICS Base culminated in a large, tedious merge operation. After testing and verification, the first release of the combined existing EPICS V3 code and the new modules (working title “EPICS V4”) was published as “EPICS 7” [2].

### pvAccess

The pvAccess (PVA) network protocol, which was introduced with EPICS 7, extends the features of the classical EPICS protocol Channel Access (CA). PVA provides the ability to transport user-defined arbitrary structures without the need to recompile client or server code. Its distin-

guishing advantage over similar industrial protocol standards is a highly efficient subscription mechanism: Both server and client have an identical copy of the complete structure. When an update is sent, only the changed elements of the structure are transferred over the network and updated on the target side.

### Initial Goals

A few main goals for EPICS 7 were identified and set from the very beginning, as part of the “EPICS V4” development:

- Compatibility of the Input/Output Controller (IOC) software. To allow a non-intrusive and gradual update process for the many existing installations, the EPICS Process Database and all interface layers below shall not change.
- Scope continuity between CA and PVA. Installations shall be able to run pvAccess as a full replacement for the classical Channel Access protocol, across all networks and layers.
- Minimal configuration. Simply migrating an installation from CA to PVA shall not require complex additional configuration.
- User-level libraries in multiple programming languages for client and server. EPICS Base shall provide well-documented user-level libraries to allow the creation of PVA server and client applications in C/C++, Java and Python.

## EVOLUTION AND ACHIEVEMENTS

### Normative Types

Two of the major factors contributing to the success of EPICS have been standard record types in the Process Database and the standard compiled-in structures of the CA network protocol. Based on these two concepts, generic EPICS applications work across many installations.

To keep that advantage, PVA includes well-defined structures, called Normative Types (NT), which are a superset of the CA standard structures.

Normative Types enable the efficient creation of generic applications and ensure the functional compatibility between CA and PVA. While pvAccess allows arbitrary data

\* Work supported in part by the U.S. Department of Energy under contracts DE-AC02-76SF00515 and DE-AC05-00OR22725.

<sup>†</sup> ralph.lange@gmx.de

<sup>‡</sup> Died 31 January 2022

# DEVELOPMENT OF LASER ACCELERATOR CONTROL SYSTEM BASED ON EPICS

Yadong Xia, Qiang Wang, Enshuo Guo, Fangnan Li, Zhen Guo, Qiangyou He, Ke Chen,  
Mengxuan Zang, Jie Zhao, Liwen Feng, Chen Lin<sup>1\*</sup>, Xueqing Yan<sup>1</sup>  
State Key Laboratory of Nuclear Physics and Technology, and Key Laboratory of  
HEDP of the Ministry of Education, CAPT, Peking University, Beijing, China  
Beijing Laser Acceleration Innovation Center, Beijing, China  
<sup>1</sup>also at Institute of Guangdong Laser Plasma Technology, Guangzhou, China

## Abstract

A proton radiotherapy device based on a petawatt (PW) laser accelerator is under constructing in Peking University, supported by the China's Ministry of Science and Technology. The control system's functionality and performance are vital for the accelerator's reliability, stability, and efficiency. The PW laser accelerator control system adopts a three-layer distributed architecture, including device control, front-end (input/output) control and central control (data management, and human-machine interface) layers. The software platform primarily uses EPICS, supplemented by PLC, Python, and Java, while the hardware platform comprises industrial control computers, servers, and private cloud configurations. The control system incorporates various subsystems that manage the laser, target field, beamline, safety interlocks, environment, synchronization, and functionalities related to data storage, display, and more. This paper presents a control system implementation suitable for laser accelerators, providing valuable insights for future laser accelerator control system development.

## INTRODUCTION

The laser accelerator is a novel mechanism that harnesses the intense interaction between ultra-powerful lasers and specific targets, driving a profusion of electrons within the target, thereby generating a potent longitudinal electric field to accelerate ions. This acceleration gradient can reach levels of  $10^3$  to  $10^6$  times that of traditional accelerators. This method has the capability to accelerate charged particles to velocities nearing the speed of light within the microscopic temporal scales of femtoseconds and spatial dimensions of micrometers. Such groundbreaking particle acceleration showcases vast potential. Compared to traditional accelerators, laser proton accelerators boast significant prospective advantages in aspects like spatial equipment requirements, installation intricacies, operational and maintenance costs, radiation protection challenges, and overall system complexity. Among the myriad applications of laser accelerators, one of the most captivating prospects is employing laser-accelerated protons for tumor radiation therapy. Peking University is in the midst of constructing a proton beam therapy system based on the petawatt (PW) laser accelerator (CLAPA-II) [1-4], which will deliver protons at an energy level of hundreds of MeVs.

\* lc0812@pku.edu.cn




Figure 1: Schematic Diagram of the Overall Layout of the CLAPA-II.

This project aims to cater to the needs of tumor treatments by pioneering a slew of key technological advancements in laser systems, target material preparation, and proton therapy systems. The primary objective is to develop an advanced apparatus for proton radiotherapy and promote its wider use worldwide.

The control system for the CLAPA-II adopts a standard distributed framework based on EPICS [5]. Specific research objectives encompass an operator interface to facilitate control, monitoring, and protection of various equipment distributed across subsystems like the laser, the optical transmission, the target field, the beamline, and the treatment room (as depicted in Fig. 1). Adhering to the designed physical objectives, the system will produce and transmit the beam, ensuring the beam parameters meet the high repetition-rate proton irradiation requirements at the terminal. Furthermore, the control system, even under extensive integration and automation, should exhibit high reliability, swift real-time responsiveness, user-friendly human-machine interfaces, holistic safety interlock protection, a database-centric information management system, comprehensive electromagnetic radiation protection, and a network-based communication system. After progressing through stages like prototype preparation, experimental prototyping, engineering prototype, and product prototype, the final system should meet medical software standards, testing requirements, and performance criteria for proton therapy equipment in terms of reliability, stability, maintainability, and interactivity. Building on this, the project aims to finalize the laser platform and its applications, transforming it into an open-access user platform.

# THE TANGO CONTROLS COLLABORATION STATUS IN 2023

T. Juerges, SKA Observatory, Jodrell Bank, United Kingdom

R. Bourtembourg, A. Götz, D. Lacoste, N. Leclercq, ESRF, Grenoble, France

G. Cuni, C. Pascual-Izarra, S. Rubio-Manrique, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain

B. Bertrand, V. Hardion, A.F. Joubert, MAXIV Sweden, Lund, Sweden

Y. Matveev, DESY, Hamburg, Germany

L. Pivetta, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy

T. Noga, M. Nabywaniec, L. Zytiniak, S2Innovation, Kraków, Poland

G. Abeillé, SOLEIL, Gif-sur-Yvette, France

T. Braun, Byte Physics e. K., Annaburg, Germany

R. Auger-Williams, T. Ives, Observatory Sciences Ltd, St Ives, United Kingdom

## Abstract

Since 2021 the Tango Controls collaboration has improved and optimised its efforts in many areas. Not only have Special Interest Group meetings (SIGs) been introduced to speed up the adoption of new technologies or improvements, the kernel has switched to a fixed six-month release cycle for quicker adoption of stable kernel versions by the community. CI/CD provides early feedback on test failures and compatibility issues. Major code refactoring allowed for a much more efficient use of developer resources. Relevant bug fixes, improvements and new features are now adopted at a much higher rate than ever before. The community participation has also noticeably improved. cppTango switched to C++14 and the logging system is undergoing a major refactoring. Among many new features and tools is jupyTango: Jupyter Notebooks on Tango Controls steroids. PyTango is now easy to install via binary wheels, old Python versions are no longer supported, the build-system is switching to CMake, and releases are now made much closer to stable cppTango releases.

## INTRODUCTION

In 1998, the European Synchrotron Radiation Facility (ESRF) submitted a paper about TANGO, a control system framework based on the paradigm of distributed objects, to ICALEPCS 1999 [1]. Today, Tango is much more than a piece of software. First, it is a well organized, very friendly and amazingly fruitful collaboration gathering eleven major institutes [2]. It is also the place of constant development, enhancement, refactoring and innovation [3].

The success of the Tango collaboration relies on its organisation, the members and developers and also on the underlying contract that governs it. The financial contribution by the core members allows the collaboration - represented by its steering committee - to oversee its own budget and finance events and technical subcontracting. The latter definitively boosts the development of Tango. It is now much more than a just framework. It is a rich and very active software ecosystem in constant evolution. The original technology still exists, but the core developers are preparing the future on a daily basis. Tango constantly improves and tries to bene-

fit from the latest technical developments while maintaining a strong backward compatibility for its users.

The present paper provides the reader with the latest news from the Tango Controls collaboration. It notably proposes a wide overview of both the organisational and the technical activity around Tango.

## CPPTANGO

### *Switch to Fixed Release Cycles*

In the past years official cppTango releases were infrequent and usually years apart. These long release cycles resulted in a number of issues:

- Users were not easily convinced to update, because the old release had “somehow worked” for “n years” and “we know the bugs”. Facilities that run in production mode are very conservative when it comes to updating core production software. They do not deploy a new release just because it is the latest version and has more features. This leads to the more serious problem of outdated installations that at some point cannot be updated anymore at all.
- The change set of a cppTango release were usually quite big, because many bug fixes, new features and additions made it into a single release. Users would see this as a mountain too big to climb, because “it has changed everywhere”. An update simply appeared to be too expensive and this contributed again to production systems running outdated cppTango versions.
- The amount of changes made it hard for the users to find the useful features that would make their lives much easier.
- Packaging of cppTango was turned into a time consuming and thus expensive problem, because dependencies had already moved on a long time ago. This re-created dependency-hell for every release and costs the cppTango team every time dearly. Distributions dropped cppTango dependencies and suddenly a substitute had to be found, investigated and adopted before a new release. This caused an incredible amount of additional pressure just before a release.
- Long release cycles bore the risk of Tango Controls

Software

Control Frameworks for Accelerator & Experiment Control

# ASYNCHRONOUS EXECUTION OF TANGO COMMANDS IN THE SKA TELESCOPE CONTROL SYSTEM: AN ALTERNATIVE TO THE TANGO ASYNC DEVICE

B.A. Ojur<sup>\*1</sup>, D. Devereux<sup>2,3</sup>, A. J. Venter<sup>1</sup>, S. N. Twum<sup>3</sup>, S. Vrcic<sup>3</sup>

<sup>1</sup>SARAO, Cape Town, South Africa

<sup>2</sup>CSIRO, Clayton, UK

<sup>3</sup>SKAO, Macclesfield, UK

## Abstract

Equipment controlled by the Square Kilometre Array (SKA) Control System will have a TANGO interface for control and monitoring. Commands on TANGO device servers have a 3000 milliseconds window to complete their execution and return to the client. This timeout places a limitation on some commands used on SKA TANGO devices which take longer than the 3000 milliseconds window to complete; the threshold is more stricter in the SKA Control System (CS) Guidelines. Such a command, identified as a Long Running Command (LRC), needs to be executed asynchronously to circumvent the timeout. TANGO has support for an asynchronous device which allows commands to be executed slower than 3000 milliseconds by using a coroutine to put the task on an event loop. During the exploration of this, a decision was made to implement a custom approach in our base repository which all devices depend on. In this approach, every command annotated as “long running” is handed over to a thread to complete the task and its progress is tracked through attributes. These attributes report the queued commands along with their progress, status and results. The client is provided with a unique identifier which can be used to track the execution of the LRC and take further action based on the outcome of that command. LRCs can be aborted safely using a custom TANGO command. We present the reference design and implementation of the Long Running Commands for the SKA Controls System.

## INTRODUCTION

A long running action, within the SKA Control System (CS) Guidelines [1], is attributed to a command that exceeds the execution time threshold of 10 milliseconds [2]. The Telescope Control System is composed of a number of sub-systems which form an intricate network of communicating components. Due to the hierarchical interaction of these components, coupled with some network and I/O bound actions, response delays are symptomatic within this distributed system. TANGO [3] device servers used to control and monitor the equipment of these components timeout on commands which run tasks longer than 3000 milliseconds. Within the SKA network, commands which execute long running tasks should be executed by clients asynchronously to avoid timeouts [4]. Additionally, device servers should delegate tasks to threads to make them responsive to other

client requests. The TANGO API enables both the implementation of an asynchronous device server and TANGO client. In addition to this solution provided by TANGO, the SKA Control System implements its own asynchronous execution of LRCs with a queue manager and a reporting mechanism to inform clients of the status of their submitted request [5, 6]. This implementation satisfies the four non-functional requirements (NFRs) relevant for handling LRCs, viz.: performance, dependability, interoperability and usability. Table 1 outlines the NFRs. The objective of this research paper is to describe the design and implementation of LRCs in the SKA Control System.

## LRC IN SKA TANGO DEVICES: DESIGN

### *An Overview of the SKA Telescope Control System*

Figure 1 depicts the normal hierarchical nature of TANGO nodes regardless of commands being LRCs.

Node A, found at the top of the hierarchy of nodes, has a process, Process X, to execute. Based on the complexity of the task, Node A divides Process X into 4 subsections namely, X1, X2, X3 and X4. Node A then delegates the four sub-processes to downstream nodes called Node A1, Node A2, Node A3 and Node A4 respectively as can be seen in Fig. 1. The input of all these sub-processes is a command together with its arguments; variables; configuration settings and state of the downstream node it is assigned to from Node A's perspective. This is indicated through the downward directed arrows connecting each downstream node with Node A. The upward directed arrows, from the perspective of Node A, indicate the values returned from the downstream nodes as well as any external and observable state of the downstream nodes. Just as Node A is responsible for dividing Process X into smaller processes it is also responsible for the aggregation of the responses received back from all the downstream nodes running the processes as well as any overhead that may come with that. After responses are aggregated, Node A can then report back to an upstream node denoted as vertical dots above Node A. As mentioned before, this architecture of disseminating and aggregating tasks is command independent and it therefore still applies within the solution of the LRCs spoken about in this paper.

The NFRs described in Table 1 together with key considerations enumerated in the CS guidelines, namely synchrony, asynchrony and concurrency were the building blocks used to design and methodically implement LRCs in

\* bojur@sarao.ac.za

# DIAMOND LIGHT SOURCE ATHENA PLATFORM

J. Shannon, C. Forrester, K. Ralphs, Diamond Light Source, Oxfordshire, UK

## Abstract

The Athena Platform aims to replace, upgrade and modernise the capabilities of Diamond Light Source's acquisition and controls tools, providing an environment for better integration with information management and analysis functionality. It is a service-based experiment orchestration system built on top of NSLS-II's Python based Bluesky/Ophyd data collection framework, providing a managed and extensible software deployment local to the beamline. By using industry standard infrastructure provision, security and interface technologies we hope to provide a sufficiently flexible and adaptable platform, to meet the wide spectrum of science use cases and beamline operation models in a reliable and maintainable way. In addition to a system design overview, we describe here some initial test deployments of core capabilities to a number of Diamond beamlines, as well as some of the technologies developed to support the overall delivery of the platform.

## INTRODUCTION AND MOTIVATION

The current data acquisition platform at Diamond Light Source (Diamond) suffers from several issues including maintenance difficulty, inflexibility and use of obsolete technology. Further motivations for introducing a new platform have been presented previously [1]. The current platform is known as "Generic Data Acquisition" (GDA) which is a large monolithic Java client-server application. It originally began development over twenty years ago.

The high level goals of Athena are to produce a platform that is maintainable, adaptable, testable, and provides beamline scientists with enhanced capabilities. It strives to leverage industry standard technologies and well tested frameworks.

In 2022 Diamond invited a number of external collaborators to review the design. The architecture was presented and the planned transition steps, as well as the research done to justify the decisions. This received a very positive response, with most feedback relating to the management of the programme rather than the technical design.

The COVID-19 pandemic highlighted the increasing need for a system which supports proper remote access for facility users. The current platform supports remote access in a limited capacity. This could be improved using a modern web based front end, significantly improving the user experience.

Cyber security is also at the heart of a new platform. The remote operation use case in particular reinforces the need for properly managed access to the platform and a robust authorisation and authentication scheme.

## Diamond II

Diamond is a currently a 3<sup>rd</sup> generation synchrotron light source with a photon energy of 3.0 GeV. A major upgrade,

named Diamond II, has just been approved to enhance the accelerator to an energy of 3.5 GeV [2]. This will allow for a much higher flux to be delivered to beamline end stations enabling a drastic increase in data acquisition rates. The "dark period" for accelerator upgrade is currently scheduled for late 2027.

Increased flux will have consequences for the computing requirements due to the increased data-rates. This is further compounded by new fast detectors which are being introduced to instruments.

Whilst all instruments will benefit and modernise from the upgrade, a key deliverable is the construction of three new "flagship beamlines". These will be the initial target for the full Athena Platform.

## BLUESKY

Bluesky [3] is set of Python libraries for experiment control and collection of data. It was developed at National Synchrotron Light Source II (NSLS-II), Brookhaven National Laboratory [4] and is currently used in several facilities around the world.

Its main features include data streaming, rich metadata, decoupling between experiment logic and hardware, and customisability.

Devices are represented as Python objects which may implement various protocols, e.g. "Readable", providing an abstraction from the underlying communication with hardware. This abstraction and interaction with a control system is handled by the Ophyd library [5].

In Bluesky the term "plan" is used for an experimental procedure. This is somewhat analogous to "scan" in other experimental orchestration frameworks. Plans are customisable and composable allowing for complex procedures to be completed. The composable nature encourages the reuse of existing plans. One feature of particular interest is the adaptive plan which allows the running plan itself to change behaviour to respond to data or conditions.

The choice of Python as the implementation language is also significant as Python plays a large role in the field of scientific research. Beamline scientists will have the opportunity to contribute directly to the experimental plan logic which is not possible with Diamond's current acquisition platform. Libraries of common functions and plans will be curated with the aim to prevent duplication and allow review and optimisation.

## Fly Scanning Enhancements in Ophyd

One of the requirements of an acquisition and scanning engine is the ability to perform continuous scans (fly scans). In these type of scans, data is collected whilst motors are moving. At Diamond this is implemented by sending a trajectory to a motion controller (PMAC) and using an FPGA

# THE KARABO CONTROL SYSTEM

S. Hauf\*, N. Anakkappalla, J. T. Bin Taufik, V. Bondar, R. Costa, W. Ehsan, S. Esenov, G. Flucke, A. Garcia-Tabares, G. Giovanetti, D. Goeries, D. Hickin, I. Karpics, A. Klimovskaia, A. Parenti, A. Samadli, H. Santos, A. Silenzi, M. A. Smith, F. Sohn, M. Staffehl, C. Youngman, European XFEL, Schenefeld, Germany

## Abstract

The Karabo distributed control system has been developed to address the challenging requirements of the European X-ray Free Electron Laser facility, which include custom-made hardware, and high data rates and volumes. Karabo implements a broker-based SCADA environment. Extensions to the core framework, called devices, provide control of hardware, monitoring, data acquisition and online processing on distributed hardware. Services for data logging and for configuration management exist. The framework exposes Python and C++ APIs, which enable developers to quickly respond to requirements within an efficient development environment. An AI driven device code generator facilitates prototyping. Karabo's GUI features an intuitive, coding-free control panel builder. This allows non-software engineers to create synoptic control views. This contribution introduces the Karabo Control System out of the view of application users and software developers. Emphasis is given to Karabo's asynchronous Python environment. We share experience of running the European XFEL using a clean-sheet developed control system, and discuss the availability of the system as free and open source software.

## INTRODUCTION

Karabo is a supervisory control and data acquisition (SCADA) system, developed to meet control and data acquisition requirements of the European X-ray Free Electron Laser (European XFEL) [1]. Development of Karabo was started in 2010, after surveying other well-known systems such as Tango [2], EPICS [3], and DOOCS [4] as possible control solutions for the planned facility in 2009. At the time of this survey, the anticipated complexity of the European XFEL, and the data volumes generated at the facility were found challenging to address with existing SCADA systems. Consequently, the development of Karabo was started. Karabo is successfully being used to control the photon systems and instrumental end stations of the European XFEL since 2017. In June 2023 it was made available to the public as free and open source software [5]. In this contribution we give a general overview of Karabo's architecture, discuss the operational benefits and drawbacks of this architecture, and conclude with an outlook on how AI-driven agents can facilitate development in the Karabo Ecosystem.

\* steffen.hauf@xfel.eu

## KARABO: CONCEPTS AND ARCHITECTURE

There are two features of the European XFEL that distinguish it from earlier light sources and have direct impact on the requirements for a control and data acquisition system at the facility: the unique time structure of the accelerator that enables MHz bunch repetition rates, and bespoke 2d imaging detectors [6] capable of resolving this time structure, resulting in data rates between 10-20 GB/s. Additionally, the facility initially rapidly changed during construction and commissioning, and nowadays instrumental setups are substantially modified depending on experimental needs, which can change with each user group on a weekly basis. For the control system the above translates to the following boundary conditions:

- the control system needs to scale and can grow alongside the facility,
- has time correlation woven into its fundamental data model,
- process data rates of tens of GByte/s at latencies of a few 10K " s,
- and cater to dynamic experimental setups that change on different time scales,

while being highly reliable and resilient to failure events. Karabo was written from scratch with these requirements in mind.

### *Karabo's Asynchronous and Event Driven Design*

Karabo communicates by asynchronously exchanging messages via a central message broker. For every Karabo installation, a broker name space, referred to as a *Karabo topic* exists, and distributed components are uniquely identifiable within this topic. *Devices* add functionality to the base system in form of pluggable software libraries, and the combination of all online and offline instances of devices in a topic constitutes the control system for this topic.

A completely event-driven publish and subscribe signal-slot messaging pattern is implemented on top of the distributed broker messaging. By subscribing to signals of another instance updates are propagated through the system without a need for polling. Any configuration update will include timing information comprised of timestamp and a unique timing identifier which facilitates correlation on a global facility level. Any number of distributed components can subscribe to a given signal which issued only once, regardless of the number of subscribers, thereby minimizing network traffic. Instance methods are registered as *slots* to make them available throughout the distributed system.

# LOG ANOMALY DETECTION ON EUXFEL NODES

A. Sulc\*, A. Eichler, T. Wilksen, DESY, Hamburg, Germany

## Abstract

This article introduces a method to detect anomalies in the log data generated by control system nodes at the European XFEL accelerator. The primary aim of this proposed method is to provide operators a comprehensive understanding of the availability, status, and problems specific to each node. This information is vital for ensuring the smooth operation. The sequential nature of logs and the absence of a rich text corpus that is specific to our nodes poses significant limitations for traditional and learning-based approaches for anomaly detection. To overcome this limitation, we propose a method that uses word embedding and models individual nodes as a sequence of these vectors that commonly co-occur, using a Hidden Markov Model (HMM). We score individual log entries by computing a probability ratio between the probability of the full log sequence including the new entry and the probability of just the previous log entries, without the new entry. This ratio indicates how probable the sequence becomes when the new entry is added. The proposed approach can detect anomalies by scoring and ranking log entries from EuXFEL nodes where entries that receive high scores are potential anomalies that do not fit the routine of the node. This method provides a warning system to alert operators about these irregular log events that may indicate issues.

## INTRODUCTION

The stability and reliability of the European XFEL facility are essential for a successful operation. To facilitate this, a network of watchdog nodes is continuously monitoring the health state of the facility's essential components. These nodes, numbering in the hundreds, act as monitoring technology, ensuring the proper functionality of crucial European XFEL accelerator elements. Within their logs lie valuable information about the health state that can signal any potential problems with specific components or parts that could impact the entire facility. Automating the costly task of monitoring these lengthy and often redundant logs becomes especially important in guaranteeing the optimal performance of all nodes. The logs contain a wealth of information concerning the system's status, encompassing error messages, anomalies, and other factors that could affect the system or its associated components. By exploiting language embedding and anomaly detection techniques on these logs, we can efficiently identify and address issues or errors at the earliest possible stage when they occur in logs. This proactive approach empowers us to pinpoint potential problems before they escalate, enabling prompt measures to be taken to resolve ongoing issues. Furthermore, it facilitates timely intervention and the implementation of preventive measures

to mitigate potential problems from arising. Monitoring the logs of the watchdog nodes by textual analysis of their logs not only provides an automated means of comprehending the European XFEL accelerator system conditions but also enables early detection and resolution of issues that would otherwise only gain significance in the event of a specific node failure.

The structure of the paper is the following: First, we summarize the related work in log anomaly detection. In the next section, we show four main steps of our approach with important justifications and examples. Lastly, we show several examples and sketch a potential future work in this field.

## RELATED WORK

A common approach to detecting anomalies in logs is to manually define rule-based systems. For example, Cinque et al. [1] and Yen et al. [2] have developed rule-based methods that scan logs for predefined patterns indicative of anomalies. However, these approaches rely heavily on expert knowledge to construct effective rules, which can be labor-intensive. To overcome this limitation, more automated techniques have emerged leveraging machine learning to discover anomalies.

With the increasing popularity of machine learning (ML) models, deep learning-based approaches gave the potential to perform a thorough log analysis under the presence of a large log corpus, often also accompanied by laboriously made labels. Long-term-short-term (LSTM) recurrent neural networks [3–5] turned out to be popular for log-anomaly detection due to its ability to handle sequential data. Recently transformers [6] were deployed in training to detect anomalies in logs [7]. In [7] they used a BERT [8] model for log-anomaly detection. However, their reliance on large training datasets and millions of parameters can limit their applicability in resource-constrained scenarios like ours. For a more comprehensive survey of ML log analysis, see [9].

Bertero et al. [10] propose an approach that treats logs as natural text and leverages word vector Word2Vec representations [11, 12] to perform automated word embedding. This technique maps words to a vector space, enabling the use of off-the-shelf classifiers for anomaly detection. A major drawback is that their approach still relies on manual labeling to train the classifier, which can be prohibitively expensive in our scenario. Additionally, they treat each log entry independently, ignoring the sequential nature of consecutive log message relationships. To mitigate the need for labeled data, other works like [13, 14] have explored unsupervised learning techniques. These methods apply text mining to logs and employ clustering approaches to identify anomalies without relying on manual labels. However, they still

\* antonin.sulc@desy.de

# HIGH AVAILABILITY ALARM SYSTEM DEPLOYED WITH KUBERNETES

J. Bellister, T. Schwander, T. Summers  
SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

To support multiple scientific facilities at SLAC, a modern alarm system designed for availability, integrity, and extensibility is required. The new alarm system deployed at SLAC fulfils these requirements by blending the Phoebus alarm server with existing open-source technologies for deployment, management, and visualization.

To deliver a high-availability deployment, Kubernetes was chosen for orchestration of the system. By deploying all parts of the system as containers with Kubernetes, each component becomes robust to failures, self-healing, and readily recoverable.

Well-supported Kubernetes Operators were selected to manage Kafka and Elasticsearch in accordance with current best practices, using high-level declarative deployment files to shift deployment details into the software itself and facilitate nearly seamless future upgrades. An automated process based on git-sync allows for automated restarts of the alarm server when configuration files change eliminating the need for sysadmin intervention.

To encourage increased accelerator operator engagement, multiple interfaces are provided for interacting with alarms. Grafana dashboards offer a user-friendly way to build displays with minimal code, while a custom Python client allows for direct consumption from the Kafka message queue and access to any information logged by the system.

## INTRODUCTION

To assist in the commissioning of the upgrade to the Linac Coherent Light Source (LCLS-II) at SLAC, an upgrade to the alarm system was proposed. In performing this upgrade there were several main priorities. First is a system which would be easy to work with for end users. Both adding new devices to be monitored, as well as interacting with alarms that are surfaced should be as straightforward as possible. We also wanted a system that had near constant uptime, as it would be monitoring multiple critical facilities across the lab. Finally, it was important to adhere to current best practices around alarm management and system deployment as much as possible. To that end we assessed the current status of various alarm management systems as well as deployment options.

The Phoebus alarm server [1] met most of our requirements so we chose to base our system on it, thus eliminating the need to reinvent the wheel for basic alarm logic and functionality. For the deployment process Kubernetes [2] was settled on for orchestration of the system. Using these technologies brings our system in line with current best practices, while meeting our goal of high availability and

performance. Further details of the deployment process follow.

## ALARM SERVER

Since the Phoebus alarm server meets most of our requirements, only a few site-specific changes were made to it prior to deployment. Two modifications of note are to allow the same process variable (PV) to appear along multiple branches of the alarm tree, and a way to make it easier for an alarm server to reload when its underlying configuration file has changed. This allows for an automated continuous deployment process in which user changes to alarm configurations can be automatically incorporated into running alarm servers.

The main functionalities of handling alarm logic and translating changes in EPICS alarm severities into Kafka [3] messages remain unchanged. As depicted in Fig. 1 below, the alarm server monitors PVs from an EPICS source, in this case a gateway. Any relevant changes are translated to Kafka messages and stored in the proper alarm topic. All clients to the system read from the Kafka queue.

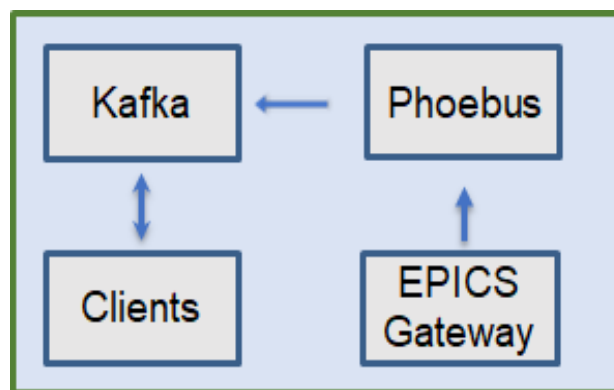


Figure 1: Simplified flow diagram of alarm data.

## ALARM LOGGER

The other main functional component of the system is the Phoebus alarm logger. This records a history of alarms and actions taken on them, and persists this history into an Elasticsearch [4] data store.

Our deployment also provides the ability to use Loki [5] as a data store instead of Elasticsearch. This is handled by using a simple python application that reads data from Kafka and pushes it to Loki. Since many systems at SLAC already use Loki for log storage, providing this flexibility allows users the choice of where to persist data, and more easily integrate it with the data storage of other applications.



# AN UPDATE ON THE CERN JOURNEY FROM BARE METAL TO ORCHESTRATED CONTAINERIZATION FOR CONTROLS

T. Oulevey\*, B. Copy†, J.-B. de Martel‡, F. Locci, S. Page, R. Rocha, C. Roderick, M. Vanden Eynden§  
CERN, Geneva, Switzerland

## Abstract

At CERN, work has been undertaken since 2019 to transition from running Accelerator controls software on bare metal to running in an orchestrated, containerized environment. This will allow engineers to optimise infrastructure cost, to improve disaster recovery and business continuity, and to streamline DevOps practices along with better security.

Container adoption requires developers to apply portable practices including aspects related to persistence integration, network exposure, and secrets management. It also promotes process isolation and supports enhanced observability.

Building on containerization, orchestration platforms (such as Kubernetes) can be used to drive the life cycle of independent services into a larger scale infrastructure.

This paper describes the strategies employed at CERN to make a smooth transition towards an orchestrated containerised environment and discusses the challenges based on the experience gained during an extended proof-of-concept phase.

## BACKGROUND

### *The CERN Technical Network*

CERN operates several distinct networks serving different purposes. Two of them will be referenced during this paper:

1. The General-Purpose Network (GPN), used for a large number of non-accelerator-related operations.
2. The Accelerator Technical Network (TN), which can be seen as a separate protected network, used solely for on-line accelerator operations.

The TN was created in the late 1990s as a private, independent physical network, dedicated to the operation of the CERN accelerator complex. In 2005, the CNIC (Computing and Network Infrastructure for Controls) working group [1] was established with the goal of reviewing, proposing and putting in place relevant network and security measures to better protect the TN. For security reasons, the provision of central IT services on the TN is an ever-growing challenge, in particular with the emergence of cloud-based licensing policies and remote hosted solutions.

\* thomas.oulevey@cern.ch

† brice.copy@cern.ch

‡ jean-baptiste.de.martel@cern.ch

§ marc.vanden.eynden@cern.ch

### *The CERN Accelerator Data Centre*

The CERN Accelerator Data Centre (ADC) is located on the French CERN site in Prévessin-Moëns, next to the CERN Control Centre (CCC) [2]. It hosts around 500 high-performance and high-availability servers, currently running the CERN CentOS7 operating system. In 2023, a project was launched to transition towards RHEL9 in 2024, as the future operating system for accelerator controls.

Since the dawn of the Large Hadron Collider (LHC) era, the ADC operates as a pure bare-metal infrastructure, each server being attributed to equipment groups in charge of the control software of dedicated accelerator sub-systems (*e.g.* cryogenics, radio frequency, beam instrumentation, etc.)

While offering a lot of flexibility to the numerous software development teams, this ADC model is no longer sustainable for reasons such as the lack of optimization and under use of the global computing power, security aspects, a plethora of operational software DevOps practices, wide-spread dependencies towards third-party software solutions, and, last but not least, a lack of agility in terms of evolution.

### *Platform Engineering*

Platform engineering in the context of Kubernetes involves designing, implementing, and maintaining the underlying infrastructure to support containerized applications. This includes configuring and optimizing Kubernetes clusters, managing container orchestration, and ensuring high availability and scalability. It enables development teams to focus on building and deploying applications efficiently within the Kubernetes ecosystem. It also allows system administrators to abstract the underlying infrastructure, making tasks like operating system updates, server hardware replacement, and network changes nearly transparent to the applications and services running on the platform. This abstraction provides a layer of insulation between the infrastructure and the workloads, promoting consistency and ease of management across diverse environments.

This approach and added value may come at the expense of increased complexity and a steeper learning curve for development, DevOps and administration teams, particularly in terms of grasping container orchestration concepts and addressing the associated challenges. System administrators must exercise particular care, to guarantee security and enforce updates, especially during the periodic CERN accelerator Technical Stops.

# DEVELOPING MODERN HIGH-LEVEL CONTROLS APIS

B. Urbaniec, L. Burdzanowski, S. Gennaro CERN, Geneva, Switzerland

## Abstract

The CERN Accelerator Controls are comprised of various high-level services that work together to provide a highly available, robust, and versatile means of controlling the Accelerator Complex. Each service includes an API (Application Programming Interface) which is used both for service-to-service interactions, as well as by end-user applications. These APIs need to support interactions from heterogeneous clients using a variety of programming languages including Java, Python, C++, or direct HTTP/REST calls. This presents several technical challenges, including aspects such as reliability, availability, and scalability. API usability is another important factor with accents on ease of access and minimizing the exposure to Controls domain complexity. At the same time, there is the requirement to efficiently and safely cater for the inevitable need to evolve the APIs over time. This paper describes concrete technical and design solutions addressing these challenges, based on experience gathered over numerous years. To further support this, the paper presents examples of real-life telemetry data focused on latency and throughput, along with the corresponding analysis. The paper also describes on-going and future API development.

## INTRODUCTION

The CERN Accelerator Control System is composed of various high-level services which work together to enable the control of the accelerator complex, experimental areas and across various supporting technical infrastructure. The information provided by these controls services is continuously accessed and modified by different software and processes to ensure optimal operation. Regular evolution of the software and hardware across the complex is needed to meet availability and performance targets. High-level APIs play a crucial role in responding to these demands, enabling developers and experts with greater ease and speed when working with software that needs to configure or interrogate the Accelerator Controls. The API became an integral part of software building blocks, reshaping how software systems are designed, implemented, integrated, and maintained.

This paper presents the journey of continuous evolution of APIs used by CERN Accelerator Controls for the needs of systems configuration and integration, from rudimentary function calls to sophisticated, feature-rich end-user interfaces. High-level APIs have been instrumental in unlocking the potential of cutting-edge technologies by providing intuitive and expressive programmatic interfaces easing the access for the end-users, and equally, when integrating complex systems. The paper describes technologies and concepts used to provide reliable, robust and scalable APIs for the centralised CERN Controls Configuration Service (CCS) [1].

## CONTROLS CONFIGURATION SERVICE

The CCS is a core component of CERN's Control system, serving as a central point for the configuration of all Controls sub-domains, in a coherent and consistent way. The CCS is used by diverse user groups, including installation teams (configuring Controls hardware), equipment experts (configuring processes and applications), and accelerator operators. Though CCS downtime does not directly impact on-going beam operation, CCS users rely on being able to interact with the service at any point in time, to verify or define appropriate configurations. As such, extended downtime periods are unacceptable. To provide the highest possible availability and quality of service, the CCS, including its API are realized as a modular system, with redundancy, monitoring, and alerting at its core. Each day, on average, more than 400 different users and processes generate 80 million requests using the CCS API to obtain or modify various configurations (Figure 1). Peak values of user traffic can reach considerably higher levels.




Figure 1: Number of requests per CCS API node.

## API REQUIREMENTS AND DESIGN

In addition to the core business requirements of the API, represented by the exposed data structures and operations, a robust API requires careful consideration of architectural, technological, and operational factors from the design stage, including:

- Scalability to meet the demands of a growing user base.
- Availability ensuring minimal downtime.
- High performance.
- Fault tolerance.
- Security and stability.
- Accessibility by non-software-developers and across different programming languages and technologies.

The overall CCS architecture is based on the REST (Representational state transfer) architectural pattern - the de-facto industry standard for HTTP-based, high-level communication between different systems. REST based APIs run on an HTTP server, are language agnostic and natively supported by popular programming languages and web browsers. Scalability and high availability are achieved by deploying multiple instances of the API on redundant physical machines to safeguard against downtime caused by hardware or network failures, therefore ensuring continuous operation. This is facilitated by the server instances being stateless and therefore not depending on any shared information, as well as isolating requests from each

# SECURE ROLE-BASED ACCESS CONTROL FOR RHIC COMPLEX\*

J. Morris, A. Sukhanov<sup>†</sup>, Brookhaven National Laboratory, Upton, NY, USA

## Abstract

This paper describes the requirements, design, and preliminary implementation of Role-Based Access Control (RBAC) for the RHIC Complex. This system is designed to protect from accidental, unauthorized access to equipment of the RHIC Complex. It also provides significant protection against malicious attacks. The role assignment is dynamic: device managers always obtains fresh roles for restricted transactions. The authentication is performed on a dedicated role server, which generates an encrypted token, based on user ID, expiration time, and role level. Device managers are equipped with an authorization mechanism which supports either Static or Dynamic assignment of permissions for device parameters. Transactions with the role server take place atomically during secure set() or get() requests. The system has small overhead: ~0.5 ms for token processing and ~1.5 ms for network round trip. A prototype version of the system has been tested at the RHIC complex since 2022. For easy transition, the access to device managers which do not have authorization mechanisms, can be done through dedicated intermediate shield managers.

## INTRODUCTION

The Control System of the RHIC complex [1] provides the operational interface to the RHIC collider and to a long chain of particle accelerators (AGS, Booster, Linac, EBIS, Tandem, CeC, LEReC), including beam injection and extraction lines and beam instrumentation systems. The number of controlled and monitored parameters exceeds 1 million.




Figure 1: RHIC Complex.

Role-Based Access Control (RBAC) is an approach that limits system access to authorized sets of users [2]. Within an organization, roles are created for various job functions. The permission to perform certain operations is assigned to specific roles. Members of staff (or other system users) are assigned particular roles, and through those role assignments acquire the permissions to perform

particular system functions. RBAC is a preventative and therefore inexpensive way to protect accelerator equipment. Other machine protection systems such as interlocks are reactive. Once triggered it can be expensive to recover operations. RBAC can prevent unauthorized users from making incorrect settings which can adversely affect accelerator equipment. RBAC can also be used to ensure machine stability during a run. Once the equipment is fine-tuned and beam is in the machine, an erroneous setting can disrupt operations for hours and valuable data can be lost. RBAC can restrict access to critical settings to a designated set of operators or system experts, reducing the likelihood of an incorrect setting. RBAC role assignments can also be used to determine who is authorized to run control applications.

## RHIC CONTROLS SOFTWARE

The RHIC Control System[3] is closed source software, developed at BNL in the 1990s. The original software was all written in C++ and C. The system architecture has stood the test of time with few changes in communications protocols. Java and Python development suites are now part of the Control System. Figure 2 shows the architecture of device control in the RHIC Control System. The accelerator equipment is controlled by Accelerator Device Object (ADO) software modules. ADOs are hosted by dedicated Front-End Computers(FECs) or by ADO Manager processes that can run on many different hardware platforms. An ADO contains a set of related control parameters (similar to EPICS PVs[4]). The communication protocol with clients is RPC[5]. The transport layer is TCPIP. The ADO handles a limited set of requests: info(), get(), set() and subscribe(). The name service, which allows clients to find ADOs of interest in the network, is provided by the ControlsNameServer (CNS).

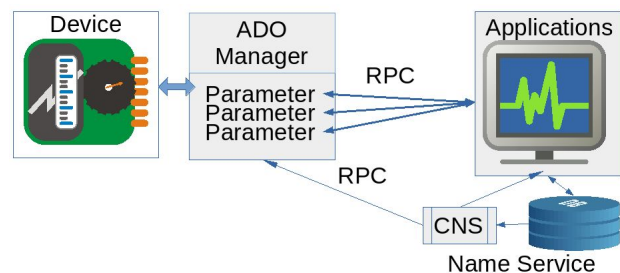


Figure 2: RHIC Controls client-server model.

## DEVICE ACCESS CONTROL AT RHIC

The device access policy in the RHIC complex has been based mainly on network restrictions and access

Software

Software Architecture & Technology Evolution

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

\* This work was supported by Brookhaven Science Associates, LLC under contract No.DE-SC0012704 with the U.S. Department of Energy.  
<sup>†</sup> sukhanov@bnl.gov

# SKA TANGO OPERATOR

M. Di Carlo\*, M. Dolci, INAF Osservatorio Astronomico d’Abruzzo, Teramo, Italy  
P. Harding, U. Yilmaz, SKA Observatory, Macclesfield, UK  
P. Osório, Atlar Innovation, Portugal  
J. B. Morgado, CICGE, Faculdade de Ciências da Universidade do Porto, Portugal

## Abstract

The Square Kilometre Array (SKA) is an international effort to build two radio interferometers in South Africa and Australia, forming one Observatory monitored and controlled from global headquarters (GHQ) based in the United Kingdom at Jodrell Bank. The software for the monitoring and control system is developed based on the *TANGO-controls* framework, which provide a distributed architecture for driving software and hardware using *CORBA* distributed objects that represent devices that communicate with *ZeroMQ* events internally. This system runs in a containerised environment managed by *Kubernetes* (k8s). k8s provides primitive resource types for the abstract management of compute, network and storage, as well as a comprehensive set of APIs for customising all aspects of cluster behaviour. These capabilities are encapsulated in a framework (Operator SDK) which enables the creation of higher order resources types assembled out of the k8s primitives (Pods, Services, PersistentVolumes), so that abstract resources can be managed as first class citizens within k8s. These methods of resource assembly and management have proven useful for reconciling some of the differences between the TANGO world and that of Cloud Native computing, where the use of Custom Resource Definitions (CRD) (i.e., Device Server and DatabaseDS) and a supporting Operator developed in the k8s framework has given rise to better usage of TANGO-controls in k8s. Keywords: kubernetes, controller, tango, operator SDK

## INTRODUCTION

The Square Kilometre Array (SKA) project has selected a software framework for the monitoring and control system called *TANGO-controls* [1], a distributed middleware for driving software and hardware using *CORBA* [2] (Common Object Request Broker Architecture) distributed objects that represent devices that communicate with *ZeroMQ* [3] events internally. The entire system runs on a containerised environment managed by *Kubernetes* (k8s) [4] with *Helm* [5] for packaging and deploying SKA software. In k8s, all deployment elements are resources abstracted away from the underlying infrastructure implementation. For example, a *Service* (network configuration), *PersistentVolume* (file-system type storage) or *Pod* (the smallest deployable unit of computing, consisting of containers). The resources reside in a cluster (a set of connected machines) and share network, storage, computing power and other resources. Namespaces in k8s create a logical separation of resources within a

shared multi-tenant environment. A Namespace enforces a separate network and set of access rights, enabling a virtual private space for contained deployment. Fundamentally, k8s uses a declarative “model” of operation that drives the system towards the desired state described by user manifests, with various controller components managing the lifecycle of the associated resources. Helm provides the concept of a chart which is a recipe to deploy multiple k8s resources (i.e., containers, storage, networking components, etc...) required for an application to run. The resources are created using templates so that the chart can adapt generic configurations to different environments (i.e., the different SKA datacentres). The SKA deployment practices include the heavy use of standardised Makefiles (i.e., for building container images, for testing, for the deployment of a chart, etc.) and *Gitlab* [6] for the CICD (continuous integration continuous deployment) practices [7],

## SETTING UP A TANGO DEVICE SERVER IN KUBERNETES

The TANGO-controls framework is middleware for connecting software processes, mainly based on the *CORBA* standard. The *CORBA* standard defines how to expose the procedures of an object within a software process with the *RPC* protocol (Remote Procedure Call). TANGO extends the definition of an object with the concept of a *Device* that represents a real or virtual device to control. This exposes commands (procedures), and attributes (i.e., state) allowing both synchronous and asynchronous communication with events generated from attributes. The software process is called *Device Server*. Figure 1 shows a module view of the aforementioned framework.

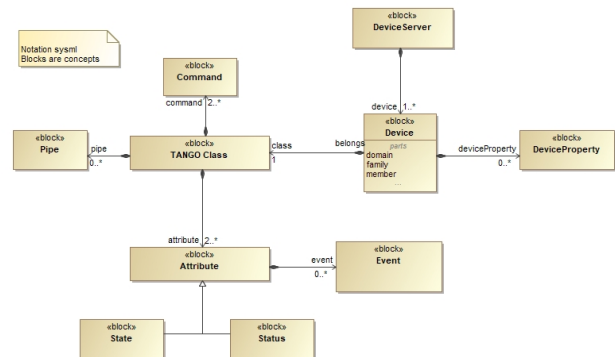


Figure 1: TANGO-controls simplified data model.

Given the TANGO-controls framework, the steps to set up a *Device Server* are the following:

\* matteo.dicarlo@inaf.it

# SYNCHRONIZED NONLINEAR MOTION TRAJECTORIES AT MAX IV BEAMLINES

P. Sjöblom\*, H. Enquist, A. Freitas, J. Lidón-Simón, M. Lindberg, S. Malki  
MAX IV Laboratory, Lund University, 224 84 Lund, Sweden

## Abstract

The motions at beamlines sometimes require components to move along non-trivial and non-linear paths. This type of motion can be achieved by combining several simple axes, typically linear and rotation actuators, and controlling them to perform synchronized motions along individual non-linear paths. A good example is the 10 meter spectrometer at MAX IV Veritas beamline, operating under the Rowland condition. The system consists of 6 linked axes that must maintain the position of detectors while avoiding causing any damage to the mechanical structure. The non-linear motions are constructed as a trajectory through energy or focus space. The trajectory changes whenever any parameter changes or when moving through focus space at fixed energy instead of through energy space. Such changes result in automated generation and uploading of new trajectories. The motion control is based on parametric trajectory functionality provided by IcePAP. Scanning and data acquisition are orchestrated through Tango and Sardana to ensure full motion synchronization and that triggers are issued correctly.

Keywords: trajectory, motion, spectrometer, control system, linked axes

## INTRODUCTION

Motion of devices is a fundamental requirement on beamlines in order to position samples, detectors, optics and other items precisely and in a timely manner [1–5]. The requirements vary vastly from being a simple on/off beam motion, continuous energy scans through linear approximation [6, 7] to even more complex ones where several axes are linked together to perform a non-linear motion where each position is dependent on others positions and where those motions should be conducted correctly in time. At MAX IV a concept of trajectories is used where each motor in a system is given a unique trajectory table to follow. Motion trajectories are used, with clear benefits, in, e.g., the SCANIA spectrometer at Balder beamline, the flight tube at CoSAXS beamline, and the monochromators at FlexPES and FinEst-BeAMS beamlines but here, we focus on the solution for Veritas Rowland spectrometer.

The 250 to 1500 eV soft x-ray beamline Veritas and its 10-meter spectrometer arm (Fig. 1) has a detector wagon that is allowed to move roughly 2.5 meters along a linear path (1.2 meters vertically) suspended on a truss of beams. Inside the truss, the vacuum system is suspended that also needs to move. In the detector wagon, the detector is suspended in three linear motors. In the near future, a second detector will be added orthogonal to the first one and a multilayer

mirror inserted to allow polarization measurements. All the involved motors have to be in a certain position described by the equations derived and described in [8] to ensure that the beam path is correct through the system.

## SPECTROMETER CONCEPT

### Rowland Circle

From a conceptual point of view the detector follows a Rowland circle [9] to maintain the focus of the detector Fig. 2. In order to maintain the Rowland circle criteria, the detector motion should follow the circle arc determined by the curved grating surface. This is achieved by using two linear motions together with a rotation. By combining the motion of two linear actuators, the detector moves along the desired circle arc. As the entire beam path must be under vacuum, a set of vacuum pipes, interconnected by bellows, are installed between the gratings and the detector, whereas the last bellow allows for a 2.5-meter displacement of the detector. This vacuum system is too heavy to passively follow the motion of the detector. Instead, it is driven by three linear actuators. It is essential to avoid mechanical stress on the system, especially the bellows, as these might be damaged. Therefore, any displacement of the detector requires all six axes to move in a synchronized manner, which keeps the vacuum system straight and stress-free, not only at the target position but also during the entire motion. Directions and strokes of the 6-axes motion system are depicted in Fig. 3, where the trajectories should place the red circles and the detectors (D1 and D2) on the red beam path, by performing linear motions marked with blue circles. The beam path differs depending on gratings used and desired energy and focus. One motion will affect other motions causing at least one red dot to move away from the correct path.

### Mechanical Limitation

Each motor is equipped with limit switches to prevent overtravel. This is however not sufficient to prevent damage to the system. One example is the mirror chamber (denoted MC in Fig. 3) in front of the detector chamber. Both of these components need to be able to move vertically, to align them correctly with respect to the beam. These motions allow much larger relative displacements than the bellow connecting the two chambers can accommodate. Thus any larger detector motion must be compensated for by also moving the mirror chamber. Another example is the two motors that move the detector horizontally and vertically. In the middle of the horizontal motion, the detector needs to pass over a part of the support structure. Around this position, the detector must be in the upper half of its vertical

\* peter.sjoblom@maxiv.lu.se

# OPEN SOURCE ETHERCAT MOTION CONTROL ROLLOUT FOR MOTION APPLICATIONS AT SLS-2.0 BEAMLINES

A. S. Acerbo\*, T. Celcer, A. E. Sandström, Paul Scherrer Institut, Villigen, Switzerland

## Abstract

The SLS-2.0 upgrade project comprises of a new storage ring and magnet lattice and will result in improved emittance and brightness by two orders of magnitude. Paired with these upgrades is a generational upgrade of the motion control system, away from VME based hardware and towards a more modern framework. For SLS-2.0 beamlines, the EtherCAT Motion Control (ECMC) open source framework has been chosen as the de-facto beamline motion control system for simple motion, analog/digital input/output and simple data collection. The ECMC framework comprises of a feature rich implementation of the EtherCAT protocol and supports a broad range of Beckhoff hardware, with the ability to add further EtherCAT devices. ECMC provides soft PLC functionality supported by the C++ Mathematical Expression Toolkit Library (ExprTk), which runs at a fixed frequency on the EtherCAT master at a rate up to the EtherCAT frame rate. This PLC approach allows for implementing complex motion, such as forward and backward kinematics of multi-positioner systems, i.e. roll, yaw, and pitch in a 5-axis mirror system. Additional logic can be loaded in the form of plugins written in C/C++. Further work is ongoing to provide flexible Position Compare functionality at a frequency of 1 kHz coupled with event triggering as a way to provide a basic fly-scan functionality for medium performance applications with the use of standardized SLS-2.0 beamline hardware. We provide an overview of these and related ECMC activities currently ongoing for the SLS-2.0 project.

## ETHERCAT MOTION CONTROL ARCHITECTURE

The EtherCAT open standard has firmly established itself as a reliable real-time fieldbus for extensively distributed and synchronized systems [1]. Diamond Light Source (DLS) [2], European Spallation Source (ESS) [3, 4], SPring-8-II [5], Paul Scherrer Institut (PSI), and others have introduced open-source solutions for the bus master in scientific installations, utilizing EtherCAT hardware for digital and analog I/O and motion control. At these institutes, EtherCAT serves mid-performance data acquisition and motion control in accelerator and beamlines applications.

The open source EtherCAT Motion Control (ECMC) [6] framework originally developed at the ESS uses the open source Etherlab master [7] to control the EtherCAT bus. ECMC interfaces with EPICS Motor Record using a model 3 driver and implements basic functionality such as positioning, homing, and limits. The ECMC framework supports advanced features including full servoloop feedback, virtual axes, interlocking, and a PLC-like scripting language

based on the C++ Mathematical Expression Toolkit Library (ExprTk) [8].

## ECMC Configuration and EPICS IOC

Configuration of EPICS IOCs using the ECMC framework is facilitated by the ECMC configuration (ECMCCFG) framework [9]: The configuration framework includes all the essential files required for setting up an EPICS IOC for EtherCAT-based motion control and DAQ. For a given EtherCAT master, individual EtherCAT slaves can be added with the addSlave script and specifying the exact model of EtherCAT module to be added. For the addSlave script, the following parameters are accepted:

- HW\_DESC: Hardware descriptor, i.e. EL1008
- SLAVE\_ID: (optional) bus position
- SUBST\_FILE: (optional) substitution file
- P\_SCRIPT (optional) naming convention prefix script
- NELM: (optional) Used for oversampling cards. Defaults to 1
- DEFAULT\_SUBS (optional) option to disable default PVs for mapped PDOs

In the above function, the HW\_DESC parameter loads EtherCAT specific parameters. Further scripts are available for adding custom (local) configuration parameters to specific slaves, modifying SDO objects, enabling diagnostics, or loading plugins from PLC. Individual axes are typically added with the addAxisYaml script, which accepts a yaml file containing configuration parameters for a single axis. Yaml configuration files are parsed with jinja2; Yaml fields for axes include:

- axis: id, mode (CSV, CSP), power on/off parameters
- var: local variables
- epics: EPICS related fields, such as PV name, EGU, PREC
- drive: motor driver parameters
- encoder: motor encoder parameters
- controller: PID loop tuning parameters
- trajectory: trajectory generation
- input: limits, homing switch, interlocks
- plc: PLC related fields
- homing: homing routine related
- softlimits: enable/disable and limit values
- monitoring: limits on velocity, tracking lag, setpoint deadband

For applications requiring highly customized axes where functionality provided by ECMC PLC is not sufficient, a plugins allowing external functions written in C/C++ to be loaded and called from PLC code. All ECMC variables are available in real time and can be passed to these external functions. For data acquisition and processing, local storage

\* alvin.acerbo@psi.ch

# THE LCLS-II EXPERIMENT CONTROL SYSTEM\*

A. Wallace<sup>†</sup>, D. Flath, M. Ghaly, K. Lauer, J. Yin, Z. Lentz, T. Johnson, R. Tang-Kong  
 SLAC National Accelerator Laboratory, Menlo Park CA, USA

## Abstract

The Linac Coherent Light Source (LCLS) has been undergoing upgrades for several years now primarily through two separate major projects: the LCLS-II (Linac Coherent Light Source II) and the LCLS-II Strategic Initiative (or L2SI project). The LCLS-II is a DOE 403.13b project responsible for upgrading the accelerator, undulators and some front-end beam delivery systems. The LCLS-II Strategic Initiative assumed responsibility for upgrading the experiment endstations to fully utilize the new XFEL machine capabilities planned to be delivered by the LCLS-II beam. Both projects included scope to design, install and commission a control system prepared to handle the risks associated with the tenfold increase in beam power we will eventually achieve. This paper provides an overview of the new control system architecture from the LCLS-II and L2SI projects and status of its commissioning.

## INTRODUCTION

Experiment Control Systems (ECS) is a division<sup>1</sup> within the LCLS directorate responsible for the X-Ray beam delivery and experiment hutch control systems. The team ensures the control system is well-designed, reliable and capable of supporting a wide range of scientific experiments conducted at LCLS. The team works closely with various stakeholders including scientific staff, other engineering disciplines and technicians to gather and understand operational requirements.




Figure 1: The increase in brightness and average power from LCLS-I to LCLS-II, with its soft x-ray undulators (SXU) and hard x-ray undulators (HXU).

\* Work supported by U.S. D.O.E. Contract DE-AC02-76SF00515

<sup>†</sup> awallace@slac.stanford.edu

<sup>1</sup> ECS is also the term used for the Experiment Control System itself.

The scope of ECS responsibilities start at the Electron Beam Dump (EBD) and continues through a Front End Enclosure (FEE), into the various experimental hutches. Currently, the team supports both LCLS-I and LCLS-II areas.

LCLS-I experiment controls is primarily based on the Experimental Physics and Industrial Control System (EPICS), integrating a wide variety of components for mechatronics and process control. Data acquisition and analysis systems are handled outside of the purview of ECS.

With the increased power and pulse energy of the new LCLS-II beam (see Figure 1), the necessity for a robust control system with a strong emphasis on machine protection became crucial. This escalation in operating requirements necessitated a different controls system design from LCLS-I. Additional key goals included increased dependency on automation for regular operations and checkouts, and optimization of beam delivery as well as anticipation of thermal effects from high-power beam. Increased machine protection requirements combined with more complex instrument designs were the primary factors influencing the LCLS-II ECS design.

## SCOPE

The ECS scope spans from integrating various mechanical and electronics components such as actuators and sensors and extends all the way to the graphical user interface (see Figure 2). This integration involves the development and implementation of advanced logic to achieve complex functionality.




Figure 2: Integrated aspects of the LCLS-II experiment system scope.

## ARCHITECTURE

The LCLS Experiment Control System (ECS) team maintains a network-based distributed controls systems built on EPICS (see Figure 3). The ECS architecture is built using a combination of real-time control platforms (in the form of PLCs), EPICS device support, high-bandwidth networking and computing infrastructure, and Python-based user interface frameworks (see Figure 4).

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# SAMBUCA: SENSORS ACQUISITION AND MOTION CONTROL FRAMEWORK AT CERN

A. Masi, O. O. Andreassen, M. Arruat, M. Di Castro, S. Danzeca, M. Donze, S. Fargier, R. Ferraro, M. Gulin, E. Van Der Bij, I. Kozsar, J. Palluel, E. Matheson, P. Peronnard, J. Serrano, E. Soria, J. Tagg, F.Vaga, CERN, Geneva, Switzerland

## Abstract

Motion control applications at CERN are often characterized by challenging requirements, such as high positioning precision (in the order of 1 ppm) in extremely radioactive environments (i.e., Total Ionizing Dose up to 30 MGray) and axes synchronization below one ms. These demanding specifications are particularly relevant for the Beam Intercepting Devices (BIDs), such as the Large Hadron Collider (LHC) collimators. With such radiation levels, reading or driving electronics need to be placed in safe areas up to several hundred meters away from sensors and actuators, while typical industrial solutions can work with cable lengths only up to a few tens of meters. A massive R&D phase lasting several years has been required to develop mechatronics solutions able to ensure the requested positioning precision in such a challenging environment. It includes the development of rad-hard actuators, new position sensors, novel reading algorithms, and actuator control solutions. The acquisition and control systems currently in operation for the LHC Collimators are based on off-the-shelf PXI cards, controllers, and chassis from National Instruments. In the frame of the control system renovation foreseen for Long Shutdown 3 (LS3), the high-performance field Sensors Acquisition and Motion Control synchronized all across the 27-kilometer length of the LHC. This synchronization achieves a positioning precision of just a few  $\mu\text{m}$  on a 30-millimeter movement range.

The same level of precision has been attained when reading the collimator position sensors, such as LVDTs and resolvers. This accomplishment can be attributed to the development of sophisticated signal-processing algorithms. It is worth noting the impressive nature of this performance, considering the extensive cable lengths between sensors/motors and the conditioning or driver electronics, which can span up to 1 kilometer [2].

Another notable development is for Piezo goniometers employed for crystal collimation in the LHC [3]. In this project, a mechatronic solution was developed to attain an angular positioning precision of 1 micro-radian (1  $\mu\text{rad}$ ) on a 20 milli-radian (20 mrad) stroke length. This level of precision is achieved through the utilization of advanced piezo actuator control methods and the implementation of Fabri-Perot based interferometry for angular measurements.

The existing control systems rely on software from the LabVIEW RT PXI platform and hardware from National Instruments, which have been effectively adapted and customized to meet the stringent reliability and availability standards required for CERN's critical missions [4]. However, as part of the comprehensive overhaul planned for

framework (SAMbuCa) has been conceived and is currently under development. It represents a complete in-house, modular, and flexible hardware and software framework able to cope with the highly demanding requirements of motion control applications at CERN, optimizing development and long-term maintenance costs. It standardizes the R&D achievements and the return of operational experience of the last 15 years on BIDs mechatronics. After LS3, more than 1200 axes at CERN will be controlled with SAMbuCA. In this paper, the hardware and software framework architectures are described in detail as well as each building block. The development and deployment plan will be also detailed.

## INTRODUCTION

Over the past 15 years at CERN, remarkable progress has been made in enhancing the precision of positioning and synchronization of motorized axes. The overall reliability and availability of the mechatronics systems have improved, as well as the ability to operate in highly radioactive environments.

One notable example is the Low-Level Control System (LLCS) for the LHC collimators [1]. In this system, more than 500 stepper motor axes are the

LHC collimator LLCS in LS3, there is a move away from the conventional "off-the-shelf" approach. This shift aims to avoid vendor lock-in and potential inaccessibility of firmware.

The proposed Sensors Acquisition and Motion Control framework (SAMbuCA) is designed to fulfil the existing and future requirements for mechatronics in the CERN Accelerator domain by considering feedback gathered through extensive operational experience with the LHC Collimator LLCS. The main objective is to deliver a modular, user-friendly, and flexible control solution for mechatronic devices. It will serve as the standard solution to reduce development time and long-term maintenance costs.

The framework features precise control capabilities of mechatronic solutions with control cables extending up to 1 kilometre, catering to a wide range of actuators including stepping, brushed, and brushless motors, as well as piezo actuators. Additionally, it offers accurate readings from an extensive array of field sensors, such as LVDTs, resolvers, potentiometers, switches, PT100 sensors, and strain gauges.

The next section will describe the SAMbuCa hardware architecture as well as its new building blocks, highlighting the improvements compared with the existing solution.



# THE SNS PLC BASED CONTROL SOLUTION FOR STEPPER MOTORS\*

D. Williams<sup>†</sup>, F. Medio<sup>‡</sup>, Oak Ridge National Laboratory, Oak Ridge, TN, USA

## Abstract

The Spallation Neutron Source (SNS) at Oak Ridge National Laboratory has been operating for over 15 years. Many electronic components are now obsolete and require replacement to assure reliability and sustainability. SNS uses stepper motors to control accelerator components throughout the facility, including the cryomodule tuners, beam scrapers, and the primary and secondary stripper foils. The original motor controls were implemented with VME controllers, custom power supplies, and various types of motor drivers. As these components became less reliable and obsolete a new controls solution was needed that could be applied to multiple motion control systems. Fast performance requirements are not crucial for these stepper motors; therefore, the Programmable Logic Control (PLC) technology was selected. The first system replaced was the RING stripper foil control system and plans are underway to replace the beam scrapers. This paper provides an overview of the commercial off-the-shelf (COTS) hardware used for the new stepper motor control at SNS. Details of the design and the challenges of converting a control system during short maintenance periods without disrupting beam operation will be covered in this paper.

## INTRODUCTION

The original hardware provided by Brookhaven National Laboratory for the stepper motors in the RING primary and secondary stripper foils was a Versa Module Europe (VME) based system with Pacific Scientific motor drivers, Acopian power supplies, OMS VME58 and VME digital I/O boards. The Experimental Physics and Industrial Control System (EPICS) interface had several parameters that were not well understood, and documentation of their functionality was minimal. The Pacific Scientific motor driver configuration parameters were stored in the hardware and retained by the onboard battery. These batteries were having end of life failures and when power was lost, the parameters had to be reloaded manually. The batteries were soldered onto the circuit boards that had to be disassembled for replacement. All the major components were obsolete and spares on-hand were running short from previous failures. With the increased risk of downtime running the current system, a new control solution was required.

## NEW PLC BASED STEPPER MOTOR CONTROLS DESIGN

Goals for the new design included using commercial off-the-shelf (COTS) hardware from established SNS standards. The SNS standard PLC is Allen-Bradley

ControlLogix [1] and the AMCI SD4840E2 stepper motor controller was selected for compatibility with our PLCs. The AMCI SD4840E2 stepper motor controller [2] combines both motor drive and controller in one Ethernet IP module which reduces the cost compared to a separate motor controller in the PLC chassis for each stepper motor. An AMCI ANA2 Ethernet resolver was also used for position readback for the chainsaw signal used to rotate a different foil on the changer. The existing stripper foil field hardware in the RING tunnel and associated field cables would be repurposed for the new design. Connectors were added to these cables to allow rollback to the old system if required during the initial production run of the new PLC system. The Acopian power supplies were replaced with two Sola 24VDC power supplies.

## Architecture of the Motor Controls

A Soft IOC replaced the VME IOC and associated hardware to communicate to the PLC for status and issue commands to the foil controls. New streamlined operator screens were developed for commanding stepper motors for the plunge and chainsaw drives. The EPICS engineering screens that allowed an operator to change the dynamics of the motor driver parameters were eliminated to prevent improper control of the motors performance that could potentially damage the foils. These are now embedded in the AMCI driver controller and are not available for an operator to change the stepper motor's performance.

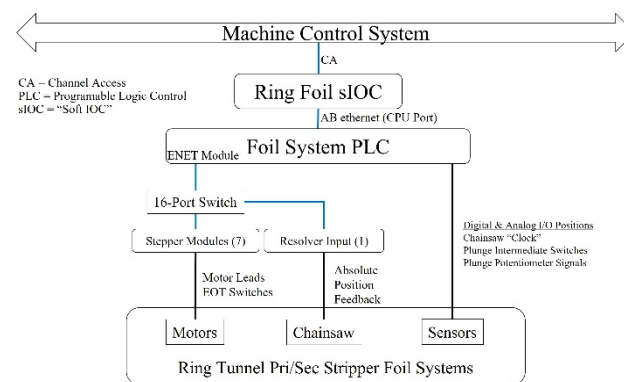


Figure 1: Control system architecture diagram.

The Soft IOC communicates over the machine Ethernet network to the PLC processor's Ethernet port using TCP/IP [3]. A separate PLC Ethernet module is connected to an unmanaged switch that is connected to the stepper motor controllers and the resolver input modules (Fig. 1). These AMCI modules are hardwired to the field devices i.e., step-per motors and resolver. The input signals for the chainsaw limit switches used to indicate which foil is inserted into the beamline are hardwired to the PLC digital input module located in the PLC chassis. A PLC analog input module is used as a potentiometer for the plunge motion of the foil.

\*This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract number DE-AC05-00OR22725.

<sup>†</sup>williamsdc@ornl.gov

<sup>‡</sup>mediofc@ornl.gov

# NEW DEVELOPMENTS ON HDB++, THE HIGH-PERFORMANCE DATA ARCHIVING FOR TANGO CONTROLS

D. Lacoste, L. Banihachemi, R. Bourtembourg, ESRF, Grenoble, France

S. Rubio-Manrique, ALBA-CELLS, Cerdanyola del Vallès, Spain

J. D. Mol, ASTRON, Dwingeloo, Netherlands

L. Pivetta, G. Scalamera, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy

J. Forsberg, MAX IV Laboratory, Lund, Sweden

T. Juerges, SKA Observatory, Jodrell Bank, United Kingdom

## Abstract

The Tango HDB++ project is a high performance event-driven archiving system which stores data with micro-second resolution timestamps. HDB++ supports many different back-ends, including MySQL/MariaDB, TimescaleDB, a time-series PostgreSQL extension, and soon SQLite. Building on its flexible design, with the latest developments supporting new back-ends is even easier. HDB++ keeps improving with new features, such as batch insertion, and by becoming easier to install or setup in a test environment, using ready to use docker images and striving to simplify all the steps of deployment. The HDB++ project is not only a data storage installation, but a full ecosystem to manage data, query it, and get the information needed. In this effort a lot of tools were developed to put a powerful back-end to its proper use and be able to get the best out of the stored data. Moreover, the latest developments in data extraction, from low level libraries to web viewer integration, such as Grafana, will be presented, pointing out strategies in use in terms of data decimation, compression and others to help deliver data as fast as possible.

## INTRODUCTION

Since about ten years, the TANGO HDB++ archiving system is developed as a collaborative project between different institutes using Tango Controls (Alba, Astron, Elettra, ESRF, INAF, MaxIV, SKAO,...).

HDB++ provides tools and components to store Tango attribute values into the database back-end of your choice, with micro-second resolution timestamps. Tools and libraries are also provided for the extraction and visualization of the stored data, for the configuration of the attributes to be stored and for the monitoring of the archiving system health. The following back-ends are currently supported: MySQL/MariaDB, PostgreSQL, TimescaleDB and SQLite. New Database back-ends can be easily added, thanks to HDB++ modular design.

## Design

The EventSubscriber Tango device subscribes to a list of Tango attribute events and stores the attribute values received with these events into a database. The ConfigurationManager Tango device helps managing the list of attributes to be stored in the historical database.

These two Tango devices, written in C++, use the libhdb++ abstraction library to decouple the interface to the database back-end from the implementation. HDB++ provides the libraries implementing the libhdb++ interface for the supported back-ends. Thanks to this design, the same tools can be used to manage and monitor the archiving system whatever back-end is used. Adding a new database back-end is quite easy, because it is just matter of creating the library implementing the libhdb++ interface for the new back-end.

## DATABASE BACK-ENDS

One of the main strengths of HDB++ is the use of proven database engines for storage back-ends. Building an abstraction layer to interact with different back-ends is a key strategy that allows to select the preferred back-end, based on performance, system footprint, preferred technology or simply in-house expertise. The HDB++ community is currently supporting five different back-ends, in production in the various institutes, facing additional requests to support new ones. A small comparison of the different back-ends is presented in Fig. 1. A skeleton project is available to developers to help supporting new back-ends, standardizing the building steps and making the integration simpler.

### *MariaDB/MySQL*

MySQL has been the database engine used for the Tango Database server since the very first releases of Tango due to its versatility and ease of installation. The fully compatible open source MariaDB variant is available in all major operating systems and Linux distributions, becoming the default option for small installations. MySQL/MariaDB also demonstrated performance and scalability on large installations, like ALBA or Elettra synchrotrons, where it is used to support the whole archiving system for accelerators and beamlines. MySQL/MariaDB installations support different schema, either using the legacy schema or the new schema introduced in HDB++. Deployment can use either one single database in a single host, database clusters or multiple databases. Using different API's like ProxySQL or PyTangoArchiving helps configuring and extracting data from the databases in a transparent way.

### *PostgreSQL/TimescaleDB*

TimescaleDB is a PostgreSQL extension for time series. It manages partitioning and offers a lot of features dedicated to

Software

Data Management

# ENHANCING DATA MANAGEMENT WITH SciCat: A COMPREHENSIVE OVERVIEW OF A METADATA CATALOGUE FOR RESEARCH INFRASTRUCTURES

C. Minotti<sup>†</sup>, A. Ashton, S. Bliven, Paul Scherrer Institute, Villigen, Switzerland  
F. Bolmsten, M. Novelli, T. Richter, European Spallation Source, Lund, Sweden  
D. McReynolds, Lawrence Berkeley National Laboratory, Berkeley, CA, USA  
M. Leorato, D. Van Dijken, MAX IV Laboratory, Lund, Sweden  
L. Schemilt, Rosalind Franklin Institute, Didcot, UK

## Abstract

In today's data-driven scientific research landscape, the management of vast amounts of data, together with low-effort search capabilities, easy access and retrieval of acquired data is crucial for enabling successful collaborations, empowering all stakeholders to contribute to data enrichment and fostering scientific advancement. Metadata catalogues, where fields are extracted from datasets and inserted into a searchable database, are pivotal in accessing scientific data in such a landscape. Data FAIRness qualities have become more and more important as an increasing number of publishing entities require transparency in published results and their provenance. Metadata catalogues help facilitate FAIR principles by enabling findability, and accessibility. With careful curation of metadata fields, they can play a vital role in interoperability.

We present SciCat a metadata catalogue designed to meet the needs of the community of scientists carrying out experiments and measurements. SciCat offers a scalable and flexible solution that empowers researchers to effectively manage, share, publish, and discover scientific datasets, thereby fostering collaboration, increasing data visibility and accelerating scientific progress.

## INTRODUCTION

Metadata is defined as the data providing information about one or more aspects of the data; it is used to summarize basic information about data that can make tracking and working with specific data easier. [1] It includes, among many, information about the source of the data, its acquisition process, responsible people and the location on a computer network where the data was created and collected.

Some metadata is auto-extracted from experimental data, while others may include unique quantitative and qualitative information produced or inferred post-experiment. This information is essential for future data utilization and may not be stored elsewhere. Metadata varies in format and standards across research fields and must capture data source dependencies. Additionally, the metadata storage process is adapted to each facility's existing infrastructure.

SciCat [2] has been developed with all these challenges in mind and the aim of being the central metadata storage

<sup>†</sup> carlo.minotti@psi.ch

solution, especially for Photon and Neutron facilities. It started as an in-kind contribution and as an open-source project between the European Spallation Source [3] and the Paul Scherrer Institut [4] within the European Photon and Neutron community. The goal was to develop a versatile metadata catalogue supporting researchers across their entire data lifecycle, as depicted in Fig. 1.

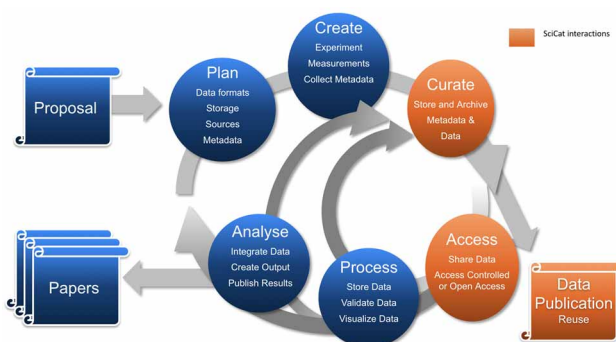


Figure 1: A typical researcher's data journey. The orange shapes are the interactions with SciCat.

Since its release, SciCat has been adopted by other facilities, including MAX IV Laboratory [5], Rosalind Franklin Institute [6], Advanced Light Source [7], the German DAPHNE project [8], Bundesanstalt für Materialforschung und prüfung [9] and the Shanghai Synchrotron Radiation Facility [10]. Additional facilities are rolling it out, for example, Synchrotron SOLEIL [11] and Deutsches Elektronen-Synchrotron [12]. The project has grown over the years with features and the support of dedicated developers and users. It now includes contributions from most of the adopting facilities. The project has recently released the latest version (4.x) which featured a complete rewrite of the backend and includes the corresponding upgrades of the frontend.

We start by introducing SciCat's core components and scalability for handling large volumes of metadata. Then, we cover search functions, metadata record creation (aka metadata ingestion), post-experiment metadata enrichment (aka data curation), and automated data management through *jobs*. The article will also discuss data publication, DOI minting, collaboration, and integration with generic and field-specific web search engines, such as the EOSC data portals [13-15], Google Dataset Search [16] and, for the latter, the PaNOSC data portal [17], developed during the PaNOSC [18] and ExPaNDS [19] EU projects.

# REFLECTIVE SERVERS: SEAMLESS OFFLOADING OF RESOURCE INTENSIVE DATA DELIVERY \*

S. Clark<sup>†</sup>, T. D'Ottavio, M. Harvey, J. P. Jamilkowski, J. Morris, S. Nemesure  
Brookhaven National Laboratory, Upton, U.S.A.

## Abstract

Brookhaven National Laboratory's Collider-Accelerator Department houses over 550 Front-End Computers (FECs) of varying specifications and resource requirements. These FECs provide operations-critical functions to the complex, and uptime is a concern among the most resource constrained units. Asynchronous data delivery is widely used by applications to provide live feedback of current conditions but contributes significantly towards resource exhaustion of FECs. To provide a balance of performance and efficiency, the Reflective system has been developed to support unrestricted use of asynchronous data delivery with even the most resource constrained FECs in the complex. The Reflective system provides components which work in unison to offload responsibilities typically handled by core controls infrastructure to hosts with the resources necessary to handle heavier workloads. The Reflective system aims to be a drop-in component of the controls system, requiring few modifications and remaining completely transparent to users and applications alike.

## BACKGROUND

The Control System of the Collider-Accelerator Department (C-AD) at Brookhaven National Laboratory provides the operational interface to RHIC and all the other accelerators in the C-AD complex. Over 77,000 Accelerator Device Objects (ADOs) provide the software interface to more than a million settings and measurements for accelerator equipment [1]. ADO servers may run on different hardware platforms, but a majority of the ADOs in C-AD are hosted on Front End Computers (FECs) with limited memory and CPU resources. Efforts to upgrade those systems are often impeded by upgrade or redesign costs, labor efforts, and scheduling concerns. As such, effective and efficient use of existing resources is paramount.

The communications protocol used by ADOs has four primary data operations: synchronous gets (blocking while retrieving data), synchronous sets (updating a set point), metadata fetches (retrieving static properties about a read-back or set point), and asynchronous gets (receiving live streams of data updates). Of these, the first three are stateless and have little impact on resource usages beyond the call context. The last, asynchronous gets (asyns), is a stateful operation that require the maintenance of information such as client identifiers, data requests, and data queues, which can consume substantial memory when numerous clients

are connected, or when high bandwidth asynchronous data is requested. Many C-AD applications establish asyns by default, which can quickly consume resources on an FEC causing a crash of the ADO software. This interrupts connections to other users, and in the worst cases interrupts operations. Async load has been attributed to FEC downtime in the past, and has caused certain FECs to be identified as low-resource relative to demand. These units must be treated with caution when interacting with hosted ADOs to prevent resource exhaustion. This can require extraneous communication between users, developers, and the control room to coordinate use, limiting both operational and development efforts.

Previous efforts to develop a data reflection system took place over the last decade [2], but fell out of use due to reliability and maintainability issues within the system. The previous reflective tools suffered from connection issues, inherent to the underlying communication layer implementation, which had the potential to place multiple systems into bad states. Remedying this required process restarts and manual intervention to restore communication between clients and devices. Since then, general purpose calculation engines have been used as stand-ins for a proper reflective system, but are limited by overhead and configuration challenges. These issues of reliability, configurability, and maintainability were identified when outlining plans for an upgraded reflective system, and a new architecture was designed from the ground up to satisfy those needs.

## REQUIREMENTS

The primary requirement for this project is straightforward: develop a system which removes asynchronous load from resource-constrained ADOs. However, the system must do so while also:

- Handling connection interruptions gracefully, allowing for quick and correct recovery from transient communication failures
- Being easy to configure and deploy, with a conspicuous way to examine connected clients and reflected devices
- Fitting seamlessly into the Controls ecosystem with minimal changes necessary to existing clients and without user intervention

## SOFTWARE ARCHITECTURE

The software architecture of this system has been designed to be highly modular, with different logical units being broken into separate processes as outlined in Fig. 1. The first

\* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy

<sup>†</sup> sclark@bnl.gov

# WHATRECORD: A PYTHON-BASED EPICS FILE FORMAT TOOL \*

Kenneth Lauer<sup>†</sup>, SLAC National Accelerator Laboratory, Menlo Park, CA

## Abstract

`whatrecord` is a Python-based parsing tool for interacting with a variety of EPICS (Experimental Physics and Industrial Control System) file formats, including V3 and V7 database files. The project aims for compliance with `epics-base` by using Lark grammars that closely reflect the original Lex/Yacc grammars.

`whatrecord` offers a suite of tools for working with its supported file formats, with convenient Python-facing data-class object representations and easy JSON (JavaScript Object Notation) serialization. A prototype backend web server for hosting IOC (Input/Output Controller) and record information is also included as well as a Vue.js-based frontend, an EPICS build system `Makefile` dependency inspector, a static analyzer-of-sorts for startup scripts, and a host of other things that the author added at whim to this side project.

## BACKGROUND

### *The Problem - and the Inspiration*

Before digging into the details of the `whatrecord` [1], `toolsuite`, let us first take a look at the problem and the inspiration behind its creation.

At the LCLS (SLAC's Linac Coherent Light Source), the accelerator and photon side control systems include approximately 3000 IOC instances in total, with hundreds of modules and dozens of versions per module.

In general, these EPICS [2] IOCs, modules, and extensions are comprised of a conglomeration of unique file formats. Some common examples of such file formats include:

- Process database files (.db)
- Database definition files (.dbd)
- Template / substitutions files
- IOC shell scripts (st.cmd)
- StreamDevice protocols (.proto)
- State notation language programs (.st)
- Gateway configuration (.pvlist)
- Access security files (.acf)
- Build system Makefiles

Additionally, facility-specific tools (centralized IOC management tools like LCLS's IOC Manager, archiver appliance automation tools, and so on) build on top of IOCs and records.

Combined, this makes for an enormous code base with a mix of these EPICS-specific file formats.

Links between these files are often implicit. Take, for example, that an EPICS IOC record has a specific record type name alongside its name in a database file (.db), an EPICS PV (Process Variable) name, in a traditional IOC,

starts with the record name defined in a database file. This PV name acts as a global identifier that allows for clients on the same network subnet to access - and potentially modify - related data.

A record is made up of fields which can contain meta-data like engineering units or user-specified descriptions, references to other records, relevant data values, and so on.

An example record instance, defining a single AI (analog input) record named `IOC:RECORD:NAME` is as follows:

```
record(ai, "IOC:RECORD:NAME") {}
```

This file does not define what the fields of the record type; that is the responsibility of the database definition file (.dbd). A simplified excerpt from a database definition file, defining a single field for the "ai" record type is as follows:

```
recordtype(ai) {  
    ...  
    field(NAME, DBF_STRING) {  
        special(SPC_NOMOD)  
        size(61)  
        prompt("Record Name")  
    }  
    ...  
}
```

Note that there is no explicit link between the database file and the database definition file: neither reference the other by filename. Rather, one can only infer the link by examining a third file, the IOC-specific IOC shell script (.cmd) file, line-by-line.

An excerpt from such a startup script could look like:

```
dbLoadDatabase("path/to/the.dbd",0,0)  
IOC_registerRecordDeviceDriver(pdbbase)  
dbLoadRecords("records.db")
```

Each line of this script includes up to one command. Each of those commands has been registered by either EPICS itself, the modules included in the IOC, or the IOC source code itself. Typically, the available commands would be found either in documentation or by executing the IOC and invoking the built-in help system. Alternatively, the most reliable fallback ends up being the source code itself.

Other direct or indirect references may be found inside fields. For example, depending on the DTYP (device type) field, the INP (input specification) field may be a custom string defined at the device support layer. Interpretation of this field requires knowledge of how these are formatted. Take `StreamDevice` [3], a generic support module for communicating with controllers that use simple byte streams for communication, for example:

\* WORK SUPPORTED BY U.S. D.O.E. CONTRACT DE-AC02-76SF00515.

<sup>†</sup> klauer@slac.stanford.edu

# SECoP INTEGRATION FOR THE OPHYD HARDWARE ABSTRACTION LAYER

P. Wegmann\*, K. Kiefer, W. Smith, L. Rossa, O. Mannix, HZB, Berlin, Germany  
M. Zolliker, PSI, Villigen, Switzerland  
E. Faulhaber, TUM FRMII, Munich, Germany

## Abstract

At the core of the Bluesky experimental control ecosystem the Ophyd hardware abstraction, a consistent high-level interface layer, is extremely powerful for complex device integration. It introduces the device data model to EPICS and eases integration of alien control protocols. This paper focuses on the integration of the Sample Environment Communication Protocol (SECoP) into the Ophyd layer, enabling seamless incorporation of sample environment hardware into beamline experiments at photon and neutron sources. The SECoP integration was designed to have a simple interface and provide plug-and-play functionality while preserving all metadata and structural information about the controlled hardware. Leveraging the self-describing characteristics of SECoP, automatic generation and configuration of Ophyd devices is facilitated upon connecting to a Sample Environment Control (SEC) node. This work builds upon a modified SECoP-client provided by the Frappy framework, intended for programming SEC nodes with a SECoP interface. This paper presents an overview of the architecture and implementation of the SECoP-Ophyd integration and includes examples for better understanding.

## INTRODUCTION

Moving and integrating research equipment between facilities with different Experimental Control Systems (ECS) can be a challenging and time-consuming process. Sample environment hardware, in particular, is usually not permanently attached to a specific experiment and is often moved both within and between research facilities. The Sample Environment Communication Protocol (SECoP) [1] has been developed under direction of the International Society for Sample Environment (ISSE) [2] to facilitate this process. It is also intended as an overarching solution for standardising communication with sample environment equipment at photon and neutron research facilities. The core design principles of the protocol (i.e. simple, inclusive, self-describing and providing rich metadata), were chosen to achieve this goal. In particular, the self-describing properties and the inclusive design philosophy behind SECoP facilitate its integration into various experimental control systems. This is also leveraged in this publication for the integration of SECoP into the hardware abstraction layer Ophyd [3], with the effect that Ophyd devices are constructed automatically upon connection, whilst retaining all metadata about the controlled hardware. A crucial application of this SECoP-Ophyd inte-

gration we introduce here is to facilitate heterogeneous configurations, integrating fixed beamline equipment supported by EPICS with mobile sample environment equipment supported by SECoP in the same Bluesky [4] environment.

## SECoP Hardware Abstraction

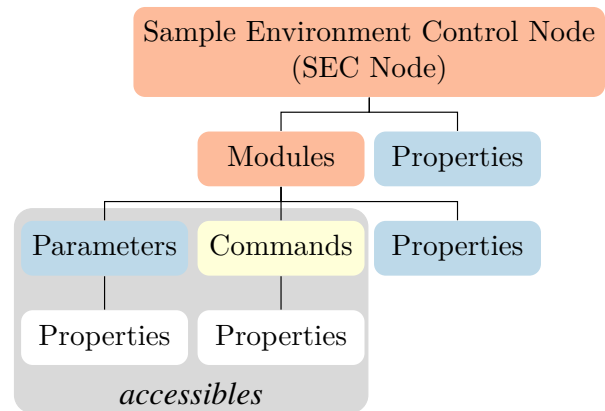


Figure 1: Schematic of SECoP hardware abstraction.

According to the SECoP specification [5], the access point for a SECoP device is a SEC node. The SEC node thereby acts as a server that allows connections from an arbitrary number of clients and gives access to the published functionalities of sample environment equipment. Within the SEC node, physical quantities of a sample environment device are logically represented by *modules*. The term *module* has been deliberately chosen to distinguish it from a device, as a sample environment apparatus is often composed of a number of modules and is only called a device as a whole. All modules are composed of *accessibles*, which can either be *Parameters* or *Commands*. While parameters provide live information on modules, commands are provided to initiate specific actions of a module. This internal structure is depicted in Fig. 1. At every level, static information of the parent entity is provided by property fields. They ensure a structured way of holding metadata regarding the SEC node and hardware it is connected to. For example, at the parameter level, the mandatory *datainfo* property holds important information about the datatype, unit and other metadata. Furthermore, the SECoP standard defines interface classes for modules, with predefined functionality and the meaning of specific mandatory accessibles and properties. Most notable here are the *Readable* and *Drivable* interface classes, enabling a standardised way of constructing readable modules such as a simple temperature sensor and drivables such as

\* peter.wegmann@helmholtz-berlin.de

# FULL STACK PLC TO EPICS INTEGRATION AT ESS

A. Rizzo\*, E. E. Foy, D. Hasselgren, A. Z. Horváth, A. Petrushenko,  
J. A. Quintanilla, S. C. F. Rose, A. Simelio, ESS, Lund, Sweden

## Abstract

The European Spallation Source (ESS) is one of the largest science and technology infrastructure projects being built today. The Control Systems at ESS are essential for the synchronisation and day-to-day running of all the equipment responsible for the production of neutrons for the experimental programs. The standardised PLC platform for ESS to handle slower signal comes from Siemens, while for faster data interchange with deterministic timing and higher processing power, from Beckhoff/EtherCAT. All the Control Systems based on the above technologies are integrated using EPICS framework. We will present how the full stack integration from PLC to EPICS is done at ESS using our standard Configuration Management Ecosystem.

## THE ESS PROJECT

The European Spallation Source (sketched in Fig. 1) is currently under construction in Lund, Sweden. The original ESS configuration was based on a 5 MW LINAC long pulse (2.86 ms) neutron source operating at 14 Hz, serving 22 neutron instruments. However, due to budget constraints, the accelerator power has later been reduced to 2 MW by decreasing the beam energy from 2 GeV to 800 MeV, and the total number of neutron instruments reduced to 15. The target station, on the other hand, is being built for the full 5 MW scope, since the reductions in the accelerator and instrument scope have been made in such a way that it can be restored at a later stage.

The reduction in accelerator power and instrument scope is largely offset by a significant improvement in the moderator design, resulting in a neutron brightness at the level of the original 5 MW design, and meaning the facility is still expected to be world leading shortly after it becomes operational. The first beam on target is expected in 2025, with user operation of the first few instruments planned for 2026 and the full 2 MW LINAC and 15 instruments operational at the end of 2027. For more details please see [1].

## ICS AUTOMATION SECTION

The scope of the Automation Section, which is part of the Hardware and Integration Group (HWI), within the Integrated Control System division (ICS) at ESS, is to coordinate the design, develop, maintain and test of PLC Control Systems, keeping the life cycle documentation updated. It is also the scope of the Automation Section to coordinate the PLC-EPICS integration [2] via the software and web tools within the Control Management, to ultimately create Input/Output Controllers (IOCs) using the ESS customized

EPICS environment called **e3** [3]. Further Automation Section responsibility is to develop and maintain, in close collaboration with stakeholders, operators, and graphic designers of the ICS software group, the relevant Operator Interfaces (OPIs), as well as the archive and alarm configuration. The generic layout architecture of a PLC-based control system is sketched in Fig. 2.

The Automation Section is composed by around 20 Automation Engineers, among ESS employees and consultants, grouped into different Work Packages. Each of them have its specific tasks and responsibilities to do PLC-based control systems integration of different parts of the ESS facility, in particular:

- LINAC
- Target
- Neutron Scattering Systems
- Conventional Facilities

## PLC HARDWARE AND DEVELOPMENT & DEPLOYMENT SOFTWARE

The ICS division has adopted different standardised hardware technologies for implementing the ESS control systems based on performance, taking into account the required signal speed which should be covered while complying to the strict reliability and availability demands at ESS. The standardisation of these technologies makes the implementation cost-efficient and maintenance relatively simple.

For mid-range performance (< 100 kHz), industry standard Beckhoff/EtherCAT [4] systems are used to implement real time fieldbus applications (e.g. motion control systems) with a good price/performance ratio.

For conventional control systems to handle slower signal (< 10 Hz), the standard PLC equipment selected is Siemens [5]; this is a cost-effective solution which addresses ESS reliability and maintainability needs.

### *Siemens TIA Portal*

The Totally Integrated Automation Portal (TIA Portal) is a software and tools package developed by Siemens, which aims to integrate multiple development tools for automation devices. It is used for programming, developing, and configuring Siemens PLCs, HMIs, and frequency inverters.

Currently all the PLC projects are gradually upgraded from version 15.1 to version 17; the possibility to support version 18 with our current control ecosystem is also under investigation.

All the Siemens licenses are stored in a centralized server in an internal DMZ network, which is part of the ESS Technical Network Zone (TN).

\* alfo.rizzo@ess.eu

# DEVELOPMENT OF THE SKA CONTROL SYSTEM, PROGRESS AND CHALLENGES

S. Vrcic, T. Juerges, SKA Observatory, Macclesfield, UK

## Abstract

The SKA Project is a science mega-project whose mission is to build an astronomical observatory that comprises two large radio-telescopes: the SKA LOW Telescope, located in the Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory in Western Australia, with the observing range 50 – 350 MHz, and the SKA MID Telescope, located in the Karoo Region, South Africa, with the observing range 350 MHz – 15 GHz. The SKA Global Headquarters is in the Jodrell Bank Observatory, near Manchester, UK. When completed, the SKA Telescopes will surpass existing radio-astronomical facilities not only in scientific criteria such as sensitivity, angular resolution, and survey speed, but also when it comes to the number of receptors and the range of the observing and processing modes. The Observatory, and each of the Telescopes, will be delivered in stages, thus supporting incremental development of the collecting area, signal, and data processing capacity, and the observing and processing modes. Unlike the scientific capability, which in some cases may be delivered in the late releases, the control system is required from the very beginning to support integration and verification. Development of the control system to support the first delivery of the Telescopes (Array Assembly 0.5) is well under way. This paper describes the SKA approach to the development of the Telescope Control System, and discusses opportunities and challenges resulting from the distributed development and staged approach to the Telescope construction.

## INTRODUCTION

The SKA Observatory [1] is an international organisation whose mandate is to build and operate two multi-purpose radio telescope arrays. The SKA Low Frequency Telescope array (in further text SKA LOW), located in the Murchison region, Western Australia, with the observing range 50 - 350 MHz, will consist of more than 131,072 log-periodic antennas organised as 512 stations; the maximum distance between two stations is 65 kilometres. The SKA Mid Frequency Telescope array (in further text SKA MID), located in the Karoo region, South Africa, with the observing range 350 MHz - 15 GHz, will comprise 197 offset-Gregorian dishes; the dishes are 15 metres in diameter, the maximum distance between two dishes is 150 kilometres. In both Telescope arrays, the receptors (stations in the SKA LOW and dishes in the SKA MID) are arranged in a dense core (with the diameter of ~1 km), and three spiral arms.

Work on the construction, which officially started in 2021, was preceded by the pre-construction activities, which ended with the successful completion of the Critical Design Review (CDR). The CDR was a requirement for

the transition to the construction phase. As a result, when the Project entered the construction phase, a comprehensive set of documentation was available, including the requirements and architecture documents, not only for the Observatory and for both Telescopes, but also for each major sub-system, including the Telescope Control System [2, 3].

This paper provides an overview of the Telescopes and Telescope Control System, describes how the development team is organised, outlines the progress made so far, and discusses some of the challenges the development team encountered during the construction and the strategy to address those challenges.

## TELESCOPE ARCHITECTURE

Figure 1 shows the major subsystems, flow of the observed data and flow of control. Figure 1 is an overview of an SKA Telescope; the receptors are different (Low Frequency Aperture Array (LFAA) and MID array of dishes) while the rest of the subsystems are very similar in the SKA MID and KA LOW.

The key Telescope subsystems that capture and process astronomical data are the receptors, Central Signal Processor and Science Data Processor. The key Telescope support systems are Synchronisation and Timing (SAT), Control System (CS), Networks and Computing Platform where software is deployed. The sub-system known as Observatory Science Operations (OSO) provides a set of tools for proposal submission and management, and for observation preparation, planning, scheduling, and execution.

Equipment and software for each SKA Telescope is deployed as follows:

- The core of the array is located at a radio-quiet site in a desert (Karoo Region, SA and Murchison region, AU)
- Receptors in the spiral arms are spread across a large area; some of the MID array dishes are located more than hundred kilometres from the core.
- The Central Processing Facility (CPF) is located near the core of the array.
- The Science Processing Centre (SPC, where Central Signal Processor and Science Data Processor are deployed, is located in a major urban centre (Cape Town and Perth).
- The Engineering Operations Centre (EOC) is located in a town relatively close to the Telescope sites (Carnarvon, SA and Geraldton, AU).



# CONAN FOR BUILDING C++ TANGO DEVICES AT SOLEIL

P. Madela<sup>†</sup>, Y. M. Abiven, G. Abeillé, X. Elattaoui, J. Pham, F. Potier  
 Synchrotron SOLEIL, Gif-sur-Yvette, France

## Abstract

At SOLEIL, our Tango devices are mainly developed in C++, with around 450 projects for building libraries and device servers for our accelerators and beamlines. We have a software factory that has enabled us to achieve continuous integration of our developments using Maven, which manages project dependencies. However, Maven is uncommon for C++. In addition, it has limitations that hinder us from supporting future platforms and new programming standards, leading us to replace it with Conan. Conan is a dependency and package manager for C and C++ that works on all platforms and integrates with various build systems. Its features are designed to enable modern continuous integration workflows with C++ and are an ideal alternative to Maven for our C++ build system. This transition is essential for the upgrade of SOLEIL (SOLEIL II), as we continue to develop new devices and update existing systems. We are confident that Conan will improve our development process and benefit our users. This paper will provide an overview of the integration process and describe the progress of deploying the new build system. We will share our insights and lessons learned throughout the transition process.

## CONTEXT

At SOLEIL [1], our software development process relies on a well-established framework that encompasses various tools and practices, as illustrated in Fig. 1. This framework is crucial for ensuring the reliable operation of our accelerators and beamlines. It is the result of many years of work aimed at automating the delivery process, as documented in previous papers [2, 3]. The key components of our existing setup are as follows.

### Deployment Process

Our deployment process involves a combination of Continuous Integration and Continuous Delivery (CI/CD) practices. Continuous integration is employed to build software artifacts automatically, ensuring that code changes are regularly integrated and tested. Additionally, we have a semi-automatic continuous delivery system in place to create packages that can be easily deployed. However, it's worth noting that the deployment of these packages currently requires manual intervention during each technical shutdown period.

### Software Factories

We maintain two separate software factories to accommodate the diverse needs of our projects. The first one is dedicated to C++ development and is responsible for

building approximately 400 Tango [4] device servers along with their respective libraries. Table 1 shows the number of C++ projects by platform. The second factory, specializing in Java development, oversees approximately 30 Tango device servers, along with libraries and graphical user interfaces.

Table 1: Number of C++ Projects

Type	Linux	Windows	Both
Application	1	0	0
Libraries	29	6	8
Devices	330	53	23

### Tools of Software Factories

To support these software factories and facilitate our development workflows, we employ a set of tools that includes Gitlab Soleil [5] for version control, Jenkins [6] for automation, and Maven [7] for dependency management.

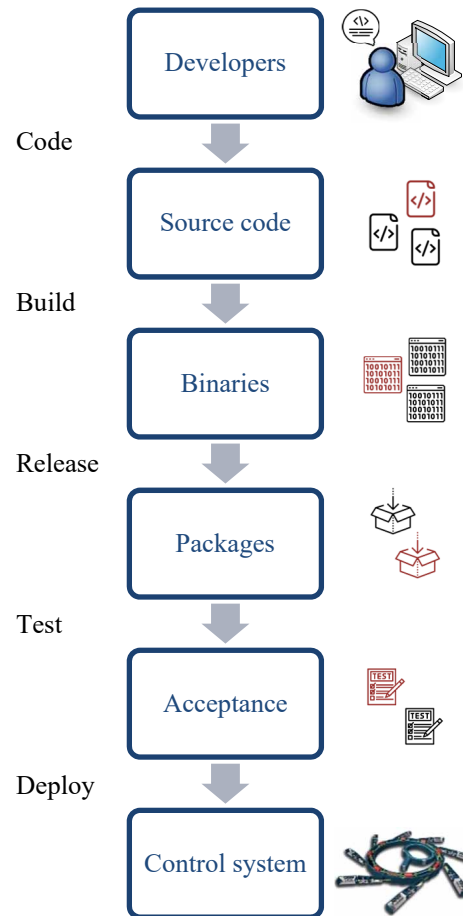


Figure 1: Software development process.

<sup>†</sup> patrick.madela@synchrotron-soleil.fr

# FAIR DATA OF PHYSICAL AND DIGITAL BEAMLINES

Gerrit Günther\*, Simone Vadilonga, Oonagh Mannix, Ovsyannikov Ruslan  
Helmholtz-Zentrum Berlin für Materialien und Energie GmbH (HZB),  
Lise-Meitner-Campus, Hahn-Meitner-Platz 1, 14109 Berlin, Germany

## Abstract

Simulations play a crucial role in instrument design, as a digital precursor of a real-world object they contain a comprehensive description of the setup. Unfortunately, this digital representation is often neglected once the real instrument is fully commissioned. To preserve the symbiosis of simulated and real-world instrument beyond commissioning we connect the two worlds through the instrument control software. The instrument control simultaneously starts measurements and simulations, receives feedback from both, and directs (meta)data to a NeXus file – a standard format in photon and neutron science. The instrument section of the produced NeXus file is enriched with detailed simulation parameters where the current state of the instrument is reflected by including real motor positions such as incorporating the actual aperture of a slit system. As a result, the enriched instrument description increases the reusability of experimental data in sense of the FAIR principles. The data is ready to be exploited by machine learning techniques, such as for predictive maintenance applications as it is possible to perform simulations of a measurement directly from the NeXus file. The realization at the Aquarius beamline at BESSY II in connection with the Ray-UI simulation software and RayPyNG API<sup>1</sup> serves as a prototype for a more general application.

## INTRODUCTION

Synchrotron beamlines are complex instruments which often comprise customized parts to enable scientific investigations that were not feasible before. The life cycle of such unique instruments starts with a simulation in which components and their arrangement to each other are optimized to achieve maximum performance. During the commissioning phase the physical beamline is compared with its digital precursor to make sure that the newly build instrument is constructed according to the simulation and meets the expected performance. Unfortunately, once the beamline is in operation, usually the simulation is neglected and the physical beamline evolves independently of its digital counterpart.

As the central authority over the physical beamline, the instrument control system orchestrates data taking and storage and, thus, is a suitable element to keep a connection to the digital beamline beyond commissioning. By combining the information of physical and digital beamline, the FAIRness in sense of the FAIR principles (findability, accessibility, interoperability, reusability) [1] of experimental and simulation data is mutually improved since they complement

each other. As found before [2, 3], the interoperability and reusability aspects of FAIR are inherent in the data and, in this case, concern experimental and simulation data as well as their relation. In this work, we report on our efforts to store beamline data in a meaningful way in that sense that the relation between experimental and simulation data is visible and distinguishable to human and machine agents. The approach is part of a wider framework which explores the next-generation experiment control and (meta)data software at HZB [3–8].

Throughout the paper, the convention of [9] is adopted to distinguish between data and metadata. Here, data refer to the primary output of physical and simulated detectors or other objects of outstanding scientific interest while metadata belong to information that helps to analyze the primary data such as the description of beamline components. If both types of data are addressed the term (meta)data is used.

## PHYSICAL BEAMLINE

The Aquarius beamline at BESSY II is currently in commissioning and, thus, represents an ideal testbed to explore (meta)data workflows between the physical instrument and its digital counterpart. Aquarius employs a next-generation experiment control system based on BlueSky which is currently developed and tested at different beamlines [5]. When starting a measurement, the experiment control system collects (meta)data of various devices and stores them in a locally accessible Mongo database which provides advanced machine-readability and search options (Fig. 1). For long-term storage and (meta)data publication, the content of the Mongo database is converted to NeXus files which are a standard data format in photon and neutron science [10]. NeXus is physically a HDF5 file format [11] whose content is arranged according to the NeXus Definition Language (NXDL), a semantic framework of predefined structure and naming convention. NeXus files are particularly suitable for storing experimental (meta)data with a dedicated instrument section to define instrument details and, thus, clearly state where (meta)data originate from.

## DIGITAL REPRESENTATION

Beamline simulations are sophisticated programs to compute beam properties along the path of photons through the instrument. They rely on appropriate methods catching the main physics and a digital representation of the beamline using characteristic numbers to describe the instrument. For example, there could be a certain method to deselect parts of the beam when passing a slit and numbers which detail the width, position and material properties of the aperture, deter-

\* gerrit.guenther@helmholtz-berlin.de

<sup>1</sup> raypyng 0.0.20, <https://pypi.org/project/raypyng/>

## ADVANCEMENTS IN BEAMLINE DIGITAL TWIN AT BESSYII\*

S. Vadilonga<sup>1†</sup>, G. Günther<sup>1</sup>, S. Kazarski<sup>1</sup>, R. Ovsyannikov<sup>1</sup>, S. Sachse<sup>1</sup>, W. Smith<sup>1</sup>  
<sup>1</sup>Helmholtz Zentrum Berlin, Berlin, Germany

### Abstract

This presentation reports on the status of beamline digital twins at BESSY II. To provide a comprehensive beamline simulation experience we have leveraged BESSY II's x-ray tracing program, RAY-UI, widely used for beamline design and commissioning and best adapted to the requirements of our soft X-ray source BESSY II. We created a Python API, RayPyNG, capable to convert our library of beamline configuration files produced by RAY-UI into Python objects. This allows to embed beamline simulation into Bluesky, our experimental controls software ecosystem. All optical elements are mapped directly into the Bluesky device abstraction (Ophyd). Thus beamline operators can run simulations and operate real systems by a common interface, allowing to directly compare theory predictions with real-time results. We will discuss the relevance of this digital twin for process tuning in terms of enhanced beamline performance and streamlined operations. We will shortly discuss alternatives to RAY-UI like other software packages and ML/AI surrogate models.

### INTRODUCTION

In the synchrotron community, simulations have long been potent tools, primarily utilized during feasibility studies, design, and the commissioning and operation of machines and beamlines. However, the untapped potential of simulations during subsequent commissioning and operational phases remains substantial. Currently, numerous facilities are focusing on digital twins, aiming to integrate these virtual replicas across various domains, including machines, beamlines, and experiments [1, 2]. This paper serves as a conduit to share our journey and outline our upcoming trajectory in the context of beamline digital twins at the BESSY II facility.

### RAYPYNG: A MODERN APPROACH TO BEAMLINE SIMULATIONS

At BESSY II, we utilize RAY-UI for simulating the beamlines, which is an X-ray tracer with an integrated graphical user interface [3]. The RAY-UI software is specifically designed for precise simulations of X-ray beamlines, enabling users to model the behavior of X-rays as they interact with optical elements like mirrors, lenses, monochromators, and other components. This simulation capability plays a crucial role in designing and optimizing beamlines, predicting experiment performance, and facilitating the commissioning and operation of the beamlines.

While a graphical user interface is helpful when setting up a beamline for the first time, more advanced simulations necessitate integration into a programming language. We opted for Python due to its extensive user community and user-friendly nature, leading to the development of RayPyNG. RayPyNG reads the beamline configuration files generated by RAY-UI and returns a Python object for each beamline component, along with methods for manipulation. As a subsequent step, we developed a simulation backend API for RAY-UI, offering a straightforward way to interact with RAY-UI through Python, read beamline configuration files, initiate simulations, and export results.

Other simulation engines are certainly possible, and we plan to explore alternatives in the future, such as RAY-X (the successor of RAY-UI, currently under development, see [4]), or machine learning surrogate models. Finally, RayPyNG provides a suite of tools and a straightforward syntax for configuring complex simulations. For more information, please refer to the online documentation. [5].

### A BEAMLINE DIGITAL TWIN

At BESSY II, an increasing number of beamlines are embracing Bluesky, a Python library that manages experiment control and scientific data collection, enabling instrument control regardless of software and hardware configurations [6–8]. A key component in this ecosystem is Ophyd, a hardware abstraction layer designed to encapsulate the control layer within a high-level interface. Ophyd enhances efficiency by consolidating individual signals into coherent logical devices.

We have developed a Python library called RayPyNG-bluesky [9], which seamlessly integrates the Python objects representing beamline components created by the RayPyNG library into Ophyd. The prerequisites are minimal: install RAY-UI on the same computer where Bluesky is installed and create a beamline file using RAY-UI. With just two lines of code, you can set up simulations within Bluesky:

```
b1 = "my_beamline.rml"  
RaypyngOphydDevices(RE=RE, rml_path=b1)
```

Ophyd devices will share the same names as defined in the XML file (see Figure 1), with the prefix *rp\_*. For example, if a mirror is named *m1* in the XML file, it will be accessible in Bluesky as *rp\_m1*. Behind the scenes, a trigger detector is created and added to the detector list for all plans using simulated devices. This detector manages the initiation of simulations and the distribution of simulation results to the various detectors specified in the Bluesky plan. This approach offers the advantage of preserving the user experience for Bluesky users. For instance, when scanning the translation along the x-axis of a mirror (M1) and checking

\* Work funded by BMBF and Land Berlin

† simone.vadilonga@helmholtz-berlin.de

# STANDARD DEPLOYMENT OF MICROTCA AT THE EUROPEAN SPALLATION SOURCE

F. Chicken, J.J. Jamroz, J.P.S. Martins  
 European Spallation Source ERIC, Lund, Sweden

## Abstract

This paper outlines the deployment strategies for more than 300 Micro Telecommunications Computing Architecture (MTCA) systems at the European Spallation Source (ESS) ERIC. These units are integral to the control systems spanning from the ion source to the instrument halls, encompassing Radio Frequency (RF), Beam Instrumentation (BI), Machine Protection (MP), and Timing Distribution (TD). As integration paths for these systems have matured, so have the deployment methods. This paper highlights the standardized deployment approaches for MicroTCA Carrier Hubs (MCHs) and Concurrent Technologies computers (CPUs), aimed at enhancing system maintainability, flexibility, and integration efficiency.

## BACKGROUND

MTCA has been selected as the field technology to implement the distributed electronics along the ESS facility. The systems are being built with:

- 9U MicroTCA.4 system for physics with RTM,
- 3U MicroTCA.4 system for physics with RTM.

The MTCA-based ESS systems that were published in detail:

- Front-end electronics in [1],
- Generic data acquisition in [2],
- TD in [3],
- BI-nBLM in [4].

An initial successful outcome was confirmed during the beam commissioning in 2022 which was presented in [5].

## Scope

The total number of MTCA systems to be deployed, once the facility is fully operational, is more than 300. Table 1 presents the ESS overview with the stakeholder breakdown.

Overview examples of the MTCA systems being built are presented in Figs. 1 and 2.

## Integration

During the initial building of MTCA [6] systems, the setup was defined by the integrator, and typically the firmware of

Table 1: The Distribution of MTCA Systems within ESS Machine

ESS $\approx$ 300			
RF	BI	TD	MP
175 x 9U	70 x 3U 15 x 9U	35 x 9U	10 x 3U




Figure 1: 9U MTCA system for TD.




Figure 2: 3U MTCA system for BI.

the MCH ran the firmware version installed by the manufacturer and the Experimental Physics and Industrial Control System [7] (EPICS) environment installed might only have a few selected libraries, and kernel drivers. This process was highly time-consuming, created many variations between crates, and was not well designed for easy updating or maintenance upon deployment. It also proved problematic for debugging issues that arose, given the array of Advanced Mezzanine Card (AMC), FPGA Mezzanine Card (FMC), and Rear-Transmission Module (RTM) being used across the different applications and stakeholders. The MCH with its connections is presented in Fig. 3.

## MCH DEPLOYMENT

The MCH model used at ESS is the NAT-MCH-PHYS board [8]. The MCH has an initial communication using an RS232 connection to a Moxa Technologies Serial-to-Ethernet device (MOXA), to allow a serial connection between the unconfigured MCH and the ESS network (presented in Fig. 4). The purpose of this is to read sufficient information from the MCH (board serial number (SN), me-

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# TOWARDS DEFINING A SYNCHRONIZATION STANDARD BETWEEN BEAMLINE COMPONENTS AND SYNCHROTRON ACCELERATORS\*

X. Serra-Gallifa<sup>†</sup>, J. Avila, ALBA Synchrotron Light Source, Barcelona, Spain  
R. Hino, ESRF, Grenoble, France  
O. H. Seeck, DESY, Hamburg  
T. Cobb, DLS, Oxfordshire, United Kingdom  
Y. Abiven, S. Zhang, SOLEIL, Gif-sur-Yvette, France

## Abstract

Standardization is a magic word in the electronics engineering jargon. Under its umbrella, it is generated the utopia of transparent integration with the rest of the parts with minimal extra effort for the software integration. But the experimental setup in a synchrotron beamline presents multiple challenges: it is highly dynamic and diverse. In the frame of LEAPS-INNOV project (\*), the Task 3 of Work Package 5 aims to define a standard for synchronization in the beamline sample environment. Their partners (ALBA, DESY, DLS, ESRF and SOLEIL) have already reached a common vision of synchronization requirements. This paper first details the participants' actual synchronization needs on their facilities. Next, the requirements foreseen for the future are outlined in terms of interfaces, time constraints and compatibility with timing systems. To conclude, we summarize the current state of the project: the hardware interfaces and the hardware platform definition. They both have been decided considering long-term availability, use of standard sub-components, and keeping the compromise between cost, development time, maintenance, reliability, flexibility and performance. This hardware architecture proposal meets the identified requirements. In the future, under the scope of LEAPS-INNOV, a demonstrator will be built, and we will work with the industry for its future commercialization.

## THE LEAPS-INNOV PROJECT

LEAPS-INNOV is a Horizon2020 EU-funded project that was granted to the League of European Accelerator-Based Photon Sources (LEAPS) consortium members. It is a pilot project focused on implementing technological developments and improving the research capabilities of the European light sources (synchrotrons and free-electron lasers) by reinforcing their partnership with industry. LEAPS-INNOV project will contribute to solving the technological challenges of the nearly 20 European light sources facilities, especially those related with the complexity increase due to the newest generation of diffraction-limited storage rings synchrotrons and X-ray FELs. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101004728.

\* Work supported by LEAPS-INNOV project has received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement No 101004728

<sup>†</sup> xserra@cells.es

The LEAPS-INNOV project is composed the following technology work packages (WP): WP1 - Project Management and Dissemination, WP2 - Development of High Throughput X-ray Spectroscopy Detector System, WP3 - SuperFlat, WP4 - NeXtgrating, WP5 - New positioning and scanning systems for speed and accuracy, WP6 - LIDs, WP7 - Data Reduction and Compression, WP8 - Industrial Innovation through Light Source, WP9 - Innovation by Co-creation to-wards Global Challenges.

## TASK 5.3: SYNCHRONIZATION BETWEEN BEAMLINE COMPONENTS

As part of the WP5, the task 3 is a self-contained project with a defined objective, deliverables and milestones.

### Objectives

The initial and main goal of task 3 of the WP5 is defining a standard protocol for synchronization in the beamline sample environment. Nonetheless, as a result of shared experiences and needs shared along the group meetings, it was detected that a standard hardware for synchronization was also needed. As a consequence, the working group decided to go beyond and to work also on a standard synchronization equipment in which the standard synchronization protocol would be nested. Both a hardware prototype and the standard protocol will be implemented in a demonstrator experiment in a light source facility at the end of the project.

### Deliverables and Milestones

For the tracking of the project the board of the work package defined some deliverables and milestones to achieve.

**Deliverable D5.1: Collated Existing Solutions and Future Requirements of Facilities for Beamline Synchronization** The first deliverable presented on May 2022 was the sum of the contributions of periodic meetings, where the representatives of participating facilities presented the equipment used for the synchronization on the machine and in the sample environment. Also, participants exposed their needs and requirements for a future synchronization standardization. Most of this work is summarized on the next section of this article.

**Milestone M5.4: Definition of the Demonstrator Experiment for Beamline Synchronization** The definition of the demonstrator as a final target of the work

# DEVELOPMENT OF A NEW TIMING SYSTEM FOR ISIS

R. A. Washington\*, ISIS, Rutherford Appleton Laboratory, United Kingdom

## Abstract

The timing system at the ISIS Neutron and Muon source has been operating in its current iteration since 2008. A Central Timing Distributor (CTD) delivers the various timing signals to equipment throughout the ISIS accelerator and these timing pulses are transmitted over RS-422 compliant timing buses. This paper will look at how the next generation of timing system for ISIS could be developed.

A new timing system for ISIS should allow for the distribution of events, triggers and timestamps over fibre optic cable, provide an increase in timing resolution and be fully backwards compatible with the current timing frame. This paper will consider a timing system based on White Rabbit technology and will investigate the results of installing this new system in parallel with the current timing system. A comparison between the systems and next steps will be discussed.

## TIMING OVERVIEW

The fundamentals of the ISIS timing system have not changed much since the first neutrons were produced in 1984 [1]. The ISIS timing system begins with an oven controlled crystal oscillator (OCXO) that provides a stable 10 MHz clock source. This 10 MHz clock is then divided down in two separate stages to produce both a 1 MHz reference for the instrument choppers, and the main 50 Hz operating frequency for the Main Magnet AC power circuit (Fig. 1).

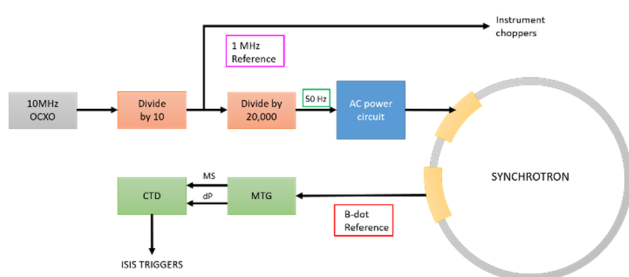


Figure 1: Overview of ISIS timing.

The magnetic field produced by the main magnet power supply (MMPS) [2] is used to define the primary ISIS timing signals, Machine Start (MS) and Delta-P (DP).

## Master Timing Generator

A B-dot sensor, located on one of the synchrotron dipole magnets, converts the magnetic field into a voltage which drives the ISIS Master Timing Generator (MTG). The function of the MTG is to lock to the B-dot waveform and produce the fundamental timing signals MS and Delta-P.

When the MMPS is turned off, the MTG seamlessly switches over to an internally generated 50 Hz waveform to continue to produce MS and Delta-P for connected equipment.

## MS and Delta-P

MS defines the start of the machine cycle and is a 2.5  $\mu$ s negative logic pulse with a period of 20 ms. Delta-P is a negative logic 200 kHz pulse train that effectively splits up the machine cycle into 4000 data points. The zeroth pulse of Delta-P (coincident with MS) is blanked off. These two signals are used as the primary timing triggers throughout the accelerator.

## Central Timing Distributor

The ISIS Neutron and Muon source is a two target facility and, as such, the timing system must be able to differentiate between pulses to the two target stations. This is all handled by the ISIS Central Timing Distributor (CTD) which defines all of the necessary timing buses and operating rates for the machine.

The CTD produces three distinct timing buses (TBs), which are termed TB0, TB1 and TB2.

TB1 contains only pulses that are intended for Target Station 1 (TS1) and, similarly, TB2 contains pulses for TS2. TB0 is the logical “or” of TB1 and TB2 and, therefore, contains all of the timing pulses in the timing frame.

Each ISIS timing bus contains the following composition of timing signals:

- Frame Start (FS), a pulse which designates the beginning of an ISIS timing frame.
- Machine Start (MS), denotes the start of a machine cycle. The nature of the MS signal depends on the timing bus. For example, TB2 exclusively contains pulses for TS2. The terminology used when referring to MS is to append the timing bus number, i.e. MS0 appears on TB0, MS1 on TB1, and MS2 on TB2.
- Gated Machine Start (GMS), is a special version of the MS signal that is issued by the CTD when the interlock chain and beam permit is complete. It is characteristically the same as MS but denotes that beam is present. The same conventions are applied for the individual timing buses as with MS.
- Delta-P (DP), a 200 kHz pulse train that splits up each machine cycle. This effectively defines the finest resolution available in the current timing system as 5  $\mu$ s.

## The ISIS Timing Frame

When ISIS was a one target machine, it was decided that the repetition rates should vary from a full rate of 50 Hz (i.e. MS/1) down to the slowest rate of MS/128. Hence the

\* robert.washington@stfc.ac.uk

# TIME SYNCHRONIZATION AND TIMESTAMPING FOR THE ESS NEUTRON INSTRUMENTS

N. Holmberg\*, T. Bögershausen, A. Pettersson, J. Petersson, T. Brys, F. Rojas, M. Olsson, T. Richter  
European Spallation Source ERIC, Lund, Sweden

## Abstract

The European Spallation Source (ESS) will be a cutting-edge research facility that uses neutrons to study the properties of materials. This paper presents the timestamping strategy employed in the neutron instruments of the ESS, to enable efficient data correlation across subsystems and between different sources of experiment data.

ESS uses absolute timestamps for all data and a global source clock to synchronize and timestamp data at the lowest appropriate level from each subsystem. This way we control the impact of jitter, delays and latencies when transferring experiment data to the data storage. ESS utilizes three time synchronisation technologies. The Network Time Protocol (NTP) providing an expected accuracy of approximately 10 milliseconds, the Precision Time Protocol (PTP) delivering roughly 10 microsecond accuracy, and hardware timing using Micro Research Finland (MRF) Event Receivers (EVR) which can reach 10 nanoseconds of accuracy. Both NTP and PTP rely on network communication using common internet protocols, while the EVRs use physical input and output signals combined with timestamp latching in hardware. The selection of the timestamping technology for each device and subsystem is based on their timestamp accuracy requirements, available interfaces, and cost requirements.

This paper describes the different methods used for a number of device types, like neutron choppers, detectors or sample environment equipment, to synchronize operations and timestamp data.

## INTRODUCTION

The European Spallation Source (ESS) will generate intense neutron beams for scientific experiments. Data from such different data producers is collected independently without synchronisation. Correlation between measurements from different devices happens via the timestamping in post processing. Accurate measurements of neutron interactions with samples require precise timestamping of events between the various instrument components, including the neutron source, choppers, sample environments, and detectors. That makes synchronisation of clocks and accurate tagging crucial for maintaining experimental data integrity and ensuring that the acquired data can be correctly interpreted and analyzed.

ESS employs a timestamping strategy that uses absolute timestamps, based on UTC, for all data and a single source oscillator used by all timing technologies. The strategy is also to ensure that each subsystem timestamps data at the lowest possible level. This eliminates the impact of jitter on

\* nicklas.holmberg@ess.eu

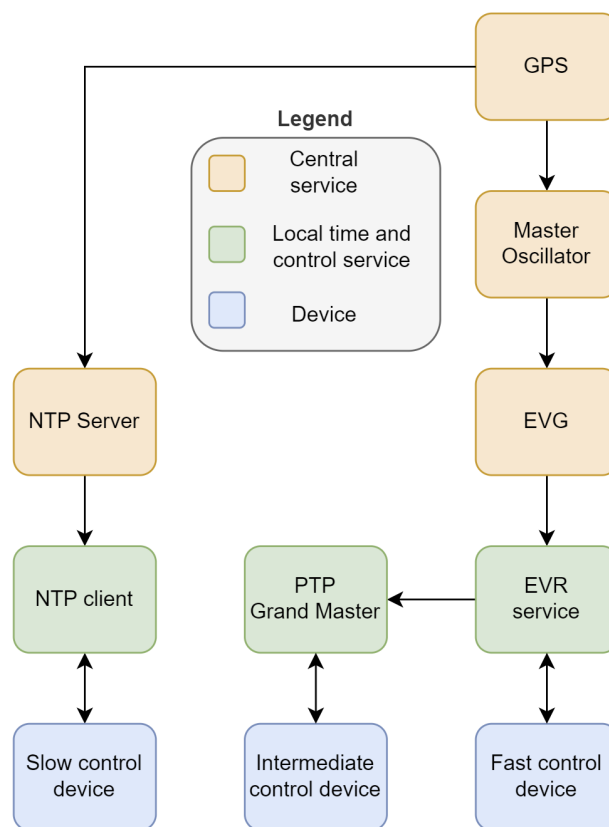


Figure 1: One central clock creates the foundation for all clocks and timestamps at the ESS facility

delays and latencies when transferring experiment data to the data storage.

The selection of the timestamping technology for each device and subsystem is based on the accuracy requirements, available interfaces, and cost requirements. In this paper, we investigate how the three different technologies, NTP, PTP and EVR, are employed, see Fig. 1. Using a few examples, we also discuss which criteria are used to select synchronization and timestamping method and how this is implemented and verified for specific devices.

### Network Time Protocol Synchronisation

Network Time Protocol (NTP) is a widely used internet-based protocol for synchronizing clocks over a computer network, including neutron instruments. NTP can achieve millisecond-level accuracy, providing satisfactory time synchronization and timestamping for many applications that do not require ultra-high precision. It can be implemented in both hardware and software, allowing for diverse applica-

## FRIB BEAM RAMP PROCESS CHECKER AT CHOPPER MONITOR\*

ZhiYong Li<sup>†</sup>, Enrique Bernal Ruiz, Masanori Ikegami, Joseph Hartford

Facility for Rare Isotope Beams (FRIB), Michigan State University, East Lansing, USA

### Abstract

Chopper in the low energy beam line is a key element to control beam power in FRIB. As appropriate functioning of chopper is critical for machine protection for FRIB, an FPGA-based chopper monitoring system was developed to monitor the beam gated pulse at logic level, deflection high voltage level, and induced charge/discharge current levels, and shut off beam promptly at detection of a deviation outside tolerance. Once FRIB beam power reaches a certain level, a cold start beam ramp mode in which the pulse repetition frequency and pulse width are linearly ramped up becomes required to mitigate heat shock to the target at beam restart. Chopper also needs to generate a notch in every machine cycle of 10 ms that is used for beam diagnostics. To overcome the challenges of monitoring such a ramping process and meeting the response time requirement of shutting off beam, two types of process checkers, namely, monitoring at the pulse level and monitoring at the machine cycle level, have been implemented. A pulse look ahead algorithm to calculate the expected range of frequency dips and rises was developed, and a simplified mathematical model suitable for multiple ramp stages was built to calculate expected time parameters of accumulated pulse on time within a given machine cycle. Both will be discussed in detail in this paper, followed by simulation results with FPGA test bench and actual instrument test results with the beam ramp process.

### INTRODUCTION

Facility for Rare Isotope Beams (FRIB) has been operating with 5 kW beam power level and in the beam power ramp up toward the design beam power of 400 kW [1]. The heavy ion beams accelerated with the driver linac are delivered to a rotating carbon target to generate secondary beams that are used for various experiments. As we increase the beam power, it is becoming more important to control transient heat load to the target to mitigate the risk of damage due to heat shock especially at restart after a beam trip. A chopper in the low energy beam line is a key element to control the beam power. The chopper deflects the beam transversely with high voltage electric field that is generated following the gate signal provided from Global Timing System (GTS), and the deflected beam is stopped at the chopper to generate the beam pulses downstream. The chopper is often used to control the beam power by controlling beam duty factor. Additionally, the chopper generates a notch in every machine cycle of 10 ms for beam diagnostics. Without the diagnostics notch, some beam diagnostics do not function properly including the

ones used for machine protection. As appropriate functioning of chopper is critical for machine protection, an FPGA-based chopper monitoring system was developed to monitor the beam gated pulse at logic level, deflection high voltage level, and induced charge/discharge current levels, and shut off beam promptly at detection of a deviation outside tolerance. Once FRIB beam power reaches a certain level, a cold start beam ramp mode which linearly ramps up the pulse repetition frequency (PRF) and pulse width (PW) becomes necessary for beam operation.

The cold start mode starts with the pulse of PW from 0.6  $\mu$ s to 2  $\mu$ s and the PRF of 2 kHz. Keeping the PW unchanged, in step one it ramps up PRF from 2 kHz to 25 kHz in 30 s with a fixed rate of 0.77 KHz/s. Keeping PRF unchanged, in step 2 it ramps up PW from its start value to 5  $\mu$ s in 30 s with a rate of 0.15  $\mu$ s/s; in step 3 it ramps from 5  $\mu$ s to 20  $\mu$ s in 40 s with a rate of 0.38  $\mu$ s/s; in step 4 ramps from 20  $\mu$ s to 39.4  $\mu$ s in 393.6 s with a rate of 0.0493  $\mu$ s/s. At last, the 25 KHz and 39.4  $\mu$ s pulse will change to 100 Hz and 9.95 ms which is called full power pulse after a transition period.

The pulses are ramped up within a repeated machine cycle (MC) time structure which starts with a diagnostic notch of 50  $\mu$ s where no beam is allowed followed by a 9.95 ms period of time where beam is allowed and ramping occurs. As the PRF increases linearly from pulse to pulse, the last pulse of a machine cycle may not fit exactly into the MC. When this is the case, the planned cycle time might be larger or less than the time left of the MC. We will call this fact as “PRF dips or rises” here. In order to keep the duty cycle constant, the PW shall rise or dip along with the PRF as it deviates from its planned value.

Beam power is adjusted using the chopper by changing the beam gated pulse [2]. The chopper monitoring system was developed to check the beam gated pulse from Global Timing System (GTS) at logic level, deflection high voltage level from High Voltage (HV) switch and induced charge and discharge current level from chopper plates [2]. The PW must not be less than 0.6  $\mu$ s because of the chopper monitor detection speed limit. It is a challenge to design a FPGA logic to be able to monitor such a ramping process and cover various cases of pulse patterns and response the fault in a timely manner.

Two types of ramp-up process checkers, a micro checker monitoring at the pulse level and a macro checker monitoring at the machine cycle level, are implemented in the design. The micro checker can detect a fault at pulse level but is not able to monitor the linear scale over an accumulated time period. The macro checker can detect a fault of linear scale over an accumulated time period but response time is slow since it has to wait for either the counts to reach beyond the expected number or for the end of the examined time period to decide if total counts is less than expected.

\* Work supported by the U.S. Dept. of Energy Office of Science under cooperative Agreement DE-SC0023633

<sup>†</sup>liz@frib.msu.edu



# MOTION CONTROLS FOR ORNL NEUTRON SCIENCE EXPERIMENTAL BEAMLINES\*

X. Geng†, A. Groff, M. Harrington, G. Taufer, and M. Pearson, ORNL, Oak Ridge, U.S.A  
G. Greene, retiree

## Abstract

This paper presents a comprehensive overview of the motion control systems employed within the neutron science user facilities at Oak Ridge National Laboratory (ORNL). The Spallation Neutron Source (SNS) and the High Flux Isotope Reactor (HFIR) at ORNL have a total of 35 neutron beam lines with numerous motors for motion control. The motion systems vary in complexity from a linear sample positioning stage to multi-axis end stations. To enhance the capabilities of these motion systems, a concerted effort has been made to establish standardized hardware and flexible software that improve performance, increase reliability and provide the capability for automated experiments. The report discusses the various motion controllers used, the EPICS-based IOCs (Input Output Controllers), high-level motion software, and plans for ongoing upgrades and new projects.

## INTRODUCTION

The Neutron Sciences Directorate at Oak Ridge National Laboratory (ORNL) maintains two independent experimental facilities providing neutron beams for user experiments: the Spallation Neutron Source (SNS), and the High-Flux Isotope Reactor (HFIR). Both facilities have multiple beamlines (19, 16) offering diverse beam properties and detector capabilities. Within each beamline experimental samples are exposed to neutron beams while measuring and recording the interactions. To expedite user time, it is desirable to automate the data acquisition process to the extent possible. Of concern here is the control and automation of sample position, orientation, and exposure within the neutron beam, to which we refer to concisely as *motion control*.

Although the experimental beamlines are quite different in configuration, sample motion control has many consistencies throughout both facilities. Our team is attempting a uniform approach to all motion control. This approach was initiated by migrating both the SNS and HFIR experimental beamlines to the Experimental Physics and Industrial Control System (EPICS) control system [1]. This effort began formally in 2012 [2, 3]. Since then, all SNS beamlines and half the HFIR beamlines have adopted EPICS systems, most utilizing the Control System Studio (CS-Studio or “CSS”) user interface [4].

## MOTION CONTROL SYSTEMS

Motion systems are used for diverse functions within SNS and HFIR beamlines. The following are typical

\*SNS is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U.S. Department of Energy

† geng@ornl.gov

applications:

- Linear and rotary sample positioning stages
- Slits sizing - both beams and detectors
- Neutron guides
- Detector positioning
- Attenuators
- Hexapods
- Shutters

However, within all systems the fundamental components are stepper and servo motors along with their associated hardware (e.g., controllers, encoders, limit switches, etc.) and software support (drivers, sequencers, user interface, etc.).




Figure 1: Control System Architecture.

## Overview

Figure 1 is a hierarchical representation of a typical motion control system used at ORNL. At the base of the diagram is the motion system itself, examples of which are listed above. (Note that the motors are embedded within these systems.) The next level shows the hardware controllers required by the motors, covered in detail below. The middle layer between hardware and user interface is the EPICS IOCs through ethernet or serial-ethernet connection. Here EPICS motor records provide driver compatibility with the motor controllers. The final layer in Fig. 1 represents the motion control software. This includes software supporting basic control signal access to high-level client applications potentially automating experimental processes over the Channel Access (CA) protocol.

## Motion Control Hardware

The stepper motors require specialized hardware interfaces called motion controllers. These devices direct the motor action according to predetermined software inputs.

# USING ARUCO CODES FOR BEAM SPOT ANALYSIS WITH A CAMERA AT AN UNKNOWN POSITION

W. Smith<sup>†1</sup>, M. Arce<sup>1,2</sup>, C.E. Jimenez<sup>1</sup>, I. Rudolph<sup>1</sup>, M. Gorgoi<sup>1,3</sup>, R.G. Wilks<sup>1,3</sup>, M. Bär<sup>1,3,4,5</sup>

<sup>1</sup> Helmholtz-Zentrum Berlin für Materialien und Energie GmbH (HZB), Berlin, Germany

<sup>2</sup> INN, CNEA-CONICET, Centro Atómico Bariloche, Argentina

<sup>3</sup> Energy Materials In-Situ Laboratory Berlin (EMIL), HZB, Berlin, Germany

<sup>4</sup> Helmholtz Institute Erlangen-Nürnberg for Renewable Energy (HI ERN), Berlin, Germany

<sup>5</sup> Friedrich-Alexander Universität Erlangen-Nürnberg (FAU), Erlangen, Germany

## Abstract

Measuring the focus size and position of an X-ray beam at the interaction point in a synchrotron beamline is a critical parameter that is used when planning experiments and optimizing beamline performance. Commonly this is performed using a dedicated ultra-high vacuum (UHV) "focus chamber" comprising a fluorescent screen at an adjustable calibrated distance from the mounting flange and a camera on the same plane as the beam. Having to install a large piece of hardware makes regular checks prohibitively time consuming, particularly on beamlines with permanently installed endstations. A fluorescent screen can be mounted to a sample holder and moved using a manipulator in the existing end-station and a camera pointed at this to show a warped version of the beam spot at the interaction point. The warping of the image is an effect of perspective related to the relative positions of the camera and the screen, which is difficult to determine and can change and come out of camera focus as the manipulator is moved. This paper proposes a solution to this problem using a fluorescent screen printed with ArUco codes which provide a reference in the image [1]. Reference points from the ArUco codes are recovered from an image and used to remove perspective and provide a calibration in real time using an EPICS ArcaDetector plugin using OpenCV [2, 3]. This analysis is presently in commissioning and aims to characterize the beam spots at the two-color beamline of the EMIL laboratory at the BESSY II light source.

## THE ENDSTATION

The Sissy-1 endstation is located at one of the focus points of the two-color EMIL beamline at BESSY II [4]. X-ray light from either soft or hard X-ray beamline can be focused on the same interaction point, allowing the continuous variation of the photon energy between 80 and 10,000 eV at one spot. The endstation comprises an UHV chamber with several fluorescence detectors, an electron analyzer, and an X-ray emission spectrometer. Samples can be brought into the chamber from the UHV backbone that connects the endstation with a variation of deposition tools, other characterization setups, and an inert-gas glovebox [5] – most prominently battery, catalyst, and solar cell related layer stacks are studied (see Fig. 1). The chamber has a 5-axis manipulator and several spare view ports. One is

equipped with a wide-angle lens to monitor the sample transfer, the others have zoom macro lenses which are used for looking at the sample for sample positioning/monitoring. The connection of the Sissy-1 endstation to the UHV backbone and to a beamline section leading to a second focus point via "pass through" geometry makes temporary replacement of the endstation with a focus chamber a practical impossibility.




Figure 1: Sissy-1 endstation connected to UHV backbone at the focus point of the two-color EMIL beamlines.

The samples that are measured in the end station can be inhomogeneous. Sometimes they can also be located inside a cell that allows the sample to be in ambient pressure conditions while the UHV conditions in the beamline and the analysis chamber can be maintained. These cells comprise X-ray transparent windows through which the X-rays must pass to reach the sample. This means that the size and shape of the X-ray beam is an important factor in the design of these cells, alignment procedures, and planning of measurement routines. In addition, knowing the size and shape of the soft and hard X-ray beams at the interaction point is critical to the verification of the beamline setup and optimization, and a crucial prerequisite for overlapping the two beam spots. Viewing these parameters in real time while adjusting the position of beamline optical elements is essential for beamline commissioning and for researchers to

<sup>†</sup> William.smith@helmholtz-berlin.de

# LImA2: EDGE DISTRIBUTED ACQUISITION AND PROCESSING FRAMEWORK FOR HIGH PERFORMANCE 2D DETECTORS

S. Debionne, A. Homs Puron <sup>†</sup>, J. Kieffer, R. Ponsard, L. Claustre, A. Götz, P. Fajardo  
 European Synchrotron Radiation Facility (ESRF), Grenoble, France

## Abstract

The continuous increase in the data rate of high-performance 2D detectors for photon science makes more and more difficult to perform data acquisition (DAQ) and online processing (ODA) on a single computer. This work presents the LImA2 framework for distributed image detector DAQ & ODA, targeting scalability and low-latency by means of efficient use of system resources, including hardware accelerators. The system controlling the PSI Jungfrau-4M detector, in operation at the ESRF ID29 Serial MX beamline is described. A dual IBM Power9 computer implementation using GPUs provide data sparsification and on-the-fly identification of diffraction peaks at a continuous rate of 500 frames/s. A LImA2 plugin for the ESRF Smartpix detector, based on the high-performance RASHPA data transfer architecture is also presented. Finally, this work includes an update on the current developments for Dectris Eiger2 support and integration in BLISS, the ESRF new beamline control system.

## INTRODUCTION

LImA [1], a Library for Image Acquisition, was developed at the ESRF to simplify and standardize of the integration of 2D detectors in synchrotron beamline experiments. By abstracting the image generation capabilities of the detectors, LImA separates the manufacturer’s Software Development Kit (SDK) from highly multi-threaded data processing algorithms. Such approach allowed the integration over the last twenty years of tens of different detectors with a unified interface, making LImA the de-facto 2D data acquisition (DAQ) reference software in the Tango synchrotron community similar to areaDetector [2] in the EPICS community.

During the last twenty-five years, the data throughput of 2D detectors at the ESRF has increased by about thirteen times every decade, exceeding the 10 GBytes/s by 2020. Assuming this trend, future detectors will surpass

100 GByte/s by 2030. Therefore, extracting the meaningful information from the raw data and reducing the data volumes becomes more and more critical for data manipulation and storage. Although specific implementations combine high-performance computing with hardware accelerators like FPGA and GPGPUs [3], such demanding tasks cannot be always performed by a single computer in a generic way. Alternatively, using a data centre cluster normally implies an additional latency, affecting online data analysis and fast experiment feedback. Finally, the original LImA architecture was not designed to run on multiple computers and implementing such functionality would be as difficult as starting a new project from scratch.

The LImA2 project was initiated with the aim of addressing all these challenges by providing a scalable, distributed data acquisition (DAQ) and low latency processing framework for high-performance, high-throughput 2D detectors. This paper presents the design considerations of LImA2, the details of the implementation and the current status of the project, including practical implementation such as the PSI/Jungfrau-4M application at the ESRF ID29 for Serial Macromolecular Crystallography (MX).

## SYSTEM TOPOLOGIES

Three main system topologies have been identified so far, shown in Fig. 1. The partial frame dispatch topology associates one DAQ/Processing computer to each independent module in a tiled array. The full frame dispatch topology is adapted to detector systems with the capability of sending different frames to different computers.

Depending on the specific needs of the application, one topology can be better adapted to the processing algorithms than the other. For instance, tomography and XPCS techniques can have optimized pipelines for partial frame dispatch topology, while radial integration of diffraction and scattering images is easier to obtain with full frame dispatch.

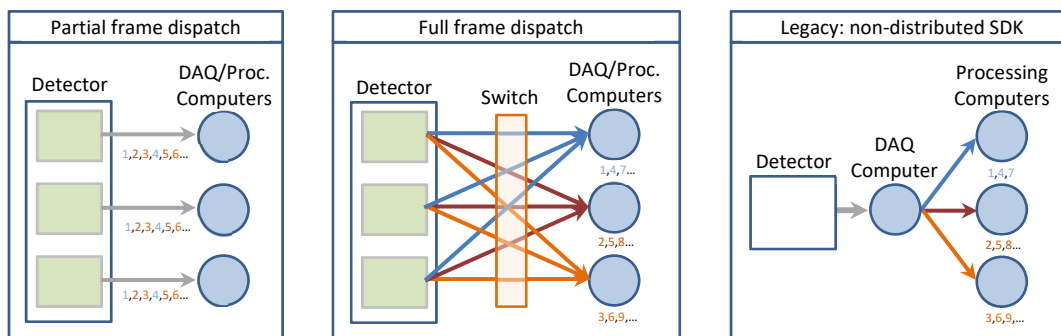


Figure 1: Currently considered LImA2 topologies.

<sup>†</sup> alejandro.homs@esrf.fr

# ROBOTIC PROCESS AUTOMATION: ON THE CONTINUITY OF APPLICATIONS DEVELOPMENT AT SOLEIL

L. E. Munoz\*, Y-M. Abiven, M. E. Couprie, M. Valléau, A. Nouredine, J. Pérez, A. Thureau, Synchrotron SOLEIL, Saint-Aubin, France

## Abstract

For some years now SOLEIL has developed and put into operation robotic applications, using 6-axis robotic arms, to automate some of its beamlines and some processes of magnetic measurements. In the last year, SOLEIL has been working on the development of two new robotic applications, having thus continuity in the development of applications using its robotic standard. This paper describes these two new applications that are being developed to automate the injection of liquid samples for the SWING beamline and to automate the mechanical and magnetic adjustment of the modules that compose an insertion device.

## INTRODUCTION

SOLEIL is a synchrotron facility which offers a large variety of experimental techniques and samples environments to its users, with its 29 beamlines and a wide energy range from THz to Hard X-ray. With the upgrade project, SOLEIL II will provide an increase in brilliance and coherence and all the beamlines will benefit from the unique properties of the new accelerator: measures at the nanometer scale, experiences up to 1,000 times more sensitive and up to 10,000 times faster. These properties will open up experimental new possibilities and also new scientific and engineering challenges to address. Automation is then considered an important piece of the upgrade project in order to achieve the previously mentioned properties.

SOLEIL has spent the last few years working on developing robotic skills to be proficient in process automation and it is especially focused on the arm robot-based automation due to the great advantages that 6-axis robots offer. The specialization includes the definition of hardware interfaces between systems, developing a flexible TANGO-based software interface as well as selecting the Stäubli brand [1]. This standardized approach was designed to guarantee proper implementation, easy deployment and efficient interventions and maintenance for technical staff. The last two robotic applications being developed with the SOLEIL robotic standard are here described. The first application aims to automate the injection of liquid samples for BioSAXS analysis of protein solutions at the SWING beamline. Hence, a pick-and-place 6-axis robot is used to automate the process of sample aspiration and sample dispensing into the measurement cell. The second application, is dedicated to the automation of the mechanical and magnetic adjustment of magnetic modules of the future insertion devices (undulator) of SOLEIL II.

\* laura.munoz-hernandez@synchrotron-soleil.fr

## PIPETTING APPLICATION

Liquid handling is a basic laboratory procedure that is done regularly across many industries: biological and chemical research, diagnostics, drug manufacturing, food quality control, among others [2]. Liquid handling could be defined as the transfer of pre-determined measured volumes of liquid from one container to another. It can be performed manually by using tools called pipettes or by using automated solutions called liquid handling workstations or simply liquid handling robots, with Cartesian robotic systems<sup>1</sup> being the most representatives of automated liquid handling systems commercially available. However these liquid handling robots are useful only for a limited range of applications, since these are devices integrated in self-contained and rigid workstations with a specified and limited workspace [3]. When Cartesian robots do not meet the application needs, arm robotic systems can also be a solution for liquid handling processes if equipped with suitable pipettes.

Pipettes come in all shapes and sizes, with a variety of different technologies incorporated into their design function. Whether mechanical or electronic, they can handle volumes ranging from a few nanoliters up to milliliters. Electronic pipettes are more precise and accurate than the mechanical ones because they use a motor to control piston movement, so they always dispense exactly the volume programmed. Pipetting protocols, including volumes and speeds, can also be pre-programmed and saved so that they are executed in the same way every time [4].

On the other hand, the SWING Beamline allows simultaneous Small-Angle X-ray Scattering (SAXS) and Wide-Angle X-ray Scattering measurements (WAXS) in the 5-16 keV energy range. SWING beamline helps answering numerous questions related to soft condensed matter, conformation of macro-molecules in solution and material sciences [5]. Currently, the beamline has a fully automated combined system, including an auto-sampler robot and online High-Performance Liquid Chromatography (HPLC), to carry out biological Small-Angle X-ray Scattering (BioSAXS) experiments. Both systems are connected to a quartz capillary cell placed within a vacuum chamber, named in this paper also as the measurement cell, see Fig. 1.

This combined system has worked perfectly on the beamline for more than a decade, however, new scientific requirements have been arisen and nowadays the beamline has to be able to inject the liquid samples as close as possible to the measurement cell in order to avoid wasting sample in the tube that connects the auto-sampler robot to the mea-

<sup>1</sup> A Cartesian robot has three linear joints (or a combination of them) that use the Cartesian coordinate system (X, Y, and Z).

# ULTRA-HIGH THROUGHPUT MACROMOLECULAR CRYSTALLOGRAPHY DATA COLLECTION USING THE BLUESKY FRAMEWORK

D. Perl<sup>†</sup>, N. Frisina, D. Oram, N. Paterson, Diamond Light Source, Didcot, UK

## Abstract

At Diamond Light Source (DLS), several Macromolecular Crystallography (MX) beamlines focus on, or include in their activities, completely automated data collection. This is used primarily for high throughput collection on samples with known or partially known structures, for example, screening a protein for drug or drug fragment interactions. The automated data collection routines are currently built on legacy experiment orchestration software which includes a lot of redundancy originally implemented for safety when human users are controlling the beamline, but which is inefficient when the beamline hardware occupies a smaller number of known states. DLS is building its next generation, service-based, Data Acquisition (DAQ) platform, Athena, using NSLS-II's *Bluesky* experiment orchestration library. The *Bluesky* library facilitates optimising the orchestration of experiment control by simplifying the work necessary to parallelise and reorganise the steps of an experimental procedure. The MX data acquisition team at Diamond is using the Athena platform to increase the possible rate of automated MX data collection both for immediate use and in preparation to take advantage of the upgraded Diamond-II synchrotron, due in several years. This project, named *Hyperion*, will include sample orientation and centring, fluorescence scanning, optical monitoring, collection strategy determination, and rotation data collection at multiple positions on a single sample pin.

## INTRODUCTION AND MOTIVATION

### Motivation

In order to take advantage of the increased beam intensity which will be available following the Diamond-II upgrade [1] and facilitate research which depends on high-throughput techniques we are developing a new unattended data collection (UDC) application. This application is named *Hyperion*, and is built on the *Bluesky* [2] library developed at Brookhaven National Laboratory. *Hyperion* is envisaged to allow the rapid collection of thousands of single-crystal X-ray diffraction datasets per day. This quantity of data is desired in order to screen thousands of candidate molecules against proteins to identify candidates for drugs or to probe structure-function relationships of those proteins.

In addition to the above, we aim for *Hyperion* to help support DLS users in the dark period during the machine upgrade by allowing us to process their samples as usual with much more limited beamtime at another facility. In the nearer future, indeed, as is already being implemented, *Hyperion* improves the existing unattended data collection (UDC) capability at Diamond beamline I03.

### Current Data Collection Software

At Diamond, for the last 20 years since the beginning of our operation, we have used the Generic Data Acquisition (GDA) [3] platform. The generic nature of this platform has allowed us to reuse software across beamlines. Data acquisition, in the context of synchrotron beamlines, refers in general to the manipulation of the position of scientific samples (whether biological or material) in an X-ray or other photon beam, and the detection of transmitted, diffracted, reflected or radiated photons from the sample. This includes techniques such as X-ray crystallography, scattering techniques, imaging techniques, and absorption, fluorescence, infrared, or Raman spectroscopy.

GDA has a monolithic architecture and is, in general, not built using contemporary software development best practices; notably, it is poorly tested. It is built in Java, and uses Jython for rapid development of experimental procedures with access to the Java objects which represent beamline hardware. Jython is limited to python 2.7, which is no longer supported. Further, we are reaching the limits of the flexibility of data collection scripts written in plain python.

### Bluesky

*Bluesky* [2] is a Python 3 library designed for experiment control applications. It was primarily developed at Brookhaven National Lab facility NSLS-II, and currently has contributors from a large collaboration of facilities including DLS.

It reduces the boilerplate needed to write scripts (termed 'plans') to run experiments and provides means to capture data and metadata in a variety of formats. It interfaces with the Experimental Physics and Industrial Control System (EPICS) through *Ophyd* [4] devices, essentially Python wrappers for collections of EPICS process variables with some additional convenience functions.

The *Bluesky* library provides for useful logical separation between hardware operations and data management. Plans only involve the steps which need to be taken with the beamline hardware in order to run an experiment, and data is managed through callback functions in separate modules. The main routine of *Bluesky* is termed the 'RunEngine', which consumes plans and translates them into instructions for *Ophyd* devices.

### Hyperion

*Hyperion* is an application revolving around a set of *Bluesky* plans written to support UDC for macromolecular crystallography. Currently it includes routines for sample centring using optical images and X-ray grid scans, capturing of optical snapshots, and rotation data collection. Additional planned features include simultaneous fluorescence data collection, collection strategy determination,

<sup>†</sup> david.perl@diamond.ac.uk

# PIEZO MOTOR BASED HARDWARE TRIGGERED NANO FOCUS CAUSTIC ACQUISITION

L. B. C. Campoi\*, L. E. P. Vecina, G. S. R. Costa, G. B. Z. L. Moreno, N. L. Archilha  
Brazilian Synchrotron Light Laboratory (LNLS),  
Brazilian Center for Research in Energy and Materials (CNPEM),  
Zip Code 13083-100, Campinas, Sao Paulo, Brazil

## Abstract

MOGNO is a micro and nano X ray tomography beamline at Sirius. It was designed to operate with a cone beam, allowing for zoom tomography experiments via the use of a set of elliptical mirrors in a Kirkpatrick-Baez (KB) system. The main source produced by the KB system has  $120 \times 120 \text{ nm}^2$ , posing a challenge on the focus evaluation system, that has to probe such focus in a timely manner.

To tackle the KB system alignment evaluation, a diagnostic comprised of a stack of three linear inertia drive piezo stages and a fluorescence detector, acquiring data via hardware-triggered mesh scans was implemented. In the piezo stack, the stages are mounted along the X (horizontal, perpendicular to the beam path), Z (along the beam path) and YZ beamline directions. Moreover, a kinematic transformation was implemented due to the fact that a stage is placed at an angle and that the beam is not aligned with the sample stage stack. Mesh scans in the XZ and YZ can be divided in two parts: hardware triggered line scan acquisition along X or Y and software triggered steps along Z between scans. In this manner, the control is done via a collection of low-level controller macros and Python scripts, such that during the scans, the piezo controllers communicate with each other and the detector via digital pulses, orchestrated by the in-house TATU (Timing and Trigger Unit) software, reducing dead time between acquisition points.

The proposed system proved to be reliable to acquire beam profiles, providing caustics in both horizontal and vertical directions, and the acquired focus caustics indicate that the main source has a size of approximately  $416 \times 480 \text{ nm}^2$  at the moment.

## INTRODUCTION

Mogno is a micro and nano tomography beamline designed to provide users with high resolution, flexibility and high data throughput, via a cone-beam geometry with high flux. To achieve the desired resolution of 120 nm in each direction, the optical system is composed by three mirrors, a pair of elliptical multilayer mirrors in a KB (Kirkpatrick-Baez) system, that is also used to select the beam's energy (22 and 39 keV at the same time or 67.5 keV) [1], and a pre-KB elliptical mirror positioned as a horizontal focusing mirror (HFM) to compensate for inhomogeneities caused by the KB's HFM in the horizontal axis, as shown in Fig. 1. The mirrors have tight alignment budgets and at the first mirror

a nano focus is already produced in the horizontal, which poses challenges on how to evaluate such focus spot [2]. A similar problem occurs while tackling the KB's system alignment, however, the KB's focus must be evaluated in both horizontal and vertical planes, requiring a system with at least three degrees of freedom (DOF).

## EXPERIMENT DESCRIPTION

To evaluate the focus, a strategy based on the knife-edge method is proposed, where a sample is scanned across the beam and the fluorescence signal generated by the interaction of the X ray beam with the sample is collected in a fluorescence detector [2]. In the case of the KB's focus, the sample used is a calibration pattern from Applied Nanotools Inc. (ANT), which has Au features deposited on a Si substrate, and the detector used was a Hitachi Vortex ME-4 SDD, with a XSPRESS3 electronic from Quantum Detectors. For evaluation of the horizontal and vertical focus directions, L-shaped features with thickness of 600 nm and varying widths (1  $\mu\text{m}$ , 500 nm, 250 nm) were available.

Before each scan, the desired sample feature was positioned approximately on the beamline focus. The data acquisition was then composed by planar mesh scans, where fast line scans in a single direction perpendicular to the beam were executed, followed by single steps along Z. The subsequent scans were executed swapping the line scan endpoints, forming a snake motion. Data processing was done as described in [2]. The focus size is then acquired at different pitch angles of the KB mirrors, which can be compared to refine the beamline resolution.

## MECHANICAL ASSEMBLY

The diagnostic device assembly designed for the evaluation of Mogno KB's focus is an iteration on the assembly used during Mogno's pre-KB mirror alignment [2]. The assembly is composed by three linear piezo stages from Physik Instrument GmbH (PI) (QMotion series), which were selected based on the requirement of long range motion (millimeter scale), with precision in the nanometer scale. The stage set selected is composed by two stages with 13 mm of range (Q-545.140 QMotion) and one with 26 mm (Q-545.240 QMotion). The stage set was assembled (from bottom up) with the Q-545.240 positioned along the beamline's Z axis, a Q-545.140 positioned along the beamline's X axis and the second Q-545.140 was positioned at a 15° angle in the YZ plane. This assembly can be seen in Fig. 2, and it was chosen both to comply with the stages' mass load

\* lucca.campoi@lnls.br

## VIDEO COMPRESSION FOR areaDetector

B. A. Sobhani, Oak Ridge National Laboratory, Oak Ridge, TN, USA

### Abstract

At neutron sources such as SNS and HFIR, neutrons collide with neutron detectors at a much lower rate than light would for an optical detector. Additionally, the image typically does not pan or otherwise move. This means that the incremental element-by-element differences between frames will be small. This makes neutron imaging data an ideal candidate for video-level compression where the incremental differences between frames are compressed and sent, as opposed to image-level compression where the entire frame is compressed and sent. This paper describes an EPICS video compression plugin for areaDetector that was developed at SNS.

### INTRODUCTION

Plugins exist to compress/decompress frames in areaDetector. This kind of compression is useful and helps to reduce size of transmitted data, but is limited by the fact it only compresses in space (the x/y coordinates of the image), but not in time. So a video consisting of 100 frames of the same image would compress to 100 multiplied by the compressed size of a single frame. Video codecs, such as H.264 and others, employ differential frame compression [1] so that a series of 100 identical images would compress to something much smaller than 100 multiplied by the size of a single compressed.

### LIBAVCODEC

In this section some terminology is briefly discussed.

#### Packets

One can think of a packet in libavcodec as being a compressed frame, with the exception that a packet does not by itself hold enough information to reconstruct a frame. A packet can only reconstruct a frame in conjunction with preceding frames. In the event of a missing packet, the video stream will experience some distortion in the subsequent frames.

#### Frames

The frame, in libavcodec, is the actual image that is displayed to the user. Note that this definition of frame is different from that used when discussing certain codecs, where I-frame, and P-frame, B-frame are more analogous to “packets” described in the previous section.

### COMPRESSION

Compression was done in areaDetector by making additions to the ADSupport module. Frame data from the NDArray is sent through to an AVContext object. Packets were then extracted and broadcast via PVAccess.

### DECOMPRESSION

Decompression code was developed as additions to the same ADSupport module. Then an ImageJ plugin was developed which is just a java wrapper to those functions. Doing it this way will allow for video decompression to be easily expanded to a variety of clients.

### CURRENT STATUS AND TESTING

At the time of writing this paper, the proof of concept was very recently demonstrated to work. The test consisted of taking images (at the NDArray level in an areaDetector plugin), producing a packet (based on the current image and the previous images that were sent), storing the packet inside a PVA stream, and decoding this PVA stream in ImageJ to get the original image.

This particular test was done with the mpeg-1 codec – this was an arbitrary choice, and since libavcodec is a general interface to many codecs it will be easy to test other codecs.

Since the proof of concept was only recently demonstrated, more thorough performance testing has not yet taken place. Controlled tests should be done to compare the size of individually compressed frames vs video compression. Additionally, there should be testing done to measure degradation of compression over time. For neutron scattering data, compression typically gets much worse over time, as noise from the neutron events get accumulates. Video compression is expected to be much more robust in this aspect – testing must be conducted to confirm it.

### ACKNOWLEDGEMENTS

Acknowledgements to Kaz Gofron and supervisor Bogdan Vacaliuc for preliminary discussions on how to implement this. Also to VULCAN beamline for allowing testing of compression using their beamline.

### REFERENCES

- [1] <https://ottverse.com/i-p-b-frames-idr-keyframes-differences-usecases/>

# NEW GENERATION QT CONTROL COMPONENTS FOR HIGH LEVEL SOFTWARE

G. Strangolino, G. Gaio, R. Passuello, Elettra Sincrotrone Trieste, Trieste, Italy

## Abstract

A new generation of Qt graphical components, namely *cumbia-qtcontrols-ng* is under development at ELETTRA. A common engine allows each component to be rendered on traditional Qt Widgets and scalable Qt Graphics Items alike. The latter technology makes it possible to integrate live controls with static SVG in order to realize any kind of synoptic with touch and scaling capabilities. A pluggable zoomer can be installed on any widget or graphics item. Apply numeric controls, Cartesian and Circular (Radar) plots are the first components realized.

## MOTIVATION

The ELETTRA synchrotron light source in Trieste, Italy, will be subjected to a major upgrade in the forthcoming years. The rationale is a substantial reduction of the emittance of the stored electron beam, targeting its levels so as to provide a diffraction limited X-ray source also in the horizontal plane. On that account, the new machine, ELETTRA 2.0, will provide intense nano-beams in the range of VUV to X-rays for the analytical study of matter with very high spatial resolution. New magnet design, innovative vacuum technology and revolutionary beam monitoring and orbit feedback systems provide at present a solid basis for the realization of the new machine. New lattice design shall be adapted to the existing ELETTRA storage ring tunnel, building and infrastructure, including the injector and the beam-lines. This implies the substitution of the twelve arcs of the existing storage ring with as many new ones of almost identical length. The new magnetic lattice reduces the present horizontal emittance of 7 nm-rad down to 0.147 or 0.212 nm-rad at 2.4 GeV, largely increasing the brilliance and coherence of the X-ray beam, whilst preserving the injection system. The implementation of the storage ring upgrade will enable ELETTRA to maintain its leadership in the context of synchrotron light sources operating within the same energy range.

ELETTRA 2.0 and the continuous challenges posed by the FERMI@ELETTRA free electron laser encouraged the development of a new generation of Qt graphical components, based on the *cumbia* [1] libraries and the Qt SVG [2] and Graphics Scene [3] technologies. The acquisition of a 55 inches 4K touch screen for the control room inspired the design of the new graphical objects with touch interaction in mind. They are offered by the new *cumbia-qtcontrols-ng* library, which flanks the *cumbia-svg* module.

<sup>1</sup> a minimum impact on the object's "paint" code is an objective, presently not completely accomplished.

## TECHNOLOGY

*cumbia-qtcontrols-ng* is a graphical library providing special twofold components capable of drawing themselves and accept interaction both within traditional Qt widgets and as scalable Qt items within a Qt Graphics Scene. The integration with *cumbia-svg* allows drawing composite *synoptic* user interfaces where static SVG items stand beside live items such as plots, gauges, numeric input objects, and so on.

### Qt Graphics View Framework

Qt Graphics View provides a surface (scene) for managing and interacting with a large number of custom-made 2D graphical items, and a view widget for their visualization, with support for zooming and rotation. Graphics View provides an item-based approach to model-view programming. Several views can observe a single scene, and the scene contains items of varying geometric shapes.

### Qt SVG

Scalable Vector Graphics (SVG) is an XML-based language for describing two-dimensional vector graphics. Qt provides classes for rendering and displaying SVG drawings in widgets and on other paint devices.

### cumbia libraries

The *cumbia* libraries are a multi threaded set of components written in C++ aimed at interfacing the lower level control system servers to the graphical user interfaces for the control room. The separate layers and plugins that make up the framework are designed to integrate with Tango [4], EPICS [5] and potentially any other distributed control system and offer an engine independent interface to the clients.




Figure 1: Main components of the *cumbia* library.

An *http* module, relying on the PUMA [6] services, allows for Qt applications running everywhere natively. This strategy has proved to be effective for developers as well as users working remotely.



# THE MICRO-SERVICES OF CERN'S CRITICAL CURRENT TEST BENCHES

P. Koziol, C. Barth, J. Fleiter, S. Hopkins, H. Reymond, C. Charrondiere, A. Ballarino,  
O. Ø. Andreassen, T. Boutboul  
CERN, Geneva, Switzerland

## Abstract

In order to characterize the critical-current density of low temperature superconductors such as niobium-titanium (Nb-Ti) and niobium-tin (Nb<sub>3</sub>Sn) or high temperature superconductors such as magnesium-diboride MgB<sub>2</sub> or Rare-earth Barium Copper Oxide REBCO tapes, a wide range of custom instruments and interfaces are used.

The critical current of a superconductor depends on temperature, magnetic field, current and strain, requiring high precision measurements in the nano Volt range, well-synchronized instrumentation, and the possibility to quickly adapt and replace instrumentation if needed. The micro-service-based application presented in this paper allows operators to measure a variety of analog signals, such as the temperature of the cryostats and sample under test, magnetic field, current passing through the sample, voltage across the sample, pressure, helium level etc.

During the run, the software protects the sample from quenching, controlling the current passed through it using high-speed field programmable gate array (FPGA) systems on Linux Real-Time (RT) based PCI eXtensions controllers (PXIe). The application records, analyzes and reports to the external Oracle database all parameters related to the test.

In this paper, we describe the development of the micro-service-based control system, how the interlocks and protection functionalities work, and how we had to develop a multi-windowed scalable acquisition application that could be adapted to the many changes occurring in the test facility.

## INTRODUCTION

The upgrade project for the Large Hadron Collider (LHC) aims to achieve higher luminosity [2], necessitating the use of superconducting magnets with higher magnetic fields, for which Nb<sub>3</sub>Sn was chosen as the conductor material of choice. It's worth noting that Nb<sub>3</sub>Sn is more brittle and challenging to manufacture into cables compared to the Nb-Ti conductor material used in the current LHC magnets.

Additionally, this high luminosity LHC project (HL-LHC) includes the installation of eight Superconducting Links (SC) comprised of high-current magnesium diboride (MgB<sub>2</sub>) cables cooled by helium gas, and connected to rare-earth-barium-copper-oxide (REBCO) cables that can operate at 60K [3].

To ensure the proper assembly of these various wires and tapes, especially for the creation of Rutherford cables for magnets and round cables for the SC Link, comprehensive characterization of these materials is crucial. The critical-

current density ( $J_c$ ) is the main parameter for assessing superconductors, and depends on various factors such as temperature, strain, and magnetic field. The most effective method for measuring the critical current is the transport method [1], in which current flows through the sample, and the voltage is measured along its length. A typical result of such tests is illustrated in Figure 1.




Figure 1: Plot of voltage vs. current for a typical superconductor sample measured at CERN.

To compare different superconductors, the critical current ( $I_c$ ) is divided by the cross-sectional area ( $A$ ) of the conductor, defining the critical current density ( $J_c$ ) [1]. The selection of the area and other methods for calculating critical current density fall beyond the scope of this article.

Historically, four small cryostats were employed to test the critical current density of superconducting strands, while one large cryostat, FRESKA (Facility for the Reception of Superconducting Cables), was dedicated to characterizing the critical current of superconducting cables.

The existing measurement systems used in these test benches had become outdated, relying on older computers running SELinux and LabVIEW® 2012. Furthermore, the need to incorporate systems for characterizing High Temperature Superconductor (HTS) tapes and the introduction of a new FRESKA2 system prompted a renovation campaign in 2020. In this paper, we will elaborate on the new demands and the developments associated with this campaign.

## RENOVATION CAMPAIGN

The diverse range of measurement equipment already in existence most of them delivered by NI, much of which has been calibrated during previous measurement campaigns, often relies on fixed and challenging-to-replace power converters. Consequently, it was determined that employing PXIe chassis as the primary platform, coupled with NI

# ROLLING OUT A NEW PLATFORM FOR INFORMATION SYSTEM ARCHITECTURE AT SOLEIL

G. Abeillé, B. Gagey, Y. M. Abiven, Synchrotron SOLEIL, Gif-sur-Yvette, France  
P. Grojean, C. Rognon, F. Quillien, V. Szyndler, EMOXA\*, Boulogne-Billancourt, France

## Abstract

SOLEIL Information System is a 20-year legacy with multiple software and IT solutions following constantly evolving business requirements. Lots of non-uniform and siloed information systems have been experienced increasing the IT complexity. The future of SOLEIL (SOLEIL II) will be based on a new architecture embracing native support for continuous digital transformation and will enhance user experience. Redesigning an information system given synchrotron-based science challenges requires a homogeneous and flexible approach.

A new organizational setup is starting with the implementation of a transversal architectural committee. Its missions will be to set the foundation of architecture design principles and to foster all projects' teams to apply them. The committee will support the building of architectural specifications and will drive all architecture gate reviews.

Interoperability is a key pillar for SOLEIL II. Therefore, a synchronous and asynchronous inter-processes communications is being built as a platform to connect existing systems and future ones; it is based both on an event broker and an API manager. An implementation has been developed to interconnect our existing operational tools (CMMS, Computerized Maintenance Management System and our ITSM portal, Information Technology Service Management). Our current use case is a brand new application dedicated to samples' lifecycle interconnected with various existing business applications.

This paper will detail our holistic approach for addressing the future evolution of our information system, made mandatory given the new requirements from SOLEIL II.

## PLUS JOURNEY

SOLEIL has launched a Proof Of Concept (POC) since early 2021 to test new technologies helping the design of the future SOLEIL Information System (IS). The main driver was to anticipate new needs for SOLEIL II [1] where we foresee a global strong requirement to answer the increasing demands from the business. Consequently, Soleil's current information system may require a deep revamping to provide a fluid integration between all components.

To achieve it, SOLEIL seeked consulting support from external experts from a French company, EMOXA, specialized to foster organisations for transforming their IS. EMOXA has primarily analysed our current IS which has confirmed our assumptions: SOLEIL has been working mostly in a silo-ed organisation that will inevitably end up with the 'spaghetti stack syndrome', i.e. resulting of point to point,

\* <https://emoxa.fr>

synchronous, ad-hoc communication links between all the interconnected systems. So a transformation becomes mandatory for SOLEIL II stakes.

Consequently, EMOXA and SOLEIL have jointly deployed a first simple use case to challenge this hypothesis and identify how to adapt and not revamp the whole IS. A small platform, named PLUS (PLaform d'Urbanisation du Synchrotron SOLEIL, or SOLEIL's PLatform for Information System Design), has been set up with Apache Kafka [2] for managing asynchronous communications and WSO2 API Manager [3] for synchronous communications (cf Fig. 1). The result is an interconnection between existing applications by aggregating data coming for our DUO tool (Digital User Office, SUN Set [4]) and a beamline acquisition system (FlyScan [5]) with a limited amount of work. The outcome was to continue the process of setting up such a platform and to extend it to real use cases.

PLUS is now in MVP (Minimum Viable Product [6]) stage (, still with the support from EMOXA company. Before going to production, we have done a gate review of the existing use cases and the expected ones. The spotted use cases will be addressed thereafter.

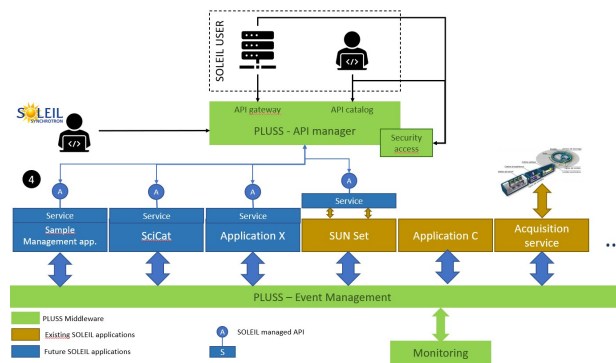


Figure 1: PLUS global overview.

During our audit, we have discovered over already heterogeneous systems, numerous people or teams with diverse development methodologies or technical backgrounds. So PLUS challenges are coming both from organisational and technical angles.

## TOWARDS SOLEIL INFORMATION SYSTEM ARCHITECTURE'S GOVERNANCE

A proficient Information System (IS) architecture does not only rely on new technologies but also leans on a human organisation. The PLUS initiative have proposed to set up a committee called CAI (Cellule d'Architecture Informatique

# UPDATE ON THE EBS STORAGE RING BEAM DYNAMICS DIGITAL TWIN

S.M. Liuzzo\*, J.-L. Pons†, N. Carmignani, L.R. Carver, L. Hoummi,  
 N. Leclercq, T. Perron, S. White, ESRF, Grenoble, France

## Abstract

The EBS storage ring control system is presently paired with an electron beam dynamics digital twin (the EBS control system simulator, EBSS). The EBSS reproduces many of the beam dynamics related quantities relevant for machine operation. This digital twin is used for the preparation and debug of software to deploy for operation. The EBSS is presently working only for the main storage ring and it is not directly connected to the machine operation but works in parallel and on demand. We present here the steps taken towards an on-line continuous use of the EBSS to monitor the evolution of not directly observable parameters such as beam optics. Keywords: digital twin, storage ring, tango, python

## INTRODUCTION

The ESRF storage ring was upgraded in 2019 to provide brighter X-rays [1]. A very tight schedule was defined for dismantling, installing and commissioning of the new storage ring. In order to cope with this strict schedule, the design and update of the whole software infrastructure had to be performed as early as possible. High-level control applications needed on the first day of commissioning, such as the magnets control, orbit correction, tune correction and beam threading algorithms [2] were tested using the EBS control system Simulator [3].

This control system simulator was strongly focused on the EBS magnets control system that needed to be completely redesigned, but included as well most of the diagnostic devices (beam position monitor (BPMs), tunes, emittances, etc..) needed for the development of the tools used for beam operation. The output values of the simulated diagnostic devices are generated from a given lattice optics model (in pyAT [4]) that is updated upon a magnetic strength or RF frequency variations in the simulated control system - as depicted in Fig. 1. This structure results in an effective beam dynamics digital twin (the EBS control system simulator, EBSS): the actions performed in the simulator are strictly identical to those performed on the real machine and the results are extremely similar or identical. The major limitations are the absence of computation of beam Lifetime, losses, injection efficiency and collective effects [5].

From the user's point of view, the simulator is identical to the production control system. For example the devices providing the information on the beam position have identical names for attributes and properties as their physical counterparts. This is the case also for tune, emittance, RF

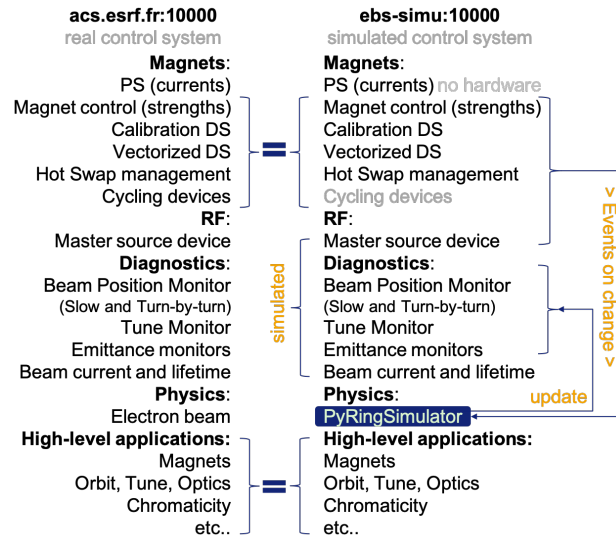


Figure 1: Real control system vs EBS simulator.

and current devices. This allows to transfer the applications from the simulator environment to the real machine by simply changing the environment variable pointing to the TANGO [6] database.

EBSS is a clone of almost all single particle beam dynamics features of the EBS storage ring control system. It allows to interact with a simulated beam and to visualize the expected behaviour of most of the relevant electron beam and lattice observables: orbit, tunes, emittances, optics and coupling.

In the past, the ability to run off-line control room applications was already available, for example via toolkits such as the Matlab-Middle-Layer [7], used in many 3rd generation light sources or specific solutions - like the one on which the Virtual XFEL [8] relies.

## STRUCTURE OF THE SIMULATOR

The core of an EBSS instance [9] is composed of more than ~4000 Tango devices compared to ~25000 in the physical accelerator complex (including the 3 accelerators - linac, booster + storage ring). Each EBSS instance runs on its own Tango control system - the associated Tango database being isolated in a dedicated physical machine. This allows multiple EBSS to run simultaneously and independently on the same host. Presently three EBSS instances are deployed at the ESRF - each one on a 16-CPU/64-Gb host. Initially the different simulators were run on individual Docker containers, but this was more complicated for the users compared to simply dedicating a virtual machine to each EBSS instance.

\* simone.liuzzo@esrf.fr

† pons@esrf.fr

# EVOLUTION OF THE LASER MEGAJOULE TIMING SYSTEM

Somerlinck, T.; CEA CESTA; F-33116 Le Barp; France

Monnier-Bourdin, D.; Hocquet, S.; Greenfield Technology; F-33130 Bègles; France

## Abstract

The Laser MégaJoule (LMJ) timing system has been fully provisioned at the end of 2022. With 15 bundles, 120 beam laser, at the end of 2023, the operational capabilities of the LMJ facility are increasing gradually until it's the full completion by 2025.

The performance of the synchronization equipment on the LMJ is picosecond class. Since 2013/2014, CEA ha continued our studies to improve performance. In 2023, CEA has introduced new features such as a 500 ps rise time at 1 V, variable width and dynamic fine tune delay calibration to improve precision.

Meanwhile, due to electronic obsolescence, a new modified prototype precise delay generator, with "new and old channels", has been tested and compared.

## INTRODUCTION

The LMJ facility is a high-power laser designed to deliver about 1.4 MJ of laser energy to target for high energy density physics experiments, including fusion experiments [1]. This energy is produced by 176 laser beams gathered in quadruplets of 4 beams. Each quadruplet is equipped with an Arbitrary Waveform Generator (AWG) that generates the desired temporal pulse shape (lasting typically 3 ns). Synchronization of LMJ's 176 laser beams is crucial to compress symmetrically the millimeter-size target in order to ignite the deuterium and tritium filled capsule. The most demanding experiences need to synchronize the quadruplets to better than 40 ps rms despite the fact that the quadruplet laser sources are separated within the building by several hundred meters. In addition to laser beams synchronization, the LMJ timing system is in charge to deliver, with the same or lower accuracy, two kinds of signals: fiducials for both temporally marked signals and plasma diagnostics, and triggers signals for manifold devices (sources, amplifiers, Pockels cells, diagnostics...).

The synchronization is therefore one of the most important components for shot experiment, from the laser sources to the target inside the chamber, as shown in Fig. 1.




Figure 1: From laser amplification to experiment in the chamber.

The error budget calculus to reach the 40 ps rms specification has showed that the requirements for the timing systems [2] were, for the most precise delay generator, less than 5 ps rms jitter between 2 outputs and 10 ps p-p drift / 1 month.

## LMJ TIMING SYSTEM

As seen previously, the LMJ requires a lot of timing channels with different accuracies. After choosing to only have 2 levels of timing system precision [3], the CEA timing system team adapt the time line with 3 reference timing triggers, SS0, SS1 and SS2, and 3 reference frequencies, Fig. 2.




Figure 2: LMJ synchronization time line.

Table 1 below summarizes the LMJ synchronization requirements and measurements [4] of the two levels of precision by equipment:

Table 1: LMJ Timing System Requirements

Classes of performance	Jitter rms	Temporal Drift peak to peak			Precision	Range
		24 hours	7 days	1 month		
Standard Precision Delay generator	<100 ps	<200 ps	<500 ps	<1 ns	<±1 ns	1 s
High Specification Measurement	<5 ps 4 ps	<6 ps 4 ps	<10 ps 4 ps	<20 ps 4 ps	<±10 ps	100 μs

## TIMING SYSTEM ARCHITECTURE

The architecture of the LMJ timing system is still unchanged:

- A LMJ master clock, the time reference, delivers an optical clock coupled with triggering data. This oscillator is stabilized with GPS connection, or rubidium oscillator to avoid long term drift.
- A passive optical distribution network sends the optical data clock signal through the whole LMJ facility.
- Two delay generators classes, one for each precision classes, receive the optical information and generate programming delays.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# EPICS INTEGRATION FOR RAPID CONTROL PROTOTYPING HARDWARE FROM SPEEDGOAT

L. Rossa\*, M. Brendike†

Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Berlin, Germany

## Abstract

To exploit the full potential of fourth generation synchrotron sources new beamline instrumentation is increasingly developed with a mechatronics approach. Implementing this raises the need for Rapid Control Prototyping (RCP) and Hardware-In-the-Loop (HIL) simulations. To integrate such RCP and HIL systems into every-day beamline operation we developed an interface from a Speedgoat real-time performance machine - programmable via MATLAB Simulink - to EPICS. The interface was developed to be simple to use and flexible. The Simulink software developer uses dedicated Simulink-blocks to export model information and real-time data into structured UDP Ethernet frames. An EPICS IOC listens to the UDP frames and auto-generates a corresponding database file to fit the data-stream from the Simulink model. The EPICS IOC can run on either a beamline measurement PC or, to keep things spatially close on a mini PC (such as a Raspberry Pi) attached to the Speedgoat machine. An overview of the interface idea, architecture, and implementation; together with a simple example will be presented.

## INTRODUCTION

Recently designed devices for research at synchrotron sources require complex mechanical and mechatronic designs, and therefore need advanced feedback control systems [1–3]. These control systems need to be constructed diligently and are often designed, simulated and tested with specific hardware for Rapid Control Prototyping (RCP) and Hardware-In-the-Loop (HIL) simulations. Suppliers of these hardware components often provide commercial products, or Microsoft Windows compatible, dynamic link libraries to interface their hardware. An open source interface to Linux based operating systems is often not available.

To still be able to use commercial RCP and HIL tools in the Linux based BESSY II beamline control environment another solution than the commercial ones is necessary. Figure 1 illustrates how an alternative solution can look like. The idea is to include an EPICS Input-Output-Controller (IOC) into the RCP and HIL environment and thus connect the System to the EPICS beamline control network.

Like the commercial products, the alternative solution should be also easy to use and to integrate. Therefore the following interface requirements are used:

- the RCP or HIL system should be integrable into a beamline via plug-and-play,

\* rossa@helmholtz-berlin.de

† maxim.brendike@helmholtz-berlin.de

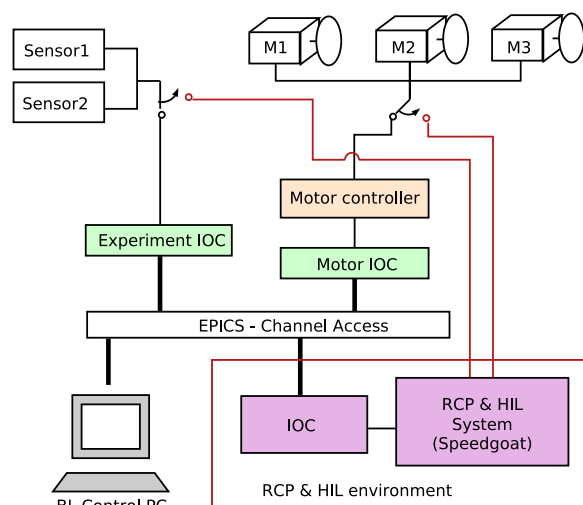


Figure 1: Integration of RCP & HIL System into beamline control environment.

- the developer behind the RCP or HIL system is not an EPICS expert,
- the beamline scientist or user is not an expert in the RCP or HIL architecture,
- the full flexibility of the RCP and HIL system should be maintained,
- the developer should not struggle to keep EPICS in sync with the RCP or HIL system.

As Fig. 1 indicates, in the RCP & HIL System block, hardware from Speedgoat GmbH - further referenced as Speedgoat PC - is used in this paper. However, the developed interface is compatible with hardware from other suppliers as long as MATLAB and Simulink are supported.

## IMPLEMENTATION

There are two possibilities to integrate the EPICS IOC into the RCP & HIL environment. One is to integrate the IOC into the Speedgoat PC and run it directly on the real-time hardware. The second is to run the IOC on separate hardware and communicate to the Speedgoat PC via a dedicated communication interface. Despite the additional hardware requirement we chose the second option. This gives more flexibility in case the Speedgoat PC needs to be replaced and doesn't bind real-time resources for EPICS communication.

The User Datagram Protocol (UDP) is used for communication between the EPICS IOC and Speedgoat PC. This protocol was selected due to its simplicity and flexibility.

# SECoP AND SECoP@HMC – METADATA IN THE SAMPLE ENVIRONMENT COMMUNICATION PROTOCOL\*

K. Kiefer<sup>†</sup>, B. Klemke, L. W. Rossa, P. Wegmann, Helmholtz-Zentrum Berlin, Berlin, Germany  
J. Kotanski, T Kracht, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany  
N. Ekström, A. Pettersson, European Spallation Source, Lund, Sweden  
G. Brandl, E. Faulhaber, A. Zaft, Forschungs-Neutronenquelle Heinz Maier-Leibnitz,  
Garching, Germany  
M. Zolliker, Paul Scherrer Institut, Villigen, Switzerland

## Abstract

The integration of sample environment (SE) equipment in x-ray and neutron experiments is a complex challenge both in the physical world and in the digital world. Different experiment control software offer different interfaces for the connection of SE equipment. Therefore, it is time-consuming to integrate new SE or to share SE equipment between facilities. To tackle this problem, the International Society for Sample Environment (ISSE) developed the Sample Environment Communication Protocol (SECoP) to standardize the communication between instrument control software and SE equipment. SECoP offers, on the one hand, a generalized way to control SE equipment. On the other hand, SECoP holds the possibility to transport SE metadata in a well-defined way. In addition, SECoP provides machine readable self-description of the SE equipment which enables a fully automated integration into the instrument control software and into the processes for data storage. Using SECoP as a common standard for controlling SE equipment and generating SE metadata will save resources and intrinsically give the opportunity to supply standardized and FAIR data compliant SE metadata. It will also supply a well-defined interface for user-provided SE equipment, for equipment shared by different research facilities and for industry. In this article will show how SECoP can help to provide a meaningful and complete set of metadata for SE equipment and we will present SECoP and the SECoP@HMC project supported by the Helmholtz Metadata Collaboration.

## INTRODUCTION

SECoP is a standardised and well documented communication protocol tailored to the needs of SE control. The main challenges for SE control and SE data/metadata collection are the diversity of available SE equipment and the flexibility of SE equipment: in some cases, the equipment is only assembled for a single experiment. In addition, SE equipment, such as magnets or cryostats, is often exchangeable between several beam lines or is provided by external users. All this complicates the collection of interoperable and reusable metadata. SECoP tackles these challenges with a hardware abstraction focussing on the

relevant physical parameters (e.g. temperature, gas flow rate, or magnetic field) and basic functionalities (e.g. reading or driving a parameter). As SECoP is a communication protocol and not a control system, the individual hardware devices with their respective drivers have to be addressed in an underlying programming layer.

## SECoP Abstraction

The abstraction in SECoP and its inherent structure is shown in Fig. 1. The Experiment Control System (ECS) is communicating via SECoP to the “Sample Environment Control Nodes” (SEC nodes). A SEC node provides access to one or several “modules”. A module represents in general a physical parameter, e.g. the hydrostatic pressure at the sample position. A module can have several “parameters” (e.g. the value and status) and “commands” (e.g. stop) associated with it. An individual parameter or command is addressed by the combination of the module name and the parameter or command name (e.g. “pressure1:value” with “pressure1” being here the module name and “value” being in this case the predefined parameter name for the main value of the module). The static information about parameters, modules and SEC nodes is stored in “properties” with predefined names and meanings (e.g. “description”).

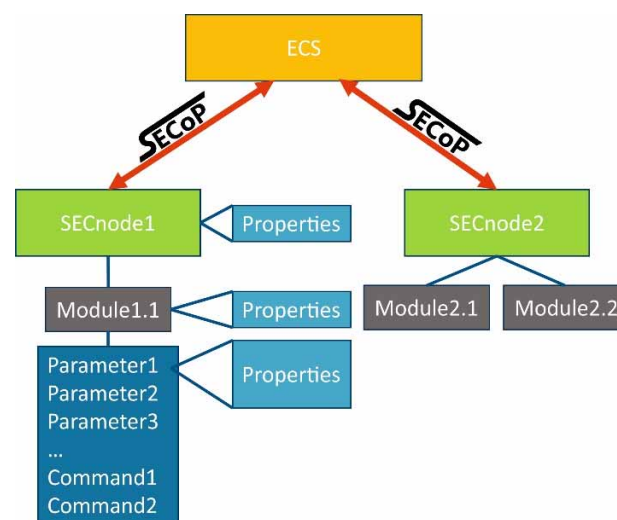


Figure 1: Abstraction and structure in SECoP. The actual SECoP protocol is used for communication between the Experiment Control Software (ECS) and the SEC nodes. For details and examples see text.

\* SECoP@HMC (ZT-I-PF-3-040) was funded by the Initiative and Networking Fund of the Helmholtz Association in the framework of the Helmholtz Metadata Collaboration project call.

<sup>†</sup> klaus.kiefer@helmholtz-berlin.de

# FULL STACK PERFORMANCE OPTIMIZATIONS FOR FAIR OPERATION

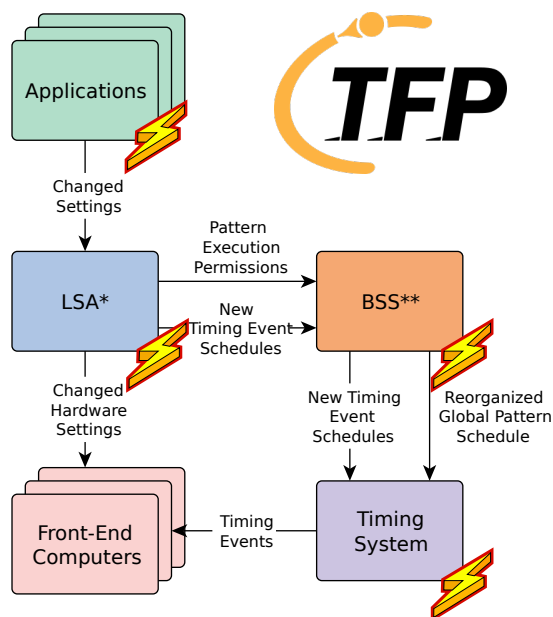
A. Schaller, H. Hüther, R. Müller, A. Walter  
GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany

## Abstract

In the last beam times, operations reported a lack of performance and long waiting times when performing simple changes of the machines' settings. To ensure performant operation of the future Facility for Antiproton and Ion Research (FAIR), the "Task Force Performance" (TFP) was formed in mid-2020, which aimed at optimizing all involved Control System components. Baseline measurements were recorded for different scenarios to compare and evaluate the steps taken by the TFP. These measurements contained data from all underlying systems, from hardware device data supply over network traffic up to user interface applications. Individual groups searched, detected and fixed performance bottlenecks in their components of the Control System stack, and the interfaces between these individual components were inspected as well. The findings are presented here.

## INTRODUCTION

At GSI, the Control System is currently being modernized, also to support the future FAIR accelerators. The new Control System layout is described in Fig. 1.



\* LHC Software Architecture (Settings Management System)  
\*\* Beam Scheduling System

Figure 1: Control System Stack at GSI and FAIR. The parts marked with a lightning bolt were optimized by the Task Force Performance.

The LHC Software Architecture (LSA) [1] internally uses a hierarchical description of physics parameters to hardware parameters [2]. Processing the parameter hierarchy has already been sped up when the first larger contexts were set up

## Hardware

## Control System Infrastructure

at GSI in 2017, for example by parallelizing the calculations done in LSA [3]. After the parallelized calculation was used in production, focus changed to implement new features like the Storage Ring Mode [4]. Since most necessary technical features are available for productive usage now, the Controls department got the mandate by management to prioritize performance improvement in 2020.

## ESTABLISHMENT OF THE TASK FORCE PERFORMANCE

The Control System stack involved in settings and timing changes of the machine should be sped up after operations reported poor performance and severely reduced usability due to long waiting times for trims (setting changes). Therefore, the Task Force Performance (TFP) was established with members from different departments and groups [5]. The project lead was taken by the chief architect of the Controls department, while the individual Controls groups delegated different developers to fully work for the TFP. Operations, as well as the different machine departments, delegated specialists to support the TFP in implementing the various scenarios like machine setup or beam manipulation.

## TIME FRAME AND MILESTONES

The first action was to plan a time frame for when the TFP could take its baseline measurements, and for when the performance optimizations should be done to perform the next beam time without interruptions. The kick-off meeting took place at end of April, while the milestone "Ready for Integration" was planned for the end of October, and "Ready for Production" was scheduled for the end of November 2020. During the first weeks, a measurement concept was established in the Control System software stack to measure the duration of its different functions.

## KICK-OFF

Operators and machine experts provided different scenarios, from simple setting changes to complex machine setup procedures, that should be optimized by the TFP. For these scenarios, baseline measurements under real production conditions were taken in June 2020, directly after the beam time 2019/2020 had ended.

## Goals

The main goal for the TFP was to substantially improve performance for the most relevant use cases, so the facility can be operated efficiently. Identifying potential performance issues that may become relevant for FAIR (i.e. the new synchrotron SIS100 and beyond) was the secondary objective.

# A DATA ACQUISITION MIDDLE LAYER SERVER WITH PYTHON SUPPORT FOR LINAC OPERATION AND EXPERIMENTS MONITORING AND CONTROL

V. Rybnikov\*, A. Sulc, DESY, Hamburg, Germany

## Abstract

This paper presents online anomaly detection on low-level radio frequency (LLRF) cavities running on the FLASH/XFEL DAQ system. The code is run by a DAQ Middle Layer (ML) server, which has online access to all collected data. The ML server executes a Python 3 script that runs a pre-trained machine-learning model on every shot in the FLASH/XFEL machine. We discuss the challenges associated with real-time anomaly detection due to high data rates generated by RF cavities and introduce a DAQ system pipeline and algorithms used for online detection on arbitrary channels in our control system. The system's performance is evaluated using real data from operational RF cavities. We also focus on the DAQ monitor server's features and its implementation.

## DAQ DATA ON-LINE ACCESS

The DOOCS [1] based FLASH DAQ [2] system (Fig. 1) has been in use since 2004. A similar XFEL DAQ [2] system has been running since 2019. The collected data is provided not only for the purpose of the LINACs operation but also for the experiments running on XFEL [3]/FLASH [4] photon beamlines. To be able to implement online (nearly real-time) control of the LINACs the data access is provided for Middle Layer (ML) servers running on the DAQ computers. The data access is done via Buffer Manager (BM) [5]. The BM provides ML servers with the synchronized data for every short in the LINAC.

A DAQ monitor server is a generic ML server that can be configured according to the user's requirements. To configure the server one needs two pieces of information:

- a set of DAQ channels which data is to access,
- a path to a Python script for the data processing.

All this information is provided via dedicated server's DOOCS properties accessible via the network.

## DAQ MONITOR SERVER IMPLEMENTATION

The DAQ monitor server is implemented as a DOOCS server. It runs on a DAQ computer and has access (via BM) to all collected DAQ data. The server can execute a Python [6] script feeding it with the required DAQ channel data.

\* vladimir.rybnikov@desy.de

## Server Design

The server is written in C++ [7]. It makes use of the Python/C API [8] (independent on the Python 3 version) to provide the data exchange between C++ and Python worlds.

## Server Configuration

The server configuration with respect to channels to be read from the DAQ BM is implemented via setting the server's property 'DAQ.REQ'. It accepts XML [9] strings. The XML string consists of a number of sections corresponding to the number of DAQ channels to read. Every section contains at least the following information:

- the DAQ channel name,
- the sub-channel number,
- channel sender source (server block name)
- and an event type mask.

All this information is used to perform the correct channel subscription at the BM.

Any DOOCS server has a configuration file for storing its property values. In addition, every property of XML type creates a corresponding text file containing the XML string content. The DAQ monitor server uses this file to set the value of the 'DAQ.REQ' property during its start-up. It

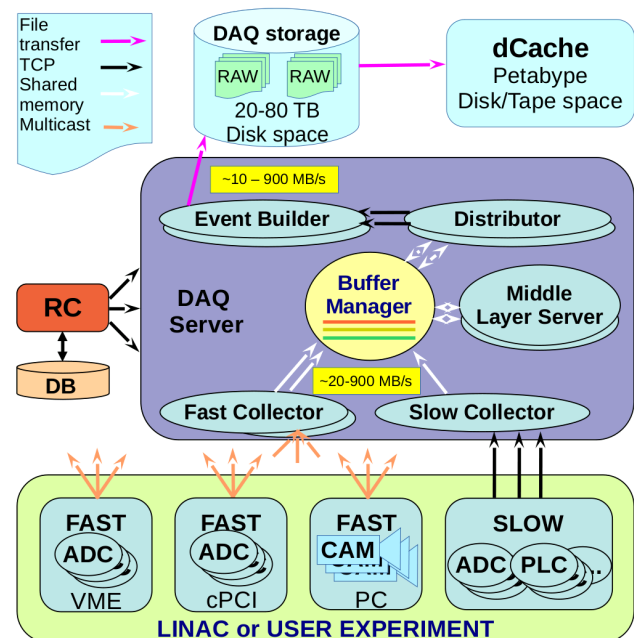


Figure 1: FLASH and XFEL DAQ architecture.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI



## THE TIMING SYSTEM FOR PETRA IV

T. Wilksen<sup>†</sup>, V. Andrei, K. Brede, H. T. Duhme, M. Fenner, U. Hurdelbrink, J. Jaeger, H. Kay, H. Lippek, F. Ludwig, M. Pawelzik, S. Ruzin, H. Schlarb, DESY, Hamburg, Germany

### Abstract

At DESY, the PETRA III synchrotron light source upgrade towards a fourth-generation, low-emittance machine PETRA IV is being pursued. The realisation of the new machine requires a complete redesign of the timing system, as the beam quality and beam control requirements will change significantly. The new timing system must generate and distribute facility-wide precise clocks, trigger signals, trigger events and beam-synchronous information. The design of the main hardware components will be based on the MTCA.4 standard, which has become a well-established platform at DESY and has successfully been in use with DESY FEL's MTCA.4-based timing systems for almost a decade now. This paper presents and discusses the PETRA IV timing system's overall concept and functionality and its hardware components' development status.

### INTRODUCTION

The PETRA IV project [1] involves the replacement of the existing 3rd generation synchrotron radiation source PETRA III with a state-of-the-art ultra-low emittance storage ring. That includes the upgrade of the storage ring infrastructure with a circumference of 2304 m, as well as the redesign of the current pre-accelerator chain and the construction of new beamlines. The 6 GeV storage ring will be operated at an RF frequency of 500 MHz as PETRA III.

The present booster synchrotron DESY II will be replaced by a new one, DESY IV, to meet the requirements of a low emittance beam injected into PETRA IV. While the LINAC II section will be kept, the gun will be upgraded, and the transfer section with the PIA accumulator will be revised. As an alternative to the DESY IV and LINAC II injector chain, developing and constructing a plasma-based injector will also be part of the PETRA IV project.

The front-end electronics for diagnostics and instrumentation for the PETRA IV accelerator will be entirely overhauled and based on the MTCA.4 standard. At the same time, the existing PETRA III control system will be changed over to a DOOCS-based one as used at the DESY FEL accelerators. That effectively leads to the necessity to redesign the entire timing and synchronization system and its controls for the storage ring as well as for the pre-accelerator chain.

### RF SYNCHRONISATION

The essential requirement of the RF synchronisation system is to provide a continuous reference RF signal generated by a unique, stable master oscillator to drive local, low-noise oscillators. The fundamental RF frequency for the PETRA IV storage ring has to be 499.6643 MHz, while the 3rd harmonic RF system requires a reference RF

frequency of 1.4989929 GHz. The DESY IV booster will have the same frequency as PETRA IV, while the PIA accumulator will be operated at 125 MHz, which will be derived from the DESY IV RF reference. The plasma-based injector will have a laser system operated at 2.9979858 GHz, but the injector itself will work at the same frequency as PETRA IV.

The RF synchronisation system will consist of three RF synthesisers in total. One for PETRA IV, one for the booster chain and one for the plasma-based injector. Redundant GPS Rubidium standards will provide a 10 MHz reference to these synthesisers. Each Rubidium standard has a GPS antenna for high reliability and is compensated for environmental impacts and component ageing. In normal operation, the synthesisers are coupled and generate a phase-synchronous frequency of 499.6643 MHz.

The PETRA IV system will supply, in addition, a 1.4989929 GHz reference for the 3rd harmonic system, while the plasma-based injector system will provide an additional 2.9979858 GHz signal for the laser system. The 499.6643 MHz will be adjustable in the range of  $\pm 1.5$  kHz, and the 3rd harmonic frequency in the range of  $\pm 4.5$  kHz. This feature will be used for non-standard operations like the aforementioned dispersion measurement. It requires decoupling the RF synthesisers while the frequency is being swept over this range.

When going back to normal operations, the frequency of the DESY IV system, resp. of the plasma-based injector system, needs to be driven in a phase-continuous way until the PETRA IV frequency is matched again. For this, the synthesiser has to be capable of Direct Digital Synthesis (DDS).

The RF references from each synthesiser provide a direct, high-quality, high-power RF reference signal to the RF systems and the central timing system components for PETRA IV and the injector chains. A direct distribution of the high-quality, low-noise and high-power RF reference signal is foreseen not only for the RF systems but also for dedicated customers, e.g. the multi-bunch feedback system and beamline experiments. This is done by using standard RF cables placed into the PETRA IV tunnel to profit from the temperature-stabilised environment.

The RF cables must be chosen to keep at the minimum the losses along the path in the tunnel. Distribution racks in the experimental supply halls, the DESY IV, PIA and LINAC II supply halls are currently planned, each having a redundant, direct connection to the central system. The distribution racks will be equipped with amplifiers and splitters for further distribution to customers, like beamline hutches and experiments. At those locations where both RF-synchronisation and fibre optical distribution are present, one system can act as a backup of the other one.

<sup>†</sup> Tim.Wilksen@desy.de

# MANAGEMENT OF EPICS IOCs IN A DISTRIBUTED NETWORK ENVIRONMENT USING SALT

E. Blomley\*, J. Gethmann, M. Schuh, A.-S. Müller  
Karlsruhe Institute of Technology, Karlsruhe, Germany  
S. Marsching, aquenos GmbH, Baden-Baden, Germany

## Abstract

An EPICS-based control system typically consists of many individual IOCs, which can be distributed across many computers in a network. Managing hundreds of deployed IOCs, keeping track of where they are running, and providing operators with basic interaction capabilities can easily become a maintenance nightmare.

At the Institute for Beam Physics and Technology (IBPT) of the Karlsruhe Institute of Technology (KIT), we operate separate networks for our accelerators KARA and FLUTE and use the Salt Project to manage the IT infrastructure. Custom Salt states take care of deploying our IOCs across multiple servers directly from the code repositories, integrating them into the host operating system and monitoring infrastructure. In addition, this allows the integration into our GUI in order to enable operators to monitor and control the process for each IOC without requiring any specific knowledge of where and how that IOC is deployed. Therefore, we can maintain and scale to any number of IOCs on any numbers of hosts nearly effortlessly.

This paper presents the design of this system, discusses the tools and overall setup required to make it work, and shows off the integration into our GUI and monitoring systems.

## ENVIRONMENT

The two accelerators KARA (Karlsruhe Research Accelerator) [1] and FLUTE (Far-infrared linac and test experiment) [2] each operate in separate, self-sufficient network environments, including DNS, DHCP, NTP and similar services. Most hosts are virtual machines running Ubuntu LTS versions for the operating system, managed via a Proxmox VE cluster [3]. Dedicated computer hardware is only in use if required. Examples are operator terminals in the control rooms, our data archiving cluster and time synchronization servers.

Most EPICS IOCs run on Ubuntu based virtual machines, as we try to avoid using dedicated hardware, such as VME crates. This requires most hardware being able to communicate via TCP or UDP, with serial communication being managed via serial-to-ethernet hardware gateways. The exception here are IOCs running directly on commercially available hardware, which in most cases is hardware for accelerator beam diagnostics, but recently also includes some power supplies.

For personal or machine protection-critical systems, PLCs are in use, which are also interfaced using EPICS IOCs via

\* edmund.blomely@kit.edu

TCP. For data archiving of EPICS process variables (PVs) a Cassandra [4] NoSQL cluster is in use. The default graphical user interface is Control System Studio (CSS). Table 1 shows an overview over the numbers of server-side components.

Table 1: Control System Components in Numbers

	KARA	FLUTE
Virtual Machines	23	11
Physical Server	8	9
Registered IP Devices	707	341
PLCs	31	11

## EPICS SOURCE CODE

Most EPICS related source code is managed via a site-specific GitLab instance [5]. While most of our EPICS related projects are not publicly visible at the moment, we are working on publishing more of the ones that are not closely tied to our infrastructure. Some custom EPICS modules are already available via our GitHub organisation *KIT-IBPT* [6].

### EPICS Distribution

As we are almost exclusively running Ubuntu Linux inside our accelerator network server infrastructure, we build and distribute our own EPICS Debian Packages (.deb). We are running EPICS 7 and only compile, build and distribute EPICS modules which are used by at least one IOC. The most important modules are asyn, StreamDevice, autosave, motor, areaDetector, s7nodave, MRF, and Open62541. We make use of GitLab Continuous Integration (CI) for applying custom patches, building each component and distributing the resulting Debian packages via an internal APT repository. The EPICS version and the available modules are exactly the same across all EPICS servers.

### IOC Structure

Typically, each IOC is represented by a repository in our GitLab instance. While we make use of the *makeBaseApp.pl* script, we also add an executable to the *iocBoot/iocLinux* folder. By default, the executable is just called *run*.

The IOCs are grouped into accelerator-specific GitLab groups. For devices which are used across both accelerators, such as certain power supplies, a group for shared IOCs exists. To simplify the IOC maintenance tasks in such cases, two different start-up files and two different executables *run\_flute* and *run\_kara* are stored, so that all other parts of the IOC code are shared. A similar approach is taken in

## EQUIPMENT LIFE-CYCLE MANAGEMENT AT EuXFEL

N. Coppola\*, B. Fernandez, N. Jardon, P. Gessler, S. Hauf, S. Huynh, M. Manetti  
European X-Ray Free-Electron Laser Facility GmbH, Holzkoppel 4, 22869 Schenefeld, Germany

### Abstract

Scientific instruments at the European X-Ray Free Electron Laser Facility (EuXFEL) comprises of a large variety of equipment, ranging from controllers, motors and encoders to valves. It is a false assumption that once a specific equipment had been procured and integrated, that no further attention is required. Reality is much more complex and incorporates various stages across the entire equipment life-cycle. This starts from the initial selection, standardization of the equipment, procurement, integration, tracking, spare part management, maintenance, documentation of interventions and repair, replacement and lastly, decommissioning. All aspects of such a life-cycle management are crucial in order to ensure safe and reliable operation across the life time of the equipment, whether it be five years, twenty years, or longer.

At EuXFEL, many aspects of the described life-cycle management are already carried out with dedicated tools. However some aspects rely on manual work, which requires significant effort and discipline.

This contribution aims to provide an overview of the requirements, and the ongoing efforts to develop and establish a complete life-cycle management at the EuXFEL.

Keywords: life-cycle, documentation, requirements, maintenance

### MOTIVATION

Many equipment, controller and controlled entities alike, are chosen and purchased, in scientific facilities without previous life-cycle assessments (see Fig. 1). The reality at many facilities shows that even temporary installations are used for 10 to 30 years or even longer.

Typical scientific and industrial instrumentation are complex and contain many sub-elements. These may be operated efficiently if all sub-components work properly and minimal efforts are needed for maintenance, safety, security and, when parts need to be exchanged or replaced, compatibility exists among hardware models and different versions.

In many cases end-users tend to request to purchase and install, and therefore support, the very same equipment which they already know, due to the fact they have used it in the past, irrespective whether or not they are completely happy of their performances or the added effort needed to introduce these elements into the control system.

This is *in principle* human and natural although not logical and moreover is not making use of technical expertise present within the support groups of the facilities.

There is also a tendency to delegate the design and selection of equipment to external companies without inviting

technical experts, present in the same facility, when making contact with these companies. These, in turn, choose equipment and the models of components optimizing for purchase costs or inserting devices they have in stock instead of investigating which type of controlling entities suit the client best, which would offer a solution with optimal purchase and operation costs simultaneously.

Up to this point in time, we have installed scientific and industrial equipment along 6 km of tunnels and in 7 experiment locations (starting around the year 2012). Some of the equipment has been running ever since, with basically 100% duty cycle.

We aim to:

- improve, develop, locate, track and maintain the devices keeping the down time to a minimum and keep historical knowledge
- develop intervention strategies (possibly w/o interference) to exchange devices near end of life-time
- keep track of integrated duty-time of devices and interventions (whether to repair or modify)
- understand which equipment has become or is going to become obsolete (and prepare for possible replacement with equivalent device(s)).

### METHODS

The European XFEL GmbH facility, in order to strengthen the efficiency of the centralized technical vetting which is

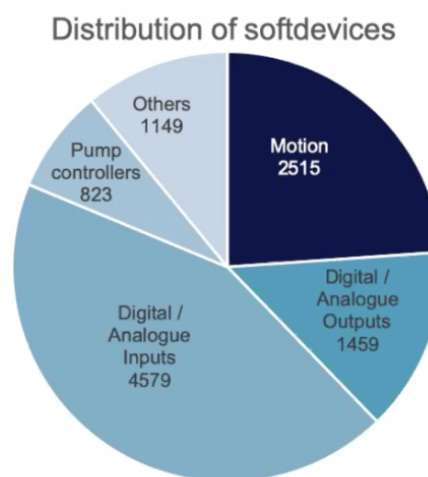


Figure 1: Numbers of devices, subdivided into categories, installed up to know at EuXFEL.

\* nicola.coppola@xfel.eu

# ADAPTABLE CONTROL SYSTEM FOR THE PHOTON BEAMLINES AT THE EUROPEAN XFEL: INTEGRATING NEW DEVICES AND TECHNOLOGIES FOR ADVANCED RESEARCH

B. Rio\*, D. Finze, M. Petrich, H. Sinn, V. Strauch, A. Trapp, R. Villanueva Guerrero, M. Dommach  
European XFEL GmbH, Schenefeld, Germany

## Abstract

The European XFEL is an X-ray free-electron laser (FEL) facility located in Schenefeld, in the vicinity of Hamburg, Germany. With a total length of 3.4 kilometers, the facility provides seven scientific instruments with extremely intense X-ray flashes ranging from the soft to the hard X-ray regime. The dimension of the beam transport and the technologies used to make this X-ray FEL unique have led to the design and buildup of a challenging and adaptable control system based on a Programmable Logic Controller (PLC). Six successful years of user operation, which started in September 2017, have required constant development of the beam transport in order to provide new features and improvements for the scientific community to perform their research activities.

The framework of this contribution is focused on the photon beamline, which starts at the undulator section and guides the X-ray beam to the scientific instruments. In this scope, the control system topology and this adaptability to integrate new devices through the PLC Management System (PLCMS) are described. In 2022, a new distribution mirror was installed in the SASE3 beam transport system to provide photon beams to the seventh and newest scientific instrument, named Soft X-ray Port (SXP). To make the scope of this paper more practical, this new installation is used as an example. The integration in the actual control system of the vacuum devices, optic elements, and interlock definition are described.

## INTRODUCTION

With a total length of 3.4 kilometres, the European free-electron laser (FEL) facility [1], sets itself apart through its scale and technical capabilities. Currently, seven scientific instruments are provided with extremely intense X-ray flashes ranging from the soft to the hard X-ray regime by accelerating electrons up to 17.5 GeV in a superconducting (SC) linear accelerator [2]. The facility's singularity is further underlined by its collaborative aspect, involving twelve shareholder countries in the project's success. It also distinguishes itself in research with a wide range of scientific disciplines, including physics, chemistry, biology, materials science, and earth science. Finally, the facility's uniqueness is underlined by the scale, the multitude of elements, and the architecture of its control system. Based on a Programmable Logic Controller (PLC), it has well over a hundred PLCs with an important variety of devices to control. Moreover, to offer an intuitive user interface, the PLCs are interfaced

with an in-house Supervisory Control and Data Acquisition (SCADA) system named Karabo [3].

From the start of user operation in September 2017 to the present year, when the commissioning of a Soft X-Ray Port (SXP) [4] instrument started, a substantial series of upgrades and enhancements have been implemented. In line with the new research requirements, the adaptability of the control system is required to keep this facility up to date, integrating new devices with cutting-edge technologies.

This paper describes the different elements constituting this adaptable control system and an illustration of device integration.

First, the PLC Management System (PLCMS), a software tool to automatically generate PLC projects, and Karabo are addressed. Then, the control system topology of the photon beamline and the device used as an example, a mirror chamber, are described. The workflow to integrate this new device into the control system is also illustrated. Finally, challenges and considerations for the future perspectives are outlined.

The scope of the paper is limited to the photon beamline, which starts at the undulator section and guides the X-ray beam to the scientific instruments. The control system concept is the same for the scientific instruments, but the topology can differ from the photon beamline control system.

## THE PLCMS

In the context of expanding hardware devices that need to be controlled and the significant use of PLCs at the European XFEL, a tool has been developed to automate the generation of PLC firmware projects [5]. This tool allows a quick generation of the PLC project to accommodate hardware modification. Developing standard components through a PLC framework has been the first main point to allow automatic generation. The framework is developed with Beckhoff TwinCAT [6] and contains different classes, known as Program Organization Units (POUs) in the IEC 61131 standard, and incorporates a wide range of standardised devices within the framework. Each of these standardized devices in the PLC framework is named soft device (SD), and each piece of hardware is linked to a corresponding instance of the SD.

The Electronic and Electrical Engineering (EEE) group at European XFEL has developed the PLC management system (PLCMS) to automate the generation of PLC projects. The PLCMS is developed using the Python 3 programming

\* benoit.rio@xfel.eu

# EVOLUTION OF CONTROL SYSTEM AND PLC INTEGRATION AT THE EUROPEAN XFEL

A. Samadli\*, T. Freyermuth, P. Gessler, G. Giovanetti, S. Hauf, D. Hickin, N. Mashayekh,  
A. Silenzi, European XFEL, Schenefeld, Germany

## Abstract

The Karabo software framework is a pluggable, distributed control system that offers rapid control feedback to meet the complex requirements of the European X-ray Free Electron Laser facility. Programmable Logic Controllers (PLC) using Beckhoff technology are the main hardware control interface system within the Karabo Control System. The communication between Karabo and PLC currently uses an in-house developed TCP/IP protocol using the same port for operational-related communications and self-description (the description of all available devices sent by PLC). While this simplifies the interface, it creates a notable load on the client and lacks certain features, such as a textual description of each command, property names coherent with the rest of the control system as well as state-awareness of available commands and properties. To address these issues and to improve user experience, the new implementation will provide a comprehensive self-description, all delivered via a dedicated TCP port and serialized in a JSON format. A Python Asyncio implementation of the Karabo device responsible for message decoding, dispatching to and from the PLC, and establishing communication with relevant software devices in Karabo incorporates lessons learned from prior design decisions to support new updates and increase developer productivity.

## INTRODUCTION

As one of the world's leading light sources, The European X-ray Free Electron Laser facility is opening up new research opportunities for scientists and industrial users by generating ultrashort X-ray flashes - 27 000 times per second and with a peak brilliance of  $5 \times 10^{33}$  photons / s / mm<sup>2</sup> / mrad<sup>2</sup> / 0.1% bandwidth [1]. Unique characteristics of the facility enable researchers to study tiny structures, ultrafast processes, extreme states, and small objects. Operating such a complex facility without a robust control system is impossible. Karabo [2] is the control system in use at European XFEL: a pluggable, distributed control system that offers rapid control feedback to meet the complex requirements of the facility. The main interface with the existing hardware infrastructure are Programmable Logic Controllers (PLC), which use Beckhoff technology. The primary objective of this project is to enhance the communication between Karabo and these Beckhoff PLCs [3].

\* ayaz.samadli@xfel.eu

## CURRENT COMMUNICATION BETWEEN KARABO AND PLC

The communication between Karabo and PLC currently uses an in-house developed protocol over TCP/IP using the same port for operational-related communications and a self-description of the PLC's configuration and functionality (the description of all devices available on the PLC) [4]. The self-description and operational-related data (e.g., Value updates) are in a custom binary format.

### *BeckhoffCom Device*

The Human-Machine Interface relies on Karabo devices, enabling users to observe and control the connected hardware devices seamlessly. The BeckhoffCom device serves as a transparent interface between Karabo devices and the PLC, and is critical to this communication flow. It distributes the updates from the PLC to the connected Karabo devices and forwards the reconfigurations and commands from the Karabo devices to the PLC. The PLC terminals are structured within what we call *soft devices*, with each *soft device* being associated with a corresponding Karabo device. BeckhoffCom's core design is based on the Model-View-Presenter (MVP) pattern, with the model managing PLC communication (including TCP communication with a view adapter, view model, and TCP view) and the presenter enforcing functionality based on functional requirements. The system defines events and methods using abstracted interfaces, which increases flexibility and testability.

### *Production Environment*

In the existing production environment at EuXFEL, there are approximately 15,000 Beckhoff devices spread across ten different scientific endstations and beam line installations. Within this infrastructure, approximately 1,000,000 Beckhoff properties are exposed to the control system, consisting of parameters such as temperature, velocity, flow, and more. The following figures offer visual insights to provide a clearer picture of the distribution of PLC devices throughout the infrastructure. Figure 1 illustrates the allocation of Beckhoff devices per instrument. The SASE2 (SA2) photon tunnel installations stand out with roughly 2,000 PLC devices, while Laser topics (LA) require fewer than 200 PLC devices.

We examine the SQS (Small Quantum Systems) to obtain a more detailed view of the distribution of various Beckhoff devices within a specific instrument. Figure 2 shows that the combination of Digital Input and Output devices makes up about 40% of the total devices, with MC2 Beckhoff Motors and Analog Inputs contributing a combined total of

# AUTOMATIC CONFIGURATION OF MOTORS AT THE EUROPEAN XFEL

F. Sohn\*, W. Ehsan, G. Giovanetti, D. Goeries, I. Karpics, K. Sukharnikov  
European XFEL GmbH, Schenefeld, Germany

## Abstract

The European XFEL (EuXFEL) scientific facility relies heavily on the SCADA control system Karabo to configure and control a plethora of hardware devices. In this contribution a software solution for automatic configuration of collections of like Karabo devices is presented. Parameter presets for the automatic configuration are stored in a central database. In particular, the tool is used in the configuration of collections of single-axis motors, which is a recurring task at EuXFEL. To facilitate flexible experimental setup, motors are moved within the EuXFEL and reused at various locations in the operation of scientific instruments. A set of parameters has to be configured for each motor controller, depending on the controller and actuator model attached to a given programmable logic controller terminal, and the location of the motor. Since manual configurations are time-consuming and error-prone for large numbers of devices, a database-driven configuration of motor parameters is desirable. The software tool allows to assign and apply stored preset configurations to individual motors. Differences between the online configurations of the motors and the stored configurations are highlighted. Moreover, the software includes a “locking” feature to prevent motor usage after unintentional reconfigurations, which could lead to hardware damage.

## INTRODUCTION

Research at the EuXFEL relies heavily on the use of more than 3000 motors to move components within the various scientific installations. Experimental setups are frequently modified to provide the best suitable infrastructure for specific scientific measurements. For this reason and, furthermore, to save resources, motors are relocated within the EuXFEL and reinstalled in various locations within the experimental setups. Typically, each motor has more than 150 configurable parameters and needs to be configured at each new location.

The necessity to reconfigure the large number of motors and their parameters leads to an extensive time investment from staff and is error-prone, if done manually. Hence, the *Motor Configurator* software tool for the automatic configuration of motors has been developed, which aims to achieve the following main goals:

- minimize time spent by staff to configure motors,
- minimize mistakes due to manual configuration of motors,
- protect against hardware damage due to accidental misconfigurations of motors.

The software tool is based on the SCADA system Karabo [1,2], which is developed at the EuXFEL. Karabo provides a high-performance, reliable, and user-friendly environment to configure and control a plethora of hardware devices, and implement high-level procedures on top of these devices. One distinctive feature of Karabo are so-called scenes, easily configurable, multi-purpose graphical user interfaces that can be shipped with Karabo-based software tools. The Motor Configurator software includes a scene to provide a user-friendly graphical user interface.

## TECHNICAL BACKGROUND

For the integration of a hardware device in Karabo a software *device class* is written, which provides access to hardware features. Within Karabo, each individual hardware device is represented in terms of an *instance* of the respective device class. For each device instance, a *configuration*, i. e. a set of *operational parameters* for the hardware device, is held by the control system.

Figure 1 shows a schematic of the IT infrastructure used to control the motorized stages in the experimental instruments at EuXFEL. Most motors at EuXFEL are connected to *programmable logic controllers (PLCs)*, one motor to one PLC terminal. The PLCs are interfaced with the Karabo SCADA system. Information about the motor hardware is held by the PLCs and forwarded to Karabo. Operational parameters and motion commands are issued within Karabo and routed via the PLCs to the motor hardware.

Within Karabo each PLC terminal receives a unique identifier, the *terminal ID*. An instance of the correct device class for the connected hardware is automatically created after information about the detected hardware has been forwarded from the PLC to Karabo. Importantly, a configuration is assigned to a terminal ID, but cannot be assigned to the connected hardware directly.

## OPERATIONAL REQUIREMENTS FOR EXPERIMENTAL SETUPS

As motors are relocated between and within experimental setups, the assignment of motors to PLC terminals is subject to changes due to motors being

- reassigned to different PLC terminals,
- added to or removed from the setup.

Even if a motor driver of the same type as the previous setup is connected to a PLC terminal, the previous configuration for the PLC terminal in Karabo might not be suitable due to a motion stage with different mechanical and electrical requirements being connected to that driver. Hence, a

\* florian.sohn@xfel.eu

# THE SUPERCONDUCTING UNDULATOR CONTROL SYSTEM FOR THE EUROPEAN XFEL

M. Yakopov<sup>†</sup>, S. Abeghyan, S. Casalbuoni, S. Karabekyan, European XFEL, Schenefeld, Germany  
A. Hobl, A. Sendner, Bilfinger Noell GmbH, Wuerzburg, Germany  
M. Gretenkord, D. Pieper, Beckhoff Automation GmbH, Verl, Germany

## Abstract

The European XFEL development program includes the implementation of an afterburner based on superconducting undulator (SCU) technology for the SASE 2 hard X-ray permanent-magnet undulator (PMU) system. The design and production of the first SCU prototype, called PRE-SerieS prOtotype (S-PRESSO), together with the required control system, are currently underway. The architecture, key parameters, and detailed description of the functionality of the S-PRESSO control system are discussed in this paper.

## PROJECT OVERVIEW

As an initial step to realize the SCU afterburner project [1], the S-PRESSO module [2] will be manufactured by Bilfinger Noell GmbH, characterized, and installed at the European XFEL facility. The project entails the incorporation of six SCU cells into the existing SASE 2 undulator line. They are planned to be installed right after the last PMU cell.

The overview of the future tunnel installation is presented in Fig. 1.

## Superconducting Undulator Cold Mass

Each SCU unit houses two pairs of two-meter-long superconducting undulator coils, with the superconducting coils-based phase shifter (PS) and Helmholtz coils (HH) in between. The whole assembly together with the number of correction and shimming coils constitutes an SCU magnet system, which is incorporated into a mechanical structure (cold mass) whose function is to keep tight mechanical tolerances to assure the quality of the magnetic field. The cold mass is enclosed in the cryostat and cooled down to 4 K by means of the second stage of three cryocoolers. The electron beam chamber (EBC) is cooled by the second stage of another three cryocoolers. The first stage of all six cryocoolers is used to keep the desired temperature of  $\sim 50$  K in the cryostat. Figure 2 shows the schematic view of the S-PRESSO magnet system, while Table 1 compiles the main parameters of S-PRESSO.




Figure 1: General overview of the SCU afterburner components installation in the tunnel including the control racks.

<sup>†</sup> mikhaile.yakopov@xfel.eu

# Voltumna Linux: A CUSTOM DISTRIBUTION FOR (EMBEDDED) SYSTEMS

L. Pivetta\*, A. I. Bogani, G. Scalamera, Elettra Sincrotrone Trieste, Trieste, Italy

## Abstract

In the last years a thorough approach has been adopted to address the ageing and the variability of control system platforms at Elettra Sincrotrone Trieste. The second generation of an in-house built operating system, named Voltumna Linux, which is based on an immutable image approach, is now ready for production, supporting a number of commercial-off-the-shelf embedded systems. Moreover, the same approach is perfectly suitable for rack-mount servers, with large memory support, that often require the inclusion of third party or closed source packages. Being entirely based on Git for revision control, Voltumna Linux brings in a number of advantages, such as reproducibility of the product, ease of upgrading or downgrading complete systems, centralised management and deployment of the user software to name a few.

## INTRODUCTION

In recent years the number of front-end machines used within Elettra control systems increased considerably. Legacy systems, such as VME single board computers based on MC680x0 and PowerPC microprocessors, were joined by x86 systems in standard 19-inch form factor, like rack-mount servers, and smaller form factor such as NUC, UPBoard, Jetway and MinnowBoard. Moreover, also embedded boards based on ARM microprocessors, such as the Beaglebone, and system-on-chip boards based on FPGA, e.g. De10-Nano, Sockit, Dinet and Arria10 have been adopted.

At Elettra, control systems are mostly based on GNU/Linux distributions, introduced over the years, many times featuring hard real-time extensions such as RTAI [1] or Xenomai [2]. Keeping the same approach for new platforms, over the time, would lead to an even more heterogeneous install base, with additional GNU/Linux distributions or versions. Different subsystems make the administration difficult; typical examples are the init-manager, available in SystemV, Upstart or systemd flavours, or the network management stack that can be based on systemd-networkd, NetworkManager etc. Many versions of system libraries would make it difficult to develop the same application on different targets, restricting the replacement of a system with a new one based on a different architecture.

Moreover, together with the new platforms, existing systems have to be maintained much longer than the typical commercial distribution support, which several times is limited to security updates.

Novel technical solutions introduced in FERMI and the prototypes developed for Elettra 2.0 required new control platforms, specifically designed; GNU/Linux is the natural

operating system choice, stated the very good know-how available in house. A common solution to deal with operating system fragmentation is to adopt automation and configuration management tools, such as Ansible or Puppet. Although valid, this approach does not cover all the requirements and cannot support all the legacy platforms in use.

## REQUIREMENTS

Based on the experience integrating and supporting control system platforms and the use foreseen for the incoming installations, the requirements to be satisfied by the operating system and software stack support for the new platforms have been identified:

1. Allow the adoption of specific versions of system components.
2. Allow to integrate third-party software, when source code is available.
3. Provide multiple levels of customization (kernel, drivers, libraries) by patching or revision control.
4. Optimise the operating system for each hardware.
5. Guarantee reproducibility, for both the operating system and the BIOS/firmware of motherboard and adapters.
6. Build, whenever possible, system configurations first.
7. Encourage software reuse making it available from initial releases.
8. Minimise platform dissimilarity, with special attention to the operating system and low level software stack.
9. Provide separate images for development and production systems.
10. Simplify working with new or low performance platforms supporting cross-compiling.

## Yocto / OpenEmbedded

The Yocto/OpenEmbedded Project [3] is an open source collaboration that provides a flexible set of tools to create custom GNU/Linux based systems for embedded hardware, regardless of the architecture. Established in 2010, it involves many hardware manufacturers, including AMD, ARM, Intel, Texas Instruments to name a few, open-source operating system vendors and electronics companies.

Within the boards in use at Elettra, Yocto/OpenEmbedded is supported by Terasic for the Sockit system-on-chip FPGA based board and by Texas Instruments for the Beaglebone.

Therefore Yocto/OpenEmbedded was the natural, more convenient, choice to base a new, custom, GNU/Linux distribution to be developed in-house.

\* lorenzo.pivetta@elettra.eu



# PARTICLE SWARM OPTIMIZATION TECHNIQUES FOR AUTOMATIC BEAM TRANSPORT AT THE LNL SUPERCONDUCTING LINAC ACCELERATORS

M. Montis<sup>†</sup>, L. Bellan, INFN-LNL, Legnaro, Italy

## Abstract

The superconductive quarter wave cavities hadron Linac ALPI is the final acceleration stage at the Legnaro National Laboratories and it is going to be used as re-acceleration line of the radioactive ion beams for the SPES (Selective Production of Exotic Species) project. The Linac was designed in '90s with the available techniques and it was one of the peak technologies of this kind in Europe at those times, controls included. In the last decade, controls related to all the functional systems composing the accelerator have been upgraded to an EPICS-based solution. This upgrade has given us the opportunity to design and test new possible solutions for automatic beam transport. The work described in this paper is based on the experience and results (in terms of time, costs, and manpower) obtained using Particle Swarm Optimization (PSO) techniques for beam transport optimization applied to the ALPI accelerator. Due to the flexibility and robustness of this method, this tool will be extended to other parts of the facility.

## LINAC DESCRIPTION

The Legnaro CW heavy ion accelerator facility is divided into two main sections: the injectors and the superconducting independent cavity Linac, known as ALPI (Acceleratore Lineare Per Ioni) [1]. Stable ion beams achieve a final output energy of about 10 MeV/u (for stable ion beams), accompanied by an average output current of roughly 100 nA. The provided ion species provided can span in a range from protons to 208Pb ions. The entire heavy ion complex is commonly denoted as TAP (Tandem-Alpi-Piave).

ALPI utilizes two injectors for stable ions: a TANDEM-type electrostatic accelerator responsible for light ion acceleration, and the PIAVE superconductive RFQ (Radio Frequency Quadrupole) [2]. The latter exploits the transition section between the normal-conductive and superconductive parts immediately after the source platform. The RFQ achieves an output energy of 587.5 keV/u.

The ALPI Linac is constructed from a series of 20 cryostats, each accommodating four Quarter Wave (QW) Cavities. These cavities necessitate individual tuning at the commencement of each ion-specific run. A pioneering European prototype developed during the 1980s and 1990s, the ALPI Linac incorporated several innovative techniques characteristic of its era. During the design phase, particular emphasis was placed on achieving a superconductive cavity acceleration field of 3 MV/m, facilitated by a 10 mm aperture diameter. In a strategic optimization of spatial resources, the ALPI period was meticulously configured to

cover a transverse focusing triplet and 2 cryostats, collectively housing 8 cavities. In the latter decade, advancements in superconductive cavity technology propelled the accelerating field of the cavities to 1.5 - 2 times the previous value of 3 MV/m.

In this approach, the energy output was significantly increased, although this was achieved at the expense of transmission efficiency, which is caused by 3 main effects: the transverse defocusing force from RF cavities, the longitudinal phase advance that can become unstable in particular conditions, and the steering effect of the QW. To counteract these phenomena (particularly pronounced within the low beta regime), Alternate Phase Focusing technique was implemented, using a  $\pm 20$ -degree synchronous phase. This adjustment led to a reduction in the longitudinal phase advance to below 120 degrees, thereby mitigating the defocusing force and managing the steering and defocusing effects. However, this solution did decrease the longitudinal acceptance of ALPI, which holds significance for the RIB acceleration facility and the overall stability of the dynamics.

## THE TRANSPORT SWARM OPTIMIZATION APPLICATION

In general, possessing a greater acceptance holds considerable significance in maintaining a robust dynamic. With the objective of enhancing the longitudinal acceptance of ALPI, Particle Swarm Optimization algorithm was chosen for Linac fine tuning.

The TAP facility is equipped with diagnostic stations containing beam profiles and Faraday Cups with a 10 mm diameter. However, these initial stations are misaligned and face various issues. Additionally, the cryostats have a significant off-axis alignment, around a few millimeters, which combined with the small apertures, makes the transverse optics optimization troublesome. Standard optics techniques such as the quadrupole shunting for steering have been only partially successful in dealing with the problem, at the price of a time-consuming tuning. To address this and to proceed to a more automatic setting of the accelerator, a dedicated python application called Transport Swarm Optimization (TSO) and based on Particle Swarm Optimization (PSO) algorithm was implemented to maximize beam transport through and automatic setting of the transverse optics.

### Particle Swarm Optimization Algorithm

PSO is a computational technique inspired by the behavior of birds or fish flocking together [3]. In PSO, a group of potential solutions (called particles) collaboratively searches for the best solution to a problem. Each particle

<sup>†</sup> maurizio.montis@lnl.infn.it

# ALPI-PIAVE BEAM TRANSPORT CONTROL SYSTEM UPGRADE AT LEGNARO NATIONAL LABORATORIES

M. Montis<sup>†</sup>, M. Giacchini, F. Gelain, INFN-LNL, Legnaro, Italy

## Abstract

During the last decade, the control system employed for ALPI and PIAVE Accelerators was upgraded to the new EPICS-based framework as part of the new standards adopted in the SPES project in construction in Legnaro. The actual control for beam transport was fully completed in 2015 and it has been in production since that year. Due to the power supply upgrade and to optimize costs and maintenance time, the original controllers based on industrial PCs were substituted with dedicated serial-over-ethernet devices and Virtual Machines (VMs). In this work we will describe the solution designed and implemented for ALPI-PIAVE accelerators.

## HARDWARE STATUS AND UPGRADE

In the original control system installation, the entire ALPI and PIAVE Accelerators were controlled by INTEL industrial PCs placed along the lines (Fig. 1) and controlled through dedicated applications based on EPICS control system framework [1-3].

Each industrial PC provided up to 9 serial interfaces to magnet power supplies. This solution used Power Over Ethernet (POE) technology as the power source to minimize the number of cables.

Magnet power supply control cards were controlled through RS232 serial communication protocol and the EPICS control application run directly inside the device: different serial communications were established simultaneously during IOC start-up and are maintained during program execution. In addition to that, other devices such as NRM and gauss meter were added to the line and remotely controlled using the industrial PCs and RS232 communication protocol.

Over the years it has been possible to observe the software part was extremely robust: only a few bugs fixed were required to optimize communications with an older version of power supply control cards and naming convention upgrade for process variables; at the same time, we discovered the hardware part was not as robust as originally designed: several hardware faults focused on POE cards and motherboards required extraordinary maintenances. The instability of these components and their availability on market brought us to force an upgrade for the entire system.

To maintain the backbone infrastructure in terms of cabling and communication protocol as much as possible, the following criteria were followed during design and implementation:

- Power Supply controls must be kept with RS-232 serial communication

- Most of the communication wires must be reused with the new configuration
- Controllers must be moved to the control system hypervisor infrastructure, using a dedicated VM



Figure 1: INTEL industrial PC used as lens controller.

- Serial communications require a dedicated serial-to-ethernet converter to guarantee the proper data exchange between VM and magnets power supplies
- The number of serial-to-ethernet converters must be optimized considering the number of devices to control and their topology

## Hardware Upgrade

According to these requirements, the following solutions and tools were used:

- To minimize the number of serial-to-ethernet converters and based on the experience matured in similar challenges, the industrial PCs were substituted with DeviceMaster systems (Fig. 2). This kind of device can be considered a server with device networking capabilities: it provides serial communications with RS-232/422/485 protocols and it supports native COM, TTY, or TCP/IP socket communications. As mentioned in the previous paragraph, all the ports available in each device have been set to work with RS-232 serial protocol.



Figure 2: DeviceMaster (serial-to-ethernet converter) used for beam transport control upgrade.

<sup>†</sup> maurizio.montis@lnl.infn.it

# ESS DRIFT TUBE LINAC CONTROL SYSTEM COMMISSIONING: RESULTS AND LESSONS LEARNED

M. Montis<sup>†</sup>, M. Giacchini, A. Baldo, L. Antoniazzi, INFN-LNL, Legnaro, Italy  
 A. Rizzo, ESS, Lund, Sweden

## Abstract

European Spallation Source (ESS) will be a neutron source using proton beam Linac of expected 5MW beam power. Designed and implemented by INFN-LNL, the Drift Tube Linac (DTL) control system is based on EPICS framework as indicated by the Project Requirements. This document aims to describe the results of the first part of the control system commissioning stage in 2022, where INFN and ESS teams were involved in the final tests on site. This phase was the first step toward a complete deployment of the control system, where the installation was composed by three sequential stages, according to the apparatus commissioning schedule. In this scenario, the firsts Site Acceptance Test (SAT) and Site Integrated Test (SIT) were crucial, and their results were the milestones for the other stages: the lessons learned can be important to speed up the future integration, calibration, and tuning of such a complex control system.

## DTL CONTROL SYSTEM ARCHITECTURE

Given the nature of the project and the involvement of multiple individuals across different levels, the design and implementation of the DTL Control System must adhere to meticulous strategies and solutions, aiming to optimize both costs and time throughout the installation campaign in Sweden. For control systems, the ESS [1] project guidelines indicate EPICS [2] as the standard for this topic.

The scope of the DTL CS is to provide the required software and hardware layers to operate the apparatus. It's important to note that not all functional sub-systems within the DTL fall under the purview of the INFN-DTL CS Group. Therefore, the control system architecture presented covered only a portion of these subsystems and all the remaining ones are designed, and implemented by ESS [3].

Figure 1 represents the schematic related to the control system architecture deployed for the DTL apparatus and it follows the standard 3-layer structure where, at the lowermost layer, encompasses all functional sub-systems of the DTL and it defines the context from which the input/output (I/O) signals originate. The intermediary layer delineates the array of controllers employed for executing the necessary logic and automation within the application, spanning both hardware and software realms. Within this layer, all EPICS Input/Output Controllers (IOCs) are tasked with running both the low-level interface applications and the high-level state machines (Control System Core). The

uppermost layer includes the suite of services offered by ESS-ERIC to facilitate routine operations of the Linac, such as Human-Machine Interface, Archiver service, and Alarms management system.




Figure 1: DTL 3-layer control system schematic.

Different functional sub-systems required different solutions for their implementation based on the design requirements and the possible technologies approved by the project. Where possible, common technologies among the systems were used to optimize development, maintenance, and costs. This approach defined the software architecture developed and the relative control system topology.

The principal technologies adopted are summarized in Table 1, while the primary details in terms of EPICS parameters are indicated in Table 2.

Table 1: DTL Functional Systems: HW, SW and Protocols

Sub-System	Technologies (HW, SW, Protocols)
High Level Control (CS Core)	<ul style="list-style-type: none"> <li>• EPICS framework</li> <li>• No dedicated hardware</li> </ul>
Tubes Thermos Sensors	<ul style="list-style-type: none"> <li>• EPICS framework</li> <li>• Beckhoff hardware</li> <li>• EtherCAT protocol</li> </ul>
Water Cooling Controls	<ul style="list-style-type: none"> <li>• Siemens PLC low-level logic and EPICS integration</li> </ul>
Tuning Motor System	<ul style="list-style-type: none"> <li>• EPICS framework</li> <li>• Beckhoff hardware</li> <li>• EtherCAT protocol</li> </ul>
Steerer System	<ul style="list-style-type: none"> <li>• EPICS framework</li> <li>• Hardware provided by the tender</li> <li>• Serial and TCP-IP communication</li> </ul>

<sup>†</sup> maurizio.montis@lnl.infn.it

# DEVELOPMENT OF THE BEAM GATE SYSTEM USING THE WHITE RABBIT AT SUPERKEKB

F. Ito<sup>†</sup>, H. Kaji, High Energy Accelerator Research Organization, Tsukuba, Japan  
Y. Iitsuka, East Japan Institute of Technology Co. Ltd., Ibaraki, Japan

## Abstract

In the SuperKEK, the beam gate is a system that turns the beam supply from the injector to the main ring on and off. Since the beam gate signal is shared with the abort system, the signal is delivered asynchronously with the operation frequency. The beam gate system, on the other hand, must start/stop each component of the accelerator working at the time of injection in a defined sequence. For reliable operation, we currently allow one blank firing to the kickers and others. A beam-gate system with WRs can stop blank firing.

## INTRODUCTION

The e-/e+ SuperKEKB collider continues to update the world's highest luminosity [1]. Figure 1 shows a schematic of the SuperKEKB accelerator, which consists of two main rings for e-/e+ and the injector. To increase the integrated luminosity, the top-up filling operation is carried out. To keep the beam current, the beam supply from the injector frequently switched between “on” and “off”. The trigger signals to each component are delivered with a precisely defined time delay from the operating frequency using the event timing system [2]. Injection is started and stopped by controlling the trigger signals to each component (electron gun, injection/excursion kicker, septum magnets, etc.) that work during beam injection by the beam gate system.




Figure 1: Schematic of the SuperKEKB accelerator.

The photo cathode RF gun used at the injector is turned on and off by opening and closing the physical shutter for the intense laser beam. To avoid laser irradiation during the shutter operation, which takes about 20 ms, the beam gate signal is input to the pulse delay module together with the trigger for laser irradiation, and the trigger signal is sent to the shutter after waiting for one laser irradiation. The beam gate signals are delivered asynchronously with the operating frequency. Therefore, as shown in Fig. 2, the time width

between the beam gate signal and the output of the trigger for the shutter varies depending on the timing of the beam gate signal. This makes it difficult to turn on/off the kicker, septum, etc. at the correct timing.

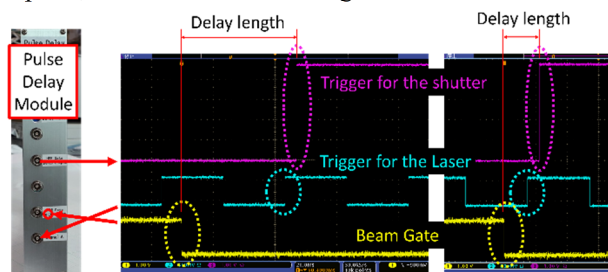


Figure 2: Waveforms of triggers for the shutter, triggers for the laser, and beam gate signals.

## TRIGGER TRANSFER USING THE WR

White Rabbit (WR) is a technology with properties such as sub-nanosecond accurate time synchronization between modules separated by kilometers, scalability, and Gigabit data communication [3]. This technology information is freely available in the Open Hardware Repository (OHWR) [4] in accordance with the CERN Open Hardware License. In this work, we used the Simple PCIexpress Carrier (SPEC), an FPGA Mezzanine Card (FMC) carrier board, and an FMC-DIO card. These were plugged into the PCIe slot of a PC (Ubuntu 20.04LTS) and set up as WR nodes.

KEK uses EPICS for the control system of the accelerator. The development to control the SPEC and FMC-DIO with EPICS has already been done at KEK and has been operated at SuperKEKB [5]. The WR nodes set up this time can send and receive time-stamped trigger events via the WR network using EPICS software (3.15.8). These WR nodes were installed in the laser hut and in the D8 Experiment Power Station (D8) where the event receivers (STD-EVE) that distribute the triggers to the kicker and others are located, respectively. FMC-DIO has 5-channel TTL signal input/output connectors and can record input/output time stamps. When the EPICS software on the WR node detects the rising or falling edge of the signal, it sends a time stamp ( $T_0$ ) when the signal change is detected, the set delay time ( $T_d$ ), and the signal status to the programmed destination node. When the destination WR node receives the trigger information, it outputs a signal at  $T_0 + T_d$ . If its own clock has already passed  $T_0 + T_d$ , it outputs the signal immediately.

## Time Synchronization

If there is an error  $\Delta T$  in time between the sender and the receiver node, the actual time at which the signal is output

<sup>†</sup> email address: fumiaki.ito@kek.jp

# INTRODUCTION OF ETHERNET-BASED FIELD NETWORKS TO INTER-DEVICE COMMUNICATION FOR RIBF CONTROL SYSTEM

A. Uchiyama<sup>†</sup>, M. Komiyama, N. Fukunishi  
RIKEN Nishina Center, Wako, Japan

## Abstract

Internet Protocol (IP) networks are widely used to remotely control measurement instruments and controllers. In addition to proprietary protocols, common commands such as the standard commands for programmable instruments (SCPI) are used by manufacturers of measuring instruments. Many IP-network-based devices have been used in RIBF control systems constructed using the experimental physics and industrial control system (EPICS); these are commercial devices designed and developed independently. EPICS input/output controllers (IOCs) usually establish socket communications to send commands to IP-network-based devices. However, in the RIBF control system, reconnection between the EPICS IOC and the device is often not established after the loss of socket communication due to an unexpected power failure of the device or network switch. In this case, it is often difficult to determine whether the socket connection to the EPICS IOC is broken even after checking the communication by ping-pong. Using Ethernet as the field network in the physical layer between the device and EPICS IOC can solve these problems. Therefore, we are considering the introduction of field networks such as EtherCAT and Ethernet/IP, which use Ethernet in the physical layer. In the implementation of the prototype system, EPICS IOCs and devices are connected via EtherCAT and Soft PLCs are run on the machine running EPICS IOCs for sequence control.

## INTRODUCTION

The control system of the RIBF accelerator facility is based on the experimental physics and industrial control system (EPICS). When it was first operated in a facility in 2002, the EPICS input/output controller (IOC) installed in the VME CPU used the VME bus as an interface between devices and was accessed via the device driver [1]. Later, with the introduction of the EPICS R3.14 series, bus connections and IP protocols were used for inter-device communication between the devices and the EPICS IOC [2]. Currently, the RIBF control system uses N-DIM, an internet protocol (IP) network-based device developed in-house, and PLCs from MELSEC, OMRON, FA-M3, and other companies that are connected via Ethernet and operated with device support using NetDev [3] and AsynDriver [4, 5]. Measurement devices such as digital multimeters are utilized to monitor the beam currents. In many cases, measurement instruments use ASCII-based communication commands called standard commands for programmable instruments (SCPIs) as the presentation layer protocol. They are operated using a Stream Device [6] with EPICS device support.

<sup>†</sup> a-uchi@riken.jp.

Generally, Ethernet, which is a physical layer and data link layer in the OSI reference model, is very convenient and easy to handle, and is currently the standard equipment in devices from many manufacturers. These devices and IOCs typically communicate through asynchronous socket connections. However, after a power failure, a reboot owing to device failure, or an unexpected power loss of a switch, the socket connection is often not restored, and the EPICS IOC must be restarted for reconnection, interrupting the smooth operation of the accelerator. In addition, the RIBF upgrade plan [7] is expected to support applications that require a fast response time, which is not possible with IP networks, while maintaining the convenience of Ethernet as a protocol. The introduction of Ethernet-based field networks to RIBF control system addresses these concerns.

## COMPARISON OF FIELD NETWORKS

Network switches for the control system are typically placed in computer, experimental, and accelerator rooms. An advantage of IP-network-based devices is that they require minimal installation of Ethernet cables. Conversely, they may be less reliable owing to their low real-time performance and slow I/O communication. However, developing a new proprietary dedicated protocol, such as the NIO [8] used in the RIBF control system, would be expensive. Therefore, we considered Ethernet-based field networks and general-purpose protocols, such as EtherCAT [9], FL-net [10], and EtherNet/IP [11].

Among these protocols, the FL-net-based interlock system has successfully been constructed for the injection of the RIKEN ring cyclotron [12]. However, FL-net requires a dedicated network and cannot be mixed with the control network, thereby reducing its convenience. EtherNet/IP is highly convenient because it does not require a dedicated network, can be connected directly to a regular IP network switch, and requires minimal Ethernet cabling.

In addition to FL-net, EtherCAT requires a dedicated network, and because it is not possible to mix TCP/IP with the network, a new cable must be installed. However, because EtherCAT master-slave communication is faster than that of general-purpose protocols, a faster response can be achieved than with EtherNet/IP and FL-net; it also enables high real-time control.

In the RIBF upgrade plan [7], accelerators have become more sophisticated, the beam intensity has increased, and high-speed interlock signal output and highly accurate synchronization control have become necessary. Therefore, we have developed a prototype system using EtherCAT. We have also constructed an EtherNet/IP prototype system which can be mixed with IP network-based switches, minimizing wiring work and providing high convenience.

# MULTI-USER VIRTUAL ACCELERATOR AT HEPS FOR HIGH-LEVEL APPLICATION DEVELOPMENT AND BEAM COMMISSIONING\*

Xiaohan Lu<sup>†</sup>, Yi Jiao, Jingyi Li, Nan Li, Cai Meng, Yuemei Peng, Peng Zhu, Gang Xu  
Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China

## Abstract

To meet the beam commissioning requirements of HEPS, a new framework named PYthon-based Accelerator Physics Application Set (Pyapas) was developed for building high level applications (HLAs) for HEPS. Based on *Pyapas*, the application development for Linac was completed in June 2022. To test the HLAs before putting them online, a multi-user virtual accelerator based on *Pyapas* was developed. It provides a virtual environment of the real machine, including virtual devices, beam parameters, errors and so on. The HLAs tested on the MUVA can be seamless applied to the real beam commissioning. This paper briefly introduces the MUVA system.

## INTRODUCTION

The High Energy Photon Source (HEPS) is a 6 GeV, 1.3 km, fourth-generation light source [1-3]. It began construction in Beijing, China in mid-2019 and is expected to become one of the world's smallest emittance light sources after completion. In order to meet the beam commissioning needs of HEPS, a new framework *Pyapas* was designed for developing the high-level applications (HLAs) [4]. To test the HLAs before putting them online, a virtual accelerator is necessary.

The virtual accelerator system not only provides a virtual hardware environment, but also provides real-time information on various beam parameters under different fault conditions. By performing beam commissioning simulations on a virtual accelerator before building the machine, it is possible to identify key factors that could potentially affect the performance of the machine under construction. Based on these findings, appropriate countermeasures can be formulated, greatly facilitating the timely completion of design goals. At the same time, the development of HLAs relies heavily on the support of virtual accelerators. A comprehensive virtual accelerator system can be used to test the usability of the application and the accuracy of the calculation and optimization results.

At HEPS, a multi-user virtual accelerator system has been developed for testing the HLAs and simulating the effects of various errors on the results of beam commissioning. The virtual accelerator is based on the *Pyapas* development framework for HLA and is designed using a client/server (C/S) architecture. It uses Ocelot [5] with custom multipole field models for physical calculations and supports error simulation for various magnet and beam instrumentation and diagnostics devices. Calculation results are sent externally through the EPICS PV channel. The multi-user virtual accelerator system was developed to

meet the needs of different users within the same network area who need to simultaneously call the virtual accelerator for software debugging and simulation research. Each user can open a unique virtual accelerator without affecting others, and can also start different virtual accelerators for different research content. The number of virtual accelerators opened is not limited. The operation of the entire virtual accelerator system can be easily switched on and off like opening an app, greatly facilitating user use. As shown in Fig.1 in the overall control system the virtual accelerator and the low-level control system are at the same level. The HLAs are capable of connecting to either the low-level control system or the virtual accelerator

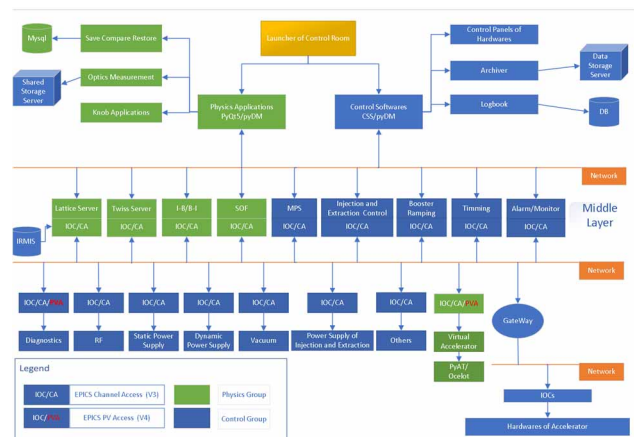


Figure 1: The control system architecture diagram of HEPS.

With the help of MUVA we completed the development of HLAs for Linac [6] and successfully applied to beam commissioning [7, 8]. This paper provides a detailed description of the design concept and implementation of the MUVA system.

## FRAMEWORK DESIGN OF MUVA

The MUVA comprises of three critical components. Firstly, a set of accurate physical computation models that can effectively and promptly compute the necessary beam parameter data based on different hardware system parameters. We use Ocelot as the basic physical calculation models and develop some new models to simulate combined magnets.

Secondly, a comprehensive virtual hardware system is designed to simulate the control and interaction scenarios with the actual hardware system. The control system of HEPS are built with EPICS, thus the key component is to setup EPICS run-time database dynamically. With *Pyapas* framework, there is a lattice file to describe the real ma-

\* Work supported by NSFC12005239

<sup>†</sup> luxh@ihep.ac.cn

# THE APPLICATION OF PYAPAS IN LINAC BEAM COMMISSIONING AT HEPS\*

Xiaohan Lu<sup>†</sup>, Yaliang Zhao, Hongfei Ji, Yi Jiao, Jingyi Li, Nan Li, Cai Meng, Yuemei Peng, Peng Zhu, Gang Xu  
Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China

## Abstract

The High Energy Photon Source (HEPS) is a 6 GeV, 1.3 km, 4th generation storage ring light source being built in Beijing, China. The development of high-level applications (HLAs) for HEPS started in early 2021. A new framework named Python-based Accelerator Physics Application Set (Pyapas) was developed for building HLAs. Based on *Pyapas*, the application development for Linac was completed in June 2022. And then the joint test with hardware system was performed, all the applications worked well in the Linac control room. Beam commissioning for the Linac began in March 9 of this year, and all the HLAs for the Linac are functioning well. This paper will present the application of *Pyapas* in linac beam commissioning.

## INTRODUCTION

The High Energy Photon Source (HEPS) is a 6 GeV, 1.3 km, fourth-generation light source [1]. It began construction in Beijing, China in mid-2019 and is expected to become one of the world's smallest emittance light sources after completion. The Linac of HEPS began beam commissioning in March 9 of this year. In order to meet the beam commissioning needs of HEPS, we began developing high-level applications (HLAs) in early 2021 and a new framework was designed for the development of HLAs [2].

As a fourth-generation light source, HEPS adopts a compact multi-bend achromat(MBA) lattice design for the storage ring [3, 4]. With MBA lattice the number of magnets in fourth-generation light sources has increased by an order of magnitude compared to existing third generation light sources. This means that the variables to be controlled have increased by one or two orders of magnitude. The error tolerances of the fourth-generation light sources are also tighter due to the ultralow emittance and stronger magnetic fields. Therefore, higher control precision and faster response times need to be considered in the HLA development. To address these issues, we have designed a new HLA framework named *Pyapas* [5], as shown in Fig. 1. It adopts a modular design philosophy and increases overall scalability. A dual-layer physical model module has been designed to meet the replaceability of online calculation models. In addition, the communication module, database module, and server module have all been specially designed to meet the needs of adjusting a large amount of parameters.

Based on *Pyapas*, we have completed the development of HLAs for the Linac [6] and successfully applied them to

beam commissioning [7], verifying the practicality and reliability of the applications. This paper will briefly introduce the design of *Pyapas* and the progress of the development of HLAs for HEPS.

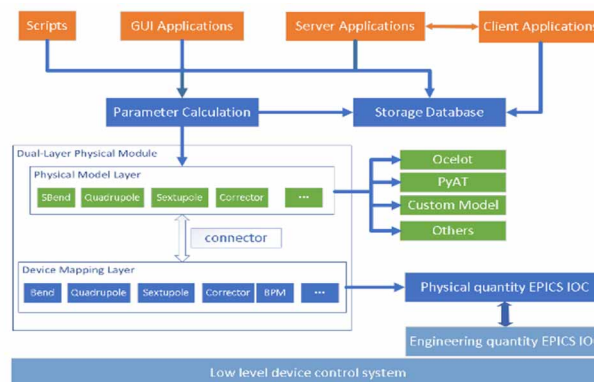


Figure 1: The complete workflow of new framework *Pyapas*.

## HLAS DEVELOPMENT FOR LINAC

To meet the beam commissioning needs of the Linac, we developed corresponding HLAs based on *Pyapas*, including orbit correction, emittance measurement, energy and energy spread measurement, phase scan, beam based alignment (BBA) and physics-based linac control application [8].

As shown in Fig. 2, the physics-based control program allows beam commissioning operators to directly adjust physical quantities such as magnetic fields, angles, and energies, enabling them to quickly find the appropriate working mode. The energy and energy spread measurement systems locate in the two energy analysis stations. AM1 and AM2 are dipole magnets which creates dispersion for measuring beam energy and energy spread. AM1 deflect beam in horizontal direction, while AM2 in vertical. According to the field of dipole magnet and the beam position in the profile monitor, the beam energy can be obtained. Then the beam energy spread can be calculated with the beam size measured by the profile monitor, as shown in Fig. 3.

For the orbit correction, one-to-one feedback correction method and global correction method based on response matrix is used. For BBA, Two methods are used to measure the BPM offset for the HEPS Linac, linear fitting method and parabolic fitting method [8], the application shows in Fig. 4. Transverse emittance is an important parameter of characterizing accelerator performance. For the main Linac, dispersion is negligible and the beam size mainly

\* Work supported by NSFC12005239

<sup>†</sup> luxh@ihep.ac.cn

# RESEARCH ON HALF HISTORICAL DATA ARCHIVER TECHNOLOGY

Xiaokang Sun\*, DaDi Zhang, Huhang Chen  
University of Science and Technology of China, HeFei, China

## Abstract

The Hefei Advanced Light Facility (HALF) is a 2.2GeV fourth-generation synchrotron radiation light source. The HALF control system is a distributed control system based on Experimental Physics and Industrial Control System (EPICS). As the essential part of the HALF control system, the Historical Data Archiving System (HDAS) is responsible to store operational data for the entire facility including the accelerator and beamlines, and provides the functions for data query and analysis. According to the estimation based on the HALF scale, approximately 25,000 EPICS PVs will be stored in HDAS, and these accumulated massive data require a dedicated database for persistent storage and management. Under the EPICS PV data scenario of HALF, a fair database test platform is designed and built to test the read-write performance of databases commonly used in the particle accelerator field. The tested objects include EPICS Archiver Appliance and the five databases MongoDB, HBase, InfluxDB, TimescaleDB, and Cassandra. The test results indicate that TimescaleDB has the fastest read performance, and  $1.4 \times 10^6$  items of data can be read per second. In the future, a TimescaleDB distributed cluster will be designed and deployed, and an HDAS prototype system will be developed based on this cluster.

## INTRODUCTION

The Hefei Advanced Light Facility (HALF) is a 2.2GeV fourth-generation synchrotron radiation light source, which is scheduled to start construction in Hefei, China in 2023. The HALF consists of an injector, a 480 meters diffraction-limited storage ring, and ten beamlines for phase I [1, 2]. The HALF control system is a distributed control system based on Experimental Physics and Industrial Control System (EPICS). As the essential part of the HALF control system, the Historical Data Archiving System (HDAS) is responsible to store operational data for the entire facility including the accelerator and beamlines, and provides the functions for data query and analysis.

The HDAS archives the EPICS PV data from the control system. PV is a structured kind of data, which includes PV ID, PV value, timestamp, severity, and other metadata. According to the estimation based on the HALF scale, the HALF control system will generate approximately 250,000 PVs in total. According to the some accelerator project cases [3-5], in order to simplify the research metrics of archived PVs, the archived PVs percentage of the HALF control system is chosen as 10%. Approximately 25,000 PVs will be stored in HDAS, which will generate tens of TB of data per year. These massive amounts of data generally require a

dedicated database for persistent storage and management. The architecture of the HDAS is shown in Fig. 1, the data archiving engine is responsible for collecting PV from the IOCs through the CA/PVA protocol, and storing the data into the database cluster. Based on the database cluster, the Web applications for data querying and analysis are developed. The database performance directly influences the speed of data processing and analysis. Therefore, it is necessary to choose a high-performance database for HDAS.

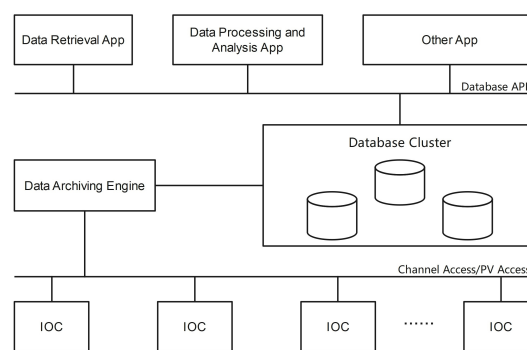


Figure 1: Architecture of the HDAS.

In order to choose a high-performance database for HDAS, in the EPICS data archiving scenario, an environment-fair test platform is designed and built to test the read-write performance of databases commonly used in the accelerator field. Additionally, research on the database software ecosystem is conducted for the development of database applications.

Different databases are adopted by the particle accelerator facilities around the world, the popular and commonly used databases are as follows: MongoDB, HBase, InfluxDB, TimescaleDB, and Cassandra. Besides, the EPICS Archiver Appliance (AA) as a dedicated EPICS data archiving tool is widely used in many particle accelerator facilities [3]. The AA is developed by a collaboration of SLAC, BNL, and MSU in 2015. The AA stores data in the form of Google Protocol Buffer files, and the multiple stages mechanism is adopted to achieve high data retrieval performance. The Beijing Electron Positron Collider II (BEPC-II) proposed a data archiving system based on MongoDB [6]. MongoDB is a database written in C++, based on distributed document storage [7]. The Hefei Light Source (HLS-II), the China ADS Front-end Demo Linac (CAFe), and the Japanese Proton Accelerator Research Complex (J-PARC) all use the HBase database for storing historical operation data [8, 9]. HBase is a column-oriented non-relational database that runs on top of the Hadoop Distributed File System (HDFS) [10]. The accel-

\* sunxk@ustc.edu.cn

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI



# THE ALARM SYSTEM AT HLS-II

Shuang Xu, Xiaokang Sun\*  
NSRL, USTC, Hefei, China

## Abstract

The control system of the Hefei Light Source II (HLS-II) is a distributed system based on Experimental Physics and Industrial Control System. The alarm system of HLS-II is responsible for monitoring the alarm state of the facility and distributing the alarm in time. The monitoring scope of the alarm system covers the front end devices and the server systems of HLS-II. The alarm distribution strategy of HLS-II is designed to overcome nuisance alarms. Zabbix is an open-source software tool used for monitor the server systems. Custom metrics are collected through external scripts. The alarm system of HLS-II provides multiple ways to notify the responsible operators, including WeChat, SMS and web-based GUI. It facilitates the operator to troubleshoot problem efficiently, so as to improve the availability of HLS-II.

## INTRODUCTION

Hefei Light Source II (HLS-II) is a vacuum ultraviolet synchrotron light source [1]. HLS-II consists of e-gun, microwave, klystron, power supply, vacuum, resonator, undulator, beam diagnosis and other front end devices. Alarm system of HLS-II is inevitable component of control system that notify operators of abnormal conditions in time. The monitoring scope of HLS-II alarm system includes the front end devices of subsystems and the server systems.

The control system of HLS-II is a distributed system based on Experimental physics and industrial control system (EPICS). Phoebus/Alarms is the latest alarm software released in EPICS community [2-4]. The alarm system for the front end devices is designed based on Phoebus/Alarms [5]. To meet the requirement of HLS-II, the alarm distribution way, such as WeChat, needs to be customized. In addition, nuisance alarms are often observed under normal circumstances due to noise and disturbance. Therefore, it is necessary to design the HLS-II alarm distribution strategy to remove nuisance alarms.

In order to achieve comprehensive monitoring of the HLS-II control system, the monitoring of the server systems is included in the alarm system. Zabbix is widely used in large scientific facilities to resolve issues related to servers and applications, such as SuperKEKB, CERN Large Hadron Collider and the NSRL facility cluster [6-8]. Thus, Zabbix is selected as the monitoring tool for the server systems.

## OVERALL ARCHITECTURE

The alarm system of HLS-II monitors the abnormal states of front end devices and server systems. The architecture of the HLS-II alarm system is shown in Fig. 1. PVs represent various attributes of the front end devices. The alarm server

monitors the PVs stored in Kafka topic and updates their alarm states in Kafka. Kafka is a distributed messaging service. It can store messages and then efficiently send all stored messages to newly connected client.

The server systems contain OPI servers, data archiver servers, database servers, web application servers and file servers, etc. Most servers are deployed on virtual machines created by VMware vSphere. The operating system of servers is Linux. Zabbix monitors metrics through Zabbix agents running on individual servers. The metrics are processed by the Zabbix server and then stored in the Zabbix database. The Zabbix Web provides graphing functionality.

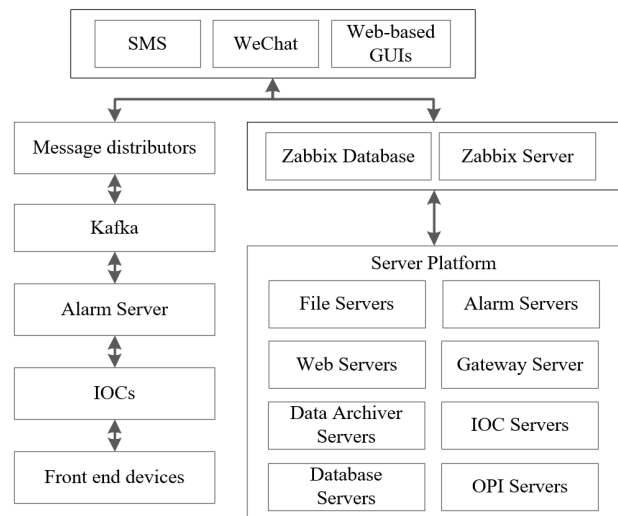


Figure 1: Architecture of the HLS-II alarm system.

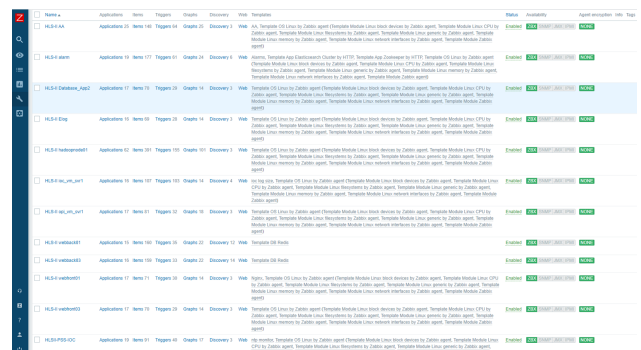


Figure 2: Screenshot of Zabbix web.

\* sunxk@ustc.edu.cn

# CONTROL SYSTEM OF THE ForMAX BEAMLINE AT THE MAX IV SYNCHROTRON

W. T. Kitka\*, S2Innovation, Kraków, Poland

V. Da Silva, V. Haghighat, Y. Li, J. Lidón-Simón, M. Lindberg, S. Malki,  
K. Nygard, E. Rosendahl, MAX IV, Lund, Sweden

## Abstract

This paper describes the design and implementation of the control system for the ForMAX beamline at the MAX IV synchrotron. MAX IV is a Swedish national laboratory that houses one of the brightest synchrotron light sources in the world, providing opportunities for cutting-edge research across a range of disciplines. ForMAX is one of the beamlines at MAX IV, designed for in-situ multiscale structural characterization from nanometer to millimeter length scales by combining full-field tomographic imaging, small- and wide-angle X-ray scattering, and scanning SWAXS imaging in a single instrument. To meet the specific demands of ForMAX, a new control system was developed using the TANGO Controls and Sardana frameworks. TANGO Controls provides a distributed control system that enables communication between devices and software, while Sardana is a Python-based software suite for controlling and coordinating data acquisition and processing. Using these frameworks allowed for the seamless integration of hardware and software, ensuring efficient and reliable beamline operation. The control system was designed to support a variety of experiments, including full-field tomographic imaging, small- and wide-angle X-ray scattering, and scanning SWAXS imaging. The system allows for precise control of the beam position, energy, intensity, and sample position. Furthermore, the system provides real-time feedback on the status of the experiments, allowing for adjustments to be made quickly and efficiently. In conclusion, the design and implementation of the control system for the ForMAX beamline at the MAX IV synchrotron has resulted in a highly flexible and efficient experimental station. TANGO Controls and Sardana have allowed for seamless integration of hardware and software, enabling precise and reliable control of the beamline for a wide range of experiments.

## MAX IV AND FORMAX INTRODUCTION

MAX IV synchrotron [1], situated in Lund, Sweden, is a facility housing a 1.5 GeV storage ring and 3 GeV storage ring designed for the generation of highly brilliant X-ray synchrotron radiation. Employing state-of-the-art insertion devices and beamlines, MAX IV enables advanced studies across diverse scientific domains. The facility's beamlines are equipped with cutting-edge instrumentation to facilitate experiments in materials science, structural biology, chemistry, and physics. The synchrotron's operational parameters and advanced optics provide researchers with an exceptional

toolset for probing materials at the atomic and molecular levels, fostering investigations into electronic structures, chemical processes, and intricate biological macromolecules. ForMAX, located at achromat 9 of the MAX IV 3 GeV ring, is a hard X-ray beamline focused on versatile structural characterization. With an emphasis on efficiency, the beamline seamlessly switches between full-field X-ray microtomography, small- and wide-angle X-ray scattering (SWAXS), and scanning SWAXS imaging. The microtomography provides non-destructive 3D mapping in the microscale range (1  $\mu\text{m}$  to 5 mm), enabling studies such as porosity characterization in forest-based materials with a temporal resolution of 1 s. SWAXS explores nanoscale structures (1 to 500 nm) for understanding biobased nanomaterials with a temporal resolution in the ms regime. Scanning SWAXS imaging generates 2D or 3D images of fibril orientation within samples, but its temporal resolution is limited due to the potential need for  $\approx 106$  individual SWAXS images for 3D reconstruction.

## FORMAX CONTROL SYSTEM

The control system governing the ForMAX beamline is a sophisticated infrastructure designed for precise and efficient experimental control. Making use of the TANGO Controls framework [2], it establishes a distributed architecture for seamless communication between diverse hardware components and software modules. Sardana [3], a Python-based software suite, assumes a pivotal role in orchestrating the control of ForMAX's instrumentation, ensuring streamlined data acquisition and processing. The integration of PandABox [4] augments system versatility, enabling adaptable control of diverse devices. IcePAPs, utilized as motor controllers, further amplify the system's capabilities, ensuring high-precision positioning and effective movement control. Taurus, Taranta and SVG synoptic are used for visualisation. The beamline synoptic panel is shown in Fig. 1.

## Optics

ForMAX's optical components crucially manipulate x-ray beam parameters, utilizing a double multilayer mirror monochromator (MLM) with W/B4C and Ru/B4C layers, and a double-crystal Si (111) monochromator (DCM), both horizontally deflecting. Dynamically bendable Kirkpatrick-Baez mirrors provide beam control, complemented by four diagnostics modules with slits, radiation safety components, and beam viewers. The system incorporates dual monochromators, their activity governed by a global variable. Their distinct operational modes requires transitions managed via the `select_mono` macro. Monochromator energy is user-

\* wojciech.kitka@s2innovation.com

# THE RF PROTECTION INTERLOCK SYSTEM PROTOTYPE VERIFICATION\*

W. Cichalewski<sup>†</sup>, W. Jalmuzna, P. Amrozik, G. Jablonski, R. Kielbik, K. Klys, R. Kotas,  
P. Marciniak, B. Pekoslawski, W. Tylman

Department of Microelectronics and Computer Science, Lodz University of Technology, Poland

B. Chase, N. Patel, P. Varghese, E. Harms, P. Prieto

Fermi National Accelerator Laboratory, Batavia IL, USA

## Abstract

The Radio Frequency Protection Interlock system plays a vital role in the LLRF-related/dependent accelerator sections Protection. Its main role is to collect information from a number of different sensors and indicators around the nearest cavities and cryomodules and provide instant RF signal termination in case of safety threshold violation. This submission describes a newly designed RFPI system tailored to the requirements of Proton Improvement Plan II (PIP-II). The proof of concept prototype of this system has been built. The paper includes also the PIP2IT (PIP-II integrated test stand) environment evaluation test results and findings as input to the next full-scope prototype design.

## INTRODUCTION

The Low Level RF (LLRF) control systems are widely used in normal and superconducting accelerators to provide optimal conditions for energy transfer from accelerating structures to the accelerated proton or electron beam. The performance of this system can determine overall beam parameters hence influence conditions of laser light or neutron beam generation or others - depending on the given infrastructure type.

The LLRF system has to ensure instantaneous correction of the amplitude and phase parameters of electromagnetic accelerating waves. In the case of superconducting resonators working as standing wave systems its main focus is the execution of the fast feedback algorithms responsible for such actions. To perform its role efficiently LLRF system has to be provided with strictly defined operating conditions for the whole accelerating structure. A system that is responsible for such conditions monitoring and system hardware protection exists in many accelerator implementations. In the case of the newly designed Proton Improvement Plan II (PIP-II) H- accelerator in Fermilab [1], such system has been named Radio Frequency Protection Interlock (RFPI)

The RFPI system's main role is to monitor various cryomodule components' status values and drop the permit signals in case any predefined safety thresholds are exceeded [2]. To gather comprehensive knowledge about subsystems' current state the RFPI has to observe several various signals from vacuum, cryogenics, temperature sensors, RF signals

leak sensors, electron pick-up, input fundamental coupler high voltage power supply, solid state amplifier status, LLRF system status and others.

For many years the RFPI system developed and implemented earlier in Fermilab provided constant protection for systems operated in various locations of local accelerator infrastructure. Together with the new design of the LLRF system, it has been also decided that the PIP-II dedicated RFPI system needs to be adapted to the specific project needs. The requirement concerning system redesign according to modern electronic standards was not the only one. A newly designed RFPI system needs to be modular and provide configurable sub-sets of input signals depending on the deployment place. Because of that not only the flexibility in system hardware configuration has to be delivered. The protection function logic needs to be also susceptible to reconfiguration and scaling. What is more, a single PIP-II RFPI system has to protect from one to four LLRF systems (and resonators) gathered in a single (or different) cryomodule.

The proposed new design fulfills the requirements concerning modularity, scalability, and reconfiguration challenges. The new solution incorporates dedicated, custom-designed signal-conditioning modules connected to the carrier board via custom FMC modules. This component contains an FPGA chip that holds and executes protection function logic. Hence configuration can be modified by removing or introducing a given conditioning submodule and adjusting the protection function when needed. This approach foreseen also the System On-Chip (SOC) chip configuration usage for better integration of the EPICS control system dedicated IOC with the RFPI infrastructure.

The first version of the system has been designed and produced in the form of a Proof of Concept (PoC) prototype [3]. This version of the system consists of custom-designed signal conditioning and dedicated interface FMC modules as well as COTS components integrated into a single RFPI instance. Although the PoC does not possess enough input channels to cover the full four cavity set-up its main goal was the verification of each input channel type conditioning/processing quality and performance.

## PROOF OF CONCEPT PROTOTYPE HARDWARE STRUCTURE

The RFPI system PoC hardware structure is illustrated in Fig. 1.

The PoC system consists of the following modules:

Hardware

Hardware Technology

\* Work supported, in part, by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under U.S. DOE Contract No. DE-AC02-07CH11359.

<sup>†</sup> wcichal@dmes.pl

# ELK STACK DEPLOYMENT WITH ANSIBLE

T. Gatsi, X. P. Baloyi, J. L. Lekganyane, R. L. Schwartz  
South African Radio Astronomy Observatory, Cape Town, South Africa

## Abstract

The 64-dish MeerKAT radio telescope, constructed in South Africa, became the largest and most sensitive radio telescope in the Southern Hemisphere and will eventually be integrated with the Square Kilometre Array (SKA).

A Control and Monitoring (CAM) system for a Radio Astronomy Project, such as MeerKAT, produces vast amounts of sensor data and event logs. Viewing and analysing this data to trace system issues require engineers and maintainers to spend significant time searching for the related events.

The ELK (Elasticsearch, Logstash, Kibana) software stack, deployed using Ansible, was implemented for the MeerKAT radio telescope in order to have the capability to aggregate system process logs and save search time.

A cluster deployment was used to ensure load balancing, high availability and fault tolerance within the MeerKAT CAM environment. This was deployed using Linux Containers (LXC) running inside a Proxmox virtual environment, with Ansible as the software deployment tool. Each container in the cluster performs cluster duties, such as deciding where to place index shards, and when to move them. Each container is also configured to be a data node, which makes up the heart of the cluster.

Logstash is used to ingest, transform and send data to Elasticsearch for indexing. The data is then made available for visualisation to the MeerKAT Operations Team and other interested users via the Kibana Graphical User Interface (GUI).

## INTRODUCTION

The MeerKAT radio telescope generates large amounts of logs daily. The ability to effectively manage, analyse, and extract insights from this data is paramount. One key solution that has emerged to address this need is the ELK cluster, comprising three core components – Elasticsearch [1], Logstash [2], and Kibana [3]. The ELK stack [4] provides a robust framework for ingesting, processing, storing, and visualising data.

In the context of the MeerKAT CAM system environment, an ELK cluster is deployed on three LXC nodes within a Proxmox virtual environment. This enables efficient display and analysis of control and monitoring system logs through Kibana, as shown in Fig. 1.

The process of deployment of this ELK cluster can be complex, requiring numerous resources and time. To make the deployment easier and faster, the automation platform from Red Hat, Ansible [5], was used within the MeerKAT CAM environment. This paper outlines the process of how this was accomplished.




Figure 1: ELK stack deployment in CAM.

## ELK SOFTWARE STACK

### Elasticsearch

At the heart of the ELK cluster lies Elasticsearch. This powerful and scalable search and analytics engine serves as the backbone of the entire setup [6]. Its primary role is to efficiently index and store data, making it searchable and retrievable in real-time.

Complementing Elasticsearch is Logstash, a versatile data processing pipeline. Logstash serves as the bridge between data sources and Elasticsearch, facilitating the ingestion of data, its transformation, and potential enrichment.

In the MeerKAT CAM environment, Elasticsearch's distributed nature allows it to be seamlessly spread across our three Ubuntu LXC container nodes. As CAM system logs are generated, they are fed into Elasticsearch via the configured Logstash server, where they are indexed and organised, forming the foundation for comprehensive log analysis.

**Indexing** Indexing entails importing data from external sources into the elasticsearch cluster. In the CAM software application, data is ingested as logs from the production system. ElasticSearch functions as a textual indexer, exclusively analysing plain text data. However, a plugin allows data storage in base64 format.

During data indexing, fields are defined, where either a built-in analyser is used, or a custom analyser is created. Specific fields are also selected to be made available in search results. Notably, indexing operations differs from typical CRUD (Create, Read, Update and Delete) operations – instead of updating or deleting data directly, Elasticsearch generates a new version of the index while marking the old version as deleted.

This point is crucial, since improper configurations can lead to indefinite data expansion due to accumulating "deleted" versions. Properly configured purging involves segmenting shards and periodically merging segments to

# SARAO SCIENCE REPOSITORY: SUSTAINABLE USE OF MeerKAT DATA

Z. Kukuma, G.L. Coetzer<sup>1</sup>, R. S. Kupa, C. Scholar  
South African Radio Astronomy Observatory, Cape Town, South Africa  
<sup>1</sup>also at University of Pretoria, Pretoria, South Africa

## Abstract

The South African Radio Astronomy Observatory (SARAO) is excited to announce the forthcoming release of its digital repository for managing astronomical data of the MeerKAT Radio Telescope. The repository, built using DSpace software, will allow researchers to catalogue and discover research data in a standardised way, while Digital Object Identifiers (DOIs) minted through the DataCite service, will ensure the unique identification and persistent citation of this research data. In this paper, we discuss the design of the repository as well as the use of DataCite for DOI minting.

## INTRODUCTION

In recent years, mechanisms for data discovery and retrieval have been established. Due to these mechanisms, new and unforeseen uses of data are being discovered [1]. With the Square Kilometre Array (SKA) radio telescope project and the new Multi-Purpose Reactor (MPR) soon coming online, organisations such as SARAO which facilitates these projects will need to make the management, analysis, publication and curation of big scientific data a priority. The recent draft of the Department of Science and Innovation (DSI) on Open Science (OS) policy, requires Open Access (OA) to both scholarly publications and scientific data [2]. It also endorses ingest, discovery and dissemination of data and metadata in a manner consistent with Wilkinson's [3] 'FAIR principles' - making data Findable, Accessible, Interoperable and Reusable (FAIR) [4]. SARAO recognises the importance of repositories and Digital Object Identifiers (DOIs) as mechanisms which can improve the FAIRness of data [5], and follows the National Research Foundation (NRF)'s vision 2030 of putting "science and research into service for a better society [6]. With this in mind, SARAO's librarian and software engineers developed a digital repository.

## OPEN SCIENCE, OPEN ACCESS AND DIGITAL OBJECT IDENTIFIERS (DOIs)

The SA draft National OS policy defines OA as "a set of principles through which research outputs are distributed online, free of cost or other access barriers" [2]. Moreover, OS includes both OA to research output and unhindered access to accompanying raw data and software used to analyse the data. OS is particularly important because it [7]:

- promotes accurate verification of research outputs;
- reduces duplication;

- promotes innovation and increased consumer choice;
- promotes citizen's trust in science;
- promotes public participation in research.

To make OS a reality, UNESCO [7] designed a set of principles to ensure that there is:

- transparency and reproducibility of research outputs;
- enhance impact of science on society;
- collaboration which transcends geography, language and resources barriers;
- solve problems of great social importance
- acknowledgement that there there does not exist a one-size-fits-all method for practising OS
- encourage pathways to practice
- sustainability by building on long-term practises, services, infrastructures and funding models.

DOIs are vital for OS, OA and FAIR data practises. A DOI is an alphanumeric used to uniquely identify objects. Objects can be articles, documents, software, networks, scientific data and products, et cetera. A resolvable DOI consists of a resolver, prefix and suffix as shown in Fig. 1.




Figure 1: Composition of a DOI (adapted from Novacescu [8]).

A DOI is assigned to a web address (URL) where an object (e.g. dataset) and its metadata can be located permanently [9]. DOIs help to ensure the integrity of digital objects - with a persistent identifier that remains unchanged even if the location changes, users can always find and access the correct version of the digital object. A DOI's value as a persistent identifier is lost if the DOI metadata is not updated when the object it references changes physical location (i.e. the URL changes). The publisher or responsible party must update the URL in the metadata record to ensure that the DOI continues to redirect users to the object. Metadata collected before data release are generally used for formal citation, thus facilitating credit and acknowledgement for data creation and usage. DOIs allow different platforms, databases and information systems to exchange information consistently and unambiguously.

Examples of softwares that encourages sustainable information systems are DuraSpace (DSpace), Figshare and Zenodo amongst others [9]. They are technical frameworks

# IMPROVING USER EXPERIENCE AND PERFORMANCE IN SARDANA AND TAURUS: A STATUS REPORT AND ROADMAP

Z. Reszela\*, J. Aguilar, M. Caixal, G. Cuni, R. Homs-Puron, E. Morales,  
M. Navarro, J. Ramos, S. Rubio, O. Vallcorba, ALBA-CELLS, Barcelona, Spain  
M. T. Núñez Pardo de Vera, DESY, Hamburg, Germany  
B. Bertrand, J. Foresberg, MAX IV Laboratory, Lund, Sweden  
M. Piekarski, SOLARIS, Krakow, Poland  
D. Schick, Max Born Institute, Berlin, Germany

## Abstract

The Sardana Suite is an open-source scientific SCADA solution used in synchrotron radiation facilities such as ALBA, DESY, MAX IV, and SOLARIS as well as in laser laboratories such as the Max Born Institute. It is formed by Sardana and Taurus - both mature projects, driven by a community of users and developers for more than ten years. Sardana provides a low-level interface to the hardware, middle-level abstractions, and a sequence engine. Taurus is a library for developing graphical user interfaces. The Sardana Suite uses client-server architecture and is built on top of TANGO.

As a community, during the last few years, on the one hand, we were focusing on improving user experience, especially in terms of reliability and performance, and on the other hand renewing the dependency stack. The system is now more stable, easier to debug and recover from a failure. An important effort was put into profiling and improving the performance of Taurus applications during startup. The codebase has been migrated to Python 3 and the plotting widgets were rewritten with pyqtgraph. In addition, we also provide new features, like for example the long-awaited Sardana configuration tools and format based on YAML which is easy and intuitive to edit, browse, and track historical changes.

Now we conclude this phase in the projects' lifetimes and are preparing for new challenging requirements in the area of continuous scans like higher data throughput and more complex synchronization configurations. Here we present the status report and the future roadmap.

## SARDANA SUITE OVERVIEW

Sardana is a Python-based scientific SCADA suite. Its primary goal is to reduce the cost and time associated with the design, development, and support of control and data acquisition systems. The suite consists of two independent projects: Taurus [1, 2] and Sardana [3, 4], and is built on top of the TANGO control system [5].

Taurus is a framework designed for creating user interfaces, including GUIs and command-line interfaces, to interact with scientific and industrial control systems, as well as other relevant data sources. GUIs are developed using PyQt.

Sardana, on the other hand, is a framework focused on the automation of experimental procedures and the control of

laboratory equipment. It includes a powerful sequencer known as MacroServer, which incorporates a versatile scan and data storage mechanism. Additionally, it features a Device Pool that defines generic interfaces for laboratory elements and implements the hardware access layer. MacroServer as the client of Device Pool use Taurus core to implement the client side part of the communication. Spock which is based on IPython serves as a centralized command-line interface application for Sardana users.

## USER EXPERIENCE

During the last years we have significantly improved the Sardana Suite user's experience with the most relevant changes listed below:

### User Interfaces

**Trends with Live and Archived Data** The `taurus_pyqtgraph` trend widget [6], introduced the archiving support in the release 0.6. This development has fundamentally transformed the way users access and interact with historical data within the widget, offering valuable context for data analysis on a single tool, while using HDB++ as a source for archived information.

We achieved this by introducing two options:

- **Loading Archived Data Once:** Users can load historical data once and visualize it based on the date time axis they are exploring, this while the trend is continuously updated with new data being added.
- **Automatic Data Loading:** This feature allows users to navigate the date time axis retrieving and displaying archived data relevant to their current view with no need to click any button. This feature ensures that users maintain an up-to-date and continuous visualization of historical data as they explore different time frames.

**Experiment Configuration** The experiments can be configured using client applications, either the `expconf` GUI tool or directly through the `Spock` CLI. To ensure consistency, any changes made in one of the clients must be immediately reflected in all others.

The server notifies the clients about configuration changes with TANGO attribute change events, then it is the client's choice on how to reflect the incoming changes. Initially, this was solved by using *pop-up* dialogs that offered `expconf` user options to accept or discard the changes. However,

\* zreszela@cells.es

## LLRF AND TIMING SYSTEM INTEGRATION AT ESS

G. Fedel\*, A. Svensson, J. J. Jamroz, J. P. S. Martins, A. Gorzawski, R. Zeng, N. Milas, A. Persson  
ESS, European Spallation Source, Lund, Sweden

### Abstract

The Low-Level Radio Frequency (LLRF) system is an important part of an accelerator as the linear accelerator from the European Spallation Source (ESS). LLRF is commonly used with many different setups depending on the aim: preparation, calibration, conditioning, commission, and others. These different setups are connected to the ESS Timing System (TS).

This paper presents the control system integration between LLRF and TS. The integration of these two systems provides different and important features as allowing different ways to trigger the RF system (synced or not to other systems), define how the RF output will be defined (based on the features of the expected beam), re-configure LLRF depending on the timing setup and more. This integration was developed on both ends, LLRF and TS, and is mostly concentrated on the control system layer based on Experimental Physics and Industrial Control System (EPICS). Dealing with the different scenarios, synchronicity, and considering all the software, hardware, and firmware involved are some of the challenges of this integration. The result of this work was used during the last two ESS accelerator commissionings in 2022 and 2023.

### INTRODUCTION

The ESS [1] is on the path to deliver the proton beam up to the target where the neutrons will be produced. The ESS linac is shown in Fig. 1. In 2022 [2] and 2023 two beam commissionings on the ESS accelerator covering up to the RF Quadrupole (RFQ) and the Drift Tube Linac (DTL) 4 were performed. To assure a stable commissioning with different setups and also allow the conditioning of different cavities in parallel a special integration between the RF systems and the timing system was needed. This integration allows the RF expert or Cavity specialist to control the RF repetition rate from one cavity independently of the others. It allows also the LLRF to have different automatic behaviors over the feed-forward correction depending on the expected beam (information delivered by the timing system).

The next sections present the details of this implementation and its challenges.

### ESS TIMING SYSTEM

The main role of the ESS TS is to generate, acquire, and distribute RF  $\approx 704.42$  MHz based timing signals, with the topology in Fig. 2. The timing hardware is based on Micro Telecommunications Computing Architecture (MTCA) and the detailed concept was described in [3]. The Timing Master (TM) is the main TS gateway for the operation control

while event receivers (EVRs) embedded within the timing distribution to machine subsystems (TDMS) perform synthesis of the received information into the timing signals required by a particular subsystem.

### RF Events

The TM produces different events that are used to trigger actions on different subsystems (TDMS from Fig. 2). Specifically for RF subsystems, there is one main event called RF Start which is used to trigger the following actions at one RF station:

- Start Modulator,
- Start LLRF (which will generate the initial RF pulse),
- Start Local Protection System (LPS) data acquisition.

Besides the RF Start event, the RF EVR is configured to look into Beam Start and Beam End events, although these events are used only as marker points for LLRF (beam envelope).

### Data Buffer

Another important piece of information that is distributed by the timing system is the data buffer information, a deterministic set of data. The data contains information about the next beam pulse, and it will be used to allow different systems to prepare themselves before the proton beam pulse is transported along the linac. The data used by a RF system for example:

- Beam Present, i.e. is there a beam expected,
- Beam Destination,
- Expected Beam Current.

The data information is provided for the subsystem (e.g. RF system) by the EVR EPICS Input/Output Controller (IOC).

### ESS RF SYSTEM AND ESS LOW-LEVEL RF

The RF systems at ESS are composed of different elements, most of which are illustrated in Fig. 3. This figure shows some of the main elements as the LLRF, pre amplifier, klystron, and modulator for a generic cavity. For different cavities, the elements can change a bit. Each of these components, responsible for producing RF for one cavity, is called an RF station.

At ESS there are two types of RF stations, the ones using modulators and klystrons, and the lower power stations, where modulators are not required. The RF systems using klystrons share the same modulator [4] (2 or 4 klystrons per modulator). From the ESS sections, the ones following this topology are RFQ, DTLs, Medium Beta Linac (MBL), and High Beta Linac (HBL). The topology of these systems is illustrated in Fig. 4. That figure presents also a new element from an RF station - EVR. The EVR is the component responsible for connecting the RF station to the ESS TS.

\* gabriel.fedel@ess.se

# CHARACTERIZING MOTION CONTROL SYSTEMS TO ENABLE ACCURATE CONTINUOUS AND EVENT-BASED SCANS

J. Petersson\*, T. Bögershausen, N. Holmberg, M. Olsson, F. Rojas, T. Richter  
European Spallation Source ERIC, Lund, Sweden

## Abstract

The European Spallation Source (ESS) is adopting innovative data acquisition and analysis methods using global timestamping for neutron scattering research. This study characterises the timing accuracy and reliability of the instrument control system by examining an integrated motion and fast detection system.

We designed an experimental apparatus featuring a motion axis controlled by a Beckhoff programmable logic controller (PLC) using TwinCAT 3 software. The encoder readback is timestamped in the PLC, which is time-synchronised with the ESS master clock via a Microresearch Finland event receiver (EVR) using Precision Time Protocol (PTP).

We repeatedly scanned the motor between known positions at different speeds. The system was characterised by correlating the position and timestamp recorded by the PLC with independent information using a fast optical position sensor read out directly by the MRF system.

The findings of this study provide a good benchmark for the upcoming experiments in neutron scattering research at ESS and should be interesting for those aiming to build similar setups.

## INTRODUCTION

The European Spallation Source (ESS) is an international collaboration to build the world's most powerful neutron source [1]. Upon completion, it will produce a high-intensity proton beam, which will be directed onto a tungsten target, thereby generating neutrons through a process known as spallation. These neutrons can then be harnessed for a wide variety of scientific investigations, providing researchers with insights into the structural and dynamical properties of materials. Applications are broad and include fields such as solid-state physics, materials science, crystallography, biology, and archaeology.

To make the most out of this research potential, ESS employs a specific approach to handling the data produced. It captures neutron data in what is referred to as event mode: Each detected neutron is recorded individually and characterised by a timestamp and a pixel identifier, tracking the exact moment and location where the neutron hit the detector [2, 3]. Compared to histogram mode, where data is accumulated over a given period and represented as a set of counts, event mode retains more detailed information. This enhanced granularity of data provides researchers with the flexibility to bin and filter data in a resolution of their choosing, thus permitting more nuanced analyses.

To make the most out of the event mode recording, all relevant parameters about the measurement need to be known with the appropriate timing accuracy. Therefore, the timing system at ESS plays an integral part in the process by synchronising operations. It consists of an Event Generator (EVG) and a series of Event Receivers (EVRs) arranged in a tree topology to distribute events, clock signals, and data [4]. The EVG creates the event clock signal, which is transmitted to the EVRs. The EVRs, in turn, decode the incoming data and generate output triggers controlled by software. The EVRs can also handle inputs which can be timestamped. In addition to the timing system hardware, the Precision Time Protocol (PTP) is utilised to synchronise clocks, thus providing accurate timestamping and synchronisation among different subsystems.

In this paper, we investigate the correlation between motion data and event data within the ESS environment. The underlying question is: Given an event's timestamp, what is the uncertainty in position for a load mounted on the linear stage for that timestamp? By investigating this, we aim to quantify the performance of the motion control systems at ESS in terms of their ability to perform accurate continuous and event-based scans.

## METHOD

In addition to the above-mentioned timing system, the experimental setup consists of a linear stage with a motor and encoder as well as a laser-based position detection system. A schematic overview is depicted in Fig. 1.

### Laser System

A laser and a laser sensor are aligned on a beam table. The laser sensor (Thorlabs S120C) is connected to a power meter (Thorlabs PM100D) that delivers an analogue output voltage [5]. The output voltage is connected to a Schmitt trigger, which can deliver a 5V TTL pulse to the EVR for timestamping when the laser beam passes through the motorised aperture.

To characterise the optical detection, an LED triggered by the EVR was placed in front of the laser sensor to determine the time resolution of the laser readout system. The resolution of the laser timestamp setup is determined by triggering the LED with a TTL output from the EVR, which has an accurate timestamp, and then timestamp the TTL pulse generated from the laser sensor on an input on the EVR, effectively measuring the round-trip of the signal. The measured standard deviation of the round-trip time is on the order of 10  $\mu$ s, which is expected since the DC bandwidth of the sensor amplifier is 100 kHz.

\* jonas.petersson@ess.eu



# TEST AUTOMATION FOR CONTROL SYSTEMS AT THE EUROPEAN SPALLATION SOURCE

Karl Vestin, Fabio dos Santos Alves, Lars Johansson, Stefano Pavinato, Kaj Rosengren, Marino Vojneski, European Spallation Source, Lund, Sweden

## Abstract

This paper describes several control system test automation frameworks for the control systems at the European Spallation Source (ESS), a cutting-edge research facility that generates neutron beams for scientific experiments. The control system is a crucial component of ESS, responsible for regulating and monitoring the facility's complex machinery, including a proton accelerator, target station, and several neutron instruments.

The traditional approach to testing control systems largely relies on manual testing, which is time-consuming and error-prone. To enhance the testing process, several different test automation frameworks have been developed for various types of applications. Some of these frameworks are integrated with the ESS control system, enabling automated testing of new software releases and updates, as well as regression testing of existing functionality.

The paper provides an overview of various automation frameworks in use at ESS, including their architecture, tools, and development techniques. It discusses the benefits of the different frameworks, such as increased testing efficiency, improved software quality, and reduced testing costs. The paper concludes by outlining future development directions.

## INTRODUCTION

The control system for a large research facility, such as the European Spallation Source (ESS), comprises thousands of different subsystems. Each of these subsystems has hundreds or even thousands of distinct configuration items, and each plays a role in fulfilling a specific function within the facility. The continued development and maintenance of such a complex system-of-systems present significant challenges in terms of verification.

The control systems at ESS are constructed using local control systems, typically based on Programmable Logic Controllers (PLCs) or high-speed data acquisition systems using Field Programmable Gate Arrays (FPGA)-based boards slotted into Micro Telecommunications Computing Architecture (MicroTCA) crates. All control systems are integrated into the unified control system using Experimental Physics and Industrial Control System (EPICS) Input Output Controllers (IOCs). All IOCs are based on reusable EPICS software modules distributed as part of the ESS EPICS Environment (e3) [1].

Manually testing and re-testing every component of the control system after each update is prohibitively time-consuming. To address this challenge, test automation is applied. Through the execution of test scripts in well-defined test environments, we can efficiently test our systems for

regressions and faults before redeploying them after an update.

At ESS, various approaches to test automation are employed, tailored to the type and application of each specific control system. This paper will provide an overview of the diverse technologies and techniques used, discuss the rationale behind their selection, and present ideas for the future expansion of test automation.

## EPICS MODULES

Most of the modules in the e3 environment are developed by the EPICS community and tested as part of the release process. ESS has developed dedicated test scripts to verify a subset of the modules. In most cases, the tests require some level of hardware emulation.

As an example, let us examine the unit test scripts for the Open Platform Communication Unified Architecture (OPCUA) e3 module [2]. The EPICS module is developed by the controls team at the International Thermonuclear Experimental Reactor (ITER) and wrapped by the ESS e3 team for integration into e3. All the unit tests are declared in a test script that is executed as a dedicated make target, essentially a defined objective for the general purpose tool for generation of executable from source code "make". This make target is executed by a process - a runner - automatically started by the ESS GitLab continuous integration function upon check-in of the changes. Figure 1 illustrates how the results are stored in GitLab and are available for all historical builds.

Hardware emulation, in this case, is achieved using a dedicated OPCUA server based on the open-source code from the open62541 project [3]. The server is built as part of building the make target. The server is then started as part of the pytest test fixture in the test script to ensure the test cases always have a server to connect and test against.

One interesting aspect of the OPCUA unit test is the utilization of libfaketime [4], which is now installed by default on our test runner machines. Using this library, the test fixture can start the OPCUA server using a different ("fake") system clock, thereby verifying that timestamping is done correctly, as specified by the Time Stamp Event (TSE) field in the record.

Before a new version of the module is released into the production environment, the e3 team verifies that the pipeline runs cleanly, including any defined unit test scripts.

For the modules in the e3 environment, work is currently underway to improve test coverage through automated testing. The ambition is to have automated tests running as part of the continuous integration process.

# CONSOLIDATION OF THE POWER TRIGGER CONTROLLERS OF THE LHC BEAM DUMPING SYSTEM

L. Strobino\*, N. Magnin, N. Voumard, CERN, Geneva, Switzerland

## Abstract

The Power Trigger Controller (PTC) of the LHC Beam Dumping System (LBDS) is in charge of the control and supervision of the Power Trigger Units (PTU), which are used to trigger the conduction of the 50 High-Voltage Pulsed Generators (HVPG) of the LBDS kicker magnets. This card is integrated in an Industrial Control System (ICS) and has the double role of controlling the PTU operating mode and monitoring its status, and of supervising the LBDS triggering and re-triggering systems.

As part of the LBDS consolidation during the LHC Long Shutdown 2 (LS2), a new PTC card was designed, based on a System-on-Chip (SoC) implemented in an FPGA. The FPGA contains an ARM Cortex-M3 softcore processor and all the required peripherals to communicate with onboard ADCs and DACs (3rd-party IPs or custom-made ones) as well as with an interchangeable fieldbus communication module, allowing the board to be integrated in various types of industrial control networks in view of future evolution.

This new architecture is presented together with the advantages in terms of modularity and reusability for future projects.

## INTRODUCTION

### LHC Beam Dumping System

The LHC Beam Dumping System (LBDS) is a critical system ensuring safe extraction of the beams from the LHC. Each counter-rotating beam is sent to its extraction channel using 15 extraction kicker magnets (MKD) and 15 extraction septa (MSD). It is then diluted by 4 horizontal (MKBH) and 6 vertical (MKBV) dilution magnets on the beam dump absorber (TDE). To allow for the rising edge of the MKD magnetic field, a particle-free Beam Abort Gap (BAG) of 3  $\mu$ s is maintained in LHC [1].

Dump Requests (DR) come from 3 different sources: the machine protection system for emergencies, the machine timing system for scheduled dumps, or the LBDS itself in case of internal failures. These spontaneously issued dump requests are synchronized with the BAG by the Trigger Synchronization Units (TSU). Synchronous dump triggers are then distributed through the Trigger Fan-Out units (TFO) to the High-Voltage Pulsed Generators (HVPG) [2].

50 HVPGs power the MKDs and MKBs. At the reception of a trigger, two redundant Power Trigger Units (PTU), each comprising a Power Trigger Controller (PTC) and two redundant Power Trigger Modules (PTM), start the conduction of two switches composed of Fast High Current Thyristors (FHCT) that discharge capacitors into the magnet. In addition,

a redundant fault-tolerant Re-Triggering System (RTS) allows the asynchronous fast re-trigger of all HVPGs if one HVPG self-triggers. Each HVPG is equipped with two Re-Trigger Boxes (RTB), that couple internal pickup signals to the redundant Re-Trigger Lines (RTL) to generate a pulse when the HVPG is triggered, and also capture the pulses on the RTLs and send them to their PTUs.

Figure 1 shows a simplified schematics of the redundant trigger distribution from the TSUs and the RTBs to the PTUs.

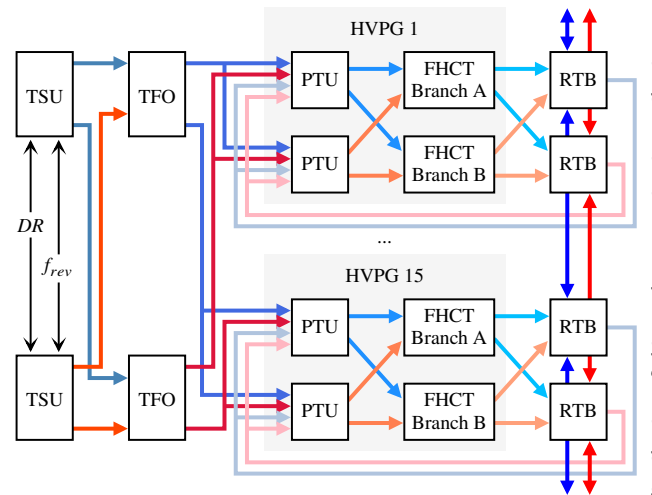


Figure 1: LBDS trigger distribution architecture.

### Power Trigger Units

A PTU is composed of several internal low and high-voltage power supplies, two PTMs operated in parallel, and one PTC to control, monitor, and interlock the PTU. All modules are connected together through a backplane, as shown in Fig. 2. Each PTU has 5 trigger inputs: two for synchronous triggers from two redundant TFOs, two for asynchronous triggers coming from the RTLs through redundant RTBs, and one software trigger input, used for test purposes. The PTU also receives a copy of the signals present at the input of one of the RTB, to be monitored by the PTC (no redundancy here: each PTU monitors one RTB).

The PTMs comprise a principal circuit generating a main trigger pulse of about 800 A for 1  $\mu$ s, and a compensation circuit maintaining a 50 A average current for a duration of about 400  $\mu$ s [3]. The principal circuit is powered from a +4 kV power supply and is controlled by 3 series-connected high-voltage IGBTs, triggered through pulse transformers by a gate driver circuit operating from a +48 V power supply. For added redundancy, each PTM is connected to a different FHCT (branches A and B in Fig. 1), in parallel with a second PTM from the other PTU.

\* lea.strobino@cern.ch

# SPS BEAM DUMP ENHANCEMENTS ON TRACKING AND SYNCHRONIZATION

N.Voumard\*, N. Magnin, P. Van Trappen, CERN, Meyrin, Switzerland

## Abstract

During CERN's long shutdown 2 (LS2) at the Super Proton Synchrotron (SPS), the SPS Beam Dump System (SBDS) was displaced from point 1 to Point and at the same time further consolidated controls-wise. This made it possible to migrate the Beam Energy Tracking system (BETS) and the Trigger Synchronization system (TSU) already operational on the LHC beam dump (LBDS) towards the SPS. The challenge encountered in this migration was the change to a Pulse-to-Pulse modulation (PPM) with much faster cycles in the SPS in comparison to the LHC. This paper describes the modification of both as well as the automatic arming sequence put in place, including the interactions with the SPS injectors, the beam revolution frequency, and the Beam Interlock System (BIS).

## BACKGROUND

### *SPS Beam Dump System (SBDS)*

Uncontrolled beam loss in the SPS will cause thermal and radiation damages of machine components as well as induced radioactivity. For this reason an internal beam dumping system was designed and installed in the seventies [1]. It uses fast pulsed magnets (kickers) to dump the beam in one SPS revolution (23.1  $\mu$ s) onto an absorber block (TIDVG). The kickers and absorber block are able to dump the beam at all momenta up to 450 GeV. The extraction system today consists of six kicker magnets which provide the beam extraction and the energy dilution 5 [2]. Three horizontal kicker magnets (MKDH) deflect the beam horizontally with a rise time of about one full SPS revolution. They provide a horizontal sweep which dilutes the beam energy across the dump block. Three vertical kicker magnets (MKDV) provide a vertical deflection and thus extraction. The short rise time of 1.2  $\mu$ s of the MKDV is responsible for the beam extraction. MKDV oscillates at flat top current which provides dispersion on the absorber block (Figure 1).

### *Trigger Synchronization Unit (TSU)*

The aim of the TSU [3, 4] is to centralize all dump requests from various clients and to synchronize them with the Beam Revolution Frequency (Frev), which allows to rise the dump kickers in the 1.44  $\mu$ s Beam Abort Gap and avoid spraying the beam around the dump region during the 1.2  $\mu$ s MKDV rise time. The TSU is a safety critical element in the SPS. The system is composed of two redundant cards in independent chassis which monitors each other and react to every discrepancy.

\* nicolas.voumard@cern.ch

Dump requests can be triggered by the following clients: Beam Energy Tracking System (BETS), Beam Interlock System (BIS), SBDS State Control and Surveillance System (SCSS), direct triggering (Early Dump), Frev loss or instability.

An important additional feature in the SBDS is that the TSU generates an Injection Permit towards the SPS Injection BIS while armed and releases it 20  $\mu$ s before issuing the actual dump trigger to avoid any possibility to inject beam after triggering and thus disarming the SBDS. The TSU also forces it User Permit to the Ring BIS during arming sequence.

Initially developed for the LHC Beam Dumping System (LBDS), the TSU system has been adapted and modified to fit to the SBDS requirements. The gateway has been reshaped by exchanging the initial PLL-based design with a less complex counter-based one for the regeneration of the Frev, because the initial LBDS design based on PLL was not adaptable to the SBDS Frev. This new gateway is compatible with both LBDS and SBDS. The TSU requires to be armed at every cycle.

### *Beam Energy Tracking System (BETS)*

The aim of the BETS [5] is to continuously survey that the SBDS MKDV and MKDH generator charge with relation to the machine momenta tracks all along the ramps, from injection to flat top. The BETS is a safety critical system which generates a dump request whenever one of the SBDS generator strengths is out of tolerance. The BETS on SBDS required modification from its implementation on LBDS, i.e allowing jumps between 14 GeV and 26 GeV beams when arming. This requires the BETS to wait for the charge of the generators to reach nominal strength during the arming sequence. On the SBDS, the BETS and TSU systems are tightly linked, while not on the LBDS. The BETS in SBDS, unlike the LBDS, must generate the references to the MKDV and MKDH high voltage power supplies (HVPS) via a dedicated DAC card. On LBDS, the generation of the generators charge references is done via the slow control Programmable Logical Controller (PLC) system that has a slow cycle time (100 ms). This is not an issue with long ramp-up of the LHC, but becomes a problem in the SPS with ramp-up of < 2 s.

## SBDS ARMING SEQUENCE

The SBDS TSU and BETS require a tight arming sequence [6] which must consider all SBDS external and internal conditions, the SPS injection transfer line TT2 BHZ377 and BZH378 dipole magnets start and abort times, SBDS generators ramp-up to injection strengths, SPS beam revolution frequency (Frev) stability, SPS Ring BIS interactions and TSU injection permit to SPS Injection BIS [7] (Fig.2).

# TOWARDS AUTOMATIC GENERATION OF FAIL-SAFE PLC CODE COMPLIANT WITH FUNCTIONAL SAFETY STANDARDS

A. Germinario\*, B. Fernández, E. Blanco, CERN, Geneva, Switzerland

## Abstract

In agreement with the IEC 61511 functional safety standard, fail-safe application programs should be written using a Limited Variability Language (LVL), that has a limited number of operations and data types, such as LD (Ladder Diagrams) or FBD (Function Block Diagrams) for safety PLC (Programmable Logic Controller) languages.

The specification of safety instrumented systems, as part of the Safety Requirements Specification document, shall unambiguously define the logic of the program, creating a one-to-one relationship between code and specification. Hence, coding becomes a translation from a specification language to PLC code. This process is repetitive and error-prone when performed by a human.

In this paper we describe the process of fully generating Siemens TIA portal LD programs for safety applications from a formal specification. The process starts by generating an intermediate model that represents a generic LD program based on a predefined meta-model. This intermediate model is then automatically translated into code.

The idea can be expanded to other equivalent LVL languages from other PLC manufacturers. In addition, the intermediate model can be generated from different specification formalism having the same level of expressiveness as the one presented in this paper: a Cause-Effect Matrix.

Our medium-term vision is to automatically generate fail-safe programs from diverse formal specification methods and using different LVLs.

Keywords: functional safety, safety instrumented systems, PLC

## INTRODUCTION

Programmable Logic Controllers (PLCs) are widely recognized as the standard for industrial process control. One of the main reason is that PLCs are robust and reliable devices that can work in harsh environments. The Mean Time Between Failures (MTBF) of PLCs are normally very high. For example, a Siemens CPU S7-1515-2 PN has a MTBF of 27.7 years (data extracted from [1]).

PLC manufacturers also provide solutions for safety systems. When in an industrial process, a particle accelerator, a machine or any other kind of system, a failure may lead to a risk for humans, the environment or a big economic and reputations loss, the project responsible must reduce the risk to the tolerable levels defined by the organization, company or regulations. Fail-safe PLCs are devices certified by organizations like TÜV SÜD [2], that have been designed to follow the IEC 61508 [3] Functional Safety standard. These

devices are used to deploy the Safety Instrumented Functions (SIF) in a way of a software program meant to reduce the risks mentioned above. The MTBF of fail-safe PLCs is also high but lower than the one of a standard PLC. For example a S7-1515F PLC has a MTBF of 24.5 years. This distinction arises from its internal architecture and increased complexity, which enables it to substantially decrease the occurrence of dangerous undetected failures when compared to a typical PLC. In discussions concerning safety-critical systems, only dangerous undetected failures are pertinent.

Fail-safe PLCs are able to detect most of their failures as stated by their Safe Failure Fraction (SFF)  $\geq 99\%$ . This means that the dangerous undetected failures are  $\leq 1\%$ . This number is so low because the safety systems internal to a safety PLC cover most of the internal failures of the controller, including hardware, operating system, firmware, etc. On the contrary, it clearly does not cover the user PLC application program (AP) for the SIFs.

The IEC 61511 standard [4] provides the guidelines to develop Safety Instrumented Systems (SIS) for industrial processes. The Clause 12 of the IEC 61511-1 focuses in the SIS AP development requirements. Among other things, it recommends the usage of Low Variability Languages (LVL) to write these programs.

Code generation is a common practice in software engineering. It has a lot of advantages in quality and efficiency like reducing the amount of coding errors introduced by the programmer, speeding up the development process, etc. For PLC programs, there are not many available tools for code generation. A good example is *PLC coder* from MathWorks [5], where ST (Structured Text) programs from the IEC61131-3 [6] can be generated automatically from Simulink models. At CERN, we use the *UNICOS* [7] framework to generate PLC programs from high level specifications.

However, most of the available tools cannot generate safety PLC programs compliant with the IEC 61511 standard. In addition, most of PLCs brands did not allow to generate safety PLC programs with external tools, import them in their programming environment and compile them as safety programs. Very recently the Totally Integrated Automation Portal (TIA portal) [8], the programming environment of Siemens PLCs, opened the door for code generation of safety PLC programs. Now source code files written in Ladder Diagrams (LD) or Function Block Diagrams (FBD) with certain restrictions to be compatible with the LVL requirements, can be imported and compiled as safety programs.

There are a very few references in the Functional Safety standards about code generation. Since this is a relative new feature in devices like PLCs, we believe that future releases of the standards will address this specific topic with more

\* andrea.germinario@cern.ch

# BEAM INSTRUMENTATION SIMULATION IN PYTHON

M. Gonzalez-Berges\*, D. Alves, A. Boccardi, V. Chariton, I. Degl'Innocenti, S. Jackson,  
J. Martínez Samblas, European Organization for Nuclear Research (CERN), Geneva, Switzerland

## Abstract

The design of acquisition electronics for particle accelerator systems relies on simulations in various domains. System level simulation frameworks can integrate the results of specific tools with analytical models and stochastic analysis. This allows the designer to estimate the performance of different architectures, compare the results, and ultimately optimise the design. These simulation frameworks are often made of custom scripts for specific designs, which are hard to share or reuse. Adopting a standard interface for modular components can address these issues. Also, providing a graphical interface where these components can be easily configured, connected and the results visualised, eases the creation of simulations. This paper identifies which characteristics ISPy (Instrumentation Simulation in Python) should fulfil as a simulation framework. It subsequently proposes a standard format for signal-processing simulation modules. Existing environments which allow script integration and an intuitive graphical interface have then been evaluated and the KNIME Analytics Platform was the proposed solution. Additionally, the need to handle parameter sweeps for any parameter of the simulation, and the need for a bespoke visualisation tool will be discussed. Python has been chosen for all of these developments due to its flexibility and its wide adoption in the scientific community. The ensuing performance of the tool will also be discussed.

Keywords: simulation, system design, digital twin, Python, performance

## INTRODUCTION

The design of electronic acquisition systems for beam instrumentation can be a lengthy process implying several steps in which the designer has to optimize different parameters. The design steps typically cover architecture definition, component selection, algorithm selection and performance estimation. For systems of a certain size, the number of parameters is not manageable directly, hence simulation tools are required to assist in producing systems that meet the specifications. A number of Python scripts were developed in the past within the CERN Beam Instrumentation (BI) group for each of the individual steps mentioned above. This paper will present the effort to have a simulation tool integrating the functionality in those scripts and extending it to cover further needs.

## SIMULATION FRAMEWORK REQUIREMENTS

Several high level requirements were identified for the simulation tool. The tool should be based on components

representing the various elements of the acquisition system (e.g. beam characteristics, filters, cables, amplifiers). Users should have the ability to extend the tool with new components without needing detailed knowledge of the tool internal workings.

Users will interact with the tool through a graphical interface, allowing them to describe simulations by connecting the different components, configuring their properties and initiating simulations. The resulting data will be stored in files, along with the simulation schematic and the parameter configuration, ensuring proper tracing of each dataset. Additionally, a visualisation tool will be provided, offering diverse views of the simulation results, complete with selection and filtering capabilities to facilitate result interpretation.

The tools and libraries employed for the implementation should be open source, encouraging collaboration and long-term sustainability. Furthermore, the definition of the simulations and their execution results should be as independent as possible of the specific tools used for their creation. This design principle will enable easy integration with other tools and facilitate the future evolution of the simulation tool.

There was already an existing set of Python scripts for partial simulations where a considerable effort had been invested. Providing methods to reuse these scripts either directly in a Python implementation or in any other compatible way would be necessary.

## CHOICE OF ENVIRONMENT

Given the requirements presented in the previous chapter, we considered several possibilities for the actual implementation. An initial idea was to develop a new tool based on Python and PyQt. These are well known in our team and there are many libraries that could be reused for the data processing, storage and visualisation. Another possible approach was to use an existing environment that could be extended to cover our simulations. We identified two possible tools: the KNIME analytics platform [1] and Kepler [2].

After a detailed technical study, it was clear it would have been possible to implement the tool with any of the three proposed solutions. However, using PyQt would have meant a larger development effort and future maintenance. In addition there would not have been any synergies with similar tools.

Out of the two existing platforms, Kepler seemed to have a smaller and less dynamic community. Additionally, the integration with Python seemed somehow more complicated than in KNIME. These points lead us to the selection of KNIME.

\* manuel.gonzalez@cern.ch

# PYTHON EXPERT APPLICATIONS FOR LARGE BEAM INSTRUMENTATION SYSTEMS AT CERN

J. Martínez Samblas\*, E. Calvo Giraldo, M. Gonzalez-Berges, M. Krupa  
European Organization for Nuclear Research (CERN), Geneva, Switzerland

## Abstract

In recent years, beam diagnostics systems with increasingly large numbers of monitors, and systems handling vast amounts of data have been deployed at CERN. Their regular operation and maintenance poses a significant challenge. These systems have to run 24/7 when the accelerators are operating and the quality of the data they produce has to be guaranteed. This paper presents our experience developing applications in Python which are used to assure the readiness and availability of these large systems. The paper will first give a brief introduction to the different functionalities required, before presenting the chosen architectural design. Although the applications work mostly with online data, logged data is also used in some cases. For the implementation, standard Python libraries (*e.g.* PyQt, pandas, NumPy) have been used, and given the demanding performance requirements of these applications, several optimisations have had to be introduced. Feedback from users, collected during the first year's run after CERN's Long Shutdown period and the 2023 LHC commissioning, will also be presented. Finally, several ideas for future work will be described.

Keywords: Python, user interfaces, data quality, data visualisation

## INTRODUCTION

The LHC is renowned for generating substantial amounts of data through particle collisions. However, it is often overlooked that a significant stream of data is generated by the numerous Beam Instrumentation (BI) systems deployed to monitor, control, and ensure the smooth operation of the accelerators.

This paper primarily focuses on two BI systems: the Diamond Beam Loss Monitors (Diamond BLMs) [1] and the Beam Position Monitors (BPMs) [2]. These large systems present serious challenges due to their extensive data production. On the one hand, despite their limited number (17 across all accelerators), Diamond BLMs can buffer millions of samples per cycle. Conversely, while BPMs individually produce less data, their vast quantity (over 1000 deployed in the LHC) contributes to a massive overall data volume.

In response to the lack of software solutions dedicated to addressing these highly demanding systems, a suite of Python applications has been developed under a set of specific mandates. Firstly, the programs must provide the flexibility to monitor all devices through a unified, user-friendly interface. Speed is also essential, not just in terms of data processing and real-time efficiency, but also in expediting system processes such as commissioning, diagnostics, and

fine-tuning, which can otherwise become tedious and time-consuming. Lastly, adopting a data-driven approach is imperative to ease maintenance and ensure scalability in the future.

## LARGE BEAM INSTRUMENTATION SYSTEMS

### *Diamond BLM System*

The LHC and SPS, along with the SPS transfer lines, are equipped with 17 Diamond BLMs. These detectors, made of diamond crystals with gold electrodes polarised at 500 V, are strategically positioned to offer a time resolution of 1.53 nanoseconds, enabling precise bunch-by-bunch loss measurements. Acquired signals are digitised at 650 MSPS, with the system supporting five parallel acquisition modes logging at 1 Hz and an on-demand mode for time windows of several milliseconds.

Diamond BLMs play a crucial role in analysing beam transfer efficiency and kicker time alignment at injection and extraction lines. In betatron collimation regions, they monitor losses throughout the entire beam cycle, which are generally associated with physical phenomena such as beam-beam interactions, electron cloud effects, and tune drifts, among others.

### *LHC BPM System*

To guarantee the safe and efficient operation of the accelerator, the LHC is armed with over 1000 BPMs distributed throughout the beam line for continuous and precise measurements of the beam's transverse position within the vacuum chamber. BPMs are electromagnetic sensors that non-destructively couple to the electromagnetic field generated by the passing beam. The analogue signals produced by these devices are first processed via analogue front-end electronics located within the LHC tunnel. Subsequently, this processed output is transmitted to the surface-level back-end electronics, where the beam position information is digitised.

Data from each BPM is acquired independently through a dedicated platform, resulting in a large network of devices streaming data at 50 Hz. To mitigate the cost and complexity of the computing infrastructure, nearby BPM acquisition boards are controlled by a shared CPU. This arrangement effectively reduces the number of logical devices recognised by the control infrastructure to 70.

The BPM system is crucial for the LHC operation. Besides continuously streaming the average beam position data, the system offers several other functionalities. For instance, it can capture positions calculated for a selected subset of LHC bunches over a specified number of consecutive turns.

Software

Software Architecture & Technology Evolution

\* javier.martinez.samblas@cern.ch

# CONTROLS OPTIMIZATION FOR ENERGY EFFICIENT COOLING AND VENTILATION AT CERN

D. Monteiro\*, N. Bunjevac, R. Barillere, I. Rühl, CERN, Meyrin, Switzerland

## Abstract

Cooling and air conditioning systems play a vital role for the operation of the accelerators and experimental complex of the European Organization for Nuclear Research (CERN). Without them, critical accelerator machinery would not operate reliably as many machines require a fine controlled thermodynamic environment. These operation conditions come with a significant energy consumption: about 12 % (75 GWh) of electricity consumed by the Large Hadron Collider (LHC) during a regular run period is devoted to cooling and air conditioning. To align with global CERN objectives of minimizing its impact on the environment, the Cooling and Ventilation (CV) group, within the Engineering Department (EN), has been developing several initiatives focused on energy savings. A particular effort is led by the automation and controls section which has been looking at how regulation strategies can be optimized without requiring costly hardware changes. This paper addresses projects of this nature, by presenting their methodology and results achieved to date. Some of them are particularly promising, as real measurements revealed that electricity consumption was more than halved after implementation. Due to the pertinence of this effort in the current context of energy crisis, the paper also draws a careful reflection on how it is planned to be further pursued to provide more energy-efficient cooling and ventilation services at CERN.

## INTRODUCTION

The European Organization for Nuclear Research (CERN) stands as a global leader in scientific research, known for its groundbreaking experiments in particle physics. At the heart of CERN's ambitious pursuits lie massive accelerators and experimental facilities, where the quest for scientific discovery demands precision and reliability. Among the critical components that enable these endeavors, cooling and air conditioning systems play an indispensable role, creating and maintaining finely controlled thermodynamic environments necessary for the functioning of advanced accelerator machinery. However, this necessity comes at a substantial electricity consumption: 12 % of energy consumed by the Large Hadron Collider (LHC) is dedicated to cooling and air conditioning. In times marked by growing environmental consciousness and energy rising costs, mitigating this significant energy footprint is paramount to aligning CERN's operations with its broader objectives. In a report issued in 2021, CERN's management has openly declared its ambition "to establish itself as the model for transparent and environmentally responsible research organisations", which

compromises pursuing "actions and technologies aiming at energy saving and reuse" [1].

In this context, the Cooling and Ventilation (CV) group, of CERN's engineering department (EN), has embarked on a series of strategic initiatives aimed at realizing substantial energy savings within its domain. These initiatives range from energy-efficient mechanical design considerations, maintenance and operation practices, as well as automation and controls strategies that seek reducing the systems' energy footprint while preserving the demanding performance requirements. This paper particularly addresses controls optimization measures, by exposing the methodology used, project examples and results achieved to date.

The paper starts by exposing a general perspective on energy saving approaches for cooling and ventilation plants. After, the controls optimization method is introduced, and two application examples of this method in the context of energy savings are discussed. A summary of the work motivation, main results, and light on future work are provided at the end.

## ENERGY SAVING MEASURES

Effective efforts for energy-efficient cooling can be addressed at several stages of a plant's life cycle. For new systems, early design considerations accounting for the energy footprint of the future plant are of paramount importance. To assist design and project engineers on this process, energy standards and regulatory frameworks can be used. Such reference documents, often created and regularly updated by professional associations with expertise in the domain - see, for instance, ASHRAE Standard 90.1 [2] - expose the best practices in terms of design, equipment selection, control and operations of the system with energy cost in view. Equally important is the realistic definition of performance requirements, which are the main driver for the plant design. The use of conservative requirements is a common practice in the sector, and this often resolves in the installation of over-sized systems that are unnecessarily energy intensive.

During the operating phase of the plant, proper maintenance and parametrization are key for a good energetic performance. An effective maintenance plan shall aim at preserving the plant's operational performance anticipated during the design phase and verified during its commissioning. As an example, increase of the pumping power is observed when plate heat-exchangers are not cleaned regularly.

Operating a plant also includes adjusting the working parameters - such as temperature or pressure set-points - according to the time-dependent user needs. It is imperative to define these parameters based on the most objective re-

\* diogo.monteiro@cern.ch

# THE EMBEDDED MONITORING PROCESSOR FOR HIGH LUMINOSITY LHC

P. Moschovakos\*, V. Ryjov, S. Schlenker, CERN, Geneva, Switzerland  
D. Ecker, University of Wuppertal, Wuppertal, Germany  
J. B. Olesen, Aarhus University, Aarhus, Denmark

## Abstract

The Embedded Monitoring Processor (EMP) is a versatile platform designed for High Luminosity LHC experiments, addressing the communication, processing, and monitoring needs of diverse applications in the ATLAS experiment, with a focus on supporting front-ends based on lpGBT (low power Giga-Bit Transceiver, a CERN-built radiation-hard ASIC). Built around a commercial System-on-Module (SoM), the EMP architecture emphasizes modularity, flexibility and the usage of standard interfaces, aiming to cover a wide range of applications and facilitating detector integrators to design and implement their specific solutions. The EMP software and firmware architecture comprises epos (the EMP operating system), quasar OPC UA servers, dedicated firmware IP cores and an ecosystem of different software libraries. This abstract outlines the software and firmware aspects of the EMP, detailing its integration with lpGBT optical interfaces, programmable logic development, and the role of the LpGbtSw library as a Hardware Abstraction Library for the LpGbt OPC UA server.

*Keywords:* system-on-module, lpGBT, OPC UA, slow control, quasar, modular

## EMBEDDED MONITORING PROCESSOR CONTEXT

### Introduction

In the context of high-energy physics experiments, especially with the High Luminosity upgrade of the Large Hadron Collider (LHC), the requirements for enhanced radiation-tolerant Detector Control Systems (DCS) are high. To address the emerging challenges, the Embedded Monitoring Processor (EMP) and the Embedded Monitoring and Control Interface (EMCI) CERN-made boards were introduced.

Designed to operate in harsh radiation environments near the detector, the EMCI [1] board serves as the frontend component, facilitating the flow of slow control data between multiple frontend electronic devices and with the experiment's control system.

The EMP [2], in contrast to the EMCI, serves as a baseboard for the TE0807 MPSoC Modules from Trenz Electronic [3] which feature the AMD Zynq [4] UltraScale+ MPSoC. This versatile platform is designed primarily for controlling and monitoring LHC experiments that utilize lpGBT-based frontends, such as EMCI. Residing in a low-radiation service area, the EMP functions as the backend that processes and exchanges frontend data with the DCS.

\* paris.moschovakos@cern.ch

Moreover, a single EMP can connect with up to 12 lpGBT-based frontend boards via optical links, interfacing them to the Supervisory Control and Data Acquisition (SCADA) system via Ethernet, as illustrated in Fig. 1.

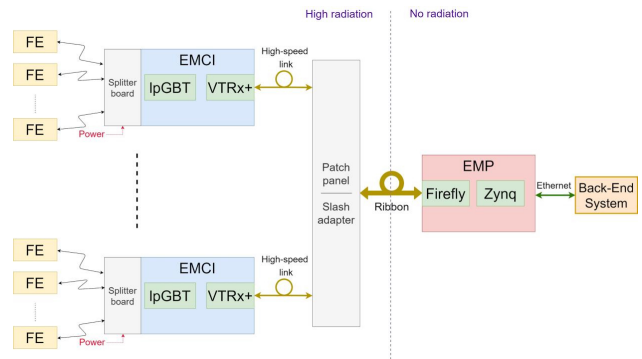


Figure 1: The EMP-EMCI system diagram.

### The Modular Design of EMP

In its broader context, the EMP is more than a hardware entity. Its full functionality is realized with the integration of tailored software for the Zynq's Processing System, firmware for the Zynq's Programmable Logic, and hardware extensions. This approach allows the EMP to serve to a wide range of monitoring and control applications.

Furthermore, the platform enables a set of functional capabilities, organized into separate application paradigms. These paradigms, discussed in detail in the following section, can coexist and function concurrently within a single EMP unit, provided certain technical conditions, like the availability of logic resources is sufficient.

### Interface Spectrum

In the context of ATLAS detector control systems, the low-power Gigabit Transceiver (lpGBT) [5] plays a crucial role, particularly its Internal Control (IC) communication channel. This channel is pivotal for the configuration, monitoring, and control of lpGBT. Elinks, electronic links that connect frontend electronics, are especially relevant in this context. These elinks are used in interfacing the frontends, as indicated in the preceding Fig. 1. Beyond its role in configuration, the lpGBT IC communication offers a wide array of I/O interfaces for slow control. In this context, the standard firmware IP core used for lpGBT communication, known as the lpGBT FPGA IP core, is used within the PL of the Zynq.

Moreover, the lpGBT facilitates data transmission using custom protocols, with employing the available elinks. To



# SELECTING A LINUX OPERATING SYSTEM FOR CERN ACCELERATOR CONTROLS

A. Radeva, F. Locci, J.M. Elyn, T. Oulevey, M. Vanden Eynden  
CERN, Geneva, Switzerland

## Abstract

Changing the operating system (OS) for large heterogeneous infrastructures in the research domain is complex. It requires great effort to prepare, migrate and validate the common generic components, followed by the specific corner cases. The trigger to change OS mainly comes from Industry and is based on multiple factors, such as OS end-of-life and the associated lack of security updates, as well as hardware end-of-life and incompatibilities between new hardware and old OS. At the time of writing, the CERN Accelerator Controls computing infrastructure consists of 4000 heterogeneous systems (servers, consoles and front-ends) running CentOS 7. The effort to move to CentOS 7 was launched in 2014 and deployed operationally 2 years later. In 2022, a project was launched to select and prepare the next Linux OS for Controls servers and consoles. This paper describes the strategy behind the OS choice, and the challenges to be overcome in order to switch to it within the next 2 years, whilst respecting the operational accelerator schedule and factoring in the global hardware procurement delays. Details will be provided on the technical solutions implemented by the System Administration team to facilitate this process. In parallel, whilst embarking on moving away from running Controls services on dedicated bare metal platforms towards containerization and orchestration, an open question is whether the OS of choice, RHEL9, is the most suitable for the near future and if not what are the alternatives?

## CONTROLS COMPUTING INFRASTRUCTURE

From a computing perspective, the CERN Control System is structured across three physical layers (Fig. 1):

1. The top (or client) tier consists of computers installed in the CERN Control Center (CCC), used by operations teams and equipment experts to run high-level graphical applications for accelerator control.
2. The middle (or business) tier, is comprised of powerful, high-availability servers, running the core control systems services which the high-level applications interact with.
3. The lower (or Front-End) tier is made of embedded computers (FECs) that execute real-time applications, interfacing with electronic boards to control and monitor accelerator devices.

Configuration aspects are handled by a central Controls Configuration Service (CCS) which is built around a relational database (CCDB) [1].

Software

Software Best Practices

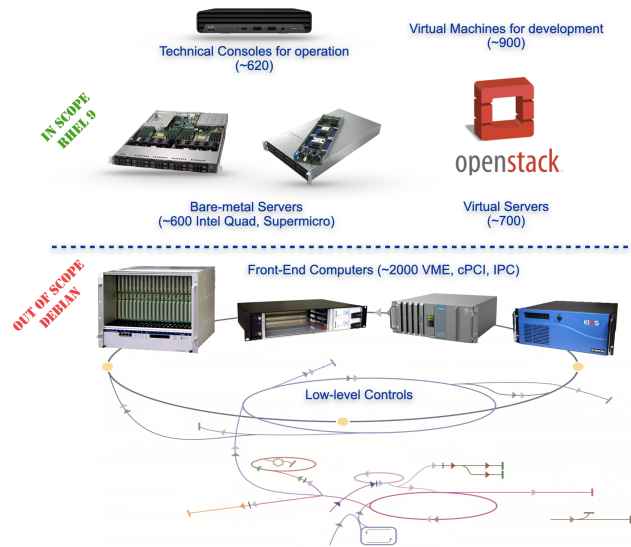


Figure 1: Computing layers in CERN's Control System.

## OPERATING SYSTEM LIFE CYCLE

To date, all computers within CERN's accelerator control system are based on a single Linux Operating System (OS). Currently, this is the CERN Community Enterprise Operating System (CC), which is based on CentOS Linux 7, which in-turn, is a downstream derivative of Red Hat Enterprise Linux (RHEL) 7. CERN's IT department closely follows the Red Hat OS life cycle (Fig. 2), prepares the corresponding CERN-specific distribution, and provides upstream support.



Figure 2: Red Hat Enterprise Linux Life Cycle.

The IT department communicates with key industry and Open Source Software Community actors for guaranteeing that the necessary specific packages are rebuilt and made available for the experts in CERN's Accelerators and Technology Sector.




Figure 3: CERN's Long-term Accelerator Schedule.

# UNIFIED SOFTWARE PRODUCTION PROCESS FOR CERN'S CRYOGENIC CONTROL APPLICATIONS

M. Pezzetti, T. Barbe, C. Fluder, T. Kubla, A. Tovar-Gonzalez  
 CERN, Meyrin, Switzerland  
 Sebastian Jan Rog, AGH, Cracow, Poland

## Abstract

The software engineering of process control system for CERN cryogenic installations is based on an automatic code production methodology and continuous integration practice. This solution was initially developed for the LHC Accelerator applications, then adapted to LHC Detectors, test facilities and non-LHC cryogenic facilities. Over the years, this approach allowed the successful implementation of many control system upgrades, as well as the development of new applications, while improving quality assurance and minimizing manpower resources. The overall complexity of automatic software production chains, their challenging maintenance, deviation between software production methods for different cryogenic domains and frequent evolution of CERN frameworks led to the system's complete review. A new unified software production system was designed for all cryogenic domains and industrial technologies used. All previously employed frameworks, tools, libraries, code templates were classified, homogenized and implemented as common submodules, while projects specific configuration were grouped in custom application files. This publication presents the new unified software production solution, benefits from shared methodology between different cryogenics domains, as well as a summary of two years of experience with several cryogenic applications from different PLCs technologies.

## INTRODUCTION

Large cryogenic systems are an integral part of CERN's accelerator complex and experimental facilities, especially the Large Hadron Collider (LHC) with its detectors [1]. The 24h a day continuous cryogenic operation during the physics campaign is mandatory to the collider's stable beam operation. As any large and complex technological system, the cryogenic installations process control programs running on an industrial control PLC must be as robust and error-free as possible. At CERN, the design and the development of the process control software is performed through the CERN UNICOS framework [2], which defines a set of conceptual objects such as valves, heaters, IOs, alarms, PID controllers, etc. Instead of writing PLC code manually, the UCPC (UNICOS Continuous Process Control) code generator creates the code based on a specification file and Python templates. For several years, CERN has developed various solutions to help with the whole lifecycle of a cryogenic process control system [3]. Nowadays, CERN's systematic use of version control (with git) and continuous integration (CI) in the engineering workflow has significantly reduced the necessary

time (from months to weeks) dedicated to develop, test and deploy a new application or an upgrade. This has allowed the production of quality and highly reliable cryogenic control and electrical systems [4].

## CERN CRYO CI SYSTEM HISTORY

The first control system using continuous integration was the 18 PLCs LHC Cryogenic Tunnel Applications, followed by cryogenic magnet test benches using the Siemens PLCs technology [3]. Additionally, a database-oriented information system was chosen for the LHC tunnel Cryogenic process control in the project's early development [5]. This choice has enabled the successful use of the continuous integration system toward a higher automatization level. It allowed to improve development speed and reduce the time in feedback loops. After a successful use of the CI in these two cases, it was progressively rolled out to all cryogenic process control systems at CERN. This includes LHC Detectors ATLAS and CMS [6] which are controlled by Schneider PLCs technology. The last category consists of several non-LHC and test facilities installations that use a mix of Siemens and Schneider PLCs. Each of these domains had their own special requirements that lead to the development of slightly different CI solutions (cf. Fig. 1). After some time, this became very hard to maintain and the system had to be redesigned.

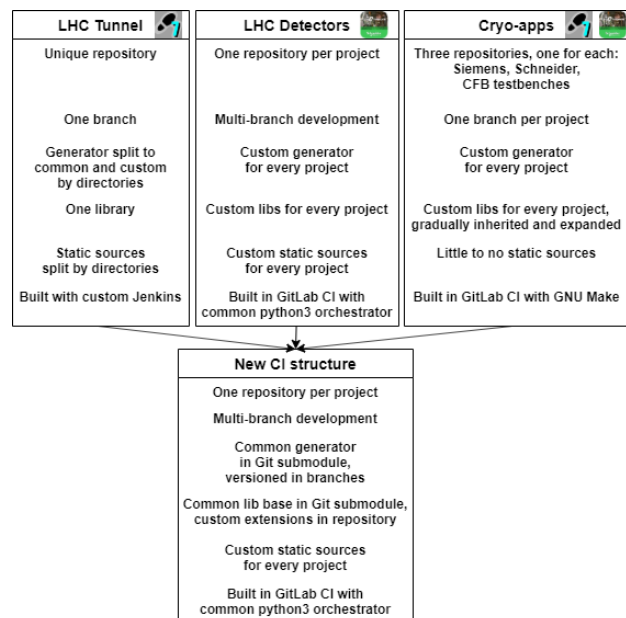


Figure 1: Evolution of CI systems.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# VISUALIZATION TOOLS TO MONITOR STRUCTURE AND GROWTH OF AN EXISTING CONTROL SYSTEM

O. Pinazza<sup>†,1,2</sup>, A. Augustinus<sup>2</sup>, P. M. Bond<sup>2</sup>, P. Chochula<sup>2</sup>,  
A. Kurepin<sup>3</sup>, M. Lechman<sup>2,4</sup>, D. Voscek<sup>2</sup>

<sup>†</sup>corresponding author; <sup>1</sup>INFN, Sezione di Bologna, Italy; <sup>2</sup>CERN, Geneva (CH);

<sup>3</sup>affiliated with an Institute Covered by a Cooperation Agreement with CERN, Geneva, (CH);

<sup>4</sup>University of the Witwatersrand, Johannesburg, South Africa

## Abstract

The ALICE experiment has been in operation for 15 years at the LHC. During its life several detectors have been replaced, new instruments installed, and readout technologies have changed. The control system has therefore also had to adapt, evolve and expand, sometimes departing from the symmetry and compactness of the original design.

ALICE is a large collaboration, where different Institutes and Universities from over 40 Countries contribute to the development of the detectors and their control systems. For the central coordination it is important to maintain the overview of the integrated control system to assure its coherence. Tools to visualize the structure and other critical aspects of the system can be of great help and can highlight problems or features of the control system such as deviations from the agreed architecture.

This paper describes how existing tools, such as graphical widgets available in the public domain, or techniques typical of scientific analysis, can be adapted and help assess the coherence of the development, revealing hidden weaknesses and highlighting the interdependence of parts of the system.

## INTRODUCTION

ALICE [1] is one of the four big experiments at the LHC. Operational since LHC started in 2007, ALICE has taken proton and heavy-ions collisions data during Run1 (2009-2013) and Run2 (2015-2018). During Long Shutdown2 (LS2, 2018-2022) ALICE has undergone an important upgrade consisting of the installation of new detectors: Inner Tracking System (ITS), Muon Forward Tracker (MFT), Fast Timing triggering detectors (FDD, FT0 and FV0); a new readout system for the main Time Projection Chamber (TPC) detector, and a brand new software model called O<sup>2</sup>, combining offline and online systems [2].

The ALICE Detector Control System (DCS) [3] has successfully accompanied these developments and adapted to new technologies and requirements, while at the same time ensuring the continuity of operations and control of different devices.

In order to verify the coherent and safe development of the different systems and the integration of new detectors, we have introduced data analysis and visualization techniques, which also enabled us to assess the interdependence and safety of the systems.

## THE ALICE CONTROL SYSTEM

The ALICE DCS is strongly based on SIMATIC WinCC Open Architecture (WinCC) [4], as well as the other CERN experiments, but it is not limited to it. The WinCC software core is running on 63 worker nodes connected to the 15 ALICE detectors, interfaced by 17 WinCC multiuser operator nodes, where experts can login to access monitoring and control interfaces; in the backend, 19 more nodes host the WinCC central services, 20 linux nodes run drivers and custom software, together with a database farm, several file servers, gateways and other specialized nodes, and several different embedded systems.

While WinCC offers a wide choice of graphical widgets allowing to develop fancy panels, to monitor the detectors' hardware and control their devices, it lacks sometimes tools and interfaces to represent the connection and interdependence between different operating system and to unveil the complexity of the overall DCS.

This type of information can help in revealing unexpected interdependencies and weaknesses, or design flaws. It is characterized by being rather static, compared to the mutability of typical DCS parameters, and can be produced whenever required, occasionally, from an external host.

The data presented here has been obtained after the LS2 upgrade. Some of the images were inspired by books on the graphic representation of hierarchical data using trees or spheres [5]. All data is extracted with *python*, Microsoft Powershell and WinCC CTRL++, and visualized with graphical libraries for *python* and *javascript*, available online [6, 7]. It will be interesting to observe the variations in time, especially when new subsystems will be added to the DCS during LS3, after the next big ALICE upgrade [8].

## PROJECTS INTERDEPENDENCE

One of the features of WinCC is the capability to interconnect systems through the DIST managers, thus allowing access to datapoints in a remote project. While the DIST mechanism is overall rather efficient, its monitoring and control is more obscure. It can happen that, under some special circumstances like a timeout during a project restart, or systems restarting in random order after a general reboot, two or more WinCC projects that are supposed to be connected, in reality are not. This situation can stay unnoticed and be difficult to emerge. To limit this kind of problems and prevent the malfunction of a system from

# TOWARDS A FLEXIBLE AND SECURE PYTHON PACKAGE REPOSITORY SERVICE

I. Sinkarenko<sup>†</sup>, P. Elson, F. Iannaccone, W. Koorn, B. Copy, CERN, Geneva, Switzerland

## Abstract

The use of third-party and internal software packages has become a crucial part of modern software development. Not only does it enable faster development, but it also facilitates sharing of common components, which is often necessary for ensuring correctness and robustness of developed software. To enable this workflow, a package repository is needed to store internal packages and provide a proxy to third-party repository services. This is particularly important for systems that operate in constrained networks, as is common for accelerator control systems.

Despite its benefits, installing arbitrary software from a third-party package repository poses security and operational risks. Therefore, it is crucial to implement effective security measures, such as usage logging, package moderation and security scanning. However, experience at CERN has shown off-the-shelf tools for running a flexible repository service for Python packages not to be satisfactory. For instance, the dependency confusion attack first published in 2021 has still not been fully addressed by the main open-source repository services.

An in-house development was conducted to address this, using a modular approach to building a Python package repository that enables the creation of a powerful and security-friendly repository service using small components. This paper describes the components that exist, demonstrates their capabilities within CERN and discusses future plans. The solution is not CERN-specific and is likely to be relevant to other institutes facing comparable challenges.

## INTRODUCTION

In recent years, the use of Python within CERN's accelerator control system has seen significant growth. The adoption of Python as a supported language beside Java triggered the rise of new software, spanning from operational and expert GUIs based on PyQt [1], high level control system APIs, offline data analysis based on PySpark, to more recent online optimisation of operations using Machine Learning [2].

Now with a community of over 500 users, it is essential to provide effective support and ensure a stable and smooth user experience, based on practices and tools discussed in [3]. Such a service is undertaken by a dedicated centralised team, where small efficiency improvements can have a significant cumulative effect due to the large number of beneficiaries. One of the encouraged practices is for users to develop Python packages, rather than scripts, (for the sake of versioning, testability, and reuse), and the wide use of virtual environments (to avoid dependency collisions),

both of which stimulate frequent installation and publishing of Python packages to a package repository. For this reason, the “Acc-Py Package Repository” is one of the most crucial services maintained by the team.

As an organisation with mission-critical, expensive, and sensitive hardware, CERN operates distinct networks separated physically and by firewall. A general-purpose network with Internet access supports everyday use, and a more restricted network without external access connects accelerator-related hardware. There is a need to install Python packages in both networks, meaning that the Acc-Py Package Repository must provide access to third-party packages from the Python Package Index (PyPI) at [pypi.org](https://pypi.org).

The use-cases of CERN's accelerator control system are surely not unique, and similar scenarios probably exist in other laboratories, and in wider industry. Therefore, this paper aims to share the ideas and implemented solutions and invites readers to contact the authors, especially in the case of an interest to contribute to further developments under an open-source license. The functionality has been presented at the Europython conference [4], which confirmed the assumption of shared challenges, and initial versions of the software have been published in [5].

## ACC-PY PACKAGE REPOSITORY

The Acc-Py Python Package Repository runs as a small set of microservices based on the FastAPI web framework. The most important ones are “simple-repository” that provides the index of packages for a package installer to consider and retrieve, and “simple-repository-upload” responsible for receiving uploaded in-house packages. A Web User Interface (Web UI) or “simple-repository-browser” provides a familiar user interface to discover packages and their metadata.

Many requirements for the API microservices are driven by the client-side de-facto tools of the Python Packaging community, namely pip as the package installer, and twine as the package uploader. In turn, this implies wheels as the standard binary package format, which is the preferred installation format, with source distributions, sdist, for other cases and fallbacks. Quality of service is continuously being improved by following the global evolution of Python packaging standards developed through the Python Enhancement Proposal (PEP) process, in a way that is compatible with both current and future versions of these client-side tools.

Despite there being a single endpoint for package downloads, internally “simple-repository” is a representation of two distinct repositories: a local one for in-house packages, and the public PyPI repository. The pip installer is pre-configured to communicate directly with the Acc-Py Package Repository service instead of attempting

# IMPLEMENTING HIGH PERFORMANCE & HIGHLY RELIABLE TIME SERIES ACQUISITION SOFTWARE FOR THE CERN-WIDE ACCELERATOR DATA LOGGING SERVICE

M.W. Sobieszek, J. Woźniak, R. Mucha, P. Sowiński,  
V. Baggiolini, C. Roderick, CERN, Geneva, Switzerland

## Abstract

The CERN Accelerator Data Logging Service (NXCALS) stores data generated by the accelerator infrastructure and beam related devices. This amounts to 3.5TB of data per day, coming from more than 2.5 million signals from heterogeneous systems at various frequencies. Around 85% of this data is transmitted through the Controls Middleware (CMW) infrastructure. To reliably gather such volumes of data, the acquisition system must be highly available, resilient and robust. It also has to be highly efficient and easily scalable, as data rates and volumes increase, notably in view of the High Luminosity LHC.

This paper describes the NXCALS time series acquisition software, known as Data Sources. System architecture, design choices, and recovery solutions for various failure scenarios (e.g., network disruptions or cluster split-brain problems) are covered. Technical implementation details are discussed, covering the clustering of Akka Actors collecting data from tens of thousands of CMW devices. The NXCALS system has been operational since 2018 and has demonstrated the capability to fulfil all aforementioned requirements, while also ensuring self-healing capabilities and no data losses during redeployments.

## INTRODUCTION

The CERN Accelerator Data Logging Service (NXCALS) stores data generated by the accelerator infrastructure and beam related devices, and in-turn makes this data available to the CERN community [1].

The NXCALS architecture is composed of three main subsystems (Figure 1): data acquisition and ingestion, data storage and compaction, and APIs or applications for data extraction. This paper focuses on the first subsystem, also known as NXCALS Data Sources (DS).




Figure 1: Overview of the NXCALS architecture.

The DS acquire data from the accelerator devices through the Controls Middleware (CMW) [2] using a subscription mechanism. One subscription corresponds to a data channel through which the device publishes data of one of its attributes (properties) to the DS. All information needed to configure the data sources and the subscriptions are managed within the central Controls Configuration Service (CCS) [3] and stored in the Controls Configuration Database (CCDB). Subscriptions are configured using a web interface (Controls Configuration Data Editor – CCDE), directly by end users (e.g. device experts, physicists, or operations teams) who know which data they want to be stored in NXCALS.

The NXCALS DS have to fulfil many requirements, but the most important is to reliably acquire huge amounts of data from accelerator infrastructure and beam related devices and transmit it to the NXCALS storage system without any data loss.

This is a surprisingly difficult requirement to fulfil. It needs a thorough analysis of the common challenges encountered when developing distributed systems.

The main concepts that must be addressed are:

- **Fault Tolerance and Resilience** - designing mechanisms to handle process failures (e.g., out of memory crashes) or network failures and partitions without service disruption.
- **Load Balancing** - distributing subscriptions evenly across processes to prevent overloading specific machines and their network interfaces.
- **Scalability** - being able to handle an increased number of subscriptions just by adding additional resources.
- **Security and Authentication** - implementing appropriate measures to protect data and prevent unauthorised access.
- **Monitoring and Debugging** - obtaining detailed information about the system performance and behaviour, especially when issues arise.

The rest of the paper will focus on how the NXCALS DS deal with the first three challenges in the list above.

## SYSTEM OVERVIEW AND HIGH-LEVEL ARCHITECTURE

Figure 2 shows a high-level overview of the DS subsystem. Conceptually, the sequence of tasks executed by a data source process is straightforward:

1. The DS begins by retrieving subscriptions metadata from the configuration database.

# A GENERIC REAL-TIME SOFTWARE IN C++ FOR DIGITAL CAMERA-BASED ACQUISITION SYSTEMS AT CERN

A. Topaloudis\*, E. Bravin, S. Burger, S. Jackson, S. Mazzone, E. Poimenidou, E. Senes  
CERN, Geneva, Switzerland

## Abstract

Until recently, most of CERN's beam visualisation systems have been based on increasingly obsolescent analogue cameras. Hence, there is an on-going campaign to replace old or install new digital equivalents. There are many challenges associated with providing a homogenised solution for the data acquisition of the various visualisation systems in an accelerator complex as diverse as CERN's. However, a generic real-time software in C++ has been developed and already installed in several locations to control such systems. This paper describes the software and the additional tools that have also been developed to exploit the acquisition systems, including a Graphical User Interface (GUI) in Java/Swing and web-based fixed displays. Furthermore, it analyses the specific challenges of each use-case and the chosen solutions that resolve issues including any subsequent performance limitations.

## INTRODUCTION

A Beam Observation System, referred to as BTV (Beam TV) at CERN, plays a crucial role in capturing beam images throughout the accelerator complex. This is accomplished by intercepting the trajectory of the beam within the vacuum chamber using a screen. When the beam particles interact with the screen, they emit visible light in direct proportion to their local intensity. Subsequently, a dedicated detector, such as a camera, can be employed to observe the resulting footprint of the beam through a specialized viewport and optical pathway [1]. Until recently, only analogue cameras have been used in the CERN accelerator complex, which were either Charge-Couple Device (CCD) cameras or vidicon tubes [2]. As CCD and vidicon camera technologies continue to become outdated, maintaining the current beam visualisation systems at CERN becomes progressively challenging. This issue is compounded by the need for extra cameras due to the growing interest in deploying beam observation systems in new areas throughout the facility.

It was therefore decided to use the GigE digital camera models from Basler [3] with Complementary Metal Oxide Semiconductor (CMOS) sensors for all new installations, as well as to gradually replace the obsolete analogue equipment in existing locations if the opportunity arises.

Digital cameras are proven to have better image quality (signal-to-noise ratio, dynamic range, sharpness) than their analogue predecessors [4]. In addition, there is no need for additional dedicated hardware for the Analogue-to-Digital Conversion (ADC) and the synchronization of the image acquisition with the beam passage.

\* athanasios.topaloudis@cern.ch

## INSTALLATIONS

Digital cameras have already been installed in several locations at CERN, thus requiring the specification of a new acquisition system. This system aims to combine legacy capabilities along with additional functionalities taking advantage of, where possible, new features offered by digital cameras compared to the analogue ones.

### *SPS Beam Dump System (SBDS)*

A BTV monitor was installed at the location of the new SBDS to ensure the safe operation of the system and its components by capturing an image of the particles before impacting the dump target. In this way, the dumped beams, including their shape and precise position in relation to the target, can be continuously monitored and subsequent injections can be inhibited in case of operational problems [5].

Since this system is linked to the safe operation of the SPS, it should always be online. Furthermore, to ensure the best possible image is provided for detecting issues with the dump, it should continuously acquire images. Then, upon receiving a timing event (e.g. that the beam was dumped) the first non-saturated image should be selected, published and stored for post-mortem analysis.

### *Advanced Wakefield Experiment (AWAKE)*

The AWAKE experiment depends strongly on its imaging systems in order to be able to operate. They are used to measure the exact shape and position of the various beams (i.e. proton, electron and laser) so that they can be properly and efficiently aligned [6]. As a consequence, the images should be constantly acquired at the laser's repetition rate (10 Hz) which is the fastest rate of the three beams. Additionally, the large sensors of AWAKE's cameras result in large images, thus large data throughput [7].

Despite the images being acquired at a rate of 10 Hz, during normal operation, only the image capturing the SPS proton extraction is of interest to the experiment. The synchronization of the image acquisition and the SPS extraction events is therefore crucial.

Finally, AWAKE is a harsh environment for the digital cameras as many are installed very close to the beam lines and are very susceptible to Single-Event-Upset (SEU) underscoring the need for a robust recovery mechanism.

### *CERN Linear Electron Accelerator for Research (CLEAR)*

Similar to AWAKE, CLEAR also depends on several cameras for its efficient operation. Since the facility is very versatile, accommodating a variety of experiments during

# BUILDING, DEPLOYING AND PROVISIONING EMBEDDED OPERATING SYSTEMS AT PSI

D. Anicic, Paul Scherrer Institut, Villigen, Switzerland

## Abstract

In the scope of the Swiss Light Source (SLS) upgrade project, SLS 2.0, at Paul Scherrer Institute (PSI) two New Processing Platforms (NPP), both running RT Linux, have been added to the portfolio of existing VxWorks and Linux VME systems. At the lower end we have picked a variety of boards, all based on the Xilinx Zynq UltraScale+ MPSoC. Even though these devices have less processing power, due to the built-in FPGA and Real-time CPU (RPU) they can deliver strict, hard RT performance. For high-throughput, soft-RT applications we went for Intel Xeon based single-board PCs in the CPCI-S form factor. All platforms are operated as diskless systems. For the Zynq systems we have decided on building in-house a Yocto Kirkstone Linux distribution, whereas for the Xeon PCs we employ off-the-shelf Debian 10 Buster. In addition to these new NPP systems, in the scope of our new EtherCAT-based Motion project, we have decided to use small x86\_64 servers, which will run the same Debian distribution as NPP. In this contribution we present the selected Operating Systems (OS) and discuss how we build, deploy and provision them to the diskless clients.

## INTRODUCTION

At PSI we operate four accelerator facilities: HIPA and PROSCAN proton facilities, and SLS and SwissFEL electron facilities. They age from few years up to several decades. A variety of hardware, computers, operating and control systems have been used during this time.

Hardware-wise mostly VME based systems are used, but also quite a lot of PC based systems and a variety of commercial small computer boxes. We also use many virtualized systems.

On the software side we have VxWorks, Windows, Scientific Linux, Red Hat Enterprise Linux, and a diversity of Embedded Linux Systems, all in several versions.

Some OS-s have been installed on local hard disk, some others are network-booted. Experience shows that network-based provisioning is simpler for updates and changes.

## NEW PLATFORMS

Three new platforms are supported:

- Zynq Ultrascale+ based computers
- x86\_64 single board computers (SBC) in CPCI-S
- Small x86\_64 servers

### *Zynq Ultrascale+ Based Computers*

The decision to use Zynq Ultrascale+ platform was taken already some time ago. Several tests were performed, and three types were envisioned, for small, medium and high

performance and/or power consumption. But because several internal groups have been involved in the development it actually resulted in almost ten different configurations. Some groups decided to go forward with ready available development boards bought directly from manufacturers, whilst others took either a Silicon-on-Chip (SoC) or Silicon-on-Module (SoM) approach, to develop their own boards. Unlike we initially assumed, this variety actually turned out not to be a problem.

### *x86\_64 SBC*

For some Zynq Ultrascale+ systems the CPCI-S bus was targeted as our crate and bus standard, providing the possibility, where needed, to create more powerful, x86\_64 based systems, as CPCI-S bus controller and also as number cruncher, if needed. Presently we have only one board type with Intel Xeon CPU, dual Ethernet connection, and 16GB of RAM.

### *Small x86\_64 Servers*

For the Motion project (motors) based on EtherCAT, our initial design involved installing operating system locally. With the operating system for x86\_64 NPP already in place, the decision was taken to re-use it also for the small servers. This increases the synergies between our groups, saves manpower and, will simplify maintenance. Although any kind of PC could be used, our great wish was that everybody would use the same small HP DL20 servers. This would simplify troubleshooting in case of problems and provide for common replacement stock.

## CHOOSING AND BUILDING OPERATING SYSTEMS

Generally, we will use Linux for all our future projects. Our wish was also to use a RealTime (RT) patched Linux kernel.

### *Zynq Ultrascale+ Based Computers*

For the initial considerations we have been using Xilinx PetaLinux SDK. This was quite convenient as a proof of concepts, but it turned out to be too complex to handle, because our different board development groups have adopted different versions. The development also took longer than planned, new hardware revisions have been coming, and the need for newer versions of SDK were occurring. It was not possible to easily satisfy all development needs by supporting several versions.

The RT kernel patch was also not available for all kernels provided by PetaLinux SDK.

For a functioning kernel, the Linux kernel device-tree, which describes hardware-to-kernel interface, must be provided for each board type. Unfortunately the device-tree

# APPLICATION DEVELOPMENT ON CPCI-S.0 HARDWARE AT PSI

I. Johnson\*, R. Biffiger, D. Felici, W. Koprek, R. Rybaniec, B. Stef and G. Theidel  
Paul Scherrer Institut (PSI), CH-5232 Villigen PSI, Switzerland

## Abstract

A Hardware and Software Toolbox is being created to accelerate the engineering of electronic components for large facility upgrades at the Paul Scherrer Institute. This Toolbox consists of modular hardware following the CPCI-S.0 standard, project base designs, firmware libraries and software packages. The goal is to simplify and accelerate developments with a set of compatible electronics, starting foundations, tools and modules.

## INTRODUCTION

Typical applications at accelerator facilities require network controlled, low-latency (1  $\mu$ s) systems. The PICMG Compact PCI Serial standard, CPCI-S.0 [1], offers an ideal platform to interconnect these real-time systems. Large facility upgrades at the PSI, SLS 2.0 and HIPA, are moving to this flexible, crate-compatible standard. This paper provides an overview of the hardware, embedded system, software and firmware developments that enable the efficient creation of applications.

## CPCI-S HARDWARE

A crate with a CPCI-S.0 compatible backplane has been designed by PSI and is being manufactured by ELMA, Fig. 1. The backplane of this crate interconnects a system slot, 8 peripheral slots, rear slots and a power backplane which houses a system monitoring card, utility cards and power supplies. Multi-gigabit capable traces in both a full-mesh topology and star topology from the system slot provide routes for high-speed communication between the front slots. Ethernet and computer-compatible buses (PCIe, USB and SATA) are in the CPCI-S.0 pinout specification to enable standard-protocol data-transfer between the front cards. These multi-gigabit links can alternative be utilized for custom protocols. There is a rear slot behind each of the 8 peripheral slots. Multi-gigabit serial links, i2c and general purpose IO create a versatile connection between the front and rear cards. Rear slots extend the real-estate for analog and digital signal processing cards.

Adhering to CPCI-S.0 standard makes the crate also compatible with both commercial off-the-shelf cards. Companies like EKF offer a wide range of CPCI-S boards. Two commonly utilized commercial cards from EKF are the SC5 Festival (a High Performance CPU Board) and the SD1-DISCO (SATA Drive Carrier Board).

More than a handful of custom PSI cards have also already been produced:



Figure 1: Picture of CPSI-S crate with 5 custom (UFC) cards inserted into the left most peripherally slots, a System Monitoring card just right of the middle and 3 redundant power supplies on the right.

- CPSI\_UFC: Universal FMC+ Carrier with a Zynq™ UltraScale+ MPSoC, HPC FMC+ slot, 48 multi-gigabit links and plenty of high speed GPIOs, Fig. 2.
- CPSI\_CIO: Communication IO card with a Zynq™ UltraScale+ MPSoC for multi-gigabit, crate-external communication with two SFP+ and two QSFP+ slots, Fig. 3.
- CPSI\_RTM\_DAC: Rear Digital to Analog Converter card which contains two 16 bit 500 Msps DAC, two SFPs and an expansion connector.
- CPSI\_RTM\_FIO: Fixed Input/Output board with one SFP+, two 500 MHz SMA and eight LEMO connectors.
- CPSI\_CM1: Crate control and monitoring card.
- Power Load Board: Card for measuring the crate's power capabilities and stability.
- Backplane Loopback Board: Card for testing all the data connection on the backplane (Full-Mesh and Star).




Figure 2: Picture of the Universal FMC+ Carrier (CPSI\_UFC).

Some of these board are already being utilized in applications, see Section APPLICATIONS.

## EMBEDDED ENVIRONMENT

Every CPSI-S custom front card contains a PSoC 62 micro-controller from Infineon. The software running on this micro-controller is responsible for the start-up and shut-

Hardware

FPGA & DAQ Hardware

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

\* ian.johnson@psi.ch



# SciLog: A FLEXIBLE LOGBOOK SYSTEM FOR EXPERIMENT DATA MANAGEMENT

K. Wakonig<sup>†</sup>, A. Ashton, C. Minotti, Paul Scherrer Institut, Villigen PSI, Switzerland

## Abstract

Capturing both raw and metadata during an experiment is of the utmost importance, as it provides valuable context for the decisions made during the experiment and the acquisition strategy. However, logbooks often lack seamless integration with facility-specific services such as authentication and data acquisition systems and can prove to be a burden, particularly in high-pressure situations, for example, during experiments. To address these challenges, SciLog has been developed at the Paul Scherrer Institut. Its primary objective is to provide a flexible and extensible environment, as well as a user-friendly interface. SciLog relies on atomic entries in a database that can be easily queried, sorted, and displayed according to the user's requirements. The integration with facility-specific authorization systems and the automatic import of new experiment proposals enable a user experience that is specifically tailored for the challenging environment of experiments conducted at large research facilities. The system is currently in use during beam time at the Paul Scherrer Institut, where it is collecting valuable feedback from scientists to enhance its capabilities.

## INTRODUCTION

Metadata is defined as the data providing information about one or more aspects of the data; it is used to summarize basic information about data that can make tracking and working with specific data easier. [1] It includes, among many, information about the source of the data, its process and responsible people and the location on a computer network where the data was created and collected.

It also covers unstructured information collection, for example, notes on a data acquisition process or keeping track of important TODOs. SciLog [2, 3] was developed specifically with this in mind, namely to improve the storing and reuse of unstructured metadata and as a consequence improve the FAIRness of data [4]. It aims at replacing legacy pen and paper experimental logbooks, often used in large-scale facilities during beamtime and often messy to consult and prone to information loss. It also supports features to monitor an experiment, by watching messages posted by the beamline to it. Each message is an atomic entry in the database, which means that every message can be decoupled from the rest of the environment. It also supports collaborative editing.

First, we will introduce the components of SciLog and a few key concepts that are useful to better comprehend the rest of the article. Then we will move to present the creation of logbooks, the search functionality, the TODOs support and the main logbook widget, which enables

displaying, adding and modifying messages. The widget can display messages ordered by date and can provide information about, for example, the time of insertion and the author.

We will address how SciLog can be scaled to accommodate high volumes of metadata and usage.

In order to maximize data dissemination, we will present the Python [5] libraries that have been developed to interact with SciLog to post and get metadata.

We will finally close the article with future directions that the community envisions for SciLog.

## DESIGN OVERVIEW

The next sections discuss the technicalities of the implementation of SciLog, including the choice of the underlying technologies and frameworks.

The SciLog stack is organised following a microservice architecture, where each service can be containerised and configured to interact with the others and the pre-existing facility infrastructure, following standard TCP [6] communication protocols, such as HTTP(s) [7] and Web-Sockets [8]. All SciLog services communicate with each other through HTTP or Web-Sockets.

The backbone of the ecosystem is the *backend* which relies on a Mongo Database [9], the connection to which must be configured as part of the setup. The *backend* is also responsible for defining the data model which formalizes the scaffolding of the metadata fields, setting the required ones and leaving great flexibility for customization.

### Data Model

The data model defines which information is stored and how it is structured, by mapping the SciLog entities with records on the database.

Each SciLog entity has a representation in the data model in the underlying MongoDB and a subset of fields is controlled by validation rules imposed by the *backend*.

The majority of entities share the same common structure and fields subset, and they differ by adding entity-specific fields or mentioning the entity type specifically in one (*snippetType*). It is often convenient to store dependencies between SciLog entities in the database, and this can be achieved using the concept of one-to-many relationships between documents in MongoDB [10].

The main SciLog entities are:

- Basesnippets
- Locations
- Logbooks
- Images
- Paragraphs
- Tasks

<sup>†</sup> klaus.wakonig@psi.ch

# PHASE-II UPGRADE OF THE CMS ELECTROMAGNETIC CALORIMETER DETECTOR CONTROL AND SAFETY SYSTEMS FOR THE HIGH LUMINOSITY LARGE HADRON COLLIDER

R. Jiménez Estupiñán<sup>†</sup>, G. Dissertori, L. Djambazov, N. Härringer, W. Lustermann, K. Stachon,  
ETH Zurich, Zürich, Switzerland  
L. Cokic\*, CERN, Geneva, Switzerland  
P. Adzic, D. Jovanovic, M. Mijic, P. Milenovic, University of Belgrade, Serbia  
On behalf of the CMS Collaboration

## Abstract

The Electromagnetic Calorimeter (ECAL) is a subdetector of the CMS experiment. Composed of a barrel and two endcaps, ECAL uses lead tungstate scintillating crystals to measure the energy of electrons and photons produced in high-energy collisions at the Large Hadron Collider (LHC). The LHC will undergo a major upgrade during the 2026-2029 period to build the High-Luminosity LHC (HL-LHC). The HL-LHC will allow for physics measurements with one order of magnitude larger luminosity during its Phase-2 operation. The higher luminosity implies a dramatic change of the environmental conditions for the detectors, which will also undergo a significant upgrade. The endcaps will be decommissioned and replaced with a new detector. The barrel will be upgraded with new front-end electronics. A Sniffer system will be installed to analyse the airflow from within the detector. New high voltage and water-cooled, radiation tolerant low voltage power supplies are under development. The ECAL barrel safety system will replace the existing one and the precision temperature monitoring system will be redesigned. From the controls point of view, the final barrel calorimeter will practically be a new detector. The large modification of the underlying hardware and software components will have a considerable impact in the architecture of the detector control system (DCS). In this document the upgrade plans and the preliminary design of the ECAL DCS to ensure reliable and efficient operation during the Phase-2 period are summarized.

## INTRODUCTION

The control systems of the CMS ECAL detector were designed [1] before the first LHC collisions in 2009. It has been maintained throughout more than 15 years of operation. During this time, multiple components have been updated or replaced by more modern ones to extend the systems lifetime across multiple data-taking periods. However, many of these components will become obsolete by the next long shutdown. Hence, the importance of evaluating and anticipating the necessary hardware and software updates. In addition to this, two major detector partitions will be decommissioned during the next long shutdown: the ECAL endcaps (EE) and the preshowerers (ES). The HL-LHC will increase its luminosity by an order of magnitude,

entailing an increasing level of radiation in the experimental area. To extend the lifetime of the ECAL crystals, the remaining ECAL barrel calorimeter (EB) will be cooled down from a nominal temperature of 18°C to 9°C. The EB supermodules will be extracted to install new front-end electronics, also bringing the opportunity to prepare them from the controls point of view. These modifications, detailed in the Technical Design Report of the CMS Barrel Calorimeters [2], will have a considerable impact on the DCS architecture. The work has been organised in 5 different projects as follows:

1. The ECAL Barrel Safety System.
2. Controls for the low voltage (LV) power.
3. Software for the high voltage (HV) power.
4. The ECAL precision temperature monitoring.
5. The Phase-2 supervisory system.

## ECAL BARREL SAFETY SYSTEM

The CMS safety systems consist of multiple interconnected components, designed to detect and mitigate potential hazards within the experiment. They form a network of Programmable Logic Controllers (PLCs), serving as the backbone of the safety infrastructure. The ECAL Safety System (ESS) is one of these components, crucial in providing safety to the ECAL detector. The design of the ESS is being currently revamped to address the requirements of the future EB, after which it will be referred as the "ECAL Barrel Safety System" (EBSS). The EBSS will be responsible for safeguarding the detector by monitoring environmental conditions and interacting with a range of devices through its interlock system. The EBSS will be built strictly with industrial components, using the latest generation of Siemens equipment, and programmed with the CMS PLC software framework. The new design includes several improvements based on the accumulated experience of more than 15 years of operations and will serve as a reference for other detectors' safety system. The most distinctive features of this system are:

- Single Siemens S7-1500 series CPU design with distributed I/O modules on a PROFINET [3] ring, ensuring high reliability and availability.
- Redundant capabilities at the level of the sensors, cable pathways and connectors.

<sup>†</sup> jraul@phys.ethz.ch

\* Previously at the University of Belgrade, Serbia.

# STREAM-BASED VIRTUAL DEVICE SIMULATION FOR ENHANCED EPICS INTEGRATION AND AUTOMATED TESTING

M. Lukaszewski\*, K. Klys, E9 Controls Limited, London, UK

## Abstract

Integrating devices into the Experimental Physics and Industrial Control System (EPICS) can often take a suboptimal path due to discrepancies between available documentation and real device behaviour. To address this issue, we introduce “vd” (virtual device), a software for simulating stream-based virtual devices that enables testing communication without connecting to the real device. It is focused on the communication layer rather than the device’s underlying physics. The vd listens to a TCP port for client commands and employs ASCII-based byte stream communication. It offers easy configuration through a user-friendly config file containing all necessary information to simulate a device, including parameters for the simulated device and information exchanged via TCP, such as commands and queries related to each parameter. Defining the protocol for data exchange through a configuration file allows users to simulate various devices without modifying the simulator’s code. The vd’s architecture enables its use as a library for creating advanced simulations, making it a tool for testing and validating device communication and integration into EPICS. Furthermore, the vd can be integrated into CI pipelines, facilitating automated testing and validation of device communication, ultimately improving the quality of the produced control system.

## INTRODUCTION

EPICS (Experimental Physics and Industrial Control System) is one of the most widely used frameworks for designing distributed control systems for large experiments such as accelerators and observatories [1].

Integrating new devices into the EPICS can be challenging. This is often attributed to the dissonance between documented device behaviours and their operations. This disparity can lead to sub-optimal integration paths, potentially compromising the robustness and reliability of the control system.

Another layer of complexity arises due to project delays, making it increasingly difficult to test the devices on time. Sometimes, specific devices can only be tested once they are delivered. This further exacerbates the challenges faced during integration, increasing the potential for integration errors and system malfunctions.

A solution to this problem could be writing device simulators and creating integration tests based on these simulators. While this may sound straightforward, crafting simulators is time-consuming. This is primarily because it necessitates extensive programming and a profound understanding of the device’s behaviour. The effort to develop a simulator can

sometimes feel counterintuitive, as the time invested only sometimes aligns with the benefits received. Furthermore, developing simulators demands a different set of networking and programming skills than those used daily by control system engineers. It also requires additional effort to ensure the simulator remains updated and consistent with the device’s documentation.

In our pursuit to streamline and enhance the integration process, we are committed to simplifying the creation of simulators. In our solution, we deliberately bypass the portion of the simulator responsible for the device’s behaviour, concentrating solely on the communication layer. This approach allows us to create specific device states by externally adjusting parameter values. Simultaneously, we can verify that the simulated device communicates these changes appropriately using the designated protocol.

By modifying parameters through a separate channel from the one used by the client, we can establish an automated test suite as part of a continuous integration process. This allows us to set up a grid of tests covering all the essential device states without coding the device’s behaviour.

## ARCHITECTURE

The software’s architectural design comprises four distinct layers delineated in Fig. 1. It employs the Golang programming language, leading to a single binary file as output. The absence of external dependencies in this configuration simplifies the distribution process and augments its compatibility, ensuring seamless execution across a diverse range of operating systems.

In the architectural design, the layers were interconnected via interfaces. Within the Golang programming language context, interfaces serve as a core way to construct complex systems. These systems comprise modular components, each of which can be independently developed and tested. Such an approach fosters flexibility and ensures that the Stream layer can be substituted in the future with an alternative layer accommodating different communication protocol formats, should the need arise.

Notably, Fig. 1 omits elements related to the reading of the configuration file in the TOML format and the logging mechanism. This intentional exclusion ensures clarity and prevents obfuscation of the primary architectural representation.

### TCP Server

The first layer under discussion is the TCP Server. It is constructed based on the standard library available in Golang. Golang, known for its efficient concurrency management and clean syntax, offers developers a solid foundation to build robust systems.

\* marcin.lukaszewski@e9controls.com

# TANGO INTEGRATION OF THE SKA-LOW POWER AND SIGNAL DISTRIBUTION SYSTEM

E. L. Arandjelovic<sup>1\*</sup>, U. K. Pedersen<sup>1†</sup>, Observatory Sciences Ltd., U.K.

J. Engelbrecht<sup>1‡</sup>, Vivo Technical, Cape Town, South Africa

D. Devereux<sup>1§</sup>, CSIRO, Australia

<sup>1</sup>also at SKA Observatory, Jodrell Bank, United Kingdom

## Abstract

The Square Kilometre Array Observatory (SKAO) is the world's largest radio telescope, currently being constructed on two sites: SKA-Low in Western Australia, and SKA-Mid in South Africa. The Power and Signal Distribution System (PaSD) is a key component of the SKA-Low telescope, responsible for control and monitoring of the electronic components of the RF signal chain for the antennas, and collecting the RF signals for transmission to the Central Processing Facility. This paper will describe how the PaSD is being integrated into the Tango-based SKA-Low Monitoring Control and Calibration Subsystem (MCCS) software, including the facility for a drop-in Python simulator which can be used to test the software.

Keywords: Tango, SKAO, Modbus, Telescope, Python.

## INTRODUCTION

SKAO represents the next generation of radio astronomy, poised to transform our understanding of the universe. Spanning two continents, Australia and South Africa, this very large international initiative boasts two cutting-edge radio telescopes. SKAO's ambitions are wide-ranging and include exploring cosmic dawn: the formation of the first stars, galaxies and black holes, investigating dark energy and the acceleration of the expansion of the universe.

## SKA-LOW

The SKA-Low is a radio telescope featuring over 130,000 log-periodic antennas, operating in the frequency range of 50 MHz to 350 MHz, with a total collecting area of 419, 000 m<sup>2</sup>. Situated in the desert of Western Australia, the antennas are grouped into 512 Field Stations, distributed in three spiral arms radiating from a central core, allowing a maximum station-to-station distance of 65 km. SKA-Low's unique design employs wire-type antennas and advanced back-end technology for efficiency at low frequencies. It operates as a mathematical telescope, processing data, applying time-delays to align the phases of signals received from a certain direction, to form virtual beams that "point" the telescope in different directions without moving parts. The signals within the beam can then be searched for transient phenomena and timed.

\* ela@observatorysciences.co.uk

† ukp@observatorysciences.co.uk

‡ jarrett@vivosa.co.za

§ drew.devereux@csiro.au

Each Field Station of the SKA-Low telescope consists of a 256 antenna-element array to capture and amplify the signal from the sky. The Power and Signal Distribution System (PaSD), designed by Curtin Institute of Radio Astronomy (CIRA) takes care of powering the antennas, collecting the RF signals, and providing local monitoring and control [1].

## PaSD ARCHITECTURE OVERVIEW

The PaSD system comprises 'SMART boxes' (SMART: Small Modular Aggregation and RFoF Trunk) which each connect directly to around 10 antennas to provide local monitoring and control, and one Field Node Distribution Hub (FNDH) per Field Station which distributes power to all the SMART boxes and provides an Ethernet-serial communications gateway, as well as additional local monitoring. Micro-controllers inside the SMART boxes and FNDH protect the equipment from damage by automatically turning off ports in response to current and temperature readings, thus separating the equipment protection concerns from any external control system.

All of the PaSD parameters that can be monitored or controlled are accessible via Modbus [2] registers in the FNDH and SMART boxes. A register map is published by the PaSD firmware manufacturers which describes the purpose and function of each register. These include read-only registers e.g. the Power Supply Unit (PSU) temperatures and output voltages, and read-write registers e.g. current trip thresholds. The first register of the FNDH and SMART boxes holds the corresponding register map revision number, which is incremented whenever this map has changed following a firmware update.

All communication to the SMART boxes is funnelled through the FNDH on a multi-drop serial bus using the Modbus ASCII protocol. A Field Node Communications Controller (FNCC) board in the FNDH handles the incoming Modbus packets, passes them on to the SMART boxes which each have a unique Modbus address, and routes responses back. The PaSD architecture is detailed in Fig. 1.

## MONITOR, CONTROL AND CALIBRATION SYSTEM

The Monitor, Control and Calibration System (MCCS) is responsible for the monitoring and control of all the local hardware on the SKA-Low Field Stations. It monitors and aggregates hardware status, making this available to the Telescope Manager (TM) which provides the observation management interface for operators and scientists [3]. It also

Software

Software Architecture & Technology Evolution

# PORTING OPENMMC TO STM32 MICROCONTROLLERS FOR FLEXIBLE AMC DEVELOPMENT

M. B. Stubbings\*, E. P. Juarez, L. Stant, Diamond Light Source, Didcot, UK  
A. A. Wujek†, CERN, Geneva, Switzerland

## Abstract

Diamond Light Source has chosen the MicroTCA platform for high performance data acquisition and controls as part of the Diamond-II 4th generation light source upgrade. One requirement is the ability to create custom Advanced Mezzanine Cards (AMCs) for signal conditioning and interlock support. To facilitate this, a Module Management Controller (MMC) is required to negotiate payload power and communications between the AMC and MicroTCA shelf. A popular open-source firmware for controlling such a device is OpenMMC, a project from the Brazillian Light Source (LNLS), which employs a modular approach using FreeRTOS on ARM microcontrollers. Initially, openMMC supported the NXP LPC series of devices. However, to make use of Diamond's existing ST Microelectronics (STM32) infrastructure, we have integrated a CERN fork of the project supporting STM32 microcontrollers into openMMC. In this paper, we outline our workflow and experiences introducing a new ARM device into the project.

## INTRODUCTION

MicroTCA is an open standard for constructing high performance computer systems in a small form factor [1]. It defines a number of hot-swappable modules, which when connected to a backplane provide power, cooling, management and user functionality. At its core are Advanced Mezzanine Cards (AMCs), which are modules that provide the processing and I/O required to implement an application. Board management and communication with the rest of the system is handled by a Module Management Controller (MMC), which is typically implemented as a low-power Microcontroller Unit (MCU) on top of the AMC.

### Electronic Keying

When an AMC is inserted into a MicroTCA shelf, the on-board MMC must pass an Electronic Keying (E-Keying) stage before it is allowed access to payload power and communications [2]. The main management module, known as the MicroTCA Carrier Hub (MCH), leads this process to determine the electronic capabilities of the inserted module and its compatibility with the crate. If the module is found to be incompatible, then it will be rejected from receiving power and unable to communicate with the rest of the system. Through this mechanism, the crate ensures that it protects itself and all Field Replaceable Units (FRUs) from mis-operation and power supply overloading.

\* michael.stubbings@diamond.ac.uk

† formerly associated with

### Diamond Light Source

At Diamond Light Source [3], we are undergoing significant changes to our infrastructure as part of the Diamond-II 4th generation light source upgrade [4]. One aspect of these changes includes the use of MicroTCA for high performance data acquisition, processing and control. In the majority of cases, AMCs will be purchased off-the-shelf from vendors. However, for a small number of high speed signal processing applications, we would like to be able to produce our own AMCs.

A major challenge presents itself when attempting to recreate the behaviour of an MMC. There are a number of complex processes, such as E-keying, that the MMC must perform in order for the AMC to operate in a MicroTCA system. Additionally, the MMC's firmware is dependant on the target controller being used and the design of the AMC. Therefore, in each application the firmware would need to be revised to account for the change in design.

Together, these issues present a significant amount of work required to support only a small number of use cases. Through this realisation, we looked for an alternative solution.

## OPENMMC

openMMC [5] is an open source, hardware independent firmware designed to carry out the operations of an MMC in a MicroTCA system [6]. Its modular architecture allows for flexible configuration of the sensors, communications and controller used for a target board. It employs the FreeRTOS [7] operating system for independent task management and advanced hardware control. Whilst it was initially released with support for the NXP LPC17xx family of chips, there is scope to port the project onto other microcontroller architectures.

### STM32 Support

A team at CERN [8] created a fork of the project that provides support for the STM32 family of microcontrollers. However, during development they faced several issues with the portability of the code so made adjustments to the core architecture to resolve this. As a result, the fork was no longer compatible with the LP17 MCUs and the flexibility was lost. At Diamond, we are interested in combining our existing ST Microelectronics infrastructure with openMMC to produce new AMCs. As such, we created this project to complete the integration work that CERN started, thereby introducing a new ARM device into the firmware.

# INTEGRATION OF BESPOKE DAQ SOFTWARE WITH TANGO CONTROLS IN THE SKAO SOFTWARE FRAMEWORK — FROM PROBLEMS TO PROGRESS

A. J. Clemens\*, Observatory Sciences Ltd., Cambridge, United Kingdom  
D. Devereux†, CSIRO, Australia

A. Magro‡, Institute of Space Sciences and Astronomy University of Malta, Malta

## Abstract

The Square Kilometre Array Observatory (SKAO) project is an international effort to build two radio interferometers in South Africa and Australia to form one Observatory monitored and controlled from the global headquarters in the United Kingdom at Jodrell Bank. The Monitoring, Control and Calibration System (MCCS) is the “front-end” management software for the Low telescope which provides monitoring and control capabilities as well as implementing calibration processes and providing complex diagnostics support.

Once completed the Low telescope will boast over 130,000 individual log-periodic antennas and so the scale of the data generated will be huge. It is estimated that an average of 8 terabits per second of data will be transferred from the SKAO telescopes in both countries to Central Processing Facilities (CPFs) located at the telescope sites.

In order to keep pace with this magnitude of data production an equally impressive data acquisition (DAQ) system is required. This paper outlines the challenges encountered and solutions adopted whilst incorporating a bespoke DAQ library within the SKAO’s Kubernetes-Tango ecosystem in the MCCS subsystem in order to allow high speed data capture whilst maintaining a consistent deployment experience.

## INTRODUCTION

The Square Kilometre Array Observatory (SKAO) represents a significant advancement in our pursuit of understanding the universe through radio astronomy. This scientific endeavor requires a complex toolchain in which each component plays a crucial role. At the core of this toolchain lies Docker [1], a well-known technology for containerization, which serves as a key element for packaging and deploying various components of the SKAO project. Additionally, Minikube [2] and Kubernetes [3] take on important roles, managing local development environments and ensuring robust production-level deployments to support the scalability and resilience necessary for a project of this scale. Helm [4], with its templating capabilities, simplifies the deployment process, while Make acts as a unifying force to streamline the intricate interactions between these components, ensuring an efficient deployment workflow.

\* ajc@observatorysciences.co.uk

† drew.devereux@csiro.au

‡ alessio.magro@um.edu.mt

Amidst this technological tapestry a significant challenge emerged: the seamless integration of third-party software, including xGPU [5] and NVIDIA’s CUDA [6] (Compute Unified Device Architecture) along with their dependencies into the framework of Tango Controls [7]. Tango Controls, the chosen control framework for the SKAO project, forms the backbone upon which our astronomical endeavors rely.

This paper documents the journey undertaken to bridge the gap between the SKAO deployment toolchain and third-party data acquisition (DAQ) software [8]. Its purpose is to provide a comprehensive account of the strategies employed, the complexities encountered, and the solutions devised during this process. The goal is to offer valuable insights, not just as a record of achievements but as a resource for engineers and scientists facing similar integration challenges.

We will explore two distinct phases: the first phase “The Quest for Data” delves into the incorporation of third party software into the SKAO’s data acquisition system. The subsequent phase “The Correlator Saga” explores the challenges and solutions encountered in the integration of these critical components.

## THE QUEST FOR DATA

In our pursuit of data acquisition the first milestone was the creation of a containerized Tango device server to drive the DAQ (data acquisition) software. This step laid the foundation for our data acquisition endeavors within the project. As we embarked on this quest we encountered a series of formidable challenges, each demanding resourcefulness and persistence to surmount.

### *Inheriting Capabilities*

The initial challenge arose from the limitations of configuring capabilities solely within the Dockerfile. While the container itself possessed the necessary capabilities, a crucial nuance emerged: the Kubernetes pods failed to inherit these essential capabilities.

To address this challenge we informed Helm about the specific capabilities required by appending them to the `securityContext::capabilities::add` field within the values file. This ensured that Kubernetes pods inherited the critical capabilities, aligning both the container and pod with the requisite permissions.

### *Selective Capability Application*

Expanding upon the prior solution, the challenge of selective capability application came to the forefront. Despite

# GATEWARE AND SOFTWARE FOR ALS-U INSTRUMENTATION \*

L. M. Russo<sup>†</sup>, A. Amodio, M. J. Chin, W. E. Norum, K. Penney, G. J. Portmann, J. M. Weber  
LBNL, Berkeley, USA

## Abstract

The Advanced Light Source Upgrade (ALS-U) is a diffraction-limited light source upgrade project under development at the Lawrence Berkeley National Laboratory. The Instrumentation team is responsible for developing hardware, gateware, embedded software and control system integration for diagnostics projects, including Beam Position Monitor (BPM), Fast Orbit Feedback (FOFB), High Speed Digitizer (HSD), Beam Current Monitor (BCM), as well as Fast Machine Protection System (FMPS) and Timing. This paper describes the gateware and software approach to these projects, its challenges, tests and integration plans for the novel accumulation and storage rings and transfer lines. Keywords: Gateware, FPGA, RFSoc, Hardware Platforms.

## INTRODUCTION

The ALS-U project is currently under development at the Lawrence Berkeley National Laboratory, in Berkeley. It consists of a new 2 GeV Storage Ring (SR), a new full-energy Accumulation Ring (AR) and new Transfer Lines between the SR and AR, while reusing the existing Injector [1].

The Instrumentation team was tasked with upgrading the systems for which it's responsible, while maintaining the network-attached device (NAD) architecture currently in use at ALS [2]. Moreover, because all systems, with the exception of the FMPS, have been in operation at ALS for years, it's important to minimize operation disruption, while upgrading the designs in an isolated and uniform way.

## HARDWARE

The Instrumentation projects maximize the usage of open-source hardware solutions in a way to improve robustness against supply chain issues, component obsolescence and collaborative nature. To that end, a dual FMC FPGA carrier board called Marble [3], developed internally and released under the CERN Open Hardware License v1.2 [4], was selected for multiple systems: Fast Orbit Feedback, Fast Machine Protection System, Event Generator, Event Fanout and Event Receiver. The board features a Xilinx Kintex-7 FPGA, DDR3 SODIMM, dual FMC slots, STM32 microcontroller for board management and a flexible clocking scheme. The board can be seen in Fig. 1.

For systems that require faster ADCs, such as the HSD, BCM and BPM, the Xilinx RFSoc SoC [5] family achieved all the requirements. Two demoboard featuring this part were selected: ZCU111 [6], featuring the Gen2 ZU28DR device, consisting of 8 12-bit 4 GSPS ADCs, 8 14-bit 6.5 GSPS




Figure 1: Marble v1.4.

DACs, and a quad-core ARM A53 processor; ZCU208 [7], featuring the Gen3 ZU47DR/ZU48DR devices, consisting of 8 14-bit 5 GSPS ADCs, 8 14-bit 9.85 GSPS DACs, and a quad-core ARM A53 processor. The boards can be seen in Fig. 2.



Figure 2: Xilinx ZCU111 (left) and ZCU208 (right).

## COMMON CODEBASE AND CONTINUOUS INTEGRATION

All Instrumentation FPGA projects are being developed using a common, reproducible and extensible build system based on Makefiles and TCL scripts from a repository called Bedrock [8]. It includes, among rules for generating testbenches and synthesis/implementation, a variety of portable Verilog modules, ranging from common digital logic designs like FIFOs, RAMs and pulse synchronizers, to more complex operations like multi-gigabit fiber protocols and DSP operations. Also, Bedrock is based on free, open-source and well-known tools and languages to reduce external dependencies, while maintaining long-term support and reducing the barrier for collaboration.

As Bedrock, Instrumentation repositories are following Continuous Integration practices, by leveraging Gitlab [9] and Gitlab runner [10] projects to provide automatic testing and a valid final bitstream at each commit. There are also plans to support Hardware-in-the-loop testing [11].

\* Work supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

<sup>†</sup> lucasrusso@lbl.gov

## EXPLORING ETHERNET-BASED CAMAC REPLACEMENTS AT ATLAS\*

K. Bunnell<sup>†</sup>, C. Dickerson, D. Novak, D. Stanton  
Argonne National Labs, Chicago, IL, USA

### Abstract

The Argonne Tandem Linear Accelerating System (ATLAS) facility at Argonne National Laboratory is researching ways to avoid an unscheduled downtime caused by the end-of-life issues with its 30 year-old CAMAC system. Replacement parts for CAMAC are difficult to obtain, causing the potential for long periods of accelerator down times once the limited CAMAC spares are exhausted.

ATLAS has recently upgraded the Ethernet in the facility from a 100-Mbps (max) to a 1-Gbps network. Therefore, an Ethernet-based data acquisition system is desirable. The data acquisition replacement requires reliability, speed, and longevity to be a viable upgrade to the facility. In addition, the transition from CAMAC to a modern data acquisition system will be done with minimal interruption of operations.

Keywords: CAMAC, Data Acquisition, VSystems, network, ioThinX

### INTRODUCTION

The ATLAS accelerator is located at the United States Department of Energy's Argonne National Laboratory in the suburbs of Chicago, Illinois. It is a National User Facility capable of delivering ions from hydrogen to uranium [1] for low energy nuclear physics research to perform analysis of the fundamental properties of the nucleus.

ATLAS uses a 15-crate CAMAC (Computer Automated Measurement and Control) Serial Highway for the majority of its critical equipment data acquisition and control which users can interact with using databases and interfaces provided by Vista Controls (VSystems). CAMAC was developed in 1967 [2], and was a cost-effective option for data acquisition and control for decades to follow [3]; however, the "end-of-life" issues with CAMAC are forcing ATLAS along with other facilities to look for alternatives [4].

The CAMAC highway requires a computer to control and monitor the system. Since the CAMAC drivers are not available to a modern Linux server, ATLAS is forced to maintain an old AlphaServer 1200 running OpenVMS. For over a decade, it has been desired to move away from both CAMAC and the AlphaServers. Therefore, the decision was made to move hardware using CAMAC to a different system that can communicate directly with a Linux server.

After exploring many options, the MOXA ioThinX 4510 was selected to begin testing at ATLAS to determine if the product is a viable replacement of CAMAC.

\* This work was supported by the U.S. Department of Energy, Office of Nuclear Physics, under Contract No. DE-AC02-06CH11357. This research used resources of ANL's ATLAS facility, which is a DOE Office of Science User Facility.

<sup>†</sup> kbunnell@anl.gov




Figure 1: Distribution of ATLAS Hardware by VSystem Database Channel.

### MOTIVATION

The CAMAC system at ATLAS is organized in a loop topology, which means that when 1 CAMAC crate (Figure 3) fails, all 15 crates lose communication with the computer. This serial highway makes up 26% of the ATLAS control system (Figure 1). Since the majority of the ATLAS CAMAC highway comprises of systems critical to operation, a failure in 1 CAMAC crate almost always guarantees a halt in operation lasting until the problem is resolved.

In recent decades, CAMAC hardware has become difficult to replace since they are no longer manufactured. The supply of spare CAMAC components at ATLAS is limited and repairing failed components is often time-intensive.

### CHALLENGES

Tracking down the use of each CAMAC channel is difficult and time consuming. Therefore, a program was created to read and organize all the databases' metadata into a spreadsheet. The spreadsheet can be easily sorted for channels that communicate with each CAMAC crate and slot to determine the scope of each piece of CAMAC hardware.

### NETWORK UPGRADE

In 2021, the entire ATLAS Ethernet network consisted of about 13 hubs with a bandwidth of 10/100Mb. From January 2022 to January 2023, these switches were gradually replaced by 1 Gbps switches. Figure 2 shows a sample of the network ping statistics between two computers on the network during this upgrade. This up-grade will play a crucial role in improving the performance of any Ethernet-based system.



# TEACHING AN OLD ACCELERATOR NEW TRICKS\*

D. Novak<sup>†</sup>, K. Bunnell, C. Dickerson, D. Santiago-Gonzalez, D. Stanton  
Argonne National Laboratory, Lemont, USA

## Abstract

The Argonne Tandem Linac Accelerator System (ATLAS) has been a National User Facility since 1985. In that time, many of the systems that help operators retrieve, modify, and store beamline parameters have not kept pace with the advancement of technology. Investigation and development of a new method of storing and retrieving beamline parameters resulted in the installation and testing of a time-series database as a potential replacement for the traditional relational databases.

## INTRODUCTION

In the operation of ATLAS [1], there is a need save and restore operating values of beamline devices. These stored profiles, which are stored on a per experiment basis, allow operators to restore the beamline to a known state. The ability to restore a similar profile, or scale a profile base on the ion species allows for more rapid tuning of the machine for a new experiment.

To save and restore beamline profiles, ATLAS operators currently use a proprietary commercial off-the-shelf (COTS) relational database (RDB) hosted on a PC. This PC-based database interacts with the control system in a limited way by querying another COTS relational database which resides on an OpenVMS [2] platform. On the OpenVMS platform, a program periodically queries a subset of the control system channels, and stores their values in a table in the OpenVMS-based RDB, for the PC to eventually retrieve and store. The PC-based RDB has been moved to End of Life (EOL) status and is no longer supported by its developer. Additionally, there is a strong desire to move off of the OpenVMS platform as well as the proprietary RDB it hosts.

## TIME SERIES DATABASE

A traditional relational database uses a collection of related tables, with each table having a key to organize rows of data, and a fixed number of columns per table. A time-series database (TSD) differs in a few key ways from a relational database. The main difference is that a TSD always uses time as its key. A nice feature of some TSDs, is that the number of columns is not fixed. Using a specific time seemed like it would be a novel way to be able to store and retrieve the value for every channel in the control system of ATLAS. This means that every control channel and every read-back channel would be stored. The beamline profile would then be a subset of the data

stored at a given point in time. This arrangement would allow secondary programs to access all of the ATLAS control system data. An operator-facing system is planned that would store meta-data about an experiment, such as ion species, experiment number, energy level, etc. This operator system would then simply point to a specific time in the TSD to recall all of the beamline parameters about that experiment.

## TESTING

InfluxDB [3] was selected due to its self-hosted Open-Source version availability as well as the simplicity of installation and setup. A program was written to periodically gather all accelerator control parameters, as well as read-back values, in the control system and store them in the time-series database. This resulted in over 13,000 distinct data points, captured at 5-minute intervals. While this seemed like a lot of data to capture, InfluxDB did not have any issue keeping up. This testing uncovered bugs in a specific part of the underlying software in the control system and has been halted until those issues can be resolved.

Graphing of the captured data is being done on Grafana [4], a self-hosted Open-Source version is available that co-exists well with InfluxDB as the back-end. Grafana made visualizing the data simple and flexible.

A second test captured 35 channels on a 1-minute cadence on the ATLAS Californium Rare Isotope Breeder Upgrade (CARIBU) [5]. This has been reliably gathering data for several months. These measurements quickly became used in regular operation of CARIBU, with physicists creating their own specialized graphs. See Figs. 1-3 Due to the success of the second test, a number of other channels have been added to the InfluxDB in an ad hoc manner in an effort to aid various teams in recovering from a power outage. See Figs. 4 and 5 for cryogenics displays.




Figure 1: CARIBU High Voltage Display.

\* This work was supported by the U.S. Department of Energy, Office of Nuclear Physics, under Contract No. DE-AC02-06CH11357. This research used resources of ANL's ATLAS facility, which is a DOE Office of Science User Facility.

<sup>†</sup> dnovak@anl.gov

## LANSCE's TIMING SYSTEM STATUS AND FUTURE PLANS\*

L. E. Walker<sup>†</sup>, B. C. Atencio, S. A. Baily, D. Fratantonio, C. D. Hatch, M. Pieck, T. Ramakrishnan, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

### Abstract

The Los Alamos Neutron Science Center (LANSCE) operates at a maximum repetition rate of 120 Hz. Timing gates are required for synchronization of the accelerator to provide beam acceleration along the LINAC and beam distribution to the five experimental areas. They are also provided to other devices with sensitive operating points relative to the machine cycle. Over the last 50 years of operations many new time sensitive pieces of equipment have been added. This has changed the demand on, and complexity of, the timing system. Further driven by equipment obsolescence issues, the timing system underwent many upgrades and revitalization efforts, with the most significant deployment starting in 2016. Due to these upgrade efforts, the timing system architecture design changed from a purely centralized system, to a distributed event-based one. The purpose of this paper is to detail the current state of the timing system, as a hybrid system with the gate events being generated from a new timing master system, while still utilizing legacy distribution and fanout systems. Upgrades to the distribution system are planned, but due to the required beam delivery schedule, they can only be deployed in sections during four-month annual maintenance cycles. The paper will also cover the off-the-shelf solutions that have been found for standardization, and the efforts towards a life cycle management process.

### INTRODUCTION

When the timing system was first deployed at the site currently known as Los Alamos Neutron Science Center (LANSCE) it was state of the art. However, that deployment was several decades ago and with any facility that has been running for a long time, the timing system is in need of an upgrade. Since the construction of LANSCE more user facilities have been added to create a total of five that are currently in use today. The timing system can produce 120 Hz timing signals, which is then divided between the five user facilities. This is done by scheduling a flavor gate with definable repetition rates and lengths for each user facility. In addition to sending beam to more users, there is an added complexity that each user facility requires their own independently timed equipment, such as fast kicking magnets, beam diagnostics, current monitors, and spill detectors requiring a variety of interacting timing signals.

The timing system at LANSCE consists of several generations of both commercial off the shelf (COTS) and custom-made components that require significant effort to maintain. A timing system modernization effort began with its conceptualization in 2007 and deployment in 2016 [1, 2].

<sup>†</sup> lwalker@lanl.gov

Within the last seven years, approximately 120 new timing Input/Output Controllers (IOCs) have been deployed throughout the facility. Great progress has been made, but there are still upgrades needed that will expand the facilities capabilities and allow the site to operate for years to come.

### CURRENT SYSTEM DESCRIPTION

Changing technologies throughout LANSCE's years of operation have led to many variations in the hardware deployed for the timing system. In 2016, the facility began an upgrade from CAMAC crates with a MicroVAX computer controller to a new timing system designed around the Micro Research Finland (MRF) timing event system [3, 4]. The MRF timing modules are available off the shelf and provide the features that are needed for the diverse timing system at LANSCE. The MRF event generators (EVGs) and event receivers (EVRs) work together to create the event-based timing structure. This upgrade is still ongoing, leaving the timing system in a hybrid state.

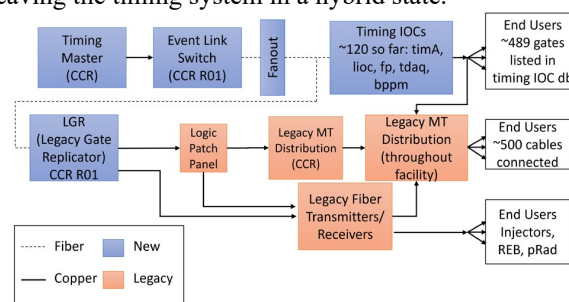


Figure 1: Hybrid state of timing system.

Figure 1 shows the current status of the timing system. The timing master, event link switch, fanout, and timing IOCs have been upgraded and deployed. The Legacy Gate Replicator (LGR) was installed to reproduce the 96 gates that were formed by the old timing system for distribution. This leaves the Logic Patch Panel, Legacy Master Timer Distribution, legacy Fiber Transmitters/Receivers systems to redesign and deploy.

### Timing Master




Figure 2: Timing Master System.

The Timing Master uses a redundant scheme, with the second system running in hot-swappable standby configuration in the event there is a failure in the primary

## LCLS-II CRYOMODULE ISOLATION VACUUM PUMP CART\*

S.C. Alverson<sup>†</sup>, D.K. Gill, S. Saraf, SLAC National Accelerator Lab, Menlo Park, California, U.S.A.

### Abstract

The Linac Coherent Light Source II (LCLS-II) Project is a major upgrade to the lab's Free Electron Laser (FEL) facility adding a new injector and superconducting linac. In order to support this new linac, a pumping scheme was needed to isolate the liquid helium lines cooling the Radio Frequency (RF) cavities inside the cryomodules from outside ambient heat as well as to exhaust any leaking helium gas.

New carts were built consisting of a mechanical backing pump, turbo-molecular pump, and several vacuum valves and gauges for this purpose. These were designed to both automate the process of pump down for the technicians via remote and local control, as well as be portable enough to be able to be moved to different locations as needed depending on the state of the vacuum in the cryomodule strings (Fig. 1).

### BACKGROUND

The primary concerns for the LCLS-II [1] cryomodule isolation region are buildup of gasses that can transfer heat to the circulating cryogenic lines as well as formation of ice which can build up causing damage [2].

- It is almost impossible to completely avoid all leakage of helium gas within the cryomodule, so it must be continuously pumped out.
- If the cryogenic lines start warming up due to heat transfer to outside atmosphere, the cryogenic fluids will expand creating more gas leakage into the isolation vacuum.
- If temperatures increase too much, a quench [3] will occur shutting off the machine and lead to downtime.

### SOLUTION

To avoid these issues, we designed a pump cart system with automated pump down sequences and interlocks. As much electronics as possible were installed remotely outside of the tunnel to avoid radiation damage and allow access for troubleshooting. Design is intended to be turn-key such that pump down and recovery operations are largely hands-off, governed by the system logic. Carts and support infrastructure are designed to be as portable as possible in case it needs to be relocated to areas with higher leak rates.

### Hardware

In order to keep the cart portable, it was designed in-house [4] to be as compact as possible so it could be easily moved to different locations in the tunnel depending on where the worst vacuum is detected. All hardware was mounted on a metal frame with casters as well as adjustable

\* LCLS is supported by the U.S. Department of Energy, Office of Science under Contract DE-AC02-76SF00515.

<sup>†</sup> alverson@slac.stanford.edu



Figure 1: Photo of LCLS-II L0B cryomodule string.

legs and brackets that can be utilized for bracing to the tunnel floor.

**Cart Devices** The cart interfaces to the cryomodule isolation vacuum volume via a VAT manual valve installed on the side of the cryomodule feed cap. Care had to be taken during installation as the cryogenic beamline is a particle-free area [5]. Hardware mounted within the cart itself (Fig. 2) consists of the following devices:

- MKS 3170037SH Convection (Pirani) Gauge (x2)
- MKS 4220014 Cold Cathode Gauge
- VAT 09140-PE24-X Pneumatic Gate Valve
- VAT 26428-KE21 Pneumatic Angle Valve
- Pfeiffer TP1 HiPace 300 Turbo Pump
- Kashiyama NeoDry 36-12 Roots Pump

A passive supplementary air reservoir was also installed on the cart to ensure that if the site air supply goes down, enough air will still be available to facilitate closing the valves.

**Controls and Infrastructure** To avoid failures due to radiation damage, all discrete electronics had to be installed in the support building high above the tunnel (Fig. 3). Spare long haul cables were installed periodically along the cryomodule strings ready to be plugged in if a cart needs to be moved. These lead up to the controls racks contained in each sector alcove where the Programmable Logic Controller (PLC) and device controllers are installed consisting of the following:

- Allen-Bradley ControlLogix<sup>®</sup> PLC
  - 1756-L83 Controller
  - 1756-EN2T Ethernet
  - 1756-IF16 Analog Input
  - 1756-OF8I Analog Output
  - 1756-IB32 Digital Input
  - 1756-OB16I Digital Output

# LCLS-II CONTROL SOFTWARE ARCHITECTURE FOR THE WIRE SCANNER DIAGNOSTICS\*

N. Balakrishnan<sup>†</sup>, J. Bong, A. Fisher, B. Jacobson, L. Sapozhnikov  
SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

The Super Conducting (SC) Linac Coherent Light Source II (LCLS-II) facility at SLAC is capable of delivering electron beam at a fast rate of up to 1 MHz. The high rate necessitates the processing algorithms and data exchanges with other similar systems to be implemented with FPGA technology. For LCLS-II, SLAC has deployed a common platform solution (hardware, firmware, software) which is used by timing, machine protection and diagnostics systems. The wire scanner diagnostic system uses this solution to acquire beam synchronous time-stamped readings, of wire scanner position and beam loss during the scan, for each individual bunch. This paper explores the software architecture and control system integration for LCLS-II wire scanners using the common platform solution.

## INTRODUCTION

For measuring the transverse electron beam profiles and emittance at various locations, wire scanners (Fig. 1) are one of the primary tools installed and commissioned for use at SLAC [1]. There are two flavors of wire scanners – a fast wire scanner and a slow wire scanner. The fast wire scanner system is comprised of a linear motor stage with an incremental linear encoder for closed loop position feedback [2]. The slow wire scanner system is comprised of a stepper motor and a linear Variable Differential Transformer (LVDT) and an incremental encoder on motor shaft for position feedback. On both the styles of wire scanner, the movable stage holds a wire card. The wire card holds thin wires (in the order of 20 nm), generally much smaller than the beam transverse size (which can range from 40 to 300 um) [3].

As this thin wire passes through the electron in the transverse direction, it intercepts the portion of the beam incident on it, creating gamma radiation captured by photodiode detectors. The intensity of these beam loss readings obtained during a scan of the wire card through multiple bunches of the beam, helps provide the cross section of the beam. This measurement depends on several parameters. For e.g.: the thickness of the wire, the repetition rate with which the electron beam is incident on the wire, the speed with which the wire moves to intercept the beam, to name a few.

The higher beam rate of LCLS-II could damage the wires [3, 4]. Thus, for higher repetition rate, the wire needs to move faster through the beam region to prevent damage while still making a reasonable

measurement. To prevent this damage, a minimum speed is defined. The wire scanner firmware is required to communicate to the Machine Protection System (MPS) if the minimum speed has not been reached prior to the wire intercepting the beam [4].

For LCLS-II, long radiation-hard optical fibres are installed from the electron gun to the beam dump as part of Beam Containment System. Beam losses generate gamma radiation that travels through the fiber. A photomultiplier tube (PMT) is installed at the end that measures this light. Signal from this PMT is used by the wire scanner.

In this paper, we will detail the software architecture of LCLS-II wire scanner system. We will also briefly describe the system for beam synchronous data acquisition for wire scanner. Wire scanner system also uses common platform solution to implement the MPS requirement and PMT signal readback. Details of this implementation will be provided in this paper.




Figure 1: Fast Wire Scanner and a Slow Wire Scanner at SLAC.

## CONTROL SYSTEM

### Software Architecture Overview

Wire scanner controls is broadly split into 4 different stack system (Fig. 2). At the lowest level, we have the motion controller which is used for trajectory planning and closed loop control. On top of this layer is the wire scanner FPGA which is used to provide the wire position and beam loss signal, beam synchronously, to EPICS. This layer also handles the MPS Fault detection. EPICS IOC layer determines the scan parameters for motion controller and Beam Loss Signal Integration Settings for FPGA. Lastly, client software such High Level Applications or User Interfaces performs the emittance measurement calculations.

\* Work supported by U.S. Department of Energy under contract number DE-AC02-76SF00515

<sup>†</sup> namrata@SLAC.Stanford.EDU

# ATCA-BASED BEAM LINE DATA SOFTWARE FOR SLAC'S LCLS-II TIMING SYSTEM\*

D. Alnajjar<sup>†</sup>, M. P. Donadio, K. Kim, M. Weaver  
 SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

Among the several acquisition services available with SLAC's high beam rate accelerator, all of which are contemplated in the acquisition service EPICS support package, resides the new ATCA Beam Line Data (BLD) service, which runs on top of SLAC's common platform software and firmware, and communicates with several high-performance systems (i.e. MPS, BPM, LLRF, timing, etc.) in LCLS, running on a 7-slot Advanced Telecommunications Computing Architecture (ATCA) crate. Once linked with an ATCA EPICS IOC and with the proper commands called in the IOC shell, it initializes the BLD FPGA logic and the upper software stack and makes PVs available allowing the control of the BLD data acquisition rates, and the starting of the BLD data acquisition. This service permits the forwarding of acquired data to configured IP addresses and ports in the format of multicast network packets. Up to four BLD rates can be configured simultaneously, each accessible at its configured IP destination, with a maximum rate of 1 MHz. Users interested in acquiring any of the four BLD rates will need to register in the corresponding IP destination to receive a copy of the multicast packet on their respective receiver software. BLD has allowed data to be transmitted over multicast packets for over a decade now at SLAC, but always at a maximum rate of 120 Hz. The present work focuses on bringing this service to the high beam rate high-performance systems using ATCAs, allowing the reuse of so many of those legacy inhouse-developed client software infrastructures.

## INTRODUCTION

The 7-slot Advanced Telecommunications Computing Architecture (ATCA) crate is used for numerous high-performance systems (HPS) at SLAC, such as the bunch charge monitor[1], bunch length monitor[1], beam position monitor[2], low-level radio frequency[3], machine protection system[4], timing system, and a few others. In all of these sub-systems, raw data is acquired, processed, timestamped, and transmitted upstream to a server where it is analyzed and exported to the network through EPICS[5].

SLAC has a set of 4 services, called Acquisition Services, used to organize timestamped data in different ways. Each way has its own use case with its set of requirements from users. Beam Synchronous Acquisition (BSA)[6] is one example of an Acquisition Service. The other ones are Beamline Data (BLD), Beam Synchronous Scalar Service (BSSS), and




Figure 1: High-performance system SW/HW overview.




Figure 2: BLD message flow overview.

Beam Synchronous Acquisition Service (BSAS).

BLD runs on top of SLAC's common platform software and firmware[7] and permits the forwarding of acquired

Hardware

Timing Systems & Synchronisation

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

\* Work supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515

<sup>†</sup> dnajjar@slac.stanford.edu

# LCLS-II ACCELERATOR VACUUM CONTROL SYSTEM DESIGN, INSTALLATION AND CHECKOUT\*

S. Saraff†, S. Karimian, S. Alverson, S. Nguyen, C. Lai,  
SLAC National Accelerator Laboratory, Menlo Park, USA

## Abstract

The LCLS-II Project at SLAC National Accelerator Laboratory has constructed a new superconducting accelerator which occupies the first kilometre of SLAC's original 2-mile-long linear accelerator tunnel. The LCLS-II Vacuum System consists of a combination of particle free (PF) and non-particle free vacuum (non-PF) areas and multiple independent and interdependent systems, including the beamline vacuum, RF system vacuum, cryogenic system vacuum and support systems vacuum.

The Vacuum Control System incorporates controls and monitoring of a variety of gauges, pumps, valves and Hiden RGAs. The design uses a Programmable Logic Controller (PLC) to perform valve interlocking functions to isolate bad vacuum areas. In PF areas, a voting scheme has been implemented for slow and fast shutter interlock logic to prevent spurious trips. Additional auxiliary control functions and high-level monitoring of vacuum components is reported to global control system via an Experimental Physics and Industrial Control System (EPICS) input output controller (IOC). This paper will discuss the design as well as the phased approach to installation and successful checkout of LCLS-II Vacuum Control System.

## SCOPE

This paper covers the LCLS-II Accelerator Vacuum Control System requirements from Injector to the Electron Beam Dump. Vacuum Control requirements for the experimental areas are not included in the scope of this paper. The LCLS-II Accelerator Vacuum System consists of multiple independent and interdependent systems, including the beamline vacuum, RF system vacuum, cryogenic system vacuum and support systems vacuum with details below:

**Beamline Vacuum System** begins at the gun, continues through the superconducting RF accelerator, bypasses the existing warm linac and then it is spread into three lines going to the primary dump, Soft Xray (SXR) and Hard Xray (HXR) lines in the Beam Switch Yard (BSY). These require ultra-high vacuum (UHV). These electron beamlines include the cryogenic beamlines (cold) and

conventional warm electron beamlines as well. The design of the vacuum interface between two neighbouring regions operating at disparate temperatures must consider the impacts of their temperatures on the vacuum environments. An additional characteristic that applies to all the cryogenic beamline and a portion of the warm electron beamline is very low particle count (also referred to as “particle-free (PF)”) cleanliness. [1]

**Cryogenic System Vacuum** consists of several copies of physically and functionally separated vacuum systems including insulating vacuum systems and sub-atmospheric helium vacuum systems. These systems are high vacuum systems with high gas loads. [1]

**RF System Vacuum** includes RF coupler vacuum and laser transport tube vacuum. The RF coupler vacuum system can be characterized as warm UHV vacuum system with particle cleanliness requirements due to the RF fields and will be treated as particle free. Laser Transport Tubes provide vacuum transport lines for the UV and IR laser beams between the Laser Room and their respective use locations and require high vacuum (HV). [1]

**Support System Vacuum** includes the various pump down carts like the Insulating vacuum roughing carts, Insulating Vacuum High Vacuum carts, UHV pump down cart and particle free UHV pump down cart. [1]

## SYSTEM DESIGN

LCLS-II Vacuum System can be separated into two parts: mechanical vacuum devices and controls vacuum devices. Mechanical devices are those that are physically part of the beam line or waveguide: vacuum valves, vacuum gauges, and vacuum pumps. These devices are specified, tested, and installed by the Mechanical Engineering and Technical Services Department (METSD). Controls devices are the remaining hardware needed to build a complete vacuum system: cables and controllers for the vacuum valves, gauges, and pumps; devices used to perform interlocking functions; and the vacuum section of the EPICS control system. These are specified, tested, and installed by the Electrical Engineering Dept [2].

\* Work supported by U.S. Department of Energy under contract number DE-AC02-76SF00515

† shweta@SLAC.Stanford.EDU

# CREATING OF HDF5 FILES AS DATA SOURCE FOR ANALYSES USING THE EXAMPLE OF ALPS IIc AND DOOCS CONTROL SYSTEM

S. Karstensen<sup>1</sup>, P. Gonzalez-Caminal<sup>3</sup>, G. Günther<sup>2</sup>, A. Lindner<sup>1</sup>, O. Mannix<sup>2</sup>, I. Oceano<sup>1</sup>,  
V. Rybnikov<sup>1</sup>, K. Schwarz<sup>1</sup>, G. Sedov<sup>1</sup>

<sup>1</sup>DESY Hamburg, Germany

<sup>2</sup>Helmholtz-Zentrum Berlin für Materialien und Energie GmbH (HZB), Germany

<sup>3</sup>Fusion for Energy (ATG Science & Technology, S.L.), Barcelona, Spain

## Abstract

A versatile and graphical HDF5 file generation project. In the realm of physical experiments, data is typically gathered through the utilization of measurement devices intricately integrated into control systems. These control systems employ diverse methods for archiving historical data. Nevertheless, users consistently grapple with a recurring challenge: how to access the data, decipher the intricacies of the control system's structure, and understand the formatting. The prevailing approach involves the creation of scripts or programs tasked with extracting data from these control systems, and if necessary, followed by the requisite data pre-processing for subsequent analysis. Typically, this pre-processing is carried out in formats known exclusively to the individual user. Moreover, the longevity of data utility often teeters on the brink when users depart or deviate from established conventions.

To address these challenges, we are actively developing software that serves a dual purpose: firstly, to provide an API for a control system (in this instance, DOOCS), and secondly, to transform the extracted data into a universally recognized format, such as HDF5, complete with all relevant metadata.

The ALPS IIc experiment serves as an ideal testbench for this software development task due to its compatibility in terms of state and timing.

## ABOUT HDF5

The Hierarchical Data Format Version 5 (HDF5) 0 is a unique high-performance technology suite that consists of an abstract data model, library, and file format for storing and managing extremely large and/or complex data collections. The technology is used worldwide by government, industry, and academia in a wide range of science, engineering, and business disciplines.

## Advantages of HDF5

- Versatile data model that can represent very complex, heterogeneous data objects and a wide variety of metadata through an unlimited variety of datatypes.
- Ready for high-speed raw data acquisition
- Portable and extensible with no limits on file size, allowing applications to evolve in their use of HDF5.

- Self-describing, requiring no outside information for applications to interpret the structure and contents of a file.
- Robust software ecosystem of open-source tools and applications for managing, manipulating, viewing, and analyzing data.
- Architecturally independent software library that runs on a wide range of computational platforms (from laptops to massively parallel systems) and programming languages (including C, C++, Fortran 90, and Java interfaces)
- Advanced performance features that allow for access time and storage space optimizations through customizable product packaging, compression, and encryption
- Long-term data archiving solution
- Self-explaining data structure

## METADATA

Metadata is structured data that contains overarching information about a resource. Metadata is used to describe the measured data (measurement data) with additional information, enabling their machine and automated processing, as well as providing insight into their origin and improving understanding. Taking metadata into account in the analysis can lead to better results.

Metadata could include:

- Process-specific information
- Images, Log Files, Links
- Hazards
- Responsible Users
- Storage Location
- Information about locations
- Permissions
- Publications
- Derived Objects
- Comments
- Hardware information
- eLogbook entries
- And many other information

# MACHINE PROTECTION SYSTEM AT SARAF

A. Gaget<sup>†</sup>, T. Joannem, V. Nadot, A. Chance, J. Dumas, F. Gougnaud, A. Lotode, S. Monnereau  
 CEA Saclay IRFU, Gif sur Yvette, France  
 E. Reinfeld, I. Shmuely, H. Isakov, A. Perry, N. Tamim, L. Weissman, SNRC, Yavne, Israel

## Abstract

CEA Saclay Irfu is in charge of the major part of the control system of the SARAF LINAC accelerator based at Soreq in Israel. This scope also includes the Machine Protection System. This system prevents any damage in the accelerator by shutting down the beam in case of detection of risky incidents like interceptive diagnostics in the beam, vacuum or cooling defects. So far, the system has been used successfully up to the MEBT. It will be tested soon for the super conducting Linac consisting of 4 cryomodels and 27 cavities.

This Machine Protection System relies on three sets: the MRF timing system that is the messenger of the “shut beam” messages coming from any devices, IOxOS MTCA boards with custom FPGA developments that monitor the Section Beam Current Transmission along the accelerator and a Beam Destination Master that manages the beam destination required. This Destination Master is based on a master PLC. It permanently monitors Siemens PLCs that are in charge of the “slow” detection for fields such as vacuum, cryogenic and water cooling system. The paper describes the architecture of this protection system and the exchanges between these three main parts.

## INTRODUCTION

SNRC and CEA collaborate to the upgrade of the SARAF accelerator to 5 mA CW 40 MeV deuteron and 35 MeV proton beams (Phase 2) at 176 MHz [1]. The CEA control team is in charge of the machine protection system (MPS), which plays a crucial role in the accelerator’s operation. It requires special attention to ensure it can shut off the beam promptly in the event of accidents, thereby preventing any potential damage. To achieve this, the MPS relies on robust hardware components, including FPGA electronic cards and PLCs. Over the years, CEA has accumulated experience with these technologies, leading them to select three primary technologies for implementing the system:

- Siemens PLC for the Beam Destination Master (BDM) part, that consists of checking the Local Control System conditions depending on the beam destination requested.
- MTCA IOxOS cards for the Section Beam Current and Transmission Board (SBCT) developed in FPGA, facilitated by the IOxOS development platform, enabling precise control of the beam status throughout the accelerator with a response time as fast as 5 μs.
- MTCA cards of Micro-Research Finland (MRF) will play a dual role within the system. They will serve as

the foundation for the timing system and will also function as the central component of the MPS, acting as the messenger for ‘shut beam’ commands throughout the entire machine.

## MACHINE PROTECTION PRINCIPLE

### Overview

In the Machine protection system, we can distinguish two types of equipment: detectors and beam stoppers. On the one hand, detectors are responsible for identifying potentially hazardous situations and will initiate a ‘shut beam’ request to prevent any damage to the machine. On the other hand, beam stoppers are designed to deactivate the beam rapidly when they receive a request from the detectors. Figure 1 illustrates the architecture of the MPS, highlighting the various detectors positioned around the three beam stoppers and illustrating their interactions.




Figure 1: Architecture of the SARAF MPS.

### Detectors

- LLRFs: The primary function of an LLRF (Low-Level Radio-Frequency) system is to manage and regulate the amplitude and phase of the electromagnetic field within the accelerating cavity. It is designed on the MTCA.4 platform and uses an EPICS driver, with the entire system being developed by Oroliia [2].
- BPMs: The role of a BPM (Beam Position Monitor) is to furnish data regarding the beam’s position, phase, and current at various locations along the accelerator line. The control is designed on the MTCA.4 platform and uses an EPICS driver, with the entire system being developed by Oroliia [3].
- NBLMs [4]: The neutron beam loss monitor (nBLM) system is based on Micromegas gaseous detector sensitivities to fast neutrons produced when beam particles

<sup>†</sup> alexis.gaget@cea.fr



# SECTOR FOCUSED CYCLOTRON POWER SUPPLY CONTROL SYSTEM UPGRADE

Xiaojun Liu, Wei Zhang, Shi An, Yun Chen, Junqi Wu, Liang Ge, Mingrui Li  
 Institute of Modern Physics, Lanzhou, China

## Abstract

The old power supply control system of SFC (Sector Focused Cyclotron) has been in operation for more than a decade. Control system architecture is centralized and equipment failure rate is getting higher and higher. The new control system uses the EPICS architecture and the hardware uses Advantech's APAX modules. The IOC runs on the APAX host and interacts with the module through API functions. The system has been running very stable for several months without failure.

## INTRODUCTION

Lanzhou Heavy Ion Research Facility (HIRFL) includes Electron Cyclotron Resonance ion source (ECR), Sector Focusing Cyclotron (SFC), Separated Sector Cyclotron (SSC), Cooling Storage Ring (CSR) main ring and experimental ring, radioactive beam line, experiment terminals and other parts [1]. The power supply control system is one of the important components of the accelerator, which can realize remote debugging, data storage and historical data analysis of the power supply. The old power supply control system adopts a centralized structure. If the main control device fails, the entire control system will not be able to operate. The new control system uses the EPICS architecture, changes to a distributed control system, and realizes control of multiple power supply types (Fig. 1).

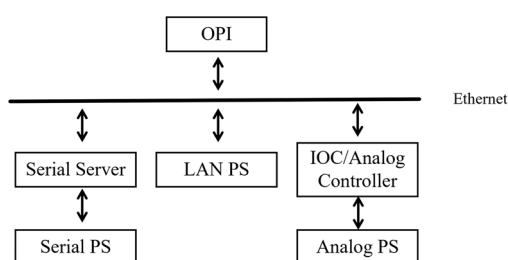


Figure 1: The structure of PS control system.

## SOFTWARE

There are two types of power supply IOC control programs, one interacts with the power supply controller through the communication interface, and the other uses the analog control module to output and input analog signals. Analog power supplies also require coefficients to make the power supply output consistent with reality. The input and output coefficients are written to the analog controller through IOC. The power supply for the communication interface only needs to send and receive data instructions through the serial port or network port.

The scanning magnet power supply of SFC T1 terminal needs to first generate a triangle wave according to the period and amplitude, and then send it to the power controller according to the power supply protocol. The control flow of all power supplies is as shown in Fig. 2.

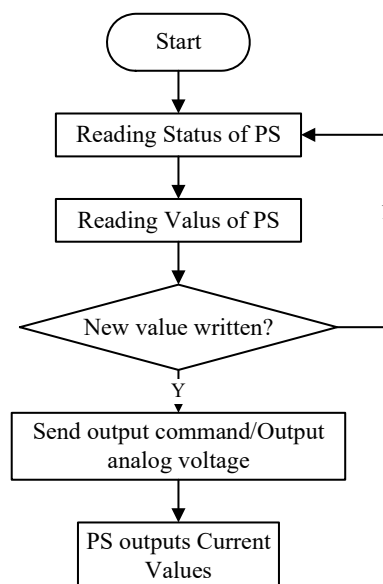


Figure 2: Control flow of all power supplies.

When the program starts running, it first reads the operating status of the digital power supply to see if there is a fault code, and then reads the output value and input value of the power supply to determine whether there are new values written. Analog power supply cannot read the status because the interface is analog signals.

## HARDWARE

The main controller of the power control system uses Advantech APAX5580, the output module is APAX5028, and the input module is APAX5017. APAX5580 adopts the sixth generation Intel® Core™ i7 processor, up to 2.6 GHz, 8 GB DDR4 memory, compact fanless design, can be used for DIN rail installation in control cabinets, providing better EtherCAT performance, supports EtherCAT line, star, and ring redundant topologies [2].

APAX5017 has 12 input channels, configurable to ±150 mV, ±500 mV, ±1 V, ±5 V, ±10 V, ±20 mA, 0~20 mA, 4~20 mA. APAX5028 has 8 output channels, which can be configured as ±2.5 V, ±5 V, ±10 V, 0~2.5 V, 0~5 V, 0~10 V, 0~20 mA, 4~20 mA. The SFC analog power supply uses ±10 V control signals. All three types of power control IOCs are in the same main controller. The control hardware is shown in Fig. 3.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# INTEGRATE EPICS 7 WITH MATLAB USING PVACCESS FOR PYTHON (P4P) MODULE

K. Kim\*, E. Williams, J. Bellister, K. Kim, M. Zelazny  
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

## Abstract

MATLAB is essential for accelerator scientists engaged in data analysis and processing across diverse fields, including particle physics experiments, synchrotron light sources, X-ray Free Electron Lasers (XFELs), and telescopes, due to its extensive range of built-in functions and tools. Scientists also depend on Experimental Physics and Industrial Control Systems (EPICS) 7 to control and monitor complex systems. SLAC has developed *Matpva*, a Python interface to integrate EPICS 7 with MATLAB. *Matpva* utilizes the PVAccess for Python (P4P) module and EPICS 7 to offer a robust and reliable interface for MATLAB users that employ EPICS 7. The EPICS 7 PVAccess API allows higher-level scientific applications to get/set/monitor simple and complex structures from an EPICS 7-based control system. Moreover, *Matpva* simplifies the process by handling the data type conversion from Python to MATLAB, making it easier for researchers to focus on their analyses and innovative ideas instead of technical data conversion. By leveraging *Matpva*, researchers can work more efficiently and make discoveries in diverse fields, including particle physics and astronomy.

## INTRODUCTION

MATLAB stands as an indispensable tool for accelerator scientists deeply immersed in the intricate realms of data analysis and processing, spanning an array of scientific domains such as particle physics experiments, synchrotron light sources, X-ray Free Electron Lasers (XFELs), and telescopic observations. Its indispensability is rooted in the vast array of built-in functions and tools it offers, empowering scientists to decipher complex data and derive meaningful insights.

In parallel, the world of scientific instrumentation and control systems relies heavily on Experimental Physics and Industrial Control Systems (EPICS) 7, a set of software tools and libraries widely used for building control systems in scientific research and industrial facilities [1]. Several attempts have been made to integrate EPICS into MATLAB, including projects like *labCA* and *MATLAB Channel Access (MATLAB CA, MCA)* [2, 3]. However, these were specifically designed for the EPICS Channel Access (CA) interface within MATLAB [4]. Consequently, they did not provide comprehensive support for EPICS PVAccess, which offers distinct advantages over CA [5]. These benefits include structured data types, enhanced security measures, efficient data transmission, and the ability to handle structured datasets, such as *NTTable* [6-8].

Efforts have been made to integrate PVAccess into MATLAB at SLAC using *epicsCoreJava* based on Java. However, this integration faced limitations because MATLAB still relies on Java 1.8.x, while the EPICS Java PVAccess support defaults to Java 11, making it incompatible with MATLAB. Moreover, Oracle JDK 8 with the premier support reached its end of life in March 2022. OpenJDK 8 is also slated to end support in November 2026.

Thus, SLAC has decided to integrate EPICS PVAccess with MATLAB using a well-supported language that aligns with their requirements, Python.

Python has emerged within the scientific community due to its numerous advantages, including readability, simplicity, versatility, a robust standard library, a vast community, compatibility, and seamless integration. As a result, numerous institutions and facilities have been transitioning to Python to leverage these benefits.

Acknowledging this shift, SLAC took a pioneering step by developing *Matpva*, a Python interface to integrate EPICS 7 with MATLAB. This work was also inspired by the *labCA*, which is the most widely used tool for integrating EPICS 3 with MATLAB. This innovative solution bridges the gap between two scientific powerhouses and provides researchers with a unified platform for their endeavors. *Matpva* utilizes Python's capabilities and integrates the PVAccess for Python (P4P) module and EPICS 7 [9]. This creates a sturdy and reliable interface for MATLAB users navigating the intricate realms of EPICS 7.

What sets *Matpva* apart, however, is its exceptional capability to streamline the cumbersome process of converting data types from Python to MATLAB, and vice versa. By doing this, *Matpva* empowers scientists to delve into their research with unwavering focus and efficiency. It serves as a catalyst for groundbreaking discoveries in the multifaceted realms of particle physics and astronomy propelling the boundaries of scientific knowledge ever further.

## IMPLEMENTATION

The PVAccess data type comprises Normative Types, which are defined as structures consisting of both required and optional fields. *Matpva* stands out for its ability to convert PVAccess Normative Types from Python to MATLAB. In Listing 1, you'll find a code snippet that demonstrates how PVAccess Normative Types can be converted into MATLAB data types in the *mpvaGet* function.

Listings 1: Snippet of *mpvaGet* function as an implementation example.

```
% Bring P4P python module into MATLAB
MatP4P = py.p4p.client.thread.Context('pva',
    pyargs('nt', false));
PV = MatP4P.get(pvname);
```

\* [ktkim@slac.stanford.edu](mailto:ktkim@slac.stanford.edu)

# CamServer: STREAM PROCESSING AT SwissFEL AND SLS 2.0

A. Gobbo<sup>†</sup>, A. Babic

Paul Scherrer Institut, Villigen PSI, Switzerland

## Abstract

CamServer is a Python package for data stream processing developed at Paul Scherrer Institute (PSI). It is a key component of SwissFEL's data acquisition, where it is deployed on a cluster of servers and used for displaying and processing images from all cameras. It scales linearly with the number of servers and is capable of handling multiple high-resolution cameras at 100 Hz, as well as a variety of data types and sources. The processing unit, called a pipeline, runs in a private process that can be either permanent or spawned on demand. Pipelines consume and produce ZMQ streams, but input data can be arbitrary using an adapter layer (e.g., EPICS). A proxy server handles requests and creates pipelines on the cluster's worker nodes according to rules. Some processing scripts are available out of the box (e.g., calculation of standard beam metrics) but users can upload custom ones. The system is managed via its REST API, using a client library or a GUI application. CamServer's output data streams are consumed by a variety of client types such as data storage, image visualization, monitoring and DAQ applications. This work describes the use of CamServer, the status of the SwissFEL's cluster and the development roadmap with plans for SLS 2.0.

## INTRODUCTION

SwissFEL is an X-ray free-electron laser facility at Paul Scherrer Institute (PSI) [1]. It delivers pulses of X-ray radiation with a duration of a few femtoseconds and operating at a repetition rate of 100 Hz [2]. The SwissFEL control system is based on EPICS [3]. Even though all data sources can be accessed via EPICS Channel Access, most of data acquired in the machine and beamlines is said to be *beam-synchronous* [4], meaning each sampled data record is tagged with a pulse identifier. Beam-synchronous data is streamed by the EPICS's IOCs using the in-house developed *BSREAD*, a library that provides data serialization and stream control over ZMQ [5].

Timing and synchronization are critical at SwissFEL. All beam synchronous sources are connected to the SwissFEL timing system [6] that distributes a unique pulse identifier for each FEL pulse. Sources send out data through a stream stamped with the pulse identifier.

Beam-synchronous data sources can be used together effectively when synchronized, meaning that records from distinct sources belong to the same pulse. The SwissFEL data-acquisition system provides tools for synchronization and thus enables online processing of beam-synchronous sources. Data channels from different sources can be efficiently aggregated and combined into new data streams.

Streams from various sources are directed to the *Data-Buffer*, which performs functions of data synchronization, dispatching and temporary storage (buffering). The Data-Buffer is an in-house development based on Java Spring [7], running in a cluster of 15 servers. Clients can request from it synchronized streams containing data from multiple sources, where all source values are aligned to the same pulse identifier. Data acquisition can be implemented either by receiving such streams or by post-retrieving data using the DataBuffer REST API, requesting a pulse range for a list of channels. Beam-synchronous scalars, waveforms and images are streamed at 100 Hz. The DataBuffer receives typically 8 GB/s of images, and 450 MB/s of scalars and waveforms, retaining data for some days.

Over the past decades stream processing [8] has played an increasingly critical role on the web and within organizations. It offers low-latency results, enabling immediate insights and faster, more informed decision-making. In the context of SwissFEL, stream processing provides effective online feedback, data reduction, filtering, and therefore resource efficiency - saving on processing, storage, and network resources. For example, the machine protection system can disable pulses, and stream processing allows for the discard of data relating to unfit pulses before processing and storage. Additionally, stream processing moves the resource intensive processing from clients to servers, while standardizing the processing algorithms.

CamServer was initially created with the goals of providing unified access to SwissFEL cameras, processing camera data, and streaming images and data forward. With time, it evolved into a generic stream processing system for SwissFEL. Development started in 2017 and features were gradually included over the years. The project is hosted on GitHub [9]. Python [10] was the language of choice for the project, allowing users to benefit from the well-known Python scientific stack to implement custom processing algorithms.

Initially CamServer executed only one standard image processing algorithm for the characterization of the beam. Later, support to custom processing scripts was added and new use-cases were then supported. More than simply sending data to visualization clients, CamServer also had the capability to send and receive processed data to and from the DataBuffer, making it a key component of the SwissFEL data acquisition system. CamServer is capable of processing generic streams (other from camera sources, such as DataBuffer streams), and can merge and synchronize streams before processing.

Initially deployed on a single server, currently CamServer runs on a Linux cluster of 13 servers and can scale further.

<sup>†</sup> alexandre.gobbo@psi.ch

# THE SKAO ENGINEERING DATA ARCHIVE: FROM BASIC DESIGN TO PROTOTYPE DEPLOYMENTS IN KUBERNETES

Thomas Juerges, SKA Observatory, Jodrell Bank, United Kingdom  
Aditya Dange, Tata Consultancy Services, Pune, India

## Abstract

During its construction and production life cycles, the Square Kilometre Array Observatory (SKAO) will generate non-scientific, i.e. engineering, data. The sources of the engineering data are either hardware devices or software programs that generate this data. Thanks to the Tango Controls software framework, the engineering data can be automatically stored in a relational database, which SKAO refers to as the Engineering Data Archive (EDA). Making the data in the EDA accessible and available to engineers and users in the observatory is as important as storing the data itself. Possible use cases for the data are verification of systems under test, performance evaluation of systems under test, predictive maintenance and general performance monitoring over time. Therefore we tried to build on the knowledge that other research facilities in the Tango Controls collaboration already gained, when they designed, implemented, deployed and ran their engineering data archives. SKAO implemented a prototype for its EDA, that leverages several open-source software packages:

- with Tango Controls' HDB++
- the Timescaledb time series database
- and Kubernetes at its core.

In this overview we will answer the immediate question "But why do we not just do, what others are doing?" and explain the reasoning behind our choices in the design and in the implementation.

Keywords: TANGO, reporting, data visualization, open source, SCADA

## INTRODUCTION

The Square Kilometre Array Observatory (SKAO) [1] with its headquarters in Jodrell Bank, UK, is currently constructing two large-scale radio telescope arrays, SKA-Mid (under construction)<sup>1</sup> [2, 3] in South Africa and SKA-Low (under construction)<sup>2</sup> [2, 4] in Western Australia.

During construction, commissioning, and operation engineers, operators, and scientists will need to inspect a range of non-scientific data with a range starting as simple as temperature and humidity to highly specific ones like a SKA-Mid

<sup>1</sup> SKA-Mid will be a radio telescope array consisting of 192 fully steerable dishes: 13 315 m SKA dishes and the already operating 6413.5 m Meerkat dishes. The receivers are sensitive at frequencies between 350 MHz and 15.4 GHz. The geographical distribution of the dishes allows for maximum baselines of 150 km.

<sup>2</sup> SKA-Low will be an aperture array radio telescope that consists of 512 stations each with 256 log-periodic dipoles. The dipoles are sensitive at frequencies between 50 MHz and 350 MHz. The configuration of the stations and their geographical distribution allow for baselines between 0.7 km and 70 km at a nominal frequency of 140 MHz.

receptor's delay coefficient. All this data is required to be handled in a uniform manner. This is where the EDA [5] software comes into the picture.

To emphasize, even during the ongoing construction phase, the already existing software that is continuously tested and integrated in Kubernetes [6] clusters and the prototype hardware generate engineering data. This data has at this moment value for debugging the software, its performance, and finding mismatches between how the hardware behaves and the described behaviour in its interface control documents. When a simple line of log cannot tell the whole story, engineering data is able to complete the picture. Engineering data has the power to provide, even independent of any logs, a realistic representation of the current state of a system.

Naturally had SKAO's initial EDA design aged somewhat in the fast moving software technology age and was ready for a fresh look. Therefore SKAO invited partner facilities in the Tango Controls collaboration to explain and demonstrate their EDAs in order to modify the existing design for the SKAO EDA with the help of the already existing knowledge about newer technologies.

The demonstrations were very enlightening and it became almost immediately clear that the existing design for SKA's EDA would really benefit from a fresh look, replacing technologies that had been shown to not work for the use case at hand as had been proven at the partner facilities and also make use of technologies that SKAO had adopted in the meanwhile.

## EARLY IDEAS, THE DESIGN AND THE ARCHITECTURE

The SKAO's EDA design was based on the Tango Controls [7] HDB++ archiver [8], consisting of the HDB++ Configuration Manager [9] and the HDB++ Event Subscriber [10], an EDA Controller and Cassandra [11] as the EDA database back end (as shown in Fig. 1).

This design focused on addressing the functional requirements and did not consider the deployment aspect. One reason was the expected shift in software deployment from bare metal deployments to cloud-based ones.

### *Evolution Over Time*

The geographical distribution of the observatory's computers and clusters already added significant complexity to deployment of, running and maintaining the software needed to operate the entire observatory with the two telescopes. This informed SKAO's decision to containerize all software that does not have to run on bare metal and also deploy

# DEVELOPING A DIGITAL TWIN FOR BESSY II SYNCHROTRON LIGHT SOURCE BASED ON EPICS AND MICROSERVICE DESIGN

W. Sulaiman Khail\*, P. Schnizer, M. Ries  
Helmholtz-Zentrum Berlin, BESSY, Berlin, Germany

## Abstract

Digital twins, i.e. theory and design tools connected to the real devices and machine by mapping of physics components to the technical correspondents, are powerful tools providing accelerators with commissioning predictions and feedback capabilities. This paper describes a new tool allowing for greater flexibility in configuring the modelling part combined with ease of adding new features. To enable the various components developed in EPICS, Python, C, and C++ to work together seamlessly, we adopt a microservice architecture, with REST API services providing the interfaces between the components. End user scripts are implemented as REST API services, allowing for better data analysis and visualization. Finally, the paper describes the integration of dash and plotly for enhanced data comparison and visualization. Overall, this workflow provides a powerful and flexible solution for managing and optimizing BESSY II digital twins, with the potential for further customization and extension to upcoming machines.

## INTRODUCTION

In recent years, the concept of digital twins has gained significant traction as virtual representations of physical systems, enabling real-time monitoring, analysis, and optimization of intricate processes [1]. This technology has found widespread application across diverse sectors, such as aerospace, smart cities, healthcare, automotive, and energy, due to its ability to enhance efficiency, reduce costs, and improve safety [1]. Within the realm of scientific research, particularly in the domain of Synchrotron Light Sources, there is a burgeoning interest in harnessing digital twins to model and monitor the intricate behavior of these complex systems.

Synchrotron Light Sources, as described by the European Synchrotron Radiation Facility (ESRF), are expansive research facilities generating intense light beams for applications in material science, biology, and drug discovery [2]. These facilities operate through intricate setups involving linear accelerators, booster rings, and storage rings, necessitating seamless coordination of thousands of components to produce high-quality light beams. Given their complexity, digital twins emerge as invaluable tools, enabling real-time modeling and monitoring. By leveraging technologies such as the Experimental Physics and Industrial Control System (EPICS) and microservice design, Synchrotron Light Sources can achieve enhanced flexibility, scalability, and reliability in their operations [3, 4].

\* waheedullah.sulaiman\_khail@helmholtz-berlin.de

This paper presents the development of digital twins for two longstanding Synchrotron Light Sources, namely BessyII and MLS (Metrology Light Source). Utilizing EPICS as the standard interface and employing MongoDB in JSON format for machine structure storage, the implementation integrates REST APIs and Python scripts to facilitate seamless communication between system components. The integration of dash and plotly further enhances data visualization and comparison, offering a robust solution for real-time monitoring and optimization of Synchrotron Light Sources (Ref 5).

In summary, this paper delves into the integration of digital twin technology within the realm of Synchrotron Light Sources, showcasing its potential for revolutionizing the monitoring and optimization processes in these complex scientific facilities. Through the utilization of established frameworks and innovative approaches, the study not only presents a comprehensive solution for current facilities but also lays the groundwork for the customization and expansion of digital twin applications in upcoming machines.

## TOWARDS AN ARCHITECTURE

Domain-driven design is based on the principle that knowledge obtained needs to be cast into business logic and surrounding components. However, this approach may be questionable if the domain of interest is subject to change.

At least one of the authors considers this approach invalid for the calculation of beam dynamics due to the field being impacted by progress in other sciences or paradigm shifts within the domain itself:

- Machines developed currently need to be controlled by complex procedures for first commissioning and periodic recommissioning (4th generation light sources).
- Machine learning requires fast evaluation of the model.
- Implementing propagation code on GPU codes makes calculations feasible for difficult tasks such as the interaction of particles within bunches or in different bunches (referred to as “collective effects”), and optimization of sextupole magnet settings to reduce beam loss (referred to as “dynamic aperture”).

Typical implementations today provide a full package dedicated to calculating beam propagation. This needs to be revisited for appropriate layering to be ready for these requirements, such as obtaining a fully differentiable model. Furthermore, the calculations need to be performed not only for values but also for objects representing dependencies.

# STATUS OF THE EUROPEAN SPALLATION SOURCE CONTROLS

T. Korhonen<sup>1</sup>, European Spallation Source ERIC, Lund, Sweden  
<sup>1</sup>Representing the Integrated Control Systems Division

## Abstract

The European Spallation Source is progressing towards completion in a few years. The control system has been through first commissioning rounds and is now in production use. While development is still going on to reach full functionality, most of the central supporting features are already in place. This paper gives an overview of the current status, the principles we have been following on the way, our use and experience with the central technologies like MTCA.4 and EPICS 7, plus an overview of the next steps. We also look at what was planned and reported in ICALEPCS 2015 and how our system of today compares with ideas of that time, and the evolution from green field project to an operating organization.

## INTRODUCTION

ESS is one of the very few true green-field projects where not only the facility but also the surrounding organisation has been built from scratch. With the project, the control system has gone through a series of growing pains. This paper will touch some of them, and how have we managed (or not) to resolve the issues that have come up. Also, we look at the overarching goals that we have tried to follow and at core principles that support these goals.

Obviously, the expectation of a new project is that the system is built to take advantage of the technological advances that are available. But this is only momentary and the real goal is to build an infrastructure that can be maintained with reasonable amount of effort, and in a way that allows not only updates but also evolution to implement features that are not known at the time of development. Experience has taught that this is a common occurrence in real life, and also defines the ultimate success of a scientific facility.

Cost efficiency is also an important goal but should be considered with the whole system lifecycle in mind. Of course, project budget and time are limited, but within these limits there are several possible ways to follow.

From these considerations the following principles emerge:

- Strive for clear standards to be followed. Standards are not meant to be set in stone forever, but to create clarity and enable collaboration between units that are often remote and disconnected.
- Support the development process so that skills of people can be optimally utilized
- Try to remove barriers towards updates so that it is easy to keep systems up to date.

Let us look at what has been done to achieve these goals.

## STANDARDIZATION

When embarking on a long project, one has to define the technology base such that it covers the project needs as well as possible. Also, one has to pay a lot of attention to the evolution capability of technologies. Controls systems are work-intense, and the work continues long after all hardware has been installed. In a sense, the work goes on as long as the facility is in use.

### Hardware Platforms

One of the earliest decisions to be made was the main hardware platforms to be used for field I/O and control.

To be useful, the selected standards should cover the majority of the expected use cases. We ended up with a three-layer strategy which has been able to cover a vast majority of use cases we have encountered so far (see Fig. 1).

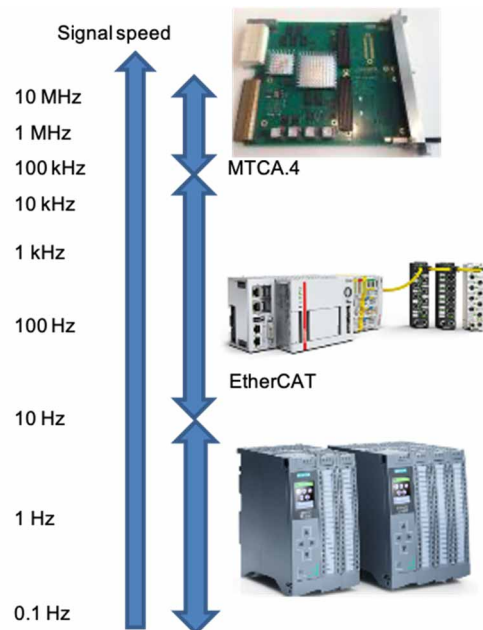


Figure 1. Hardware standards and performance levels.

Hardware costs are directly visible, unlike software and that leads often to closer scrutiny than for software projects. There is also the temptation to lower the upfront cost instead of considering the lifecycle of the decisions.

### MTCA.4

At the point of ESS decision, MTCA [1] could still be considered as an emerging standard. Pioneering work had been done at DESY to develop the MTCA.4 standard to a useable level, by investing time in creating the standard, but also developing products together with industrial partners.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

# CONTROLS AT THE FERMILAB PIP-II SUPERCONDUCTING LINAC\*

Dennis J. Nicklaus, Pierrick Hanlet, Charles King, Daniel McArthur, Richard Neswold  
Fermilab, Batavia, IL, USA

## Abstract

PIP-II is an 800 MeV superconducting RF linear accelerator under development at Fermilab. As the new first stage in our accelerator chain, it will deliver high-power beam to multiple experiments simultaneously and thus drive Fermilab's particle physics program for years to come. In a pivot for Fermilab, controls for PIP-II are based on EPICS instead of ACNET, the legacy control system for accelerators at the lab. This paper discusses the status of the EPICS controls work for PIP-II. We describe the EPICS tools selected for our system and the experience of operators new to EPICS. We introduce our continuous integration / continuous development environment. We also describe some efforts at cooperation between EPICS and ACNET, or efforts to move towards a unified interface that can apply to both control systems.

## PIP-II INTRODUCTION

PIP-II is a new superconducting radio frequency linear accelerator under construction at Fermilab. It will produce a pulsed 800MeV high intensity pulsed H- beam to supply the needs of Fermilab Long Baseline Neutrino Facility experiments and more. The Linac is followed by a beam transport line to deliver the beam to the Booster. Key components of the Linac include:

- Ion Source
- Low Energy Beam Transport
- RF Quadrupole to accelerate to 2.1 MeV
- Medium Energy Beam Transport with a chopper system to form the required bunch structure
- One 162.5 MHz Half Wave Resonator Cryomodule
- Two 325 Single Spoke Resonator Cryomodule (SSR1)
- Seven 325 Single Spoke Resonator Cryomodule (SSR2)
- Eleven Low Beta 650 MHz cryomodules
- Four High Beta 650 MHz cryomodules

Of course, there is a considerable amount of instrumentation and beam controlling magnets in between the variety of cryomodules, as well as infrastructure to supply the cryogenics required to operate the superconducting cavities at 2K.

## CONTROLS OVERVIEW

EPICS has been selected as the central control system for PIP-II. However, the nature of the EPICS toolkit means that there are still many choices to be made for key elements of the system. Furthermore, while PIP-II is a new machine without legacy controls to hinder the adoption of

EPICS, it will feed into the Fermilab Booster accelerator which is controlled by Fermilab's legacy ACNET [1] control system. While it would be technically possible for the two control systems to simply co-exist in our control rooms, operation efficiency is going to require that there is a fair amount of cooperation and unity between PIP-II's EPICS and the rest of the complex's ACNET.

## Modernization Cooperation

Fermilab is also beginning a major modernization effort for the legacy control system, known by the project acronym ACORN (Accelerator Controls Operations Research Network) [2]. ACORN is looking at improving all facets of the control system, from the software frameworks to fieldbus hardware to user interfaces. The later schedule of the modernization effort means that PIP-II has a policy of not depending on any result or software that ACORN will produce. It also adds a slight mental burden to PIP-II controls development because we don't want to make any decisions that are drastically conflicting with likely ACORN directions. However, PIP-II is able to synergistically work with ACORN in some respects, taking advantage of research done by ACORN. One instance of this was when both projects wanted to evaluate software technologies for web-based controls applications. PIP-II began prototyping after ACORN selected a combination of Dart language and Flutter for application user interface building.

## EPICS COMPONENTS

### IOCs

PIP-II will have many IOCs (Input Output Controllers) built off EPICS base as the data collection end of the control system. These IOCs will be built with EPICS version 7 to enable the PV Access protocol. In a major change for Fermilab, we don't foresee any VME-based embedded computers running a real time operating system. While it is an architecture that our hardware and software engineers are very comfortable with, the price/performance of these embedded computer cards has become extremely poor in recent years, so we no longer see a use case for these systems. We expect that most of our IOCs will be "soft" IOCs, server based and communicating with field hardware over ethernet. Many of the subsystems will be FPGA (including "system on a module" (SOM) style subsystems). We may eventually embed an IOC into some of these subsystems, but at this moment, most of our plans call for the IOC to be on a remote server, using some other (non-PV Access) Ethernet protocol to communicate with the FPGA systems.

For higher bandwidth instrumentation systems, e.g., BPMs (Beam Position Monitors), BLMs (Beam Loss Monitors), current monitors, etc. we are planning to use MicroTCA bus-based systems. But here again, since the MicroTCA bus and the data acquisition cards can support

\* This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

# SKA PROJECT STATUS UPDATE

N. P. Rees\*, SKA Observatory, Jodrell Bank, UK  
on behalf of the The SKA Software Collaboration†

## Abstract

The SKA Project is a science mega-project whose mission is to build the world's two largest radio telescopes with sensitivity, angular resolution, and survey speed far surpassing current state-of-the-art instruments at relevant radio frequencies. The Low Frequency telescope, SKA-Low, is designed to observe between 50 and 350 MHz and will be built at Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory in Western Australia. The Mid Frequency telescope, SKA-Mid, is designed to observe between 350 MHz and 15 GHz and will be built in the Meerkat National Park, in the Northern Cape of South Africa. Each telescope will be delivered in a number of stages, called Array Assemblies. Each Array Assembly will be a fully working telescope which will allow us to understand the design and potentially improve the system to deliver a better scientific instrument for the users. The final control system will consist of around 2 million control points per telescope, and the first Array Assembly, known as AA0.5, is being delivered at the time of ICALEPCS 2023.

## INTRODUCTION

This is the third in a series of SKA status papers presented at ICALEPCS meetings and together they provide a picture of the evolution of the SKA Software. At the time of the 2021 paper formal construction had just commenced, and we were in the process of issuing the first contracts. This paper covers the evolution since this time, but it starts with a brief historical overview.

## HISTORY

### Early Years

While the SKAO was officially created as an international observatory on 15 January 2021, the concept dates back to the early 1990's when astronomers proposed the idea of tracing the history of the Universe by mapping its most abundant element: Hydrogen (see, for example, Wilkinson [1]). Specifically, they proposed to observe the H1 emission line from its rest frequency of 1420 MHz to red-shifts of more than 10, thereby observing it from just after the beginning of the universe to the present day.

To achieve this goal required a telescope of unparalleled sensitivity, with a collecting area approaching a square kilometre, and a frequency range from less than 100 MHz to a few GHz. Of course, once these basic parameters were described, scientists realised this instrument would be capable of observing a huge number of other phenomena. This

significantly enhanced the science case and scientific interest in the project, but it also complicated the design, and necessitated the construction of two telescopes - one optimised for mid frequencies between 0.35 and 15 GHz (being built in South Africa) and the other between 50 and 350 MHz (being built in Australia).

The main technical difference between the two telescopes is that at the higher frequencies, it is still economic to provide the collecting area with traditional dishes - in the baseline design (known as Array Assembly 4) there are 133 15 m diameter dishes that will be integrated with the existing 64 13.5 m MeerKAT dishes on the same site to create an array of 197 dishes.

For the SKA-Low telescope, it is uneconomic to build dishes large enough that will be efficient at collecting at the longer wavelengths (50 MHz is about 6 m), so the SKA-Low design has 512 "stations" each of which is a collection of 256 phased wire log-periodic antennas. Each station is about 40 m in diameter, and is roughly equivalent to a dish antenna of the same diameter, but is much cheaper to build.

The project has always been seen as international in scope, and in 2013 it was agreed to establish an Inter-governmental Organisation to manage construction and operate the telescopes. This was about the same time when detailed design began, and so from 2013 to the start of construction in 2021 development followed two paths - one technical to deliver the design and the other governance, to deliver the organisation and funding structure.

### Pre-Construction

Up until 2019 the technical development was known as "pre-construction" and international consortia developed detailed designs on different aspects of the telescopes. These culminated in a series of Critical Design Reviews coordinated by the SKA Organisation. Between the completion of these CDR's and the formal start of construction was a "bridging" phase. In software, the design phase was far from ideal because it was largely a paper exercise, and the design consortia did limited prototyping, and consequently limited practical design validation. This changed in the bridging phase, where we pivoted to the development of processes with the adoption of the Scaled Agile Framework [2], and code creation, but we are still evolving the design as we learn about the system.

### Construction

At the time of the last paper construction had only just commenced. The plan is to build the two telescopes simultaneously, and each is to be delivered in stages, as successively larger and more powerful interferometers. Each stage is known as an "Array Assembly (AA)", and the basic parame-

\* Nick.Rees@skao.int

† See Appendix



# THE CONTROLS AND SCIENCE IT PROJECT FOR THE SLS 2.0 UPGRADE

A. W. Ashton<sup>†</sup>, E. Zimoch, X. (Marie) Yao, S. Fries, H-H. Braun  
 Paul Scherrer Institute, Switzerland

## Abstract

On the 30<sup>th</sup> of September 2023 the dark time for the Swiss Light Source (SLS) facility began. SLS 2.0, a project to upgrade the storage ring and selected beamlines now enters the construction phase. With the advent of the next generation of synchrotron light sources, called diffraction-limited storage-rings (DLSRs), that yield an emittance and brightness improved by up to two orders of magnitude, it has become equally imperative to upgrade the SLS to accommodate the new developments. The storage ring is undergoing its upgrade in 2023/2024 with a planned reduction of emittance by a factor of 40, before the facility returns to user operations in 2025.

## INTRODUCTION

The SLS [1] (Fig. 1), operational since 2001, has remained one of the leading examples of third-generation storage-ring technology for two decades. However, the increasing scope and impact of the uses of synchrotron light sources in almost all areas of the natural and engineering sciences, improvements in source and instrument technology generally, and the advent of diffraction-limited storage-rings (DLSRs), highlighted that the SLS would benefit from undergoing a comprehensive upgrade to remain competitive and perform cutting-edge science [2].



Figure 1: The main entry to the SLS building on the morning of the 30<sup>th</sup> of September 2023, just prior to dark time.

The SLS 2.0 upgrade project will provide a significant increase in brightness by replacing the current magnet lattice of the storage ring by a new multi-bend achromat (MBA) magnet structure. This, combined with advanced hardware and instrumentation, will enhance the performance of all techniques currently practiced at the SLS by up to three to four orders of magnitude in some cases, while heralding, on the one hand, new and game-changing sources and, on the other, new and innovative techniques.

Table 1: Main SLS Specifications after SLS 2.0 Upgrade

Parameter	Value
Circumference	288m
Total straight length	~80m
Beam current	400mA
Source point position [shifts]	< 70mm
Lattice type	7 bend achromat
Emittance	157pm
Energy	2.7 GeV

The upgrade focuses on the transformation of the storage ring lattice to MBA technology and the upgrade of the selected beamlines and end-stations to take full advantage of the increased brightness of the machine (Table 1).

The accelerator upgrade will incorporate new technologies in the SLS, such as superconducting magnets, which will extend the X-ray spectra to higher energies. In addition, permanent magnet material will be used for many of the ambient temperature magnets, resulting in an important reduction in power consumption. A crucial feature for the MBA design is a significant reduction in the vacuum chamber cross-section, achieved thanks to the use of a non-evaporable getter coating on the chamber surface. One truly original feature of the new lattice is the use of “reverse-bend” magnets, which play an important role in the brightness increase. The schedule for the upgrade and its key milestones are in Table 2.

Table 2: SLS 2.0 upgrade timetable. (\*a total budget of 130MCHF and ~4% for controls and IT, excluding existing staff resources.)

Milestone	Date
1 Definition of ring lattice	30/06/2020
2 Beamline and positions defined	30/09/2020
3 SLS 2.0 funding secured*	01/01/2021
4 Ready for dark time	17/10/2022
5 Start of dark time	30/09/2023
6 Tunnel closure	20/12/2024
7 First beam available	01/05/2025
8 Start of user operation on first beamlines	01/08/2025
9 Start of shutdown phase 2	21/12/2025
10 Re-start user operation	01/06/2026

The upgrades of the accelerator and the beamlines and PSI’s leading role in the development of complementary technology (e.g. insertion-device design, x-ray detectors,

<sup>†</sup> alun.ashton@psi.ch

# HOW ACCURATE LASER PHYSICS MODELING IS ENABLING NUCLEAR FUSION IGNITION EXPERIMENTS\*

K. P. McCandless<sup>†</sup>, R. H. Aden, A. Bhasker, R. T. Deveno, J. M. Di Nicola, M. A. Erickson, T. E. Lanier, S. A. McLaren, G. J. Mennerat, F. X. Morrissey, J. Penner, T. Petersen, B. Raymond, S. E. Schrauth, M. F. Tam, K. C. Varadan, L. J. Waxer  
Lawrence Livermore National Laboratory (LLNL), Livermore, USA

## Abstract

This last year we achieved an important milestone by reaching fusion ignition at Lawrence Livermore National Laboratory's (LLNL) National Ignition Facility (NIF), a multi-decadal effort involving a large collaboration. The NIF facility contains a 192-beam 4.2 MJ neodymium-doped phosphate glass laser (around 1053 nm) that is frequency converted to 351 nm light. To meet stringent laser performance required for ignition, laser modeling codes including the Virtual Beamline (VBL) and its predecessors are used as engines of the Laser Performance Operations Model (LPOM). VBL comprises an advanced nonlinear physics model that captures the response of all the NIF laser components (from IR to UV and nJ to MJ) and precisely computes the input beam power profile needed to deliver the desired UV output on target. NIF was built to access the extreme high energy density conditions needed to support the nation's nuclear stockpile and to study Inertial Confinement Fusion (ICF). The design, operation and future enhancements to this laser system are guided by the VBL physics modeling code which uses best-in-class standards to enable high-resolution simulations on the Laboratory's high-performance computing platforms. The future of repeated and optimized ignition experiments relies on the ability for the laser system to accurately model and produce desired power profiles at an expanded regime from the laser's original design criteria.

## INTRODUCTION

Here at NIF we have now officially repeated the achievement of fusion ignition in a laboratory setting - once in December 2022 and this past summer in July 2023. This feat was long in coming and relies on many teams working together to push the frontiers of high energy density science, laser performance and target design. This paper will explore some of the details and challenges on the laser performance team and highlight a few cases where we have tightened up the accuracy to deliver better quality laser pulses at the target.




Figure 1: Artist rendition of 192 laser beams of NIF entering the target hohlraum to start heating the walls and cause an X-ray bath which begins compressing the fuel pellet (grey sphere) triggering a runaway fusion ignition process. (left) In a fusion reaction, the nuclei of hydrogen deuterium and tritium are forced together by extremes of temperature and pressure and fuse to form a helium nucleus. In the process some of the mass is converted to energy and released as neutrons (right).

## BACKGROUND

At its core, fusion ignition refers to the point at which a controlled fusion reaction becomes self-sustaining, releasing an immense amount of energy through the fusion of atomic nuclei, the same process that powers the sun (see Fig. 1). Achieving and studying fusion ignition in a controlled environment has several key implications: providing a clean and abundant energy source, advancement in our understanding of fundamental physics, plasma behaviours and high-energy processes, and it is a key factor in avoiding a return to underground nuclear weapons testing as a stable fusion-based scientific platform is needed to assure the nuclear weapons stockpile in the United States remains safe, secure, and reliable [1].

At NIF we use a laser driven technique to achieve fusion ignition by a process termed "Inertial Confinement Fusion". ICF is a method of achieving controlled nuclear fusion by using intense pulses of energy to compress and heat a small target containing fusion fuel. In the context of ICF, the term "inertial confinement" refers to the fact that the fusion fuel is compressed and heated through the inertia of the material surrounding it, rather than by using external pressure. The energy released from the fusion reactions can potentially be harnessed for various applications including electricity generation and survivability experiments (see Fig. 1)

\* This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344

<sup>†</sup> mccandless2@llnl.gov LLNL-CONF-854863

# A DIGITAL TWIN FOR NEUTRON INSTRUMENTS

S. Nourbakhsh\*, Y. Le Goc, P. Mutti  
Institut Laue-Langevin, Grenoble, France

## Abstract

Research infrastructures and facility users have manifested an increasing interest in data from virtual experiments. Simulated data are an important ingredient for instrument scientists to develop and optimize current and future instruments, for external users to train in the usage of the instrument control system (ICS), for scientists in quantifying and reducing instrumental effects when analysing acquired data. Furthermore large sets of simulated data are also a necessary ingredient for the development of surrogate models for faster and more accurate simulation, reduction and analysis of the data.

The development of a Digital Twin (DT) of an instrument can answer such different needs with a single unified approach wrapping in a user-friendly envelop the knowledge about the instrument physical description, the specific of the simulation packages and their interaction, and the high performing computing setup.

In this article we will present the general architecture of the DT prototype developed at the Institut Laue-Langevin (ILL) in the framework of the PaNOSC European project in close collaboration with other research facilities (ESS, European XFEL). The communication patterns (based on ZeroMQ) and interaction between the NOMAD control system, simulation software (McStas), instrument description and configuration, process management (Cameo) will be detailed.

The adoption of FAIR Principles (Findability, Accessibility, Interoperability, and Reusability) for data formats and policies in combination with open-source software make it a sustainable project both for development and maintenance in the mid and long-term.

## INTRODUCTION

The development of a DT at the ILL can play a significant role in better preparation of experiments ahead of time and hence either reduce the amount of beam time needed for each experiment or an improvement in the acquired data with better suited instrument settings for the specific experiment. The objective is to enrich the offer of user tools with the possibility to run a virtual experiment with an ILL instrument. ILL's DT is designed to be used by users with no knowledge about simulations, providing a user experience almost unchanged with respect to what they get when running a real experiment. All technicalities are hidden from the user behind the familiar ICS interface where both the instrument configuration and acquisition workflow are set up. Simulated data are provided by the ICS in the same format as the real data, with no additional step to be taken before

\* nourbakhsh@ill.fr

undergoing the reduction and analysis. A further requirement, in order to have a long term sustainability of the DT, is to use of state-of-the-art simulation software adopted by the neutron community and to make publicly available the instrument description. The simulation experts' feedback and contribution can play a major role in the improvement of the accuracy of the simulations, while keeping orthogonal the development and maintenance of the DT implementation.

The DT of an instrument would be able to serve for several purposes: training new ILL users to the ICS or current users to a particular instrument configuration settings and capabilities; studying and optimizing of instrument settings for specific figure of merit; improving analyses with better understanding of some background sources and uncertainties; or enriching proposals for demanding beam time with results from simulated data with the specific instrument.

## DIGITAL TWIN IMPLEMENTATION

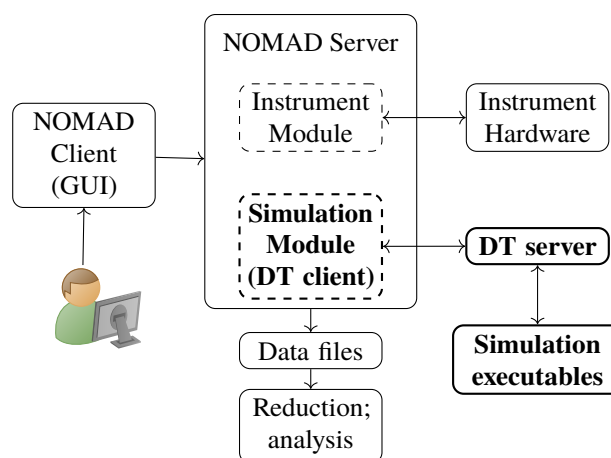


Figure 1: Overview of the logical elements and their communication links. Highlighted in bold are the DT specific parts.

The DT is composed by three logical components: a simulation executable provided by the chosen simulation software, a client application programming interface (API) used by the ICS (acting as a DT client) and a central unit linking those two parts (DT server). An overview of the logical elements and communication links in the ILL's architecture is shown in Fig. 1. NOMAD [1] is the instrument control software developed and used at ILL for the data acquisition and instrument control. Its core is represented by a C++ server that communicates with the instrument hardware with specific modules. The interaction with the user happens via NOMAD client (JAVA application). The integration of the DT into NOMAD is achieved adding a module to the ICS to communicate with the DT server and replacing the com-

# A PHYSICS-BASED SIMULATOR TO FACILITATE REINFORCEMENT LEARNING IN THE RHIC ACCELERATOR COMPLEX\*

L. K. Nguyen<sup>†</sup>, K.A. Brown<sup>1</sup>, M.R. Costanzo, Y. Gao, M. Harvey, J. P. Jamilkowski, J. T. Morris, V. Schoefer, Collider-Accelerator Dept., Brookhaven National Laboratory, Upton, NY, USA  
<sup>1</sup>also at ECE Dept., Stony Brook University, Stony Brook, NY, USA

## Abstract

The successful use of machine learning (ML) in particle accelerators has greatly expanded in recent years; however, the realities of operations often mean very limited machine availability for ML development, impeding its progress in many cases. This paper presents a framework for exploiting physics-based simulations, coupled with real machine data structure, to facilitate the investigation and implementation of reinforcement learning (RL) algorithms, using the longitudinal bunch-merge process in the Booster and Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory (BNL) as an example. Here, an initial fake wall current monitor (WCM) signal is fed through a noisy physics-based model simulating the behavior of bunches in the accelerator under given RF parameters and external perturbations between WCM samples; the resulting output becomes the input for the RL algorithm and subsequent pass through the simulated ring, whose RF parameters have been modified by the RL algorithm. This process continues until an optimal policy for the RF bunch merge gymnastics has been learned for injecting bunches with the required intensity and emittance into the Relativistic Heavy Ion Collider (RHIC), according to the physics model. Robustness of the RL algorithm can be evaluated by introducing other drifts and noisy scenarios before the algorithm is deployed and final optimization occurs in the field.

## INTRODUCTION AND MOTIVATION

Interest in machine learning (ML) for use at particle accelerator facilities has rapidly expanded over the years, and accelerator scientists and engineers engaging in ML continue to identify and deliver on important applications [1]. However, the time and resources needed for ML development and model training can be prohibitive. At Brookhaven National Laboratory (BNL), for example, the bunch merge process in the Booster and Alternating Gradient Synchrotron (AGS) rings has been identified as a potentially high-reward area for ML optimization due to the criticality of good bunch merging to operations. Despite this, competing priorities and lack of machine availability have impeded investigations. In particular:

1. Booster & AGS are part of the accelerator chain for multiple programs, and they are often in operational use when not supplying the Relativistic Heavy Ion Collider (RHIC) with beam. Real machine time for ML development is therefore very hard to come by.

2. Any machine downtime is generally allotted to maintenance and/or needed repairs. Meanwhile, part of the ML development cycle is purely software-related (e.g., debugging). This makes real machine time for ML development expensive both in terms of financial costs and opportunity costs.




Figure 1: Real mountain range data for a bunch merge in Booster.

Some ML approaches, such as reinforcement learning (RL), do not learn machine parameters, but rather environments—making them amenable to other development paths. The importance and under-explored benefits of RL for accelerator optimization problems has already been recognized by the community [1]. We therefore pursued a framework for investigating RL optimization by replacing the accelerator, attendant diagnostics, and controls with a physics-based simulator mimicking real machine data structures and communications. Such a simulator was created for Booster and AGS using the Python programming language.

## ENVIRONMENT TO SIMULATE

Bunches merge in Booster for injection into AGS, and bunches merge in AGS for injection into RHIC. To diagnose a merge, a wall current monitor (WCM) is used. The WCM generates a voltage vs. time signal in response to passing bunches. Subsequent signal traces are stacked on an oscilloscope to create a so-called mountain range plot; see Fig. 1 for an example from Booster. A certain number of accelerator periods separate each trace, and this number can vary depending on the merge. It is typically between 100 and 200 turns.

In addition to time between acquisitions, captured mountain range data typically varies with respect to oscilloscope

\* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.

<sup>†</sup> lnguyen@bnl.gov

# PYTHON LIBRARY FOR SIMULATED COMMISSIONING OF STORAGE-RING ACCELERATORS\*

L. Malina<sup>†</sup>, I. Agapov, E. Musa J. Keil, and B. Veglia

Deutsches Elektronen-Synchrotron, 22607 Hamburg, Germany

N. Carmignani, L. R. Carver, L. Hoummi, S.M. Liuzzo, T. Perron and S. White  
ESRF, Grenoble, France

T. Hellert, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

## Abstract

Simulations of the commissioning procedure became vital to the storage-ring lattice design process. The achievable tolerances on lattice imperfections, such as equipment misalignments or magnet gradient errors, would, without correction, prohibit reaching the design parameters. We present a Python library which includes an extensive set of error sources in the accelerator lattice and provides a variety of correction algorithms to commission a storage ring. The underlying beam dynamics simulations are performed with pyAT. This project builds upon previous works and expands them in the direction of realistic control room experience and software maintainability. The performance is demonstrated using example commissioning studies, and further development plans are discussed.

## INTRODUCTION

Fourth-generation storage ring (SR) synchrotron light sources aim at improving the hard X-ray source brightness with a reduction of their horizontal electron beam emittance, obtained with the extensive use of Multi-Bend Achromat (MBA) schemes [1–5]. Such compact lattices inherently feature strong focusing elements and strong nonlinear magnets, yielding highly nonlinear beam dynamics. The combination of strong nonlinearities and compactness, which leave less space for correction and compensation, makes these schemes highly sensitive to alignment and magnet strength errors. Operability of newly built or upgraded ultra-low emittance SR is conditioned to accurate understanding of imperfections, analysis of turn-by-turn beam data prior to the application of efficient beam orbit and optics correction methods during beam measurements [6–8].

## Commissioning Simulations

Simulations have become indispensable in this process, as they allow for a meticulous examination of the complex interplay between various factors that can affect the performance of these cutting-edge accelerators [9–17]. In response to this challenge, we have developed a Python library that offers a comprehensive suite of error sources and a wide array of correction algorithms. This library can help operators and accelerator physicists to commission storage rings effectively.

The aim is a start-to-end simulation of the machine commissioning beginning from first-turn trajectory correction, progressing to orbit correction, and culminating in lattice correction and coupling adjustment. Commissioning simulations enable the development of efficient and well-structured procedures for beam commissioning. These procedures encompass beam injection, acceleration, and beamline setup. Next-generation synchrotron radiation sources often represent significant investments, and minimising downtime is crucial. Commissioning simulations help identify potential bottlenecks, inefficiencies, or issues that could lead to extended downtime. By addressing these concerns in advance, facilities can reduce their so-called “dark time”.

In this paper, we present a new computational tool developed to meet the requirements of 4<sup>th</sup> generation machines and its application to the PETRA IV, ALS-U and ESRF-EBS SR lattices. The Python Simulated Commissioning, shortened in pySC, leverages modern open-source scientific libraries such as NumPy [18] and SciPy [19] and bases its lattice manipulation and simulations on the Python-based Accelerator Toolbox (pyAT) [20, 21], which gains momentum both in usage and support. We aim to conduct realistic commissioning simulations for electron storage rings, encompassing a wide range of error sources while accurately modelling beam diagnostics. The simulation allows for tests of the applications to achieve stored beam and adjust properties of the electron beam, such as orbit, tune and optics. The project builds mainly up on the *Toolkit for Simulated Commissioning* (SC) [11, 22] (written in Matlab and already used in several laboratories) and the tools developed at the ESRF [23], which are based on pyAT. Following a discussion on synergies of the two projects, translation of the SC library and later inclusion of ESRF’s tools, most importantly analytical formulas for response matrices, was considered the most efficient further development. The development began with translating the SC library into Python with the support of the Artificial Intelligence (AI) tool, ChatGPT [24].

## TRANSLATION BY CHATGPT

In the translation, ChatGPT’s engine code-davinci-002 was used. Due to limitations on the length of prompts, the SC code was first stripped of all comments on separate lines. It was submitted to the engine along with a description of the input (Matlab code), the action (conversion) and the requested output (Python code) followed by the first lines of the assumed Python code, i.e. `import numpy` and `matplotlib`

\* This project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No. 871072

<sup>†</sup> lukas.malina@desy.de

# REACT AUTOMATION STUDIO: MODERN SCIENTIFIC CONTROL WITH THE WEB

W. D. Duckitt. Stellenbosch University, Stellenbosch, South Africa

J. K. Abraham. iThemba LABS, Cape Town, South Africa

G. Savarese, D. Marcatò. INFN Legnaro National Laboratories, Legnaro, Italy

## Abstract

React Automation Studio is a progressive web application framework that enables the control of large scientific equipment through EPICS from any smart device connected to a network. With built-in advanced features such as reusable widgets and components, macro substitution, OAuth 2.0 authentication, access rights administration, alarm-handling with notifications, diagnostic probes and archived data viewing, it allows one to build modern, secure and fully responsive control user interfaces and overview screens for the desktop, web browser, TV, mobile and tablet devices. A general overview of React Automation Studio and its features as well as the system architecture, implementation, community involvement and future plans for the system is presented.

## INTRODUCTION

React Automation Studio (RAS) is a progressive web application (PWA) framework that enables the control of large scientific equipment through the Experimental Physics and Industrial Control System (EPICS) [1].

It was born out of the need for cross-platform and cross-device applications that deliver an instantaneous user experience for the web, desktop, on a mobile phone and on tablet devices.

The initial open-source release [2] addressed many of the challenges faced in creating EPICS applications for these devices. Since then, RAS has been through another three major revisions and is now at V4.0.2.

It supports advanced features such as reusable widgets and components, macro substitution, Open Authorization (OAuth) 2.0 authentication, access rights administration, alarm-handling with notifications, diagnostic probes, loading and saving of process variable (PV) data and archived data viewing.

This enables one to build modern, secure and fully responsive control user interfaces and overview screens for the desktop, web browser, television (TV), mobile and tablet devices, such as the examples show in Figs. 1, 2, 3 and 4.

A general overview of RAS, its features as well as the system architecture, implementation, community involvement and future plans for the system is presented.

## SYSTEM OVERVIEW

React Automation Studio (RAS) is a PWA framework that enables the control of large scientific equipment through EPICS. It is presented as a Git [3] mono-repository [4] with multiple microservices that are container-

Software

User Interfaces & User Experience

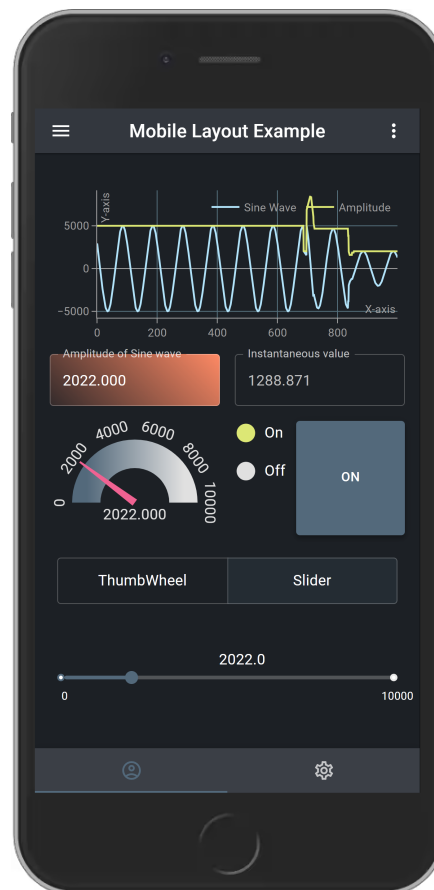


Figure 1: An example of real-time mobile user interface created with RAS.

ized with Docker [5] and orchestrated with Docker Compose [6].

A high level block diagram of the system, showing the information flow between the primary microservices available in the mono-repository is shown in Fig. 5. Each of these microservices are discussed below:

### pvServer

This is the Python [7] process variable server (pvServer). It is based on Flask-SocketIO [8] web application framework and the PyEpics [9] framework to serve the EPICS PVs to clients.

Communication between clients and the pvServer occurs between the data connection wrapper in the client components through Socket-IO [10] and REST application programming interfaces (API). The pvServer handles EPICS Channel Access (CA), authentication, authorisation and

FR2BC001

1643

# A LEAN UX APPROACH FOR DEVELOPING EFFECTIVE MONITORING AND CONTROL USER INTERFACES: A CASE STUDY FOR THE SKA CSP.LMC SUBSYSTEM

V. Alberti\*, INAF-OATs, Trieste, Italy

C. Baffa, E. Giani, G. Marotta, INAF-OA Arcetri, Firenze, Italy

M. Colciago, I. Novak, Cosylab Switzerland, Brugg, Switzerland

G. Brajnik, University of Udine and IDS, Udine, Italy

## Abstract

The Central Signal Processor Local Monitor and Control (CSP.LMC) is a software component that allows the flow of information and commands between the Telescope Manager (TM) and the subsystems dedicated to signal processing, namely the correlator and beamformer, the pulsar search and the pulsar timing engines. It acts as an adapter by specialising the commands and associated data from the TM to the subsystems and by exposing the subsystems as a unified entity while monitoring their status. In this paper, we approach the problem of creating a User Interface (UI) for such a component. Through a series of short learning cycles, we want to explore different ways of looking at the system and build an initial set of UIs that can be refined to be used as engineering UIs in the first Array Assembly of the Square Kilometre Array. The process heavily involves some of the developers of the CSP.LMC in creating the dashboards, and other ones as participants in informal evaluations. In fact, the opportunities offered by Taranta, a tool to develop web UIs without needing web-development skills, make it possible to quickly realise a working dashboard that can be promptly tested. This also supports the short feedback cycle advocated by a Lean UX approach and maps well in a bi-weekly sprint cadence. In this paper, we will describe the method and present the results highlighting strengths and pain points where faced.

## INTRODUCTION

The SKA Observatory (SKAO) occupies a prominent role in the framework of modern, ambitious scientific projects [1]. The two telescopes that comprise the observatory will provide the scientific community with powerful magnifying glasses to observe the universe at radio frequencies. Being a complex system, the SKA project adopted a staged programme to incrementally deliver the telescope and reduce the risk of not identifying issues promptly, therefore providing the scientific community with a sub-optimal instrument. The goal of the first stage, Array Assembly 0.5 (AA0.5), which is planned to last until the end of next year, is to demonstrate a working implementation of the architecture and supply chain as early as possible. The system being verified consists of 4 dishes, 6 stations, and all the necessary infrastructure positioned in the Karoo and Murchinson deserts, respectively. The telescope software delivery is coordinated by

leveraging the Scaled Agile Framework (SAFe) [2], whose 3-month heartbeat<sup>1</sup> helps synchronise the work developed by about thirty teams. As systems grow in complexity, the need for clear, intuitive, and efficient user interfaces becomes of paramount importance, especially in sectors where precision and rapid responses are critical. In the context of growing agile methodologies, Lean UX has emerged as an approach that blends product and interface design processes into iterative cycles. It integrates feedback loops and constant iteration, ensuring that the UI evolves in tandem with user needs and system requirements [3]. This study presents an exploration of Lean UX principles applied to the CSP.LMC subsystem's Monitoring and Control UIs. Key to this implementation was the utilisation of Taranta [4], a tool that allows for creating web-based GUIs for the TANGO Control System [5] devices. The Lean UX method had to be tuned to take into consideration the specific context of the study: the final user developed the dashboards himself, thanks to the opportunity given by Taranta; the team involved isn't a UI team but rather the control software development team; the selected approach limited the initial scope of the dashboards, with plans to increase it later on. Through this paper, we aim to detail our journey, methodologies employed, and lessons learned. Our hope is that, by sharing our experiences, we can offer insights and facilitate the use of Lean UX principles and methods in the context of other complex control systems and teams developing them.

## CSP.LMC

The Central Signal Processor (CSP) is a core component of the SKA software, responsible for processing data received from the antennas in order to enable further scientific analysis [6]. It comprises three primary instruments, hereafter referred to as "subsystems", that are devoted to specific data processing:

- The Correlator and Beam Former (CBF) processes raw antenna data to produce visibility<sup>2</sup>.
- The Pulsar Search (PSS) identifies potential candidates for pulsar research.

<sup>1</sup> SAFe bases its cadence on periodic events among which there are Program Increments (PIs): a PI is a time frame during with an Agile Release Trains plans, and progressively releases a working system. the length of a PI can vary between 8-12 weeks, but in the SKA project case it has been extended to 3 months.

<sup>2</sup> Visibilities are complex flux measurements in spatial frequency space

\* valentina.alberti@inaf.it

# TARANTA PROJECT - UPDATE AND CURRENT STATUS

Y. Li\*, V. Hardion, M. Eguiraun, J. Forsberg, M. Leorato  
MAX IV Laboratory, Lund, Sweden  
D. Trojanowska, M. Gandor, S2Innovation, Kraków, Poland  
M. Canzari, INAF-OAAB, Teramo, Italy  
V. Alberti, INAF-OATs, Trieste, Italy  
H. Ribeiro, Atlar Innovation, Portugal  
A. Dubey, Persistent Systems, Pune, India

## Abstract

Taranta, developed jointly by MAX IV Laboratory and SKA Observatory, is a web based no-code interface for remote control of instruments at accelerators and other scientific facilities. It has seen a great success in system development and scientific experiment usage. In the past two years, the panel of users has greatly expanded. The first generation of Taranta was not able to handle the challenges introduced by the user cases, notably the decreased performance when a high number of data points are requested, as well as new functionality requests. Therefore, a series of refactoring and performance improvements of Taranta are ongoing, to prepare it for handling large data transmission between Taranta and multiple sources of information, and to provide more possibilities for users to develop their own dashboards. This article presents the status of the Taranta project from the aspects of widgets updates, packages management, optimization of the communication with the backend TangoGQL, as well as the investigation on a new python library compatible with the newest python version for TangoGQL.

In addition to the technical improvements, more facilities other than MAX IV and SKAO are considering to join Taranta project. One workshop has been successfully held and there will be more in the future. This article also presents the lesson learned from this project, the road map, and the GUI strategy for the near future.

## INTRODUCTION

A web-based Graphic User Interface (GUI) application tailored for big research facilities represents a pivotal tool in advancing scientific research and experimentation within this specialized field. In this academic context, we will introduce a web-based GUI application that is involved in big research facilities and plays a critical role in facilitating and enhancing various aspects of scientific experiments and data visualization.

Taranta serves as a bridge between a Tango ecosystem and the researchers who utilize its resources. It provides an intuitive and accessible platform through web browsers, simplifying experiment setup, data acquisition, and post-processing tasks. It helps the end users to build a user oriented software. This article will explore the improvement on the fundamental components of Taranta in the synchrotron

and radio telescope context, emphasizing their roles in device debugging, beamline control, and experiment management.

Taranta, initially named Webjive, was started in 2018 [1] at MAX IV shortly thereafter joined by SKAO. This web application aims at providing a no-code interface for end-users to easily build up their own GUI with minimal coding skills [2]. In 2019, this project was renamed to Taranta and in 2021, officially entered in the Tango Controls Collaboration [3] portfolio. Since then, Taranta has undergone a continuous process of development and exploration, with the aim of incorporating new features and optimizing the user experience. As the spectrum of user requirements continues to broaden, the expansion of current context of Taranta becomes foreseeable.

After these years of development and learning from users' experiences, Taranta version 2.x was released. This article delves into the limitations have been seen on previous release of Taranta (version 1.x) and the improvements achieved in Taranta version 2.x, specially the enhanced frontend performance in handling large number of data points. The improvements on data transmission and widgets are also highlighted, affording researchers the ability to seamlessly monitor equipment, tune parameters, and greater control over experimental processes.

## FRONTEND IMPROVEMENT

### *Refactoring on Data Management*

The refactoring on the data transmission and management is one of the biggest improvement of this second version. Taranta currently consists of two views: Devices and Dashboards. In Devices, a generic view of all devices available on the specified Tango database is provided, where users can navigate through the device tree and interact with a selected device. An example of the device viewer is shown in Fig. 1. One can read the device's current state and interact with its attributes and commands, and access the user actions on this tango device through Taranta.

In the Dashboards view, a variety of widgets can be utilized to create a graphical user interface for different purposes on an empty canvas. It allows users to easily develop and customize a dashboard and can immediately run the dashboard to interact with Tango devices. Once created, the dashboards can be saved, shared within the same user group and exported to a json file.

\* yimeng.li@maxiv.lu.se



# MICRO FRONTENDS - A NEW MIGRATION PROCESS FOR MONOLITHIC WEB APPLICATIONS

A. Asko\*, S. Deghaye, E. Galatas, A. Kustra, C. Roderick, B. Urbaniec, CERN, Geneva, Switzerland

## Abstract

Numerous standalone web applications have been developed over the last 10 years to support the configuration and operation of the CERN accelerator complex. These applications have different levels of complexity, but they all support hundreds of users for essential activities. A monolithic architecture has been utilised so far, tailoring the standalone applications to specific accelerator needs.

The global GUI technology landscape continues to evolve quickly, with most GUI technologies typically reaching end-of-life within 1-to-5 years. Keeping up-to-date with technologies presents a major challenge for the GUI application maintainers, with larger monolithic applications requiring long migration cycles which impede the introduction of new functionalities during the migration phase.

To tackle the above issues within the CERN Controls domain, a new Micro Frontend architecture has been introduced and is being used to gradually migrate a large and complex AngularJS-based web application to Angular. This paper introduces the new generic architecture, which is not tied to any specific web framework. The development workflow, challenges, and lessons learned so far will be covered. The differences of this approach, particularly when compared to monolithic application technology migrations, will also be discussed.

## INTRODUCTION

Since its inception, web development has been characterized as a rapidly evolving field, continuously advancing with fresh tools and solutions for organizations and businesses. Over time, numerous web applications have been created to facilitate the control, monitoring, and configuration of various components within CERN's accelerator complex. In numerous instances, given the intricacies inherent to a setting like CERN, the development of these applications has extended over long periods, often spanning years.

Given the volatile nature of web technologies and the industry as a whole, these applications would ultimately rely on already outdated technologies, shortly after they were deemed feature-complete. In turn, this would trigger a significant migration effort, transitioning to the latest "standard" web technology stack, a process that could also span several years. Simultaneously, fresh requirements emerging from the accelerator complex sometimes require the development

of new applications, sometimes using yet another technology stack.

The development of these applications in this manner, can lead to being trapped in a repetitive cycle, endeavoring to deliver the necessary functionalities while also remaining aligned with the swiftly progressing market technology landscape. This raises the question, why bother to try to align with modern technologies? There are two main reasons for this:

1. To profit from the global industry investment in modern, supported technologies and deliver solutions that meet user requirements and user experience expectations.
2. To use technologies that recent generations of engineers are both familiar with and aspire to use, and in turn provide them with transferable skills. This aligns with one of CERN's core missions, which entails training and preparing the upcoming generations of engineers to confront the challenges that will inevitably arise in the future.

Despite the aforementioned reasoning, it was clearly desirable to find an alternative approach. Rather than constantly embarking on repetitive, technology-driven, lengthy re-development upgrades, a path was sought to be able to introduce new technologies in a more gradual manner. The aim being to avoid lengthy full application re-development cycles and be able to add new functionality and improvements to existing applications, leveraging new technology in the process. The solution has been to adopt a new architectural approach, known as Micro Frontends, gradually replacing outdated, monolithic applications. Leveraging this new development style has the potential to redefine the landscape of web applications for accelerator controls.

## OBJECTIVES AND GUIDING PRINCIPALS

Currently, CERN accelerator controls rely on more than 30 web applications that have diverse levels of complexity and are based on various technology stacks. To reach a more maintainable situation, it is imperative to tackle this heterogeneity and strive for standardization. Several web application frameworks are currently in use, including AngularJS (which reached end-of-life in January 2022), Angular, and VueJS. This presents an additional challenge for development teams, requiring competence in multiple frameworks and their unique approaches or methodologies. The pursuit of standardization on a single framework will facilitate the concurrent migration of multiple applications, enable more flexible allocation of development teams, and promote the reuse of common solutions across the overall application portfolio.

\* anti.asko@cern.ch,  
stephane.deghaye@cern.ch,  
epameinondas.galatas@cern.ch,  
ajob.kustra@cern.ch,  
chris.roderick@cern.ch,  
bartek.urbaniec@cern.ch

# MAGNET INFORMATION MANAGEMENT SYSTEM BASED ON WEB APPLICATION FOR THE KEK e-/e+ INJECTOR LINAC

M. Satoh<sup>†, 1</sup>, Y. Enomoto<sup>1</sup>, High Energy Accelerator Organization (KEK),  
Accelerator Laboratory, Tsukuba, Japan

T. Kudou<sup>1</sup>, Mitsubishi Electric System & Service Co., Ltd (MSC), Tsukuba, Japan  
<sup>1</sup>also at The Graduate University for Advanced Studies (SOKENDAI),  
Department of Accelerator Science, Tsukuba, Japan

## Abstract

The KEK injector LINAC provides e-/e+ beams to four independent storage rings and a positron damping ring. An accurate information management system of accelerator components is very important since it is used for the development of a beam tuning model. In particular, an incorrect magnet database could cause a large deterioration of the quality of beam emittance. At LINAC, a text-based database has long been used for the device information management of magnet systems. It comprises several independent text files that contain master information for generating EPICS database files and other configuration files required for many LINAC control software programs. In this management scheme, it is not easy for common users except a control software expert to access and update any information. For this reason, a new web application-based magnet information management system is developed with the Angular framework and PHP scripts. In the new system, the magnet information can be easily extracted and modified using any web browser by any user. In this paper, we report the new magnet information management system alongside future prospects.

## INTRODUCTION

The KEK e-/e+ injector LINAC [1-3] established the simultaneous top-up injection into five rings in 2019 to support both SuperKEKB particle collider experiments with the DR, LER and HER rings [4], and photon science experiments at the PF ring and PF-AR. Figure 1 shows the layout of LINAC and the lepton accelerator complex. At the most downstream part of LINAC, the third beam switching yard, one LINAC beam line branches to the beam transport lines with the deflecting pulsed bending magnet, as shown in Fig. 2. We succeeded in improving the efficiency of SuperKEKB collision experiments by more than 200% after the introduction of simultaneous top-up injections [5]. This noble injection scheme became indispensable because the beam lifetime of the SuperKEKB ring is quite short, especially at the LER positron ring, which was less than 10 min in 2021. The basis of this operation arrangement, LINAC gradually improved the injection performance and contributed to the achievement of the world-record collision luminosity of SuperKEKB [6].

For the SuperKEKB project, significant updates have been made to the magnet system at LINAC. In 2017, around 120 pulsed magnets were installed together with a new magnet control system [7, 8]. With a total length of

700 m, the injector LINAC currently employs approximately 600 DC and pulsed magnets for the beam operation.




Figure 1: Layout of the KEK e-/e+ injector LINAC and the lepton accelerator complex.



Figure 2: Photograph of the third beam switching yard at the KEK e-/e+ injector LINAC. The LINAC beam line branches to the beam transport lines with the deflecting pulsed bending magnet.

As a tool for magnet system information management, a database using simple text files has been used for many years since the beginning of the KEKB project. The information contained within this database, such as magnet names, corresponding power supply specifications, and excitation curve parameters, has been used for the many different control system software programs based on the Experimental Physics and Industrial Control System (EPICS) [9]. However, the modification of this database is not very easy particularly for non-control software experts. To improve the usability of database handling, a new web application with a relational database was developed at LINAC.