Proceedings of Smart and Sustainable Built Environment Conference Series

SASBE 2026 aims to encourage the international exchange of innovative ideas between researchers from academia and industry. In addition to knowledge dissemination, the conference offers a valuable platform for professional networking, particularly benefiting university professors, graduate students, and postdoctoral researchers.

Research Article

Digital Twins for Mission-Critical Infrastructures: A BIM-Based Framework for Real-Time Monitoring and Decision Support

Viviana Vaccaro[1], Lavinia Chiara Tagliabue[1], Robert Birke[1], Marco Aldinucci[1], Silvia Meschini[1]

[1]Department of Computer Science University of Turin, Italy
Correspondence: viviana.vaccaro@unito.it, laviniachiara.tagliabue@unito.it, robert.birke@unito.it, marco.aldinucci@unito.it, silvia.meschini@unito.it

Copyright: Copyright: © 2026 by the authors. ISSN: 2979-1804

SASBE is an open-access proceedings distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0).
View this license's legal deed at https://creativecommons.org/licenses/by/4.0/

Received: dd/mm/yyyy
Revised: dd/mm/yyyy
Accepted: dd/mm/yyyy
Published: dd/mm/yyyy

Volume: xx
Issue: xx
Pages: pp-pp

Abstract

Digital Twins are increasingly adopted to support the operational management of complex and energy-intensive infrastructures, yet their practical implementation often remains fragmented across heterogeneous platforms and proprietary monitoring systems. In data center environments, this fragmentation limits the interpretability of real-time data and weakens the connection between operational measurements, spatial context, and asset-level decision-making. This paper presents a BIM-based Digital Twin architecture developed for the HPC4AI (High-Performance Computing for Artificial Intelligence) data center, hosted at the University of Turin's Computer Science Department, within the framework of the DYMAN project, which aims to develop advanced solutions to reduce the impact of thermal energy management in high-performance computing installations and data server rooms. The proposed approach integrates real-time IoT monitoring with Building Information Modelling through a modular layered architecture that decouples sensing, time-series data management, and digital representation. Sensor data are stored in a time-series database and synchronized into the BIM environment via automated Dynamo workflows and embedded Python scripts, where shared parameters enable semantic mapping to building elements. This supports near real-time spatial contextualization, rule-based alerting, and native BIM visualisation without reliance on external dashboards. The case study demonstrates how the BIM model evolves from a static geometric artifact into an operational interface capable of supporting situational awareness, diagnostics, and integration with predictive analytics and energy optimization strategies. The paper discusses system architecture, data synchronization mechanisms, and transferability to other mission-critical infrastructures, highlighting the role of data-driven integration strategies in advancing scalable and interoperable Digital Twins for smart and sustainable built environments.

Keywords: Digital Twin, BIM–IoT Integration, Data Centers, Real-Time Monitoring, Facility Management, Data-Driven Operation, Model Predictive Control.

Highlights

- A modular BIM–IoT architecture enables real-time monitoring of HPC data centers natively in Revit.
- BIM-integrated predictive control optimizes cooling setpoints with MAE = 0.39°C.
- BIM evolves from static artifact to active decision-support interface for mission-critical DTs.

Cite this article:
Lastname, F. M., & Lastname, F. M. (Year). Title of the article. Title of the Journal, Volume (Issue), Page range. https://doi.org/10.66408/sasbe.2026.YourPAPERid

1 Introduction

In recent years, the concept of Digital Twin (DT) has emerged as a key paradigm for managing complex and mission-critical infrastructures, where real-time awareness, system reliability, and energy efficiency must be balanced. In smart and sustainable cities, DTs connect physical assets, sensor networks, and data-driven models, enabling continuous monitoring and informed decision-making (Jafari et al., 2023; Rasheed et al., 2020; Zhang et al., 2021; Qi & Tao, 2018). However, the fragmentation of current monitoring practices, particularly in high-complexity infrastructures, limits diagnostic activities and anomaly localization, making them dependent on expert knowledge rather than integrated, interpretable representations.

Data centers are among the most critical of such systems, acting as the backbone of digital services while contributing significantly to urban energy demand (IEA, 2023). Current monitoring relies on heterogeneous platforms and proprietary Building Management Systems (BMS), which provide high-frequency data but lack spatial and semantic contextualization.

Previous research on the HPC4AI facility, a HPC data center at the University of Turin, focused focused on the development of a dynamic thermal and energy simulation model capable of reproducing the behaviour of the data hall under varying IT loads and climatic boundary conditions. That modelling phase enabled scenario-based evaluation of cooling system configurations, free-cooling potential, and energy performance indicators, providing insight into optimization strategies under controlled assumptions (Vaccaro et al., 2025). The present work represents a methodological progression: by coupling IoT data streams with BIM and extending the framework with a closed-loop MPC engine, the proposed DT transitions from descriptive monitoring to active, energy-aware decision support.

The main contributions are:

- the definition of a modular, layered DT architecture that decouples sensing infrastructure, data management, and digital representation, supporting scalability and interoperability;
- the development of an automated BIM–IoT synchronization workflow based on open protocols, enabling near real-time integration of time-series sensor data into the BIM environment;
- the introduction of a parameter-based sensor-to-BIM mapping strategy ensuring stable spatial and semantic contextualization of operational data;
- the implementation of a closed-loop Model Predictive Control (MPC) engine based on a calibrated first-order Resistance-Capacitance (RC) thermal model, fully integrated in the Dynamo workflow and validated against real sensor data stored in InfluxDB (Mean Absolute Error MAE = 0.39°C, Mean Absolute Percentage Error MAPE = 1.73% at a 60-minute prediction horizon);
- the demonstration, through a real-world case study, of how a BIM model can evolve into an operational DT supporting situational awareness, diagnostics, and energy optimization.

2 Background and Related Work

2.1 Digital Twin for the Built Environment

In the built environment domain, Digital Twin has progressively evolved from a generic notion of digital representation toward an operational paradigm centred on continuous data exchange between physical assets and their digital counterparts. Operational Digital Twins can be understood as digital systems that are persistently synchronized with physical entities through real-time or near real-time data flows, enabling monitoring, analysis, and decision support during operation rather than only during design or commissioning phases.

This distinguishes them from static BIM models, which provide geometric and semantic descriptions but remain disconnected from actual operational state, and from physics-based simulations, which allow the exploration of system behaviour under predefined scenarios but typically rely on offline computation and predefined assumptions. Operational DTs integrate both paradigms, combining structured digital representations with live data streams to enable descriptive, diagnostic, and, when extended with predictive modules, prescriptive capabilities.

2.2 BIM–IoT Integration and Data-Driven Approaches

BIM has been widely adopted as a central information repository throughout the lifecycle of buildings and infrastructure. In recent years, increasing attention has been devoted to the integration of BIM with IoT technologies, with the aim of extending BIM from a static documentation tool toward a data-driven operational platform (Pinmas et al., 2022). A substantial body of research has explored BIM–IoT integration for facility management, smart buildings, and indoor environmental monitoring (Pasini et al., 2016). Proposed solutions range from middleware-based architectures to cloud platforms and custom APIs linking sensor measurements to BIM elements.

Two recent works are particularly relevant for positioning the present contribution. Harode et al. (2023) proposed a functionality-based DT architecture for healthcare facility management, integrating BIM geometry with real-time IoT telemetry through a cloud-based pipeline (Power BI, VCAD, Azure) and visualizing data via external dashboards. Muñoz Pavón et al. (2025) developed a BIM-based DT platform for campus management, linking BIM to custom web applications, SQL databases, and IoT sensors for space reservation, occupancy monitoring, and HVAC control.

The present work differs from both approaches along three methodological axes. First, whereas Harode et al. rely on Power BI as an external central platform, decoupling the BIM model from the live data layer, the framework proposed here uses the BIM authoring environment (Autodesk Revit) as the primary operational interface, embedding sensor data directly into BIM shared parameters and leveraging native visualisation mechanisms without dependence on third-party dashboards. Second, in contrast to the general-purpose campus platform of Muñoz Pavón et al., the proposed approach is purpose-designed for mission-critical HPC infrastructure and incorporates domain-specific alerting logic aligned with ASHRAE guidelines for IT environments, and Power Usage Effectiveness (PUE) monitoring. Third, the present framework extends beyond monitoring and alerting to integrate a closed-loop MPC engine within the Dynamo workflow, enabling active, predictive control of cooling systems directly within the BIM environment.

Despite these advances, automation in existing implementations remains frequently limited to data visualization, with little support for semantic mapping, rule-based interpretation, or predictive control within the BIM environment. Addressing these limitations requires workflows that prioritize modularity and interoperability, enabling continuous data integration and active decision support while preserving BIM integrity.

3 Methodology and System Architecture

3.1 Design Principles and Objectives

The development of an operational DT for complex and energy-intensive infrastructures requires a methodological approach that balances integration, flexibility, and long-term maintainability. The proposed framework is therefore based on a modular architecture that separates sensing, data

management, and digital representation, enabling integration without creating dependency. Although implemented using specific software tools, these represent operational instances of a transferable methodological structure designed to avoid tight coupling between system layers.

Within this layered architecture, BIM assumes the role of semantic and spatial interface of the DT. It is not responsible for data acquisition, storage, or advanced processing; instead, it provides the structured environment in which operational information is contextualized. By linking measurements to explicitly defined building elements, systems, and assets, the BIM model establishes the spatial and functional relationships necessary to interpret data in relation to physical layout and operational hierarchies, enabling coherent interpretation, aggregation, and visualisation within a unified spatial framework.

The separation of layers ensures independent evolution of components: changes in sensing technologies, monitoring platforms, or databases affect only their respective layers, provided the data interface remains consistent. The methodology is conceptually independent of specific technologies and transferable to alternative BIM tools, databases, and sensing infrastructures.

Based on these principles, the methodology pursues four main objectives: (1) spatial contextualization of operational data, ensuring persistent association between real-time measurements and building elements; (2) automation of data integration and updating processes, enabling near real-time synchronization with minimal manual intervention; (3) interoperability and extensibility, through modular workflows, open communication protocols, and parameter-based mapping strategies; (4) closed-loop decision support through a Model Predictive Control engine that computes and communicates optimal cooling actions within the BIM environment, based on a physics-informed thermal model continuously fed by live sensor data.

3.2 Layered Digital Twin Architecture

Consistent with the principles outlined above, the proposed DT framework adopts a five-component layered architecture.

The **physical sensing layer** acquires operational and environmental data through a heterogeneous network of sensors, reflecting the physical state of the infrastructure, such as thermal conditions, humidity, and other operational indicators relevant to energy efficiency and reliability.

The **data management layer** handles the storage, retrieval, and historical organization of sensor data in InfluxDB, a time-series database optimized for high-frequency, time-stamped measurements.

The **integration and synchronization layer** mediates between the database and the BIM environment. It retrieves, preprocesses, and validates sensor data before synchronizing them with the digital model according to predefined mapping rules.

The **digital representation layer** contextualizes and visualizes operational data through shared parameters, schedules, and native BIM tools, evolving the model from a static artifact into a continuously updated operational interface.

A fifth functional module, the **MPC engine**, operates at the interface between data management and digital representation, consuming real-time sensor data and writing predictive control outputs back into BIM shared parameters, closing the control loop and extending the DT toward prescriptive decision support.

Figure 1 provides a schematic overview of the five-layer architecture and the data flows between components, including the lateral MPC module and its bidirectional interaction with the data management and digital representation layers.

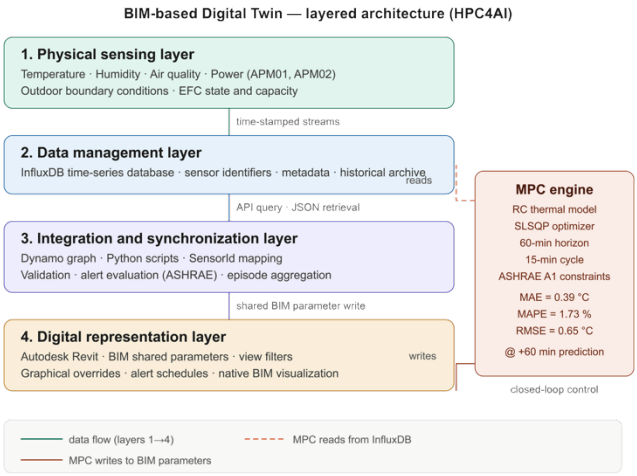


Figure 1. Five-layer architecture of the proposed BIM-based Digital Twin framework. The MPC engine module operates across layers 2–4, closing the control loop between InfluxDB sensor data and BIM shared parameters.

3.3 Data Synchronization and Automation Workflow

The interaction between the data management layer and the digital representation layer is operationalized through an automated Dynamo-based workflow with embedded Python scripts. API-based queries dynamically retrieve sensor data using unique sensor identifiers and configurable temporal parameters. Each synchronized value is associated with its timestamp, enabling detection of stale or missing data. Validation routines filter invalid measurements before writing results into predefined shared BIM parameters, preserving semantic consistency and enabling downstream alerting, visualization, and MPC integration. By automating the entire synchronization process, the workflow minimizes manual intervention and supports repeatable, near real-time updates of the DT.

3.4 Limitations and Transferability Considerations

While the proposed framework demonstrates practical feasibility in a real-world operational context, several limitations warrant explicit acknowledgement.

The current implementation relies on Autodesk Revit as the BIM authoring environment and Dynamo as the visual programming interface. In many industry contexts, contractors may deliver BIM assets in exchange formats such as IFC or Navisworks (.nwd/.nwc) rather than native Revit files. In such scenarios, the Dynamo-based synchronization workflow would require adaptation. Possible mitigation strategies include: (a) re-implementing the synchronization logic using IFC-compliant tools (e.g., ifcOpenShell, Speckle); (b) maintaining a lightweight Revit receiver model populated from IFC geometry via automated import pipelines; or (c) decoupling the parameter-writing logic from the BIM authoring tool and targeting an open data exchange format at the middleware layer. The methodology's conceptual independence from specific tools is intended precisely to support such adaptations.

A second limitation concerns the requirement that BIM families be enriched with shared parameters prior to deployment. In contexts where BIM deliverables are contractually specified without operational parameter schemas, this enrichment may become the responsibility of the owner or

facility manager, adding effort and cost. This challenge reflects a broader gap in current BIM contractual practice regarding operational DTs, and aligns with ongoing standardization efforts under ISO 19650.

Finally, the current validation is based on a single primary case study (HPC4AI data center). Preliminary work is underway applying a comparable framework to a museum environment, where different operational requirements (heritage conservation, visitor management, indoor climate control) and a distinct BIM maturity context provide a complementary validation scenario. Results from this second deployment will be reported in forthcoming work.

4 Sensor–BIM Mapping and Logic

4.1 Parameter-Based Strategy

Persistent association between physical measurements and digital counterparts is achieved through a parameter-based mapping strategy. Each sensor is identified by a unique ID stored as a shared BIM parameter attached to the relevant model elements (sensor families, racks, or spatial zones), ensuring stable mapping regardless of geometric or Level of Detail (LoD) changes. This decouples sensing logic from geometric representation: sensors may be explicitly modelled when spatial placement is relevant, or logically associated without dedicated geometry, while preserving semantic linkage at the parameter level.

The parameter-driven mapping enhances extensibility: new sensors, variables, or data sources can be incorporated by extending the parameter schema without modifying the synchronization workflow or the geometric model, providing a scalable foundation for data-intensive and evolving operational environments.

4.2 Alerting and Interpretation Logic

Beyond data synchronization, an operational DT must support actionable interpretation. A rule-based alerting logic operates on synchronized sensor data, translating measurements into diagnostic signals configured directly within Dynamo and associated Python routines, allowing alert criteria to be tuned for different assets or operating conditions while keeping the BIM model stable. Individual threshold violations are aggregated into temporally coherent alert episodes, defined by configurable persistence criteria, to mitigate transient fluctuations. Conditions such as missing data, delayed updates, or static readings are treated as dedicated alert types, ensuring continuous monitoring of both environmental performance and sensor reliability. Unlike conventional dashboards, this approach evaluates alerts directly in relation to BIM elements, embedding alert states within the Digital Twin's parametric structure and enabling integrated querying, aggregation, and historical analysis within the same spatial environment.

4.3 Native BIM-Based Visualisation

The digital representation of operational conditions is implemented using native BIM visualisation mechanisms, such as view filters, graphical overrides, and schedules.

Filters are defined based on shared parameters storing live measurements and alert states, enabling dynamic classification of elements according to their operational condition. Dedicated views can be

configured for specific monitoring domains, such as temperature, humidity, or alert status, allowing users to focus on relevant aspects.

Graphical overrides may be applied directly to sensor families or propagated to higher-level BIM elements, such as racks or zones, enabling visual emphasis of critical conditions. This improves interpretability in dense and complex environments, where spatial configuration is central to effective diagnosis and intervention, by allowing users to perceive operational states directly within their architectural and system context. This not only facilitates rapid anomaly localization but also provides diverse stakeholders, including facility managers, energy specialists, and designers, with access to a shared and interpretable DT representation, without relying on dedicated visualisation platforms.

5 Case Study: HPC4AI Data Center

5.1 Facility Overview

The proposed Digital Twin framework was deployed and validated on the HPC4AI (High-Performance Computing for Artificial Intelligence) data center, located at the Computer Science Department of the University of Turin (Fig.2). The Tier-III–oriented facility supports high-performance computing and AI workloads, with an installed capacity of approximately 250 kVA and 16 server racks arranged in a centrally contained hot aisle configuration to optimize airflow and cooling efficiency.
Cooling is provided by roof-mounted adiabatic chillers operating under multiple modes, including indirect free cooling, evaporative cooling, and mechanical compression, prioritizing low-energy strategies when climatic conditions permit, contributing to reduced energy consumption and lower environmental impact (Vertiv, 2021). Electrical monitoring extends from the main transformer to rack-level power distribution units, enabling detailed energy analysis and real-time estimation of performance indicators such as Power Usage Effectiveness (PUE) (University of Turin, 2024a).



Figure 2. Interior view of the HPC4AI data center.

5.2 IoT Infrastructure and Data Sources

The monitoring infrastructure is structured as a multi-layered sensing ecosystem combining building-level supervision, rack-level monitoring, and distributed environmental sensing.

At system level, a BMS supervises power distribution and cooling equipment, continuously tracking electrical loads, system states, air temperature, relative humidity, and dew point to ensure safe operating conditions and prevent thermal stress or condensation risks.

Complementing BMS data, a distributed network of environmental and air quality sensors is deployed throughout the data hall and adjacent spaces, measuring temperature, humidity, differential pressure, particulate matter (PM1–PM10), and gaseous pollutants (CO_2, tVOC, O_3, NH_3, CO, NO_2) (University of Turin, 2024a). This additional layer increases spatial resolution and supports airflow analysis, equipment health assessment, and indoor environmental quality evaluation. All measurements are streamed to a centralized time-series database, which stores time-stamped data along with sensor identifiers and metadata, serving as the authoritative source for real-time access and historical analysis.

5.3 BIM Model and Digital Twin Deployment

The BIM model was developed from as-built documentation and refined to support operational DT functionalities at a moderate LoD, balancing geometric fidelity and model manageability. Custom parametric BIM families were developed to represent the different categories of sensors deployed within the facility, including rack-level environmental sensors, air quality sensing stations located across the cold aisle (Fig.3), hot aisle, control room, and adjacent corridor, an on-site weather station providing outdoor boundary conditions, and roof-mounted adiabatic chillers monitored as components of the thermal management system. Each family is enriched with shared parameters storing unique sensor identifiers, live measurement values, timestamps, and diagnostic flags, acting as data endpoints for continuous injection of operational information.

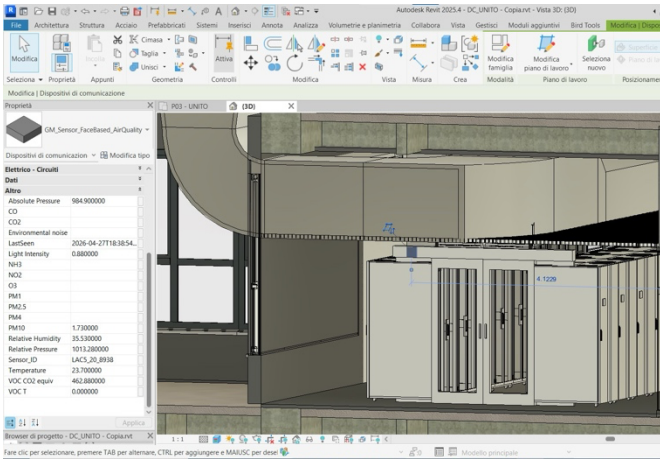


Figure 3. Property panels of a sensor family.

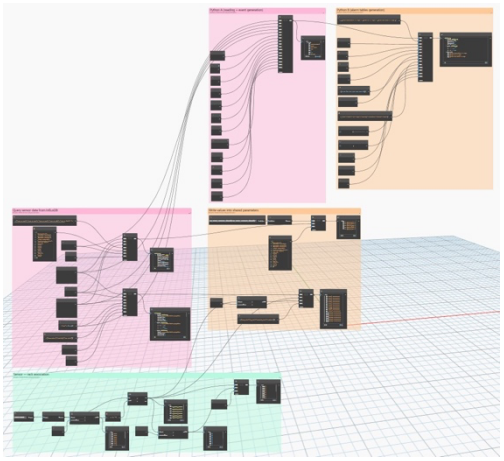


Figure 4. Dynamo workflow overview.

5.4 Dynamo-Based Integration and Operational Logic Configuration

The synchronization workflow described in Section 3 was operationalized through a modular Dynamo graph organized into functional stages reflecting the layered architecture of the DT. As shown in Fig. 4, the graph is articulated into visually distinct and logically independent blocks covering sensor-to-BIM mapping, data acquisition and preprocessing, validation and alert aggregation, and parameter persistence within the BIM environment. This separation of responsibilities improves maintainability and transparency, supporting incremental extension of monitored variables, substitution of the time-series backend without affecting mapping logic, and integration of additional analytical components.

In the **first phase** (element identification and mapping), Revit elements are filtered by category and family type to isolate sensor instances. A shared parameter (SensorId) embedded in each family serves as the primary key linking BIM elements to their counterparts in InfluxDB. Two aligned lists, BIM elements and associated identifiers, are extracted to guarantee deterministic writes during parameter

updates: this alignment ensures that each measurement is written back to the correct BIM instance without ambiguity, even when sensors are added, relocated, or replaced.

In the **second phase** (data acquisition and preprocessing), embedded Python scripts construct authenticated HTTP requests to InfluxDB using SensorId as tag key. The query dynamically integrates the sensor identifier, relevant measurement fields (temperature, relative humidity, dew point), a configurable temporal horizon (e.g., rolling 24-hour window), and a predefined aggregation interval (e.g., 5-minute resolution). Retrieved JSON data are parsed into temporally aligned datasets; null values and inconsistent entries are filtered prior to synchronization, and consistency between tag keys in Dynamo and the database schema is explicitly enforced to prevent parameter mismatches.

In the **third phase** (validation and alert evaluation), validated measurements undergo rule-based evaluation against ASHRAE guidelines for IT environments (ASHRAE, 2021). The alert taxonomy is summarized in Table 1. Individual detections are aggregated into temporally coherent episodes to mitigate transient anomalies and false positives, with retention policies pruning obsolete records outside the defined temporal horizon.

In the **fourth phase** (parameter persistence and visualization triggering), processed measurements and alert states are written to shared BIM parameters via a dedicated Python routine iterating over the aligned element–identifier lists. Parameter updates automatically trigger native BIM visualization mechanisms: view filters, graphical overrides (Fig. 4), and alert episode schedules (Fig. 5).

Table 1. Alert taxonomy and triggering rules implemented in the BIM–IoT workflow (ASHRAE basis).

Condition	Threshold Rule	Severity	Window	ASHRAE Basis
High Temperature	T ≥ 27°C (≥30°C instant)	CRIT	5’	Class A1 recommended / allowable limit
Rapid Temp. Variation	ΔT ≥ 3°C	WARN	Instant	Hot-aisle containment best practice
Statistical Deviation	T ≥ +4°C vs rack median	WARN	5’	Historical data distribution
Humidity Out of Range	RH < 20% or > 80%	WARN→CRIT	10’–30’	ASHRAE ESD / condensation limits
High Dew Point	DP > 16°C	CRIT	5’	ASHRAE condensation risk threshold
Sensor Inactivity	No data received	CRIT	15’	Mission-critical monitoring continuity

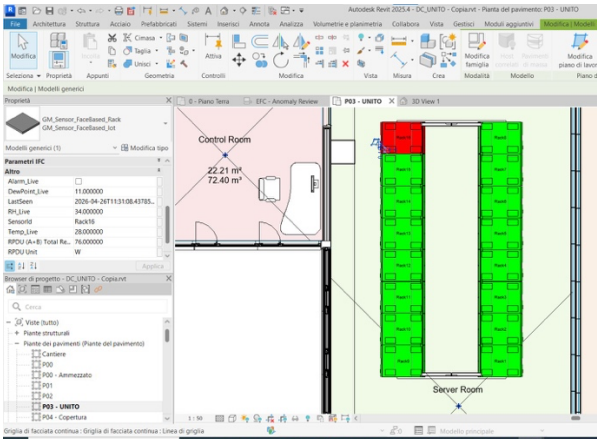


Figure 5. Rack-level BIM visualization driven by propagated sensor measurements.

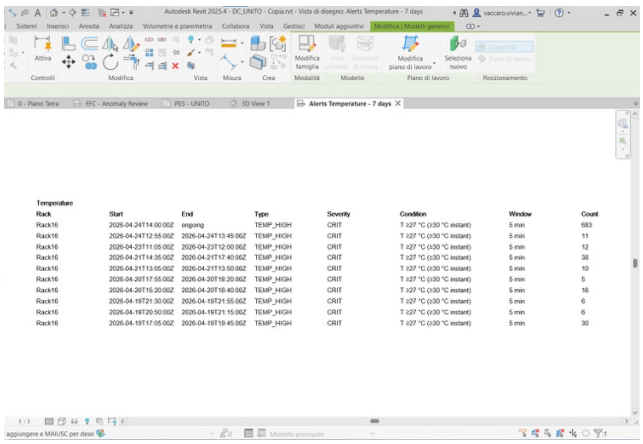


Figure 6. Episode-based alert reporting

5.5 Model Predictive Control Integration

Building upon the synchronization architecture, a closed-loop MPC module has been implemented as an additional functional group within the Dynamo workflow, entirely developed through embedded Python scripts executed within the Dynamo visual programming environment. The module adopts a data-driven modelling approach: the thermal RC model parameters are identified from real operational data streamed from the facility's Building Management System (BMS) and stored in InfluxDB, rather than relying on design-stage assumptions or offline calibration alone. This ensures that the model captures the actual thermal behaviour of the HPC4AI cold aisle as it evolves under real IT loads and climatic boundary conditions. The MPC engine operates on a first-order thermal Resistance-Capacitance (RC) model of the HPC4AI cold aisle:

$$C_{th} \times dT/dt = k_{leak} \times P_{IT} + Q_{roof} - Q_{EFC} \quad \text{(Eq. 1)}$$

where C_th = 421 Wh/°C is the thermal capacitance of the server room (identified from the physical geometry — floor area 146 m^2, height 2.91 m, prefabricated reinforced concrete slab — and validated on historical BMS data), k_leak = 0.0993 is the hot-aisle containment leakage coefficient (calibrated on 571 real BMS samples streamed from InfluxDB over four weeks of continuous operation; cross-validated against Vertiv EFC technical specifications, coinciding with the physical estimate of 0.10), Q_roof is the roof slab thermal load (U = 0.50 W/m^2K, A = 146 m^2, consistent with a prefabricated concrete slab with minimal insulation layer as per as-built documentation), and Q_EFC is the cooling power delivered by the adiabatic chillers, read in real time from the BMS telemetry.

The MPC Dynamo group is structured into three Python script nodes executed in sequence at each 15-minute cycle. The first node reads current sensor values from InfluxDB (cold aisle and hot aisle temperature, IT load from APM01+APM02, EFC state and capacity, outdoor boundary conditions, and the dynamic EFC setpoint — 24°C at the time of validation, read in real time at each cycle). The second node solves a Sequential Least Squares Programming (SLSQP) optimization over a 60-minute prediction horizon (N = 12 steps, dt = 5 min), minimizing a cost function that combines thermal deviation from the dynamic setpoint and cooling energy consumption, subject to ASHRAE A1 envelope constraints (18°C ≤ T ≤ 27°C) and an asymmetric penalty that penalizes only upward threshold violations. The third node writes the resulting optimal cooling setpoint, predicted temperature trajectory, and energy KPIs (partial PUE, COP, EFC utilization rate) to predefined shared BIM parameters of the corresponding Revit elements — EFC units and aisle sensor families. The entire workflow is schedulable automatically via Dynamo Multiplayer (the server-based Dynamo execution environment) combined with Windows Task Scheduler, enabling unattended periodic execution. Validation metrics are reported in Table 2.

Table 2. MPC validation metrics on real InfluxDB data (sensor LAC5_20_8938, cold aisle).

Metric	Value	Dataset	Notes
MAE	0.39°C	34 samples, 8 h window	+60 min prediction horizon
RMSE	0.65°C	34 samples, 8 h window	+60 min prediction horizon
MAPE	1.73%	34 samples, 8 h window	+60 min prediction horizon
Annual validation	6,001 samples	250 days historical data, hourly	Full-year dataset
Calibration	571 samples	4 weeks real InfluxDB data	Cross-validated vs. physical estimate (k=0.10)

6 Conclusions and Future Work

The proposed BIM-based DT framework for the HPC4AI data center demonstrates the practical feasibility of integrating real-time operational monitoring and predictive control into a semantically structured digital environment. The modular layered architecture, linking sensing infrastructure, time-series data management, and BIM representation without tight coupling, enables the model to evolve from a static design artifact into an active operational interface supporting monitoring, diagnostics, and structured alert management. The closed-loop MPC engine extends the DT toward active, energy-aware decision support. By advancing energy-aware operational control within a native BIM environment, the proposed framework contributes to the broader agenda of sustainable and resilient built environments promoted by the BDTA ecosystem.

Building upon previous dynamic modelling activities (Vaccaro et al., 2025), the real-time data backbone established by this framework provides the necessary foundation for future coupling between simulated and measured performance. In this perspective, dynamic simulation and operational monitoring can interact bidirectionally: real-time measurements support continuous calibration and validation of the simulation model, while predictive simulations inform adaptive control strategies and energy optimization scenarios.

Future work will focus on three directions: (1) integration of dynamic simulation outputs with monitored data for systematic model calibration and performance gap analysis; (2) incorporation of machine learning for workload-aware cooling optimization and anomaly detection; and (3) extension toward sustainability-oriented indicators for dynamic evaluation of energy efficiency and environmental impact. Validation across additional deployment contexts — including the Italian MaCA (Museo A come Ambiente) museum Living Lab in Turin, where sensing infrastructure, data management, BIM preparation, and automated synchronization are being introduced in a controlled sequence (Associazione A come Ambiente, 2024) — will confirm generalizability beyond the data center domain.

Acknowledgements
This work has been supported by the DYMAN project funded by the European Union - European Innovation Council under G.A. n. 101161930.

Funding
This research was funded by the European Union — European Innovation Council, DYMAN project, Grant Agreement n. 101161930.

Data Availability Statement
Sensor data used in this study are partially accessible through the HPC4AI Building Management System portal (University of Turin, 2024a) and the Centraline HPC4AI system (University of Turin, 2024b).

Conflicts of Interest
The authors declare no conflict of interest.

References

ASHRAE. (2021). Thermal Guidelines for Data Processing Environments (5th ed.). Atlanta, GA: ASHRAE.

Associazione A come Ambiente. (2024). Relazione e Bilancio al 31 dicembre 2024 (Progetto di Bilancio). Associazione A come Ambiente, Turin, Italy.

Harode, A., Thabet, W., & Dongre, P. (2023). A tool-based system architecture for a digital twin: A case study in a healthcare facility. Journal of Information Technology in Construction (ITcon), 28, 107–137. doi: 10.36680/j.itcon.2023.006

International Energy Agency (IEA). (2023). Electricity 2024: Analysis and Forecast to 2026. Paris, France: IEA. Retrieved from https://www.iea.org/reports/electricity-2024, Last Access: May 2026.

International Organization for Standardization (ISO). (2018). ISO 19650-1: Organization and digitization of information about buildings and civil engineering works, including building information modelling (BIM) — Information management using building information modelling — Part 1: Concepts and principles. ISO.

Jafari, M., Kavousi-Fard, A., Chen, T., & Karimi, M. (2023). A review on digital twin technology in smart grid, transportation system and smart city: Challenges and future. IEEE Access, 11, 17471–17484. doi: 10.1109/ACCESS.2023.3241588

Muñoz Pavón, R., García Alberti, M., Arcos Álvarez, A. A., & Jerez Cepa, J. (2025). BIM-based Digital Twin development for university campus management. Case study ETSICCP. Expert Systems with Applications, 262, 125696. doi: 10.1016/j.eswa.2024.125696

Pasini, D., Mastrolembo Ventura, S., Rinaldi, S., Bellagente, P., & Ciribini, A. L. C. (2016). Exploiting Internet of Things and Building Information Modeling framework for management of cognitive buildings. Paper presented at the IEEE Second International Smart Cities Conference (ISC2), Trento, Italy. doi: 10.1109/ISC2.2016.7580861

Pinmas, P., Sooksatra, S., Wuttidittachotti, P., & Sutheebanjard, P. (2022). Towards digital twin data center using building information modeling and real-time data sensing. Paper presented at the 37th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC), Phuket, Thailand. doi: 10.1109/ITC-CSCC55581.2022.9894916

Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison. IEEE Access, 6, 3585–3593. doi: 10.1109/ACCESS.2018.2793265

Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital twin: Values, challenges and enablers from a modeling perspective. IEEE Access, 8, 21980–22012. doi: 10.1109/ACCESS.2020.2970143

University of Turin. (2024a). HPC4AI Building Management System (BMS). Retrieved from https://frontend.hpc4ai.unito.it:8080/BMS/, Last Access: May 2026.

University of Turin. (2024b). Centraline HPC4AI. Retrieved from https://centraline.di.unito.it/, Last Access: May 2026.

Vaccaro, V., Birke, R., Tagliabue, L. C., Rabellino, S., & Aldinucci, M. (2025). HPC4AI@UNITO: A use case for datacenter digital twin. In Proceedings of the 33rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP) (pp. 499–505). Torino, Italy: IEEE. doi: 10.1109/PDP63540.2025.00076

Vertiv. (2021). University of Torino: High-performance computing and sustainability with Vertiv Liebert® EFC [Case study]. Retrieved from https://prod.vertiv.cn/globalassets/documents/case-studies/vertiv-university-of-torino-cs-en-emea-web_357267_0.pdf, Last Access: May 2026.

Zhang, Y., Yu, Z., & Zeng, D. (2021). A digital twin approach for predictive maintenance of data center cooling systems. Energy and Buildings, 231, 110591. doi: 10.1016/j.enbuild.2021.110591

Disclaimer/Publisher's Note

The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and do not reflect the views of the Architecture, Buildings, Construction and Cities (ABC2) Journal and/or its editor(s). ABC2 Journal and/or its editor(s) disclaim any responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.