ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

ENERGY-OPTIMAL PREDICTIVE CONTROLLER FOR CHILLER-BASED COOLING PLANTS OF ACCELERATOR FACILITIES

D. Monteiro*, N. Bunijevac, CERN, Geneva, Switzerland
A. La Bella, L. B. de Giuli, L. Nigro, Politecnico di Milano, Milan, Italy

Abstract

Cooling plants are energy-intensive systems which enable ideal thermodynamical conditions for the smooth operation of accelerator facilities. Among this class of plants, chilled water production systems, used to bring cooling water from high to low temperatures, are commonly found. The automatic control system of such plant is usually complex due to the large number of correlated control inputs available, rendering it particularly challenging to minimize its energy consumption when solely relying on conventional control methods (e.g., PID and/or if-based logic). In 2024, CERN's Engineering department partnered with the university Politecnico di Milano to develop an energy optimal model-predictive controller (MPC) for one of the critical chilled water production plants used for the cooling of CERN's flagship accelerator, the Large Hadron Collider. This 18-month collaboration is well underway, and this paper explores the motivational and organizational aspects of the project, as well as highlights the technical solution proposed, the challenges faced to date, and how these were overcome. Simulation-based results are presented for a detailed performance comparison between MPC and the currently used rule-based logic controller. Finally, the architecture of controls and operator interface for the MPC deployment in the real plant are discussed, in view of extending this optimal control solution to numerous similar systems at CERN.

INTRODUCTION

In the follow-up of previous successful initiatives for energy savings by controls optimization [1], the Engineering department of the European Organization for Nuclear Research (CERN) has decided to pursue further this endeavor. This time, the team decided to focus on chiller-based cooling plants, which are widely present in the laboratory, and are accountable for a significant part of the electricity consumption of the cooling and ventilation (CV) equipment of the organization. As a reference, CV plants account for up to 12% of the consumption of the flagship accelerator of the organization, the Large Hadron Collider (LHC).

Chiller-based cooling plants are particularly energy intensive plants as they incorporate high-power equipment, such as chillers and pumps, and are usually subject to high cooling loads. As an example, the overall installed cooling capacity of chiller-based cooling plants at CERN is estimated at 85 MW[1]. These plants, which are widely used in accelerator and industry facilities, are the default mechanism to produce chilled water to cool down equipment and enable air-conditioning in technical facilities. Not less important, from the industrial control's perspective, they compose a very interesting problem as the available control variables are highly correlated, and hence challenging to manage in parallel. As detailed later in the paper, these plants have historically been controlled by a rule-based strategy combined with Proportional-Integrative-Derivative (PID) controllers. This strategy does not incorporate any kind of formal optimization, being for energy savings or any other type of performance-based metric.

This paper focuses on the work developed to date on the development of a new energy-optimal control solution for chiller-based plants at CERN. In the remaining of the paper, we delve into the foundation and organization of this research collaboration, describe the basic understanding of a generic chiller-based cooling plant, and provide a high-level description of both the legacy control strategy and the new energy-optimal control. In addition, the methodology and key design considerations for the development of the energy-optimal controller are highlighted. Finally, simulation-based comparative results on energy performance are presented, as well as some concluding remarks on the definitive deployment of the new solution in the case-study plant and plants alike.

A SUCCESSFUL RESEARCH COLLABORATION

At the core of the project discussed in this paper, there is a research collaboration gathering several scientific partners: CERN's Engineergin deparment, CERN's Knowledge Transfer group, and the Systems & Control group of Politecnico di Milano (PoliMi). As introduced, the objective set for the project was the development of an energetically optimized control system using a chilled water plant at CERN as case-study. The challenge posed by the topic was the key motivation to bring together and leverage the expertise of each participating partner: CERN contributing primarily with its experience on innovation and research project management, controls optimization for energy savings, and year's of field-experience operating chiller-based cooling plants; PoliMi with a key contribution to the solution development given its long track record of research in advanced control methods for energy systems, including techniques as Model Predictive Control (MPC), and its application to real-world problems.

The project, which is planned to last 18 months (July 2024 to December 2025), was organized by a well-defined set of sequential tasks and deliverables, each with a pre-

* diogo.monteiro@cern.ch

[1] To put in perspective, the usual capacity of a nuclear reactor is 1000 MW [2].

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

defined responsible party and deadlines, as summarized in Table 1. The project's technical team has been holding a remote meeting every 2 weeks for an optimal follow-up of the progress.

So far the project is well on-track and the team is about to complete the integration of the new solution with the control system of the real plant (task 6). For a couple of months the solution will be running in shadow-mode, i.e. the control algorithm is feed by real field data, but the control action proposed is not enforced. This conservative deployment strategy allows to safely collect key feedback on the solution, perform any fine-tuning required, and support the training of the plant operators. At the time of writing, the team has collected promising simulation-based results, which are discussed later in the paper.

The on-time and fruitful collaboration environment in the team have been the pillars of the project's success to date. The joint research is expected to culminate with the definitive commissioning of the new control algorithm in the real plant by the end of 2025.

CHILLER-BASED COOLING PLANTS

Chiller-based cooling plants have the mission of stepping down the temperature of a fluid, most commonly water, which is used to cool equipment or spaces within accelerator facilities. The key equipment class in these plants are chillers, an equipment capable of transferring heat between two mediums through a compression/expansion refrigeration cycle. As depicted in Fig. 1, chillers enable the heat extraction from cooling water, in the evaporator coil, by transferring it to a medium in the condenser coil, typically water or air. A particular interesting feature of chillers is that the mediums on each side may be at different temperature levels (e.g. 6 °C to 12 °C in evaporator coil and 28 °C to 36 °C in condenser coil).

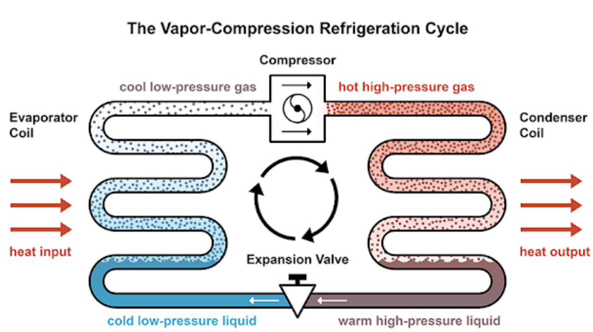

Figure 1: Refrigeration cycle in chillers [3].

Figure 2 depicts the usual scheme of a chiller-based cooling plant. The plant is composed by three main parts: *production*, where chilled water is generated by the one or more chillers, *distribution*, equipped with pump groups to supply chilled-water to the one or more loads[2], and, finally, the *storage* with the so-called buffer tank serving as interface between the two other parts. The *distribution* side foresees two types of loads: the first type (Load_i^1) includes a 3-way motorized valve which enables mixing between the chilled water supply from the *production* side and the load's return water, while the second (Load_i^2) is directly fed by water at the outlet temperature of the buffer tank ($\tilde{T}_\text{L}^\text{in}$).

CERN operates about 40 chiller-based cooling plants for the cooling of its heat-intensive and geographically sparse experimental infrastructure. These plants serve heat loads in the range of 100 kW to 7000 kW, with an overall installed cooling capacity of 85 MW. Given the criticality of these plants, their control system logic has to be reliable. A particular challenge relates to the time-variant nature of heat loads, what requires a constant evaluation and adaptation of the working configuration of the plant. In the following section, we delve into the control system design, drawing an evolution from the rule-based to the newly proposed energy-optimal strategy.

FROM RULE-BASED TO ENERGY-OPTIMAL CONTROL

From the load's perspective, the control system of a chiller-based cooling plant is expected to operate such that the inlet temperature of the load ($\tilde{T}_{\text{L},i}^\text{in}$) is within a required range. This is the baseline requirement of any client with cooling needs. In addition, the system must also ensure the capacity to provide a certain mass flow of chilled-water ($w_{\text{L},i}^\text{in}$) which set following the system design requirements. The systems considered in the scope of this project are equipped with constant speed pumps, hence pump speed control is disregarded in this research. The controls' logic is then left with three types of controlled variables, which are denoted in black in Fig. 2 and detailed as follows:

- $\delta_{\text{CH,i}}$: ON/OFF status of the *i*-th chiller [boolean]
- $T_{\text{CH},i}^\text{out}$: Output temperature of the *i*-th chiller [°C]
- $\tilde{T}_{\text{L},i}^\text{in}$: Inlet temperature of the *i*-th type-2 load [°C]

It should be noted that the outlet temperature of each chiller is regulated by its own independent controller, while the inlet temperature of the type-2 loads is maintained by a dedicated PID loop acting on the associated three-way valve.

Rule-based Controls, the classical approach

In the rule-based strategy, the inlet temperature of type-2 loads ($\tilde{T}_{\text{L},i}^\text{in}$) is set to a constant value. In some rare cases, the plant operator manually switches between a summer and winter setpoint. It is stressed, however, that commonly chilled-water clients accept this setpoint to be actively managed within an agreed range.

As for the outlet temperature of each chiller ($T_{\text{CH},i}^\text{out}$), this is again set to a constant year-round setpoint. In some instances, season-dependent setpoints are considered, with operators being in charge of manually changing it once or twice per year. Once more, the admissible range of the chiller's outlet

[2] Clients of the plant who make use of chilled water for their cooling needs

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

Table 1: Project Task Breakdown with Responsibilities and Deadlines

Task	Activity Summary	Responsible Party	Deadline
1	Plant Selection + Meeting at CERN	CERN	31/07/24
2	Digital Twin Development	PoliMi	31/01/25
3	Assessment of Rule-based Controller (Simulation) + Intermediate Report	PoliMi + CERN	15/03/24
4	Development of Energy-Optimal Controller	PoliMi	30/05/25
5	Assessment of Energy-Optimal Controller (Simulation)	PoliMi	30/07/25
6	Integration of Energy-Optimal Controller in CERN's Industrial Control System	CERN	30/09/25
7	Assessment of Energy-Optimal Controller (Real) + Final Report	PoliMi + CERN	31/12/25

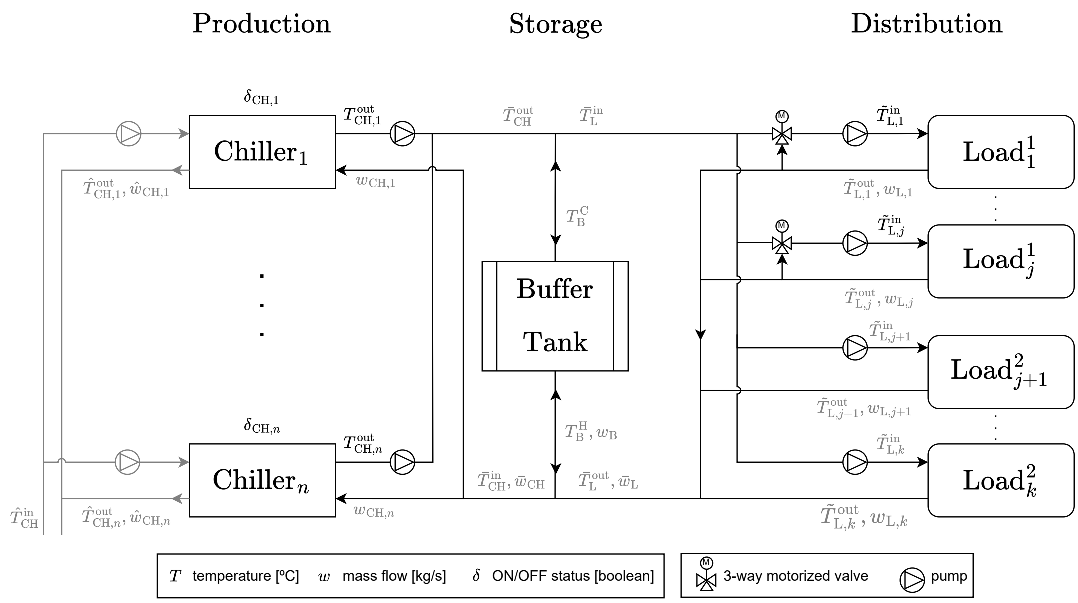

Figure 2: Scheme of chiller-based cooling plant.

temperature setpoint is not actively exploited by the control system.

The last control variable to manage is the ON/OFF status of each chiller ($\delta_{\mathrm{CH,i}}$). This is usually controlled by a complex staging logic built on the experience of system operators and control engineers. The objective of this logic is to ensure that the number of chillers operating can face the time-variant cooling loads of the plant. At CERN, this staging logic is based on continuously evaluating the water temperature and flow readings at key locations of the process, ensuring that these remain within safety limits. If limits are exceeded, the control logic updates the number of running chillers. In practice, CERN's rule-based logic converges in four if-based rules, combined with an operator-defined chiller priority list, which are used to sequentially increment or decrement the number of running chillers. These rules, which are detailed by equations (1)–(4), include a set of parameters (w_1, w_2, T_1, T_2) which must be tuned by the plant operators according to their experience.

Increment Rules:

$$\bar{w}_{\mathrm{L}} < \bar{w}_{\mathrm{CH}} + w_1 \qquad \text{for} \quad 1200\,\mathrm{s} \tag{1}$$

$$\bar{T}_{\mathrm{L}}^{\mathrm{in}} > T_1 \qquad \text{for} \quad 1200\,\mathrm{s} \tag{2}$$

Decrement Rules:

$$\bar{w}_{\mathrm{CH}} - w_{\mathrm{CH},i^*} > \bar{w}_{\mathrm{L}} - w_2 \qquad \text{for} \quad 1200\,\mathrm{s} \tag{3}$$

$$\bar{T}_{\mathrm{CH}}^{\mathrm{in}} - \bar{T}_{\mathrm{L}}^{\mathrm{in}} < T_2 \qquad \text{for} \quad 1200\,\mathrm{s} \tag{4}$$

with w_{CH,i^*} being the water flow of next chiller scheduled to stop according to the priority list.

A New Energy-Optimal Approach with MPC

The classical control strategy described before does not incorporate any kind of formal optimization, instead it relies only on the real-time monitoring of certain key process variables to ensure the plant fulfills its main mission: continuous provision of chilled water below a given temperature. One may argue, however, that this is insufficient in times that demand the sustainable operation of energy-intensive plants. In addition, CERN's ambition to reduce its environmental impact [4] calls for an alternative strategy, one that fulfils its cooling mission while minimizing the energy consumption.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

Model Predictive Control (MPC) is a control method that determines the control action by optimizing a predefined cost function — in this work, the energy consumption of the plant. The optimization accounts not only for the current system state, but also for predicted future states, anticipated disturbances[3], and both control and process constraints, all within a defined prediction horizon. Such an approach results in a predictive controller, in contrast to reactive strategies such as the rule-based controller.

The architecture of MPC-type controllers includes a plant model (or *digital twin*) and an optimization module. The former is a digital representation of the dynamics of the plant, allowing the simulation of the plant's behaviour without impacting the real system. As for the latter, it is the module in charge of finding the so-called optimal control sequence within the defined prediction horizon.

The energy-optimal MPC controller applied to the chiller-based cooling plant problem in hands, can be formulated in discrete-time k as follows:

$$\min \sum_{k=0}^{h} \mathcal{J}[k] \tag{5}$$

$$\underset{\forall k \in \mathcal{H}_k}{\textit{subject to}} \quad \mathcal{J}[k] = \sum_{i=1}^{n} \mathrm{P}^{\mathrm{e}}_{\mathrm{CH},i}[k], \quad \forall i \in \mathcal{N} \tag{6}$$

$$x[k+1] = \mathcal{M}(x[k], u[k], d[k]) \tag{7}$$

$$\delta_{\mathrm{CH},i}[k] \in \{0, 1\}, \quad \forall i \in \mathcal{N} \tag{8}$$

$$T^{\mathrm{out}}_{\mathrm{CH,min}} \le T^{\mathrm{out}}_{\mathrm{CH},i}[k] \le T^{\mathrm{out}}_{\mathrm{CH,max}}, \quad \forall i \in \mathcal{N} \tag{9}$$

$$\tilde{T}^{\mathrm{in}}_{\mathrm{L},i,\min} \le \tilde{T}^{\mathrm{in}}_{\mathrm{L},i}[k] \le \tilde{T}^{\mathrm{in}}_{\mathrm{L},i,\max}, \quad \forall i \in \mathcal{L} \tag{10}$$

with $\mathrm{P}^{\mathrm{e}}_{\mathrm{CH},i}$ as the electricity power draw by the i-th chiller, $\mathcal{J}$ as the controller's cost function, and $\mathcal{M}$ as the system model which approximates the plant dynamics. The vectors $x[k], u[k], d[k]$ are, respectively, the plant's states, control inputs, and disturbances. $\mathcal{H}_k = \{0, \ldots, h\}$, $\mathcal{N}$, and $\mathcal{L}$ are the vectors identifying the optimization step, chiller indexes, and load indexes, respectively. The parameters $T^{\mathrm{out}}_{\mathrm{CH,min}}, T^{\mathrm{out}}_{\mathrm{CH,max}}$ define the allowable range of the chillers' output temperature, whereas $\tilde{T}^{\mathrm{in}}_{\mathrm{L},i,\min}, \tilde{T}^{\mathrm{in}}_{\mathrm{L},i,\max}$ set the admissible working range of the inlet temperature of the i-th type-2 load. As shown in the formulation above, MPC extends beyond the sole minimization of energy consumption. It enables further adjusting the system's behaviour by introducing any mathematically formulated process constraint. For example, in the case-study applied in the project, there was the desire to incorporate four additional requirements to enhance the safety and reliability of the control logic:

1. Limit the rate of change of the inlet temperature of the loads:

$$|\tilde{T}^{\mathrm{in}}_{\mathrm{L},i}[k] - \tilde{T}^{\mathrm{in}}_{\mathrm{L},i}[k-1]| \le \dot{\tilde{T}}_{\mathrm{L},i} \tag{11}$$

   with $\dot{\tilde{T}}_{\mathrm{L},i}$ as the parameter defining the maximum rate of change.

2. Limit the maximum temperature in the buffer tank in favor of a cooling safety margin, may a running chiller stop unexpectedly:

$$T_{\mathrm{B}}[k] \le \min_{\forall i \in \mathcal{L}} \tilde{T}_{\mathrm{L},i,\min}, \quad \forall k \in \mathcal{H}_k \tag{12}$$

3. Limit the frequency of change of the chiller status, to reduce machine wear and potentially increase the chiller's lifetime:

$$\sum_{s=k-S}^{k} \delta_{\mathrm{CH},i}[s] \le 1, \quad \forall i \in \mathcal{N}, \quad \forall k \in \mathcal{H}_k \tag{13}$$

   with S as the minimum number of consecutive steps with one chiller's status change at most.

4. Bound the cooling power requested to each chiller

$$P^{c}_{\mathrm{CH},i,\min} \le P^{c}_{\mathrm{CH,i}}[k] \le P^{c}_{\mathrm{CH},i,\max} \tag{14}$$

$$P^{c}_{\mathrm{CH},i}[k] = c_{p,w} \cdot w_{\mathrm{CH},i}[k] \cdot (\bar{T}^{\mathrm{in}}_{\mathrm{CH}}[k] - \bar{T}^{\mathrm{out}}_{\mathrm{CH},i}[k]) \tag{15}$$
$$\forall i \in \mathcal{N}, \quad \forall k \in \mathcal{H}_k$$

   with $P^{c}_{\mathrm{CH,i}}[k]$ as the cooling power delivered by the i-th chiller at step k, and $c_{p,w}$ as the isobaric specific heat capacity of water.

An important advantage of the proposed control strategy over its rule-based counterpart lies in the ability to easily and formally define behavioural constraints, such as those described above. Implementing equivalent requirements through condition-based ('if–then') rules would be highly complex, if not impractical. It is stressed, however, that the optimization end-result and potential gains depend significantly on the flexibility allowed by the defined constraints.

DEVELOPMENT OF ENERGY-OPTIMAL CONTROLLER

Digital-Twin

The development of the optimal controller started with the creation of the plant's *digital twin*. Chiller-based plants can be decomposed in five classes of components: chillers, loads, buffer tank, 3-way valves, and piping joints, which can be modeled according to different methods. The strategy selected was a mixed approach between first-principles (i.e. physical equations) and data-driven models, according to the dynamics knowledge, complexity, and data availability for each component. A second order choice relates to the usage of a static or dynamic model. In practice, if the transient behavior of the component has reduced importance, the former is preferred for simplicity. Key additional assumptions which allowed to streamline the modeling process without compromising the overall model reliability were:

- Pumps were assumed as an integral component of either chillers or loads. This means that the water mass-flow and heat-load introduced by the pump is not explicitly decoupled from mass-flow and thermal balance of the associated chiller or load.

[3] Disturbance refers to any external signal to the plant which has an effect on its dynamics (e.g. heat loads introduced by the clients)

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

- Piping joints were assumed to be heat-loss free given that pipes are insulated.

On this basis, the modeling principles adopted for each of the five classes of elements composing the plant are summarized in Table 2.

Table 2: Modelling Strategies Used for the Digital-Twin Development

	Model Type	
Element Class	**Physics-Based**	**Data-Driven**
Piping Joints	x (static)	
Buffer Tank	x (dynamic)	
3-Way Valves	x (static)	
Chillers		x (dynamic)
Loads		x (dynamic)

All physics-based models were built with simple thermal energy and mass balances. For the chillers and loads, a data-driven modeling approach was adopted due to the large number of unknowns, such as their internal hydraulic structure and regulation strategies, as well as the availability of extensive historical data. When the available data was insufficient, these elements were operated under specific configurations to collect additional measurements. The development of the data-driven model begins with defining the relevant input and output variables. The output variables were straightforward to identify, as they correspond to those required by the overall model, whereas the selection of input variables required a more elaborate correlation analysis. For the case study in hands, the correlation analysis turned out to be very useful as it confirmed that two loads had a strong correlation with the external air temperature, while that was not true for the remaining. As for the data-driven modeling method, the choice felt for Neural Nonlinear Autoregressive Exogenous (NNARX) [5], a modern modeling strategy based on Recurrent Neural Networks (RNNs) which is well adapted to the modeling of dynamical systems.

The *digital-twin* implementation was done with the programming language Python™ [6] and composed as an optimization problem within the CasADi open-source framework [7]. The overall plant model validation was performed by a combination of visual inspection and detailed comparison between historical-data measurements and equivalent estimates by the digital model. As an example, Fig. 3 provides two of the comparisons performed. Ultimately, the *digital-twin* reliability was deemed reasonable for the accuracy required and, therefore, appropriate to progress with the optimizer development.

Optimizer

The optimizer is the second key component of the energy-optimal controller: it leverages the digital-twin in order to compute the optimal control sequence within a prediction horizon, given a certain goal, a set of constraints, and a series of predicted disturbances. The optimization problem

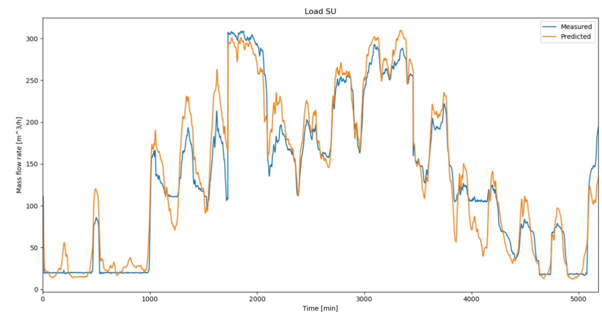

(a) Mass-flow rate of load.

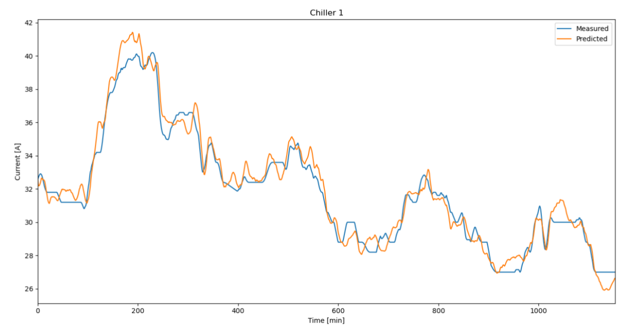

(b) Absorbed current of chiller.

Figure 3: Digital-twin validation examples.

considered, as formulated in Eqs. (5–15), is classified as a mixed-integer nonlinear problem (MINLP) given the nonlinear constraints and presence of one integer decision variable [Eq. (8)]. A practical simplification decision was to decouple the optimization problem in two in order to isolate the integer from the two continuous decision variables:

- **Real Time Optimizer (RTO)**: upper layer optimization problem integrating only the integer decision variable ($\delta_{\mathrm{CH},i}$) and running at a frequency equal to the operator-defined maximum frequency of change of the chiller status [see. Eq. (13)]. Running frequency: every 30 minutes, in this study;

- **Model Predictive Controller (MPC)**: lower layer optimization problem incorporating only the continuous decision variables ($T^{\mathrm{out}}_{\mathrm{CH},i}$ and $\tilde{T}_{\mathrm{L},i}$) and using directly the latest solution provided by the RTO to set the running chillers. Running frequency: every 5 minutes, in this study;

Similarly to the *digital-twin*, the two-layer optimizer was implemented in Python with casADi framework. The solver used is the Ipot [8], an open-source package for large-scale nonlinear optimization.

INTEGRATION OF ENERGY-OPTIMAL CONTROLLER WITH PLC-BASED CONTROL SYSTEM

This section presents the strategy followed for the integration of the energy-optimal controller within the plant's

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

control system. An important factor in the architecture design is that the energy-optimal controller can not run in a Programmable Logical Controller (PLC), i.e. the field-level controller in which rule-based logic is implemented, because this type of hardware does not support Python code. In addition, the MPC requires connection to internet services for weather forecast downloading what is not possible at the equipment level network of CERN, according to the organization's network security policies. To overcome these limitations, the architecture proposed (Fig. 4) foresees a multi-layer setup for the complete integration of the energy-optimal controller. The legacy rule-based control logic is preserved in the PLC code, allowing it to serve as a fallback in the event of an unexpected interruption of the energy-optimal controller or a communication failure with adjacent layers.

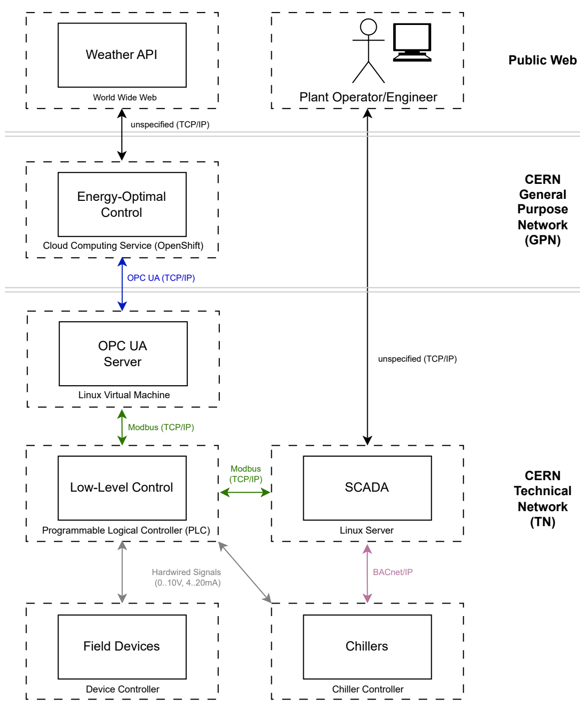

Figure 4: Controls architecture with energy-optimal controller.

Finally, the operator-interface to tune the parameters required by the energy-optimal controller is added to the existing SCADA[4] system of the plant (Fig. 5).

SIMULATION RESULTS

This section discusses the key findings resulting from the simulation-based comparison of both the rule-based controller and energy-optimal controller when subject to the plant's heat-load profiles registered between April and October 2024. The period of analysis selected covers the summer season and the running period of the largest accelerator at CERN, corresponding to the period of highest heat-load charge. The parameterization used for the rule-based strategy is the same as the one used in the real plant, whereas the energy-optimal controller parametrization was performed as thought to be used following the input provided by the plant operator. Table 3 provides an overview of the monthly energy consumption resulting from each control strategy. On average, the energy-optimal controller results in 11.4% in energy savings over the simulated months. When extending the results obtained to the complete year, conservatively assuming 10 MWh of energy savings for the non-simulated months, the total savings are 174 MWh, i.e. about 10% less when compared to the rule-based counterpart. Economically, the savings figure stand at 28 k€ on a yearly-basis[5].

Table 3: Electricity Consumption Analysis

	Electricity Consumption (MWh)		**Reduction**	
	Energy-Optimal	**Rule-Based**	**MWh**	**%**
April	112	128	16	12.5
May	133	146	13	8.9
July	197	222	25	11.3
August	192	220	28	12.7
September	136	156	20	12.8
October	119	131	12	9.2
Full Year (estimate)	**1549**	**1723**	**174**	**10.1**

CONCLUSION

This paper discusses a new energy-optimal control strategy for chiller-based cooling plants of accelerator facilities, focusing on a concrete case study at CERN. The research project which resulted from a successful collaboration between CERN and Politecnico di Milano, included the development of a digital-twin and design of a new plant controller based on Model Predictive Control techniques. Among others, it is shown that moving from the legacy rule-based control to an energy-optimal controllers presents several advantages:

- 10% yearly energy-savings;
- Enhancement in plant's operational reliability by adding control on important process variables not considered before (e.g. buffer tank temperature);
- Potential increase of chiller's lifetime by incorporating easy management of key functional metrics (e.g. frequency of start/stops, limitation of minimum and maximum cooling capacity);

[4] Supervisory Control and Data Acquisition

[5] Electricity cost considered was 160€/MWh [9]

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

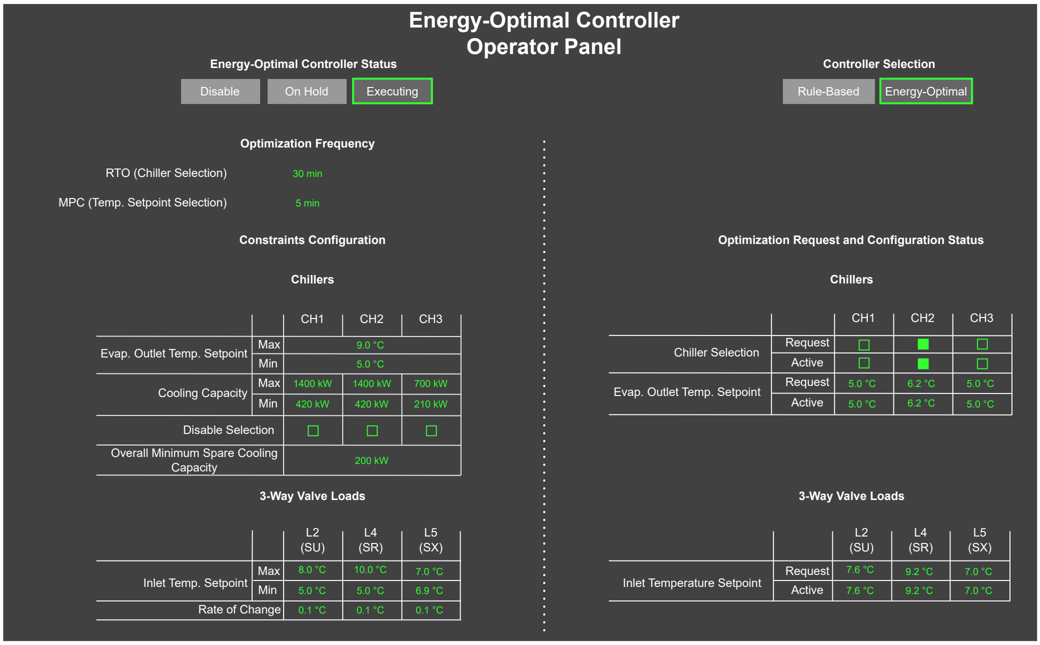

Figure 5: Operator panel for the energy-optimal controller.

It is noted, however, that implementing and operating an energy-optimal controller as the one developed requires that controls engineers become familiarized with model-based control techniques, in particular MPC. From the operator's perspective, once the operating principle is understood, the management and parametrization of the controller is straightforward and considered more intuitive than the rule-based counterpart.

Considering the successful results to date, the project team is motivated to bring the project to completion by deploying the new controller in the real plant environment before December 2025. In parallel, the team is evaluating how this energy-optimal control solution can be extended to other chiller-based cooling plants at CERN, knowing that the tools developed and experience accumulated with the first deployment shall significantly reduce the time-to-deployment in future projects.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the fruitful collaboration among the committed professionals and doctoral students at CERN and Politecnico di Milano. The work presented in this paper has greatly benefited from enriching discussions and contributions to the project's deliverables, particularly those of Alessio La Bella, Laura Boca di Giuli, Lorenzo Nigro, and Nikolina Bunijevac. The continuous support of the plant operators, Alexandre Marion and Florent Beaufils, has also been of paramount importance. Finally, the author expresses sincere appreciation to CERN's Knowledge Transfer group and the Cooling and Ventilation group of CERN's Engineering Department for their technical, organizational, and financial support.

REFERENCES

[1] D. Monteiro, R. Barillère, N. Bunijevac, I. Rühl, "Controls Optimization for Energy Efficient Cooling and Ventilation at CERN", in *Proc. ICALEPCS'23*, Cape Town, South Africa, Oct. 2023, pp. 1465–1469. doi:10.18429/JACoW-ICALEPCS2023-THPDP062

[2] Average Generation Capactiy of a Nuclear Reactor, https://www.energy.gov/ne/articles/nuclear-power-most-reliable-energy-source-and-its-not-even-close

[3] Vapor-Compression Refrigeration Cycle, https://www.buildingenclosureonline.com/blogs/14-the-be-blog-building-enclosure/post/90307-vapor-compression-refrigeration-cycle

[4] CERN's main objectives for the period 2021-2025, https://home.cern/resources/brochure/cern/cerns-main-objectives-2021-2025

[5] F. Bonassi, J. Xie, M. Farina, and R. Scattolini, "An Offset-Free Nonlinear MPC scheme for systems learned by Neural NARX models", in *2022 IEEE 61st Conf. on Decision and Control (CDC)*, Dec. 2022, pp. 2123–2128. doi:10.1109/cdc51059.2022.9992362

[6] Python Programming Language, https://www.python.org/

[7] CasADi: open-source tool for nonlinear optimization and algorithmic differentiation, https://web.casadi.org/

[8] Ipopt: open source software package for large-scale nonlinear optimization, https://coin-or.github.io/Ipopt/

[9] Electricity prices for non-household consumers, Eurostat, doi:10.2908/NRG_PC_205