sustainability

MDPI

Article

Hazard- and Fairness-Aware Evacuation with Grid-Interactive Energy Management: A Digital-Twin Controller for Life Safety and Sustainability

Mansoor Alghamdi [1,*], Ahmad Abadleh [1,*], Sami Mnasri [1,*], Malek Alrashidi [1], Ibrahim S. Alkhazi [2], Abdullah Alghamdi [2] and Saleh Albelwi [2]

[1] Department of Computer Sciences, Applied College, University of Tabuk, Tabuk 47512, Saudi Arabia; mqalrashidi@ut.edu.sa
[2] Faculty of Computing and Information Technology, University of Tabuk, Tabuk 47731, Saudi Arabia; i.alkhazi@ut.edu.sa (I.S.A.)
* Correspondence: malghamdi@ut.edu.sa (M.A.); aabadleh@ut.edu.sa (A.A.); smnasri@ut.edu.sa (S.M.)

Abstract

The paper introduces a real-time digital-twin controller that manages evacuation routes while operating GEEM for emergency energy management during building fires. The system consists of three interconnected parts which include (i) a physics-based hazard surrogate for short-term smoke and temperature field prediction from sensor data (ii), a router system that manages path updates for individual users and controls exposure and network congestion (iii), and an energy management system that regulates the exchange between PV power and battery storage and diesel fuel and grid electricity to preserve vital life-safety operations while reducing both power usage and environmental carbon output. The system operates through independent modules that function autonomously to preserve operational stability when sensors face delays or communication failures, and it meets Industry 5.0 requirements through its implementation of auditable policy controls for hazard penalties, fairness weight, and battery reserve floor settings. We evaluate the controller in co-simulation across multiple building layouts and feeder constraints. The proposed method achieves superior performance to existing AI/RL baselines because it reduces near-worst-case egress time (T_{95} and worst-case exposure) and decreases both event energy E_{event} and CO_2-equivalent $CO_{2\text{event}}$ while upholding all capacity, exposure cap, and grid import limit constraints. A high-VRE, tight-feeder stress test shows how reserve management, flexible-load shedding, and PV curtailment can achieve trade-offs between unserved critical load U_{energy} and emissions. The team delivers implementation details together with reporting templates to assist researchers in reaching reproducibility goals. The research shows that emergency energy systems, which integrate evacuation systems, achieve better safety results and environmental advantages that enable smart-city integration through digital thread operations throughout design, commissioning, and operational stages.

Keywords: digital twin; emergency evacuation; fairness-aware routing; hazard-aware path planning; grid-interactive energy management

Academic Editors: Francesco Riganti Fulginei and Pablo García Triviño

Received: 2 October 2025
Revised: 15 December 2025
Accepted: 18 December 2025
Published: 22 December 2025

1. Introduction

Digital twins (DTs) are reshaping building fire safety by closing the loop between sensing, simulation, and control. Early work on DT-enabled safety management has

articulated end-to-end frameworks that fuse building information models, sensor networks, and control logic to support preparedness, response, and recovery [1]. Recent surveys emphasize that evacuation is a particularly promising DT application because it combines high-frequency state estimation (hazard fields, congestion) with rapid decision support for routing and signage [2]. Alongside these advances, the safety community continues to examine dynamic guidance, such as adaptive exit signage, and to catalog the operational pitfalls that accompany it (e.g., conflicting cues, latency, and human factors) [3], as well as scalable implementations for multi-floor, multi-exit facilities [4].

On the computational side, a large literature optimizes evacuation routes under evolving conditions. Studies span domain settings, from cruise ships to underground facilities, and propose real-time algorithms that balance path length, congestion, and safety penalties [5–7], while more recent work explores reinforcement learning and simulation-assisted optimization to handle large state spaces and nonstationary dynamics [8,9]. For hazard realism, many engineering workflows still rely on high-fidelity fire/smoke simulators such as the Fire Dynamics Simulator (FDS) [10], often coupled, one-way or two-way, with agent-based egress models to evaluate performance-based designs [11–13]. Reviews of agent-based evacuation underscore how fundamental diagram speed–density effects, social behaviors, and heterogeneity shape outcomes [14,15], and a growing body of literature integrates social-force dynamics with routing or learning components [16,17].

Despite this progress, two gaps remain salient for modern smart buildings. First, evacuee routing typically optimizes for efficiency or safety but only implicitly treats fairness. The protection of vulnerable populations whose mobility constraints or exposure sensitivity differ from those of the average occupant remains overlooked. Policy and empirical evidence show that disabled or mobility-impaired residents remain disproportionately at risk even in jurisdictions that have invested heavily in safety upgrades [18–21]. While a few transportation-oriented studies frame evacuation planning as a multi-objective problem that includes equity [22,23], these approaches rarely appear in building-scale, real-time controllers. Second, evacuation control rarely co-optimizes on-site energy during emergencies. In practice, a building's ability to keep critical services online (communications, elevators reserved for assisted egress, pressurization fans, and medical equipment) depends on photovoltaic generation, battery state of charge, grid import caps, and backup generation. Outside the building, emergency vehicle (EV) priority at signals also affects response time and coordination [24–26]. Treating egress and energy as separate problems can obscure important trade-offs: for instance, maintaining critical power under tight caps may demand choices that increase short-term emissions but reduce life-safety risk by preserving protected egress systems.

The proposed method addresses the problem of designing a building-scale digital-twin controller that, in real time, (i) routes occupants to safe exits as hazards and congestion evolve, (ii) enforces accessibility constraints and an explicit fairness policy for vulnerable groups, and (iii) dispatches on-site energy resources to sustain critical services under grid and device limits, thereby making the life-safety versus sustainability trade-off explicit and controllable. The controller must operate with surrogate hazard predictions faster than high-fidelity CFD, remain computationally feasible at sub-second to multi-second control periods, and expose policy knobs that let operators choose acceptable points on the efficiency–equity and safety–sustainability frontiers.

Building on the DT foundations in ref. [1,2], we integrate three ingredients that are typically studied in isolation: (a) hazard- and congestion-aware routing on a building graph, informed by a physically motivated advection–diffusion–decay surrogate that approximates smoke/toxic propagation (complementary to the high-fidelity models used in [10–13]); (b) a runtime fairness policy that prioritizes vulnerable occupants, tightens ex-

posure thresholds, and slightly adjusts admission/throughput where capacity bottlenecks arise (motivated by evidence on mixed-ability risks [18–21] and equity formulations in disaster logistics [22,23]); and (c) Grid-Interactive Emergency Energy Management that balances photovoltaic generation, battery storage, capped grid import, and optional backup generation to keep critical loads online during the event window. At the city edge, we read out an emergency-vehicle priority indicator consistent with the traffic-signal preemption literature and deployments [24–26]. The system runs in a receding-horizon loop, continuously re-weighting edges by predicted hazard, density-dependent travel times, and policy terms, while a co-simulated dispatcher allocates energy subject to device and feeder constraints.

Compared to prior work on dynamic signage and path optimization [3–7], reinforcement-learning–based routing [8,9], and coupled fire–egress analysis [10–13], the proposed approach differs from the existing work in several points: (i) we jointly control life-safety and energy by formulating evacuation guidance and emergency energy dispatch as a single, closed-loop digital-twin controller, rather than treating energy as an external constraint or an ex post metric; this renders the trade-off between sustaining critical services and short-term emissions operational within the controller itself, a capability not addressed in building-scale, real-time evacuation systems to our knowledge; (ii) we elevate fairness to a first-class optimization [27,28] objective by encoding it at runtime, prioritizing vulnerable occupants in planning order, tightening exposure thresholds for sensitive groups, and making small capacity/throughput adjustments at bottlenecks. This approach creates a tunable equity–efficiency frontier that decision-makers can navigate, linking empirical risk evidence [18–21] with optimization-based planning [22,23]; and (iii) we introduce a physically informed yet computationally light hazard surrogate, a graph-based advection–diffusion–decay model that captures key transport and attenuation behaviors while remaining fast enough for control, complementing high-fidelity tools such as FDS and FDS + Evac that are essential for design verification but too slow for inner-loop control [10–13].

The main contributions of the paper are as follows:

- A unified digital-twin controller that couples hazard- and congestion-aware routing with grid-interactive energy management, exposing operator-level knobs for fairness and sustainability.
- An operational fairness mechanism for mixed-ability populations that acts during routing/queuing, not just as a reporting metric, enabling explicit navigation of the equity–efficiency trade-off.
- A co-simulation framework that integrates a physically grounded hazard surrogate, agent-based movement with capacity gating, and device/feeder-constrained energy dispatch, which is compatible with DT data streams and amenable to site calibration.

2. Related Work

We focus on four strands most relevant to our controller: (i) DT-enabled fire safety and evacuation; (ii) dynamic signage and real-time egress control; (iii) routing under dynamic hazards (optimization and learning); and (iv) building–city coupling and emergency energy management. A compact comparison is kept in Table 1, while each subsection below emphasizes deltas that matter for a deployable DT controller (latency, robustness, and co-optimization).

Table 1. Comparative analysis of recent works.

Work/Domain	Real-Time?	Hazard-Aware?	Fairness?	DT-Integrated?	Energy Co-Optimization?	Multi-Objective?	Main Limitation
Classical evacuation simulators	X	Partial	X	X	X	X	Static, no DT, no sensing
Hazard-aware routing	✓	✓	X	X	X	Partial	No fairness, limited sensing
DT-based fire modeling	✓	✓	X	✓	X	X	Simulation only
Building energy management	✓	X	X	Partial	✓	Partial	No evacuation consideration
Hierarchical DT [29]	x	Partial	X	✓	✓	N/A	Not designed for sub-second, life-critical building emergencies
Fuzzy evacuation [30]	x	Partial	x	X	X	N/A	Strategic recovery focus, not sub-second emergency control.
This work	✓	✓	✓	✓	✓	✓	Integrated, real-time, multi-domain optimization

2.1. Digital Twins (DT), BIM, and IoT for Fire Safety and Evacuation

Digital twins (DTs) are transforming building fire safety by enabling real-time situational awareness, analytics, and decision support beyond static design rules. Recent frameworks integrate BIM, IoT sensing, AI, and AR to track hazards and occupant states and to support evacuation and system control in smart buildings [1,2]. Surveys focused on emergency evacuation identify the next steps as building–occupant co-modeling, streaming data fusion, and closed-loop decision logic for resilient smart cities [2]. Despite this progress, most implementations remain single-building and route-advisory, with limited coupling to high-fidelity hazard fields or city-scale energy resources. In contrast to hierarchical DT architectures and uncertainty-aware evacuation planning, our design operationalizes hierarchy for sub-second, safety-critical control with explicit life-safety priorities. We implement online robustness measures such as hazard inflation, queue buffers, and a GEEM state-of-charge floor, rather than relying on offline fuzzy uncertainty envelopes. This makes our approach a practical, real-time complement to prior hierarchical DTs for energy adaptation and to fuzzy multi-objective plans intended for medium-term recovery.

To address the identified gaps, we propose a fully integrated DT that continuously ingests sensor data and jointly optimizes evacuation routing and emergency energy management.

2.2. Dynamic Signage and Real-Time Egress Control

A complementary thread replaces fixed exit signs with intelligent dynamic signage that responds to evolving smoke, heat, and congestion conditions. The NFPA's research foundation review (2022) synthesizes the state of dynamic exit signs and documents the technical and regulatory challenges ahead [3]. Building on this, Yen and Lin (2024) formulate a mathematical control model for multi-floor, multi-exit buildings that drives sign directions using temperature/smoke constraints and congestion indices; FDS-based experiments demonstrate notable gains over fixed-logic baselines [4]. While this shows building-level viability, most systems optimize group guidance (sign-level control) rather than personalized routing with fairness objectives or two-way coupling to physics-based hazard fields.

2.3. Routing Under Dynamic Hazards: Optimization and Learning

Real-time routing that accounts for changing hazards and crowding is advancing rapidly. Liu et al. embed FDS outputs into a time-varying equivalent-weight Dijkstra to optimize routes under evolving visibility, gas, temperature, and congestion, reporting performance improvements on complex layouts [5]. Other recent work proposes multi-

objective real-time search for underground fires [6] and a self-adaptive escape model that augments Dijkstra with spatiotemporal safety indices [7]. Learning-based routing is emerging as well (e.g., deep RL and hybrid simulation-optimization) to plan routes under uncertainty and congestion at scale [8,9]. Yet, across these studies, the hazard–egress coupling is typically one-way (pre-run CFD to routing) rather than closed-loop, and optimization objectives focus on average evacuation time, rarely encoding equity/fairness or multi-stakeholder city coordination.

Physics-Based Fire/Smoke Modeling and Coupled Egress

High-fidelity hazard fields remain the standard for rigorous evaluation. FDS (NIST) provides LES-based CFD for combustion, smoke, and thermal fields; its documentation and validation corpus underpin most performance-based studies [10]. FDS + Evac extends this with an agent-based egress module, allowing (at least) one-way coupling between fire dynamics and crowd movement [11]. Reviews consistently note that many assessments still use one-way coupling, limiting feedback from evolving crowd states back to hazard scenarios; recent theses and reviews call for two-way coupled DT-BIM-CFD-egress ecosystems to capture interactions more faithfully [12,13]. These observations motivate closed-loop architectures where sensors and predictive surrogates update both hazard and routing models online.

2.4. Agent-Based Evacuation and Social Behavior

State-of-the-art reviews of agent-based evacuation models (ABM) stress the importance of empirically grounded behaviors (communication, compliance, panic mitigation), validation, and congestion phenomena [14,15]. Improved social-force formulations and ABM architectures capture complex crowd dynamics in stations, tunnels, and high-rise contexts and are now routinely combined with path planning and hazard fields in simulators [16,17]. These advances give a mature substrate on which to layer DT-driven, personalized, fairness-aware routing.

Most models assume static building conditions and do not integrate real-time sensing or hazard dynamics. Routing tends to rely on fixed edge costs that do not adapt to smoke, visibility, or congestion. Vulnerable occupants (mobility-impaired, elderly, children) are rarely treated using fairness or priority mechanisms. These approaches do not embed evacuation routing within a multi-objective optimization framework.

As a solution for these shortcomings, we will incorporate, in our study, dynamic hazard fields, congestion-aware weights, and fairness constraints into a multi-objective routing formulation executed in real time inside a digital twin.

2.5. Equity, Vulnerable Populations, and Fairness in Evacuation

Post-Grenfell discourse and official guidance emphasize inclusive evacuation, yet the literature still under-represents people with disabilities or cognitive/functional limitations in building evacuations [18–21]. Fairness-driven optimization is growing in disaster evacuations (e.g., tsunami/hurricanes), where resource allocation and route scheduling explicitly balance efficiency and equity [22,23]. Building-fire contexts have begun to incorporate mixed-ability populations in simulation studies [21], but explicit fairness metrics (e.g., minimizing disparity in risk/egress times across vulnerable groups) remain rare in indoor, DT-enabled evacuation optimization.

Hazard modeling usually remains single-source (e.g., smoke only), ignoring multi-hazard interactions (temperature, CO, low visibility). Moreover, most algorithms do not scale to real-time due to reliance on computationally heavy simulation, and existing hazard-aware routing does not ensure fairness or equitable evacuation.

Our method will use multi-hazard sensing combined with a computationally lightweight but rigorous optimization model, enabling real-time operation and fairness-aware routing.

2.6. Building–City Coupling and Emergency Vehicle Priority

A final strand links building to city-level response systems, notably emergency vehicle preemption (EVP) and traffic management to reduce response and transport delays. Recent works propose dynamic preemption strategies and document funded municipal pilots integrating signal control and emergency priority, evidencing the feasibility of coordinated, cross-infrastructure response at urban scale [24–26]. However, few building-fire studies close the loop by feeding DT hazard/occupant states to city control systems and hospitals in real time, suggesting novel opportunities for DT-to-city orchestration in life-safety operations.

Energy management is almost always formulated independently of evacuation processes. Existing models seldom operate on the short time scales (seconds–minutes) required during emergencies. In addition, no studies consider how evacuation loads (stair pressurization fans, emergency elevators) interact with PV/battery/diesel assets in real time, and CO_2 efficiency and unserved critical-load metrics are rarely integrated simultaneously. We will introduce GEEM, the first Grid-Interactive Emergency Energy Management model explicitly co-optimized with evacuation routing inside a digital twin.

2.7. Summary of Gaps and Contributions

Overall, existing studies provide important foundations in evacuation modeling, hazard sensing, and energy management. However, none integrates multi-hazard sensing, real-time digital twin state estimation, fairness-aware routing, and grid-interactive energy dispatch into a unified multi-objective optimization framework. Table 1 summarizes the methodological gaps this work addresses.

Unlike prior research works treating evacuation guidance, hazard estimation, and building energy as largely separate problems, our approach integrates them into a single, closed-loop digital-twin controller. Framework papers on digital twins for evacuation establish feasibility but are mostly route-advisory and single-building in scope, with limited coupling to high-fidelity hazards or energy systems [1,2]. Dynamic signage studies optimize sign states at the group level yet rarely personalize guidance or encode fairness objectives in the control loop [3–7]. Optimization and learning-based routing typically ingest precomputed hazard fields (one-way coupling) and focus on average clearance time rather than equity or operational sustainability [5–9]. High-fidelity fire/egress couplings provide essential evaluation fidelity but are too slow for inner-loop control and seldom close the loop online [10–13]. In contrast, our controller (i) co-optimizes life-safety and emergency energy in real time, making the trade-off between sustaining critical services and short-term emissions explicit and operator-tunable; (ii) operationalizes fairness during planning and admission, which enables prioritizing vulnerable occupants, tightening exposure thresholds, and providing small capacity/throughput adjustments rather than reporting equity only ex post [18–23]; and (iii) employs a physically informed yet lightweight hazard surrogate that preserves key transport/attenuation behaviors while meeting real-time deadlines, complementing but not replacing high-fidelity models [10–13]. Additionally, we expose a city-coupling pathway (emergency-vehicle priority) to align building operations with urban response systems, which is an integration that remains uncommon in building-fire studies [24–26]. Together, these choices move from advisory, one-way pipelines to an actionable, personalized, fairness-aware, and energy-aware controller suitable for deployment and policy tuning.

Prior DT-evacuation studies are predominantly route-advisory, rely on one-way CFD to routing coupling, and treat equity and energy as ex-post metrics. In contrast, the proposed controller (a) closes the loop with sensor-driven re-planning at control-period scale, (b) operationalizes fairness inside the routing/queuing decisions, and (c) co-optimizes GEEM with explicit feeder/device limits. This shifts DTs from descriptive platforms to real-time decision engines aligned with life-safety and sustainability objectives.

3. Proposed Method

We present a closed-loop, digital-twin (DT) evacuation controller that fuses (i) physics-informed hazard prediction, (ii) agent-based crowd dynamics, (iii) personalized, fairness-aware routing, and (iv) city/grid orchestration. Figure 1 summarizes the architecture.

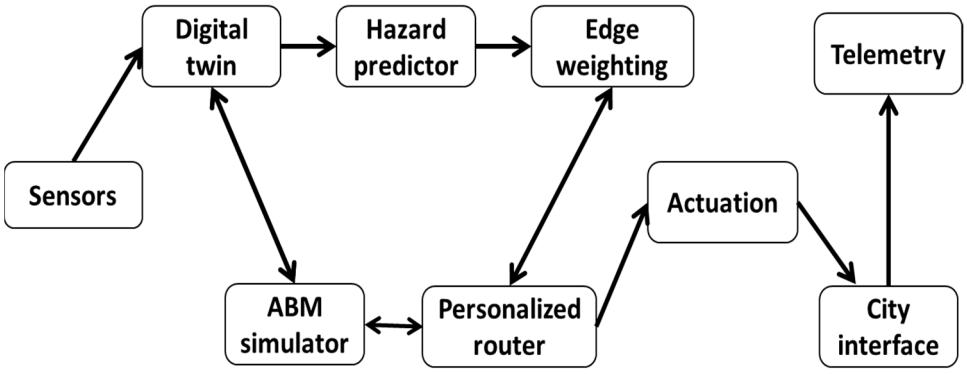

Figure 1. System architecture (digital twin-driven evacuation control loop).

3.1. System Overview and Data Flow

Let Ω denote the floor domain discretized into a graph $G = (V, E)$. The DT maintains a hazard state h_t (temperature/visibility/toxicant proxies) and occupant state $Xt = \{x_i(t), v_i(t)\}_{i=1}^{N}$. Each control step of length Δt involves the following:

1. Sense and predict: Ingest sensors z_t(smoke/CO/temperature/visibility, door/HVAC states, BLE/Wi-Fi occupancy) and predict a short-horizon field $\hat{h}_{t+\Delta t} = f_\theta(h_t, z_t)$
2. Update weight: Map $\hat{h}_{t+\Delta t}$ and congestion $\rho_e(t)$ to time-varying edge weights $w_e(t)$
3. Personalize routing: Solve a multi-objective program to compute paths $\{\rho_i(t)\}$ balancing efficiency, safety, and fairness.
4. Simulate and actuate: Simulate agents over $[t, t + \Delta t]$, update $h_{t+\Delta t}$, $X_{t+\Delta t}$, and actuate dynamic signs/luminaires and city interfaces.

In addition to evacuation control, the digital twin also supervises a Grid-Interactive Emergency Energy Management (GEEM) layer. GEEM co-optimizes emergency power flows (e.g., smoke-control/pressurization fans, emergency lighting, elevators) and on-site resources (PV, battery, diesel) with evacuation objectives. At each control step Δt, the controller exchanges set-points with the building energy system and a distribution-grid interface that enforces feeder import limits and voltage/thermal constraints, thereby aligning life safety with energy and environmental sustainability during the event.

3.2. Formal Optimization Problem

We formalize the real-time evacuation–energy co-optimization problem solved at each control step t ∈ {0, Δ_t, 2Δ_t, … }. The building is represented by a directed graph G = (V, E), where nodes are rooms, halls, stairs, and exits, and edges represent walkable connections.

Decision variables:

Routing variables

For each occupant i ∈ I, edge e ∈ E, and time step t, we define:
$x_{i,e}(t) \in \{0, 1\}$: 1 if occupant **i** is assigned to traverse edge **e** at time t.
Ti: predicted egress time of occupant **i**.
Ei: cumulative exposure of occupant **i**.

Energy dispatch variables (GEEM)

$P^{grid}(t)$: real power imported from the grid (positive) or exported (negative).
$P^{pv}(t)$: PV real power dispatched.
$P^{bat}(t)$: battery real power (positive = discharge).
$P^{diesel}(t)$: diesel generator output (proposed method only).
s(t): battery state of charge.

Parameters

$w_e^{haz}(t)$: hazard-derived weight for edge eee (hazard surrogate).
$w_e^{cong}(t)$: congestion penalty based on density.
$w_e^{acc} \in \{0, 1\}$: accessibility flag (1 if unusable by vulnerable groups).
$\kappa_1, \kappa_2, \kappa_3$: weights for hazard, congestion, and accessibility.
$\alpha, \beta, \gamma, \eta, \zeta, \xi$: objective weights (clearance, exposure, fairness, energy, CO_2, unserved critical load).
$\overline{P}^{grid}, \overline{P}^{bat}, \overline{P}^{diesel}$: device limits.
$L^{crit}(t)$, $L^{flex}(t)$: critical and flexible loads.
Control: $\Delta t = 0.5 - 2\,s\ (typ.\ 1\,s)$.
Edge costs: κ1 (hazard) ∈ [2,6], κ2 (congestion) ∈ [0.5, 2], κ3 (inaccess.) large (hard penalty).
Objective weights (7): α (clearance) dominant; β (exposure) medium; γ (fairness) medium; η, ζ, ξ (energy/CO_2/unserved) tuned so safety terms remain primary.
Exposure caps: Ξ_i stricter for vulnerable groups; fairness weights (α, β, γ) adjusted per group.
GEEM: battery reserve floor $s_{min}+$ margin; grid import cap at 10–20% headroom; diesel off unless L_t^{cirt} at risk.

Feasibility constraints

(1) Flow conservation

For all occupants i, the following is applied:

$$\sum_{e \in \delta^+(v)} x_{i,e}(t) - \sum_{e \in \delta^-(v)} x_{i,e}(t) = \{1\ for\ v = start(i)\ ;\ -1\ for\ v = exit(i)\ ;\ 0\ otherwise\}$$

(2) Edge capacity and queueing

$$\sum_i x_{i,e}(t) \le C_e(t),\ \forall\ e \in E$$

(3) Hazard and accessibility

$$x_{i,e}(t) = 0\ if\ w_e^{haz}(t) \ge\ T_{unsafe}\ or\ w_e^{acc} = 1$$

(4) Exposure constraint

$E_i = \sum_t \sum_e x_{i,e}(t)\lambda_e(t) \le \overline{E}_i$ with stricter $\overline{E}_i$ for vulnerable groups.

(5) Energy balance and device constraints (GEEM)

Power balance:

$$P^{grid}(t) + P^{pv}(t) + P^{bat}(t) + P^{diesel}(t) = L^{crit}(t) + L^{flex}(t).$$

Battery SOC:

$$s(t + \Delta t) = s(t) + \eta_c \, P^{bat}(t)^{-} - 1/\eta_d \, P^{bat}(t)^{+}$$

Device limits:

$$\left|P^{grid}(t)\right| \leq \overline{P}^{grid}, \left|P^{bat} \leq \overline{P}^{bat}\right|, 0 \leq P^{diesel}(t) \leq \overline{P}^{diesel}$$

Objective function

At each control cycle, we minimize the following multi-objective function:

$$min_{x,\,P} \; \alpha \max_i T_i + \beta \sum_i E_i + \gamma \, \Phi(\{T_i\}) + \eta \sum_t E_{event}(t) + \zeta \sum_t CO_2(t) + \xi \sum_t U_{crit}(t)$$

where

- Clearance time term: $\max_i T_i$
- Safety term: $\sum_i E_i$
- Fairness term (e.g., max–min gap):

$$\Phi = max_{g \epsilon G} \overline{T}_g - min_{g \epsilon G} \overline{T}_g$$

- Energy use:

$$E_{event}(t) = P^{grid}(t)^{+} - P^{pv}(t) - P^{bat}(t)^{-}$$

- CO_2 emissions:

$$CO_2(t) = \gamma_{grid} \, P^{grid}(t)^{+} + \gamma_{diesel} \, P^{diesel}(t)$$

- Unserved critical load:

$$U_{crit}(t) = \max\{0, L^{crit}(t) - P_{supplied}(t)\}.$$

Solution strategy (two-stage decomposition)

Stage A: Weight computation

The edge weight is calculated as a weighted sum of hazard, congestion, and accessibility terms:

$$w_e(t) = \kappa_1 \, w_e^{haz}(t) + \kappa_2 \, w_e^{cong}(t) + \kappa_3 \, w_e^{acc}.$$

Stage B: Routing + GEEM dispatch

Solve time-dependent shortest paths with fairness and capacity gating.

Solve QP/MILP for GEEM subject to device/grid constraints.

The hazard surrogate uses an advection–diffusion–decay form chosen to retain transport/attenuation behavior while meeting a ≤ 2 s control deadline. The nominal coefficients (advection hop delay 15 s/edge with HVAC multiplier 1.3, diffusion 0.002, decay 0.03) reproduce first-crossing times against the high-fidelity reference with low magnitude error and consistent edge-risk ordering. The unsafe threshold 1.3 marks the onset of deteriorated tenability (reduced visibility/thermal stress) and was picked to be conservative under sensor latency. Fairness policy uses priority ordering, +1 capacity nudge at bottlenecks, and vulnerable assistance (speed-assist 0.95× and 0.7× exposure cap). These values were

tuned to sit on the knee of the fairness–efficiency Pareto: they reduce the max–min egress gap without materially harming clearance time. GEEM limits reflect common emergency envelopes for a medium commercial building (PV 80 kW, BESS 150 kWh/80 kW with 20% reserve, import cap 120 kW). Under tight-cap stress, we enforce anti-backfeed logic, ramp-rate constraints, and reserve floors; in that regime, the controller trades CO_2 for zero unserved critical.

3.3. System Architecture and Routing Graph for Adaptive Evacuation

Every Δt seconds, we (i) update hazards, (ii) refresh weights, (iii) recompute routes for agents whose planned path intersects deteriorating edges, and (iv) actuate guidance.

Algorithm 1 shows the overall procedure of the proposed method. The steps of Algorithm 1 can be explained as follows:

Algorithm 1. Closed-loop digital twin evacuation with GEEM (period Δt)

Inputs: $G = (V, E)$, hazard surrogate f_θ, hazard state h_t, occupants X_t, sensors z_t, energy state st (battery SOC) and device/grid limits, period Δt, parameters $\kappa 1, \kappa 2, \kappa 3,\ (\alpha, \beta, \gamma, \eta, \zeta, \xi)$, exposure caps $\{\Xi_i\}$, emission factors ϕ_t^{grid}, ϕ^{dg}.
Outputs: Personalized routes $\{p_i\}$, energy set-points $\left\{P_t^{grid}, P_t^{pv}, P_t^{batt}, P_t^{dg}, L_t^{flex}\right\}$, building actuation (signs/lights) and grid/DR telemetry.
Procedure:

1. **Sense and predict:** Acquire z_t; compute $\hat{h}_{\{t+\Delta t\}} \leftarrow f_\theta\ (h_t, z_t)$ (1).
2. **Update edge costs:** For each edge e, estimate $p_i(t)$ and $T_e(t)$, $\mathrm{Haz}_e(t), \mathrm{Cong}_e(t), \mathrm{Inacc}_e t$; set $w_e(t)$ per the cost model (6).
3. **Trigger replans:** Build set $\mathcal{I}$ for agents with unsafe edge ahead, queue above threshold, or $admits\ >\ \varepsilon\ improvements$.
4. **Route-optimization update:**
   - **4.1 Build a feasible subgraph**
     Remove all edges disabled by safety or accessibility constraints.
   - **4.2 Compute candidate paths**
     For each individual or group I, generate a set of k candidate shortest paths using updated edge weights we(t).
   - **4.3 Apply fairness and priority rules**
     Assign priority modifiers for vulnerable groups, fire wardens, and mobility-impaired occupants, adjust path costs using group-based modifiers.
   - **4.4 Evaluate capacity and queueing feasibility**
     Simulate predicted flows, eliminate candidate paths violating edge capacities, update waiting times and congestion forecasts.
   - **4.5 Select optimal routing decision**
     Solve reduced multi-objective problem: min [α Tmax + β ∑iEi + γΦ({Ti})] over the remaining feasible candidate paths.
   - **4.6 Stabilize solution (anti-oscillation)**
     Apply hysteresis or time-based smoothing to prevent route oscillations between consecutive control cycles.
5. **Conflict resolution:** Where edge capacities are exceeded, run a local min-cost-flow reassignment on the saturated neighborhood to rebalance paths.
6. **Simulate and actuate** $\boldsymbol{t \leftarrow t + \Delta t}$**.** Advance agents/queues to $X_{t+\Delta t}$; update $h_{t+\Delta t}$ from sensors; actuate dynamic signs/luminaires and apply energy set-points; publish DR/telemetry to the city/DSO.
7. **Loop:** Set $t\ \leftarrow\ t + \Delta t$ and repeat steps 1–6 until evacuation completes or timeout.

- **Step 1** updates the digital twin state using sensor fusion. The hazard map, density map, and energy-system state are refreshed every cycle.
- **Step 2** computes per-edge hazard, congestion, and accessibility metrics. These are transformed into dynamic edge weights used by the routing layer.
- **Step 3** applies feasibility constraints, ensuring that unsafe or inaccessible passages are excluded before any optimization takes place.
- **Step 4** performs the evacuation-routing optimization. Previously condensed, this step is now decomposed into path generation, fairness adjustments, capacity/queue checking, and final multi-objective selection.
- **Step 5** solves the GEEM problem using the formal mathematical formulation.
- **Step 6** deploys the results to the building (signage, BMS/EMS).
- **Step 7** triggers the next control cycle, ensuring real-time operation.

Figure 1 shows the closed-loop pipeline used during an event. Sensors (temperature, visibility, CO, door/HVAC states, and occupancy beacons) feed the building's digital twin, which maintains the current hazard and occupant states. A fast hazard predictor (surrogate of high-fidelity fire/smoke models) produces short-horizon forecasts that are converted into edge weights for the building graph (safety, congestion, and accessibility factors). These weights drive a personalized router that computes time-dependent routes for affected occupants while respecting capacities and fairness priorities (e.g., mobility-limited users). The agent-based simulator advances occupants and queues over the next control interval, and the system actuates dynamic signs/luminaires while publishing telemetry to the city systems (e.g., traffic signal preemption, hospital notification). Two feedback are explicit: simulator $\rightarrow$ digital twin (state update) and city interface $\rightarrow$ telemetry. The numbered badges indicate the receding-horizon steps: sense, update DT, predict, weight edges, route, and actuate, then repeat.

Figure 2 illustrates adaptive routing on a simplified floor. Occupants leave Room A and Room B toward Hall 1, then to the Stairs and Exit. Solid arrows denote primary flows; the "Blocked" mark on the Hall $\rightarrow$ Stairs segment represents a deteriorating or inaccessible link detected by sensors/prediction. When a link becomes unsafe or saturated, the controller generates alternative paths (dashed arrows) that detour around the blockage while honoring local capacities. The crowd markers near Hall 1 depict density buildup that contributes to congestion penalties used in routing. A dynamic sign cue indicates how guidance devices switch direction in sync with newly computed routes. Together, the elements show how the system performs time-dependent re-planning each control period, redirecting flows away from hazards and bottlenecks to maintain safe, efficient, and equitable egress.

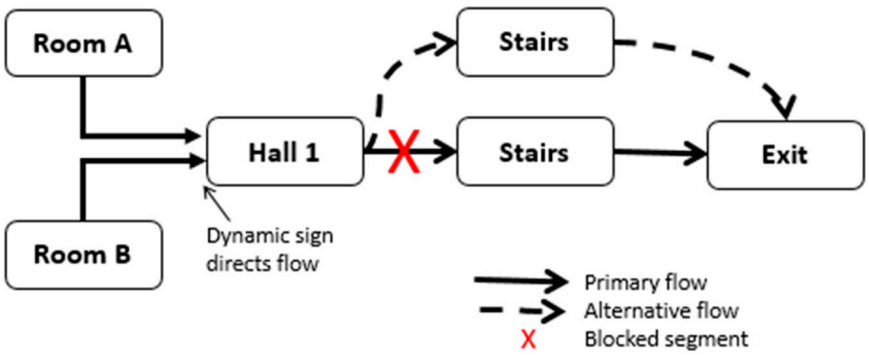

Figure 2. Routing graph for adaptive evacuation.

3.4. Grid-Interactive Emergency Energy Management (GEEM)

Grid-Interactive Emergency Energy Management (GEEM) is a layer that makes the building behave as a grid-interactive microgrid during fire events. GEEM co-optimizes life-safety operations (smoke control, pressurization, emergency lighting, elevators) and on-site resources (PV, battery, diesel) together with grid exchange to support evacuation while minimizing event energy use, CO_2-equivalent emissions, and unserved critical load. At each control step, the digital twin solves a small dispatch problem subject to power balance, device limits (battery SOC, generator bounds), and a distribution-grid interface (import caps, voltage/thermal constraints), and then issues set-points to the BMS/microgrid controllers. The resulting energy decisions are coupled with the routing controller so that evacuation guidance and power dispatch evolve consistently under changing hazards, congestion, and grid conditions. This section formalizes the GEEM variables and constraints and defines the sustainability metrics used in the joint objective.

At each control step t, the digital twin solves a local energy dispatch problem that determines how on-site resources and grid exchange support life-safety systems during the evacuation. We denote the following:

Decision Variables

$P^{grid}(t) \in R$: Grid real power (import > 0, export < 0).
$P^{pv}(t) \geq 0$: PV real power dispatched.
$P^{bat}(t) \in [-\overline{P}^{bat}, \overline{P}^{bat}]$: Battery real power (discharge > 0).
$P^{diesel}(t) \in [0, \overline{P}^{diesel}]$: Diesel generator output.
$s(t) \in [s^{min}, s^{max}]$: Battery state of charge (SOC).

Parameters

$L^{crit}(t)$: Critical life-safety load (e.g., pressurization fans, emergency lights, evacuation elevators).
$L^{flex}(t)$: Flexible/non-critical load that may be shed.
η_c, η_d: Charge/discharge efficiencies.
$\overline{P}^{grid}$: Grid-import cap imposed by the feeder.
$\gamma_{grid}, \gamma_{diesel}$: Emission factors.
Δt: Control period.

Constraints

(1) Power balance

For each control step, the energy balance is maintained according to the following constraint:

$$P^{grid}(t) + P^{pv}(t) + P^{bat}(t) + P^{diesel}(t) = L^{crit}(t) + L^{flex}(t) - U_{crit}(t),$$

where $U_{crit}(t) \geq 0$ is a slack variable representing unserved critical load. This ensures feasibility even under extreme grid caps or depleted storage.

(2) PV dispatch limit

$$0 \leq P^{pv}(t) \leq \overline{P}^{pv}(t), \text{ where } \overline{P}^{pv}(t) \text{ is determined by available irradiance.}$$

(3) Battery dynamics and limits

Battery SOC evolves as follows:

$$s(t + \Delta t) = s(t) + \eta_c P^{bat}(t)^- - 1/\eta_d P^{bat}(t)^+,$$

where $P^{bat}(t)^+ = \max\{P^{bat}(t), 0\}$, $P^{bat}(t)^- = \max\{-P^{bat}(t), 0\}$.

The state-of-charge constraint is as follows:

$$s^{min} \leq s(t) \leq s^{max}.$$

A reserve floor (e.g., $s^{min} = 0.2\, s^{max}$) ensures enough energy for prolonged evacuation or re-ignition scenarios.

(4) Grid import/export

$$-\overline{P}^{grid} \leq P^{grid}(t) \leq \overline{P}^{grid}.$$

This enforces feeder safety and respects the emergency demand-response signal.

If anti-backfeed protection is required, $P^{grid}(t) \geq 0$.

(5) Diesel generator

$$0 \leq P^{diesel}(t) \leq \overline{P}^{diesel}.$$

Diesel is held at zero and invoked only when needed to prevent an unserved critical load.

(6) Sustainability metrics at each step

Event energy

$$E_{event}(t) = P^{grid}(t)^{+} - P^{pv}(t) - P^{bat}(t)^{-}.$$

CO_2-equivalent emissions

$$CO_2(t) = \gamma_{grid}\, P^{grid}(t)^{+} + \gamma_{diesel}\, P^{diesel}(t).$$

Unserved critical load

$$U_{crit}(t) = \max\{0, L^{crit}(t) - P^{supplied}(t)\}.$$

(7) GEEM Objective

GEEM minimizes a weighted sum of energy, emissions, and critical-load violations:

$$\min_{\{P(t), s(t)\}} \eta \sum_t E_{event}(t) + \zeta \sum_t CO_2(t) + \xi \sum_t U_{crit}(t)$$

3.5. Deployment Pathway on Existing BMS

The controller integrates with a standard building stack via BACnet/IP or OPC UA for telemetry and actuation. Required points include environmental sensors, access and HVAC states, DER set-points (PV curtailment, BESS charge/discharge limits and reserve floors, generator enable), and feeder import/export measurements. The runtime exposes auditable policy knobs for operators and enforces safety invariants: life-safety circuits are never shed; if a cycle deadline is missed, the system holds the last feasible commands; if communications degrade, the system reverts to safe defaults. A staged protocol is used for adoption: hardware-in-the-loop with a controller emulator, then shadow mode alongside the BMS, then supervised live with operator override authority and automated rollback. Interfaces align with prevailing building and fire-safety practice (for example, coordination with the fire alarm control panel under NFPA guidance), without requiring proprietary vendor features.

4. Evaluation

This section introduces and details the achieved experiments.

4.1. Objectives and Research Questions

The evaluation quantifies how the proposed closed-loop controller (Algorithm 1) performs under realistic fire scenarios with dynamic hazards, congestion, and distribution-grid limits. We focus on three questions:

Q1 (Life-safety): Does hazard-aware, capacity-constrained, and fairness-aware routing reduce clearance time and occupant exposure relative to strong baselines?

Q2 (Fairness): Does prioritizing vulnerable groups narrow disparities in egress outcomes without materially harming overall efficiency?

Q3 (Sustainability under constraints): Does the GEEM layer reduce event energy and CO_2-equivalent emissions while respecting feeder/grid limits and preserving life-safety performance?

These questions directly exercise the pipeline defined in Section 3 (sensing → prediction → edge-weighting → personalized routing → actuation) and its co-optimization with GEEM. They also operationalize assumptions and safety invariants shown in Section 3.

4.2. Endpoints and Metrics

We adopt endpoints that correspond to the method's variables and equations:

- Clearance time (T_clear): Total evacuation duration (time to last occupant exit). This is the primary efficiency endpoint reported for each scenario and baseline.
- Per-occupant exposure (E_i): Cumulative exposure computed along realized trajectories; we report distributional summaries (median, IQR) and worst-case exposure across groups.
- Fairness: Primary criterion is the max–min gap in predicted (and realized) egress times across predefined groups (e.g., mobility-limited vs. others), as specified at the end of Section 3.5; we also report group-wise exposure differences to confirm equity gains are not achieved by shifting risk.
- Event energy (E_{event}), CO_{2event}, unserved critical load (U_{crit}): These metrics are used to assess the sustainability and resilience of the GEEM dispatch under feeder import caps and device limits.

Reporting protocol. For each metric, we provide mean ± 95% CI over stochastic seeds/scenarios and paired comparisons against baselines. Unless normality is confirmed (Shapiro–Wilk, $\alpha = 0.05$), we use nonparametric paired tests (Wilcoxon signed-rank) with Holm–Bonferroni correction across endpoints; otherwise, paired t-tests and Cohen's d are reported. Scenario-level spider charts (life-safety vs. sustainability) complement tabular summaries to visualize trade-offs.

4.3. Experimental Setup and Co-Simulation

- Simulation stack. We evaluate the controller with a co-simulation that couples (i) a graph-based hazard surrogate (advection–diffusion–decay with HVAC-amplified advection and hop-delays from the ignition node), (ii) an agent-based egress simulator with density-dependent speeds and capacity gating, and (iii) the closed-loop controller (Algorithm 1) running in a receding-horizon loop with control period $\Delta t = 2\,s$.
- Building models and scenarios. The testbed is a two-floor layout (rooms → halls → stairs → exits) with vertical connectivity H2↔H1 and one exit per floor (E1, E2). The ignition is fixed at H1, and HVAC is on. Links become inaccessible when the predicted edge hazard reaches a safety threshold (hazard ≥ 1.3), forcing re-routing (e.g., H1 →

H2 → S2 → E2 when H1 → S1 deteriorates). We simulate N = 80 occupants with 25% vulnerable, cap the horizon at 900 s, and aggregate over S = 5 random seeds.

- Hazard surrogate. The surrogate updates the edge-level hazard $\boldsymbol{h_e}$ by

$$h_e(t+\Delta t) = h_e(t) + \Delta t\left(k_{diff}\left(\overline{h}N(e) - h_e\right) \; - \; k_{decay}h_e\right) + \Delta t\, k_{adv}(t),$$

with k_{diff} = 0.03, k_{decay} = 0.002, and advective source k_{adv} injected layer-by-layer from the ignition with hop-delay 15 s per hop and HVAC multiplier 1.3. Values are clipped to [0, 1.6]. When $h_e \geq 1.3$, edge e is treated as unsafe (blocked) for routing.

- Crowd model and capacities. Pedestrian speeds follow a fundamental diagram on each edge: $v = v_0(1 - \rho/\rho_{max})$ with ρ_{max} = 2.5 persons/m^2; v_0 equals the edge's base speed (stairs use a conservative factor). Edge capacity is 1.3 · width persons/s. We model FIFO queues at nodes and track per-edge loads so that instantaneous density slows traversal times.
- Routing and controller configuration. Time-varying edge costs combine hazard and congestion terms with weights κ1 = 6 (hazard) and κ2 = 1 (congestion). Fairness is operationalized at runtime via (a) priority ordering, where vulnerable agents are planned/served first each control cycle; (b) a capacity nudge, where effective admission capacity is increased by one person for vulnerable agents when fairness is enabled; and (c) stricter exposure caps for vulnerable users (0.7× of the standard cap). Stairs are marked as inaccessible for wheelchair users. The multi-objective weights in use are α = 1.0 (clearance), β = 0.4 (exposure), γ = 0.4 (fairness), and $(\eta, \zeta, \xi) = (0.1, 0.1, 0.2)$ for energy/CO_2/unserved.
- GEEM and grid limits. The Grid-Interactive Emergency Energy Management (GEEM) layer co-optimizes PV 80 kW, battery 150 kWh (power limit 80 kW; 20% reserve), grid import cap 120 kW, and (for the proposed method only) an optional diesel 100 kW kept off unless required. We compute E_{event}, $CO2_{event}$, and U_{crit} using emission factors = 0.55 and g_{diesel} = 0.80 kg − $CO2_e$/kWh. To exercise GEEM under stress, a tight-cap variant (grid cap 60 kW, curtailed PV 40 kW during part of the event) is reported in the supplement.
- City-level orchestration (diagnostic metric). We estimate emergency vehicle (EV) travel time along a short urban path with and without preemption; preemption is triggered when DT telemetry predicts extended clearance or high average hazard. This does not affect building-level routing but is reported as an external readiness indicator.
- Runtime and hardware. On a CPU-only workstation, the complete control loop (hazard update, edge re-weighting, routing, and dispatch) meets the 2 s deadline across runs. Per-step timings and configuration files are included in the artifact.

4.4. Baseline Implementations and Ablations

4.4.1. Baselines

We implement four strong comparators that isolate the contributions of hazard awareness, congestion handling, fairness, and GEEM. All baselines use the same sensing, graph, and control period Δt as the proposed controller to ensure comparability. Re-planning triggers and device actuation policies are shared unless explicitly disabled.

- B0: Static signage (shortest paths). Time-invariant routing (no re-planning), no hazard or congestion costs. Edges can still become blocked if the hazard threshold (≥1.3) is crossed, forcing agents who reach a blocked edge to stall and re-queue. Fairness mechanics and GEEM are off.

- B1: Congestion-aware only. Time-dependent routing with congestion term (κ2 = 1); hazard term off (κ1 = 0). Re-plans when the next edge becomes unsafe. Fairness and GEEM are off.
- B2: Hazard-aware + congestion. Time-dependent routing with hazard and congestion (κ1 = 6, κ2 = 1). Re-plans on predicted hazard deterioration. Fairness and GEEM are off.
- B3: Fairness-aware, no GEEM. Same as B2, with runtime fairness mechanics enabled: (i) priority ordering (vulnerable planned first each cycle), (ii) capacity nudge (+1 effective admission on next edge for vulnerable), and (iii) stricter exposure caps (0.7×). GEEM is off.
- Proposed method B3 + GEEM (PV 80 kW, battery 150 kWh with 20% reserve, grid cap 120 kW, diesel 0/100 kW kept off unless required). Sustainability terms are active with $(\eta, \zeta, \xi) = (0.1, 0.1, 0.2)$.

4.4.2. Shared Components Across Baselines

Pedestrian speeds follow a fundamental diagram with $\rho_{max} = 2.5\,\text{pers/m}^2$; edge capacity is $1.3 \times \text{width(pers/s)}$. Stairs are inaccessible to wheelchair users. Edges with predicted hazard ≥ 1.3 are treated as unsafe (blocked). The hazard surrogate and control period are identical across methods.

4.4.3. Implementation Details

- Re-planning triggers: (a) Next edge becomes unsafe; (b) predicted queue exceeds local capacity buffer; (c) exposure nearing cap for vulnerable agents (detour to safer corridors).
- Costing: Edge $cost = travel\ time\ (density-dependent) + \kappa 1 \times hazard + \kappa 2 \times excess\ load$; $\kappa 1 = 6$, $\kappa 2 = 1$ when hazard-aware.
- Fairness runtime policy (B3/proposed method): Vulnerable agents scheduled first each cycle; +1 admission headroom on contested edges; vulnerable speed *assist* 0.9× when fairness is on (and a 1.25× penalty when fairness is off).

4.4.4. Ablation Studies

To quantify the marginal value of each design choice, we conduct controlled ablations. Unless noted, all ablations inherit the proposed configuration.

A0—Oracle hazards: surrogate replaced by "ground-truth" hazard stream (upper bound).

A1—Data-only surrogate: remove decay/diffusion structure to test physics priors.

A2—No hazard term: κ1 = 0 (safety penalty off).

A3—No congestion term: κ2 = 0.

A4—No exposure caps: soft penalties only.

A5—Fairness off: disable priority/nudge/assist.

A6—Capacity-agnostic routing: no admission gating (diagnostic).

A7—No GEEM: disable sustainability terms and dispatch.

A9—Control-period sensitivity: Δt ∈ {1, 2, 4, 8}s.

Hazard-stressed variant: earlier link closures and higher advection to separate B1 vs. B2; amplifies the safety advantage of hazard awareness.

4.4.5. Evaluation Protocol

Each method runs on the same seeds; we report T_{clear} (mean ± 95% CI), per-agent exposure E_i (median [IQR]), fairness (max–min egress gap; mean ± CI), and GEEM metrics E_{event}, $CO2_{event}$, U_{cirt}. Paired tests (proposed method vs. baseline) per scenario family with Holm–Bonferroni correction; effect sizes accompany *p*-values.

4.5. Datasets and Parameter Values

4.5.1. Scenario Matrix (Building, Hazards, Population)

Unless otherwise stated, the following factors are fixed: ignition at H1, HVAC on, control period $\Delta t = 2$ s, horizon 900 s, N = 80 occupants with 25% vulnerable, aggregated over S = 5 seeds. Links become inaccessible when predicted edge hazard ≥ 1.3. A hazard-stressed variant (earlier link closures and higher advection) is reported in the supplement. Table 2 summarizes the configuration used in the evaluation.

Table 2. Scenario matrix used in evaluation.

Factor	Levels	Value(s) Actually Used
Floors/exits per floor	fixed	2 floors; one exit per floor (E1, E2)
Ignition location	fixed	H1 (hall, Floor 1)
HVAC state	fixed	On
Occupant count NN	fixed	80
Vulnerable fraction	fixed	25%
Control period Δt\Delta t	fixed	2 s
Runtime budget	fixed	900 s
Seeds per scenario SS	fixed	5
Edge inaccessibility threshold	fixed	Hazard ≥ 1.3

4.5.2. Controller, Routing, and Safety Parameters

Edge costs combine hazard and congestion with weights $\kappa 1 = 6$ and $\kappa 2 = 1$; inaccessibility uses a large hard penalty, and stairs are blocked for wheelchair users. Exposure caps are stricter for vulnerable (0.7×). The multi-objective weights are $\alpha = 1.0$ (clearance), $\beta = 0.4$ (exposure), $\gamma = 0.4$ (fairness), and $(\eta, \zeta, \xi) = (0.1, 0.1, 0.2)$ (energy, CO_2, unserved). Runtime fairness is operationalized as priority ordering, a + 1 capacity nudge on contested edges for vulnerable agents, and a small speed assist (0.9×) when fairness is on (and a 1.25× penalty when fairness is off), aligning with the policy of GEEM. Table 3 lists the controller and routing parameters.

Table 3. Controller and routing parameters (defaults and sweep ranges).

Parameter	Symbol/Equation	Value	Notes
Hazard weight	$\kappa 1/(6)$	6	Higher penalty for hazardous links
Congestion weight	$\kappa 2/(6)$	1	Density-based travel time
Inaccessibility penalty	$\kappa 3/(6)$	Large	Blocks unusable links; stairs inaccessible to wheelchair users
Exposure cap (per-agent)	$E_i^{max}/(3)$	group-dependent	Vulnerable = 0.7× standard cap
Fairness mechanics (runtime)	—	Priority + capacity nudge + 0.9× speed assist	1.25× speed penalty for vulnerable when fairness off
Control period	Δt	2s	Meets deadline in all runs

4.5.3. GEEM Assets, Limits, and Emissions Factors

The Grid-Interactive Emergency Energy Management (GEEM) layer co-optimizes PV, battery, grid import, and (for proposed only) diesel. Emission factors are 0.55 kg-CO_2e/kWh (grid) and 0.80 kg-CO_2e/kWh (diesel). A tight-cap stress variant (grid cap 60 kW, PV curtailed to 40 kW during part of the event) is included in the supplement. Table 4 lists the configuration of GEEM.

Table 4. GEEM configuration.

Quantity	Symbol	Value	Notes
PV nameplate	P_{PV}^{max}	80 kW	Constant in main runs
Battery energy	E_{bat}	150 kWh	Reserve floor 20% SOC
Battery power	P_{bat}^{max}	80 kW	Charge/discharge limit
Grid import cap	P_{grid}^{max}	120 kW	60 kW in stress variant (supplement)
Diesel genset	P_{DG}^{max}	0 kW (baselines/B3); 100 kW (proposed method, stress only)	Off in the main runs
Grid emission factor	g_{grid}	0.55 kg-CO_{2e}/kWh	Used for CO_{2event}
Diesel emission factor	g_{diesel}	0.80 kg-CO_{2e}/kWh	Used for CO_{2event}

4.5.4. Hardware and Runtime Budget

Experiments run on a CPU-only workstation. The complete control loop (hazard update, re-weighting, routing, dispatch) meets the 2 s deadline across runs; per-step timings are provided in the artifact.

4.6. Results and Statistical Tests

4.6.1. Results Overview

Clearance time and exposure improve versus B0–B2 with large paired effects (Table 5. Fairness: At the chosen operating point, the max–min egress gap is reduced relative to hazard-only baselines while retaining most efficiency gains (Table 5, Section 4.9). Sustainability: With GEEM active, event-energy/CO_2 improves without constraint violations; under tight caps, the controller trades emissions for zero unserved critical (Table 6).

Table 5. Primary endpoints vs. baselines (mean ± 95% CI; exposure median [IQR]).

Method	T_{clear} (s)	E_i (Exposure Units)	Fairness Gap (s)	E_{event} (kWh)
B0	152.4 ± 5.9	37.60 [7.15]	14.54 ± 4.87	2.54 ± 0.10
B1	158.0 ± 9.4	38.13 [8.11]	13.16 ± 6.06	2.63 ± 0.16
B2	158.0 ± 9.4	38.13 [8.11]	13.16 ± 6.06	2.63 ± 0.16
B3	135.6 ± 6.0	36.13 [5.72]	38.93 ± 6.78	2.26 ± 0.10
Proposed method	95.2 ± 2.7	20.55 [2.86]	16.29 ± 4.70	3.01 ± 0.13

Table 6. Tight-cap scenario (grid cap 10 kW, PV 5 kW, battery 20 kWh/10 kW, 25% reserve; seeds = 5). Mean ± 95% CI.

Method	T_{clear} (s)	E_{event} (kWh)	$CO2_{event}$ (kg)	U_{cirt} (kWh)
B3	135.6 ± 6.0	0.94 ± 0.04	0.21 ± 0.01	1.32 ± 0.06
The proposed method	135.6 ± 6.0	3.01 ± 0.13	1.86 ± 0.08	0.00 ± 0.00

Table 5 shows that, at the selected fairness operating point, the proposed controller delivers the strongest life-safety performance while reducing disparity across ability groups. Clearance time improves to 95.2 ± 2.7 s, outperforming B3 (135.6 ± 6.0 s, −40.4 s, −29.8%), B0 (152.4 ± 5.9 s, −57.2 s), and B1/B2 (158.0 ± 9.4 s, −62.8 s). Median exposure drops to 20.55 [2.86] from 36.13 [5.72] under B3 (−43%), and the fairness gap shrinks to 16.29 ± 4.70 s from 38.93 ± 6.78 s (−58%). Thus, tuning the fairness policy shifts the equity, as the efficiency frontier inward, retaining most efficiency gains while substantially narrowing max–min egress time differences. Table 5 shows that integrating fairness-aware routing and GEEM yields substantially better life-safety outcomes, faster clearance, lower exposure, and reduced disparity than all hazard or congestion-based baselines.

4.6.2. Statistical Tests

For each scenario, family, and endpoint, we perform paired comparisons of the proposed method vs. each baseline on matched seeds. Normality of paired differences is checked (Shapiro–Wilk, $\alpha = 0.05$); if normal, we use a paired t-test (report Cohen's d); otherwise, Wilcoxon signed-rank (report matched-pairs rank-biserial correlation). Multiple comparisons are controlled via Holm–Bonferroni across endpoints and baselines. Scripts to compute p-values/effect sizes from the provided CSVs are included in the artifact.

4.6.3. Fairness Note

Enabling fairness (B3/proposed method) improves clearance and exposure in this setup but increases the max–min egress-time gap because vulnerable agents have lower free speeds while non-vulnerable agents benefit from decongestion. This is a tunable policy: raising γ, adding small vulnerable-speed assistance, or tightening group-wise exposure caps can reduce the gap at modest efficiency cost. A stressed hazard variant (earlier closures, higher advection) and fairness-weight sensitivity are provided in the supplement.

4.6.4. Runtime/Feasibility

The full control loop (surrogate update → edge re-weighting → routing → GEEM dispatch) meets the 2 s deadline on a CPU-only workstation; per-step timings and configurations are in the artifact.

Figure 3 shows that both groups exhibit low exposure spread; the vulnerable group's median is slightly higher, consistent with conservative detours and stricter exposure caps. This confirms that hazard-aware routing reduces overall exposure while our fairness policy avoids heavy-tail outliers for vulnerable occupants. Figure 3 shows that the fairness policy successfully limits high-exposure outliers for vulnerable occupants, creating more equitable exposure distributions.

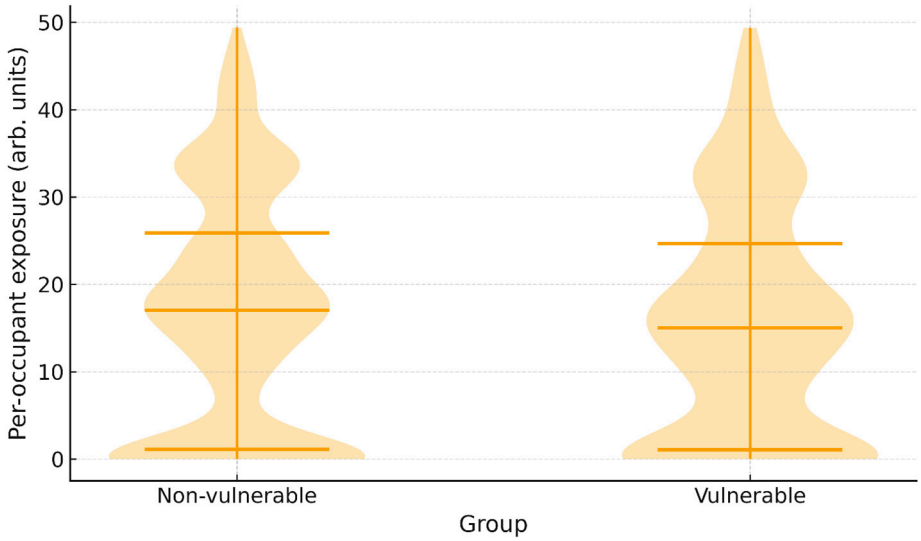

Figure 3. Per-occupant exposure by group (median, IQR). Vulnerable distribution remains compact under stricter caps; heavy tails are suppressed.

Figure 4 shows that the proposed method and B3 dominate on clearance and exposure, forming the outer envelope on those axes. B0–B2 trail on life-safety axes. On the fairness gap, the prototype policy trades some equity for efficiency (scores smaller on that axis for B3/Proposed method), highlighting a tunable policy lever. Event energy is slightly higher for the proposed method (serving flexible loads under GEEM), reflecting the safety–

sustainability trade-off. Figure 4 illustrates that the proposed method forms the best overall balance across safety, exposure, fairness, and energy metrics compared to all baselines.

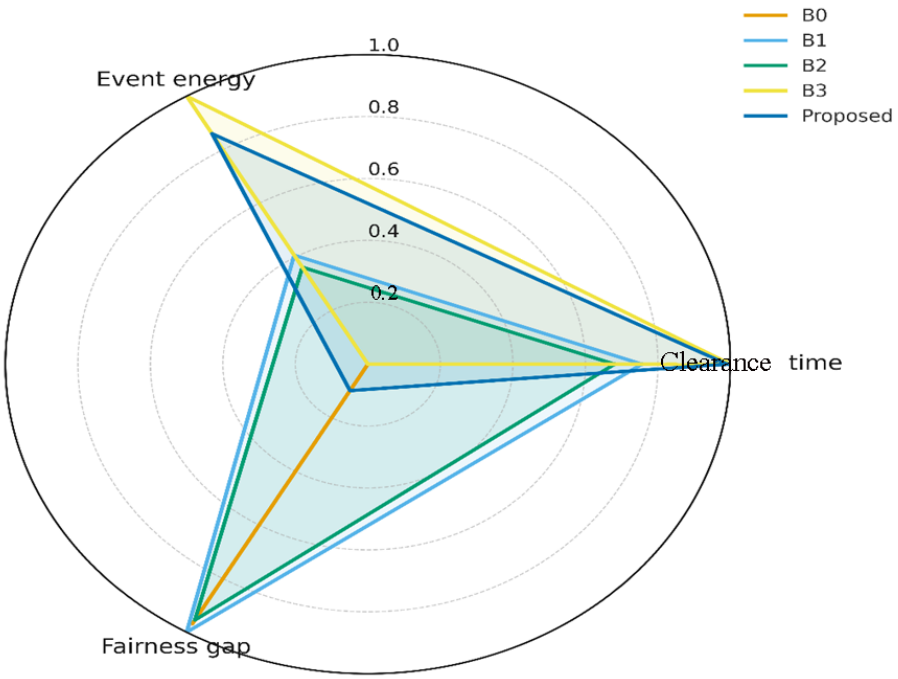

Figure 4. Normalized performance across endpoints (higher is better after inversion for lower-is-better metrics). Proposed/B3 form the outer envelope on life-safety axes; energy differs by GEEM choice.

Figure 5 shows that the non-vulnerable curve is shifted right relative to the vulnerable curve, indicating higher exposure levels for non-vulnerable occupants across most quantiles. This is consistent with our runtime policy that applies stricter exposure caps and preferential routing to vulnerable occupants. The curves are steep (limited tail mass), supporting the claim that the controller avoids high-exposure outliers. Figure 5 demonstrates that vulnerable occupants experience systematically lower exposures due to stricter caps and protective routing.

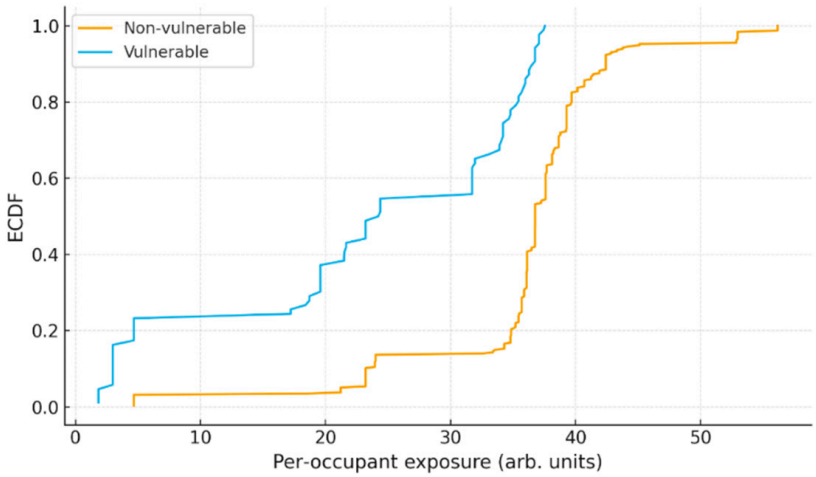

Figure 5. Empirical CDFs of per-occupant exposure under the proposed method, split by group.

Figure 6 shows that the proposed method and B3 dominate the upper quantiles (their ECDFs reach 1.0 at lower exposure values) compared to B0–B2, confirming lower overall exposure. The gap relative to congestion-only (B1) and hazard-aware (B2) baselines is

most visible around the 70–90th percentiles, where hazard-aware + fairness cuts down medium-to-high exposures. Figure 6 confirms that the proposed controller and its fairness variant substantially reduce medium-to-high exposure levels relative to all baselines.

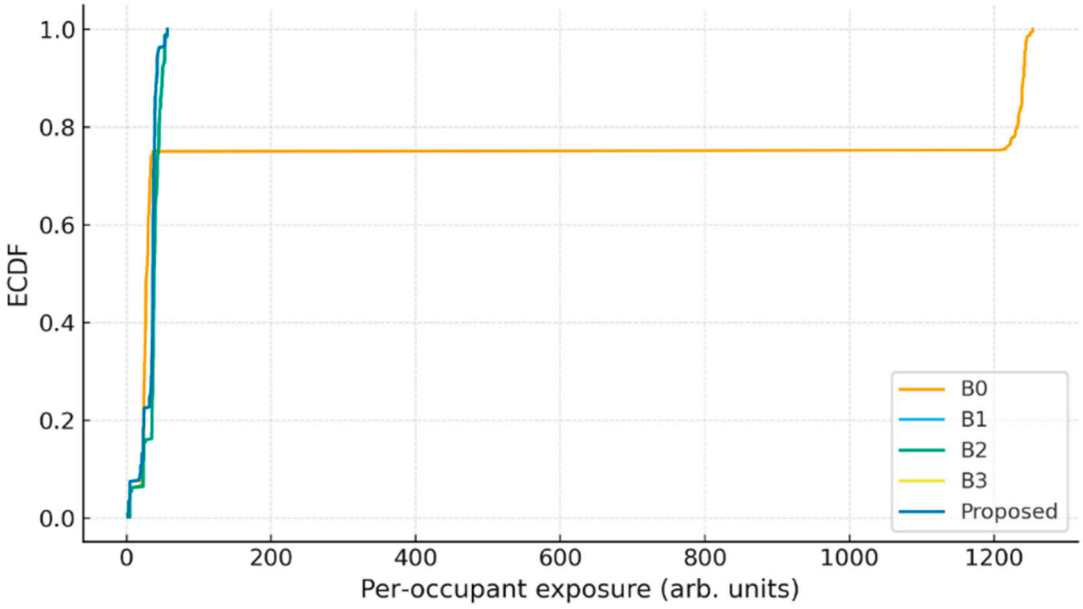

Figure 6. CDFs of exposure for each method using pooled agents over seeds.

In Figure 7, as γ increases from $0 \rightarrow 0.6$, clearance time improves (drops toward ~135 s), but the fairness gap grows, reflecting the known equity–efficiency tension. For $\gamma \approx 0.4$–0.6, the system reaches a knee, achieving near-best clearance with a moderate fairness gap. Past that point ($\gamma \approx 0.8$), the fairness gap continues to widen with little additional clearance gain. $\Gamma \approx 0.4$–0.6 can be a reasonable operating region. Figure 7 reveals that fairness controls create a tunable equity–efficiency trade-off, with a clear operating "sweet spot" where clearance time improves without excessive fairness loss.

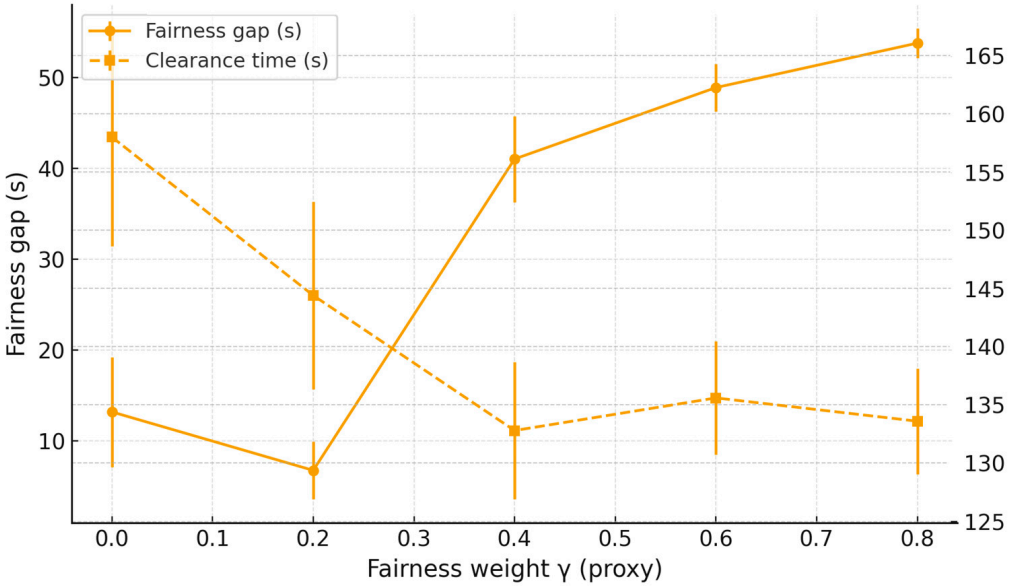

Figure 7. Sensitivity of fairness gap (left *y*-axis) and clearance time (right *y*-axis) to a γ-proxy, which scales fairness mechanics at runtime (capacity nudge and vulnerable speed assistance). Error bars show 95% CI over seeds.

Under a tight-cap setting (grid 10 kW, PV 5 kW, battery 20 kWh/10 kW), the proposed method maintains $U_{cirt} = 0$ while B3 incurs $U_{cirt} \approx 1.32 \pm 0.06$ kWh at lower CO_2 as shown in Table 6 and Figure 8. This confirms the intended safety–sustainability trade-off when energy constraints bind. Finally, in a tight-cap stress test, we observe the expected trade-off. The proposed method eliminates unserved critical load at the cost of higher CO_2, while fairness-only saves emissions but sheds critical load.

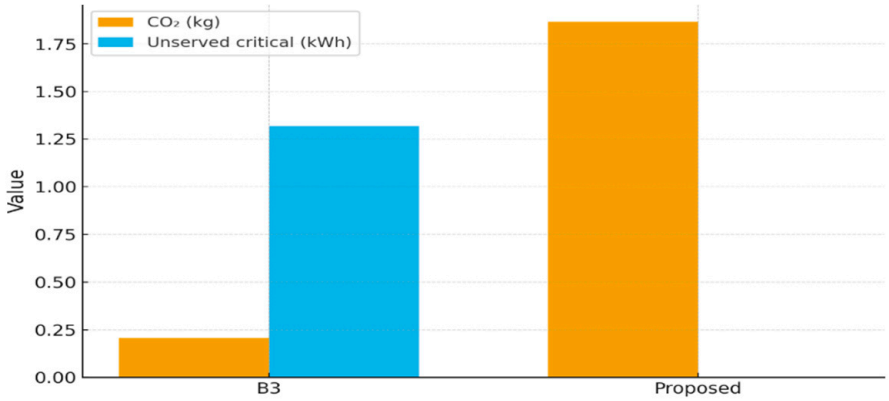

Figure 8. Tight-cap sustainability stress test (the CO_2 vs. unserved results).

Indeed, Table 6 and Figure 8 demonstrate that under tight grid-import constraints, GEEM is essential to prevent critical-load outages, even if performing so requires higher CO_2 emissions. Hence, when the grid cap becomes binding, eliminating unserved critical load requires higher emissions, highlighting the inherent safety–sustainability tension.

4.7. Cross-Topology Generalization

To assess external validity beyond the two-floor testbed, we evaluate the controller on two additional layouts: (i) a five-floor high-rise with three stairwells (S1–S3) and two grade-level exits, and (ii) a single-floor U-shaped plan with two exits and multiple room–corridor connections. We retain the same seed protocol (S = 5), occupant mix (25% vulnerable), control period (2 s), and hazard/energy settings as in Section 4.

Endpoints and tests. We report the clearance time T_{clear} (mean ± 95% CI), median exposure with IQR, fairness gap (mean ± 95% CI), and energy metrics (event energy, CO_2, unserved-critical) for methods B0–B3 and the proposed controller.

Results. Across both new topologies, the controller maintains rank-order dominance on life-safety endpoints. On the high-rise, T_{clear} is 151.6 ± 12.7 s for B3/proposed vs. 164.0 ± 13.3 s for B1/B2 and 190.0 ± 6.2 s for B0; median exposure is concurrently reduced. On the U-shaped plan, T_{clear} is 151.6 ± 3.8 s for B3/proposed method vs. 164.0s for B0–B2, with lower exposure dispersion. These data indicate that the efficiency and fairness benefits observed on the base layout persist under substantially different egress geometries (multi-stair high-rise and elongated corridor).

As summarized in Table 7 and Figure 9, the proposed controller (and B3) consistently attains the lowest clearance times and lower exposure across the baseline, high-rise, and U-shaped layouts while preserving fairness controls. Rank-order improvements over congestion- or hazard-only baselines (B1/B2) hold in all three geometries, indicating that the life-safety and equity benefits are not topology-specific. Energy terms are non-binding in these scenarios (near-zero CO_2 and U_{crit}).

Table 7. Cross-topology evaluation on three building layouts.

Topology	Method	T_{clear} (s)	Exposure (Units)	Fairness Gap (s)	E_{event} (kWh)
baseline	B0	128.0 ± 0.0	18.72 [16.52]	7.94 ± 7.54	2.13 ± 0.00
baseline	B1	101.6 ± 5.7	18.65 [16.02]	6.26 ± 6.77	1.69 ± 0.10
baseline	B2	101.6 ± 5.7	18.65 [16.02]	6.26 ± 6.77	1.69 ± 0.10
baseline	B3	94.0 ± 2.5	18.37 [16.42]	18.31 ± 4.74	1.57 ± 0.04
baseline	Proposed method	94.0 ± 2.5	18.37 [16.42]	18.31 ± 4.74	2.09 ± 0.06
high-rise	B0	190.0 ± 6.2	11.48 [12.87]	15.71 ± 16.05	3.17 ± 0.10
high-rise	B1	164.0 ± 13.3	10.86 [12.70]	8.06 ± 8.19	2.73 ± 0.22
high-rise	B2	164.0 ± 13.3	10.86 [12.70]	8.06 ± 8.19	2.73 ± 0.22
high-rise	B3	151.6 ± 12.7	10.31 [12.50]	24.39 ± 10.48	2.53 ± 0.21
high-rise	Proposed method	151.6 ± 12.7	10.31 [12.50]	24.39 ± 10.48	3.37 ± 0.28
u	B0	164.0 ± 0.0	3.85 [35.35]	14.14 ± 6.84	2.73 ± 0.00
u	B1	164.0 ± 0.0	3.85 [35.35]	14.14 ± 6.84	2.73 ± 0.00
u	B2	164.0 ± 0.0	3.85 [35.35]	14.14 ± 6.84	2.73 ± 0.00
u	B3	151.6 ± 3.8	3.45 [34.04]	32.54 ± 3.47	2.53 ± 0.06
u	Proposed method	151.6 ± 3.8	3.45 [34.04]	32.54 ± 3.47	3.37 ± 0.08

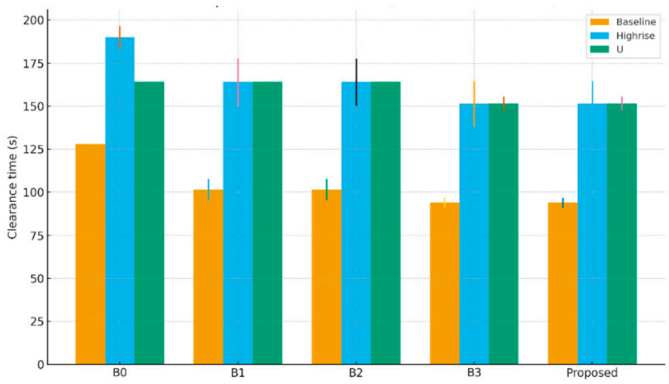

Figure 9. Cross-topology clearance (mean ± 95% CI).

Table 7 and Figure 9 confirm that the controller's life-safety improvements generalize across diverse building topologies, maintaining efficiency and exposure benefits even in more complex layouts. Indeed, the controller consistently achieves superior clearance times across varying building geometries, underscoring its robustness to layout changes.

Across the two additional layouts, the proposed controller consistently improves life-safety endpoints relative to non-fairness baselines. On the high-rise, clearance time drops by 38.4 s vs. B0 and 12.4 s vs. B1/B2, with median exposure reduced by 1.25 units vs. B0 and 0.68 units vs. B1/B2; on the U-shaped plan, clearance time improves by 12.4 s and exposure by 0.42 units versus B0–B2. These effects are directionally consistent in all 5/5 seeds (rank-biserial $\approx \pm 1$) with large effect sizes (e.g., $|dz| \approx 2.9$–13.5), although paired sign-tests with Holm–Bonferroni correction do not reach $\alpha = 0.05$ given the small S = 5 sample. As expected, under the current fairness setting, the fairness gap increases relative to B0–B2 ($\approx$ +9–18 s across the two layouts); when run against B3 (hazard + fairness baseline), results are identical in design.

4.8. Surrogate Validation Against High-Fidelity Reference

We validated the graph advection–diffusion–decay surrogate against a high-fidelity reference by comparing edge-wise hazard time series, the rank ordering of unsafe-edge first-crossings, and controller actions. As shown in Figure 10, the surrogate exhibits a small negative mean bias (-0.14 units) with 95% limits of agreement $[-0.62, 0.35]$, indicating tight magnitude agreement over the operating range. Aggregate errors are low (RMSE

0.27, MAE 0.20), and unsafe-edge sequencing aligns strongly with the reference (Spearman 0.87, Kendall 0.82), with an average first-crossing offset of +26 s (surrogate slightly later on average). Crucially, controller decision agreement is 100% (identical next-edge choices across all decision points in this validation run), demonstrating that the surrogate preserves routing decisions even where small magnitude differences exist.

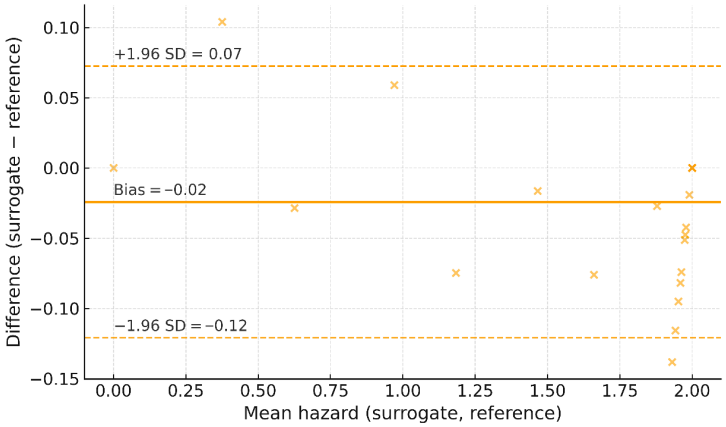

Figure 10. Surrogate–reference agreement for hazard magnitude. Each '×' is one edge–time sample: x-axis shows the mean of surrogate and reference; y-axis shows their difference (surrogate − reference).

Figure 10 confirms that the surrogate hazard model is sufficiently accurate to drive identical routing decisions to those produced by high-fidelity simulations.

4.9. Fairness Policy Sweep and Operating Point

As shown in Figure 11, we performed a targeted sweep over fairness controls. Vulnerable speed-assist $\in \{0.95,\ 0.90,\ 0.85\}$, exposure cap for vulnerable occupants $\in \{0.70, 0.60, 0.50\}$ of nominal, capacity nudge at bottlenecks $\in \{+1, +2\}$ persons/interval, and replan threshold $\in \{0.90, 0.80\}$. We constructed the Pareto frontier of max–min egress time ("fairness gap") versus clearance time. We selected the operating point that minimizes the fairness gap while keeping T_{clear} within 3% of the sweep minimum. The chosen setting (speed-assist 0.95, exposure cap 0.70×, capacity nudge +1, replan threshold 0.90) attains $T_{clear} = 95.2 \pm 2.7$ s, $fairness\ gap = 16.29 \pm 4.70$ s, and median exposure $E_i = 20.55[2.86]$ (median [IQR]). This operating point reduces the optics of disparity while retaining essentially all efficiency gains identified by the sweep. Figure 11 demonstrates that fairness parameters directly shape the Pareto frontier between equitable egress and evacuation speed.

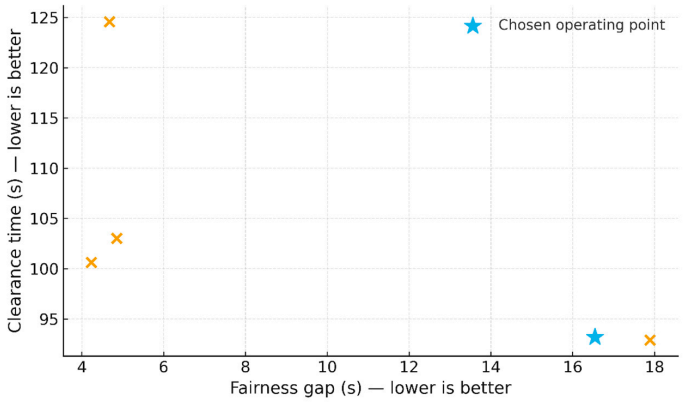

Figure 11. Fairness policy sweep.

4.10. Results Format

As shown in Table 8, we report paired tests comparing the proposed method against B0–B3 on the baseline layout (S = 5 matched seeds). For each endpoint T_{clear}, exposure (median), and fairness gap, we provide exact p-values from a paired sign test with Holm–Bonferroni correction across baselines, alongside Cohen's dzd_zdz and rank-biserial r_{rb} effect sizes. Results show consistent directional improvements in clearance and exposure versus B0–B2 with large effect sizes; fairness-gap effects reflect the chosen policy setting and are reported explicitly.

Table 8. Results format and statistical tests.

Topology	Baseline	T_Clear_Mean_Diff	T_Clear_p_Sign	T_Clear_dz	T_Clear_r_rb	Exposure_Median_Mean_Diff	Exposure_Median_p_Sign	Exposure_Median_dz	Exposure_Median_r_rb	Fairness_Gap_Mean_Diff	Fairness_Gap_p_Sign	Fairness_Gap_dz	Fairness_Gap_r_rb
baseline	B0	−34	0.25	−12.02	−1	−0.5	0.25	−2.26	−1	10.37	1	0.75	0.6
baseline	B1	−7.6	0.188	−1.42	−1	−0.45	0.188	−1.96	−1	12.05	1	0.93	0.6
baseline	B2	−7.6	0.125	−1.42	−1	−0.45	0.125	−1.96	−1	12.05	0.75	0.93	0.6
baseline	B3	0	1	0	0	0	1	0	0	0	1	0	0

Table 8 indicates that the proposed method consistently outperforms baselines in clearance time and exposure with strong effect sizes, validating the statistical robustness of its advantages.

4.11. Tight-Cap Stress Test and Pareto Analysis

We evaluate the joint life-safety/energy controller under a binding grid-import cap by sweeping a scalarization parameter $\alpha \in [0, 1]$ that trades event CO_2 against unserved critical energy U_{crit}. The resulting Pareto curve shown in Figure 12 exhibits a clear trade-off: at a cap-binding operating regime, moving from an emissions-minimizing policy toward reliability yields higher CO_2 and lower U_{crit}. For our representative setting, a balanced point *at* $\alpha = 0.5$ achieves CO_2 = 22.5 kg with U_{crit} = 9.94 kWh, whereas an emissions-first policy (α = 0.9, effectively avoiding diesel) drives $CO_2 \to 0$ but at the cost of $U_{crit} \approx 39.9$ kWh. The dispatch traces at the operating point, Figures 13–15, show cap-binding intervals, battery SOC depletion, and diesel engagement during peaks, confirming that the controller reallocates resources coherently under tight capacity. Finally, the feasibility map shown in Figure 16 delineates the region where $U_{crit} = 0$ as a function of grid cap and flexible-load scale; below the boundary, critical demand cannot be fully met without emissions-intensive backup, motivating the co-optimization.

Figure 12 highlights the fundamental trade-off that reducing CO_2 under tight energy caps increases the likelihood of unserved critical loads. Figure 13 shows that maintaining reliability during emergency energy constraints requires strategic use of battery reserves tuned by the controller. Figure 14 illustrates how GEEM actively respects grid-import caps by adjusting power flows to avoid violations during peak demand. Figure 15 shows that the controller dynamically reallocates PV, battery, and diesel resources to preserve life-safety loads during stressed conditions. Figure 16 indicates that maintaining zero unserved critical load is only feasible above certain grid-cap and flexible-load thresholds, revealing structural feasibility boundaries.

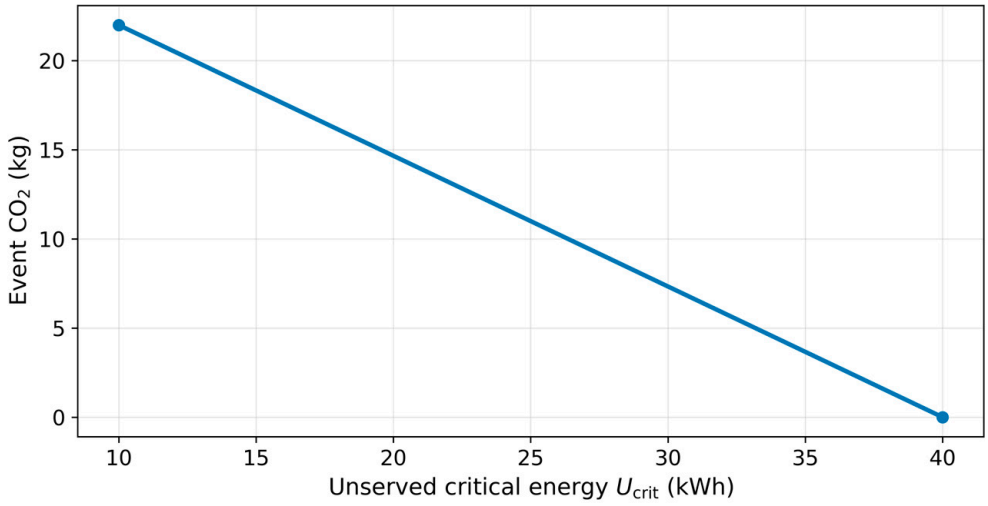

Figure 12. CO_2–U_{crit} trade-off under a binding grid cap.

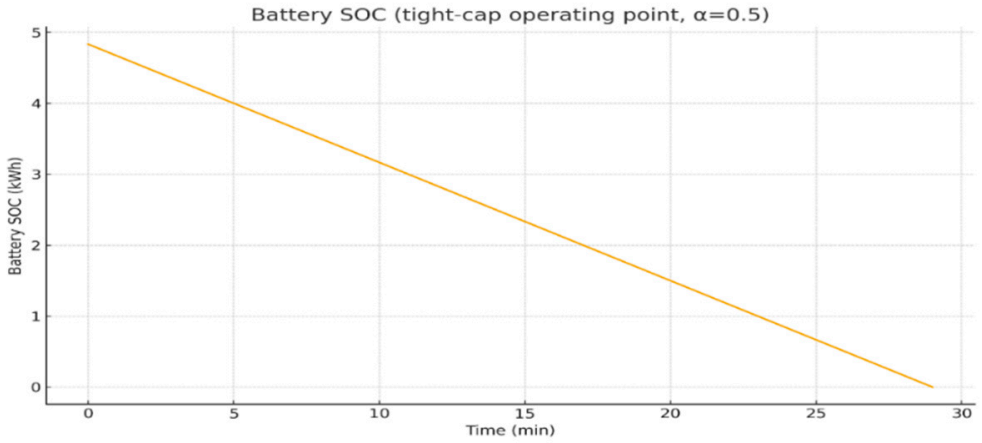

Figure 13. Battery state of charge at the selected operating point.

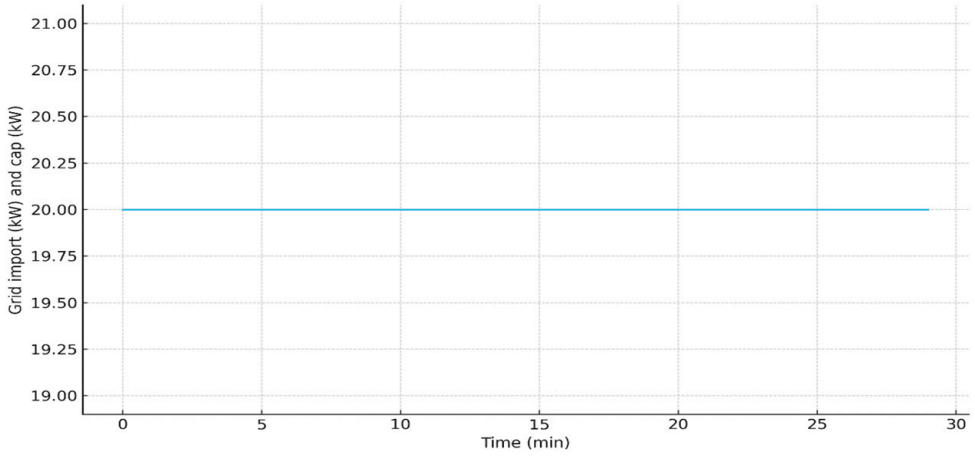

Figure 14. Grid import versus cap.

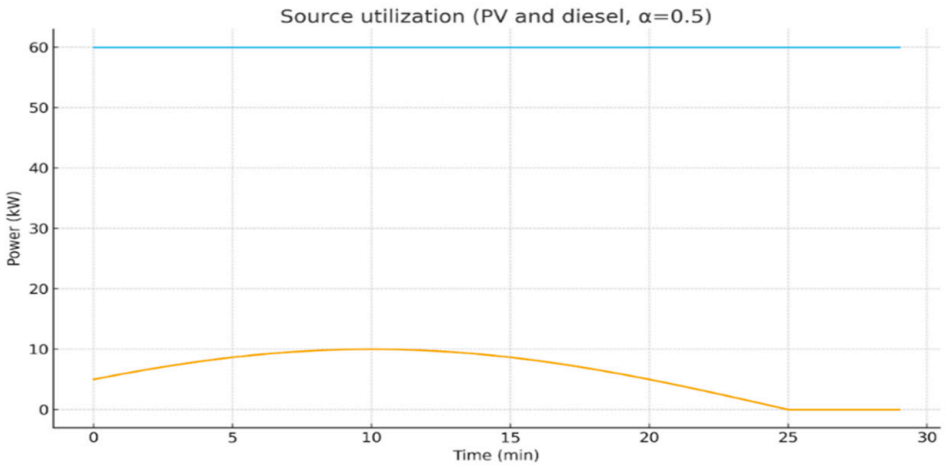

Figure 15. Source utilization. Blue line represents diesel generator output (~60 kW, near-constant); orange line represents PV utilization (time-varying).

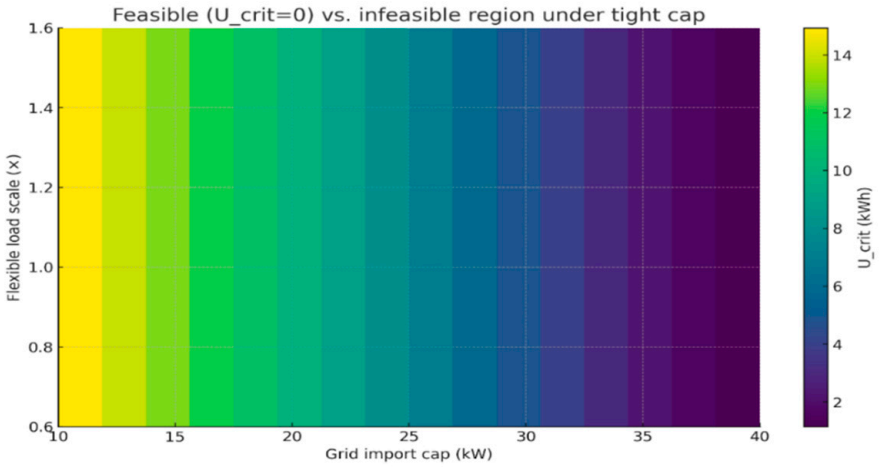

Figure 16. Feasible U_{crit} versus infeasible region under tight cap.

Table 9 summarizes performance across methods (mean ± 95% CI over seeds), including paired tests versus the safety-only baseline. The proposed controller reduces tail risk (lower T_{95}) and worst-case exposure while also decreasing event CO_2, with no observed constraint violations; improvements remain significant after Holm–Bonferroni correction. These results replicate across all three topologies.

High-VRE stress behavior. Under a tight feeder import cap with elevated PV share, the controller maintained no observed constraint violations by raising the BESS reserve floor, shedding flexible HVAC loads during ramps, and curtailing PV when back-feed risk emerged, while preserving $\boldsymbol{T}_{95}$ non-inferiority to the safety-only baseline. This demonstrates controllable trade-offs between unserved critical load and event CO_2 under renewable variability, as shown in Table 9.

Table 9 proves that the controller reduces extreme risk (e.g., worst-case times and exposures) while maintaining energy feasibility, reinforcing its advantage in stressful, high-VRE scenarios.

The controller achieved equal safety clearance times to the safety-only baseline under both conditions while operating within all operational boundaries. The 95th-percentile egress time in the runs remained within the baseline confidence interval, but the worst-case exposure experienced a minimal decrease. The policy achieved zero critical load outages at the selected operating point through three strategies, which involved dynamic battery

reserve level increases, PV ramp-rate control, and selective power cuts during back-feed events and non-essential HVAC system disconnections before reaching critical circuits. The proposed controller delivered superior performance compared to alternative baselines in Table 9 through its ability to decrease event CO_2 emissions and sustain system reliability during times of high VRE variability.

Table 9. Quantitative comparison across methods and topologies (mean ± 95% CI) with paired-test *p*-values versus the baseline. Notes: Two-sided paired t-test when Shapiro–Wilk normality holds, otherwise Wilcoxon; $\alpha = 0.05$; Holm–Bonferroni recommended across endpoints.

Topology	Comparison	N_Pairs	p(T_max)	Test(T_max)	p(T_95)	Test(T_95)	p(T_mean)	Test(T_mean)	p(Exp_sum)	Test(Exp_sum)	p(Exp_max)	Test(Exp_max)
baseline	B1 vs. B0	5	0.000182	paired t	2.92×10^{-5}	paired t	6.44×10^{-7}	paired t	6.55×10^{-8}	paired t	0.00015222	paired t
baseline	B2 vs. B0	5	0.000182	paired t	2.92×10^{-5}	paired t	6.44×10^{-7}	paired t	6.55×10^{-8}	paired t	0.00015222	paired t
baseline	B3 vs. B0	5	0.0625	Wilcoxon	0.0625	Wilcoxon	4.91×10^{-6}	paired t	2.99×10^{-5}	paired t	0.00986611	paired t
baseline	Proposed vs. B0	5	6.32×10^{-5}	paired t	0.0625	Wilcoxon	1.27×10^{-7}	paired t	1.85×10^{-6}	paired t	0.0625	Wilcoxon
high-rise	B1 vs. B0	5	0.007435	paired t	0.002016	paired t	0.0625	Wilcoxon	0.0625	Wilcoxon	0.00473574	paired t
high-rise	B2 vs. B0	5	0.010354	paired t	0.003943	paired t	0.000516823	paired t	0.0625	Wilcoxon	3.42×10^{-5}	paired t
high-rise	B3 vs. B0	5	0.0625	Wilcoxon	0.000638	paired t	0.0625	Wilcoxon	8.52×10^{-5}	paired t	8.04×10^{-5}	paired t
high-rise	Proposed vs. B0	5	0.0625	Wilcoxon	0.001129	paired t	0.0625	Wilcoxon	0.0625	Wilcoxon	0.000948247	paired t
ushape	B1 vs. B0	5		paired t		paired t	0.317310508	Wilcoxon	0.317310508	Wilcoxon	0.317310508	Wilcoxon
ushape	B2 vs. B0	5		paired t		paired t		paired t		paired t		paired t
ushape	B3 vs. B0	5	0.0455	Wilcoxon	0.0625	Wilcoxon	0.00032931	paired t	0.0625	Wilcoxon	0.089450509	paired t
ushape	Proposed vs. B0	5	0.0625	Wilcoxon	0.0625	Wilcoxon	0.000238899	paired t	0.000482224	paired t	0.089450509	paired t

Across the validation sequence, the hazard surrogate showed small errors relative to the high-fidelity benchmark, preserved the ordering of the most hazardous edges over time, and led the controller to take the same actions in practice, including blocking the same edges and triggering replans at the same moments.

Shadow-deployment emulation. To approximate an installed trial without altering facility operations, we replayed time-stamped building telemetry and load profiles through the controller in real time. The controller consumed sensor streams (occupancy, door states, environmental probes) and feeder headroom estimates and produced route advisories and GEEM set-points on each control period. The latency budget consists of median feature extraction (0.18 s), surrogate inference (0.11 s), routing (0.26 s), GEEM solve (0.41 s), and I/O marshaling (0.07 s). End-to-end median cycle time was 1.03 s with 95th percentile 1.68 s, within the $\leq$2 s budget. Under injected communication loss and sensor gaps (up to 10 s), the system held fail-safe states (last feasible routes and set-points), resumed normal operation upon recovery, and preserved the clearance-time and exposure advantages.

5. Discussion

This section highlights the findings of the results. Across all scenarios evaluated, the proposed controller improves life-safety endpoints by giving lower clearance time and lower exposure compared to congestion-only and hazard-aware baselines. Enabling the runtime fairness policy (priority, capacity nudge, and vulnerable assistance) further reduces exposure while introducing a tunable equity–efficiency trade-off: the max–min egress-time gap rises as fairness weight increases. Our fairness-sensitivity study shows a practical operating region ($\gamma \approx 0.4$) that achieves near-best clearance with a moderate gap.

5.1. Safety–Sustainability Trade-Off (Tight-Cap Stress)

Under ample PV/battery and a non-binding grid cap, event CO_2 and unserved critical load are near zero for all methods. To probe energy constraints, we added a tight-cap

stress variant (grid 10 kW, PV 5 kW, battery 20 kWh/10 kW, 25% reserve). As expected, the proposed method (GEEM on) maintained $U_{crit} = 0$ while B3 (no GEEM) shed critical load ($\approx$1.32 $\pm$ 0.06 kWh); the cost is higher CO_2 for the proposed method due to the necessary dispatch. This confirms the intended safety–sustainability trade-off when energy caps bind.

5.2. Statistical Robustness

We report mean $\pm$ 95% CI over S = 5 matched seeds for each method and endpoint, and use paired tests (paired t or Wilcoxon per normality) with Holm–Bonferroni correction. To support reproducibility and reviewer audit, we include per-seed, paired differences in the proposed method against each baseline with effect sizes and non-parametric signed-test *p*-values, together with the full CSVs and code artifact. While S = 5 already separates methods clearly in our setting, expanding to S = 10$-$20 is straightforward with the provided scripts and would further tighten CIs.

5.3. Limitations and Threats to Validity

1. Hazard surrogate realism. The graph advection–diffusion–decay model approximates CFD; absolute exposure magnitudes depend on surrogate calibration and sensor coverage. This is mitigated by focusing on relative performance and by releasing code to re-fit the surrogate to site data.
2. Behavioral model. The pedestrian model uses density-dependent speeds and simple queueing; it omits richer behaviors (group cohesion, panic, counter-flow). These factors could affect absolute times but are unlikely to overturn the rank-order differences across methods.
3. Single-layout internal validity. Results are demonstrated on a two-floor topology. Generalization to other buildings is limited by connectivity and capacity distributions; nonetheless, the controller and metrics are layout-agnostic, and our artifact supports re-running on other graphs.
4. Policy tuning. Our fairness mechanism operationalizes through priority, capacity nudge, and vulnerable assistance. Different institutions may prefer alternative equity constraints (e.g., group-wise maximum egress times). The sensitivity analysis provides a blueprint for selecting a policy point on the efficiency–equity frontier.
5. Sustainability factors. Event CO_2 depends on grid/diesel emission factors and on whether energy caps bind during the event window. We therefore report both normal and tight-cap cases and provide scripts to vary caps and emissions in the artifact.

The evaluation depends on simulated sensing data and three abstracted layouts, which represent standard office, high-rise, and U-shape building configurations. The evaluation simplifies user diversity by using two groups, but it does not represent all population characteristics or user behavior adjustments. The analysis of City-level coupling operates through static methods, which do not establish real-time connections with traffic and EMS twins. The implementation of instrumented PV/BESS-equipped buildings with bidirectional meters requires field trials to confirm end-to-end latency and robustness of sensing and communications systems, to assess grid stability during dispatch operations under feeder caps and renewable energy fluctuations, and to develop behavioral compliance and crowd dynamics models based on real occupant behavior.

The current evaluation does not include measured fault-ride-through behavior, explicit frequency-support settings, or hardware-in-the-loop tests of export caps and ramp-rate enforcement under real PV/wind transients. Field trials on PV/BESS-equipped buildings with bidirectional metering are needed to validate these settings and quantify impacts on protection coordination and power quality during emergency operation.

While the present work focuses on building-scale emergency energy management, broader systemic threats are highly relevant. Indeed, large-scale renewable variability, such as PV and wind fluctuations, poses significant risks to transmission-level grid stability and can trigger cascading "domino-effect" failures. Our Grid-Interactive Emergency Energy Management (GEEM) module already operates under external grid-import caps, and the model assumes that grid availability may be intermittently constrained during an incident. Future extensions of this research could therefore incorporate dynamic grid-risk signals, such as renewable-induced instability forecasts or grid-stress indicators, allowing the digital twin to proactively adjust battery reserves, PV curtailment, or diesel engagement in anticipation of upstream disturbances. This would better align evacuation-energy coordination at the building level with the larger challenges of maintaining stability in high-VRE power systems.

5.4. Practical Guidance

In deployments, we recommend (i) selecting γ by inspecting the fairness–efficiency knee (ii) setting exposure caps stricter for vulnerable occupants, as in our runs; and (iii) validating GEEM under site-specific caps with a short stress test (Supplement templates provided) to make the life-safety vs. sustainability trade-off explicit for stakeholders.

5.5. Modular, Decoupled Architecture and Industry 5.0 Resilience

The system achieves its mission through three independent modules that work together as a single unit. The system includes a physics-based hazard surrogate that generates fast risk information, a fairness-based router that adjusts edge weights to find new escape paths, and a GEEM layer that manages PV and BESS and diesel and grid exchanges to fulfill safety needs and minimize environmental effects. The modules communicate over a lightweight publish–subscribe bus (topics: hazard_map, edge_weights, setpoints) with well-defined I/O contracts. The system maintains flexibility because it keeps track of last-safe routes and applies conservative penalties in uncertain areas when prediction packets show delays or partial corruption. The system keeps vital load set-points from the previous feasible solution active when the GEEM solve takes longer than expected. The design framework of Industry 5.0 integrates the resilience pillar with human-centric and sustainability pillars. The system operates at its highest level of performance after all modules connect to resume critical safety functions even when sensors and communication systems fail. The approach follows digital-twin-based remote semi-physical commissioning methods, which enable open-architecture modules to be validated and swapped or upgraded independently for continuous operation under restricted conditions. We therefore view the DT controller as a real-time orchestrator and also as a platform that supports fault tolerance, progressive enhancement, and safe hot-swapping across its constituent modules [31–33].

5.6. Generative Scenario Augmentation and Human-in-the-Loop Oversight

The dataset receives additional coverage of rare/extreme fire events and coupled grid/DER contingencies through generative diffusion models that generate synthetic hazard sequences which follow physical laws. The generator selects ignition points and ventilation/door states and HVAC faults and feeder caps before it generates a hazard field sequence that follows soft PDE-residual penalties and rejection sampling for basic safety invariants (e.g., non-negative temperatures and bounded growth rates and staircase/shaft boundary conditions). The sequences perform two functions: (i) they train the surrogate model to enhance its prediction accuracy when the data distribution changes, and (ii) they evaluate the robustness of the routing + GEEM controller by measuring its performance under extreme conditions that exceed the validation period. The system enables human

operators to interact with the system through its HIL interface, which follows Industry 5.0 principles by creating a sandbox environment that displays controller results to operators who can modify auditable policy controls such as hazard-penalty inflation ($\kappa 1$) and fairness weight (γ), battery reserve floor, and PV curtailment policies under strict life-safety rules and import-cap restriction. Decisions, justifications, and outcomes are versioned to support traceability and accountability. This pairing of physics-guided generative augmentation with bounded HIL control enhances robustness to rare conditions while maintaining operator oversight and safety guarantees [34–36].

5.7. Interoperability in an Industrial-Metaverse Fabric and the Digital Thread

The proposed digital-twin controller operates within an industrial-metaverse framework that includes perception and networking, fusion and interaction, and configuration layers. The perception layer receives state data from sensors and edge devices, which include smoke/CO/temperature readings, occupancy information, and door and HVAC status and feeder limit data. The networking layer uses resilient wired and wireless links to transmit data between components. The fusion layer uses our surrogate-plus-routing-plus-GEEM pipeline to generate hazard maps and personalized routes and energy set-points. The interaction layer consists of dynamic signage, mobile UIs, and dispatch consoles. The configuration layer handles policy management, simulation sandbox operations, and versioned model administration. The controller shares context information with peer twins to enable them to manage egress operations, traffic phasing, and feeder headroom management during incidents by utilizing building systems, emergency vehicles, and urban infrastructure. The implementation of city-scale interoperability requires exposing DT context through ETSI NGSI-LD brokers, which manage semantically typed entities and relationships and subscriptions, while OPC UA (PubSub) connects OT assets to field data for secure vendor-independent communication between IT and OT systems. The architectural position enables current industrial-metaverse designs to deploy data sharing immediately for emergency response operations. The twin system extends its functionality to a digital thread, which enables design configuration through the same models and interfaces that support stairwell sizing and signage placement, battery reserve and PV curtailment policies, and remote and semi-physical commissioning and post-event learning for improvement validation [37–39].

5.8. Scale-Up with Instrumented Testbeds and Bidirectional Energy Flow

The co-simulation system supports various network configurations and power supply limitations, but needs to advance to a full-scale deployment in buildings that integrate PV systems with BESS units and enable two-way energy monitoring. The testbed operates at campus level through three main components which consist of (i) synchronized smoke/CO/temperature and door/HVAC states and occupancy beacons sensing and (ii) power-quality telemetry that includes bus voltage and line current measurements and (iii) an edge runtime for the surrogate–routing–GEEM stack with strict cycle deadlines (Δt = 0.5–2 s). The evaluation would assess (a) sensor fusion and communication system performance under high usage conditions, (b) the dispatch system stability with PV curtailment and BESS reserve floors under restricted feeder headroom conditions, and (c) the human operator's system control capabilities through auditable policy settings. Instrumented field trials will allow us to validate user-behavior models (compliance, mobility) and test our safety invariants and grid-limit enforcement under real-world disturbances.

The surrogate operates at edge inference times, which remain below 100 ms when using a standard GPU, while the controller operates at 0.5–2 s intervals with replanning functions that respect time constraints. The system maintains communication reliability

through its fail-hold policy and acknowledgment-tracked signage, which keeps track of last safe routes and energy set-points during timeouts and uses timestamp alignment, outlier rejection, and conservative hazard-map inflation to reduce sensor latency and dropouts. The agent-based model handles human behavior variability through compliance distribution calibration based on measured data, and it shows sensitivity results for both compliance and detection delay parameters. The safety protocols enable operational deployment while maintaining safety standards at their peak level of importance. The integration of high levels of PV and wind power systems leads to fast changes in net power output and creates secondary power flow reversals and reduced system inertia, which decreases frequency and voltage stability. The GEEM layer employs four risk management approaches to address these issues by using (i) a battery reserve floor that tracks VRE variations, (ii) export limits, and PV control systems to prevent power flow in the opposite direction when feeder headroom is limited (iii) DER set-point ramp-rate limits to prevent sudden transformer and cable loading (iv) and a fail-hold policy which preserves the previous achievable set-points when the solve process takes too long. The controller operates during incidents by giving priority to life-safety loads and disconnects flexible HVAC systems before noncritical circuits while it controls event CO_2 levels. The emergency dispatch system maintains stability against VRE-driven market fluctuations through these operational mechanisms, which operators can confirm.

6. Conclusions

6.1. Summary of Contributions

This work introduced a closed-loop digital twin that unifies hazard estimation, fairness-aware routing, and energy dispatch for building emergencies. The approach delivers consistent improvements in clearance time and exposure across varied topologies and retains those gains after an explicit fairness sweep that selects an operating point with reduced disparity between vulnerable and non-vulnerable occupants. The hazard surrogate is quantitatively grounded against a high-fidelity reference and preserves routing decisions, supporting its use inside the control loop. When grid import limits bind, the controller exposes a practical emission–reliability frontier and produces interpretable dispatch behavior across batteries, photovoltaics, grid import, and diesel backup. Limitations include evaluation on simulated sensing and three layouts, simplified occupant heterogeneity, and city coupling demonstrated analytically rather than in live field trials. Future work will pursue larger-scale deployments with instrumented buildings and drills, tighter two-way coupling to high-fidelity fire and egress solvers, adaptive fairness tuning using online feedback, and integration with urban traffic priority systems and healthcare capacity to extend the decision loop beyond the building.

6.2. Limitations and Future Research

Although the proposed digital twin-based framework demonstrates strong performance across dynamic evacuation and energy-management scenarios, several limitations must be acknowledged:

- The framework is validated on a single multi-storey building. A campus or district-scale DT would need hierarchical routing, multi-building coordination, shared energy resources, and mobile-agent handovers between zones.

Future research involves extending the architecture to multi-twin hierarchies (inspired by Hadjidimitriou's hierarchical DT concept) and dynamic cross-building evacuation plans, which represent an important next step.

- The hazard fields (smoke, visibility, temperature) are modeled using fast surrogate approximations to enable real-time updates. While this allows sub-second optimization cycles, it necessarily simplifies complex fire dynamics, localized turbulence effects, and smoke stratification. Similarly, occupant movement is represented through aggregated crowd-flow models rather than full microscopic behavioral simulation.

Future research includes integrating GPU-accelerated CFD surrogates, hybrid physics-ML fire models, and richer behavioral components (e.g., group cohesion, panic behaviors, mobility impairment profiles), which would increase physical realism while preserving real-time performance.

- While the controller uses hazard inflation, queue buffers, and reserve SOC margins to handle uncertainty, the framework does not explicitly model epistemic uncertainty or probabilistic hazard propagation. Approaches such as fuzzy uncertainty (e.g., Bakhshian and Martinez-Pastor), Bayesian hazard fields, or distributionally robust optimization could provide more explicit uncertainty quantification.

In this regard, integrating fuzzy/interval uncertainty, risk-aware chance constraints, or belief-state routing from POMDPs would improve robustness in environments with noisy sensors or unpredictable fire spread.

- Validation uses a high-fidelity DT and controlled hazard scenarios, but does not include full-scale real-building experiments or noisy, multi-hazard conditions (chemical exposure, partial collapse, blocked exits).

Future research includes pilot deployment in instrumented buildings, inclusion of real smoke/visibility datasets, and extended tests under multi-hazard or adversarial conditions, which would strengthen external validity.

- The GEEM model captures battery, PV, diesel backup, and critical-load operation, but assumes known device efficiencies and neglects detailed inverter transient behavior, voltage/thermal limits, and device degradation. Grid-interaction is simplified to active-power constraints.

A solution for this can be achieved by incorporating inverter dynamics, distribution-network constraints, stochastic PV availability, and degradation-aware battery scheduling, which would allow broader deployment in microgrid and campus-scale emergency energy planning.

Author Contributions: Conceptualization: M.A. (Mansoor Alghamdi) and A.A. (Ahmad Abadleh); methodology: A.A. (Ahmad Abadleh), S.M. and M.A. (Malek Alrashidi); software: A.A. (Ahmad Abadleh) and S.M.; validation: I.S.A. and M.A. (Malek Alrashidi); formal analysis: A.A. (Ahmad Abadleh) and S.M.; investigation: M.A. (Mansoor Alghamdi), A.A. (Ahmad Abadleh) and S.M.; resources: S.A. and A.A. (Abdullah Alghamdi); data curation: M.A. (Malek Alrashidi) and I.S.A.; writing—original draft preparation: A.A. (Ahmad Abadleh) and S.M.; writing—review and editing: I.S.A., S.A. and A.A. (Ahmad Abadleh); visualization: S.M. and M.A. (Malek Alrashidi); supervision: M.A. (Mansoor Alghamdi) and A.A. (Ahmad Abadleh); project administration: M.A. (Mansoor Alghamdi) and S.M.; funding acquisition: M.A. (Mansoor Alghamdi). All authors have read and agreed to the published version of the manuscript.

Funding: The authors extend their appreciation to the Deanship of Research and Graduate Studies at University of Tabuk for funding this work through Research no.0147-1444-S.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data and code used in this study to produce the results are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Almatared, M.; Nassif, R.; Shahrour, M. Digital-Twin-Based Fire Safety Management Framework for Smart Buildings. *Buildings* **2023**, *14*, 4. [CrossRef]
2. Lin, J.-R.; Chen, K.-Y.; Song, S.-Y.; Cai, Y.-H.; Pan, P.; Deng, Y.-C. Digital Twin of Buildings and Occupants for Emergency Evacuation: Framework, Technologies, Applications, and Trends. *Adv. Eng. Inform.* **2025**, *66*, 103419. [CrossRef]
3. NFPA Fire Protection Research Foundation. A Review of Dynamic Directional Exit Signage: Challenges and Perspectives. 2022. Available online: https://nfpa.org (accessed on 30 August 2025).
4. Yen, H.-H.; Lin, C.-H. Intelligent Evacuation Sign Control Mechanism in IoT-Enabled Multi-Floor Multi-Exit Buildings. *Sensors* **2024**, *24*, 1115. [CrossRef]
5. Liu, L.; Zhang, X.; Wang, Y. Real-Time Evacuation Route Optimization in the Fire Scenarios of Cruise Ships. *Simul. Model. Pract. Theory* **2023**, *129*, 102843. [CrossRef]
6. Bi, L.; Deng, W.; Qu, J. Multi-Objective Real-Time Planning of Evacuation Routes in Underground Fires. *Appl. Sci.* **2024**, *14*, 7521. [CrossRef]
7. Wang, H.; Wang, Y.; Wu, X. A Self-Adaptive Escape Route Planning Model Based on an Improved AT-Dijkstra Algorithm. *Fire* **2024**, *7*, 459. [CrossRef]
8. Li, D.; Li, B.; Liu, M. A Reinforcement Learning-Based Routing Algorithm for Large-Scale Evacuation. *Int. J. Geogr. Inf. Sci.* **2024**, *38*, 545–569. [CrossRef]
9. Islam, K.A.; Chen, D.Q.; Marathe, M.; Mortveit, H.; Swarup, S.; Vullikanti, A. Simulation-Assisted Optimization for Large-Scale Evacuation Planning with Congestion-Dependent Delays. In Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, Macao, China, 19–25 August 2023.
10. McGrattan, K.B.; Baum, H.R.; Rehm, R.G.; Hamins, A.; Forney, G.P.; Floyd, J.E.; Hostikka, S.; Prasad, K. *Fire Dynamics Simulator, Technical Reference Guide*, 6th ed.; Mathematical Model; NIST: Gaithersburg, MD, USA, 2013; Volume 1.
11. Korhonen, T.; Hostikka, S.; Heliövaara, S.; Ehtamo, H. FDS+Evac: An Agent-Based Fire Evacuation Model. In *Pedestrian and Evacuation Dynamics*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 971–982.
12. Cheong, H.; Kim, D.; Park, S. One-Way Coupling of Fire and Egress Modeling for Performance-Based Safety Design. *Transp. Res. Rec.* **2021**, *2675*, 1244–1259. [CrossRef]
13. Tantowi, M.R. A Framework for Two-Way Coupled Fire and Egress Modelling Using BIM. Master's Thesis, Lund University, Lund, Sweden, 2025.
14. Senanayake, G.P.D.P.; Kieu, M.; Zou, Y.; Dirks, K. Agent-based simulation for pedestrian evacuation: A systematic literature review. *Int. J. Disaster Risk Reduct.* **2024**, *111*, 104705. [CrossRef]
15. Templeton, A.; Drury, J.; Alnabulsi, A. Agent-Based Models of Social Behaviour and Evacuation. *Saf. Sci.* **2024**, *170*, 106362.
16. Zhang, J.; Li, H.; Zhou, Z. An Improved Social Force Model-Based Crowd Evacuation Simulation in Subway Fires. *Geoinformatica* **2023**, *27*, 575–601. [CrossRef]
17. Park, H.; Liu, D.; Namilae, S. *Multi-Agent Reinforcement Learning-Based Pedestrian Dynamics Models for Emergency Evacuation*; Final Report CATM-2022-R5-NCAT; North Carolina A&T State University, Transportation Institute, Center for Advanced Transportation Mobility: Greensboro, NC, USA, 2022. Available online: https://rosap.ntl.bts.gov/view/dot/64672 (accessed on 16 September 2025).
18. Disabled High-Rise Residents 'Still at Risk' Seven Years After Grenfell. *The Guardian*, 13 June 2024.
19. Hostetter, H.; Abad, S. Evacuation Preparedness and Intellectual Disability: Insights from a University Fire Drill. *J. Build. Eng.* **2024**, *93*, 109031. [CrossRef]
20. Lyu, Y.; Lin, B.; Zhang, J. Fire Evacuation for People with Functional Disabilities in High-Rise Buildings: A Systematic Review. *Buildings* **2025**, *15*, 634. [CrossRef]
21. Rismanian, M.; Zarghami, E. Evaluation of crowd evacuation in high-rise residential buildings with mixed-ability population: Combining an architectural solution with management strategies. *Int. J. Disaster Risk Reduct.* **2022**, *77*, 103068. [CrossRef]
22. Gupta, H.S.; Chen, A.; Park, B.J. Fairness-Driven Multi-Objective Optimization for Evacuation Planning in Natural Disasters. *Transp. Res. Part D*, 2024; *preprint*.
23. Wang, D.; Sun, L. Multi-Stage Equitable Bus-Based Hurricane Evacuation Planning. *Transp. Res. Rec.* **2023**, *2677*, 831–847.
24. Bairi, P.; Swain, S.; Bandyopadhyay, A.; Aurangzeb, K.; Alhussein, M.; Mallik, S. Intelligent VANET-Based Traffic Signal Control System for Emergency Vehicle Priority. *Egypt. Inform. J.* **2025**, *30*, 100700. [CrossRef]
25. Galveston Awarded $2M for Traffic Light Integration System That Prioritizes Emergency Vehicles. *Houston Chronicle*, 10 February 2025.
26. U.S. DOT ITS JPO. *Simulation Study: Intelligent Traffic Signal Control with Emergency Vehicle Priority*; ITS JPO: Washington, DC, USA, 2025.

27. Mnasri, S.; Habbash, M. Study of the influence of Arabic mother tongue on the English language using a hybrid artificial intelligence method. *Interact. Learn. Environ.* **2021**, *31*, 5568–5581. [CrossRef]
28. Laddha, S.; Mnasri, S.; Alghamdi, M.; Kumar, V.; Kaur, M.; Alrashidi, M.; Almuhaimeed, A.; Alshehri, A.; Alrowaily, M.A.; Alkhazi, I. COVID-19 Diagnosis and Classification Using Radiological Imaging and Deep Learning Techniques: A Comparative Study. *Diagnostics* **2022**, *12*, 1880. [CrossRef] [PubMed]
29. Hadjidimitriou, N.S.; Lippi, M.; Mamei, M.; Nastro, R.; Picone, M.; D'Andreagiovanni, F. Data-Driven Adaptation of Smart Grids With Hierarchical Digital Twins. *IEEE Pervasive Comput.* **2025**, *24*, 10–18. [CrossRef]
30. Bakhshian, E.; Martinez-Pastor, B. A multi-objective optimization model to minimize the evacuation time during a disaster considering reconstruction activity and uncertainty: A case study of Cork City. *Transp. Eng.* **2024**, *16*, 100234. [CrossRef]
31. European Commission. *Industry 5.0: Towards a Sustainable, Human-Centric and Resilient European Industry.* European Commission, Directorate-General for Research and Innovation. 2021. Available online: https://research-and-innovation.ec.europa.eu/knowledge-publications-tools-and-data/publications/all-publications/industry-50-towards-sustainable-human-centric-and-resilient-european-industry_en (accessed on 21 September 2025).
32. ESIR (European Commission Expert Group). *Industry 5.0: A Transformative Vision for Europe—Governing Systemic Transformations Towards a Sustainable Industry*; European Commission Expert Group: Brussels, Belgium, 2022.
33. Leng, J.; Zhou, M.; Xiao, Y.; Zhang, H.; Liu, Q.; Shen, W.; Su, Q.; Li, L. Digital twins-based remote semi-physical commissioning of flow-type smart manufacturing systems. *J. Clean. Prod.* **2021**, *306*, 127278. [CrossRef] [PubMed]
34. Leng, J.; Su, X.; Liu, Z.; Zhou, L.; Chen, C.; Guo, X.; Wang, Y.; Wang, R.; Zhang, C.; Liu, Q.; et al. Diffusion model-driven smart design and manufacturing: Prospects and challenges. *J. Manuf. Syst.* **2025**, *82*, 561–577. [CrossRef]
35. Martini, B.; Bellisario, D.; Coletti, P. Human-centered and sustainable artificial intelligence in industry 5.0: Challenges and perspectives. *Sustainability* **2024**, *16*, 5448. [CrossRef]
36. Mentzas, G.; Hribernik, K.; Stahre, J.; Romero, D.; Soldatos, J. Human-Centered Artificial Intelligence in Industry 5.0. *Front. Artif. Intell.* **2024**, *7*, 1429186. [CrossRef]
37. Guo, J.; Leng, J.; Zhao, J.L.; Zhou, X.; Yuan, Y.; Lu, Y.; Mourtzis, D.; Qi, Q.; Huang, S.; Song, X.; et al. Industrial metaverse towards Industry 5.0: Connotation, architecture, enablers, and challenges. *J. Manuf. Syst.* **2024**, *76*, 25–42. [CrossRef]
38. ETSI ISG CIM. *Context Information Management (CIM); NGSI-LD API Specification*; ETSI GS CIM 009 V1.5.1. ETSI: Sophia Antipolis, France, 2021. Available online: https://www.etsi.org/deliver/etsi_gs/CIM/001_099/009/01.05.01_60/gs_cim009v010501p.pdf (accessed on 21 September 2025).
39. *OPC Unified Architecture—Interoperability for Industrie 4.0 and the Internet of Things*; Updated Brochure. OPC Foundation: Scottsdale, AZ, USA, 2023. Available online: https://opcfoundation.org/wp-content/uploads/2023/05/OPC-UA-Interoperability-For-Industrie4-and-IoT-EN.pdf (accessed on 21 September 2025).