Results in Engineering 32 (2026) 113471

Contents lists available at ScienceDirect

Results in Engineering

journal homepage: www.sciencedirect.com/journal/results-in-engineering

Research paper

Integrated fault detection and supervisory compensation for hydronic radiator faults in district-heated buildings

Nima Monghasemi *, Stavros Vouros, Konstantinos Kyprianidis, Amir Vadiee

Department of Engineering Sciences, Mälardalen University, 721 23, Västerås, Sweden

ARTICLE INFO

Keywords:
District heating
Hydronic radiator
Fault detection
Fault compensation
Building energy simulation

ABSTRACT

Hydronic radiator faults in district-heated buildings can degrade comfort, alter return-water temperatures, and increase unnecessary heat use. Although many fault detection and diagnostics studies identify anomalies, fewer close the loop by translating diagnosis into corrective action. This paper develops an integrated fault detection and compensation framework that couples unsupervised anomaly detection with signature-based diagnosis and supervisory setpoint compensation. One-Class Support Vector Machine, Isolation Forest, and Local Outlier Factor models were trained only on fault-free seasonal simulation data and embedded in the EnergyPlus Python API loop for causal, timestep-by-timestep operation. The framework was evaluated using baseline, faulty, and compensation-enabled scenarios, followed by severity sweeps for three representative hydronic faults: a stuck-closed valve, a zone-level over-delivery setpoint-bias proxy for a stuck-open valve, and a system-level supply-temperature setpoint bias. The results show that the achievable supervisory recovery is fault dependent. For the stuck-open proxy, the framework recovered 68.8% of fault-induced discomfort degree-hours and reduced hydronic energy use by 15% relative to the uncompensated fault. For the supply-setpoint bias, recovery reached 57.2%, although the compensation increased the return-water temperature. For the stuck-closed fault, recovery was limited to 38.6%, demonstrating the hydraulic ceiling of supervisory compensation. The severity sweep identified detection floors, saturation regions, and recovery limits. These findings support a fault-specific strategy in which selected control-type faults may be candidates for supervisory compensation, subject to deployable release logic and safety validation, whereas physically constrained faults are prioritized for maintenance.

1. Introduction

District heating (DH) is widely recognized as a key infrastructure for decarbonizing the urban heat supply. The European Union aims for climate neutrality by 2050 and a 55% reduction in emissions by 2030. DH is central to this, with its share of final heat consumption projected to rise from 15% in 2020 to 31% by 2050 [1]. The evolution of DH technology through successive generations, from steam-based first-generation systems to low-temperature fourth- and fifth-generation concepts, reflects a persistent drive toward lower operating temperatures and higher energy efficiency [2]. Fourth-generation systems are characterized by supply temperatures of 60–70 °C and an emphasis on integrating renewable and waste heat sources. However, most existing networks still operate at supply temperatures between 80 and 100 °C [3]. Their transition toward lower temperatures presents both hydraulic and thermal challenges at the network, substation, and building levels [4].

Achieving low return temperatures on the end-user side is a critical prerequisite for realizing temperature reductions. The return temperature is a key indicator of the overall system efficiency. It dictates network heat losses, influences production performance, and determines the feasibility of integrating low-grade heat sources [5]. Notably, the return temperature cannot be directly controlled at the network level. Instead, it is a consequence of how effectively the building-level heating system, including radiators and substations, extracts heat from the supply water [6]. The sensitivity of return temperatures to substation-level decisions is further illustrated by Li et al. [7], who showed that switching between different control strategies in a dual-temperature district heating substation can shift the seasonal average return water temperature by several K within the same network. This makes building-level hydronic systems an operational bottleneck in district heating performance. Faults in customer installations, whether physical degradations such as valve stiction and flow restrictions or control logic errors, can

* Corresponding author.
E-mail addresses: nima.monghasemi@mdu.se (N. Monghasemi), stavros.vouros@mdu.se (S. Vouros), konstantinos.kyprianidis@mdu.se (K. Kyprianidis), amir.vadiee@mdu.se (A. Vadiee).

https://doi.org/10.1016/j.rineng.2026.113471
Received 2 June 2026; Received in revised form 5 September 2026; Accepted 9 October 2026
Available online 12 October 2026
2590-1230/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Nomenclature

Variables

c_p	specific heat capacity ($J\,kg^{-1}\,K^{-1}$)
DDH	discomfort degree-hours (°C·h)
E_{hyd}	delivered hydronic heating energy (kWh)
e_T	temperature error (°C)
$\bar{e}_{T,2h}$	two-hour rolling mean of e_T (°C)
$H(t)$	fault-window indicator function (–)
$\dot{m}$	hot-water mass flow rate ($kg\,s^{-1}$)
$\dot{m}_{ref}$	fixed fault-free reference mass flow rate ($kg\,s^{-1}$)
$N_{persist}$	persistence filter length (timesteps)
$N_{release}$	clear-streak counter for compensation release (timesteps)
N_{retry}	consecutive unmatched-signature counter for diagnosis reset (timesteps)
$r_{\dot{m}}$	normalized mass flow ratio (–)
T	temperature (°C)
$\bar{T}_{ret}$	average return water temperature (°C)
$\bar{T}_{zone}$	average zone air temperature (°C)
T_{SP}	supply-temperature setpoint (°C)
$T_{sup,baseline}$	fault-free supply-water temperature setpoint (°C)
$T_{sup,fault}$	supply-water temperature setpoint under the injected fault (°C)
$T_{set,schedule}$	nominal scheduled zone heating setpoint (°C)
T_{target}	compensation target temperature (°C)
$\mathcal{T}_{fault}$	set of occupied timesteps within the fault window
$\mathcal{T}_{heat}$	set of heating-active timesteps within the fault window
$\dot{T}_{zone}$	zone temperature rate of change (°C per timestep)
Δt_h	simulation timestep in hours (h)
Δt_s	simulation timestep in seconds (s)
ΔT_{bias}	stuck-open setpoint bias (°C)
ΔT_{hw}	hot-water supply–return differential (K)
ΔT_{supply}	system-level supply-temperature setpoint bias (K)

Greek Symbols

α	exponential ramp gain (–)
φ	flow-restriction fraction (–)

Derived indicators

ΔE_{base}	change in delivered hydronic heating energy relative to the fault-free baseline (%)
ΔE_{fault}	change in delivered hydronic heating energy relative to the uncompensated fault (%)
ΔT_{ret}	shift in the flow-weighted return-water temperature relative to the fault-free baseline (K)
t_{90}	analytic time for the supply setpoint ramp to complete 90% of a step change (min)

Subscripts

baseline	fault-free reference condition
current	current value at present timestep
fault / faulty	uncompensated fault condition
FDC	condition with fault detection and compensation active
hw	hot water
max	maximum value
new	updated value after compensation
out	outdoor
ref	fixed fault-free reference value
ret	return water
schedule	scheduled value
set	heating setpoint
SP	supply-temperature setpoint
sup	supply water
zone	zone air

Abbreviations

DDH	discomfort degree-hours
IF	Isolation Forest
DH	district heating
LOF	Local Outlier Factor
EMS	energy management system
ML	machine learning
FDC	fault detection and compensation
OCSVM	One-Class Support Vector Machine
FDD	fault detection and diagnostics
API	application programming interface
FPR	false-positive rate
BEMS	building energy management system
FTC	fault-tolerant control
DHN	district heating network
HVAC	heating, ventilation, and air conditioning
LSTM	long short-term memory
DOE	U.S. Department of Energy
PCA	principal component analysis
NARX	nonlinear autoregressive with exogenous inputs
SHGC	solar heat gain coefficient
SHAP	Shapley additive explanations
HIL	hardware-in-the-loop
RL	reinforcement learning

degrade network performance through various mechanisms. These include elevated return temperatures, compressed supply–return differentials, and inadequate heat delivery under flow restrictions [8].

A survey of Swedish DH utilities found that leakages and faults in building internal heating systems were the most frequently reported fault categories [9]. Compared to air-side HVAC systems, fault detection and diagnostics (FDD) for hydronic and district-heating systems relies on a smaller and less standardized evidence base. This disparity stems largely from lower digitalization levels and the subtle thermal signatures of hydronic faults, yet their prevalence remains a critical operational concern [10–13]. Studies indicate that up to 74% of buildings connected to district heating networks contain faults that elevate return temperatures [14]. An analysis of one year of hourly meter readings from 135 substations in two Swedish district heating networks similarly found that approximately three quarters exhibited at least one fault or sub-optimal behavior, including anomalous heat-load patterns, low average temperature differences, and poor substation control [15]. More recently, Leiria et al. [16] segmented and clustered end-user substation fault behavior in Danish smart-meter data, reinforcing the evidence that customer-side faults are widespread. These findings are especially consequential in the context of fourth-generation systems, which operate with narrower temperature margins and are therefore more sensitive to customer-side disturbances than previous-generation systems. Such faults primarily impact building occupants by causing discomfort through inadequate or excessive heating. At the network level, they compress the supply–return temperature differential and increase distribution flow rates, undermining the energy and exergy efficiency that low-temperature operation aims to achieve.

This context motivates linking fault detection directly to corrective action at the building level. In this study, supervisory compensation

denotes a bounded, temporary control action that mitigates the comfort and energy consequences of a diagnosed fault by modulating the supply-temperature setpoint.

The aim of the study is to develop and evaluate an integrated framework that combines unsupervised anomaly detection, signature-based diagnosis and fault-specific supervisory compensation for hydronic radiator faults in a district-heated building. The framework is tested under isolated single-fault conditions, focusing on detection performance, comfort and energy recovery, and the practical limits that appear as fault severity changes.

2. Literature review and research gap

This section reviews the literature on fault detection and diagnosis for HVAC and district heating systems and then identifies the research gap addressed by the present study.

2.1. Data-driven fault detection and diagnosis

The growing availability of operational data from building automation systems and smart metering infrastructure has motivated substantial research on data-driven FDD for HVAC and DH systems in recent years. Several recent surveys have mapped this expanding landscape of research. Chen et al. [10] reviewed data-driven FDD methods for building HVAC systems, documenting a shift from rule-based commercial tools to machine learning approaches that require little a priori knowledge of physical processes. Similar trends were evident in the district heating (DH) domain. Van Dreven et al. [11] surveyed 57 studies on intelligent FDD in DH networks, while Neumayer et al. [12] focused on the methods and datasets used for fault and anomaly detection in DH substations. Across these surveys, a consistent finding is the predominance of unsupervised methods, which are driven by the scarcity of labeled fault data in operational systems. Labeling faults in real buildings is time-consuming, costly, and typically limited to those occurring during a monitored period [17].

Within this unsupervised paradigm, various techniques have been applied. Soultanzadeh et al. [17] developed a principal component analysis (PCA)-based unsupervised FDD method transferable across different light-commercial building configurations, whereas Lee et al. [18] combined Gaussian Mixture Models with long short-term memory (LSTM) autoencoders for proactive radiator fault diagnosis using vibration signals. In the DH substation context, Hermans et al. [19] showed that the Gaussian error assumption underlying conventional anomaly detection methods breaks down in practice and proposed quantile regression as a more robust, data-conditioned alternative. At the network scale, van Dreven et al. [20] introduced a shared nearest-neighbor clustering approach that approximates the relative DH network topology, thereby enabling localized anomaly detection without requiring the processing of sensitive geographical data.

Deep learning methods have also gained traction in FDD for HVAC. Zhang et al. [21] conducted a systematic review with meta-analysis covering 47 eligible studies and found convolutional neural networks to be the most frequently applied architecture; however, they noted that the majority of studies relied on simulation or laboratory data, leaving real-world deployment largely untested. A related limitation of supervised and deep learning methods is their dependence on comprehensive labeled datasets and their restricted transferability across buildings with different system configurations. Transfer learning has emerged as a strategy to mitigate this problem. Feng et al. [22] proposed an attention-empowered domain adaptation method for HVAC sensor fault diagnosis that achieved a cross-building accuracy of 93%. Chen et al. [23] applied transfer learning with interpretable Shapley additive explanations (SHAP)-based analysis for cross-system chiller fault diagnosis, demonstrating that fine-tuning strategies outperformed non-transferred models. Data-driven digital twin concepts have been explored for real-time anomaly surveillance at the DH network level. Bahlawan et al. [24] coupled a physics-based DHN simulation model with an optimization algorithm that detected and identified both thermal and hydraulic pipe faults, including water leakages and anomalous heat and pressure losses, across a university campus district heating network. Losi et al. [25] employed nonlinear autoregressive with exogenous inputs (NARX) neural

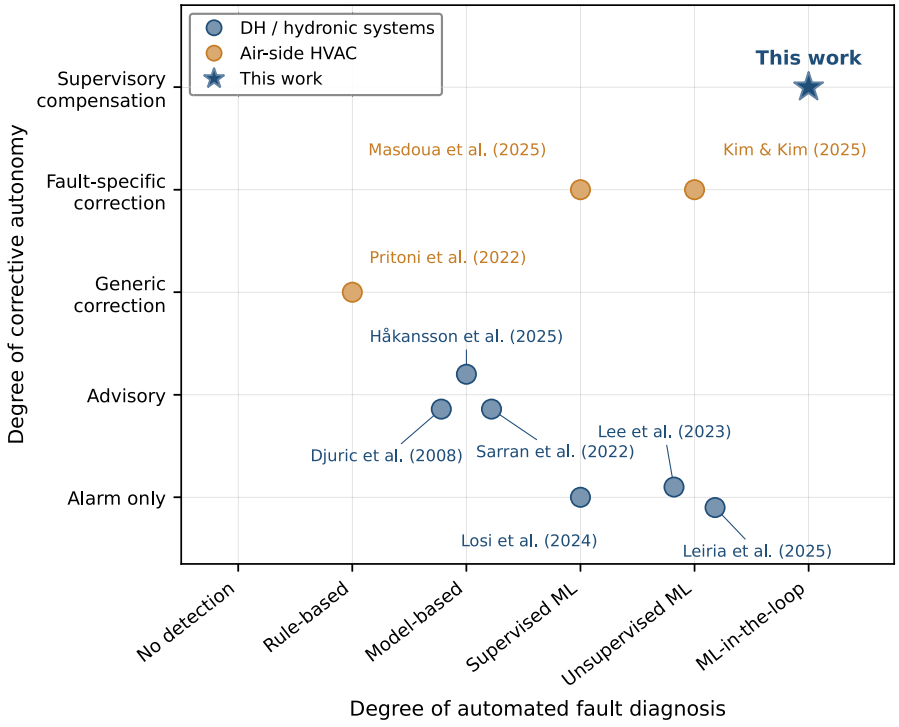

Fig. 1. Literature positioning map for fault diagnosis and corrective autonomy in building energy systems. Positions are qualitative and intended for literature comparison rather than quantitative scoring.

networks trained on healthy operational data to simulate the expected DH network behavior and detect deviations using a threshold-based criterion. Rätz et al. [26] combined multiple regression models with a knowledge-based diagnosis stage, circumventing the need for labeled fault data in DH substations.

2.2. Hydronic heating fault detection

Beyond the air-side HVAC domain, a smaller but established body of work has addressed FDD for hydronic heating and district-heating installations directly. Djuric et al. [27] combined BEMS data with a calibrated simulation model and an optimization tool to diagnose faults in a real hydronic heating system, including an erroneous outdoor-temperature measurement, incorrect operating schedules, and flow imbalance. Sarran et al. [28] developed gray-box models and virtual sensors for the continuous commissioning of hydronic floor heating, identifying underperforming loops and estimating advisory supply-temperature reductions. This activity is also reflected at the sector level. The IEA DHC Annex TS4 reviews automatic FDD in substations and end-user heating systems [13], while the ongoing Annex TS9 addresses the data integration and interoperability needed for digital applications, including FDD [29]. Nevertheless, the closest study-level contributions remain focused mainly on diagnosis, continuous commissioning, and advisory intervention rather than closed-loop supervisory compensation.

2.3. Automated correction

Despite the breadth and growing maturity of these detection and diagnostic methods, a fundamental limitation persists in nearly all reported systems. Unsupervised, supervised, and deep-learning-based approaches are predominantly designed to identify the occurrence of a fault and, in some cases, to classify its type or estimate its severity; however, they stop short of closing the loop with automated corrective actions. The fault signal is typically passed to a human operator or maintenance dispatch process, with no embedded supervisory strategy to compensate for the detected fault in real time. This decoupling of detection from response implies that the energy penalty and occupant discomfort caused by the fault persist until manual intervention occurs, a delay that can span hours or days in operational district heating networks.

In parallel with detection research, a separate stream of work has addressed the compensation and correction of operational faults in building and district heating systems. At the network level, Guelpa et al. [3] reviewed strategies for reducing supply temperatures in existing district heating networks, including the exploitation of network oversizing, substation upgrades, and demand-side flexibility measures; however, these interventions target steady-state design conditions rather than real-time fault events. At the building level, Håkansson et al. [30] proposed a model-assisted hydronic balancing method that uses operational sensor data to evaluate radiator flow rate performance and compute optimal valve reconfigurations, addressing the chronic problem of unbalanced flow distribution in multi-zone heating systems. In the broader HVAC domain, Pritoni et al. [31] conducted one of the few field studies on automated fault correction, deploying seven auto-correction algorithms integrated with commercial FDD platforms across four buildings and three building automation systems. The study demonstrated that faults, such as incorrect schedules, unreleased overrides, and suboptimal setpoint resets, could be corrected autonomously. Complementing this, Markus et al. [32] showed that virtual meters can project the energy savings achievable by correcting HVAC control faults before any intervention, providing operators with a quantitative basis for prioritizing fault correction. Stecher et al. [33] underscored that the foundational step of creating properly labeled fault datasets for district heating substations remains an open challenge, with no standardized feature sets or labeling conventions. This limitation has impeded the transition from passive anomaly flagging to actionable fault-management.

Recently, fault-tolerant control (FTC) concepts originating from the process control and aerospace domains have been adapted for HVAC systems. An early example is the work of Du and Jin [34], who applied multi-level principal component analysis models together with joint angle plots and a reconstruction scheme to detect, isolate, and recover from multiple sensor faults in variable air volume systems. Masdoua et al. [35] developed an active FTC architecture that combined a convolutional N-BEATS fault classifier with a double deep Q-network controller, thereby enabling a reinforcement learning (RL) agent to reconfigure its control policy in real time upon the detection of actuator or sensor faults in a simulated heating system. Kim and Kim [36] demonstrated a closed-loop fault management framework for air-handling units in which virtual-sensor-assisted autoencoder detection triggers a transactive control strategy that adjusts zone temperature setpoints in proportion to fault severity, achieving approximately 10% energy cost savings while maintaining occupant comfort within an agent-based hardware-in-the-loop (HIL) simulation.

2.4. Research gap and literature positioning

Although these recent studies represent significant advances in linking fault detection with automated corrective actions, they share certain limitations with conventional compensation methods. Many of these strategies remain agnostic to the specific fault and apply generic setpoint adjustments without diagnosing the underlying cause, and several address air-side HVAC systems rather than hydronic networks. Consequently, there remains a need to integrate online detection using machine learning with supply temperature compensation that relies on diagnostics and addresses specific faults in hydronic systems at the building level. Previous studies have largely left unexplored how the supervisory recovery potential fundamentally differs between errors based on control and faults constrained by physical processes across varying severity levels.

Fig. 1 synthesizes the literature by mapping representative studies along two dimensions: diagnostic automation and corrective autonomy. Existing research on district heating and hydronic systems has focused on fault detection without corrective action [16,18,25] or model-assisted analysis and advisory reconfiguration without automated closed-loop action [27,28,30]. Although integrated fault-tolerant architectures are emerging for air-side HVAC and standalone heating loops [35,36], they do not address the temperature dynamics of the district heating networks. As illustrated, among the hydronic heating studies, few combine online data-driven diagnosis with fault-specific closed-loop supervisory action.

This study addresses this gap by developing and evaluating an integrated fault detection and compensation (FDC) framework in which unsupervised machine learning models operate inside the building energy simulation loop, enabling online detection that drives fault-specific supervisory compensation of the supply temperature. The detection methodology builds on an earlier study by the authors [37], which demonstrated that a One-Class Support Vector Machine (OCSVM) model trained on fault-free radiator data could identify valve and thermostatic radiator valve faults of different severity levels. A subsequent study [38] integrated a rule-based detection proxy with a supervisory compensation module, revealing a dichotomy in the supervisory recovery potential of physically constrained versus control-based faults. The present work extends both studies by replacing the rule-based proxy with three trained unsupervised machine learning (ML) models, expanding the fault taxonomy to include a system-level supply-setpoint bias fault, and conducting a parametric severity analysis that maps the detection floors, compensation saturation thresholds, and boundaries of supervisory controllability across all three fault types.

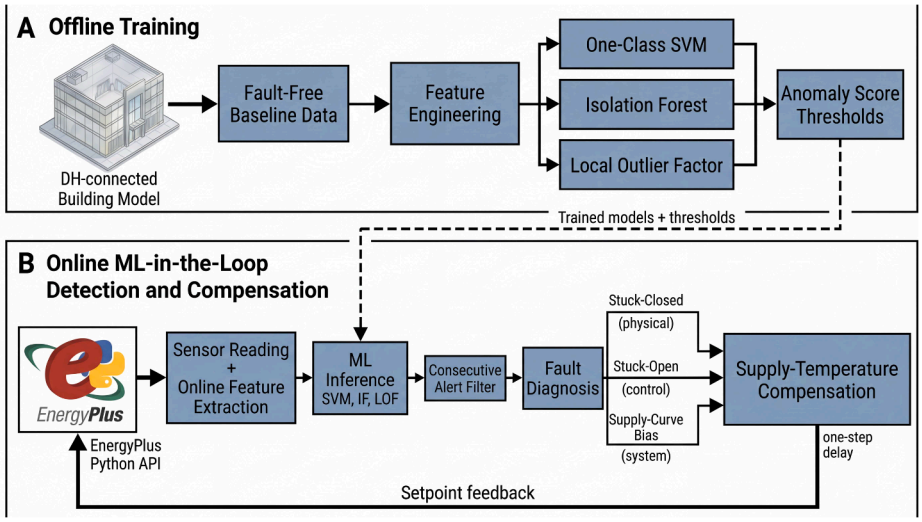

Fig. 2. System architecture of the integrated FDC framework: (A) offline model training and (B) online closed-loop detection, diagnosis, and compensation.

2.5. Contributions and paper outline

This study addresses the gap between passive fault detection and active corrective autonomy by proposing an integrated fault detection and compensation (FDC) framework. The specific contributions are as follows:

- An ML-in-the-loop control architecture that embeds unsupervised anomaly-detection models inside the EnergyPlus simulation loop via the Python API, enabling online fault detection that drives signature-based diagnosis and supply-temperature compensation across three fault types spanning physical and control failure mechanisms at both zone and system levels.
- A characterization of supervisory recovery limits across three fault types and five severity levels, establishing that the two tested control-type faults are recoverable through supply-temperature modulation whereas physically constrained faults require maintenance dispatch, and identifying the detection floors and recovery ceilings that bound the framework's operational envelope.

The remainder of this paper is organized as follows. Section 3 describes the simulation testbed, fault taxonomy, ML training pipeline, and integrated FDC architecture. Section 4 presents the detection tuning, scenario comparison results, dynamic fault responses, compensation-parameter sensitivity, severity analysis and compound-fault stress tests. Section 5 interprets the findings, discusses limitations, and identifies directions for future work. Section 6 summarizes the key conclusions and their implications for operational fault management in DH-connected buildings.

3. Methodology

The integrated FDC framework operates in two stages, as shown in Fig. 2. In the offline stage, fault-free baseline data from a full heating season were filtered to retain only occupied periods with active heating and were then transformed into six features. Three unsupervised anomaly detection models were trained on this feature set, and their anomaly score thresholds were calibrated at multiple percentiles, producing a library of detection configurations that were subsequently optimized per fault type.

In the online stage, the trained models were embedded inside the EnergyPlus simulation loop via the Python application programming interface (API). At each timestep, the runtime acquires the sensor readings, reconstructs the feature vector using a rolling buffer, and evaluates all three models to produce anomaly scores. These scores were compared with calibrated thresholds to generate binary flags that must persist for a configurable number of consecutive timesteps before the diagnosis proceeds. Once a sustained anomaly is confirmed, a signature-based diagnosis module classifies the fault as being stuck-closed, stuck-open, or having a supply-setpoint bias. The framework then executes a compensation action by modulating the temperature setpoint, increasing it for under-delivery faults, or decreasing it for over-delivery. The compensated setpoint was fed back into the simulation at the next timestep, thereby completing the closed loop with a one-step causal delay.

The framework was evaluated through a controlled three-scenario comparison for each fault type and a parametric severity sweep across five fault severity levels.

Table 1
Simulation testbed and building characteristics of the modified model.

Item	Specification
Building model	U.S. DOE Retail Standalone (modified)
Location	Västerås, Sweden (ASHRAE Climate Zone 6A)
Simulation period	1 October – 30 April
Simulation engine	EnergyPlus V25.1.0 + Python API
Timestep	10 min
Primary zone	Core retail
Additional zones	Back space, front retail
Heating system	Water-based baseboard heaters
Ventilation	Decoupled; energy recovery ventilators
Baseline supply temperature	60 °C
Conditioned floor area	2294 m^2
Thermal zones	5 (all conditioned; no plenums)
Exterior-wall / roof U-value	0.48 / 0.36 $W\,m^{-2}\,K^{-1}$
Ground-floor U-value	0.05–0.42 (mean 0.11) $W\,m^{-2}\,K^{-1}$
Window U-value / SHGC	3.24 $W\,m^{-2}\,K^{-1}$ / 0.39
Window-to-wall ratio	11.0% overall; 39.1% south façade

3.1. Simulation testbed

The simulation testbed modifies the U.S. Department of Energy (DOE) Retail Standalone commercial reference building [39] by replacing the original rooftop HVAC with a centralized hydronic heating plant. Water-based baseboard heaters emulate a district heating connection, and ventilation is fully decoupled to ensure that indoor air quality control does not interfere with fault injection dynamics. The baseboard heaters serve as EnergyPlus [40] surrogates for hydronic radiator terminals, as both are water-to-air convective emitters governed by the same heat transfer equations in the simulation engine.

Table 2
Fault taxonomy.

Fault type	Nature / scope	Parameter	Reference severity	Severity levels
Stuck-closed	Physical, zone	φ (flow fraction)	0.30	0.10, 0.20, 0.30, 0.50, 0.70
Stuck-open	Control, zone	ΔT_{bias} [°C]	+2.0	+0.5, +1.0, +2.0, +4.0, +6.0
Supply-setpoint bias	Control, system	ΔT_{supply} [K]	−8	−2, −5, −8, −12, −15

As detailed in Table 1, the core retail zone served as the primary subject for fault injection, with two additional zones instrumented to enable cross-zone impact assessment for the supply-setpoint bias fault.

The envelope U-values reported in Table 1 were extracted from the EnergyPlus envelope summary of the modified model. Ground-floor conductance varies across the five F-factor floor constructions, and the area-weighted mean is therefore reported alongside the full range. Skylights above the core retail zone have an assembly U-value of 3.79 $\text{W m}^{-2}\,\text{K}^{-1}$ and a solar heat gain coefficient of 0.50. Over the heating season, the fault-free baseline delivered 166.8 MWh of whole-building hydronic space heat, corresponding to a heating intensity of 72.7 kWh m^{-2} across the 2294 m^2 conditioned floor area. Metered directly at the district heating interface, this figure represents space heating alone and excludes domestic hot water, auxiliary electrical consumption, and cooling. While this seasonal benchmark characterizes the entire building, the fault investigations in subsequent sections focus primarily on the hydronic branch supplying the core retail zone, which accounts for 1600 m^2 or 70% of the conditioned space.

3.2. Fault taxonomy

Radiator-side faults can be categorized by their nature, distinguishing between physical degradation and control logic errors, as well as their scope, which differentiates between impacts on a single zone and those affecting the entire building. Physical faults impose hard constraints on heat delivery that a supervisory action cannot fully alleviate, whereas control faults distort the demand signal while preserving hydraulic paths. Three representative faults were selected to span this taxonomy, as summarized in Table 2.

A stuck-closed valve represents a partial mechanical obstruction that restricts the hot water mass flow rate through the radiator. The obstruction is parameterized by a flow-restriction fraction φ, defined as the ratio of the maximum permissible flow under fault conditions to a fixed fault-free reference flow

$$\varphi = \frac{\dot{m}_{\text{max,fault}}}{\dot{m}_{\text{ref}}}, \qquad \varphi \in \{0.10,\ 0.20,\ 0.30,\ 0.50,\ 0.70\} \tag{1}$$

where $\dot{m}_{\text{ref}} = 1.98\ \text{kg s}^{-1}$ is the maximum branch flow observed in the gated fault-free training dataset and $\dot{m}_{\text{max,fault}}$ is the enforced upper limit on the mass flow rate. This same reference constant defines the normalized mass flow ratio, $r_{\dot{m}} = \dot{m}/\dot{m}_{\text{ref}}$, evaluated throughout the detection and diagnosis stages. Because peak unrestricted branch flow reaches approximately 2.68 kg s^{-1} during the colder January evaluation window, $r_{\dot{m}}$ represents a fixed-reference scaling factor rather than a bounded fraction, allowing values to rise above unity under high heating demand. At $\varphi = 0.30$, only 30% of this reference flow capacity is available, causing the zone temperature to drop below the heating setpoint. Under milder restrictions (e.g., $\varphi \geq 0.50$), sufficient flow remains for the return temperature to reflect the actual heat extraction performance. The expected thermal signature is a sustained negative temperature error that coincides with an anomalously low flow. Because the constraint is hydraulic, the compensation strategy is to elevate the supply water temperature to maximize heat transfer within the restricted flow. The fault was activated during the occupied hours for a two-day window. Each severity level requires a dedicated EnergyPlus input file in which the EMS actuator enforces a restricted mass flow rate.

A stuck-open valve arises from a control malfunction that keeps the radiator valve open beyond the actual heating requirements. This fault is implemented as a positive bias on the zone heating setpoint during the occupied hours as follows:

$$T_{\text{set,fault}}(t) = T_{\text{set,schedule}}(t) + \Delta T_{\text{bias}} \times H(t),$$
$$\Delta T_{\text{bias}} \in \{+0.5,\ +1.0,\ +2.0,\ +4.0,\ +6.0\}\ ^{\circ}\text{C}, \tag{2}$$

where $H(t)$ is unity during the fault window and zero otherwise. Because the EnergyPlus zone thermostats use a dual-setpoint definition, the inactive upper cooling setpoint was raised in parallel, solely to preserve a valid deadband above the biased heating setpoint. The testbed contains no cooling equipment, and no cooling energy was delivered in any simulation. This setpoint-bias implementation is a control-level proxy for over-delivery faults. It reproduces the sustained heat over-delivery, zone overheating, and compressed supply–return differential that a mechanically stuck-open valve would cause, but it acts through the control path rather than the hydraulic path. A mechanically seized valve would hold the flow high irrespective of the control signal, whereas the proxy raises the effective demand, so the two can differ in their flow signatures under part-load conditions. The resulting over-delivery manifests as persistent overheating, elevated energy consumption, and a compressed supply–return differential. Once diagnosed, the supervisory layer corrects the erroneous setpoint and reduces the supply water temperature. The fault was injected via the Python API over the same two-day window as the stuck-closed case.

The supply-setpoint bias extends the taxonomy to a system-level failure affecting all zones simultaneously. DH substations typically regulate the supply temperature through a weather-compensated curve that maps outdoor temperature to a supply-temperature command. Several realistic failure mechanisms can bias this command such as a drift in the outdoor-temperature sensor, a disconnected sensor replaced by an erroneous fallback value and an incorrect reset-curve slope or intercept. Faulty outdoor-temperature measurements of this kind have been documented in a real hydronic heating system by Djuric et al. [27]. In the present testbed, the baseline supply temperature is a constant 60 °C, and the fault is injected as a constant negative offset:

$$T_{\text{sup,fault}}(t) = T_{\text{sup,baseline}} + \Delta T_{\text{supply}},$$
$$\Delta T_{\text{supply}} \in \{-2,\ -5,\ -8,\ -12,\ -15\}\ \text{K}, \tag{3}$$

where $T_{\text{sup,baseline}} = 60\ ^{\circ}\text{C}$. This constant offset is a controlled abstraction that reproduces the thermal consequence of the reset failures listed above while isolating the effect of fault severity. For instance, at $\Delta T_{\text{supply}} = -8$ K, the supply temperature drops to 52 °C, reducing the thermal driving force across all radiators. The key diagnostic distinction from the stuck-closed fault lies in the scope, where a supply-setpoint bias depresses the temperatures across multiple zones simultaneously, whereas the zone-level flow rates may appear normal. The compensation strategy also involves boosting the supply temperature; however, because the constraint is informational rather than hydraulic, the supervisory action can recover the thermal performance more effectively. In the simulation, the depressed supply temperature was replicated by directly overriding the hot-water loop temperature schedule using the Python API. A 48-hour fault window was used to capture the full diurnal heating cycle, including overnight periods.

Table 3
Engineered features for anomaly detection.

Feature	Definition	Physical meaning
Temperature error, e_T	$T_{\text{zone}} - T_{\text{set,schedule}}$	Instantaneous comfort deviation
Flow ratio, $r_{\dot{m}}$	$\dot{m} / \dot{m}_{\text{ref}}$	Normalized mass flow rate
Rolling temperature error, $\bar{e}_{T,2\text{h}}$	2-hr rolling mean of e_T	Comfort trend indicator
Zone temperature rate, $\dot{T}_{\text{zone}}$	ΔT_{zone} (10-min first difference)	Thermal inertia proxy
Outdoor temperature, T_{out}	Outdoor dry-bulb temperature	Environmental context
Hot-water differential, ΔT_{hw}	$T_{\text{hw,inlet}} - T_{\text{hw,outlet}}$	Heat extraction rate

3.3. Offline training

The detection component employs unsupervised anomaly detection models trained exclusively on fault-free baseline data. By learning the boundaries of normal operation, these models flag emerging deviations without requiring prior examples of specific faults or manually constructed fault models. The methodology for constructing the baseline features was originally developed in [37] and this study retains the same feature set but retrains the models on the testbed.

Training data were generated from a fault-free simulation spanning a full heating season (October–April), yielding approximately 30,500 timesteps at a 10-minute resolution. Two filters were applied to retain only the periods of active thermal regulation. The building must be occupied, defined by a setpoint $T_{\text{set}} \geq 18$ °C, and the heating system must be active, as indicated by a mass flow rate $\dot{m} > 0.001$ kg/s. Six features were then engineered from the raw sensor signals, as summarized in Table 3. Throughout the framework the temperature error e_T is evaluated against the nominal scheduled heating setpoint $T_{\text{set,schedule}}$, read from the thermostat schedule, and never against the fault-biased control setpoint. The stuck-open bias is injected through the zone temperature-control actuator, which overrides the thermostat without altering the schedule, so the same definition of e_T serves the detection features, the diagnostic signatures and the discomfort metric. The temperature error and flow ratio characterize the operating point. A two-hour rolling mean of the temperature error smooths the transient fluctuations over a 12-step window. The zone temperature rate captures the thermal inertia effects that distinguish slow-onset faults from rapid-load transients. The outdoor temperature provides an environmental context, and the hot water differential reflects the instantaneous heat extraction rate.

The filtered and feature-engineered data were standardized to zero mean and unit variance, and three unsupervised anomaly-detection models were trained on the fault-free feature set, namely, One-Class Support Vector Machine (OCSVM) [41], Isolation Forest (IF) [42], and Local Outlier Factor (LOF) [43]. Applying the occupancy and heating gates yielded 16,259 training timesteps, from which all standardization statistics and the reference flow rate were established. During online execution, daily morning buffer warm-up leaves 13,289 scored baseline timesteps for seasonal false-positive evaluation. The detectors were parameterized using OCSVM with $\nu = 0.05$, IF with 200 estimators, and LOF with 20 neighbors. Anomaly score thresholds were calibrated against the training score distributions, with the first-percentile threshold retained across all three fault types. The three models were selected to provide complementary decision boundaries by combining margin-based, tree-based, and density-based detectors, ensuring that different anomaly geometries in the feature space could be captured. In addition to the three individual models, a majority-vote ensemble of their thresholded outputs was included as a fourth candidate detector during post-hoc selection. The best combination of model, threshold, and temporal persistence filter was then selected independently for each fault type through a post-hoc optimization sweep.

3.4. Online detection, diagnosis, and compensation

The online stage embeds the trained models inside the EnergyPlus simulation loop using the Python API. At each timestep, the six features were reconstructed from the current sensor readings using a rolling buffer for the two-hour average. Inference proceeds only when the building is occupied, the heating system is active, and the buffer is fully populated. The reconstructed feature vector is evaluated by the selected detector, which is a majority vote across all three models for the stuck-closed and stuck-open faults or a single model for the supply-setpoint bias fault, to produce an anomaly score or a consensus flag. This output is compared with the optimized threshold to generate a binary detection flag. A temporal persistence filter requires N_{persist} consecutive flags before the diagnosis proceeds, suppressing transient false alarms. All detection components remain fixed during online operation. The feature scaler, the three anomaly models, the calibrated anomaly-score thresholds, and the persistence settings are loaded once at the start of each simulation and are not retrained, re-tuned, or otherwise adapted online.

Once the persistence threshold is met, the controller classifies the fault using physically motivated signatures evaluated in a fixed-priority order, as illustrated in Fig. 3. Beginning with a persistent anomaly trigger, the logic evaluates each signature sequentially. At the highest priority (P1), the stuck-closed fault is identified by a unique combination of low flow ($r_{\dot{m}} < 0.6$) and sustained underheating. If this signature does not match, the evaluation proceeds to the stuck-open fault (P2), which is characterized by overheating under normal flow. A dual-path check allows the rolling average to capture sustained overheating, whereas the instantaneous error provides a rapid path for diagnosis confirmation. The supply-setpoint bias fault, denoted as P3, is reached through two converging entry paths. These include a direct supply temperature depression check and multi-zone under-delivery confirmation, both of which require a normal flow. The scope distinction between zone- and system-level signatures separates the supply-setpoint bias diagnosis from the stuck-closed signature. If no signature matches (P4) after N_{retry} consecutive attempts, set to three timesteps, the fault is classified as unknown, and the detection streak is reset. Each successful classification also determines the corresponding compensation target temperature, linking the diagnosis directly to the corrective action.

Upon successful diagnosis, the compensation module ramps the supply temperature setpoint toward a fault-specific target as follows:

$$T_{\text{SP,new}} = T_{\text{SP,current}} + \alpha\left(T_{\text{target}} - T_{\text{SP,current}}\right), \qquad \alpha = 0.2, \tag{4}$$

where $T_{\text{target}} = 70$ °C for under-delivery faults and $T_{\text{target}} = 50$ °C for the stuck-open fault. The targets of 70 and 50 °C were chosen as bounded supervisory actions on either side of the 60 °C baseline supply temperature; they were not optimized and therefore represent predefined, bounded fixed-rule settings. The framework is therefore not formulated as an online multi-objective optimizer. Instead, comfort restoration, heating energy, and return-water temperature are evaluated jointly as outcome criteria. The sensitivity of the resulting trade-off to the fixed

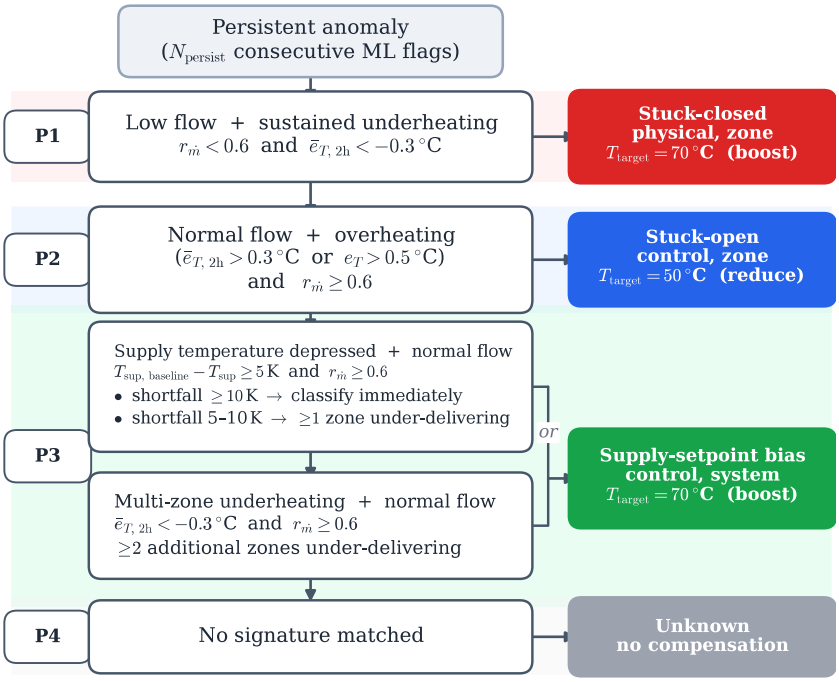

Fig. 3. Signature-based fault diagnosis ladder. Priority labels (P1–P4) indicate the evaluation order; unmatched signatures fall through to the next level.

targets and ramp gain is quantified in Section 4.5. For the stuck-open fault, the compensation additionally restores the nominal scheduled zone setpoint by removing the injected bias, so that the supervisory action simultaneously corrects the control fault and lowers the supply temperature. The same update drives the setpoint back to the baseline upon release. For the stuck-open and stuck-closed faults, compensation release requires $N_{\text{release}} = 6$ consecutive no-flag timesteps before the setpoint reverts to the baseline. The diagnosis at timestep N drives the compensation at timestep $N + 1$, reflecting the causal delay of a physically realizable controller.

The ramp gain α was selected heuristically rather than through formal optimization. In terms of the remaining gap to the target, $g_k = T_{\text{target}} - T_{\text{SP},k}$, the setpoint progression follows a discrete first-order decay with the closed-form solution

$$g_{k+1} = (1-\alpha)\, g_k, \qquad g_n = (1-\alpha)^n g_0 \tag{5}$$

For $\alpha = 0.2$ at the 10-minute timestep, approximately 90% of a requested setpoint change is completed within 110 min and 95% within 140 min. This bounded trajectory avoids abrupt step changes at the plant loop while achieving most of the adjustment within roughly two hours. To assess sensitivity to the ramp rate, alternative gains are evaluated alongside the reference setting.

For the supply-setpoint bias fault, successful compensation boosts the supply temperature above the depressed fault level to the 70 °C target, which also removes the under-supply signature that triggered detection. However, this corrected state can still trigger the detector. Because the elevated supply temperature lies outside the fault-free training distribution, the model may flag the compensation itself as an anomaly. Without a safeguard, the detector would release compensation, the depressed supply temperature would reassert, and the system would cycle between activation and release. To suppress compensation cycling under the supply-setpoint bias fault, an asymmetric release logic extends the clear-streak requirement to 12 consecutive timesteps (120 min) and maintains compensation across the active fault window. In the simulation testbed, this fault-window hold serves as an idealized benchmark to isolate the performance of continuous compensation from feedback-induced cycling, whereas a physical deployment would approximate this behavior using an observable substitute such as a duration-locked hold. This logic sustains a single uninterrupted compensation block throughout the remainder of each fault window.

3.5. Experimental design and performance indicators

The FDC framework was evaluated through four experimental campaigns. In the controlled scenario comparison, three scenarios were simulated for each fault type: fault-free baseline, faulty run with compensation disabled, and fully FDC-enabled run. All scenarios used identical weather data, occupancy schedules, and internal gains. A single baseline run served all the fault types. In the parametric severity sweep, each fault type was tested across five severity levels, with detection-only and compensation-enabled simulations per level. The primary evaluation considers isolated, abrupt faults under single-fault conditions. While healthy operational and weather variations are captured in the baseline training data, concurrent equipment or control faults are not. Under these nominal conditions, the fixed-priority diagnostic ladder produces a deterministic outcome, although this rule hierarchy does not ensure diagnostic uniqueness when multiple disturbances occur simultaneously. To evaluate behavior under concurrent disturbances, two compound-fault stress tests examine simultaneous faults in a diagnosis-only mode without engaging compensation. In addition, parameter sensitivity sweeps evaluate three compensation targets alongside three ramp gains ($\alpha \in \{0.1, 0.2, 0.4\}$) at reference severities. Across the baseline scenario, severity sweeps, parameter analyses, and compound-fault stress tests, the evaluation comprises 45 distinct EnergyPlus simulations in total.

All detection metrics are computed over gated timesteps, defined as occupied periods with active heating following feature warm-up. Discomfort degree-hours are summed across occupied fault-window timesteps, and delivered hydronic energy is evaluated over active heating periods.

Five thermal and energy outcome indicators were computed over the fault injection window. Discomfort degree-hours (DDH) quantify occupant discomfort as the cumulative absolute temperature deviation from the heating setpoint:

$$\text{DDH} = \sum_{t \in \mathcal{T}_{\text{fault}}} \left| e_T(t) \right| \cdot \Delta t_{\text{h}} \tag{6}$$

where $\mathcal{T}_{\text{fault}}$ is the set of occupied timesteps within the fault window and $\Delta t_{\text{h}} = 1/6$ h. The absolute-value formulation captures both underheating and overheating in a single metric. For comfort evaluation, $e_T(t)$ is computed relative to the nominal scheduled heating setpoint, $T_{\text{set,schedule}}$, rather than the fault-biased control setpoint used to inject the stuck-open fault; the overheating produced by the setpoint bias therefore registers as discomfort.

The delivered hydronic heating energy was estimated using the hot-water mass flow rate and temperature differential as follows:

$$E_{\text{hyd}} = \sum_{t \in \mathcal{T}_{\text{heat}}} \dot{m}(t)\, c_p\, \max\big(\Delta T_{\text{hw}}(t),\, 0\big)\, \frac{\Delta t_{\text{s}}}{3.6 \times 10^6} \qquad [\text{kWh}] \tag{7}$$

where $c_p = 4180$ J/(kg K), $\Delta T_{\text{hw}} = T_{\text{hw,inlet}} - T_{\text{hw,outlet}}$, $\Delta t_{\text{s}} = 600$ s, and $\mathcal{T}_{\text{heat}}$ is the set of heating-active timesteps within the fault window. This quantity represents the thermal energy transferred across the monitored hydronic branch serving the core retail zone, which covers the majority of the conditioned area. The reported return temperatures are likewise branch return temperatures, whose shifts indicate the direction, rather than the network-level magnitude, of the corresponding district heating return-temperature effects.

The average return water temperature is the flow-weighted mean of the hot water outlet temperature during active heating:

$$\bar{T}_{\text{ret}} = \frac{\sum_{t \in \mathcal{T}_{\text{heat}}} \dot{m}(t)\, T_{\text{hw,outlet}}(t)}{\sum_{t \in \mathcal{T}_{\text{heat}}} \dot{m}(t)} \tag{8}$$

Under normal flow conditions, lower return temperatures indicate better heat extraction and are beneficial for district heating network performance [5]. However, this interpretation requires that the mass flow rate is operationally meaningful; under severe flow restrictions, the return temperature signal may not reflect the actual network-side performance.

DDH recovery quantifies the fraction of fault-induced discomfort recovered by the FDC.

$$\text{Recovery}_{\text{DDH}} = \frac{\text{DDH}_{\text{faulty}} - \text{DDH}_{\text{FDC}}}{\text{DDH}_{\text{faulty}} - \text{DDH}_{\text{baseline}}} \times 100\ \% \tag{9}$$

A recovery of 100% indicates the full restoration of baseline comfort. Values exceeding 100% or below 0% are possible if the compensation over-corrects or worsens the conditions. The average zone air temperature $\bar{T}_{\text{zone}}$ during the fault window is reported as a complementary indicator of the thermal state experienced by occupants.

To evaluate relative changes during parameter sweeps, percentage shifts in delivered hydronic heating energy relative to the fault-free baseline and the uncompensated fault are defined as

$$\Delta E_{\text{base}} = 100\, \frac{E_{\text{FDC}} - E_{\text{baseline}}}{E_{\text{baseline}}}, \qquad \Delta E_{\text{fault}} = 100\, \frac{E_{\text{FDC}} - E_{\text{fault}}}{E_{\text{fault}}} \tag{10}$$

The corresponding effect on hydronic conditions is quantified by the shift in flow-weighted return-water temperature

$$\Delta T_{\text{ret}} = \bar{T}_{\text{ret,FDC}} - \bar{T}_{\text{ret,baseline}} \tag{11}$$

where positive values denote an increase relative to the reference scenario.

Detector accuracy is evaluated against known fault ground truth using standard binary classification metrics, including precision, recall, F1-score, and false-positive rate. Detection latency is measured using two conventions, representing the duration from fault onset to the first raw vote and to the persistence-confirmed alarm. Post-hoc optimization identified the detector model, threshold, and persistence length that maximized the F1-score while restricting the seasonal false-positive rate below 5%.

4. Results

The results are organized into seven subsections. Section 4.1 identifies the best detection configuration for each fault type. Section 4.2 compares the three scenarios for all three faults. Section 4.3 examines the time-resolved behavior of each fault with compensation. Section 4.4 synthesizes the supervisory recovery spectrum. Section 4.5 evaluates trade-offs between thermal comfort, heating energy, and return-water temperature across different compensation targets and ramp gains. Section 4.6 presents the parametric severity analysis and Section 4.7 evaluates the diagnostic framework under two scenarios with concurrent faults.

4.1. Detector selection

The post-hoc optimization swept 168 configurations per fault type, combining four detector variants consisting of the three individual models and their majority-vote ensemble, seven calibrated anomaly-score thresholds, and six temporal persistence lengths spanning one to six timesteps. Each configuration was evaluated on both the fault-injection and fault-free baseline logs, and the combination that maximized the F1-score subject to a false-positive rate below 5% was selected. Because detector parameters were selected and evaluated on the same fault records, the reported metrics reflect the best achievable benchmark across the evaluated configurations rather than performance on unseen data. Once established at the reference severity, each detector configuration was held fixed across all severity levels without further re-tuning. Table 4 lists the resulting configurations and detection performance metrics.

The results revealed that the optimal detector configuration was strongly fault-dependent. The stuck-closed valve fault produces a pronounced thermal signature by combining a sustained negative temperature error with an anomalously low flow. The majority-vote ensemble achieves an F1-score of 0.90 and detects this fault with perfect recall within 20 min of onset. Although LOF matches this recall on its own, it triggers a much higher false-positive rate during normal operation. The ensemble is therefore preferred because it preserves perfect detection without increasing false alarms across the full season.

The stuck-open fault presents a more diffuse anomaly pattern, as the overheating signal builds gradually and is distributed across multiple features rather than being concentrated in a single variable. Majority voting across all three models suppresses the higher individual false-positive rates, yielding an F1 of 0.86 with a detection latency of 100 min. Twelve configurations exceeded F1 = 0.80, suggesting a broader operating plateau than the stuck-closed case.

The supply-setpoint bias is the most challenging to detect. Because the depressed supply temperature affects all zones uniformly, the resulting thermal signatures overlap with the normal low-load conditions. The best configuration achieved a high recall but low precision, meaning that the fault was reliably detected but at the cost of frequent false alarms during normal operation. None of the 168 evaluated configurations exceeded an F1-score of 0.52. The reported detection latency of 550 min was measured from the midnight fault onset. Because the detector operates only during occupied hours with active heating, the effective latency from the first heated timestep is approximately 80 min.

From an operational standpoint, these results suggest that a single detection configuration cannot effectively serve all fault types. The tested zone-level physical fault produced a strong localized signature that enabled low-latency detection with high recall, whereas system-level control faults with diffuse thermal impacts overlap with normal operating conditions and achieve substantially lower precision. Despite this, the low-precision supply-setpoint bias detection still enables the compensation loop to engage because high recall ensures that the fault is reliably identified. In turn, the subsequent signature-based diagnosis filters out most false triggers prior to control actuation, even though occasional false actions during normal operation cannot be ruled out entirely. The per-fault-type detector selection adopted here serves a research characterization purpose by isolating the best achievable detection performance for each fault mechanism. An operational deployment would require either a universal detector configuration, a parallel

Table 4
Best-performing detector configuration and metrics per fault type.

Fault type	Model	$N_{persist}$	Precision	Recall	F1-score	FPR (%)	Confirmed latency (min)
Stuck-closed	Majority vote	3	0.82	1.00	0.90	0.18	20
Stuck-open	Majority vote	4	0.81	0.92	0.86	0.18	100
Supply-setpoint bias	LOF	6	0.36	0.98	0.52	2.14	550

Table 5
Scenario comparison results for the three fault types.

Fault type	Scenario	DDH (°C·h)	Delivered heat (kWh)	$\bar{T}_{ret}$ (°C)	$\bar{T}_{zone}$ (°C)	Recovery (%)
Stuck-closed	Baseline	0.07	710.8	54.8	22.00	N/A
	Faulty	9.86	627.5	46.8	21.45	N/A
	FDC	6.08	660.3	53.9	21.66	38.6
Stuck-open	Baseline	0.07	710.8	54.8	22.00	N/A
	Faulty	30.55	1029.1	54.1	23.69	N/A
	FDC	9.57	874.8	50.5	22.52	68.8
Supply-setpoint bias	Baseline	14.60	1316.9	54.5	21.25	N/A
	Faulty	21.02	1279.9	47.3	21.06	N/A
	FDC	17.35	1306.9	60.5	21.17	57.2

detector bank with arbitration logic, or an automated model selector, which remains an open design problem.

4.2. Scenario comparison results

Table 5 summarizes the key performance indicators for the three scenarios across all fault types. The single baseline run is shared across faults because the fault-free reference is independent of the fault that will be later injected. For zone-level faults, the fault window spans two occupied days, whereas the supply-setpoint bias window covers 48 h, including overnight periods; therefore, the absolute DDH values are not directly comparable across fault types. The supply-setpoint bias baseline DDH is substantially larger than the zone-level baseline because the 48-hour window captures early morning warm-up transients on both days, where the heating setpoint steps up before the heating system activates. During the daytime, the baseline DDH contribution was negligible. This reflects the scheduled setback recovery of the building rather than a modeling artifact.

The results reveal a fault-type dichotomy in terms of supervisory recovery potential. The two control-type faults achieved substantially higher DDH recovery than the physically constrained stuck-closed fault, confirming that the achievable recovery depends on whether the supervisory controller retains authority over the fault mechanism.

Among the two zone-level faults, the stuck-open fault produced the largest absolute discomfort impact and strongest compensation response. The FDC achieved a 68.8% DDH recovery, bringing the mean zone temperature from 1.69 °C above setpoint down to within 0.52 °C of baseline. The residual overheating reflects the detection transient during which no corrective action is active. Hydronic energy consumption drops by 15% relative to the uncompensated scenario, though it remains elevated above the fault-free baseline because the compensation must sustain a reduced supply temperature to counteract the biased setpoint. The compensated return temperature falls 4.2 K below the baseline value, a secondary benefit for district heating network performance arising from the reduced supply temperature.

The stuck-closed valve fault achieved 38.6% DDH recovery. The restricted flow path limits heat delivery regardless of how high the supply temperature is raised, and the supply temperature boost partially compensates for the hydraulic bottleneck. The mean zone temperature improved modestly but remained below the setpoint. The faulty scenario depresses the return temperature well below the baseline, and the FDC raises it closer to the baseline value by increasing the supply temperature entering the flow-restricted radiator. Under milder flow restrictions, where sufficient flow is maintained, the return temperature response may carry greater network-side significance. The energy deficit caused by flow restriction was only partially recovered.

The supply-setpoint bias achieved 57.2% DDH recovery, placing it between the two zone-level faults. The impact of this fault on comfort at the reference severity was moderate compared with that of zone-level faults. The principal value of this fault type in the evaluation lies in two features absent from the zone-level faults, specifically, its system-wide scope and the network-side trade-off introduced by the compensation. Once compensation is engaged, it effectively restores the supply temperature toward the target, bringing the zone temperature close to the setpoint during the compensated portion of the fault window. The aggregate recovery is limited primarily by the long detection transient, rather than by a fundamental physical constraint. However, the compensation introduces a network-side trade-off by raising the return temperature 6.0 K above the baseline value because the unrestricted hydraulic path allows the elevated supply temperature to drive excess heat through the radiators. This contrasts with the stuck-closed case, where the flow restriction naturally limits heat transfer and maintains return temperatures close to the baseline under compensation.

4.3. Fault-specific dynamic responses

The following subsections examine the time-resolved behavior of the FDC framework for each fault type by comparing the baseline, faulty, and compensated trajectories during the fault injection window.

4.3.1. Stuck-closed valve fault

Fig. 4 presents the dynamic response during the two-day stuck-closed fault. The fault clamped the hot-water mass flow rate to 0.59 kg/s at the start of each occupied period, against an unrestricted flow that reaches approximately 2.7 kg/s over the same window. Flow restriction is visible as an abrupt drop, where both the faulty and compensated traces fall to the same restricted level, confirming that the physical constraint is identical in both cases and that the FDC does not attempt to manipulate the flow path.

Immediately following fault onset at the start of occupancy, the anomaly scores exceeded their thresholds and triggered detection flags. Satisfying the three-step persistence filter enabled a confirmed stuck-closed diagnosis within 20 min, after which compensation was engaged at the subsequent timestep. This rapid response curtails the initial thermal drop, where the zone temperature briefly falls toward 20 °C before the supply-temperature boost begins to take effect.

The compensation module then ramps the supply temperature setpoint toward the 70 °C target. The ramp produces a rapid initial adjust-

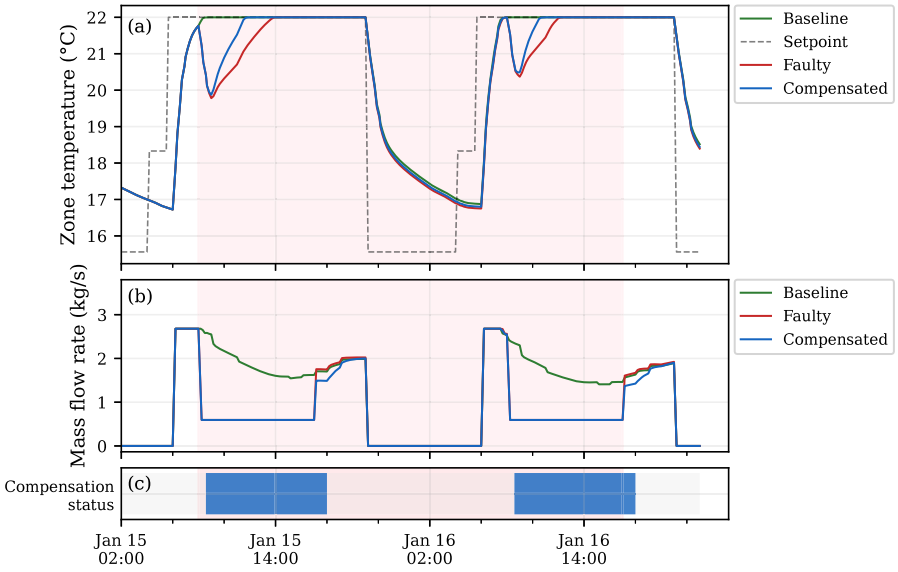

Fig. 4. Dynamic response during the stuck-closed valve fault at the reference severity ($\varphi = 0.30$) (a) zone air temperature with the heating setpoint, (b) mass flow rate, and (c) compensation state. The shaded bands mark the two occupied fault-injection periods.

ment, followed by asymptotic convergence, reaching the ceiling by mid-afternoon. The elevated supply temperature drives a larger temperature differential across the flow-restricted radiator, extracting more heat per unit flow, despite the reduced mass flow rate. The zone temperature recovers progressively, reaching the setpoint approximately 4.2 h after fault onset. The pattern repeated on January 16, with a slightly faster recovery, likely because the building's thermal mass retained some of the previous day's additional heat input.

A secondary effect emerged in the hydraulic domain. Under fault conditions without compensation, the return water temperature drops well below the baseline: the restricted mass flow reduces the total heat delivered to the zone, while the longer residence time through the emitter increases the temperature drop per unit mass, yielding a lower outlet temperature. When the supply temperature was elevated and the zone recovered to the setpoint, the return temperature rose to a steady-state average close to the baseline value during the afternoon. This is a counterintuitive outcome, where the FDC raises the return temperature closer to the baseline by increasing the supply temperature.

For stuck-closed faults, the FDC framework primarily functions as a diagnostic and triage tool rather than a full corrective controller. It correctly identifies the fault within 20 min and partially mitigates the comfort impact; however, the physical flow restriction limits the achievable recovery. In a real building, this diagnosis should trigger a maintenance dispatch to clear the valve obstruction, whereas the supply temperature boost provides interim relief.

4.3.2. *Stuck-open proxy fault*

Fig. 5 presents the dynamic response during the two-day stuck-open fault. A +2.0 °C bias was added to the occupied heating setpoint, raising it from 22.0 to 24.0 °C and inducing overheating. Without compensation, the zone temperature rises to approximately 24.0 °C and remains elevated throughout the fault window, while the hydronic mass flow rate stays high in response to the inflated heating demand.

The compensated response exhibits a qualitatively different dynamic from that of the stuck-closed case. Rather than a single detection–ramp–recovery sequence, the controller enters a sustained cycling pattern, visible as a sawtooth oscillation in panel (a) and as discrete activation blocks in panel (c).

After the fault was activated at 08:00, the zone temperature began to rise above the intended setpoint. As the overheating accumulates in the rolling two-hour features, the majority-vote anomaly detector asserts, and after satisfying the four-step persistence requirement, the signature-based diagnosis identifies a stuck-open condition and engages compensation. The compensated state simultaneously restores the effective heating setpoint to 22.0 °C and ramps the supply temperature downward toward the 50 °C compensation target. The zone temperature returned to approximately 22.0 °C within roughly two hours.

At this point, the controller largely removed the dominant thermal signature of the fault. With the temperature error near zero, the majority-vote trigger clears, and after six consecutive no-vote steps, the persistence filter releases compensation. The supply temperature then ramps back toward 60 °C, and the setpoint bias becomes effective again. The zone temperature climbs, typically reaching 22.5–23.3 °C before the detector re-triggers and a new compensation episode begins.

This engage–release–re-engage cycle repeats four to five times per occupied day, with each cycle lasting approximately 1.5–2 hours. Compensation is active for approximately 51% of the fault window timesteps, implying that for approximately half of the fault period, the supervisory correction is not engaged. The sawtooth oscillation in panel (a) directly reflects this.

This contrasts with the stuck-closed case, where compensation remained engaged for the fully occupied period. The difference is that the stuck-closed flow restriction persists regardless of the supply temperature adjustment, keeping the anomaly signal visible. For the stuck-open fault, successful compensation removed the overheating signature that triggered detection, causing the observed cycling.

Despite this cycling, the FDC achieved a 68.8% DDH recovery. The residual discomfort arises from the initial detection transient on each day, overheating spikes during each release interval, and thermal inertia that delays the temperature response after compensation re-engagement. Hydronic energy consumption dropped by 15% relative to the uncompensated scenario, although it remained elevated above the baseline because the release intervals temporarily re-enabled the elevated heating demand. The compensated return temperature fell below the baseline value, a network-side benefit produced by the reduced supply temperature during the active compensation phases.

The stuck-open fault is thus substantially recoverable through supervisory action, with the FDC recovering more than two-thirds of the fault-induced discomfort while reducing energy use and lowering the return temperature. A fixed-duration hold or fault-window-aware release guard could reduce cycling and likely improve recovery.

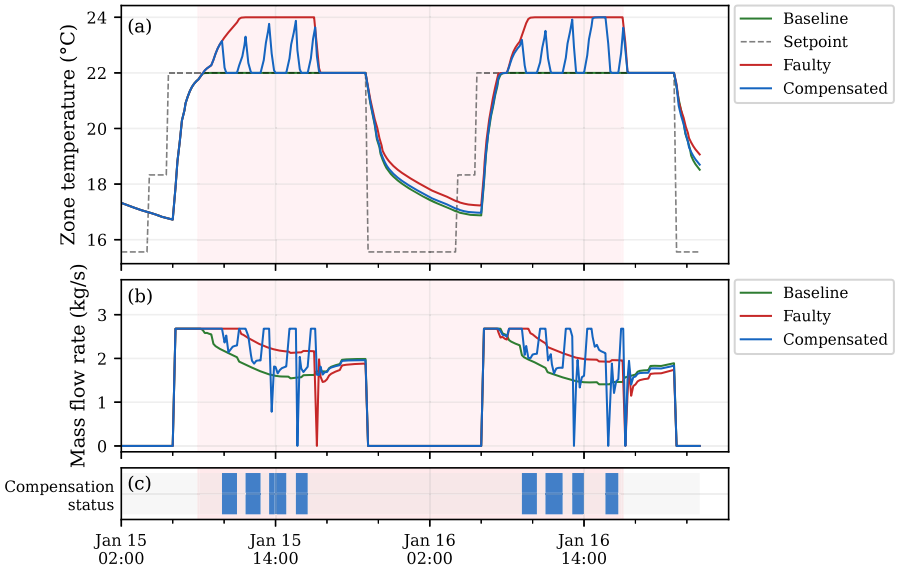

Fig. 5. Dynamic response during the stuck-open proxy fault at the reference severity ($\Delta T_{\text{bias}} = +2.0\,°\text{C}$) (a) zone air temperature with the heating setpoint, (b) mass flow rate, and (c) compensation state. The shaded bands mark the two occupied fault-injection periods.

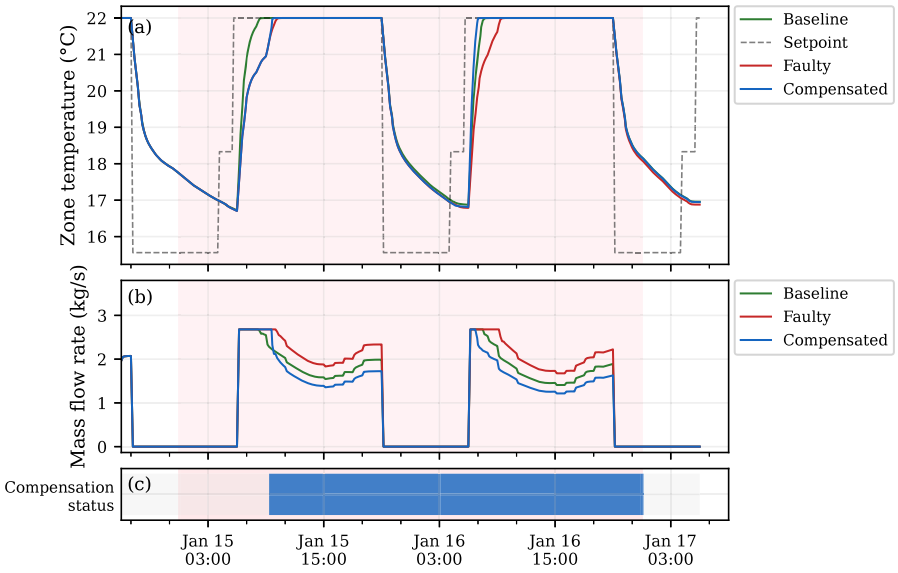

Fig. 6. Dynamic response during the supply-setpoint bias fault at the reference severity ($\Delta T_{\text{supply}} = -8\,\text{K}$) (a) zone air temperature with the heating setpoint, (b) mass flow rate, and (c) compensation state. The shaded band marks the 48-hour fault window.

4.3.3. *Supply-setpoint bias fault*

Fig. 6 presents the dynamic response during the 48-hour supply-setpoint bias fault. The fault depresses the supply water temperature from 60 to 52 °C, affecting the entire building simultaneously. Unlike zone-level faults, which alter either the flow path or demand signal in a single zone, this fault reduces the thermal driving force across all radiators.

The impact of the fault on comfort was moderate. Under the faulty scenario, the depressed supply temperature slowed the morning warm-up, with the most pronounced shortfall occurring between 06:00 and 10:00 when the thermal demand peaked. However, at 52 °C the supply temperature still provides sufficient capacity to maintain setpoint once the zone reaches steady state. The primary network-side consequence is a depressed return temperature, which reflects the reduced thermal driving force across the radiators.

The compensated response demonstrates the effect of the asymmetric release logic. The LOF detector flags the anomaly within 30 min of the first occupied timestep, and the persistence filter confirms the detection. Compensation is activated approximately 80 min after the start of occupied heating. The diagnosis identifies the fault as a supply-setpoint bias based on the concurrent conditions of a large supply temperature deviation and multi-zone under-delivery.

Once activated, the asymmetric release logic sustains compensation for the remainder of the fault window, preventing the cycling that characterizes the stuck-open response. Panel (c) confirms the existence of a single continuous compensation block. The supply temperature setpoint ramps toward 70 °C, and the zone temperature recovers to the setpoint within 20 min of compensation onset. On January 16, because compensation is already active from the overnight carry-over, the zone reaches setpoint approximately one hour earlier than the faulty scenario, eliminating the morning warm-up penalty entirely on the second day.

The aggregate DDH recovery was 57.2%. The unrecovered discomfort is concentrated in the first morning of the fault before compensa-

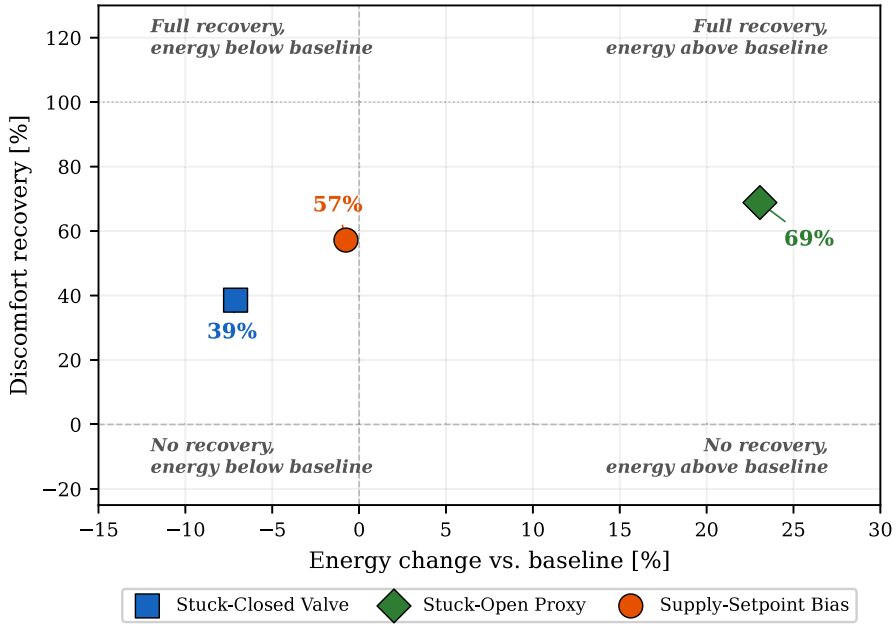

Fig. 7. Supervisory recovery spectrum across the three fault types at the reference severities, targets, and gain.

tion is engaged. However, this compensation introduces a temperature penalty. The elevated supply temperature drives a large temperature differential across radiators that are not flow-restricted, and the return water temperature rises above both the baseline and fault scenarios. For the supply-setpoint bias fault, the hydraulic path is unrestricted, and the full supply-temperature boost passes through the radiators with relatively low thermal resistance. From a district heating network perspective, this elevated return temperature is counterproductive, suggesting that a fault-specific compensation target lower than 70 °C would be more appropriate.

4.4. Supervisory recovery spectrum

The three fault-specific analyses revealed a supervisory recovery spectrum governed by two independent factors: the physical nature of the fault and the feedback interaction between compensation and detection. Fig. 7 maps each fault type in the two-dimensional space of DDH recovery versus energy change, relative to the fault-free baseline.

The stuck-open fault achieves the highest DDH recovery but incurs the largest energy penalty, placing it in the upper-right region of the spectrum. The energy cost arises because the setpoint bias persists during the release intervals of the compensation cycle. Each time the compensation disengages, the biased setpoint resumes, and the heating system returns to full demand until the detector re-triggers. The result is alternating periods of correction and unrestricted overheating, which collectively raise energy consumption above the level required for sustained compensation.

The supply-setpoint bias achieves 57.2% recovery at a near-zero energy penalty, placing it in the upper-left region. The extended release guard tied to the fault schedule eliminates cycling and enables steady-state compensation for most of the fault window. The residual discomfort was concentrated in the detection transient on the first morning rather than being distributed across repeated cycling intervals.

The stuck-closed valve fault achieved the lowest recovery and the largest below-baseline delivered heat, placing it in the lower-left region of the spectrum. This energy reduction occurs because the mechanical flow restriction limits heat delivery below nominal demand, an effect that the supply-temperature boost can only partially mitigate. This below-baseline heat use does not represent energy efficiency, but rather an energy deficit associated with unmet thermal load, indicating that comfort and energy outcomes must be evaluated jointly. The underlying physical constraint therefore imposes a recovery ceiling that supervisory setpoint adjustments cannot overcome.

Together, these three tested cases suggest a broader interpretation of supervisory recovery. The decisive factor is not simply whether a fault is physical or control-based but whether the supervisory controller retains sufficient authority over the thermal consequences of the fault. Two conditions must be met to ensure effective compensation. First, the controller must have a physical pathway to counteract the thermal effects of the fault. Second, the compensation must not mask the detection signal to the point of destabilizing sustained correction.

The supply-setpoint bias fault satisfies both conditions through unrestricted hydraulics and an extended release guard that maintains continuous compensation throughout the fault window. In this simulation environment, this scheduled hold serves as an idealized benchmark for the performance that a practical, duration-locked timer could achieve. This suggests a design principle for integrated FDC system development. Faults where compensation suppresses the anomaly signature should employ fixed-duration holds or timed release logic rather than relying on continuous anomaly re-detection. For physically constrained faults, the primary value of the framework shifts from corrective control to diagnostic triage, where the rapid detection latency and correct fault classification provide actionable information for maintenance dispatch, whereas the supply temperature boost offers partial interim comfort relief.

4.5. Compensation-parameter sensitivity

To make the trade-offs between thermal comfort, energy use, and network return temperature explicit, the compensation target was swept over three values per fault type at the reference severities: 60, 65, and 70 °C for the two under-delivery faults, and 50, 55, and 60 °C for the stuck-open proxy. All detection, diagnosis, and release settings remained unchanged. The performance of each run is evaluated relative to the fault-free baseline using DDH recovery, the change in delivered hydronic heating energy, and the return-water temperature shift ΔT_{ret}. Table 6 reports the numerical results and Fig. 8 maps both the target sweep and the ramp-gain sweep in the recovery–energy plane.

For the stuck-closed fault, recovery increases monotonically with the target. The 60 °C case represents a lower bound where commanded sup-

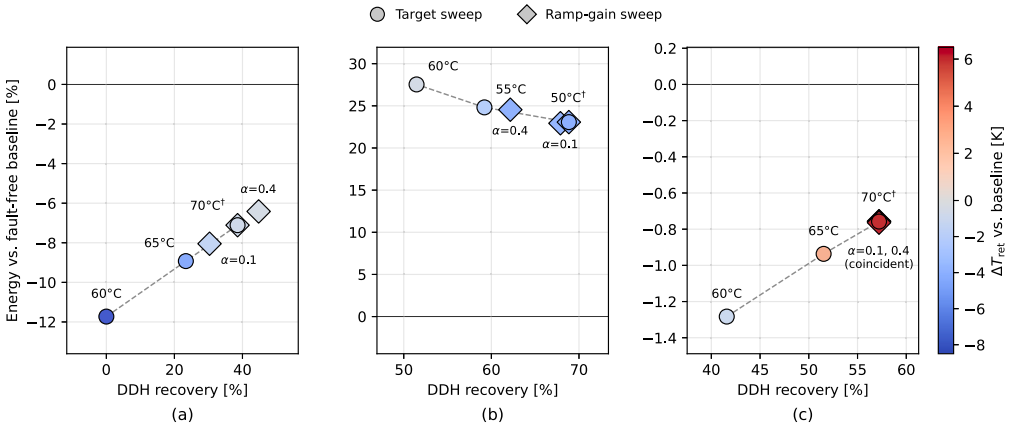

Fig. 8. Compensation-parameter sweep mapped in the recovery–energy plane for (a) stuck-closed valve, (b) stuck-open proxy, and (c) supply-setpoint bias. Marker color encodes the flow-weighted return-temperature shift relative to the baseline.

Table 6
Compensation-target sweep at reference severities.

Fault type	Target (°C)	Recovery (%)	ΔE_{base} (%)	ΔE_{fault} (%)	ΔT_{ret} (K)
Stuck-closed	60	0.0	−11.7	0.0	−8.0
	65	23.4	−8.9	+3.2	−4.5
	70†	38.6	−7.1	+5.2	−0.9
Stuck-open	50†	68.8	+23.1	−15.0	−4.2
	55	59.2	+24.8	−13.8	−2.2
	60	51.5	+27.5	−11.9	−0.6
Supply-setpoint bias	60	41.6	−1.3	+1.6	−1.1
	65	51.5	−0.9	+1.9	+2.5
	70†	57.2	−0.8	+2.1	+6.0

†Reference configuration

Table 7
Ramp-gain sweep at reference targets.

Fault type	α (–)	t_{90} (min)	Recovery (%)	ΔE_{base} (%)	ΔT_{ret} (K)
Stuck-closed	0.1	220	30.3	−8.0	−1.6
	0.2†	110	38.6	−7.1	−0.9
	0.4	50	44.8	−6.4	−0.5
Stuck-open	0.1	220	67.9	+22.9	−3.8
	0.2†	110	68.8	+23.1	−4.2
	0.4	50	62.2	+24.5	−4.1
Supply-setpoint bias	0.1	220	57.2	−0.8	+5.7
	0.2†	110	57.2	−0.8	+6.0
	0.4	50	57.2	−0.8	+6.1

†Reference configuration

ply equals baseline, reproducing the uncompensated fault trajectory. Delivered energy remains below baseline across all targets due to the hydraulic restriction rather than actual savings. Higher targets restore the depressed return temperature toward baseline, moving ΔT_{ret} from −8.0 to −0.9 K. Under a comfort-first criterion, the 70 °C reference is preferred because it yields the highest recovery and brings return temperature closest to baseline, whereas the lower energy use at lower targets simply reflects an unmet load deficit.

The stuck-open sweep separates the corrective action into setpoint restoration and supply-temperature modulation. At a 60 °C target, supply temperature is not reduced, and the controller acts through setpoint restoration alone. This action recovers 51.5% of fault-induced discomfort, representing roughly three-quarters of the 68.8% achieved at the 50 °C reference. Lowering the target to 50 °C contributes approximately 17 additional percentage points of recovery while reducing the energy penalty from +27.5% to +23.1% and deepening the return-temperature reduction from −0.6 to −4.2 K. The 50 °C reference setting therefore performs best across all three criteria. The number of engage–release cycles remains identical across targets, confirming that detection–compensation feedback governs cycling rather than the target setpoint.

For the supply-setpoint bias, the release guard maintained uninterrupted compensation across all targets, keeping the controller active for 81% of the fault window. Recovery rises from 41.6% to 51.5% and 57.2% at 60, 65, and 70 °C, while ΔT_{ret} shifts from −1.1 K to +2.5 K and +6.0 K. Energy change remains within 1.3% of baseline throughout. Target selection depends on the operational priority given to comfort recovery versus network return-temperature constraints. Restoring the supply to the 60 °C baseline eliminates return-temperature overshoot but slows morning warm-up recovery. In contrast, the intermediate 65 °C target captures roughly 90% of maximum comfort recovery while cutting return-temperature overshoot by more than half.

To complement the target sweep, the ramp gain was evaluated across $\alpha \in \{0.1, 0.2, 0.4\}$ at reference targets, corresponding to 90% convergence times of approximately 220, 110, and 50 min. Table 7 summarizes the results. Because both sweeps share the reference configuration ($\alpha = 0.2$ at the reference target), its circle and diamond markers overlap in Fig. 8.

The influence of the gain mirrors the compensation topology of each fault. For the stuck-closed fault, whose persistent hydraulic signature sustains uninterrupted compensation blocks, recovery increases monotonically with the gain, because a faster ramp delivers the supply-temperature boost earlier within each occupied fault period. For the stuck-open fault, the dependence is non-monotonic, and the reference gain is the best of the swept values. The number of engage–release cycles and the active fraction are unchanged across gains, so the differences arise from the ramp-back dynamics during the released intervals. At $\alpha = 0.1$, the supply temperature decays back toward the baseline so slowly that the released intervals retain a partial passive correction, whereas at $\alpha = 0.4$ the rapid rebound fully re-exposes the setpoint bias between cycles, degrading both comfort and energy. For the supply-setpoint bias, performance is largely insensitive to the ramp gain. Comfort recovery and energy change remain identical across all three values because unrecovered discomfort occurs primarily during the initial detection delay before compensation activates. The ramp speed only slightly affects the network side, shifting the return-temperature increase from 5.7 to 6.1 K.

The reference targets and ramp gain ($\alpha = 0.2$) are retained across subsequent evaluations to ensure consistent comparisons. Although the

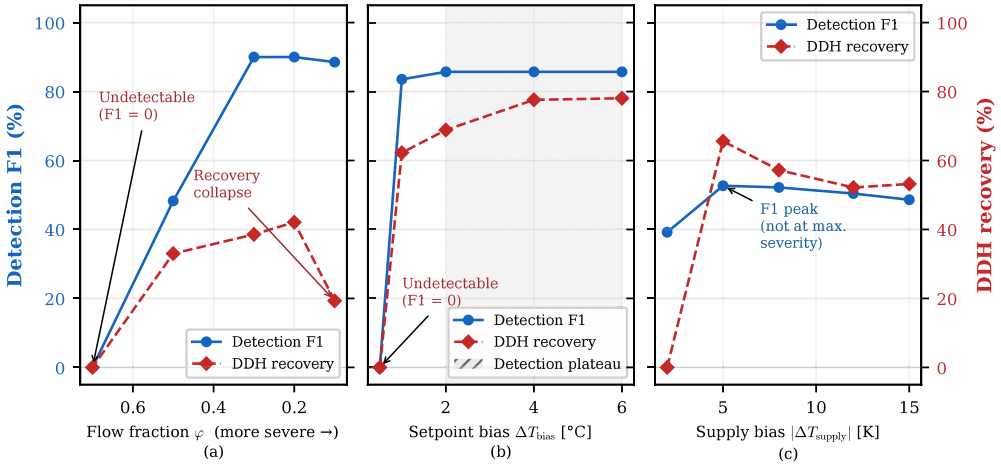

Fig. 9. Parametric severity sweep for (a) stuck-closed valve, (b) stuck-open proxy and (c) supply-setpoint bias.

stuck-closed fault benefits from a faster ramp, $\alpha = 0.2$ provides an effective compromise across all three fault types.

4.6. Fault severity sensitivity

The parametric severity sweep extends the single-point scenario comparison by evaluating the detection and compensation performance across five severity levels per fault type. Fig. 9 presents the results as three panels with the detection F1-score and DDH recovery plotted against increasing fault severity. Full numerical detection and compensation results are provided in Supplementary Tables S1 and S2, respectively.

Four key findings emerged from the severity analysis, each revealing a practical limitation of the FDC framework.

For the stuck-closed fault, the DDH recovery does not increase monotonically with severity. Recovery peaks at $\varphi = 0.20$ and collapses at $\varphi = 0.10$, which is the most severe restriction. At such extreme flow limitations, the supply temperature boost provides an insufficient additional thermal driving force. The recovery also falls to zero at $\varphi = 0.70$, but for a different reason. At that restriction the zone remains within 0.06 °C of its setpoint, the ensemble detector never confirms an alarm, and compensation is therefore never engaged. This is a detection floor rather than an actuation limit, and it mirrors the floor observed for the mildest stuck-open bias. The peak at $\varphi = 0.20$ is the best-performing tested severity, at which the interplay between fault impact and actuator authority is most favorable.

The mildest stuck-open bias (+0.5 °C) is entirely undetectable, as the overheating signal lies within the normal variability of the zone temperature and does not produce anomaly scores that exceed the threshold under any persistence setting. Compensation is never activated, and DDH recovery is zero. Detection becomes effective at +1.0 °C, establishing a practical detection threshold for this fault type.

Detection F1 for the supply-setpoint bias fault peaks at a moderate bias and declines at higher severity levels. This counterintuitive trend occurs because very large supply temperature depressions push the heating system into a regime where thermal residuals resemble cold-weather normal operation rather than a distinct fault signature. The fixed anomaly score threshold, calibrated for mild-to-moderate deviations, loses separability at extreme biases. The recall remained high across all severity levels, but the precision degraded because the detector could not distinguish the extreme-fault regime from the normal cold-load conditions. A similar detection floor exists at $|\Delta T_{\text{supply}}| = 2$ K, where the detector achieves only marginal performance and compensation does not engage; both become operational at $|\Delta T_{\text{supply}}| = 5$ K.

The stuck-open fault exhibited a detection saturation band, within which F1, precision, and recall were identical across consecutive severity levels. The stuck-closed fault instead shows a sharper transition, holding perfect recall down to $\varphi = 0.30$, falling to partial recall at $\varphi = 0.50$, and losing detection entirely at $\varphi = 0.70$. This plateau indicates that the detector reaches its performance ceiling within a moderate severity band and cannot differentiate between individual severity levels within that range. From a practical standpoint, this is acceptable because the system reliably detects all faults above the detection floor with equal confidence. However, the plateaus also imply that the current feature set and threshold calibration provide no severity estimation capability, a limitation relevant for maintenance prioritization, where distinguishing between different restriction levels would inform the urgency of intervention.

The severity curves refined the supervisory recovery spectrum established in the scenario comparison. The stuck-closed fault achieves its peak recovery at a moderate restriction level but collapses at the extremes, confirming that the physical compensation ceiling is most binding when the fault is most severe. The stuck-open fault is recoverable across a wide severity range, with recovery increasing monotonically and saturating at higher biases. The supply-setpoint bias fault achieved moderate recovery for biases above the detection threshold, with a mild decline at extreme biases, reflecting counterintuitive detection degradation.

4.7. Compound-fault stress tests

Two test scenarios evaluate the diagnostic framework under concurrent faults. In each scenario, two faults were injected simultaneously during the daytime heating window (January 15–16, 08:00–17:00). The system operated in a diagnosis-only mode, logging anomaly detection and signature classifications without activating compensation. Case C1 combines the stuck-closed restriction ($\varphi = 0.30$) with the supply-setpoint bias ($\Delta T_{\text{supply}} = -8$ K), representing two mutually reinforcing under-delivery mechanisms. In contrast, case C2 pairs the stuck-open bias ($\Delta T_{\text{bias}} = +2.0$ °C) with the same supply-setpoint bias, creating two partially canceling faults. Alongside the runtime ladder, all three diagnostic signatures were evaluated in parallel at each alarmed timestep. Each step was classified as a single supported signature, multiple agreeing signatures, conflicting signatures that inhibit compensation, or an unmatched anomaly. Table 8 summarizes detection and diagnostic results for both cases. Each compound scenario adopts the detector configuration calibrated for its constituent zone-level fault, with all diagnostic signature thresholds held constant. These stress tests characterize how the diagnostic ladder behaves under concurrent disturbances, evaluating whether overlapping hydraulic and thermal symptoms compromise classification.

Table 8
Compound-fault stress tests in diagnosis-only mode.

	Case C1 Stuck-closed + supply bias	Case C2 Stuck-open + supply bias
Detection		
Gated fault-window steps	108	108
Alarmed steps	108	83
Alarm coverage (%)	100.0	76.9
First raw vote (min)	0	170
isolated fault	0	70
Persistence-confirmed (min)	20	200
isolated fault	20	100
Runtime ladder verdicts (steps)		
Stuck-closed	104	–
Stuck-open	–	77
No verdict	4	6
Parallel signature evaluation (% of alarmed steps)		
Single signature supported	90.7	19.3
Agreeing signatures	–	–
Conflicting signatures	–	80.7
Unmatched signatures	9.3	–

Indented rows show isolated-fault latencies for comparison. Dashes denote states that did not occur.

In case C1, the two under-delivery faults compound thermally, deepening the average temperature deficit from 0.5 °C for the isolated restriction to 1.2 °C. Because this amplified deficit does not obscure the underlying low-flow signature, alarms remain prompt and continuous throughout the fault window, with the persistence-confirmed alarm arriving at the same 20 min latency as for the isolated restriction. The runtime ladder consistently classifies the condition as a stuck-closed fault on 104 of the 108 alarmed steps. In the parallel per-step evaluation, stuck-closed is likewise the sole supported signature on 90.7% of alarmed steps, while the remaining 9.3% match no signature. The supply-setpoint bias is never independently supported because its rule strictly requires normal flow. Consequently, the diagnosis is incomplete but action-consistent, as both constituent faults call for a supply-temperature boost.

Case C2 represents the opposite regime, where the two mechanisms partially cancel. This mutual masking moderates the zone overheating and substantially delays initial detection. Once alarms trigger, the fixed-priority ladder commits to stuck-open on 77 of 83 alarmed steps because the positive temperature error satisfies an earlier priority rule before the depressed supply signature is evaluated. Acting on this verdict would mistakenly command a supply-temperature reduction when the supply is already 8 K below baseline. Parallel evaluation mitigates this failure mode by revealing that opposing actions (reduce versus boost) are supported simultaneously across 80.7% of alarmed timesteps, which would invoke a safety hold in a compensation-enabled implementation. Only 19.3% of steps support a single signature, occurring before the secondary supply deficit is confirmed, while no timesteps exhibited agreeing signatures. Because a supervisory controller would still execute an incorrect action during this initial window, parallel arbitration alone cannot guarantee diagnostic uniqueness.

Outside the fault windows, the signature layer would have blocked most false compensation. Across both runs, the diagnostic logic dismissed nearly 88% of out-of-window alarms because no physical signature was matched. Only isolated timesteps satisfied the supply-bias criteria, demonstrating that downstream signature verification substantially curtails unwarranted control actions without eliminating false interventions entirely. Overall, the stress tests depict the single-fault operational scope of the framework. While reinforcing faults maintain action-consistent control decisions, partially canceling faults represent the primary vulnerability of fixed-priority ladders. Parallel signature evaluation identified conflicting actions on 80.7% of the alarmed timesteps in case C2. In a compensation-enabled deployment, these conflicting states could withhold action to reduce the risk of erroneous compensation without requiring additional sensors or models. However, an arbitration layer alone cannot guarantee diagnostic uniqueness because 19.3% of the alarmed steps support only the stuck-open signature before the supply-bias signature becomes observable. Practical systems would therefore require conflict checking alongside a brief confirmation delay before altering setpoints, leaving compound-fault compensation to future work.

5. Discussion

The results establish that the achievable supervisory recovery is governed by two conditions: the controller must have a physical pathway to counteract the fault, and compensation must not destabilize its own detection signal. The three fault types span a spectrum across these requirements. The stuck-closed valve violates the first condition, as mechanical flow restriction imposes a strict hydraulic ceiling that supply-temperature elevation cannot breach; recovery peaks at moderate restriction and collapses under severe throttling. The stuck-open fault satisfies the physical actuation condition but destabilizes detection, where successful compensation suppresses the overheating signature and triggers feedback cycling that caps recovery. The supply-setpoint bias fault avoids this feedback instability through an extended hold across the active fault window, sustaining continuous compensation. Consequently, release logic design is as consequential for supervisory recovery as the detection algorithm itself. Tailoring release logic to fault mechanisms can improve recovery without elevating false-positive risks, addressing a central challenge in active fault-tolerant control [44].

5.1. Operational implications for fault-specific triage

For control-type faults, the supervisory controller can provide immediate autonomous relief until maintenance is conducted. However, this requires a deployable release strategy that operates without knowing the fault window. For physically constrained faults, the framework functions primarily as a diagnostic tool; rapid classification provides actionable data for priority maintenance dispatch, while supply-temperature boosting offers partial, interim comfort relief. The return-temperature overshoot observed for the supply-setpoint bias fault demonstrates that generic compensation targets are not appropriate for all fault types, as restoring supply temperature to its baseline rather than an arbitrary maximum eliminates this penalty. The target sweep quantifies this recommendation. Restoring the commanded supply temperature to its 60 °C baseline avoids return temperature overshoot ($\Delta T_{ret} = -1.1$ K) at 41.6% recovery, whereas an intermediate 65 °C target retains 51.5% recovery, while reducing the overshoot from +6.0 to +2.5 K. These findings provide a quantitative basis for structured substation fault prioritization, addressing an operational need recognized independently [45].

5.2. Performance contextualization relative to prior work

A direct numerical comparison with prior fault-handling studies is not feasible because the literature is split across fundamentally different system domains, evidence bases, and evaluation metrics. As mapped in Table 9, prior work bounds the present contribution within two distinct peer families: hydronic and district-heating studies that stop at diagnosis or advisory intervention [27,28,30], and automated-control studies conducted outside district heating [31,35,36]. While the latter demonstrate automated correction in air-side HVAC or standalone loops, their metrics (such as energy cost savings or schedule overrides) cannot be compared directly with thermal comfort recovery or return-water temperatures.

Against this literature landscape, the present framework provides a closed-loop benchmark specifically designed for DH-connected build-

Table 9
Positioning of the framework against representative fault-handling literature.

Study	Domain & testbed	Corrective action	Primary outcome
Hydronic and district heating systems			
Djuric et al. [27]	DH building (field)	None (advisory)	Quantified fault energy/comfort penalties
Sarran et al. [28]	DH floor heating (field)	Advisory setpoint reset	Simulated 8.6/6.5 K supply/return reductions
Håkansson et al. [30]	Radiator system (field)	Advisory rebalancing	Reduced zone temperature variation
Automated correction and fault-tolerant control			
Pritoni et al. [31]	Air-side HVAC (field)	Auto-correction routines	Field-corrected overrides and reset schedules
Masdoua et al. [35]	Heating loop (simulated)	Active FTC (RL agent)	Restored zone comfort under actuator faults
Kim and Kim [36]	Air-handling unit (HIL)	Transactive setpoint	10% energy cost savings under stuck damper
Proposed framework			
Present study	DH building (simulated)	Closed-loop supervisory	38.6–68.8% comfort recovery; ΔT_{ret} evaluated

ings. Evaluating delivered hydronic energy alongside return-water temperatures reveals system-level trade-offs omitted in air-side literature, showing that supervisory compensation can either reduce network return temperatures or impose thermal penalties depending on the underlying fault mechanism.

5.3. Toward safe online adaptation

All detection models, thresholds, and persistence filters remain frozen during online operation. Adapting the detector while a fault may be active risks absorbing the anomalous signature into the baseline reference, progressively blinding the system. However, operating with fixed parameters exposes seasonal sensitivities. A monthly baseline analysis indicates that false-positive rates concentrate during peak winter heating, reaching approximately 8% for the LOF single-model configuration while the majority-vote ensemble remains at or below 1% across all months. Because fixed thresholds cannot distinguish extreme seasonal loads from anomalies, guarded adaptation is warranted.

Rather than adapting online, detector updates should occur through controlled offline stages to maintain operational stability. Retraining data should be collected only during confirmed normal operating periods without active alarms, supervisory compensation, or pending maintenance tasks. Candidate models can then be benchmarked offline against historical fault records, replacing the active detector only if they match or improve detection accuracy, with previous models archived for immediate rollback. Applied on a seasonal schedule, this guarded process could reduce winter false alarms, lower detection limits, and account for gradual drift from equipment aging, envelope retrofits, and occupancy changes. This architecture is outlined to guide practical implementation, but has not yet been tested experimentally.

5.4. Limitations and future work

The evaluation entails several idealized simulation assumptions, including deterministic schedules, noiseless sensing, an ideal district heating connection, and single-building training without cross-building transfer. Baseboard heaters also serve as radiator surrogates, which may not capture the full thermal inertia and convection dynamics of older panel or cast-iron radiators. The normalization reference $\dot{m}_{ref}$ is inherited from the preliminary training season, so the flow-ratio feature is not a bounded fraction. Furthermore, the supply-setpoint bias release logic relies on a simulation-observable fault hold, meaning the reported 57.2% recovery represents an upper bound on what causal release could achieve. The principal campaigns examine single, abrupt faults; compound-fault tests serve only as an initial diagnosis-only probe, while gradual faults remain unexamined. The stuck-open fault is likewise implemented as a setpoint-bias proxy rather than a hydraulic valve failure, which may alter part-load flow signatures relative to a mechanically seized valve. Furthermore, detector configurations were selected post hoc with frozen thresholds, and compensation targets are predefined, bounded fixed-rule settings rather than optimized policies. The reported metrics therefore reflect favorable post-hoc selections, while compensation results provide a non-optimized benchmark that fault-specific tuning could improve upon. Despite these boundaries, the 45 simulations across severity sweeps and compound stress tests support the qualitative recovery dynamics and operational ceilings. Given the disparity between simulation and field environments [46], empirical validation in real buildings remains essential [47].

The future directions map directly onto these limitations. Field validation in instrumented district heating substations is the most critical step to assess sensor noise and unmodeled hydraulic dynamics. Addressing compensation cycling under over-delivery faults motivates developing fixed-duration or timer-based release logic that operates without simulation oracles. Extending the architecture to multi-zone detection and zone-specific compensation will facilitate deployment in larger buildings. Guarded offline retraining offers a practical path for adaptive threshold calibration to track operational drift. In addition, incorporating parallel signature evaluation with an explicit abstain state and a temporal confirmation hold will resolve conflicting symptoms during concurrent faults.

6. Conclusions

This study developed and evaluated an integrated fault detection and compensation framework that embeds unsupervised machine learning inside the EnergyPlus simulation loop, enabling online detection to drive signature-based diagnosis and supervisory setpoint compensation in district-heated buildings. The findings demonstrate that optimal detection is inherently fault-dependent. Localized physical faults produce distinct thermal signatures that enable rapid detection with high accuracy, whereas system-level control faults generate diffuse deviations that overlap with normal operating conditions, requiring multi-model consensus or multi-zone verification to avoid false alarms.

Achievable supervisory recovery is governed by hydraulic actuator authority and feedback stability between compensation and detection. While control-type faults are substantially recoverable, mechanical flow restrictions impose an unbreachable hydraulic ceiling. Correcting over-delivery faults suppresses the underlying anomaly signature, inducing feedback cycling that caps net recovery unless release logic is decoupled from instantaneous alarms. Furthermore, parametric and sensitivity sweeps show that supervisory compensation involves inherent trade-offs, where boosting supply temperatures restores comfort but risks elevating network return-water temperatures, while severe restrictions cause recovery to collapse. Under concurrent disturbances, parallel signature evaluation with an explicit abstain state and temporal confirmation is necessary to prevent conflicting control actions.

Within the tested scenarios, these findings support a fault-specific operational strategy. Selected control-type faults are viable candidates for immediate supervisory mitigation to reduce discomfort and unneces-

sary energy use while maintenance is scheduled, subject to deployable release logic and safety validation. Conversely, physically constrained faults require priority maintenance dispatch, with setpoint modulation providing only partial, interim relief. Supervisory control and maintenance intervention are complementary measures; distinguishing between their operational boundaries allows building operators and district heating utilities to manage customer-side faults without compromising network efficiency.

CRediT authorship contribution statement

Nima Monghasemi: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Stavros Vouros:** Writing – review & editing, Supervision; **Konstantinos Kyprianidis:** Writing – review & editing, Supervision, Conceptualization; **Amir Vadiee:** Writing – review & editing, Supervision, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used Paperpal (https://paperpal.com/) to improve the readability, grammar, and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Funding

This work was supported by the Swedish Energy Agency (project No. 52686-1).

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Amir Vadiee reports financial support was provided by Swedish Energy Agency. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors extend their sincere appreciation to the project developers affiliated with the NEEDs project (innovative approaches for energy efficiency improvement for positive or self-balanced districts; project No. 52686-1).

Data availability

The original data presented in the study are openly available in Zenodo at https://doi.org/10.5281/zenodo.21869564, with the development repository at https://github.com/NimaMngh/hydronic-fdc-energyplus (accessed on 3 September 2026).

Supplementary material

Supplementary material associated with this article can be found in the online version at 10.1016/j.rineng.2026.113471.

References

[1] S. Yao, J. Wu, M. Qadrdan, A state-of-the-art analysis and perspectives on the 4th/5th generation district heating and cooling systems, Renew. Sustain. Energy Rev. 202 (2024) 114729. https://doi.org/10.1016/j.rser.2024.114729

[2] L.M. Dang, L.Q. Nguyen, J. Nam, T.N. Nguyen, S. Lee, H.-K. Song, H. Moon, Fifth generation district heating and cooling: a comprehensive survey, Energy Rep. 11 (2024) 1723–1741. https://doi.org/10.1016/j.egyr.2024.01.037

[3] E. Guelpa, M. Capone, A. Sciacovelli, N. Vasset, R. Baviere, V. Verda, Reduction of supply temperature in existing district heating: a review of strategies and implementations, Energy 262 (2023) 125363. https://doi.org/10.1016/j.energy.2022.125363

[4] M. Capone, E. Guelpa, V. Verda, Exploring opportunities for temperature reduction in existing district heating infrastructures, Energy 302 (2024) 131871. https://doi.org/10.1016/j.energy.2024.131871

[5] H.İ. Tol, H.B. Madessa, Return-temperature reduction at district heating systems: focus on end-user sites, Energies 17 (19) (2024) 4901. https://doi.org/10.3390/en17194901

[6] P. Lauenburg, J. Wollerstrand, Adaptive control of radiator systems for a lowest possible district heating return temperature, Energy Build. 72 (2014) 132–140. https://doi.org/10.1016/j.enbuild.2013.12.011

[7] T. Li, W. Chen, X. Li, B. Wang, W. Shi, Investigation of control strategies for dual-temperature district heating substations with two absorption heat pumps and two heat exchangers, Energy Convers. Manag. 309 (2024) 118431. https://doi.org/10.1016/j.enconman.2024.118431

[8] Q. Yang, R. Salenbien, K.M. Smith, M. Tunzi, Identifying untraced faults associated with high return temperatures from heating systems in buildings connected to district heating networks, Energy 309 (2024) 133097. https://doi.org/10.1016/j.energy.2024.133097

[9] S. Månsson, P.-O. Johansson Kallioniemi, M. Thern, T. Van Oevelen, K. Sernhed, Faults in district heating customer installations and ways to approach them: experiences from Swedish utilities, Energy 180 (2019) 163–174. https://doi.org/10.1016/j.energy.2019.04.220

[10] Z. Chen, Z. O'Neill, J. Wen, O. Pradhan, T. Yang, X. Lu, G. Lin, S. Miyata, S. Lee, C. Shen, R. Chiosa, M.S. Piscitelli, A. Capozzoli, F. Hengel, A. Kührer, M. Pritoni, W. Liu, J. Clauß, Y. Chen, T. Herr, A review of data-driven fault detection and diagnostics for building HVAC systems, Appl. Energy 339 (2023) 121030. https://doi.org/10.1016/j.apenergy.2023.121030

[11] J. van Dreven, V. Boeva, S. Abghari, H. Grahn, J. Al Koussa, E. Motoasca, Intelligent approaches to fault detection and diagnosis in district heating: current trends, challenges, and opportunities, Electronics 12 (6) (2023) 1448. https://doi.org/10.3390/electronics12061448

[12] M. Neumayer, D. Stecher, S. Grimm, A. Maier, D. Bücker, J. Schmidt, Fault and anomaly detection in district heating substations: a survey on methodology and data sets, Energy 276 (2023) 127569. https://doi.org/10.1016/j.energy.2023.127569

[13] IEA DHC Annex TS4, Guidebook for the Digitalisation of District Heating: Transforming Heat Networks for a Sustainable Future, Final Report of DHC Annex TS4, Technical Report, International Energy Agency Technology Collaboration Programme on District Heating and Cooling, Frankfurt am Main, 2023. ISBN 978-3-89999-118-9, https://www.iea-dhc.org/the-research/annexes/2018-2024-annex-ts4

[14] S. Månsson, M. Thern, P.-O. Johansson Kallioniemi, K. Sernhed, A fault handling process for faults in district heating customer installations, Energies 14 (11) (2021) 3169. https://doi.org/10.3390/en14113169

[15] H. Gadd, S. Werner, Fault detection in district heating substations, Appl. Energy 157 (2015) 51–59. https://doi.org/10.1016/j.apenergy.2015.07.061

[16] D. Leiria, H. Johra, J. Anoruo, I. Praulins, M.S. Piscitelli, A. Capozzoli, A. Marszal-Pomianowska, M.Z. Pomianowski, Is it returning too hot? Time series segmentation and feature clustering of end-user substation faults in district heating systems, Appl. Energy 381 (2025) 125122. https://doi.org/10.1016/j.apenergy.2024.125122

[17] M.B. Soultanzadeh, M. Nik-Bakht, M.M. Ouf, P. Paquette, S. Lupien, Unsupervised automated fault detection and diagnosis for light commercial buildings' HVAC systems, Build. Environ. 267 (2025) 112312. https://doi.org/10.1016/j.buildenv.2024.112312

[18] J.-G. Lee, D.-H. Kim, J.H. Lee, Proactive fault diagnosis of a radiator: a combination of Gaussian mixture model and LSTM autoencoder, Sensors 23 (21) (2023) 8688. https://doi.org/10.3390/s23218688

[19] C. Hermans, J. Al Koussa, T. Van Oevelen, D. Vanhoudt, Fault detection for district heating substations: beyond three-sigma approaches, Smart Energy 16 (2024) 100159. https://doi.org/10.1016/j.segy.2024.100159

[20] J. van Dreven, A. Cheddad, S. Alawadi, A.N. Ghazi, J. Al Koussa, D. Vanhoudt, SHEDAD: SNN-enhanced district heating anomaly detection for urban substations, in: Proceedings of the 9th International Conference on Fog and Mobile Edge Computing (FMEC 2024), Malmö, Sweden, 2024, pp. 130–137. https://doi.org/10.1109/FMEC62297.2024.10710205

[21] F. Zhang, N. Saeed, P. Sadeghian, Deep learning in fault detection and diagnosis of building HVAC systems: a systematic review with meta analysis, Energy AI 12 (2023) 100235. https://doi.org/10.1016/j.egyai.2023.100235

[22] B. Feng, Q. Zhou, J. Xing, Q. Yang, Y. Chen, Z. Deng, Attention-empowered transfer learning method for HVAC sensor fault diagnosis in dynamic building environments, Build. Environ. 250 (2024) 111148. https://doi.org/10.1016/j.buildenv.2023.111148

[23] L. Chen, G. Li, J. Liu, L. Liu, C. Zhang, J. Gao, C. Xu, X. Fang, Z. Yao, Fault diagnosis for cross-building energy systems based on transfer learning and model interpretation, J. Build. Eng. 91 (2024) 109424. https://doi.org/10.1016/j.jobe.2024.109424

[24] H. Bahlawan, N. Ferraro, A. Gambarotta, E. Losi, L. Manservigi, M. Morini, C. Saletti, P.R. Spina, M. Venturini, Detection and identification of faults in a district heating network, Energy Convers. Manag. 266 (2022) 115837. https://doi.org/10.1016/j.enconman.2022.115837

[25] E. Losi, L. Manservigi, P.R. Spina, M. Venturini, Data-driven approach for the detection of faults in district heating networks, Sustain. Energy Grids Netw. 38 (2024) 101355. https://doi.org/10.1016/j.segan.2024.101355

[26] M. Rätz, R. Streblow, D. Müller, Severe fault diagnosis in digitized district heating network substations: a machine learning regression approach, in: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems (e-Energy '24), Singapore, 2024, pp. 555–556. https://doi.org/10.1145/3632775.3661971

[27] N. Djuric, V. Novakovic, F. Frydenlund, Heating system performance estimation using optimization tool and BEMS data, Energy Build. 40 (8) (2008) 1367–1376. https://doi.org/10.1016/j.enbuild.2007.12.006

[28] L. Sarran, K.M. Smith, C.A. Hviid, C. Rode, Grey-box modelling and virtual sensors enabling continuous commissioning of hydronic floor heating, Energy 261 (2022) 125282. https://doi.org/10.1016/j.energy.2022.125282

[29] IEA DHC Annex TS9, Digitalisation of district heating and cooling: improving efficiency and performance through data integration, 2024, (IEA Technology Collaboration Programme on District Heating and Cooling). Project period 2024–2028. Accessed 2 August 2026, https://www.iea-dhc.org/the-research/annexes/2024-2028-annex-ts9

[30] H. Håkansson, M. Önnheim, J. Sjöberg, M. Jirstrand, Model-assisted hydronic balancing in residential heating systems using operational sensor data, Energy Build. 348 (2025) 116464. https://doi.org/10.1016/j.enbuild.2025.116464

[31] M. Pritoni, G. Lin, Y. Chen, R. Vitti, C. Weyandt, J. Granderson, From fault-detection to automated fault correction: a field study, Build. Environ. 214 (2022) 108900. https://doi.org/10.1016/j.buildenv.2022.108900

[32] A.A. Markus, J. Bursill, H.B. Gunay, B.W. Hobson, Projecting and estimating HVAC energy savings from correcting control faults: comparison between physical and virtual metering approaches, Energy Build. 328 (2025) 115169. https://doi.org/10.1016/j.enbuild.2024.115169

[33] D. Stecher, M. Neumayer, A. Ramachandran, A. Hort, A. Maier, D. Bücker, J. Schmidt, Creating a labeled district heating data set: from anomaly detection towards fault detection, Energy 313 (2024) 134016. https://doi.org/10.1016/j.energy.2024.134016

[34] Z. Du, X. Jin, Tolerant control for multiple faults of sensors in VAV systems, Energy Convers. Manag. 48 (3) (2007) 764–777. https://doi.org/10.1016/j.enconman.2006.09.007

[35] Y. Masdoua, M. Boukhnifer, K.H. Adjallah, Active fault-tolerant control based on DDQN architecture applied to HVAC system, Trans. Inst. Meas. Control 47 (16) (2025) 3303–3312. https://doi.org/10.1177/01423312241273767

[36] I. Kim, W. Kim, Fault detection and diagnosis and fault-tolerant control for air-handling unit via agent-based hardware-in-the-loop simulation, J. Build. Eng. 111 (2025) 113251. https://doi.org/10.1016/j.jobe.2025.113251

[37] N. Monghasemi, A. Vadiee, S. Vouros, K. Kyprianidis, Radiator fault detection in a multi-energy source building using unsupervised learning techniques, in: Proceedings of the 37th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2024), Rhodes, Greece, 2024, pp. 2470–2481. https://doi.org/10.52202/077185-0213

[38] N. Monghasemi, A. Vadiee, Evaluation of an integrated fault detection and compensation for building-level radiator systems in district heating, in: Proceedings of the 17th International Conference on Applied Energy (ICAE2025), 66, Bangkok, Thailand, 2025, p. 12258. https://doi.org/10.46855/energy-proceedings-12258

[39] K. Field, M. Deru, D. Studer, Using DOE commercial reference buildings for simulation studies, in: Proceedings of SimBuild 2010: 4th Conference of IBPSA-USA, New York, NY, USA, 2010, pp. 85–93. NREL/CP-550-48588, https://publications.ibpsa.org/proceedings/simbuild/2010/papers/simbuild2010_TS02B-02-Field.pdf.

[40] D.B. Crawley, L.K. Lawrie, F.C. Winkelmann, W.F. Buhl, Y.J. Huang, C.O. Pedersen, R.K. Strand, R.J. Liesen, D.E. Fisher, M.J. Witte, J. Glazer, EnergyPlus: creating a new-generation building energy simulation program, Energy Build. 33 (4) (2001) 319–331. https://doi.org/10.1016/S0378-7788(00)00114-6

[41] B. Schölkopf, J.C. Platt, J. Shawe-Taylor, A.J. Smola, R.C. Williamson, Estimating the support of a high-dimensional distribution, Neural Comput. 13 (7) (2001) 1443–1471. https://doi.org/10.1162/089976601750264965

[42] F.T. Liu, K.M. Ting, Z.-H. Zhou, Isolation-based anomaly detection, ACM Trans. Knowl. Discov. Data 6 (1) (2012) 3:1–3:39. https://doi.org/10.1145/2133360.2133363

[43] M.M. Breunig, H.-P. Kriegel, R.T. Ng, J. Sander, LOF: identifying density-based local outliers, in: Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data, Dallas, TX, USA, 2000, pp. 93–104. https://doi.org/10.1145/335191.335388

[44] Y. Zhang, J. Jiang, Bibliographical review on reconfigurable fault-tolerant control systems, Annu. Rev. Control 32 (2) (2008) 229–252. https://doi.org/10.1016/j.arcontrol.2008.03.008

[45] E. Guevara Bastidas, S. Faulstich, H. Dittmer, M. Neumayer, G. Sakthivel Mohan, K. Sercan-Calismaz, F. Hosenfelder, T. Glenewinkel, K. Fischer-Florschütz, A. Cadenbach, Prioritisation of faults in district heating substations: towards predictive maintenance and optimised operation, Energy 333 (2025) 137210. https://doi.org/10.1016/j.energy.2025.137210

[46] J. Huang, J. Wen, H. Yoon, O. Pradhan, T. Wu, Z. O'Neill, K.S. Candan, Real vs. simulated: questions on the capability of simulated datasets on building fault detection for energy efficiency from a data-driven perspective, Energy Build. 259 (2022) 111872. https://doi.org/10.1016/j.enbuild.2022.111872

[47] Y. Lei, C. Fan, H. He, Y. Xie, Enabling efficient cross-building HVAC fault inferences through novel unsupervised domain adaptation methods, Build. Environ. 272 (2025) 112678. https://doi.org/10.1016/j.buildenv.2025.112678