Knowledge graph-enabled rule-based fault detection and diagnosis for air handling units

Laura Block[a,b], Martin Rätz[a], Dirk Müller[a]

[a] *RWTH Aachen University, E.ON Energy Research Center,*
Institute for Energy Efficient Buildings and Indoor Climate, Germany
[b] *laura.block@eonerc.rwth-aachen.de, CA*

Abstract:

The high prevalence of faults in air handling units (AHUs) significantly compromises building operational performance, increases costs, and reduces occupant comfort. Consequently, reliable fault detection and diagnosis (FDD) methods are needed. While traditional rule-based FDD methods are simple and interpretable, their scalability in real-world applications is limited by the need for manual configuration and the difficulty of modeling complex systems. Knowledge graphs (KGs) enable the automated configuration of FDD methods by representing the physical and logical relationships between components. Therefore, this paper presents a method for automated configuration of a rule-based FDD based on a KG. To achieve this, a KG representing the components and relationships of a real-world AHU is used. The AHU topology is automatically extracted by querying the KG to serve as the foundation for the FDD system. Faults are detected and diagnosed through a hierarchical sequence of rule-based checks applicable to any water-to-air heat exchanger in the AHU. The methodology is validated using real-world operational data from an AHU preheater. The results demonstrate that the KG-enabled FDD system reliably detects and diagnoses faults. Furthermore, KG integration enables the automated reuse of rules for recurring patterns and streamlines time-series data mapping, reducing manual configuration effort. Overall, KG integration is a promising approach to advance the automated design and generation of FDD.

Keywords:

Air Handling Units; Fault Detection and Diagnosis; Knowledge Graphs; Semantic Web Technologies.

1. Introduction

Building operational energy consumption accounts for approximately 30 % of global energy use [1]. Air handling units (AHUs) are common components of heating, ventilation, and air conditioning systems in buildings. Approximately 40 % of AHUs operate with faults on any given day [2]. Faults have a substantial impact on energy consumption and occupant thermal comfort [3]. Consequently, implementing fault detection and diagnosis (FDD) and the subsequent maintenance actions leads to energy savings. FDD for building energy systems (BES) is still in an early stage [4]. BES and their control are highly complex and heterogeneous [5, 6]. Therefore, the installation and configuration of monitoring systems, such as FDD, by leveraging building automation systems (BAS) are often labor-intensive and time-consuming [6]. Developing FDD rules requires a comprehensive understanding of the building, including equipment characteristics, control logic, and linked datapoints [6, 7]. Due to a lack of semantic information, domain knowledge must be manually collected and interpreted [8]. To address this, Pritoni et al. [6] propose a novel approach to FDD deployment that generates rules and visualizations directly from standardized data models. The application of standardized ontologies via a semantic knowledge graph (KG) can significantly reduce implementation costs and facilitate broader FDD adoption [6]. This is because KGs provide a machine-readable representation of the BES, its metadata, and underlying relationships. Bampoulas et al. [5] highlight a trend towards automated FDD processes, emphasizing ontologies as a key instrument for structuring knowledge in building energy management systems. The following section provides a brief overview of recent advancements of ontologies for the BES domain and FDD applications.

In the building domain, the Building Information Modeling (BIM) standard is designed to represent a building's geometry, system, and component data, but it lacks the ability to incorporate control logic and is limited in integrating operational data [9]. To enable digital transformation of building automation, researchers propose applying W3C-standardized Semantic Web technologies [10, 11]. Initiatives such as *Project Haystack* [12] and the *Brick Schema* [13] have emerged to address this problem. Both schemas offer a unified structure

and vocabulary for building semantic metadata, facilitating the exchange of shared models among application providers [7]. Nevertheless, they lack formal standardization and were not designed to capture the detailed topology of equipment and building systems required for advanced analytics [6]. The *223p* ontology [14], which follows the *ASHRAE Standard 223P*, was developed to provide a consistent representation of building semantic metadata [6], standardizing the semantic modeling of interconnected fluid flows. Building upon these standardized semantic frameworks, recent research increasingly adopts KGs to represent complex BES in FDD applications.

Several studies extend existing semantic schemas. Hwang et al. [15] developed *FSBrick*, an extension of the *Brick schema* designed to model faults, fault severities, and fault-symptom relationships in HVAC systems, thereby minimizing manual adjustments upon system reconfiguration. Similarly, Li et al. [16] introduced an ontology specifically tailored to fault detection in BES that integrates expert-derived semantic rules to detect operational, control, and equipment faults, successfully validating the approach in an industrial building. Blechmann et al. [17] proposed a *Brick*-based ontology combined with generic semantic rules derived from air handling performance assessment rules [18]. Demonstrated on real-world data, the approach achieved a 99.3% accuracy rate in identifying various injected faults. To achieve scalable automation, Pruvost et al. [19] implemented a semantic expert system utilizing KGs and logical reasoning to autonomously configure a monitoring system based on building metadata. Their approach correlates usage patterns and device performance data to generate tailored recommendations for facility managers. Implementation across multiple European buildings resulted in an energy consumption reduction of over 12 %. Pruvost and Wilde [20] utilized ontologies to translate building metadata into self-configuring algorithms capable of identifying faults and energy waste. Their prototype showed that the manual effort required to maintain a complex BAS system is reduced. Apart from rule-based systems, Hosseini Gourabpasi and Nik-Bakht [21] integrated sensor data streams into BIM via the *AFDDOnto* ontology, employing artificial neural networks and support vector machines for fault detection. Yin [22] combined graph neural networks and Bayesian inference to localize faults based on structural and sensor deviations, achieving a diagnostic accuracy of 93.1 % trained on expert-validated logs and historic maintenance records.

Existing KG-based FDD frameworks predominantly rely on ontologies that map structural components and sensors. For in-depth fault diagnosis and root cause analysis in BES, integrating control logic and directional hydraulic connections is crucial. Therefore, this paper employs multiple ontologies to integrate control logic and directional fluid flows into the KG, utilizing a hierarchical structure for fault diagnosis. Furthermore, validation based on real-world operational data remains scarce. Thus, the developed approach is evaluated on a real-world operational dataset from a test hall's air handling unit.

2. Methodology

In the following, the use case, the KG, and the methodology for the KG-enabled FDD are presented.

2.1. Use case

To evaluate the developed method, real-world data from a central air handling unit (AHU) is used. The AHU is located in the test hall of E.ON Energy Research Center in Aachen, Germany. The test hall and the AHU are shown in Fig. 1. The AHU supplies jet nozzles that feed one part of the test hall ($630\,\mathrm{m}^2$) and ceiling induction diffusers that feed an office area ($140\,\mathrm{m}^2$). While the test hall is also equipped with concrete core activation, this work focuses solely on the central AHU. A detailed overview of the building and its energy system is provided by Schraven et al. [23].

(a) *Test hall of the E.ON Energy Research Center in Aachen, Germany*

(b) *Central air handling unit in the test hall*

Figure 1. *Overview of the use case.*

Figure 2 provides a schematic representation of the central AHU, which consists of four dampers, two fans, a heat recovery unit, a steam humidifier, two heating coils, and one cooling coil. This work focuses on the heating and cooling coils, as they are connected to their respective sources via identical mixing circuits. The setup is shown schematically in Fig. 3 and the locations of the temperature sensors are indicated. To maintain concise nomenclature, variables are denoted as follows: T for temperature, $\dot{V}$ for volume flow, n for speed, M for operating mode, and u for mixing valve position. Subscripts specify the medium (Air or W for water), location (Sup for supply, Ret for return, Prim for primary), and control setpoints (Set). Additionally, the subscript $C \in \{\mathrm{PH}, \mathrm{CO}, \mathrm{RH}\}$ designates the specific component: preheater (PH), cooler (CO), or reheater (RH). One exception is the setpoint of the pump ($P_{Set,C}$) as the modulated variable can change with the pump's operating mode. One of the heating coils serves as a PH, which supports the RH when outdoor temperatures drop below 8 °C. The PH pump operates at a constant speed, while the water-side supply temperature ($T_{\mathrm{W,Sup},PH}$) is controlled by a PID controller that modulates the mixing valve position (u_C). An unusual control strategy is applied to the second heating coil (RH) and the cooling coil (CO). The valve opening regulates the water-side supply temperature ($T_{\mathrm{W,Sup,C}}$), while the pump speed controls the air-side return temperature ($T_{\mathrm{Air,Ret},C}$). This strategy is designed to minimize pump energy consumption, as detailed by Teichmann et al. [24]. The supervisory control strategy is implemented on a transparent and customizable programmable logic controller (PLC). Some control signals are transmitted directly to the components via BACnet, while others are routed through a proprietary manufacturer's PLC. Because the manufacturer's PLC can overwrite these signals and its internal logic remains inaccessible to the authors, certain assumptions regarding the underlying control logic cannot be verified. Timeseries data, including sensor measurements, actuator states, and AHU operating modes, are stored in a database and can be accessed via API. In this study, the developed methodology is evaluated using data from the preheater.

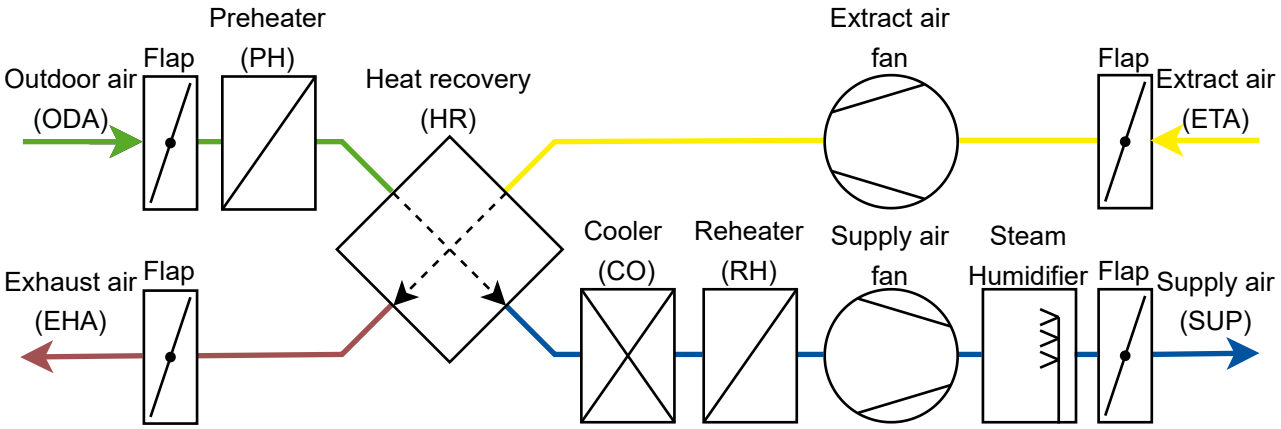

Figure 2. *Schematic representation of the central air handling unit adapted from Schraven et al. [23].*

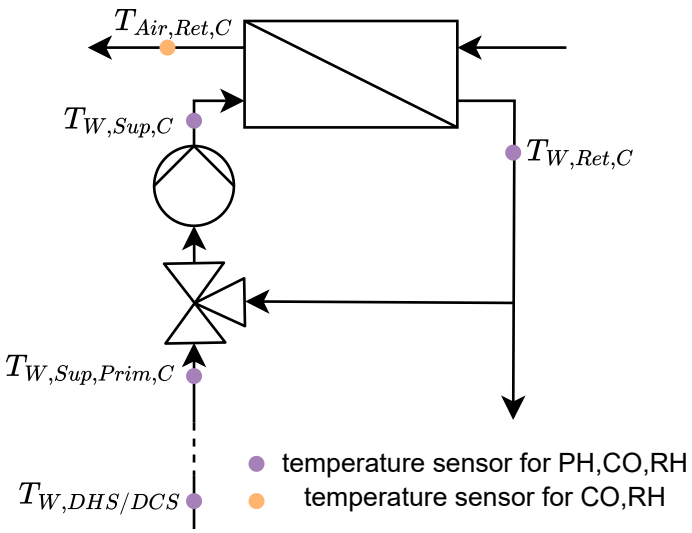

Figure 3. *Schematic representation of the investigated heating and cooling coils and their hydraulic mixing circuit.*

2.2. Knowledge graph

To represent the AHU topology, equipment, and data points, we utilize a comprehensive, manually constructed KG based on multiple ontologies. The *Brick Schema* [13] is employed to model equipment, its operating mode, corresponding sensors, and timeseries references. Since *Brick Schema* does not natively support directional fluid flow, *223p* [14] is used to map fluid connections (e.g., *s223:hasConnectionPoint*). The use of *223p* allows for the distinction between primary and secondary hydraulic circuits, which is essential for detecting specific faults. Moreover, controllers are represented using the *VDI 3814* ontology [25]. The *VDI 3814* ontology is currently developed at the E.ON Energy Research Center as a digital representation of the VDI 3814 Sheet 3.1 [26]. The ontology enables the representation of building related functions in accordance with the VDI standard. The semantic dependencies between an observed measurement, the controller, and the consequential actuator commands are modeled. Following its creation, the KG is enriched by inference based on 223p to derive implicit from explicit knowledge. The usage of the different ontologies for modeling the KG is shown in Fig. 4 exemplary for fragments from the preheater.

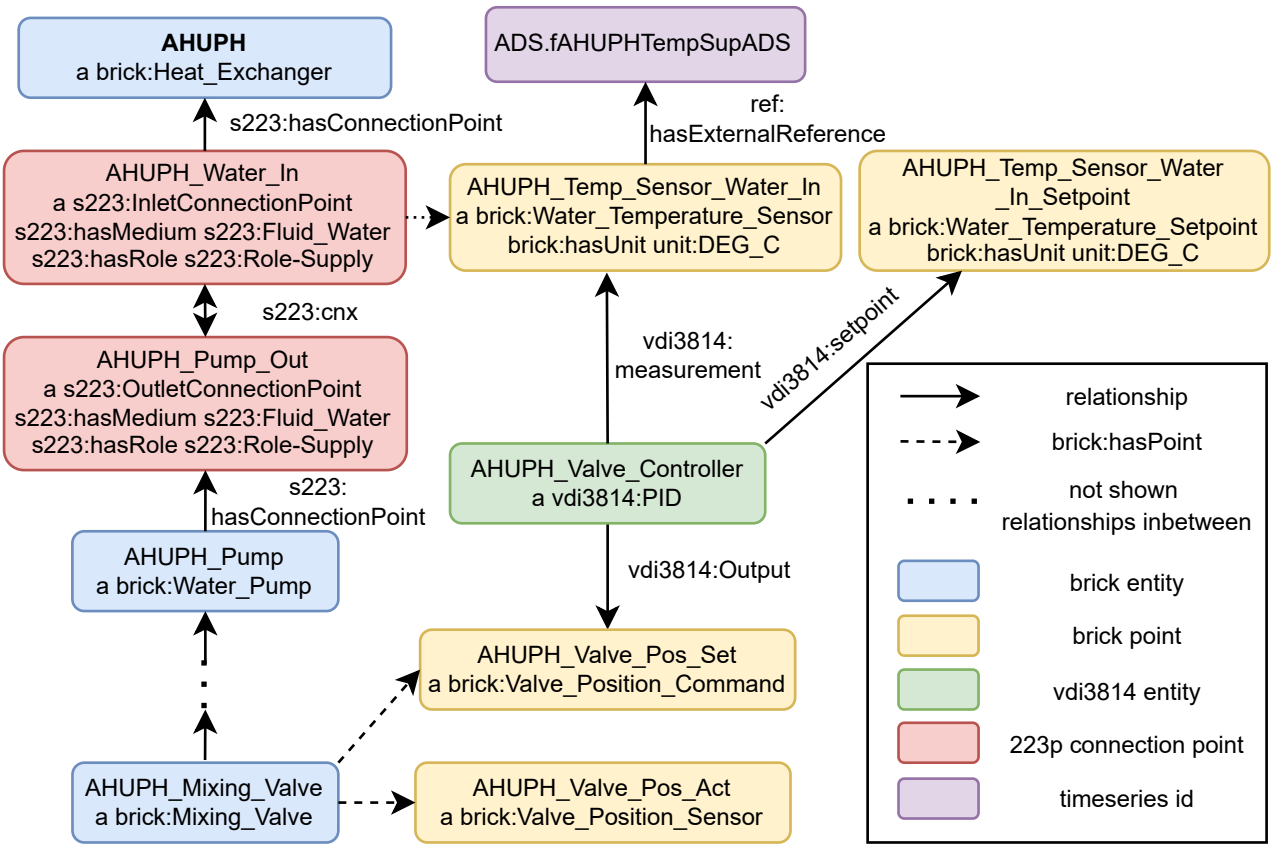

Figure 4. *Excerpt of the knowledge graph of the preheater showing the usage of and the relations between the different ontologies.*

2.3. Fault detection and diagnosis

This section describes the proposed FDD methodology, which consists of two steps: dynamic topology discovery and an automated FDD Loop. The goal of the dynamic topology discovery is to locate mixing circuit patterns that match the automated FDD setup. The automated FDD Loop detects faults and isolates the most initial fault for a time step. Both steps are implemented in Python.

During the dynamic topology discovery, the KG is queried via SPARQL queries to extract the physical and logical structure of the AHU. The discovery process originates at the root node (e.g., *brick:AHU*). The *brick:hasPart* predicate is used to discover connected heating and cooling coils classified as *brick:Heat_Exchanger*. Because the heat recovery (HR) system differs structurally from the heating and cooling coils, it is filtered out of the detection loops during the SPARQL request. After the identification of the heat exchangers, the corresponding hydronic actuators (*brick:Water_Pump*) and flow modulators (*brick:Mixing_Valve*) are located by cross-referencing queries. By exploiting the *brick:hasPoint* relationship, sensor identifiers like *brick:Pump_Speed_Sensor*, *brick:Temperature_Sensor*, and *brick:Valve_Position_Sensor* are localized. Through the *223p* connectivity nodes (*s223:hasConnectionPoint*), specific temperature sensors (*brick:Temperature_Sensor*) are mapped to their respective locations (e.g., supply/return) and topological boundaries. As a final step of topology discovery,

the controllers (*vdi3814:PID*) for the pumps and valves are queried to bind the controller outputs with actuation commands (*brick:Valve_Position_Command*). Regulated data points are discovered originating from the controller to map the complete control loop. Finally, the *ref:hasExternalReference* relationship from *Brick* is used to query and map specific timeseries IDs for each relevant node. The dynamically collected variables and the corresponding timeseries IDs serve as the foundation for the subsequent FDD process.

The automated FDD loop begins by pulling timeseries data for all discovered nodes from the database for a specified period. If a heat exchanger is discovered in the topology, a hierarchical sequence of rule based checks for faults is applied, as illustrated in Fig. 5. Each block in the figure represents a specific fault type. The arrows indicate a cascading logic. A subsequent fault type is only evaluated if the preceding fault in the hierarchy has been detected for that specific time step. Moreover, a fault check is only performed if the corresponding node and its linked timeseries data were successfully identified during the topology discovery phase. For fault diagnosis, the systematic fault checks originate at the symptom level (e.g., incorrect regulated air temperature) and progress toward the identification of root causes. The corresponding rules for these fault types are specified in Tab. 1. While the hierarchical order and rules presented here are derived from use-case-specific control logic and hydraulic configurations, they are designed for interoperability across different topologies. This is achieved through the dynamic mapping and a consistent hierarchical structure for every discovered heat exchanger. The final output of the automated FDD loop is a set of all prevalent fault types identified for each time step.

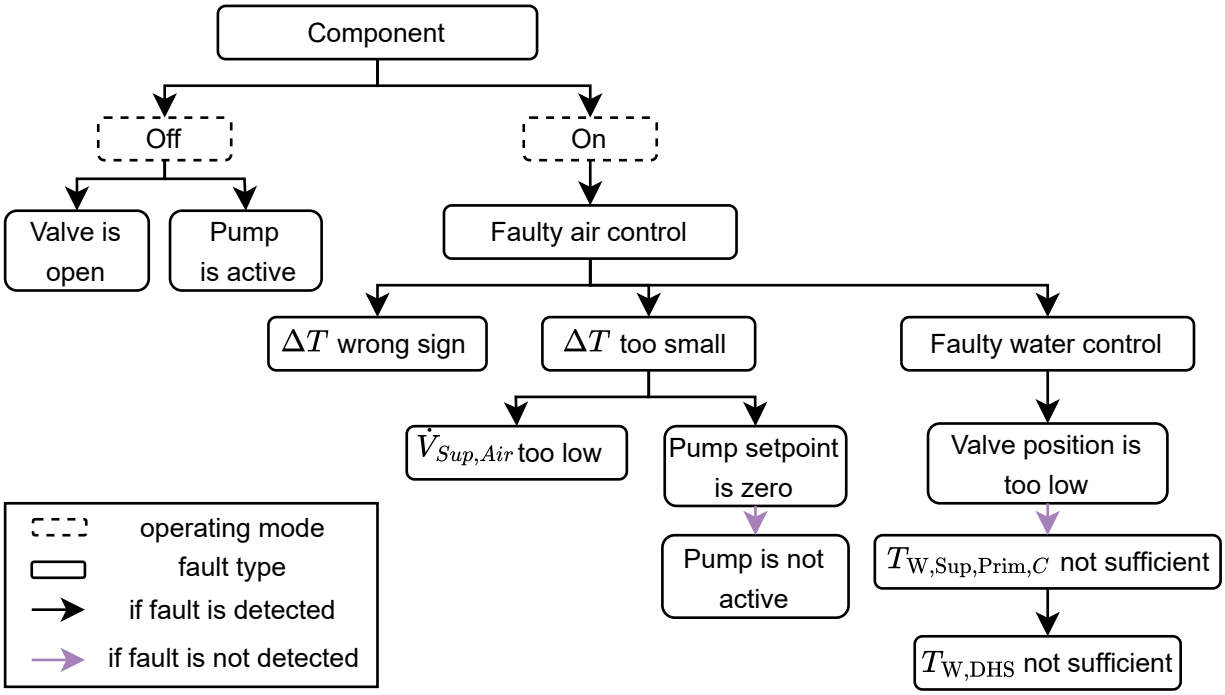

Figure 5. *Hierarchical order of the fault types for the components preheater, cooler, and reheater. The rules are derived from the control logic and the hydraulic circuit.*

Table 1. *Detection rules for faults types.*

Fault type	Rules
Faulty air control	$T_{\mathrm{Air,Ret,Set},C} - 1\ \mathrm{K} < T_{\mathrm{Air,Ret},C} < T_{\mathrm{Air,Ret,Set},C} + 1\ \mathrm{K}$
ΔT too small	$\mid T_{\mathrm{W,Sup},C} - T_{\mathrm{W,Ret},C} \mid < 0.5\ \mathrm{K}$
ΔT wrong sign	$T_{\mathrm{W,Ret},CO} - T_{\mathrm{W,Sup},CO} < 0, T_{\mathrm{W,Sup},C} - T_{\mathrm{W,Ret},C} < 0\ \forall C \in [\mathrm{PH}, \mathrm{RH}]$
Faulty water control	$T_{\mathrm{W,Sup,Set},C} - 1\ \mathrm{K} < T_{\mathrm{W,Sup},C} < T_{\mathrm{W,Sup,Set},C} + 1\ \mathrm{K}$
$\dot{V}_{Sup,Air}$ too low	$\dot{V}_{Sup,Air} < 0.1$
Pump setpoint is zero	$P_{\mathrm{Set},C} = 0\ \%$
Valve position is too low	$u_{\mathrm{Set},C} < 0.1\ \% \vee\ u_C < 0.1\ \%$
Pump is not active	$n_{\mathrm{Pump},C} < 0.1 \wedge \dot{V}_{\mathrm{Pump},C} < 0.1$
$T_{\mathrm{W,Sup,Prim},C}$ not sufficient	$T_{\mathrm{W,Sup,Set,CO}} > T_{\mathrm{W,Sup,Prim,CO}}, T_{\mathrm{W,Sup,Set},C} < T_{\mathrm{W,Sup,Prim},C}\ \forall C \in [\mathrm{PH}, \mathrm{RH}]$
$T_{\mathrm{W,DHS}}$ not sufficient	$T_{\mathrm{W,Sup,Set,CO}} > T_{\mathrm{W,DCS}}, T_{\mathrm{W,Sup,Set},C} < T_{\mathrm{W,DHS}}\ \forall C \in [\mathrm{PH}, \mathrm{RH}]$
Valve is open	$u_{\mathrm{Set},C} > 0.1\ \% \vee\ u_C > 0.1\ \%$
Pump is active	$P_{\mathrm{Set},C} > 0.1\% \vee\ \ n_{\mathrm{Pump},C} > 0.1 \vee \dot{V}_{\mathrm{Pump},C} > 0.1 \vee M_{\mathrm{Pump},C} \neq 0$

3. Results and discussion

This chapter presents the results of the FDD for the preheater of the AHU as a exemplary case study. Thereafter, the advantages of the KG integration are discussed.

3.1. Fault detection and diagnosis results for the preheater

The performance of the FDD is evaluated for the preheater of the AHU as a proof of concept. Figure 6 shows a heatmap of all fault types detected for the preheater from the 01.10.2025 to the 31.12.2025 and the operating mode of preheater. Multiple faults are detected during this period. The fault *valve is open* is frequently detected during off mode periods of the preheater. In the end of October 2025, three fault types occur simultaneously. Two of these, *Faulty water control* and *Valve position is too low* are hierarchically linked (cf. Fig. 5) and are therefore analyzed together for this period. Similarly, the fault types ΔT *too small* and ΔT *wrong sign* are also evaluated together, as they typically occur concurrently. The fault type *Faulty water control* is nearly always detected while the preheater is in heating mode. These fault types are analyzed below for selected time intervals in detail to verify the reliability of the fault detection and diagnosis. Due to space restrictions the fault type *insufficient* $T_{\mathrm{W,Sup,Prim,C}}$ remains unaddressed in this analysis as it is the least insightful fault type in this period.

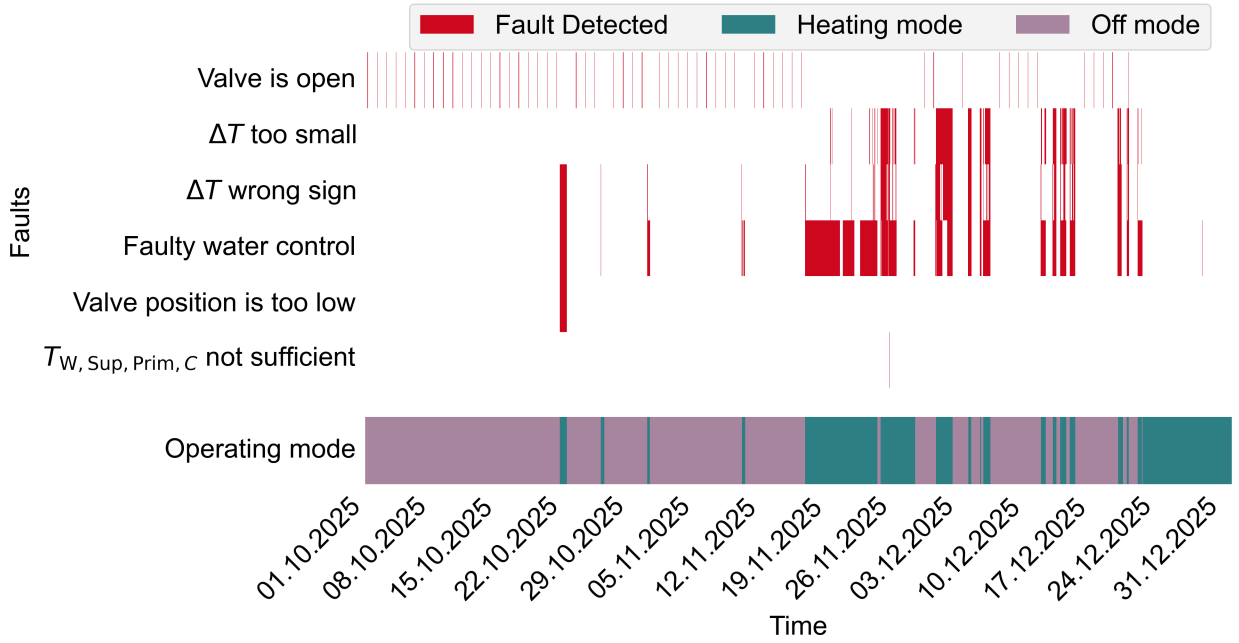

Figure 6. *Heatmap of all fault types of the preheater occurred between the 01.10.2025 and the 19.12.2025 and operating mode of preheater. To facilitate visualization, data is aggregated into 1-hour bins.*

For the hierarchically connected fault types *Faulty water control* and *Valve position is too low* the water-side supply temperature ($T_{\mathrm{W,Sup,PH}}$), the corresponding setpoint ($T_{\mathrm{W,Sup,Set,PH}}$), the valve position ($u_{\mathrm{PH}}$), the corresponding setpoint ($u_{\mathrm{Set,PH}}$), the preheater pump speed ($n_{\mathrm{Pump,PH}}$), the supply fan speed ($n_{\mathrm{Sup,Fan}}$) and the operating mode are shown in Fig. 7. The water-side supply temperature fails to track its setpoint. If the preheater is in heating mode the water-side supply temperature should follow the setpoint. During the faulty period, indicated by the red background, this is not the case. Moreover, the actual valve position does not align with the command sent by the supervisory control and the supply fan speed remains at zero throughout the faulty period, indicating that the AHU is not in operation. On the other side, the preheater pump is active ($n_{\mathrm{Pump,PH}} > 0$), which implies that water is circulating through the mixing circuit while the fan is off. The supply temperature increases upon pump activation and oscillates above the setpoint for the duration of the fault. This abrupt temperature rise could indicate valve leakage during pump operation. Furthermore, the temperature remains high because there is no airflow through the AHU to dissipate the heat. While the supervisory control sets the pump setpoint directly via BACnet, valve position and supply fan setpoint are forwarded to the manufacturer's PLC and set there. Consequently, an alternative explanation is that the manufacturer's PLC may have overridden the setpoint and due to a communication issue during the faulty period this change is not logged. In summary, the fault is correctly detected, as the behavior clearly deviates from normal AHU operation. The FDD successfully identifies the initial fault type in the hierarchical order. However, the definitive root cause remains an assumption based on this analysis.

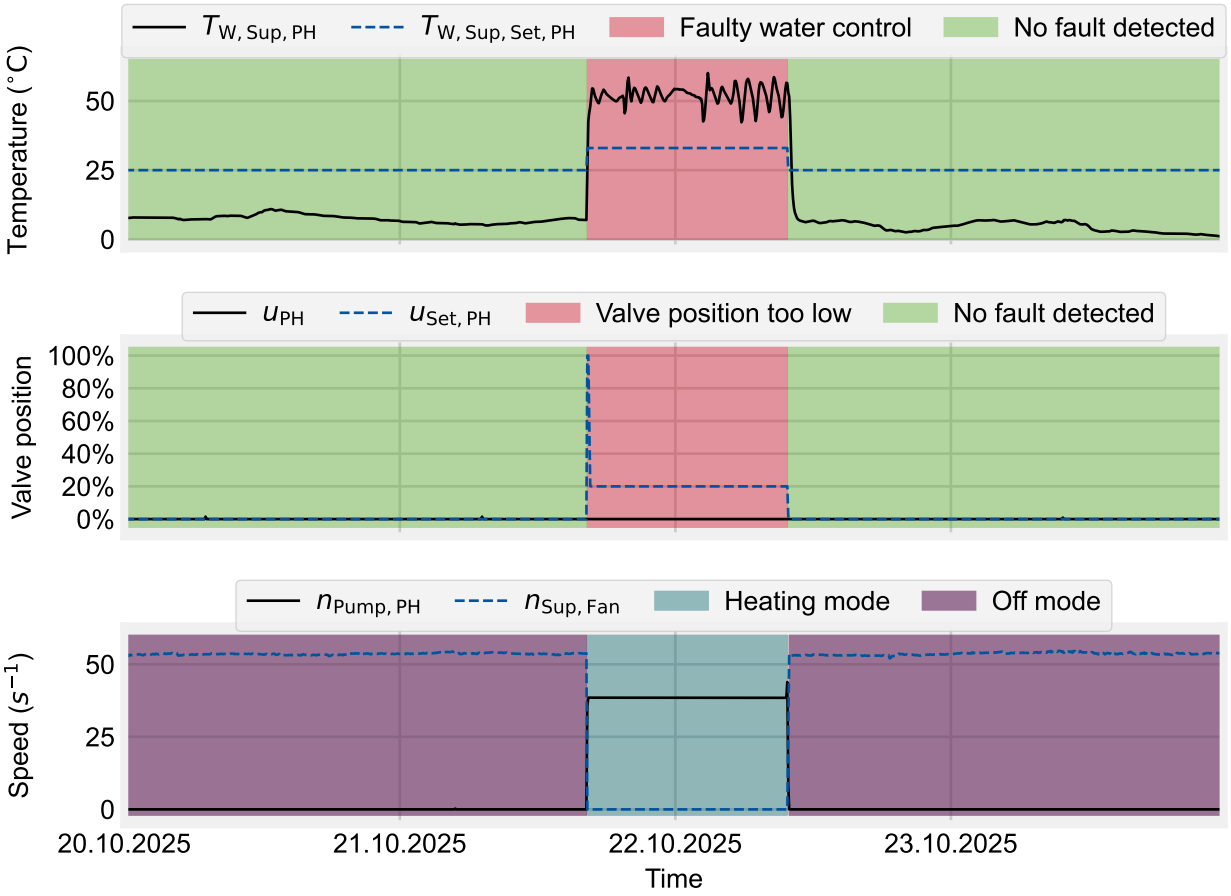

Figure 7. *Fault detection and diagnosis results for the connected fault types "Faulty water control" and "Valve position too low" for the time period from 20.10.2025 to 23.10.2025.*

The fault type *valve is open* is analyzed for an exemplary period. In Fig. 8 the actual valve position (u_{PH}), the valve setpoint ($u_{\mathrm{Set,PH}}$), the detected fault type and the operating mode are shown from the 08.10.2025 to the 10.10.2025. Throughout this period, the preheater remains in off mode. The actual valve position is mostly zero. Except for every day at 9:00 am, where the actual valve position is 10 %. According to the supervisory control logic, both the actual valve position and the setpoint should remain at zero. The setpoint for the valve, which is zero, is sent from the supervisory control to the manufacturer's PLC. The assumption is therefore that the manufacturer's PLC overwrites the setpoint every day at 9:00 am. Due to a lack of access to the internal logic of the manufacturer's PLC, this assumption cannot be verified. However, this periodic opening likely serves as a scheduled maintenance routine to ensure the actuator remains functional after extended periods of inactivity, such as during the summer months. An alternative cause could be voltage fluctuations at the PLC terminal triggered by other scheduled signals, though this is considered less probable. While the FDD correctly identifies this behavior as a fault based on the supervisory control rules, it most likely represents an intended maintenance action. To prevent such maintenance patterns from being incorrectly detected as faults, all control information and logic must be made fully accessible within the KG.

Figure 9 presents the analysis of the fault type *Faulty water control*. The supply temperature ($T_{\mathrm{W,Sup,PH}}$), the corresponding setpoint ($T_{\mathrm{W,Sup,Set,PH}}$), the valve position ($u_{\mathrm{PH}}$), the corresponding setpoint ($u_{\mathrm{Set,PH}}$) and the operating mode of the preheater are shown from the 16.11.2025 - 17.11.2025. During this interval, the operation mode switches from off to heating mode. The fault occurs during the heating mode. The supply temperature oscillates around the setpoint during the heating mode. When the fault is detected, the difference between the supply temperature and the setpoint exceeds 1 °C. While the valve position accurately tracks its setpoint, the setpoint itself oscillates. This indicates that the PID controller regulating the supply temperature is not optimally tuned. A conclusion supported by the fact that the fault persists throughout every heating period. As the PID parameters can be adjusted to resolve these oscillations, the fault is considered correctly detected and diagnosed, with the root cause attributed to the temperature control logic.

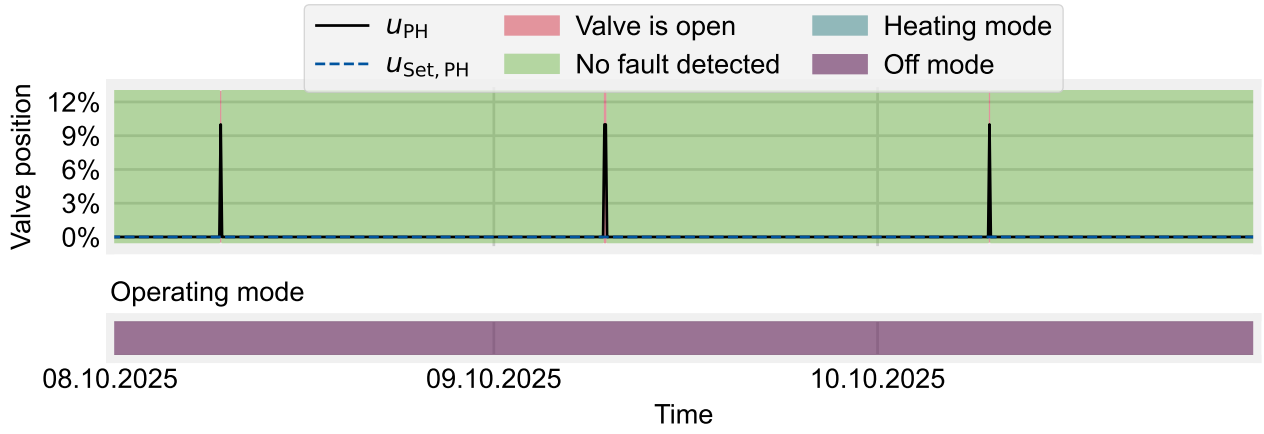

Figure 8. *Fault detection and diagnosis results and operation mode for the fault type "Valve open" for the time period from 08.10.2025 to 10.10.2025.*

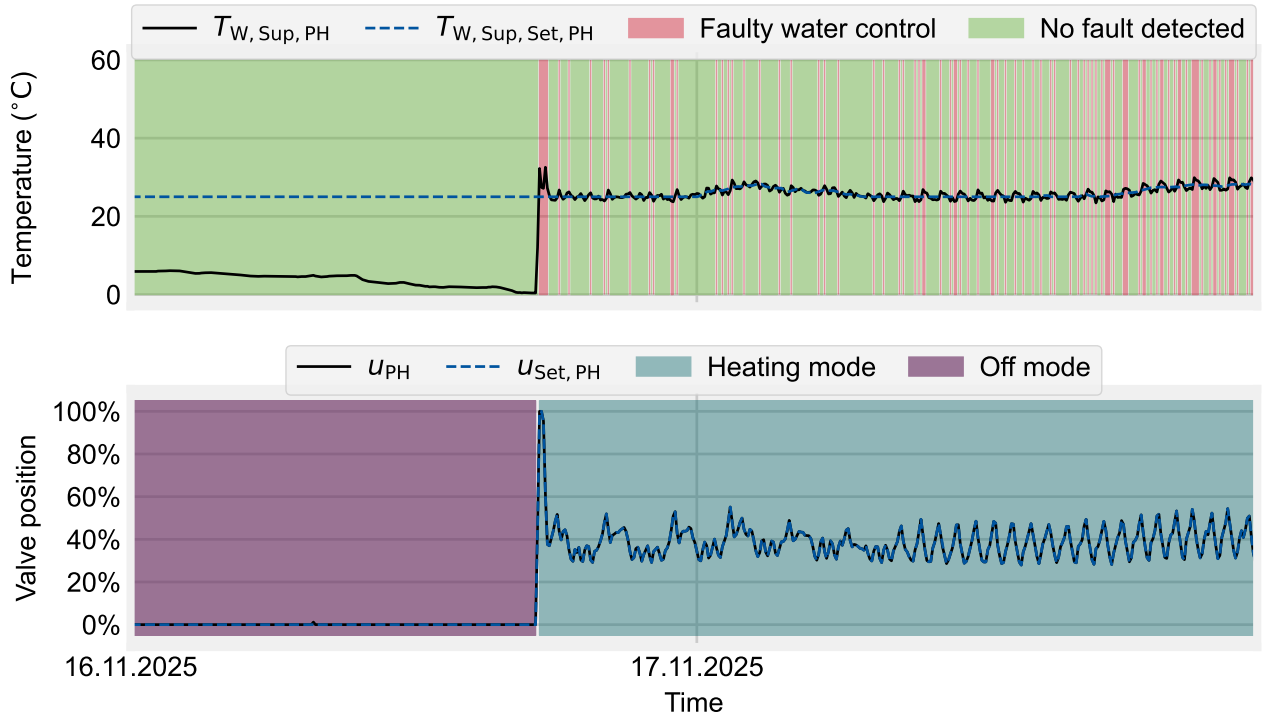

Figure 9. *Fault detection and diagnosis results and operation mode for the fault type "Faulty water control" for the time period from 16.11.2025 to 17.11.2025.*

The detection of the fault types ΔT *too low* and ΔT *wrong sign* is reviewed in the following. In Fig. 10 the supply temperature ($T_{\mathrm{W,Sup,PH}}$), the return temperature ($T_{\mathrm{W,Ret,PH}}$), their difference (ΔT_{PH}), and the operation mode are shown from the 30.11.2025 until the 02.12.2025. For clarity in visualization, these faults are not marked individually, as they frequently occur simultaneously. While the faults are being detected, the supply and return temperature is nearly identical, resulting in a ΔT that is either close to zero or negative. Physically, this would imply no heat is being transferred or the preheater is cooling the air instead of heating it. However, since the preheater only activates when the outdoor air temperature is below 8 °C and the water supply temperature remains consistently above 20 °C during this mode, cooling is physically impossible. Consequently, the measured temperatures must be faulty. To verify this hypothesis, the temperature sensors were calibrated in a calibration bath. Significant offsets were identified and subsequently adjusted starting from the 18.02.2026. Results that include data from the period following this calibration are presented to demonstrate the corrected behavior. Therefore, Fig. 11 shows the heatmap of the fault types ΔT *too low* and ΔT *wrong sign* between the 01.01.2026 and the 15.03.2026 and the operating mode of preheater. Following the sensor calibration on the 18.02.2026, both fault types are not detected again. This confirms that these faults were caused by incorrect

calibration rather than underlying issues with heat transfer performance. Thus, the fault is correctly detected by the methodology before the calibration. Moving forward, both fault types can be interpreted as reliable indicators of faulty sensor calibration, further advancing the diagnostic capabilities of the FDD method.

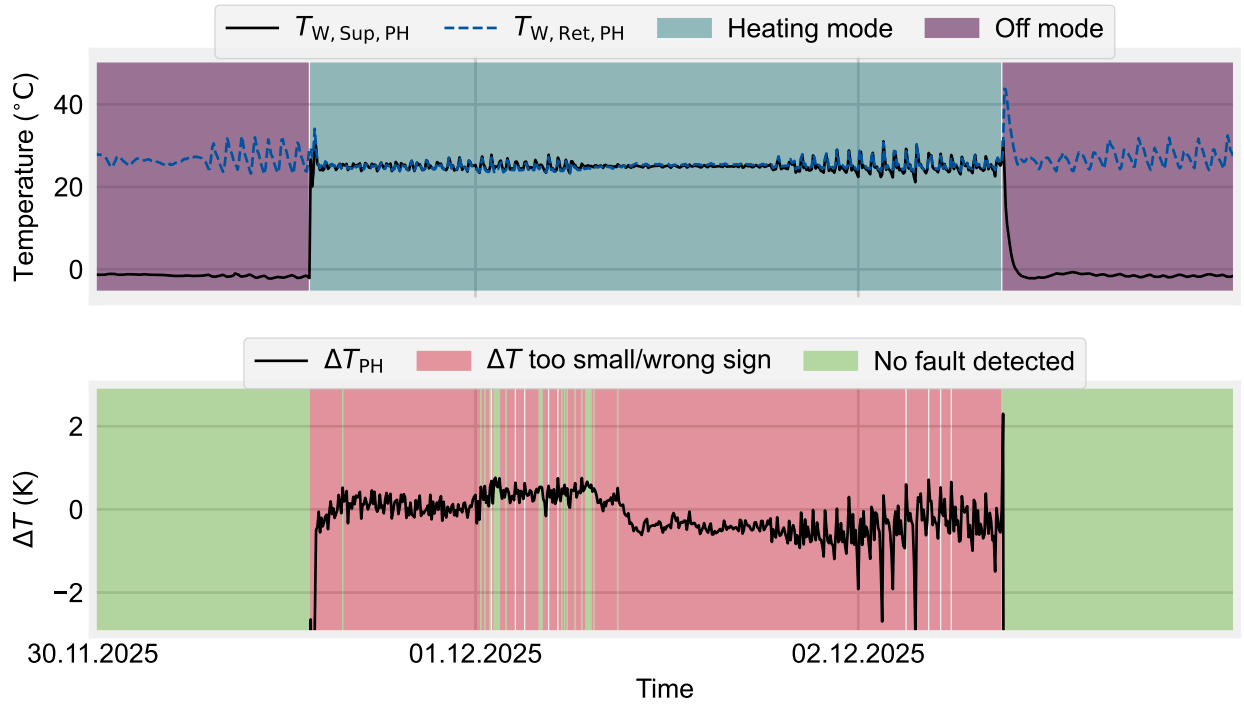

Figure 10. *Fault detection and diagnosis results and operation mode for the fault types "ΔT too low/wrong sign" for the time period from 30.11.2025 to 03.12.2025.*

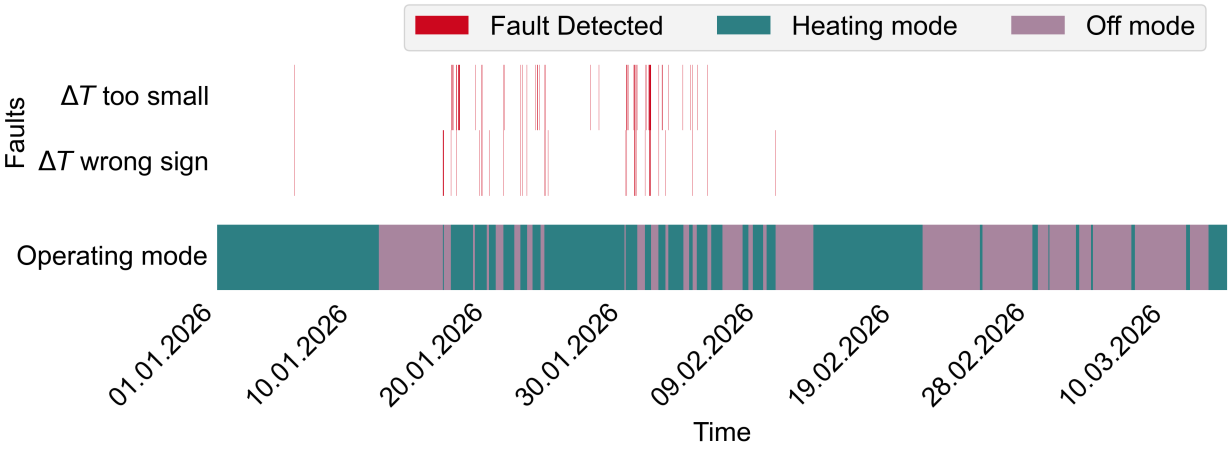

Figure 11. *Heatmap of the fault types "ΔT too low/wrong sign" of the preheater occurred between the 01.01.2026 and the 15.03.2026 and operating mode of preheater.*

Overall, the evaluation of the FDD for the preheater shows promising results. The detected faults successfully reveal operational inefficiencies, incorrectly calibrated sensors, and system failure. Consequently, the proposed FDD methodology can effectively support maintenance and monitoring of AHU components.

3.2. Knowledge graph integration

By adding the semantic KG, uncertainties regarding naming conventions and minor structural variations between the heat exchangers can be handled effectively. Discrepancies in sensor measurement naming schemes are resolved through semantic mapping. Moreover, the KG integration demonstrates that the rules and hierarchical

order modeled for a single mixing circuit can be applied across different heat exchangers, mostly based on topological inference. Nevertheless, some manual adjustments remain necessary, as the supervisory control of the AHU is highly customized, for instance, the supply water setpoint changes based on the operating mode for all heat exchangers. The current SPARQL wrappers are closely coupled with this specific supervisory logic, which may limit transferability to alternative topologies. To ensure broader interoperability, the identification of novel hydraulic structures via SPARQL wrappers must be further automated. In conclusion, while the KG integration already facilitates more automated FDD deployment, future efforts should aim towards fully automated FDD design and generation.

4. Conclusion and outlook

The proposed rule-based semantics-driven FDD reliably uncovers inefficiencies, miscalibrated sensors, and equipment faults, demonstrating the effectiveness of the proof of concept for the preheater use case. The method shows significant potential to enhance maintenance and condition monitoring. Although the hierarchical order successfully identifies initial fault types, these typically correspond to the most initial detectable indicator of a fault based on the available sensor measurements rather than the underlying root causes. Consequently, determining the underlying root cause still require expert analysis in most cases. For more robust and reliable diagnosis, proprietary manufacturer control strategies should be integrated directly into the KG. Future work should expand this methodology to include remaining AHU components and other building energy systems, potentially linking the KG with other FDD approaches, such as machine-learning-based methods.

This study also demonstrates how KG integration can advance automated FDD design. The modeling effort for the rule based FDD across individual heat exchangers is reduced through the KG integration, and generalizability of the approach between similar equipment has been validated. However, the current setup must be tested across diverse heat exchanger configurations and alternative control strategies to further assess its robustness. Moreover, the topology discovery process should be modified to handle various different hydraulic connections. By employing the proposed methodology, the primary modeling effort shifts from manual FDD logic design toward KG development. This shift emphasizes the necessity for automated and standardized KG generation and validation. In future research, the methodology will be evaluated using automatically generated KGs to ensure the advantages of semantic integration are realized.

Acknowledgments

We gratefully acknowledge the financial support by the Federal Ministry for Economic Affairs and Energy (BMWE), promotional reference 03EI4081A and 03EN1117A.

This research was funded by CETPartnership, the Clean Energy Transition Partnership under the 2022 joint call for research proposals, co-funded by the European Commission (GA N°101069750) and with the funding organizations detailed on `https://cetpartnership.eu/funding-agencies-and-call-modules`.

Nomenclature

n	speed, $1/s$
T	temperature, °C
u	valve position, -
$\dot{V}$	volume flow rate, m^3/s

Greek symbols

Δ	difference

Subscripts and superscripts

AHU	air handling unit
BAS	building automation systems
BES	building energy system
CO	cooler
DH	district heating
DCS	district cooling station
DHS	district heating station
EHA	exhaust air
ETA	extract air
FDD	fault detection and diagnosis
KG	knowledge graph
PH	preheater

PLC	programmable logic controller
RH	reheater
Ret	return
Set	setpoint
Sup	supply
W	water

References

[1] IEA. *Tracking Clean Energy Progress 2023*. en. July 2023. URL: https://www.iea.org/reports/tracking-clean-energy-progress-2023 (visited on 03/24/2026).

[2] E. Crowe, Y. Chen, H. Reeve, D. Yuill, A. Ebrahimifakhar, Y. Chen, L. Troup, A. Smith, and J. Granderson. Empirical analysis of the prevalence of HVAC faults in commercial buildings. In: *Science and Technology for the Built Environment* 29.10 (Nov. 2023), pp. 1027–1038. ISSN: 2374-4731. DOI: 10.1080/23744731.2023.2263324. URL: https://doi.org/10.1080/23744731.2023.2263324 (visited on 07/08/2025).

[3] Y. Li and Z. O'Neill. An innovative fault impact analysis framework for enhancing building operations. In: *Energy and Buildings* 199 (Sept. 2019), pp. 311–331. ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2019.07.011. URL: https://www.sciencedirect.com/science/article/pii/S0378778819308266 (visited on 03/24/2026).

[4] S. P. Melgaard, K. H. Andersen, A. Marszal-Pomianowska, R. L. Jensen, and P. K. Heiselberg. Fault Detection and Diagnosis Encyclopedia for Building Systems: A Systematic Review. In: *Energies* 15.12 (2022). ISSN: 1996-1073. DOI: 10.3390/en15124366. URL: https://www.mdpi.com/1996-1073/15/12/4366.

[5] A. Bampoulas, K. Raj, A. Parthiban, M. Saffari, and E. Mangina. Ontology-driven energy management in smart buildings: A comprehensive review of methodologies, tools, and challenges. In: *Energy and Buildings* 352 (Feb. 2026), p. 116817. ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2025.116817. URL: https://www.sciencedirect.com/science/article/pii/S0378778825015476 (visited on 03/24/2026).

[6] M. Pritoni, M. Wetter, L. Paul, A. K. Prakash, W. Huang, S. Bushby, P. Delgoshaei, M. Poplawski, A. Saha, G. Fierro, M. Steen, J. Bender, and P. Ehrlich. Digital and Interoperable: The future of building automation is on the horizon. What's in it for me? en. In: (Aug. 2024). URL: https://escholarship.org/uc/item/191333wd (visited on 01/16/2026).

[7] M. Pritoni, D. Paine, G. Fierro, C. Mosiman, M. Poplawski, A. Saha, J. Bender, and J. Granderson. Metadata Schemas and Ontologies for Building Energy Applications: A Critical Review and Use Case Analysis. en. In: *Energies* 14.7 (Jan. 2021). Publisher: Multidisciplinary Digital Publishing Institute, p. 2024. ISSN: 1996-1073. DOI: 10.3390/en14072024. URL: https://www.mdpi.com/1996-1073/14/7/2024 (visited on 03/13/2026).

[8] G. Lin, H. Kramer, V. Nibler, E. Crowe, and J. Granderson. Building Analytics Tool Deployment at Scale: Benefits, Costs, and Deployment Practices. en. In: *Energies* 15.13 (Jan. 2022). Publisher: Multidisciplinary Digital Publishing Institute, p. 4858. ISSN: 1996-1073. DOI: 10.3390/en15134858. URL: https://www.mdpi.com/1996-1073/15/13/4858 (visited on 03/13/2026).

[9] N. Luo, M. Pritoni, and T. Hong. An overview of data tools for representing and managing building information and performance data. In: *Renewable and Sustainable Energy Reviews* 147 (Sept. 2021), p. 111224. ISSN: 1364-0321. DOI: 10.1016/j.rser.2021.111224. URL: https://www.sciencedirect.com/science/article/pii/S1364032121005116 (visited on 03/30/2026).

[10] S. Runde and A. Fay. Software Support for Building Automation Requirements Engineering—An Application of Semantic Web Technologies in Automation. In: *IEEE Transactions on Industrial Informatics* 7.4 (Nov. 2011), pp. 723–730. ISSN: 1941-0050. DOI: 10.1109/TII.2011.2166784. URL: https://ieeexplore.ieee.org/document/6019030/ (visited on 03/24/2026).

[11] D. Ruikar, C. Anumba, A. Duke, P. Carrillo, and N. Bouchlaghem. Using the semantic web for project information management. en. In: *Facilities* 25.13/14 (Oct. 2007). Ed. by E. Finch, pp. 507–524. ISSN: 0263-2772. DOI: 10.1108/02632770710822607. URL: http://www.emerald.com/f/article/25/13-14/507-524/112617 (visited on 03/24/2026).

[12] *Project Haystack*. URL: https://www.project-haystack.org/.

[13] B. Balaji, A. Bhattacharya, G. Fierro, J. Gao, J. Gluck, D. Hong, A. Johansen, J. Koh, J. Ploennigs, Y. Agarwal, et al. Brick: Towards a unified metadata schema for buildings. In: *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*. 2016, pp. 41–50.

[14] *Open223*. URL: https://github.com/open223.

[15] M. Y. Hwang, B. Akinci, and M. Berges. FSBrick: An information model for representing fault-symptom relationships in HVAC systems. In: *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. BuildSys '23. New York, NY, USA: Association for Computing Machinery, Nov. 2023, pp. 69–78. ISBN: 979-8-4007-0230-3. DOI: 10.1145/3600100.3623729. URL: https://dl.acm.org/doi/10.1145/3600100.3623729 (visited on 01/09/2026).

[16] T. Li, Y. Zhao, C. Zhang, K. Zhou, and X. Zhang. A semantic model-based fault detection approach for building energy systems. In: *Building and Environment* 207 (Jan. 2022), p. 108548. ISSN: 0360-1323. DOI: 10.1016/j.buildenv.2021.108548. URL: https://www.sciencedirect.com/science/article/pii/S0360132321009410 (visited on 03/25/2026).

[17] S. Blechmann, H. Görigk, R. Streblow, and D. Müller. Semantics-based expert system for fault detection in air handling units. en. In: ().

[18] J. Schein, S. T. Bushby, N. S. Castro, and J. M. House. A rule-based fault detection method for air handling units. In: *Energy and Buildings* 38.12 (2006), pp. 1485–1492. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2006.04.014. URL: https://www.sciencedirect.com/science/article/pii/S0378778806001034.

[19] H. Pruvost, A. Wilde, and O. Enge-Rosenblatt. Ontology-Based Expert System for Automated Monitoring of Building Energy Systems. en. In: *Journal of Computing in Civil Engineering* 37.1 (Jan. 2023). Publisher: American Society of Civil Engineers, p. 04022054. DOI: 10.1061/(ASCE)CP.1943-5487.0001065. URL: https://ascelibrary.org/doi/10.1061/%28ASCE%29CP.1943-5487.0001065 (visited on 03/24/2026).

[20] H. Pruvost and A. Wilde. Semantic Provisioning of IoT Devices for Autonomous Fault Detection Services. In: *Proceedings of the 11th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. BuildSys '24. New York, NY, USA: Association for Computing Machinery, Oct. 2024, pp. 307–311. ISBN: 979-8-4007-0706-3. DOI: 10.1145/3671127.3698791. URL: https://dl.acm.org/doi/10.1145/3671127.3698791 (visited on 03/24/2026).

[21] A. Hosseini Gourabpasi and M. Nik-Bakht. BIM-based automated fault detection and diagnostics of HVAC systems in commercial buildings. In: *Journal of Building Engineering* 87 (June 2024), p. 109022. ISSN: 2352-7102. DOI: 10.1016/j.jobe.2024.109022. URL: https://www.sciencedirect.com/science/article/pii/S2352710224005904 (visited on 03/24/2026).

[22] H. Yin. Knowledge Graph-Enhanced Control and Diagnosis in Smart Building Energy Systems. In: *2025 3rd International Conference on Artificial Intelligence and Automation Control (AIAC)*. Oct. 2025, pp. 498–502. DOI: 10.1109/AIAC68175.2025.11332243. URL: https://ieeexplore.ieee.org/document/11332243/ (visited on 03/24/2026).

[23] M. Schraven, A. Kümpel, M. Baranski, and D. Mueller. A comprehensive building HVAC design for application of model-predictive control: Experiences and challenges of construction, commissioning, and operation in a real-world scenario. In: (Apr. 2023). DOI: 10.18154/RWTH-2023-03811.

[24] J. P. Teichmann, J. Oppermann, P. Matthes, D. Müller, P. Mathis, and A. Kümpel. *Reducing energy consumption of air handling units by optimized pump control*. de. Tech. rep. RWTH-2018-224864. Conference Name: Roomvent&Ventilation; 2018. SIY Indoor Air Information Oy, 2018. URL: https://publications.rwth-aachen.de/record/726351 (visited on 01/16/2026).

[25] M. Berktold, M. Rätz, and D. Müller. Towards Automated Planning and Deployment of Building Automation Systems Using Semantic Web Technologies and IFC. Unpublished manuscript. 2026.

[26] VDI Verein Deutscher Ingenieure e.V. *VDI 3814 Blatt 3.1 - Gebäudeautomation (GA) - GA-Funktionen - Automationsfunktionen*. Tech. rep. Ausgabedatum: 2019-01. Berlin: DIN Media GmbH, Jan. 2019. URL: https://www.vdi.de.