ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-MO3A006

ENERGY CONSUMPTION OPTIMISATION USING ADVANCED CONTROL ALGORITHMS

F. Ghawash[1], B. Schofield*, E. Blanco, CERN, Geneva, Switzerland
[1] 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

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

feedback necessary to deal with disturbances and model mismatch. As it is based on solving an optimization problem, the approach is very broad, as many possible optimization problems could be posed for a given control problem. The approach is directly applicable to multivariable systems, with many measurements and actuators, and constraints can be taken into account explicitly.

The key challenges with designing a model predictive controller are the development of the model, and the formulation of the optimization problem. The class of optimization problem that results dictates the type of solver that must be used.

Receding Horizon Control

The core behaviour of MPC, namely iteratively solving a finite horizon optimization problem, is sometimes called Receding Horizon control, as the finite horizon continually advances at each new time step. The principle is illustrated in Fig. 1. Although the core feature of the MPC scheme is

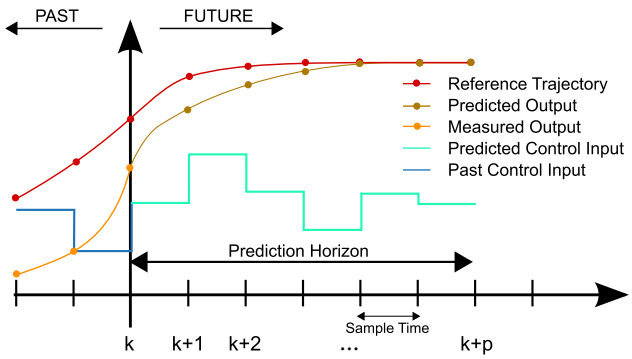

Figure 1: Receding horizon control principle [2]. At each sample time, new measurements are taken and the problem is solved again. Only the first of the computed control outputs are applied to the process.

the use of a dynamic plant model to enable prediction of future output, it is important to note that the scheme is very well suited to situations in which future values of other plant information, such as external inputs or disturbances, are known or can be predicted. In the case of AHUs, one such external input is the outside air temperature, a prediction for which can be reliably obtained from meteorological services.

Open vs Closed Loop

A subtle but important aspect of MPC is that it does not produce a control law, as in a feedback rule to apply to the outputs or states of the process, but rather simply a sequence of control inputs over the control horizon. This is essentially an open loop control signal. Feedback is achieved through the iterative nature of the strategy.

MODELING OF AIR HANDLING UNITS

In this section, simple control orientated models of the different components of an AHU will be derived, primarily based on mass and energy balances. For more details about the model derivation see [3].

Air Handling Unit

Figure 2 illustrates the structure of the air handling unit considered in this article. The unit contains a variable speed fan, heating and cooling coils, and dampers to control the airflow. The aim of the AHU is to maintain the air temperature in the zone, and to provide a certain amount of fresh air from the exterior of the building in order to maintain air quality.

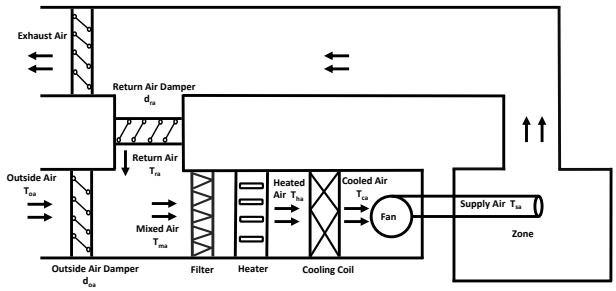

Figure 2: Air Handling Unit.

Zone Model

The zone air temperature mainly depends on the supply temperature from the AHU, outside air temperature and the heat load on the building, which is typically time varying. Assuming the air in the zone is fully mixed and the density of the air remains constant, the zone energy balance leads to (1). Similar models have been reported in [4–6].

$$C_z \frac{dT_{za}}{dt} = \dot{m}_{sa}^{in} C_{pa} T_{sa} - \dot{m}_{sa}^{out} C_{pa} T_{za} + \alpha (T_{oa} - T_{za}) + q(t) \quad (1)$$

Here C_z represents the zone thermal capacitance. T_{za}, T_{sa}, T_{oa} are the zone, supply and outside air temperature respectively. C_{pa} is the specific heat capacity of air, α is the heat transfer coefficient and $q(t)$ represents the time varying heat load. It must be noted that, C_z, α and $q(t)$ must be identified from operational data. The mass balance for the zone gives (2).

$$\dot{m}_{sa}^{in} = \dot{m}_{sa}^{out} \quad (2)$$

where $\dot{m}_{sa}^{in}$ and $\dot{m}_{sa}^{out}$ represents the mass flow rate of supply air entering and leaving the zone. Figure 3 shows the zone model performance compared to operational data.

Dampers and Mixing Chamber Model

In an AHU, dampers play a key role of modulating the flow of outside air and return into the mixing chamber. The mixing of the air flows in correct proportions can be used to regulate the mixed air temperature. Due to the negligible heat interaction with the surrounding, the mixing process is assumed to be adiabatic. The energy balance leads to (3).

$$C_m \frac{dT_{ma}}{dt} = \dot{m}_{oa} C_{pa} T_{oa} + \dot{m}_{ra} C_{pa} T_{ra} - \dot{m}_{sa} C_{pa} T_{ma} \quad (3)$$

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

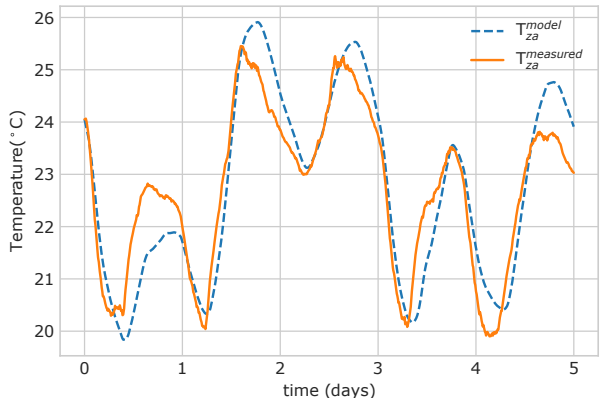

Figure 3: Zone model validation.

where C_m is the mixing chamber thermal capacitance. Under the steady state conditions (3) simplifies to:

$$T_{ma} = \frac{\dot{m}_{oa}T_{oa} + \dot{m}_{ra}T_{ra}}{\dot{m}_{sa}} \tag{4}$$

where T_{ma} and T_{ra} are mixed and return air temperatures. $\dot{m}_{oa}$ and $\dot{m}_{ra}$ represent the mass flow rate of outside and return air respectively. The mass balance leads to (5).

$$\dot{m}_{sa} = \dot{m}_{oa} + \dot{m}_{ra} \tag{5}$$

The outside and return air dampers are usually designed in a matched configuration and are operated in unison. For instance, when the outside air damper closes, the return air damper opens (and vice-versa). The dampers are sized such that, the decrease in the mass flow rate of one is equally matched by the increase in the flow of the other. Hence, the total mass flow rate of supply air roughly remains constant irrespective of the mixing ratio. Under such an operational configuration, the mass flow rate of supply air becomes independent of the dampers' position and depends only on the fan speed. The relation can be approximated by the linear function $\dot{m}_{sa} = a_0\omega + a_1$. Under the matched dampers configuration, (4) can be rewritten as follows:

$$T_{ma} = \frac{\phi(d_{oa})\dot{m}_{sa}T_{oa} + (1-\phi(d_{oa}))\dot{m}_{sa}T_{ra}}{\dot{m}_{sa}} \tag{6}$$

where $\phi(d_{oa})$ represents the installed flow characteristics and d_{oa} is the opening of the OA dampers. Note that, ϕ and $1-\phi$ are the fraction of the mass flow rate of supply air contributed by the outside and return air dampers respectively. In the case of linear flow characteristics (installed) of the dampers [4–6], the fraction of the mass flow rate of air contributed by the damper is proportional to the damper's opening (i.e $\phi = d_{oa}$). In practice, the dampers exhibit non-linear flow characteristics. Depending on the type of the dampers (parallel or opposed blades), the inherent flow characteristic (ϕ_{inhe}) of the dampers can be approximated using a third order polynomial.

$$\phi_{inhe}(d_{oa}) = b_0 d_{oa}^3 + b_1 d_{oa}^2 + b_2 d_{oa} \tag{7}$$

It is important to note that, the inherent non-linear flow characteristics are measured under the laboratory conditions (i.e with a constant differential pressure across the dampers). However, when the dampers are installed in the system, the installed flow characteristics can significantly deviate from the inherent flow characteristics depending on the authority[1] of the dampers. The installed flow characteristics are given by the following equation [7]:

$$\phi = \phi_{inhe}\sqrt{\frac{1}{(1-d^{auth}) \times \phi_{inhe}^2 + d^{auth}}} \tag{8}$$

where d^{auth} denotes the authority of the damper, which must be estimated from operational data. The performance of the mixing chamber model against operational data is shown in Fig. 4.

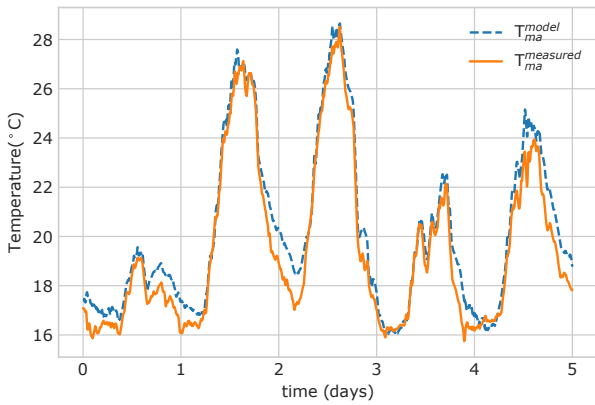

Figure 4: Mixing chamber model validation.

Cooling Coil Model

When a supply temperature lower than the mixed air temperature is required, the cooling coil valve is adjusted to provide adequate cooling of the mixed air. Under suitable assumptions, the mass and energy balances are used to derive a semi-empirical model of the cooling coil to effectively represent the dynamics of off coil temperature of the air [8].

$$\frac{dT_{ca}}{dt} = -c_1\dot{m}_{sa}(T_{ca} - T_{ha}) - \frac{c_2\dot{m}_{sa}^l}{1 + (\frac{\dot{m}_{sa}}{\dot{m}_w})^l}(T_{ca} - T_w) \tag{9}$$

T_{ca} and T_w are the cooled air and chilled water temperature respectively. c_1, c_2 and l are the empirical parameters of the model. Although, the optimization formulation ignores the dynamics of the cooling coil, the model can be effectively utilized in closed loop simulations. The power consumption associated with the generation of chilled water consumed by the cooling coils is given by the following equation [4, 5].

$$P_c = \frac{\dot{m}_{sa}C_{pa}(T_{ha} - T_{ca})}{\eta_c COP_c} \tag{10}$$

where η_c is the efficiency of the cooling coil and COP_c represents the coefficient of performance of the cooling plant.

[1] The authority of the damper is defined as the pressure drop across the fully open damper divided by the total pressure drop in the system.

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

Heater Model

Depending on the zone heating requirement, a supply air temperature higher than the mixed air temperature might be required. This can be achieved by heating the mixed using an electric heater, or heating coil. In this case study, a heating coil is used. The modeling follows that outlined above for the cooling coil. Again, for the MPC formulation, the heating coil dynamics are not required for the model; only the energy consumption is considered. This is given by:

$$P_h = \frac{\dot{m}_{sa} C_{pa} (T_{ha} - T_{ma})}{\eta_h COP_h} \tag{11}$$

where η_h is the efficiency of the heating coil and COP_h represents the coefficient of performance of the heating plant.

Fan Heat Loss Model

When the fan is placed in the path of the airflow, a fraction of the heat generated due to motor losses is absorbed by the air passing through the motor. The increase in the temperature of the air can be modeled as follows:

$$T_{sa} = T_{ca} + \frac{\beta P_f(\omega)}{\dot{m}_{sa} C_{pa}} \tag{12}$$

where T_{sa} is the supply air temperature, β represents the fraction of heat absorbed by the air. A fairly accurate approximation of fan power consumption can be obtained using a quadratic polynomial:

$$P_f(\omega) = g_0 \omega^2 + g_1 \omega + g_2 \tag{13}$$

FORMULATION OF THE MODEL PREDICTIVE CONTROLLER

In this section, a non-linear optimization program is formulated which achieves the desired control objectives while minimizing the total energy consumption of the actuators.

Objectives and Constraints

The main objective is to regulate the zone temperature within the desired range while minimizing the energy usage of the actuators. To this end the objective function is constructed by summing the power consumption values of the actuators based on the models presented in the previous section. In order to make use of 'free' heating or cooling when available depending on outside temperatures, a quadratic cost on the deviation from an ideal zone setpoint can be added. This term can be given a low weight such that it is dominated by the power consumption terms.

For the constraints, typically only actuator ranges are given as 'hard' constraints. For process variables which should be controlled in a range, so-called 'slack' variables are used, which provide a margin of error for the process variable to violate the constraint. These slack variables are added to the cost function with corresponding large weights. This approach is taken to avoid infeasibility of the optimization problem when process variables exceed the desired constraints, for example at plant startup.

The cost function of the optimization problem is shown in (14*a*), while the constraints are shown in (14*b*) to (14*p*)

$$\min_{\mathbf{d},\boldsymbol{\omega},\boldsymbol{\epsilon},\mathbf{T}} \sum_{k=0}^{N} P_f(\omega_k) + P_h(T^{ha}_{k+1}) + P_c(T^{ca}_{k+1}) + r_{sp}(T^{za}_k - T^{za}_{sp})^2 + \mathbf{r}^T \boldsymbol{\epsilon}^2 \tag{14a}$$

$$\text{s.t} \quad T^{za}_{k+1} = f_{za}(T^{za}_{,k}, T^{sa}_k, T^{oa}_k, \omega_k), \forall k \in \mathcal{H}_N \tag{14b}$$

$$T^{ma}_{k+1} = f_{ma}(T^{za}_k, T^{oa}_k, d^{oa}_k), \forall k \in \mathcal{H}_N \tag{14c}$$

$$T^{ca}_{k+1} - \epsilon_{ha} \leq T^{ha}_{k+1} \leq T^{max}_{ha} + \epsilon_{ha}, \ \forall k \in \mathcal{H}_N \tag{14d}$$

$$T^{min}_{ca} - \epsilon_{ca} \leq T^{ca}_{k+1} \leq T^{ha}_{k+1} + \epsilon_{ca}, \ \forall k \in \mathcal{H}_N \tag{14e}$$

$$T^{sa}_{k+1} = f_{fan}(T^{ca}_k), \ \forall k \in \mathcal{H}_N \tag{14f}$$

$$T^{min}_{sa} - \epsilon_{sa} \leq T^{sa}_{k+1} \leq T^{max}_{sa} + \epsilon_{sa}, \ \forall k \in \mathcal{H}_N \tag{14g}$$

$$T^{min}_{za} - \epsilon_{za} \leq T^{za}_{k+1} \leq T^{max}_{za} + \epsilon_{za}, \forall k \in \mathcal{H}_N \tag{14h}$$

$$d^{min}_{oa} \leq d^{oa}_k \leq d^{max}_{oa}, \ \forall k \in \mathcal{H}_N \tag{14i}$$

$$\omega^{min} \leq \omega_k \leq \omega^{max}, \ \forall k \in \mathcal{H}_N \tag{14j}$$

$$-\dot{d}_{oa} \leq d^{oa}_{k+1} - d^{oa}_k \leq \dot{d}_{oa}, \ \forall k \in \mathcal{H}_N \tag{14k}$$

$$-\dot{\omega} \leq \omega_{k+1} - \omega_k \leq \dot{\omega}, \ \forall k \in \mathcal{H}_N \tag{14l}$$

$$-\dot{T}^{ca} \leq T^{ca}_{k+1} - T^{ca}_k \leq \dot{T}^{ca}, \ \forall k \in \mathcal{H}_N \tag{14m}$$

$$-\dot{T}^{ha} \leq T^{ha}_{k+1} - T^{ha}_k \leq \dot{T}^{ha}, \ \forall k \in \mathcal{H}_N \tag{14n}$$

$$T_{za}[0] = T^{init}_{za}, \ T^{ma}[0] = T^{init}_{ma}, T_{ha}[0] = T^{init}_{ha}, \tag{14o}$$

$$T_{ca}[0] = T^{init}_{ca}, \ \epsilon_{sa} > 0, \epsilon_{za} > 0, \ \mathbf{r} > 0 \tag{14p}$$

where $\boldsymbol{\epsilon} = (\epsilon_{sa}, \epsilon_{za}, \epsilon_{ca}, \epsilon_{ha})$ are the slack variables for the process value constraints, and $\mathcal{H}_N = \{0, \cdots, N\}$ is the prediction horizon. Here f_{za}, f_{ma} and f_{fan} represent the zone, mixing chamber, and fan models respectively. The minimization is performed over the vectors of control variables and system states, namely $\mathbf{d} = \{d^{oa}[0], \cdots, d^{oa}[N]\}$, $\boldsymbol{\omega} = \{\omega[0], \cdots, \omega[N]\}$, $\mathbf{T} = \{T_{za}[1], \cdots, T_{za}[N+1], T^{ma}[1], \cdots, T^{ma}[N+1], T_{ha}[1], \cdots, T_{ha}[N+1], T_{ca}[1], \cdots, T_{ca}[N+1], T_{sa}[1], \cdots, T_{sa}[N+1]\}$. The equality constraints (14*b*), (14*c*), and (14*f*) are used to ensure that the trajectories of the state variables must evolve as per the defined models. (14*g*), (14*h*), (14*i*), (14*j*) represents the maximum and minimum allowable values for the supply air temperature, zone temperature, outside air damper position, and fan speed.

(14*k*), (14*l*), (14*m*), (14*n*) are used to impose the rate of change constraints for the actuators. (14*o*), (14*p*) represent the initial conditions of different variables. It must be noted that T_{ca} is the set point for the cooled air temperature. The set point can be achieved by a local controller by adjusting the opening of the cooling valve. The constraint of the maximum cooling capacity of the cooling coil is imposed as a lower bound in (14*e*). Moreover, the upper bound in (14*e*) ensures consistency in the cases when no cooling is required. A similar constraint is imposed for the heating coil in (14*d*). Such a formulation can intrinsically take into account different operational modes of the AHU. For example, when heating is required, the optimization program will choose the heated air temperature (T_{ha}) to be higher than the mixed air temperature, while setting the cooled air temperature (T_{ca}) equal to the heated air temperature. In case cooling is required, the heated air temperature is set to mixed air temperature (T^{ma}) while the cooled air temperature (T_{ca})

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

set point is chosen. In free cooling mode, the heater and cooling valve remain off and the optimization program will set the heated as well as the cooled air temperature to mixed air temperature. Hence, the desired supply temperature is maintained by modulating the damper positions.

IMPLEMENTATION

The optimization problem presented in (14a) to (14p) constitutes a Nonlinear Programming (NLP) problem. Implementing a solver for this class of problem on an standard PLC is not realistic, both in terms of performance of the hardware and the development time required. There exist controllers which allow the use of higher level languages and external libraries, but the use of such controllers was not an option in this pilot project. A review of industrial solutions for the implementation of advanced controls can be found in [9].

Solving the MPC Problem

To solve the optimization problem, it was decided to use CasADi [10] (via its Python interface) with IPOPT [11] as the solver. Python was chosen as it allows straightforward implementation of both optimization problem formulation and the related signal processing and data acquisition. An additional advantage of Python is its portability. For the purposes of the pilot implementation described here, the choice of location for running the MPC algorithm was driven primarily by the focus on assessing the stability and robustness of the control algorithm, rather than the controller's reliability. Initially it will run on the SCADA server, but can be moved to the control layer when entering wider scale operation.

Interfacing With Industrial Controllers

The existing control system is implemented using a pair of Siemens S7-315 PLCs. It was decided to keep the existing control logic, and extend it with the possibility of switching to MPC control. Additional PID controllers were added for the heating and cooling valve controls, and tracking calculations implemented to allow seamless transition between the existing control strategy and MPC. In order to interface the PLC with the MPC application, OPC UA was used. The MPC application includes an OPC UA client which reads process values from the PLC, and writes back actuator commands and local controller setpoints for the fan, dampers and valves respectively. The PLC monitors the communication via a watchdog and automatically switches over to the legacy control scheme if communication is lost. The architecture of the controls communication is shown in Fig. 5.

Supervision

For the pilot project, it was decided not to expose any MPC-specific tuning parameters in the supervision system. Instead, the MPC is parameterized with the same values as used by the existing control system, including minimum, maximum and ideal zone setpoints, supply air setpoint limits,

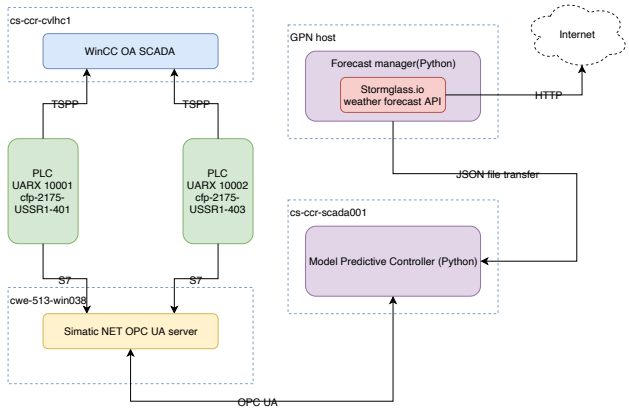

Figure 5: Architecture for the pilot deployment. All communication between the PLC and the Python application is OPC UA.

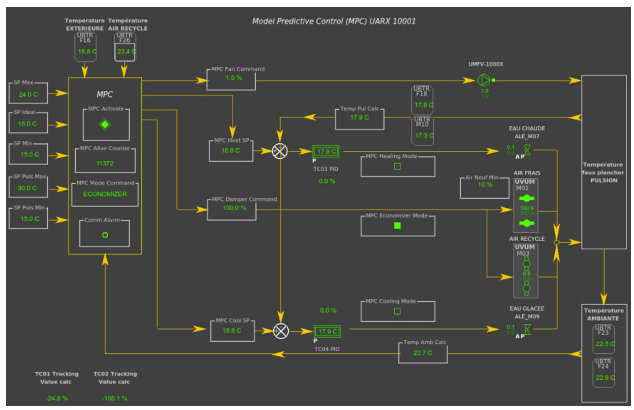

Figure 6: MPC synoptic in WinCC OA.

and actuator limits. A new synoptic was developed, shown in Fig. 6, to provide operators with an overview of the signals provided to and obtained from the MPC application.

VIRTUAL COMMISSIONING USING A PROCESS SIMULATOR

In order to reliably develop and test the MPC controller, it was decided to create a process simulator using the multi-domain modeling and simulation tool EcosimPro. The tool allows first-principles physical models to be created and simulated. The model for the AHU in this case study is shown in Fig. 7. Model parameters were fitted using historical plant data, and the completed model validated against separate sets of historical data. The development of the model is covered in [12].

Models created in EcosimPro can be compiled into executables providing a number of possible interfaces, including Python and OPC UA. For initial validation, the MPC controller was connected directly to the model using the Python interface. Later, in order to test the integration with the PLC application, a more complete testbed was set up using PLCs. In this case, a simulation driver application written in Python was created to interface the EcosimPro model with the PLCs

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

via OPC UA. The MPC application then interfaces to the PLCs as in the production architecture. A schematic of this setup is shown in Fig. 8.

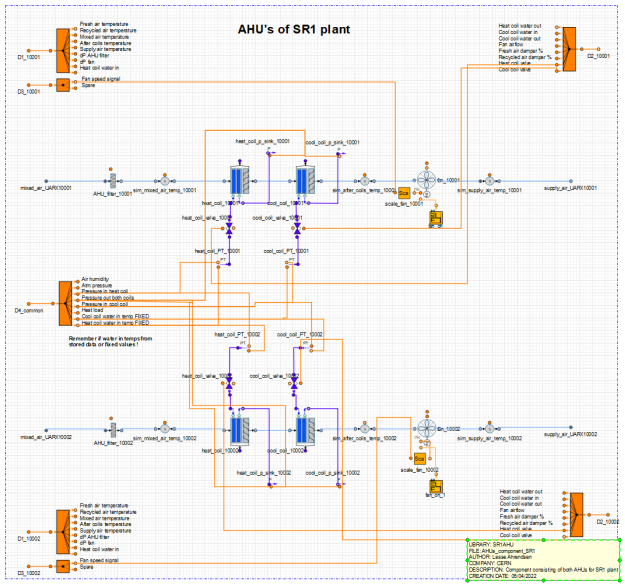

Figure 7: AHU model in EcosimPro.

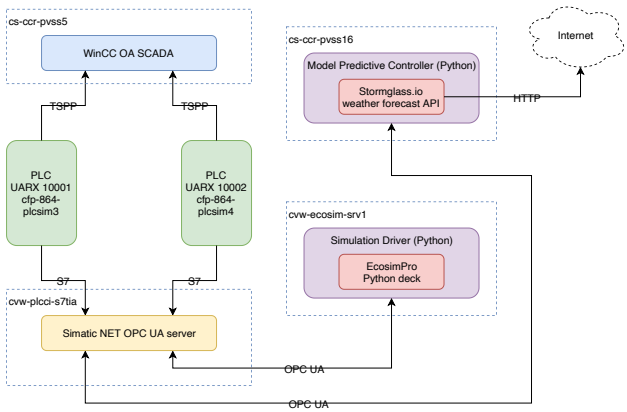

Figure 8: Architecture of the virtual commissioning system. The simulation driver and the MPC both communicate with the PLC via OPC UA. The simulation can be run either in historic mode, where past weather data used, or in 'live' mode based on current conditions.

RESULTS

Comparative Simulation Studies

Before progressing to deployment on the real plant, it was necessary to demonstrate the likely effectiveness of the MPC controller in terms of energy efficiency. To do this, the simulation model was fed with historic data (primarily outside air temperature) for a selection of 30 day periods, chosen to cover a wide range of operational conditions. The MPC controller was applied to the simulation, and the energy consumption calculated. This was then repeated with the legacy (cascaded PID) controls. The reason for simulating the legacy controls rather than using the historic data directly was to provide a more accurate comparison; in this way any modelling errors are common to both controllers. The results of this study are shown in Table 1.

The results indicate that the MPC controller should be able to achieve energy savings of around 20 % compared to the existing controller. In fact, it is likely that further savings could be made by adjusting operational limits. In the study, the upper and lower bounds of the zone and supply air temperatures were the same in the MPC and legacy cases. However, as MPC allows much tighter control around limits, these limits could be relaxed in order to gain additional energy savings.

Table 1: Energy Consumption (kWh) Comparison Between Simulated MPC Control and Actual Controls for a Number of 30-Day Periods in 2021-2022

Region	Energy Consumption		
	MPC	Actual	Savings
Summer	3178	4008	20.7 %
Autumn	1185	1557	23.9 %
Spring	942	1064	11.4 %

Operational Results

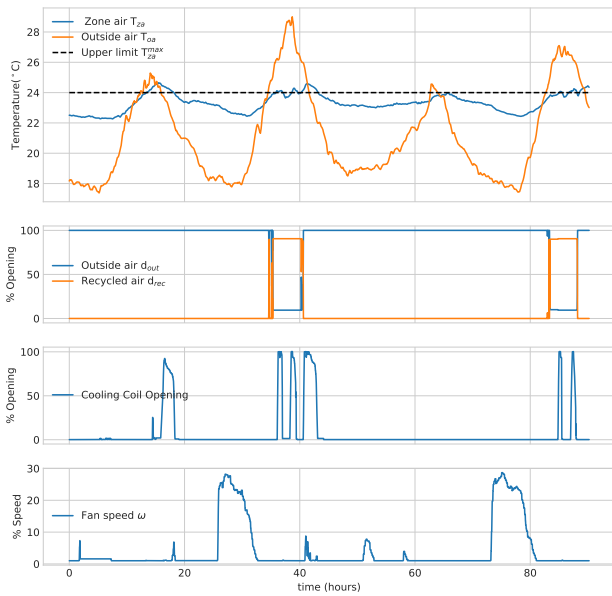

Figure 9: MPC in operation, September 2023. Of particular interest is the use of the fan during night time to reduce the zone temperature prior to warm days. This is made possible by the use of weather forecast data in the optimization problem.

After satisfactorily showing that MPC should be capable of achieving energy savings as well as good control performance in simulation, the deployment of the system on the process was approved. The building on which the MPC would be deployed has two AHUs. It was decided to deploy

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

MPC on one of them, and to leave the existing controls on the other AHU. In operation, the MPC behaves as expected, allowing the zone temperature to change freely within the accepted range, and then controlling quite aggressively when it reaches the limits. Operational data from a number of days in September 2023 is shown in Fig. 9. A key aspect to note is the use of the fan during nighttime prior to warm days; this is enabled by the use of weather forecast data. Making a direct comparison of the energy performance of the MPC scheme compared to the previous controls is difficult. One possible comparison is to calculate the cooling power supplied by the AHU, and compare with the power consumed. This are plotted in Fig. 10, where it can be seen that after the deployment of the MPC (on September 13th), the power consumed by the AHU decreases considerably. A portion of this decrease in energy consumption can be attributed to changing weather conditions, with cooler outside air temperatures after the deployment. To make a further comparison, it is possible to make the same calculation for the other AHU operating in the same building. This is shown in Fig. 11. Here is can be seen that both the supplied cooling power and consumed power are largely constant. Although there are some mitigating factors in this comparison, it is clear that MPC can enable very efficient operation of the AHUs.

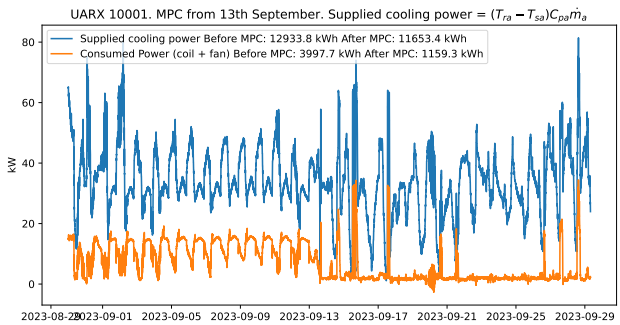

Figure 10: Calculated supplied cooling power and consumed power for AHU before and after MPC deployment.

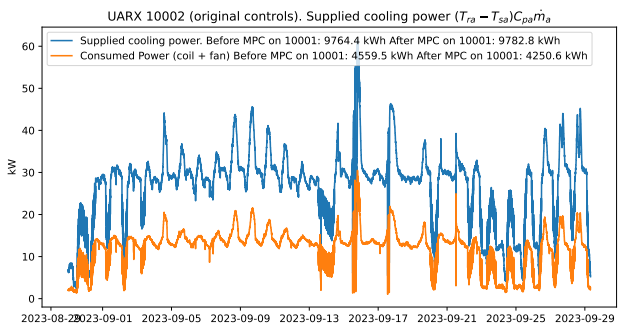

Figure 11: Calculated supplied cooling power and consumed power for second AHU with original controls.

CONCLUSION

Although a longer period of data collection from the deployed MPC controller will be required in order to obtain statistically reliable results on the achieved energy savings, the initial results, together with the possibility of changing operation settings thanks to MPC's tighter control, indicate that energy savings are indeed to be expected. In this context it should be noted that the existing controller is in fact a very recent design, and is in turn considerably more efficient than older control strategies used in a large number of other AHUs at CERN. For this reason it is likely that MPC could enable very considerable energy savings if applied to large numbers of these installations. Although the focus of this article has been the use of MPC for reduction of energy consumption, it is also the case that MPC can provide superior controller performance. The existing control system of the plant presented here exhibits some dynamic issues, linked to the fact that the setpoint of the outer controller is linked to the output of the inner controller in the cascade, with the intention of reducing energy consumption. This leads to oscillation under certain conditions. The MPC scheme presented here does not exhibit such behaviour. The time spent developing the detailed process model was clearly well invested, as the MPC implementation required essentially no commissioning time when deployed in production. Thorough testing in the virtual commissioning phase (made possible by the model), and the perfomance of the model itself are to thank for this.

ACKNOWLEDGEMENTS

The authors would like to thank the Cooling and Ventilation group (EN-CV) at CERN for their close cooperation in this project. Special thanks go to Diogo Monteiro. Lasse Ahrendsen's work on development of the EcosimPro was invaluable to the testing and virtual commissioning phase. The development of the MPC was carried out in a joint Ph.D. project between CERN and the Norwegian University of Science and Technology (NTNU), supervised by Professor Morten Hovd.

REFERENCES

[1] L. Pérez-Lombard, J. Ortiz, and C. Pout, "A review on buildings energy consumption information", *Energy and buildings*, vol. 40, no. 3, pp. 394–398, 2008.

[2] A basic working principle of model predictive control, `https://commons.wikimedia.org/wiki/File:MPC_scheme_basic.svg`

[3] F. Ghawash, M. Hovd, B. Schofield, and D. Monteiro, "Model predictive control of air handling unit for a single zone setup", in *2022 IEEE International Symposium on Advanced Control of Industrial Processes (AdCONIP)*, 2022, pp. 158–163. `doi:10.1109/AdCONIP55568.2022.9894128`

[4] Y. Ma, A. Kelman, A. Daly, and F. Borrelli, "Predictive control for energy efficient buildings with thermal storage: Modeling, stimulation, and experiments", *IEEE control systems magazine*, vol. 32, no. 1, pp. 44–64, 2012.

[5] A. Kelman and F. Borrelli, "Bilinear model predictive control of a hvac system using sequential quadratic programming", *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 9869–9874, 2011.

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-M03A006

[6] G. Lymperopoulos, P. M. Papadopoulos, P. Ioannou, and M. M. Polycarpou, "Distributed adaptive control of air handling units for interconnected building zones", in *2020 American Control Conference (ACC)*, IEEE, 2020, pp. 4207–4212.

[7] C. Federspiel, "Damper authority estimation and adaptive flow control", in *Proceedings of Clima 2000 Conference, held Brussels, August 30th to September 2nd*, 1997.

[8] G.-Y. Jin, P.-Y. Tan, X.-D. Ding, and T.-M. Koh, "Cooling coil unit dynamic control of in hvac system", in *2011 6th IEEE Conference on Industrial Electronics and Applications*, IEEE, 2011, pp. 942–947.

[9] A. Patil *et al.*, "Leveraging local intelligence to cern industrial control systems through edge technologies", in *19th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALPECS)*, 2023.

[10] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control", *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019. doi:10.1007/s12532-018-0139-4

[11] A. Wächter and L. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming", *Mathematical programming*, vol. 106, pp. 25–57, 2006. doi:10.1007/s10107-004-0559-y

[12] L. Ahrendsen, "HVAC system modeling using EcosimPro: A digital twin for advanced control studies and virtual commissioning simulations.", Presented 15 Jun 2022, 2022. https://cds.cern.ch/record/2866476