# 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$^3$ 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

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M01BC002

### *Networks, Servers and Control Rooms*

The network infrastructure is based on a dual star redundant configuration using single mode fiber cables covering the whole 1000 by 400 meters of the ITER platform. Point-to-point copper cables are used for critical nuclear safety functions and magnet quench protection. The infrastructure includes 175 racks of active network equipment and 160 passive connection points close to the local control systems. There is a total of more than 800 inter-building cables and 2500 client connections. The network infrastructure supports multiple Ethernet based networks for different purposes: control and monitoring, time synchronization, distributed real-time control, and high throughput data acquisition. There are dedicated networks for protection and safety.

As the construction of the control building is not yet complete, a temporary infrastructure has been set up to support the integration, commissioning and operation of delivered local control systems. Temporary servers and control rooms have also been installed. The control building with the main control and server rooms and the final infrastructure are expected to go into operation in 2024.

### *Software*

The software stack is based on Siemens STEP7 and TiA Portal development environment for slow controllers and EPICS for fast controllers. Hard real-time control is provided by FPGAs interfacing the fast controllers. EPICS Channel Access and pvAccess are used as communication protocols for control and monitoring, while a dedicated publish subscribe multicast UDP protocol is used for distributed real-time control, UDP/TCP for high throughput data acquisition and Precision Time Protocol (PTP) for time synchronization. The top layer providing the human machine interface (HMI) is based on Control System Studio. High level operation applications software (supervision, automation, configuration, plasma control and data handling functions) runs on central virtual servers. Data is streamed to HDF5 based back-end storage accessible on the intranet. Unified Data Access (UDA [4]) provides a common API for all archived data independent of source. Dedicated applications using this API allow easy access to both archived and live data. The latter has proven to be extremely useful during commissioning, allowing for example the use of a 4G smart phone to monitor real-time data in the field.

For diversity and licensing reasons, protection and safety use WinCC-OA in place of EPICS, while critical nuclear safety functions use hardwired logic without any software.

Early standardization and software distribution have been effective to enforce this architecture and associated technologies for all in-kind contributions. This enforcement is critical in realizing our vision statement: “The ITER control system performs the functional integration of the ITER plant and enables integrated and automated operation.”

## INTEGRATION STATUS

Steady state electrical network and building services were the first in-kind systems to be delivered starting in 2018. In the following years, additional electrical load centers and buildings have been supplied and integrated.

The electrical distribution network consists of four 400 kV to 22 kV transformers connected to the grid. The 22 kV is fed to 3 medium (6 kV) and 6 low (400 V) voltage load centers distributed over the site and operated 24/7. In addition, it supplies a 15 kV network supporting construction and administrative buildings. Figure 1 shows the top level HMI giving a geographical overview of electrical distribution and instant power consumption. The states of the electrical components and the energization of high voltage cables are indicated using animated color coding. To date the electrical distribution network is more than 50% complete.

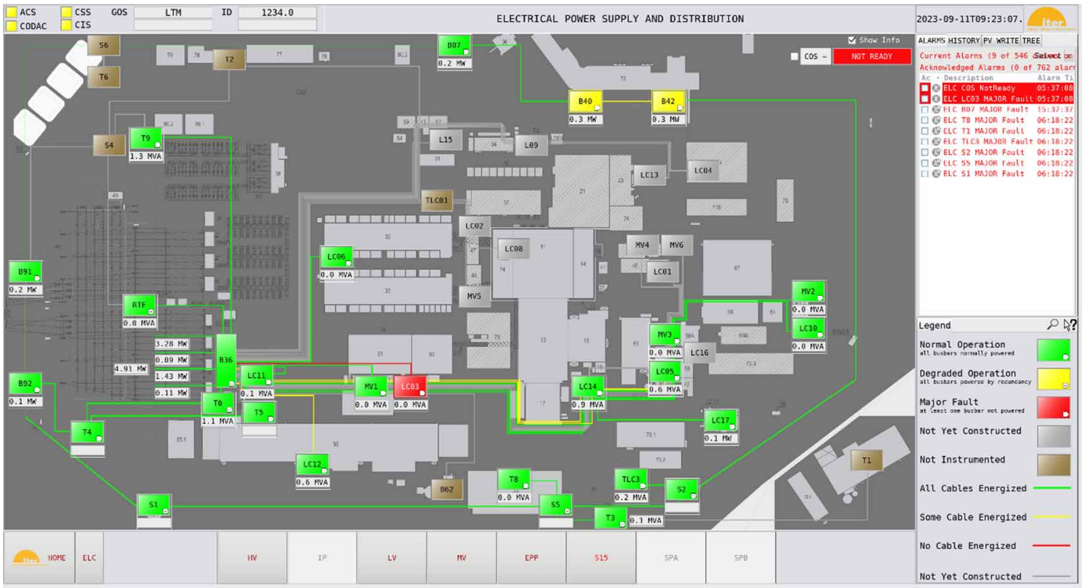

```
Current Alarms (9 of 546
Acknowledged Alarms (0 of 762 alarr
Ac  Description               Alarm Ti
    ELC COS NotReady          05:37:08
    ELC LC03 MAJOR Fault      05:37:08
    ELC B07 MAJOR Fault       15:37:37
    ELC T8 MAJOR Fault        06:18:22
    ELC T1 MAJOR Fault        06:18:22
    ELC TLC3 MAJOR Fault      06:18:22
    ELC S2 MAJOR Fault        06:18:22
    ELC S5 MAJOR Fault        06:18:22
    ELC S1 MAJOR Fault        06:18:22
```

Figure 1: Electrical Distribution top level Human Machine Interface (screenshot).

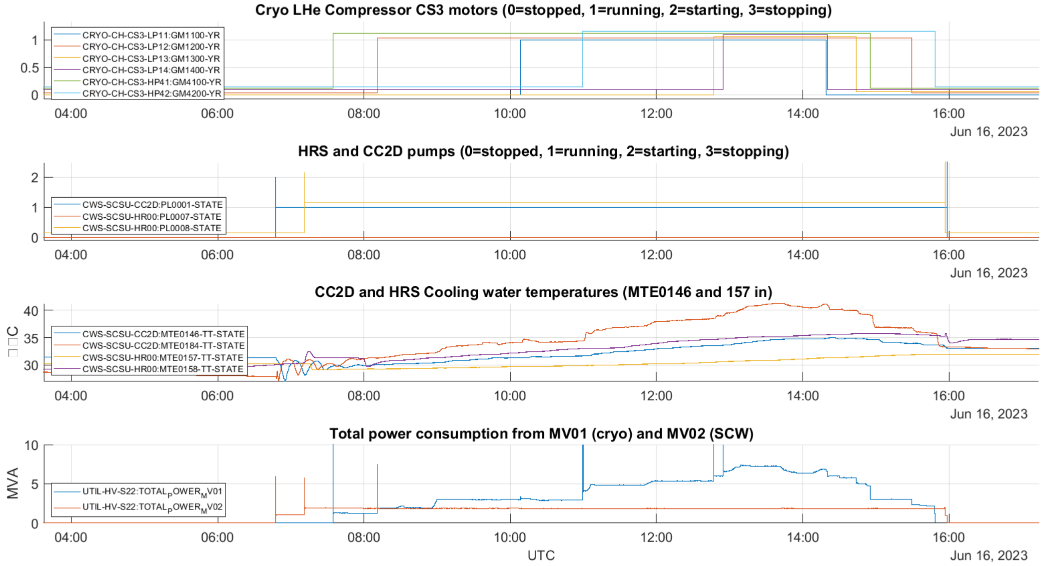

Figure 2: Data from cryoplant commissioning run with six compressors in parallel.

Building services provide: the environmental conditions required by the systems installed in a building, fire protection and local distribution of electricity, liquid and gas. These services are now in operation for half the auxiliary buildings.

The cooling water system started commissioning in 2021 and consists of a heat rejection system with cooling towers, several component cooling loops and chilled water distribution. Large pumps and pipes with flow rates of thousands of kg per second distribute the water all over the site. Cooling water is a pre-requisite for many other systems e.g., magnet power supplies, cryoplant and additional heating systems.

The cryoplant started commissioning in 2022 and is responsible for producing and supplying liquid nitrogen and liquid helium to the superconducting magnets, cryopumps and thermal shield. In June 2023, a major milestone was achieved by running six water cooled liquid helium compressors in parallel, consuming 7.3 MVA. Figure 2 shows the data associated with this achievement, extracted from the data archive. It is a good example of plant system dependencies and the software integration accomplished across multiple systems delivered by many different suppliers.

At time of writing the pulsed power electrical network is being energized, providing 66 kV and 22 kV using three 400 kV transformers connected to the grid. This is the main power source for the coil power supplies as well as additional heating systems and is tightly coupled with the reactive power compensation and harmonic filtering system. The latter improves the power quality and protects the grid by limiting voltage variations on the 400 kV busbars.

All systems described so far are industrial systems based on Siemens programmable logical controllers (PLC) running up to 100 Hz. Recently, the first system making use of ITER standards for fast control applications, reactive power compensation, has been integrated. It is based on industrial computers running RT Linux, input/output over PCI express links, makes use of FPGAs for hard real-time control and protection functions, and interfaces with the high-performance networks for time synchronisation, distributed real-time control and high throughput data acquisition. In the coming year, electron cyclotron heating and magnet power supplies will be integrated, using similar technologies. Integration will be followed by commissioning. This will complete all supporting systems outside the Tokamak complex. It is expected that integration and commissioning inside the Tokamak complex will start in 2025 with vacuum, fueling, magnets, first wall, additional heating and diagnostics for measuring the plasma parameters.

## *Metrics*

A useful metric to measure size and progress is the number of integrated EPICS process variables. This is easy to obtain from the software repository and can be predicted for future systems based on design data and system complexity. To start system commissioning, the electronics racks must be energized and the process variables integrated. The commissioning schedule then provides the timeline. Figure 3 shows the evolution of integrated process variables. We expect close to 4,000,000 process variables for first plasma and we have achieved about 25% of that scope today.

Another representative metric is the number of energized electronics racks. According to engineering data, a total of 1500 is expected at first plasma. Figure 4 shows the evolution of this metric, which follows the same trajectory as seen in Fig. 3.

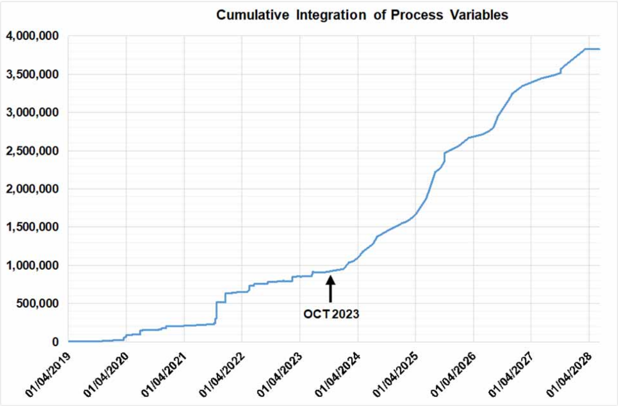

Figure 3: Evolution of integrated process variables.

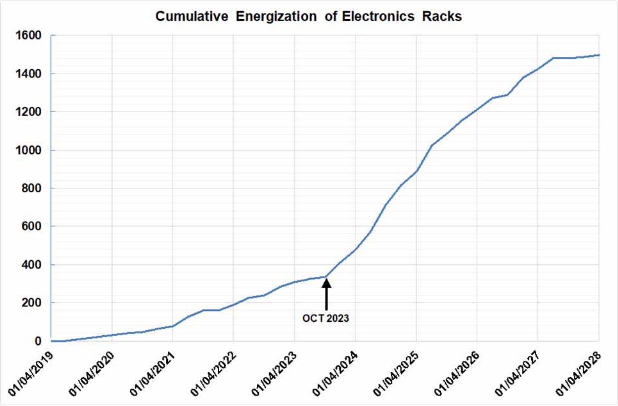

Figure 4: Evolution of electronics racks energization.

Other metrics include archived data and instances of software deployments. The cumulative amount of archived data from the integrated systems, principally comprising slow 24/7 time series, has reached 6 TB. During integration, commissioning and, to a lesser extent, operation, bugs and improvements are identified. Tracked using tickets these result in new software versions being deployed. Typically there are a total of two new deployments every week.

## INTEGRATION PROCESS

Standardization of the integration process, illustrated in Fig. 5, is crucial. The pre-requisites are the installation of the hardware and the delivery of documentation and software source code from the supplier. Hardware installation includes racks and all cabling as well as third party electrical legal inspection to get permission to energize. As a first step, an inventory of the deliverables is carried out. This includes an inspection of the software source code and its packaging, then deployment on a test platform [5]. Subsequently the central services are configured: human machine interface, data archiving and alarm handling. When power is available in the field, the systems electronics racks are energized and the networks configured. The networks between the central infrastructure, the local system and the field are established, the controllers are configured and field signals verified. Finally, functional verification of the local system is performed. At this point the system is declared ready for process commissioning.

The architecture is such that new systems can be introduced in pre-defined places, e.g. HMI navigation, while protecting systems already in operation. Software is delivered to the repository with standardized unit naming. Central services are segregated to minimize the impact on systems in operation while inserting a new system. At the same time the integrated control system provides the means for all the systems to act in unison by using standard protocols, archives, alarms, HMIs, configuration and orchestration.

As an increasing number of systems are being added and operated by separate groups, access control becomes important. Any terminal connected anywhere on the network has read access to all data, but control actions can only be executed by authorized people. This is enforced by using personal accounts and EPICS gateway access lists. Three conditions are required for control access: a user belongs to authorized group: the control action is executed from an authorized terminal and the ITER machine is in a state allowing the action. All control actions are logged, including the executing user, for post-mortem analysis.

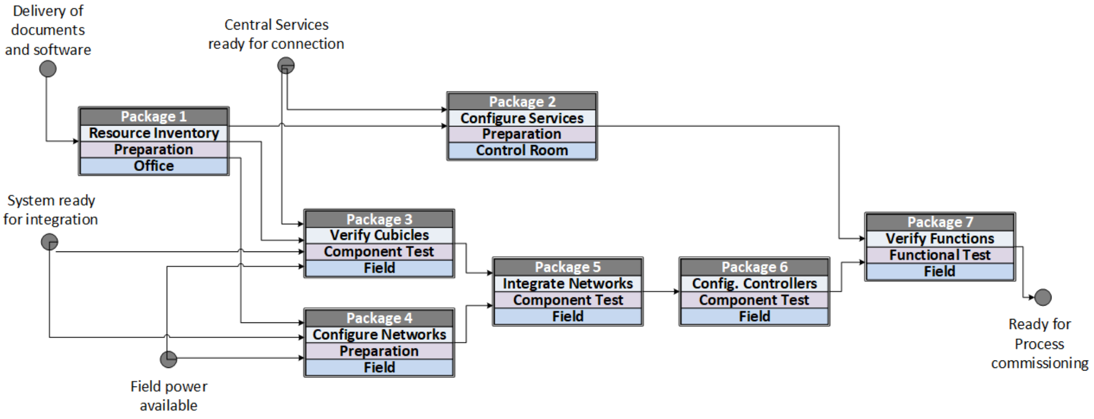

Figure 5: Integration process.

The Control System Studio front-end and its configuration is deployed to the operator terminal. To guarantee that all terminals have the same configuration, an automatic installation on all terminals is executed once per week. This guarantees that all updates over the last week are available on all terminals. A similar approach is used for the EPICS gateway's configuration to ensure all access lists are up to date.

During integration, commissioning and early operation of a first-of-a-kind system, understanding of how it behaves increases with time. This knowledge is captured in documented operational procedures, which gives opportunity for automation. The high-level supervision and automation software provides a generic behavior tree based sequencer [6], which can easily be configured to automate such operational procedures. The first use case of this approach has been the startup and shutdown of the cooling water loops, which involved validation of pre-requisites and sequences of commands to different systems. Early interactions between system operator and software developer provided high quality feedback and enabled incremental improvement of the software.

Another important topic is the management of all configuration parameters. Supplied systems as delivered often require manual setting of each configuration parameter via an HMI. This cannot be sustained in a large integrated system, where the configuration of systems needs to be managed consistently to deliver the required performance. The high-level configuration framework software to address this is in early development. It is planned to have a first use case in the next months. This will be for the reactive power compensation system, which has hundreds of configuration parameters.

The most challenging control problem for ITER is distributed real-time plasma control. A software package called Real-Time Framework (RTF) [7] provides common services and capabilities to build real-time applications. Early stubs of this software are being deployed for testing some fast control system and early versions have also been deployed on KSTAR Tokamak now operating in Korea.

### *Software Configuration Control*

Robust configuration control of all software source code is essential in the production environment. This is always challenging, as agility is required during commissioning. The applied model is illustrated in Fig. 6.

The supplier delivers a tagged software version to the repository and requests deployment using the ticketing system. The administrator extracts the software, performs packaging and deploys the RPM on the control model [5] for quality checks. If successful, the administrator deploys the RPM on the production target, otherwise they inform the supplier using the ticketing system. The tester performs tests on the target and registers any issues and/or change requests, again via the ticketing system. The supplier addresses the issues/change requests and submits a new tagged version to the repository and the workflow is repeated. More rigid quality control is applied to early iterations. A historical inventory of all deployments is maintained. As an example, the reactive power compensation unit has undergone 48 iterations over the last two years.

During time critical commissioning or in the case of some emergency, agility is achieved by assigning the same person to multiple roles and exceptionally allowing hot fixes. Any irregularities thus introduced are rectified afterwards using the same workflow.

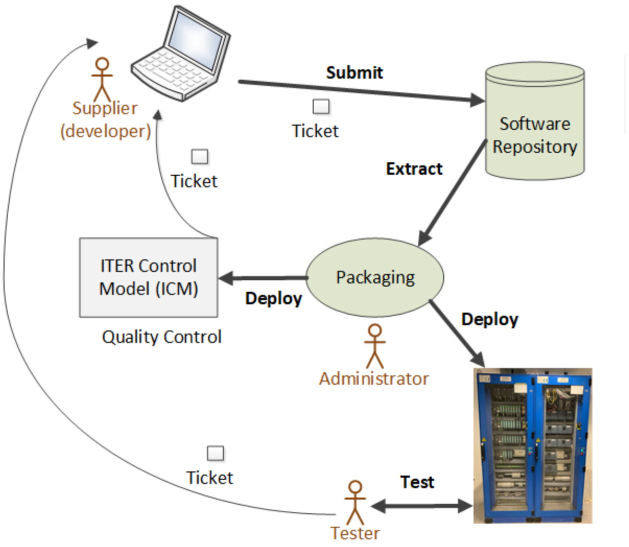

Figure 6: Software configuration control and deployment workflow.

## LESSONS LEARNED

While integrating all these systems, a strong focus has been kept on streamlining and standardizing the integration process as outlined above. The pre-requisites are (1) energization of the local control system racks after passing electrical legal inspection, (2) physical connection to the central infrastructure and (3) delivery of the source code of the local control system to the central software repository. Systems are delivered by many different suppliers and contractors. There is resistance to satisfying these pre-requisites and a preference for working in isolation. In-kind procurement with distributed decision making and financial constrains make integration challenging. We have applied a mix of sticks and carrots to resolve this issue.

Although the standards and guidelines for construction of electronics racks are quite exhaustive, delivered hardware often does not meet the requirements set by French legislation. This has been discovered during the third party electrical legal inspection and has required hardware modification before energization was authorized. This problem has certainly been underestimated and has resulted in delays and extra cost.

With each new supplier there is resistance to delivering their source code to the software repository. Even if such requirements are clearly stated, the procurement model with many intermediate contractors makes it difficult to enforce. Only after many discussions and at an extremely late stage in the process do we receive the source code. We encounter the same problem for documentation. After the correct procedure has been established once subsequent deliveries from the same supplier are smoother.

Delivered software source code is often poorly tested and of inferior quality. Considerable internal resources must be allocated to rectify this. Although some resources have been allocated for this purpose in planning, the actual resource requirement is larger by a factor of about 3.

Delivered software code is often not structured as required. It is necessary to pre-process delivered software before it can be packaged and deployed on the target. This has been automated by scripting and is no longer a significant problem.

The best results are obtained if a technical team with all stakeholders and a common goal can be established early in the integration process. By demonstrating added value, engagement and commitment, trust can be earned. Good examples of added value are the central data archive and its access tools, remote real-time access and notifications, access to other systems data and dealing with deployments of central services. We have seen multiple cases where such technical teams have been established. When this approach fails it is usually for non-technical reasons imposed by politics and/or financial and contractual issues.

In some cases the delivered system has been declared not fit for purpose and the supplier has been unwilling or unable to fix it. In those cases a formal non-compliance process has been followed to transfer the scope, with funding, to the central team and the system has been reworked. Such a process is tedious and involves lengthy discussions and arguments between the central team and ITER members/suppliers. As an example, this has been the case for compressed air and demineralized water production systems.

In some other cases, an early agreement has been reached to transfer the scope, with funding, during the design phase. This has been successful for the cooling water system and initiated for the vacuum master control system.

In general, when done early enough, scope transfer has always been successful and less expensive for the project. Unfortunately, the ITER members and suppliers have not realized this in time.

## NEXT STEPS

In the coming year, integration of coil power supplies, electron cyclotron heating power supplies and gyrotrons will start. In parallel, the high-level software will be further developed and deployed and the final infrastructure, with networks and main control room, will be put into operation.

In the following years, integration will start in the Tokamak complex covering vacuum vessel, cryostat, magnets, first wall, vacuum, fuelling, Tokamak cooling water, cryolines, additional heating transmission lines/launchers and diagnostics.

In parallel, obsolescence management becomes increasingly important due to the long lifetime of the project. The first activities to address difficulties with availability of spare parts have started and detailed obsolescence management plans are being developed.

Expanding the system will lead to new discoveries and feedback from experience will be continuously integrated back into the deployed systems. Incrementally, the integrated control system will become larger, more mature, more reliable and more automated.

## CONCLUSION

The ITER control system is in production with close to one million integrated EPICS process variables. This corresponds to approximately 25% of the scope towards first plasma. More and more high-level software is going into production to facilitate integrated operation.

Integration of many in-kind local systems has been and remains challenging. The most successful mitigations have been establishment of joint technical teams or scope transfer. However, the project organization makes both these mitigation actions difficult to achieve. Nevertheless, we continue to work towards the functional integration of the ITER plant to enable integrated and automated operation.

## ACKNOWLEDGEMENTS

The work presented in this paper is the result of the effort of everyone in the Controls group of the ITER Organization. We acknowledge the contribution of the seven ITER members and all their contractors.

The views and opinions expressed herein do not necessary reflect those of the ITER Organization.

## REFERENCES

[1] Plant Control Design Handbook, ITR-23-009, `https://www.iter.org/technical-reports`

[2] F. Di Maio *et al.*, “CODAC Core System, the ITER software distribution for I&C”, in *Proc. ICALEPCS’11*, Grenoble, France, Oct. 2011, paper MOPKS029, pp. 227-230.

[3] A. Wallander *et al.*, “News from ITER controls – a status report”, in *Proc. ICALEPCS’11*, Grenoble, France, Oct. 2011, paper MOBBAUST01, pp. 1-4.

[4] R. Castro *et al.*, “ITER Unified Data Access (UDA)”, in *11th IAEA Technical Meeting on Control Data Acquisition and Remote Participation for Fusion Research,* Greifswald, Germany, May 2017.

[5] M. Prokopas *et al.*, “ITER Control Systems Model: A Full-Scale Simulation Platform for the CODAC Infrastructure”, in *11th IAEA Technical Meeting on Control, Data Acquisition and Remote Participation for Fusion Research,* Greifswald, Germany, May 2017.

[6] W. Van Herck *et al.*, “Automated operation of ITER using behaviour tree semantics”, in *Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021, pp. 628-630. `doi:10.18429/JACoW-ICALEPCS2021-WEPV006`

[7] W. Lee *et al.*, “Real-time framework for ITER control systems”, in *Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021, pp. 45-49. `doi:10.18429/JACoW-ICALEPCS2021-MOBL02`