2026 European Conference on Computing in Construction
Corfu, Greece
July 12-15, 2026

EDGE COMPUTING FOR FEDERATED LEARNING–ENABLED ENERGY MANAGEMENT AND MONITORING SYSTEMS

Mustapha Habib, Youssef Elomari, and Qian Wang
Building Technology and Design, Department of Civil and Architectural Engineering,
KTH Royal Institute of Technology, 11428 Stockholm, Sweden.

Abstract

Despite growing digitalization of energy management systems (EMS) in buildings and local networks, integrating data-driven monitoring and control remains challenging. Edge computing offers enhanced local processing that can be embedded in existing EMS infrastructures. Within the European HYSTORE project, two edge-computing use cases were developed: a network-level platform enabling coordinated EMS operation using building thermal mass estimates, and a building-level controller supporting advanced monitoring and control of thermal energy storage. Both architectures are evaluated for design, integration, and potential for federated learning, enabling secure, efficient, and privacy-preserving multi-agent control.

Introduction

The growing demand for energy-efficient and sustainable building management systems (BMS) has driven significant advancements in technology. Traditional BMSs often rely on centralized computing architectures, which, while effective for basic monitoring and control, fall short in handling the complexities of modern building energy systems. With the rise of distributed energy resources, renewable energy integration, and dynamic grid interactions, there is an increasing need for more responsive, decentralized solutions. Edge computing has emerged as a promising approach, enabling real-time data processing and decision-making closer to the source of data generation. This paradigm not only reduces latency but also enhances system resilience, making it an ideal candidate for addressing the evolving challenges in building management. Combined with artificial intelligence (AI) and the internet of things (IoT), edge computing offers the potential to revolutionize the way energy is monitored, optimized, and controlled in smart buildings.

A few contributions were proposed to employ edge computing-driven advanced services to upgrade building management. The authors in (C. H. Wong et al., 2021) discuss how AI systems can facilitate real-time analysis, optimization, and enhancement of building services. They demonstrate how AI can improve energy efficiency, enhance occupant satisfaction, mitigate risks, and reduce costs in building management and control systems. However, the authors also highlighted some limitations regarding some applications, including restricted access to large datasets, high initial investment costs, and the need for specialized expertise. Authors in (L. A. Husein et al., 2022) review the integration of IoT and AI technologies for improving BMSs, emphasizing the need for greater academic attention to these applications. While the review highlights the potential of IoT-AI methods, it notes that few studies have implemented these technologies in the building sector, and their full potential remains unrealized. The authors call for more research to address these gaps and advance the adoption of IoT and AI in building energy management. G. Carvalho et al. examine the use of AI, particularly machine learning (ML), to enable computation offloading in edge computing environments (G. Carvalho et al., 2020). AI and ML techniques are shown to address the challenges of computation offloading in these systems, offering more accurate, adaptable, and portable decision-making compared to traditional methods. The authors categorize various ML algorithms and provide an in-depth analysis of their application in Vehicular Edge Computing Networks, exploring the benefits and limitations of offloading tasks in edge computing systems. Key challenges discussed include determining when and where to offload tasks, managing decision-making costs, and minimizing the risk of errors that could undermine edge computing's advantages. A. Bourechak et al. explore the intersection of AI and edge computing across various IoT application domains, providing a comprehensive analysis of AI's role in edge-based applications across eight key areas (A. Bourechak et al., 2023). The paper includes a qualitative comparison of the reviewed works, an evaluation of the challenges and limitations of AI model development and implementation in edge environments, and a discussion on how AI can address these obstacles. In (G. Walia et al., 2024), authors present a detailed review of resource management challenges in Fog/Edge computing, categorizing these into areas such as resource provisioning, task offloading, resource scheduling, service placement, and load balancing. The paper examines AI-based and non-AI-based solutions, evaluating their performance metrics, limitations, and challenges. It highlights future research opportunities, including integrating emerging technologies like serverless computing, 5G, blockchain, digital twins, and quantum computing to enhance Fog/Edge frameworks for IoT applications. Challenges discussed include the complexity of managing heterogeneous resources and workload requirements, as well as the limitations of current non-AI-based solutions. In another study, B. Muniandi at el. discuss AI-driven energy management systems for smart buildings, emphasizing their ability to optimize energy use, enhance operational efficiency, and promote sustainability (B. Muniandi et al., 2024). Key functionalities include predictive analytics, adaptive control of building systems, and integration with renewable energy sources. These systems also enable smart buildings to participate in demand response programs, reducing costs and supporting grid stability. However, limitations such as data privacy and security concerns, as well as scalability and interoperability challenges across diverse building portfolios, are also underlined.

This study investigates two real-world applications in which edge computing offers significant potential to enhance energy management systems (EMSs) and building energy performance monitoring. The work covers the design, assembly, and commissioning of the proposed edge-computing solutions in two operational scenarios: a micro-district heating and cooling (micro-DHC) network and a residential building. Except the federated learning (FL) application for indoor temperature estimation in a micro-DHC–connected building, this study does not delve into the detailed mathematical formulation of each use case. The focus is instead on outlining and discussing the range of application scenarios in which FL-oriented edge-computing solutions can have a practical and impactful role.

Edge computing solutions: concepts and roles

This section presents the technical specifications of the two edge-computing concepts developed within the framework of the European HYSTORE project, along with the engineering details related to their design and integration into the existing systems.

Engineering basis of the proposed edge solutions

Solution 1: micro-district heating and cooling

Problem statement:

Isolated energy networks such as micro-district heating and cooling (micro-DHC) systems often suffer from limited coordination between buildings, which typically operate independently due to the lack of effective inter-building communication and system-level monitoring. These limitations become critical when complex thermal energy storage technologies, such as thermochemical (TCM) and phase-change-material (PCM) systems (see Fig. 1), are integrated, as they require advanced observer-based estimation and control strategies beyond the capabilities of conventional building management systems (BMSs). To address these challenges, the Montserrat micro-DHC demonstration within the HYSTORE project introduces a centralized edge-computing platform that enables network-wide data access and the deployment of advanced algorithms for state-of-charge estimation and model predictive control, supporting coordinated operation and improved energy management at the system level.

Solution 2: residential building substation

Problem statement:

In this situation, the purpose of the edge computing solution is to upgrade the computing capability of the BMS so it can handle the complex dynamics of eventually installed energy systems by incorporating advanced algorithms for control and monitoring.

With the current system, there has been an issue with the central air-handling unit (AHU) that uses a central heat pump system for heating. In extreme weather conditions, the AHU sometimes faces some challenges in ensuring acceptable thermal comfort. One of the potential solutions was to boost the temperature coming from the central HP system when it is overloaded, because it is providing heat to the other two buildings as well. In the HYSTORE project, the solution is partially or fully achieved by integrating an edge-driven sensible/latent energy system to boost the temperature in the AHU heating battery.

Edge-upgraded BMS:

The edge-driven solution for this use case is to develop two industrial panels, low- and high-level control. Each one is:

Low-level control:

- Direct communication of the LTES sensors.
- Multi-protocol enabled by Python programming.
- Handle the translation of the control commands to a set of actions on the hydronic system actuators.

High-level control:

- Long-term data storage.
- Complex AI-based and advanced control.
- Extend the control to be the energy management of the entire microgrid.

Within the framework of the HYSTORE project, the primary objective of the edge-driven solution is to enable high-level control of a hybrid sensible/latent thermal energy storage, PCM-based LTES (see Fig. 2) system through the implementation of observer-based model predictive control (MPC). By leveraging direct access to the system database, the edge platform can also coordinate the operation of the LTES with a heat pump in the Live-In Lab microgrid, thereby exploiting on-site photovoltaic generation and battery energy storage systems (BESS). The functional flowchart of the proposed edge-based solution to upgrade BMS functionality for indoor heating, which incorporates latent thermal energy storage (LTES), is presented in Fig. 3.

On-site installation, commissioning, and algorithms

This section details the installation procedures for each edge-computing solution and discusses the commissioning challenges encountered during deployment.

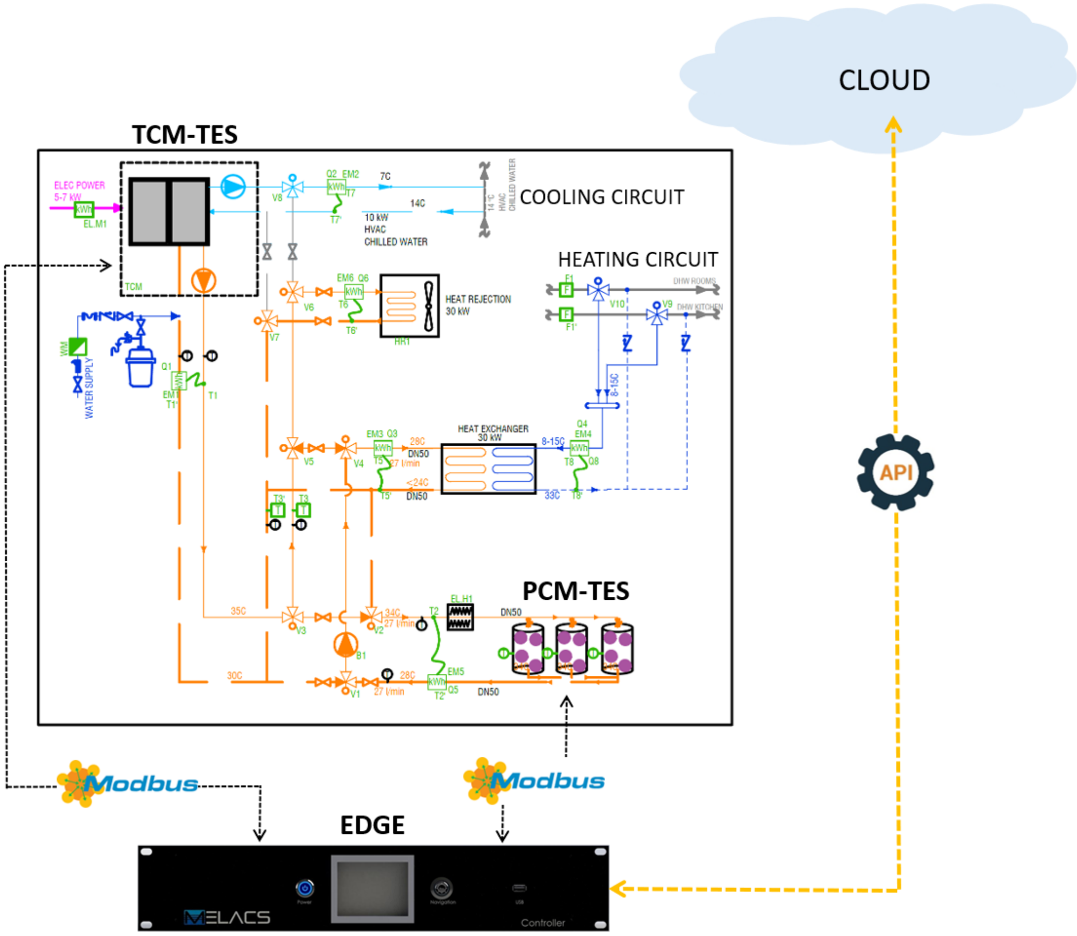

Figure 1: Edge solution as a high-level control for two LTES technologies: PCM and TCM at Montserrat micro-DHC.

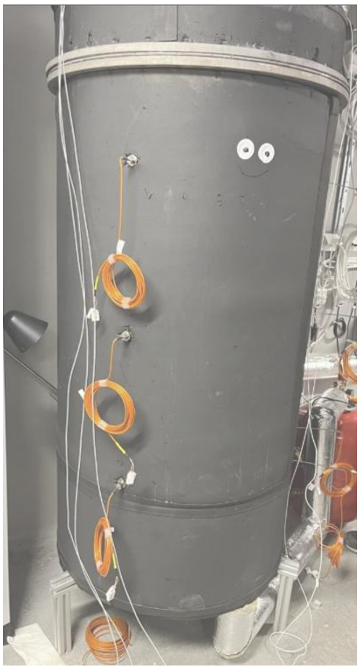

Figure 2: Installation of temperature sensors in the PCM-TES at the KTH live-in lab substation

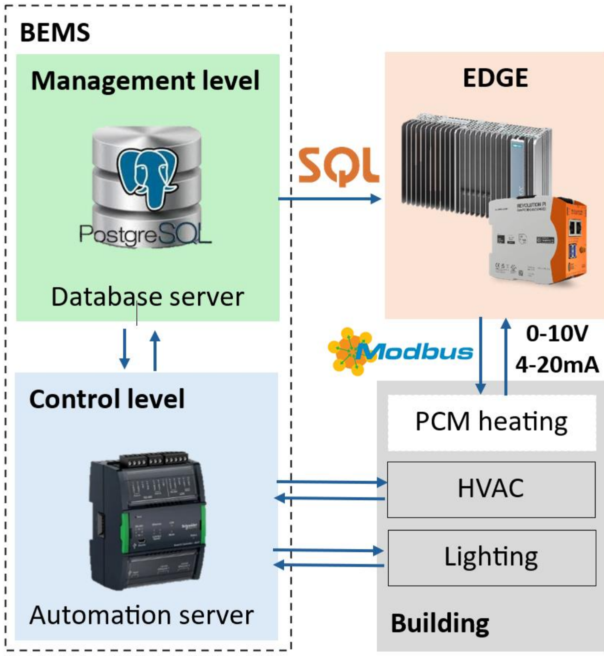

Figure 3: Edge solution as a high-level control for PCM LTES at KTH live-in lab

Solution 1

In Montserrat micro-DHC, a hotel building was selected for the installation of the two thermal energy storage technologies (TCM and PCM), as well as for deploying the second edge-computing solution (see Fig. 1). This choice was driven by the decision to exclude the sensors associated with the two LTES technologies from the central BMS database; as a result, direct physical communication with the devices was adopted. Communication with the various sensors and meters within the micro-DHC, including indoor and HVAC-related measurements from the hotel building, is achieved via the Modbus protocol over the local area network (LAN). Consequently, the edge computer was assigned a static IP address to ensure reliable identification and accessibility by other devices on the network.

Due to time constraints during the on-site commissioning phase of the HYSTORE project, full local communication with all network devices was not possible. To address this limitation, secure shell (SSH) access to the edge computer was configured remotely, allowing continued development and testing of the communication interfaces using Python.

Figure 4: Edge computing solution 1 installed in the hotel substation in Montserrat

Algorithms

In this real-world case, the edge computing platform implemented in Solution 1 serves two primary functions:

- Role 1: Centralized monitoring of the micro-DHC network.
- Role 2: High-level control of complex thermal energy storage (TES) systems.

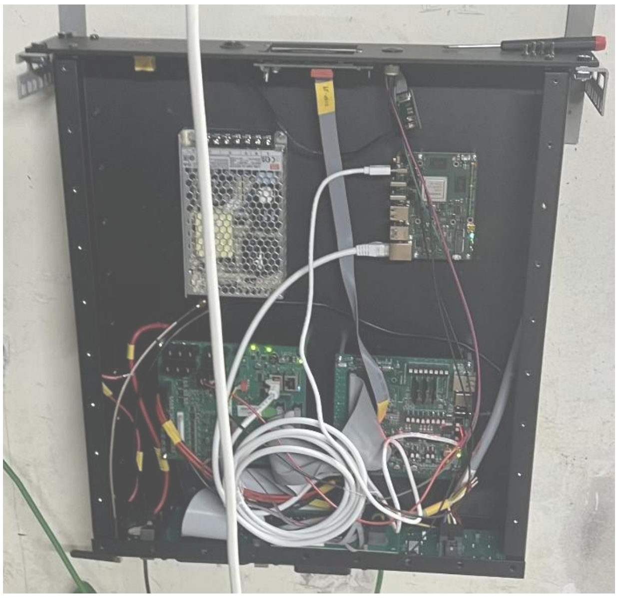

Figure 5: Components of the edge computing solution 1 installed in the hotel substation in Montserrat. It incorporates a GPU-accelerated controller with additional communication cards designed for industrial applications

For Role 1, several functionalities can be envisaged, including anomaly detection systems (ADS), centralized energy management systems (EMS), and coordinated demand-side management (DSM). The presence of a central computing entity communicating with all network components—such as HVAC systems, sensors, meters, and databases—enables a system-level perspective of energy flows and performance. This holistic view supports both network-wide energy management and detailed assessment of building and HVAC performance in heating and cooling generation. Within this scope, the present study focuses on the mathematical formulation of one representative function, which is presented at the end of this section.

For Role 2, a data-driven SoC estimation is implemented using a state observer. The algorithm combines Koopman theory with a linear Luenberger observer (KOOPMAN-LSO) to estimate internal PCM temperatures within the LTES based solely on the heat transfer fluid (water) outlet temperature. This approach has been validated in simulation using real operational data for multiple charging and discharging cycles of compact PCM TES, enabling on-site deployment without the need for intrusive PCM temperature sensors, which are often impractical in real applications.

A similar observer-based strategy will be developed to estimate the SoC of the thermochemical material (TCM) TES, which is co-located with the PCM-TES in the same building substation (see Fig. 1). Accurate SoC estimates allow the edge computer to optimally schedule charging and discharging of both LTES systems, maintaining thermal comfort while minimizing energy consumption. Two candidate control strategies will be tested for this purpose:

- Model predictive control (MPC)
- Reinforcement learning (RL)

Solution 2

In the residential building substation at the KTH Live-In Lab, the installation closely followed the approach used for Solution 1. A static IP address was assigned to the edge panel to enable connection to the LAN and allow real-time data queries from the system database via SQL. This database collects all operational data from the electrical systems, HVAC units, and indoor sensors across the microgrid. Since the edge panel was primarily tasked with managing the charging process of a PCM-TES, it was equipped with a Modbus coupler to allow direct physical wiring of the LTES sensor data to the controller. Additional analog and digital input/output cards were installed to interface with analog sensors and to control actuators, including ON/OFF switches and regulation valves. All remaining data acquisition and control of the heat pump are performed via Modbus TCP/IP.

Algorithms

As stated earlier, the planned role for the edge Solution 2 is to control the charging process of a hybrid latent/sensible using ground-source HP. Two main scenarios will be applied:

- A rule-based control (RBC), where a set of operational modes will be activated if certain conditions are met.
- MPC.

.

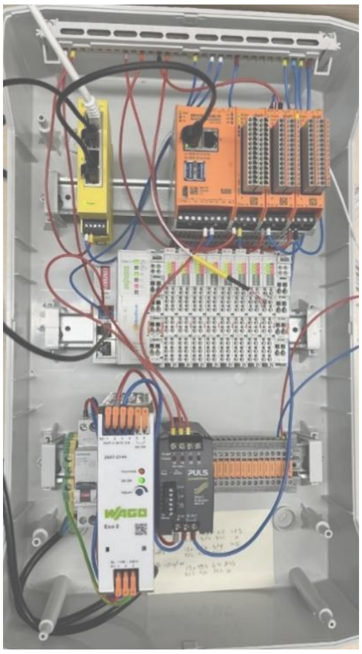

Figure 6: Components of the edge computing solution 2 (low-level control) meant for managing the PCM-TES and acts as an FL client. It incorporates a Python-enabled controller, IO cards, and Modbus coupler with the capacity of integrating 32 4-wire RTD temperature sensors needed to trach the PCM-TES SoC

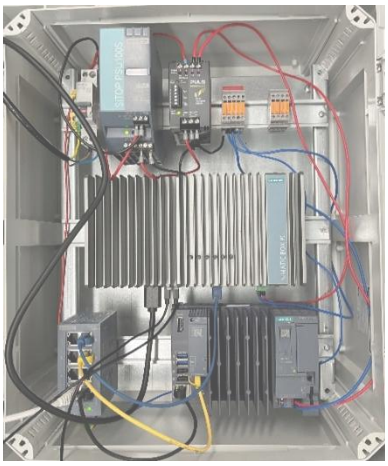

Figure 7: Components of the edge computing solution 2 (high-level control) meant for managing the KTH live-in lab microgrid, along with the different FL clients. It incorporates a high-capability industrial computer for running AI tasks and a PLC to enable communication with the field devices

For the rule-based control (RBC) strategy, a set of predefined operational modes is established, such as heating via the heat pump desuperheater, heating from the PCM-TES, combined operation, PCM-TES charging, and PCM-TES discharging. Each mode corresponds to a specific set of control actions executed by the edge panel through its embedded I/O modules. In this initial stage, the state-of-charge (SoC) of the PCM-TES is estimated using a simple physics-based formulation.

As operational data accumulates during RBC operation and is stored in a time-series database on the high-level control panel, a linear data-driven model of the PCM-TES is subsequently identified. This model enables the implementation of a linear model predictive control (MPC) scheme to maintain the existing heating functionality while achieving cost-effective operation based on dynamic electricity price forecasts and on-site photovoltaic (PV) power generation forecasts.

Federated learning concept

Role 1: microgrid management:

The KTH Live-In Lab can be regarded as a microgrid consisting of a local energy network serving three student residential buildings. The system integrates multiple energy resources, including renewable generation (PV, PVT, and wastewater heat recovery), energy storage systems (battery energy storage systems, sensible thermal energy storage, and geothermal storage), and diverse HVAC technologies, resulting in a sector-coupled energy system that includes heat pumps and air handling units. When operating in island mode, maintaining occupant comfort while optimizing energy utilization requires coordinated control across the three buildings, taking into account heating and domestic hot water demands, the state of charge of BESS and LTES units, renewable energy availability, and dynamic electricity tariffs.

Within this framework, each controllable HVAC system is associated with a federated learning client, implemented at the low-level edge computing layer (an example is shown in Fig. 5). These clients locally train and update models for load prediction, renewable energy forecasting, and BESS/TES state-of-charge estimation. Only the learned model parameters are transmitted to a central FL server (Fig. 6), which uses them to perform network-level optimization aimed at maximizing renewable energy utilization, preserving thermal comfort, and sustaining islanded operation whenever technically feasible. If required, the server can propagate updated model parameters back to the clients using FedAvg or other aggregation and fitting methods.

Role 2: Indoor temperature / thermal mass prediction:

One of the main advantages of deploying a central monitoring computer in an islanded energy network, such as the Montserrat micro-DHC is its ability to operate as an FL server that coordinates multiple clients distributed across substations and buildings. Each client locally collects indoor temperature and HVAC operational data, trains its own thermal model, and transmits only the learned model parameters to the server. In this study, an RC-model-based observer is developed at each client to estimate indoor temperature dynamics in real time. Only the linear RC model coefficients and the associated observer gains are shared with the FL server, while raw measurements of indoor temperature, heating demand, and HVAC operation remain strictly local, preserving data privacy. Depending on the control objective, the FL server can exploit the received models in two ways. First, it can use the individual observer-corrected models to compute real-time optimal setpoints for each building, enabling coordinated yet decentralized energy management. Second, the server can aggregate the observation-based RC models to construct a global network-level thermal representation, which is then used to derive optimal boiler-side control actions, such as supply temperature and pump operation. This dual use of federated, observation-based models allows the micro-DHC to improve both building-level comfort control and system-level energy efficiency without centralized data collection. The building management system (BMS) collects at each building the following measurements:

- Outdoor temperature T_{out}
- Supply air temperature T_{sup}
- Humidity RH (can be used indirectly)
- Air quality (CO_2 ppm → occupancy proxy)
- Setpoint cooling $T_{sp,c}$
- Setpoint heating $T_{sp,h}$

The dynamic model in this case is based on the state

$$x = T_r$$

Where T_r: room air temperature. Therefore, the continuous-time equation is:

$$C_{eq}\dot{T}_r = \frac{T_{out} - T_r}{\mathrm{R}_{eq}} + Q_{HVAC} + Q_{\mathrm{int}} \tag{1}$$

Where R_{HVAC} is effective HVAC thermal resistance and Q_{int} internal gains (occupancy + equipment). Thanks to the availability of the air quality sensor data, the heat source emissions related to the occupants (Qin) can be approximated as follows:

$$Q\mathrm{int} = \alpha.(CO_2 - CO_{2baseline}) \tag{2}$$

Humidity sensor data can be incorporated as a regression term in HVAC heat flow:

$$Q_{HVAC} = \beta_1.(T_{\mathrm{sup}} - T_r) + \beta_2.RH \tag{3}$$

$$\beta_1 = R_{HVAC}^{-1} = \begin{cases} k_c & T_r > T_{sp,c} \\ k_h & T_r < T_{sp,h} \\ 0 & deadband \end{cases}$$

The input vector is:

$$u = \begin{bmatrix} T_{out} \\ T_{sup} \\ RH \\ \mathrm{CO_2} \end{bmatrix} \tag{4}$$

Output:

$$y = T_r \tag{5}$$

Continuous-time state model is:

$$A = -\left(\frac{1}{C_{eq}\,\mathrm{R}_{eq}} + \frac{\beta_1}{C_{eq}}\right) \tag{6}$$

$$B = \begin{bmatrix} \dfrac{1}{C_{eq}\,\mathrm{R}_{eq}} & \dfrac{\beta_1}{C_{eq}} & \dfrac{\beta_2}{C_{eq}} & \dfrac{\alpha}{C_{eq}} \end{bmatrix} \tag{7}$$

Discrete state-space representation is:

$$x_{k+1} = A.x_k + B.u_k \tag{8}$$

Parameters to be identified:

$$\left\{C_{eq},\, \mathrm{R}_{eq},\, R_{HVAC},\, \alpha,\, \beta_1,\, \beta_2\right\}$$

Since no RC model can perfectly capture building energy performance under changing physical conditions, this study proposes incorporating real-time model correction using sensor measurements. A linear Luenberger observer is employed to fuse measured data with model predictions, thereby mitigating the effects of sensor noise, modeling inaccuracies, and external disturbances (e.g., weather variability, occupancy changes, internal heat gains, and unmodeled HVAC dynamics). The indoor temperature estimate can then be formulated as follows:

$$\begin{aligned}\hat{x}_{k+1} &= A.\hat{x}_k + B.u_k + L(y_k - \hat{y}_k) \\ \hat{y}_k &= C.\hat{x}_k\end{aligned} \tag{9}$$

Local learning objective (client optimization)

Each client independently solves a local prediction-error minimization problem**:**

$$\min_{\theta^i} \sum_{k=1}^{T} \left\| y_k^{(i)} - \hat{y}_k^{(i)}(\theta^{(i)}) \right\| \tag{10}$$

subject to physical constraints on RC parameters (stability, positivity) and local data only.

Federated learning formulation (server-side)

Let: $\theta^{(i)} = \{A^{(i)}, B^{(i)}, L^{(i)}\}$ be the model and observer parameters learned locally. At the communication round r, each client sends $\theta_r^{(i)}$ to the central server

Server aggregation (FedAvg-style)

Federated averaging (FedAvg) computes a weighted average of client parameters:

$$\theta_{global} = \frac{\sum_{i=1}^{N} n_i \theta^{(i)}}{\sum_{i=1}^{N} n_i} \tag{11}$$

n_i is the weight for client $\boldsymbol{i}$, typically number of samples used to train the local model. These aggregated coefficients at the FL server side can be used to represent a global heating demand dynamic of the whole network so it is feasible to optimize the heating energy extracted from the main boilers.

- Let $\theta^{(i)} = \{A^{(i)}, B^{(i)}, L^{(i)}\}$.
- One can compute a global A and B as:

$$A_{global} = \frac{\sum_{i=1}^{N} w_i A^{(i)}}{\sum_{i=1}^{N} w_i}, \quad B_{global} = \frac{\sum_{i=1}^{N} w_i B^{(i)}}{\sum_{i=1}^{N} w_i} \tag{12}$$

- Weights w_i are to be chosen proportional to building thermal mass or typical heat demand.
- No need for L_global, since the observer is only relevant for estimating individual building states.

In the first case, the FL server does not attempt to aggregate building models into a global network representation. Instead, it uses each building's RC model with observer-corrected states to compute real-time optimal setpoints locally. The server sends these setpoints back to the respective clients, allowing each building to regulate its indoor temperature efficiently while respecting occupant comfort and local constraints. Here, the server acts as a coordinator and optimizer, leveraging up-to-date state information to suggest adjustments, but the boiler and network operation are effectively treated as secondary, since the focus is on individual building performance rather than total network-level energy optimization. In the second case, the FL server aggregates the observer-corrected RC models from all connected buildings to construct a global network model. Using this network-level model, the server can simulate how changes in supply temperature, flow rates, or pump operation affect the entire micro-DHC system. With this approach, the server can compute optimal boiler-side control actions that minimize total energy consumption, prevent peak loads, and maintain network stability, while still ensuring buildings receive the heating they need. The observer-based global model allows the server to incorporate real-time deviations in building states, making the control decisions robust to disturbances and heterogeneous building dynamics.

Simulation results

This section details the preliminary simulation results for the edge application of Solution 1. Due to a more streamlined installation and commissioning process, Solution 1 reached the testing phase ahead of Solution 2. Figure 8 illustrates the performance of an RC model in predicting indoor temperatures for a hotel building integrated with a Micro-DHC. Even in its base form – without observer-based corrections – the RC model demonstrates strong predictive accuracy. Future enhancements will involve real-time state observer corrections utilizing micro-DHC heat exchanger data, which will eventually be integrated into an optimal control framework, such as Model Predictive Control (MPC)

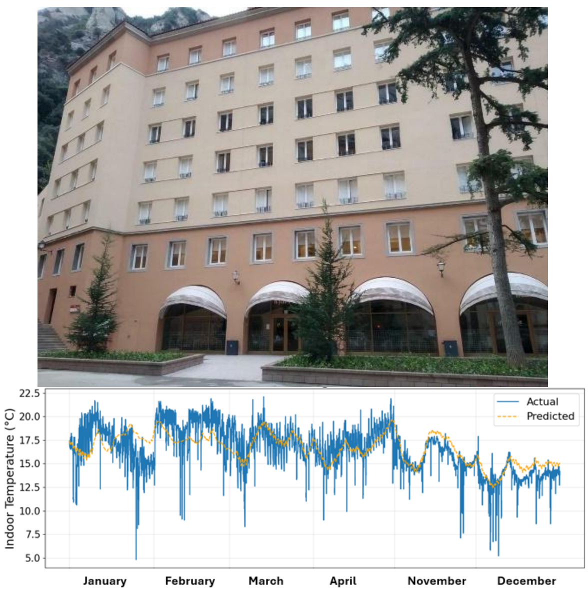

Figure 8: Indoor temperature prediction using trained RC model: (top) the view of the Abat Oliba hotel concerned by the study, (bottom) the seasonal prediction performance for 2023

Fig. 9 illustrates another application featuring an advanced data-driven observer implemented for LTES monitoring. This technique utilizes the KOOPMAN theory to model the LTES, leveraging collected datasets to construct a linear representation within a lifted observation space. The linearity of this space facilitates the use of a classical Luenberger Observer, which compensates for low-quality measurement data and estimates unmeasured PCM temperature.

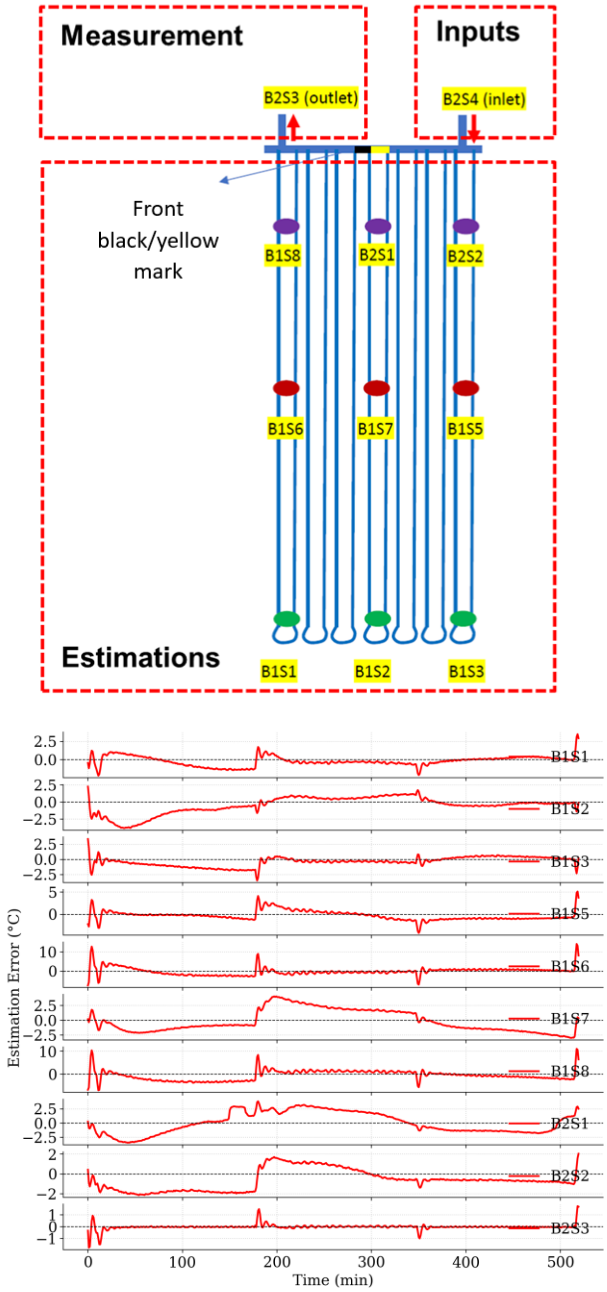

Figure 9: Edge-hosted observer for PCM temperature estimation of LTES (top)sensor placement in the LTES during KOOPMAN model development, (bottom) estimation error for all PCM temperatures

Conclusions

This paper presented two real-world edge-computing implementations developed within the HYSTORE project, targeting coordinated energy management in micro-DHC and residential microgrid environments. In the Montserrat micro-DHC, a centralized edge platform enables network-level monitoring and advanced control of complex thermochemical and PCM thermal energy storage systems using observer-based state estimation and anomaly detection. At the building level, the KTH Live-In Lab implementation demonstrates how similar edge concepts support PCM-TES operation and coordination with heat pumps, photovoltaic generation, and battery storage. Together, these cases show how edge computing can overcome coordination and control limitations of conventional EMSs, enabling observer-driven, privacy-preserving, and scalable control of complex energy systems in isolated energy networks.

Acknowledgments

This work was funded by the European Commission under Horizon Europe (Grant 101096789).

References

C. H. Wong et al, IOP Conf. Ser.: Earth Environ. Sci. 881 012044, 2021, DOI 10.1088/1755-1315/881/1/012044

Lama Alhaj Husein, Imad Alsyouf, Emad Mushtaha and Ahmad Alzghoul, Towards High-Performance Buildings using IoT and AI technologies: A Comprehensive Review, Proceedings of the International Conference on Industrial Engineering and Operations Management Nsukka, Nigeria, 5 - 7 April, 2022.

G. Carvalho, B. Cabral, Vasco Pereira, Jorge Bernardino, Computation offloading in Edge Computing environments using Artificial Intelligence techniques, Engineering applications of artificial intelligence, 2020.

Amira Bourechak, Ouarda Zedadra, Mohamed Nadjib Kouahla, Antonio Guerrieri, Hamid Seridi and Giancarlo Fortino, At the Confluence of Artificial Intelligence and Edge Computing in IoT-Based Applications: A Review and New Perspectives, Sensors 2023, 23, 1639. https://doi.org/10.3390/s23031639

G. Walia, Mohit Kumar, S. Gill, AI-Empowered Fog/Edge Resource Management for IoT Applications: A Comprehensive Review, Research Challenges, and Future Perspectives, EEE Communications Surveys and Tutorials, 2024.

Balakumar Muniandi, Purushottam Kumar Maurya, CH Bhavani, Shailesh Kulkarni, Ramswaroop Reddy Yellu, Dr. Nidhi Chauhan. AI-Driven Energy Management Systems for Smart Buildings, Power System Technology, Volume 48 Issue 1, March 2024.