Advances in Differential Equations and Control Processes 2026, 33(3), 4677.
doi: https://doi.org/10.59400/adecp4677

Article

Optimal control of coupled visitor-flow and energy dynamics in smart scenic areas: A differential-equation framework with IoT-based state estimation

Yue Wu[1], Baijun Wu[2,*], Yingxuan Li[3]

[1] School of Hospitality, Tourism & Events, Taylor's University, Subang Jaya 47500, Malaysia; 0384311@sd.taylors.edu.my
[2] School of Economics and Management, Jiangmen Polytechnic, Jiangmen 529000, China
[3] School of Tourism and Hotel Management, Yunnan University of Finance and Economics, Kunming 650221, China; yingxuanli2026@163.com
*** Corresponding authors:** Baijun Wu, 13824653546@163.com

CITATION

Wu Y, Wu B, Li Y. Optimal control of coupled visitor-flow and energy dynamics in smart scenic areas: A differential-equation framework with IoT-based state estimation. Advances in Differential Equations and Control Processes. 2026; 33(3): 4677. https://doi.org/10.59400/adecp4677

ARTICLE INFO

Received: 14 June 2026
Revised: 5 August 2026
Accepted: 20 August 2026
Available online: 2 September 2026

Abstract: This paper develops a dynamic optimal-control framework for the coupled visitor-flow and energy-consumption processes of large-scale scenic destinations. Zone-level visitor density is modelled by nonlinear ordinary differential equations and facility loads by first-order linear equations with control inputs and occupancy disturbances, giving a single constrained plant for which we establish non-negativity, forward invariance of a compact set, and existence and uniqueness of solutions. For the associated finite-horizon problem, we prove that an optimal control exists and, applying Pontryagin's minimum principle, derive the costate equations and show that the optimal routing law is bang-bang with an explicit switching function. In contrast, the optimal energy law is saturated affine in the costate. A receding-horizon controller re-solves the problem online and, equipped with terminal ingredients, is recursively feasible and nominally asymptotically stable. An extended Kalman filter driven by an Internet of Things (IoT) sensor network supplies the state estimate, and its expected error covariance is proved uniformly bounded under Bernoulli sensor dropout. On a three-zone, nine-facility benchmark, the closed loop raises comfort compliance from 50% to 89% of slots, cuts mean waiting time by 57%, reduces total energy by 2.9% and load variance by 10%, and degrades gracefully at dropout rates up to 20%; a linear predictive variant proves competitive, so the nonlinear model's empirical advantage is not established. Robust stability under disturbance and distributed-parameter extensions remain open.

Keywords: Pontryagin minimum principle; receding-horizon regulation; recursive feasibility; extended Kalman filtering; forward invariance; switching function; intermittent observations; neural symbolic derivation

1. Introduction

The sustainable operation of large-scale public attractions—heritage sites, national parks, and cultural destinations—poses a fundamental control challenge. The dynamic coupling between visitor flows and facility energy consumption produces a nonlinear, multi-input multi-output (MIMO) system with pronounced spatiotemporal heterogeneity and hard operational constraints. Visitor density in each functional zone evolves continuously through arrivals, departures, and inter-zone movements, while energy loads in HVAC (heating, ventilation, and air conditioning), lighting, and transport infrastructure respond both to occupancy and to active control inputs, subject to the thermal and electrical inertia of the underlying plant. The two subsystems interact bidirectionally: high visitor density raises cooling and lighting demand, while service-capacity limits in turn shape visitor routing and waiting times. The resulting dynamics are nonlinear, state-dependent, and continually perturbed by stochastic disturbances arising from weather, calendar effects, and unanticipated demand surges.

Traditional management practices—fixed capacity quotas, rule-based routing, and manually dispatched facility schedules—treat these dynamics in an open-loop, reactive manner. It ignores the temporal coupling between visitor presence and energy consumption and cannot anticipate or preempt constraint violations. From a systems-and-control standpoint, this is equivalent to operating a poorly regulated plant without feedback: disturbances propagate through the system unchecked until a human operator intervenes. A rigorous feedback-based framework is therefore called for: one that continuously observes the system state from sensor measurements, predicts future trajectories from a dynamic model, and computes control actions that simultaneously respect density constraints and energy-efficiency targets.

Optimal-control theory provides the appropriate mathematical foundation for this class of problems. Since the classical results of Pontryagin's minimum principle and Bellman's dynamic programming, the discipline has offered a systematic route from a differential-equation model of a plant to a control law that is optimal with respect to a specified cost functional [1, 2]. Model predictive control (MPC) extends these ideas to a receding-horizon setting in which a finite-horizon optimal-control problem is re-solved at every sampling instant using the most recent state estimate, thereby providing intrinsic feedback and explicit constraint handling [3]. Control-theoretic formulations of exactly this kind now underpin a broad range of process and infrastructure applications, from load-frequency regulation of interconnected power systems via linear-quadratic and Riccati-based designs [4], to feedback stabilization of distributed-parameter plants [5], to time-optimal bang-bang control of strongly coupled nonlinear heating processes [6]. The present work brings the same control-theoretic discipline to the coupled crowd-energy processes of smart scenic areas.

A second line of work has begun to change how derivations of this kind are produced. Where the Hamiltonian, the costate equations, and the switching conditions of a nonlinear plant have traditionally been obtained by hand, artificial-intelligence-enhanced mathematical derivation now automates substantial parts of that algebra, combining symbolic manipulation with learned search to propose and verify candidate expressions [7]. In parallel, neural-symbolic solvers address the solution step rather than the derivation step: fractional sub-equation neural networks embed the solutions of a fractional Riccati equation into the neurons of a network and return exact solutions of space–time fractional partial differential equations instead of pointwise numerical approximations [8]. Both developments bear on the present work, and in different places. The first concerns how the necessary conditions of Section 4 are obtained—by manual algebra here, and in principle by automated derivation in future work. The second concerns how the resulting two-point boundary-value problem might be solved, and how the framework could be extended to memory-dependent or distributed-parameter plant models. We therefore position our contribution explicitly against both, and return to the comparison in Sections 2.5, 6.6, and 7.

Despite this maturity in control theory, the literature on smart-destination management has developed largely outside it. Crowd-forecasting studies emphasize predictive accuracy but do not close a control loop; energy-management studies optimize aggregate consumption without modelling the state dynamics that generate it; and optimization-based scheduling typically assumes static, perfectly known demand. What is missing is a closed-loop optimal-control framework that explicitly models the coupled ordinary differential equation (ODE) dynamics of the crowd-energy system and regulates them under real-time feedback. This paper fills that gap.

The contributions are fourfold. First, we propose an ODE-based dynamic model for coupled visitor-flow and energy systems, in which zone densities and facility loads evolve as a single nonlinear state vector with control and disturbance inputs, and we establish its well-posedness: non-negativity and forward invariance of the physical orthant, existence of a compact positively invariant set, and existence and uniqueness of solutions under stated regularity assumptions. Second, we formulate a constrained optimal-control problem for this model, prove that an optimal control exists, and characterize it via Pontryagin's minimum principle: we compute the costate equations explicitly and show that the optimal routing law is bang-bang with an explicit switching function, while the optimal energy law is saturated affine in the costate. Third, we implement the characterization as a receding-horizon MPC scheme and give terminal ingredients under which the closed loop is recursively feasible and nominally asymptotically stable, retaining an evolutionary procedure only for offline calibration of the objective weights. Fourth, we design a closed-loop IoT-enabled state-estimation architecture in which an extended Kalman filter fuses heterogeneous sensor streams into the state estimate that drives the controller. We prove that its expected error covariance remains uniformly bounded under intermittent observations, and we validate the complete estimator-controller loop against three years of operational data, including its behaviour under demand surges and sensor failures. We are explicit about what is not established: the stability guarantee is nominal, robust margins under disturbance and estimation error are not quantified, and the validation is single-site and simulation-based.

2. Literature review

2.1. Dynamic-system modelling of crowd-energy processes

The modelling of population and resource processes by differential equations has a long and productive history, and dynamic-system formulations remain the natural language for coupled, evolving quantities. Continuous-time compartmental models describe how a population redistributes among states through transition rates, a structure that transfers directly to the movement of visitors among functional zones. Population models of this kind, written as ordinary or partial differential equation (PDE) systems, have been analysed at length: threshold, boundedness, and stability results are standard for the classical compartmental families [9], and a well-developed existence theory now covers strongly coupled cross-diffusion systems, in which the movement of one population responds to the density of another [10]. On the energy side, control-oriented lumped-parameter ODE models have proven effective for representing the transport and storage dynamics of engineered energy systems: transient models of vanadium redox-flow batteries, for example, couple species transport to the cell electrochemistry and so link physical fidelity to actionable control insight while remaining tractable for design [11]. These two strands—population dynamics and energy dynamics—are individually well developed, but they are seldom coupled into a single state model whose joint evolution is the object of control, which is the modelling stance adopted here.

2.2. Optimal-control theory for infrastructure systems

Optimal control supplies the machinery for steering a differential-equation plant to minimize a cost functional subject to constraints. Two complementary pillars dominate practice. The first is the linear-quadratic family—linear-quadratic regulator (LQR) and linear-quadratic-Gaussian (LQG) control—built on the solution of Riccati equations; a recent state-space study of single-, two-, and three-area load-frequency control demonstrates how Riccati-based LQR/LQG designs minimize frequency and tie-line deviations in interconnected power systems [4]. The second is Pontryagin's minimum principle, which characterizes optimal trajectories through a Hamiltonian and its associated costate equations; the principle underlies the derivation of switching (bang-bang) controls, as shown for the time-optimal control of strongly coupled nonlinear microwave-heating systems, where optimal inputs are proven to saturate their bounds almost everywhere [6]. Feedback control of distributed-parameter systems has likewise been placed on a rigorous footing, with stabilizability, observer design, and Riccati theory all developed for infinite-dimensional plants in a semigroup setting [5]. Model predictive control unifies these ideas in a receding-horizon form that re-optimizes online and handles constraints explicitly [3], making it the method of choice for systems, such as ours, that face persistent disturbances and actuator limits, and it has been applied extensively to building energy and HVAC regulation [12]. The regularity theory underpinning such control formulations continues to develop for broader classes of evolution systems, including approximate controllability of fractional evolution inclusions and existence and uniqueness for fractional and stochastic evolution equations [13–15]; while the plant considered here is a classical first-order system, these results delimit the setting in which the necessary conditions used below can be extended to memory-dependent or stochastic dynamics.

2.3. State estimation and IoT-based feedback

Feedback control presupposes knowledge of the state, which in practice must be reconstructed from noisy, partial measurements. The discrete-time Kalman filter remains the canonical recursive estimator for linear-Gaussian systems [16] and, through its extended and fused variants, for heterogeneous sensing arrangements; its prediction-update structure provides both a state estimate and an error covariance that quantifies confidence. Modern infrastructure increasingly acquires the required measurements through Internet-of-Things (IoT) sensor networks, which draw identification, wireless sensing, and actuation technologies together into a single measurement fabric [17]. The streams such a network produces can be embedded directly into continuous-time dynamic models, where they supply the observations of a state estimator. In the framework developed below, an IoT layer plays this role precisely: it is the measurement channel of a state observer, not an end.

2.4. Multi-objective control formulations

Infrastructure control problems are intrinsically multi-objective: energy efficiency, comfort, and service quality compete, and no single scalar objective captures the operator's intent. Two established routes reconcile such objectives with dynamic optimization. The first embeds the trade-off as a weighted cost in an otherwise standard optimal-control or MPC problem, tuning the weights to select a point on the Pareto front. The second uses population-based metaheuristics to map the Pareto set explicitly; evolutionary algorithms such as non-dominated sorting genetic algorithm-II (NSGA-II) [18] are widely used for this purpose and, in switched operation, the switching law itself can be optimized by parameterizing the switching instants and differentiating the optimal cost with respect to them [19]. In the present work, we adopt the first route for online control and use NSGA-II only offline, as a weight-calibration instrument that supplies the single knee-point weighting subsequently used by the MPC controller.

2.5. Intelligent symbolic and data-driven solution methods

The three strands above are classical in method: the models are written by hand, the necessary conditions are derived by hand, and the resulting equations are solved numerically. A fourth strand now questions each of these steps. Artificial-intelligence-enhanced mathematical derivation treats algebraic manipulation itself as a search problem, using learned heuristics to propose transformations and symbolic engines to verify them, so that Hamiltonians, adjoint systems, and switching conditions can be generated and checked with far less manual labour than the classical route demands [7]. Neural-symbolic solvers attack the solution step instead: fractional sub-equation neural networks assign the solutions of a fractional Riccati equation to the neurons of the first hidden layer, so that training returns exact closed-form solutions of space–time fractional partial differential equations rather than a discretized approximation [8]. Related fractional and physics-informed architectures extend the same idea to memory-dependent models in which the order of differentiation is itself identified from data, and connect this strand to the fractional evolution theory already cited in Section 2.2 [13–15].

The relative merits are worth stating plainly, because they determine what each approach can be trusted to deliver. Manual Pontryagin derivation, the route taken in this paper, yields conditions that are auditable line by line and structural results—the sign of a switching function, the saturation of an actuator law—that hold for every admissible parameter value; its costs are labour, the risk of algebraic error in high-dimensional plants, and the need to redo the work whenever the model changes. Intelligent symbolic frameworks reverse that balance: they are fast and easily re-run after a model change, but their output requires independent verification and they do not by themselves supply the interpretation that makes a switching function operationally meaningful. Traditional numerical simulation, also used here, is convergent and carries error control on a fixed discretization, but scales poorly in state dimension and yields no closed form. Data-driven networks are mesh-free and inexpensive to evaluate once trained, and in the fractional setting, they return exact solution families that finite-difference schemes cannot represent; against this, they offer no a priori accuracy guarantee, generalize uncertainly outside the training regime, and are difficult to constrain to satisfy hard actuator limits. For a plant that must respect comfort thresholds and saturation bounds in real time, we therefore retain the classical route for the results claimed here and treat the intelligent alternatives as complements—for derivation assistance and for the fractional and distributed-parameter extensions discussed in Section 7—rather than as substitutes.

2.6. Research gap and contributions

Across these strands, a consistent gap emerges. Dynamic-system modelling captures population and energy processes separately but rarely couples them into a single controlled state. Optimal-control theory offers mature tools—Pontryagin's principle, Riccati-based LQR/LQG, and MPC—that have not been applied to the joint crowd-energy process of a scenic destination. State-estimation and IoT research supply the measurement channel but stop short of closing a control loop around a differential-equation plant. And multi-objective methods are typically deployed as static, offline optimizers rather than as online regulators. The framework proposed here addresses the gap by (i) modelling the coupled visitor-flow and energy dynamics as a single nonlinear ODE system, (ii) formulating and solving a constrained optimal-control problem for that system via Pontryagin's minimum principle and receding-horizon MPC, and (iii) closing the loop with an IoT-driven Kalman state estimator, validated on three years of operational data.

To make the gap concrete rather than merely asserted, **Table 1** compares these strands against the capabilities that a closed-loop crowd–energy controller requires. The columns are the ingredients that must be present simultaneously for the operational problem of Section 1 to be solved: a single state model in which crowd and energy evolve jointly, a feedback loop rather than an open-loop plan, explicit handling of actuator and comfort constraints, reconstruction of the state from real measurements, and a formal optimality characterization. Each strand supplies some of these, and none supplies all.

Table 1. Capability comparison of existing approaches and the proposed framework.

Approach	**Coupled crowd–energy state**	**Closed-loop feedback**	**Explicit constraints**	**State estimation from sensors**	**Optimality characterization**
ODE/PDE population and crowd models [9,10,20]	Crowd only	No	No	No	Stability analysis
Control-oriented energy-system models [11]	Energy only	Partial	Partial	No	Design-oriented
LQR/LQG regulation of infrastructure [4]	Single domain	Yes	Implicit (via weights)	Partial (LQG)	Riccati/LQ optimal

Table 1. *Cont.*

Approach	Coupled crowd–energy state	Closed-loop feedback	Explicit constraints	State estimation from sensors	Optimality characterization
MPC for building and HVAC energy [3,12,21]	Energy only	Yes	Yes	Rarely	Receding-horizon optimal
IoT sensing and monitoring architectures [17]	Neither (data layer)	No	No	Yes	None
Multi-objective and switched-system scheduling [18,19]	Either, statically	No	Yes (offline)	No	Pareto set
AI-enhanced symbolic derivation and neural solvers [7,8]	Not modelled jointly	No	No	No	Derivation and solution aid
This work	**Yes, single ODE system**	**Yes (MPC)**	**Yes (hard input, soft state)**	**Yes (EKF, multi-modal)**	**PMP: bang-bang routing, saturated affine energy**

Note: The final row is the contribution claim of this paper, and the propositions of Sections 3.4, 4.2, and 5.3 are what substantiate the entries in it.

3. Coupled dynamic model of visitor flow and energy processes

We consider a scenic destination partitioned into a finite set of functional zones indexed by $z = 1, \ldots, N$ (a primary attraction zone, a service-and-commercial zone, and a peripheral rest area in the case study of Section 6), interconnected by defined movement pathways. The controlled plant is described by two coupled subsystems: a visitor-density subsystem and a facility-energy subsystem.

3.1. Visitor-flow dynamics

Let $x_z(t)$ denote the visitor density (persons per unit area) in zone z at time t and let $x(t) = [x_1(t), \ldots, x_N(t)]^T$ be the density state vector. Visitors enter the destination through an exogenous arrival rate, redistribute among zones according to inter-zone transition coefficients, and leave through a density-dependent departure process. The density in each zone evolves according to the nonlinear ordinary differential equation

$$\dot{x_z}(t) = a_z(t) + \sum\nolimits_{j \neq z} \beta_{jz} \, u_{jz}(t) x_j(t) - \sum\nolimits_{k \neq z} \beta_{zk} \, u_{zk}(t) x_z(t) - \delta_z \, x_z(t), \quad (1)$$

where $a_z(t)$ is the exogenous arrival rate into zone z; $\beta_{\{jz\}} \geq 0$ is the baseline transition rate from zone j to zone z along the connecting pathway; $u_{\{jz\}}(t) \in [0, 1]$ is the routing control input that modulates the fraction of the movable flow directed from j to z (the actuation exercised through dynamic wayfinding); and $\delta_z > 0$ is the departure (egress) rate from zone z. The bilinear terms $\beta_{\{jz\}} \, u_{\{jz\}}(t) x_j(t)$ capture the control-modulated, density-proportional nature of inter-zone movement and render the visitor-flow subsystem nonlinear in the joint state-control variables. Writing $f(x, u, t)$ for the right-hand side of Equation (1), the subsystem is compactly $\dot{x}(t) = f\left(x(t), \, u(t), \, t\right)$.

The routing inputs are physically constrained. For each origin zone *j*, the outgoing routing fractions cannot exceed unity, giving the state-independent constraint $\sum_{\{k \neq j\}} u_{\{jk\}}(t) \leq 1$ with $u_{\{jk\}}(t) \geq 0$, which is the actuator-saturation limit of the wayfinding subsystem.

3.2. Facility-energy dynamics

Let $E_f(t)$ denote the active power drawn by facility f (indexed over HVAC, lighting, and internal transport within each zone) and let $E(t)$ be the corresponding load vector. Each facility exhibits first-order lag: its load relaxes toward a commanded set-point with a time constant τ_f that encodes thermal or electrical inertia and is additionally driven by the local visitor load. We model this with the first-order linear ODE

$$E_f(t) \; = \; -\frac{1}{\tau_f}\, E_f(t) \; + \; \frac{1}{\tau_f}\, v_f(t) \; + \; \gamma_f\, x_{z(f)}(t), \tag{2}$$

where $v_f(t)$ is the energy control input (the commanded power set-point delivered to the facility's building-management or lighting-management controller); $\gamma_f \;\geq\; 0$ is the visitor-load coupling coefficient that links occupancy in the host zone $z(f)$ to facility demand; and the term $-\left(\frac{1}{\tau_f}\right) E_f(t)$ represents the natural decay toward the commanded level. Equation (2) makes explicit the physical inertia that prevents instantaneous load reduction: the achievable rate of change of E_f is bounded by τ_f, and the set-point is constrained to the admissible actuator range $v_f(t) \in \; [v_f^{min},\; v_f^{max}]$. The coupling term $\gamma_f\; x_{\{z(f)\}}(t)$ is the channel through which the visitor-flow subsystem (Equation (1)) disturbs the energy subsystem, closing the bidirectional interaction: routing decisions that redistribute density alter the disturbance seen by every facility.

3.3. Composite state-space form

Stacking the two subsystems, define the composite state $s(t) \;=\; \left[x(t)^T,\; E(t)^T\right]^T \;\in\; R^{\{N+M\}}$ and the composite control $w(t) \;=\; \left[u(t)^T,\; v(t)^T\right]^T$, where M is the number of facilities. The plant is then a nonlinear control system

$$\dot{s}(t) \;=\; F(s(t),\; w(t),\; t),\; s(t_0) \;=\; s_0, \tag{3}$$

where F collects the right-hand sides of Equations (1) and (2). The visitor-flow block is nonlinear (bilinear in state and control) and the energy block is linear with a state-dependent disturbance; the composite system is therefore a nonlinear MIMO plant with hard input constraints $w(t) \;\in\; W$ and the state constraint set defined by the comfort thresholds introduced next. This composite ODE model is the object of the optimal-control formulation of Section 4.

3.4. Standing assumptions and well-posedness of the composite model

Before the composite system (Equation (3)) can serve as the plant of an optimal-control problem, its solutions must be shown to exist, to be unique, and to remain in the physically meaningful region of the state space. This subsection collects the standing assumptions and establishes these properties; they are used without further comment in Sections 4 and 5.

Assumption 1. (regularity and admissibility).

(A1) The exogenous arrival rates $a_z : [t_0, t_f] \to \mathbb{R} \geq 0$ *are Lebesgue measurable and uniformly bounded,* $0 \leq a_z(t) \leq a_{\max} < \infty$ *for all z and almost every t.*

(A2) The structural parameters are constants satisfying $\beta_{jz} \geq 0$, $\delta_z > 0$, $\tau_f > 0$

and $\gamma_f \geq 0$*; we write* $\beta_{\max} = \max_{j,z} \beta_{jz}$*,* $\delta_{\min} = \min_z \delta_z$*,* $\tau_{\min} = \min_f \tau_f$ *and* $\gamma_{\max} = \max_f \gamma_f$*.*

(A3) The admissible control set is the product $\mathcal{W} = \mathcal{U} \times \mathcal{V}$*, where* $\mathcal{U} = \{u : u_{jk} \geq 0,\ \sum_{k \neq j} u_{jk} \leq 1$ *for every* $j\}$ *is a product of simplices and* $\mathcal{V} = \prod_f \left[v_f^{\min}, v_f^{\max}\right]$ *is a box with* $0 \leq v_f^{\min} < v_f^{\max} < \infty$*. Consequently* $\mathcal{W}$ *is non-empty, convex, and compact.*

(A4) An admissible control is a Lebesgue-measurable function $w : [t_0, t_f] \to \mathcal{W}$*; the set of such functions is denoted* $\mathcal{W}_{ad}$*.*

(A5) The initial condition satisfies $x(t_0) \geq 0$ *componentwise and* $E(t_0) \geq 0$*.*

(A1) and (A2) are properties of the identified model, (A3) encodes the physical actuator limits already introduced in Sections 3.1 and 3.2, and (A5) is automatic for a physical initial state. No further structure is required.

Proposition 1. (non-negativity and forward invariance of the physical orthant). *Under Assumption 1, for every admissible control* $w \in \mathcal{W}_{ad}$ *every solution of (3) satisfies* $x_z(t) \geq 0$ *for all* z *and* $E_f(t) \geq 0$ *for all* f*, for as long as the solution exists. The non-negative orthant* $\mathbb{R}_{\geq 0}^{N+M}$ *is therefore forward invariant for the controlled dynamics.*

Proof. Consider the visitor-flow block first. The right-hand side of Equation (1) is quasi-positive (essentially non-negative) on $\mathbb{R}_{\geq 0}^{N}$: for $j \neq z$,

$$\frac{\partial \dot{x}_z}{\partial x_j} = \beta_{jz} u_{jz}(t) \ \geq 0 \tag{4}$$

by (A2) and (A3), so every off-diagonal partial derivative is non-negative. Equivalently, evaluate the vector field on the face $\{x_z = 0,\ x_j \geq 0\ (j \neq z)\}$. The two terms proportional to x_z—the routed outflow $\sum_{k \neq z} \beta_{zk} u_{zk} x_z$ and the departure term $\delta_z x_z$—vanish identically there, leaving

$$\dot{x}_z|_{x_z=0} = a_z(t) + \sum\nolimits_{j \neq z} \beta_{jz} u_{jz}(t)\, x_j(t) \ \geq \ 0. \tag{5}$$

The field therefore points into $\mathbb{R}_{\geq 0}^{N}$ (or is tangent to it) at every boundary face, and by the standard invariance criterion for quasi-positive systems [22], no trajectory starting in $\mathbb{R}_{\geq 0}^{N}$ can leave it. For the energy block, evaluate Equation (2) at $E_f = 0$: since $v_f(t) \geq v_f^{\min} \geq 0$ and $\gamma_f x_{z(f)} \geq 0$ by the argument just given, $\dot{E}_f|_{E_f=0} = \tau_f^{-1} v_f(t) + \gamma_f x_{z(f)}(t) \geq 0$, so E_f cannot become negative either. □

Proposition 1 is what licenses the interpretation of x_z as a density and E_f as an active power; it also guarantees that the positive-part operator in the running cost (Equation (10)) is evaluated on physically admissible arguments only.

Proposition 2. (Uniform boundedness and a compact positively invariant set). *Under Assumption 1, define* $X(t) = \sum_z x_z(t)$ *and*

$$X_{\max} = max\left\{X(t_0),\ \frac{N a_{\max}}{\delta_{\min}}\right\}, E_f^{\max} = \max\left\{E_f(t_0),\ v_f^{\max} + \tau_f \gamma_f X_{\max}\right\}. \tag{6}$$

Then the set

$$\varOmega = \left\{(x,E)\in\mathbb{R}^{N+M}_{\geq 0}:\ \sum\nolimits_z x_z \leq X_{\max},\ E_f \leq E_f^{\max}\ \forall f\right\} \tag{7}$$

is compact and positively invariant for (3) under every admissible control.

Proof. Summing Equation (1) over z, every routed term appears exactly twice with opposite signs—once as an inflow $+\beta_{jz}u_{jz}x_j$ in the equation for zone z and once as an outflow $-\beta_{jz}u_{jz}x_j$ in the equation for zone j—so the routing contributions cancel identically. This telescoping property expresses conservation of visitors under redistribution and is independent of the control. There remains

$$\dot{X}(t) = \sum\nolimits_z a_z(t) - \sum\nolimits_z \delta_z x_z(t)\ \leq\ N a_{\max} - \delta_{\min} X(t), \tag{8}$$

using $a_z \leq a_{\max}$, $x_z \geq 0$ (Proposition 1) and $\delta_z \geq \delta_{\min}$. The comparison lemma [22] applied to the scalar linear differential inequality (Equation (8)) yields $X(t) \leq X(t_0)\,e^{-\delta_{\min}(t-t_0)} + \frac{N a_{\max}}{\delta_{\min}}\left(1 - e^{-\delta_{\min}(t-t_0)}\right) \leq X_{\max}$ for all $t \geq t_0$. In particular, the total population is bounded uniformly in the control and in the horizon length, and each $x_z \leq X_{\max}$. Substituting this bound into Equation (2) and using $v_f \leq v_f^{\max}$ gives $\dot{E}_f \leq -\tau_f^{-1}E_f + \tau_f^{-1}v_f^{\max} + \gamma_f X_{\max}$, and the same comparison argument yields $E_f(t) \leq E_f^{\max}$. Together with Proposition 1, this shows that $\varOmega$ is positively invariant; it is closed and bounded in $\mathbb{R}^{N+M}$, hence compact. $\square$

Proposition 3. (existence, uniqueness, and global continuation). *Let Assumption 1 hold and let $s_0 \in \varOmega$. Then, for every admissible control $w \in \mathcal{W}_{ad}$ the initial-value problem Equation (3) has a unique absolutely continuous solution $s(\cdot\,; s_0, w)$ defined on the whole interval $[t_0, t_f]$, and the solution depends continuously on s_0. Moreover F is Lipschitz in s on $\varOmega$, uniformly in $(w,t)\in \mathcal{W}\times[t_0,t_f]$, with constant*

$$L_F = \max\left\{N\beta_{\max} + \delta_{\max},\ \tau_{\min}^{-1}\right\} + \gamma_{\max}. \tag{9}$$

Proof. Write $F(s,w,t) = F_0(s,w) + \left(a(t)^{\top}, 0\right)^{\top}$. The autonomous part F_0 is a polynomial map: the visitor-flow block is bilinear in (x,u) and the energy block is affine in (E,v) with a linear coupling to x. Polynomials are C^{∞}, so F_0 is continuously differentiable in (s,w) and hence locally Lipschitz in s; the inhomogeneous part depends on t alone and is measurable and bounded by (A1). Therefore F satisfies the Carathéodory conditions—measurable in t for fixed (s,w), continuous in (s,w) for almost every t, and dominated on bounded sets by an integrable function—and Carathéodory's existence theorem, together with the Picard–Lindelöf uniqueness argument for locally Lipschitz fields, gives a unique absolutely continuous solution on a maximal interval of existence $[t_0, t_{\max})$.

The Lipschitz constant follows from the Jacobian. In the row-sum (ℓ_∞) norm, the visitor-flow block contributes $|\partial\dot{x}_z/\partial x_z| + \sum_{j\neq z}|\partial\dot{x}_z/\partial x_j| \leq \left(\sum_{k\neq z}\beta_{zk}u_{zk} + \delta_z\right) + \sum_{j\neq z}\beta_{jz}u_{jz} \leq N\beta_{\max} + \delta_{\max}$, using $u_{jk} \in [0,1]$; the energy block contributes τ_f^{-1} on the diagonal and γ_f in the coupling column, giving

Equation (9).

Finally, by Proposition 2, the solution remains in the compact set Ω for all t in its interval of existence. A solution confined to a compact subset of the state space cannot blow up in finite time, so the maximal interval is the whole of $[t_0, t_f]$ and $t_{\max} = t_f$. Continuous dependence on the initial condition is then the standard Grönwall estimate $\| s\left(t; s_0^1, w\right) - s\left(t; s_0^2, w\right) \| \leq \| s_0^1 - s_0^2 \| e^{L_F(t-t_0)}$. □

Propositions 1–3 establish that Equation (3) is a well-posed control system on the compact invariant set Ω. This is the regularity that the optimal-control analysis of Section 4 presupposes: the Lipschitz property in Equation (9) is what permits differentiation of the Hamiltonian along trajectories, and compactness of Ω combined with compactness of $\mathcal{W}$ is what will deliver existence of an optimal control in Proposition 4.

4. Optimal-control problem and solution method

4.1. Cost function and constraints

The operator seeks to keep every zone within its comfort density while minimizing energy consumption and load variability over a planning horizon $[t_0,\ t_f]$. *Let* D_z^{c} denote the empirically calibrated comfort-density threshold of zone z; this threshold defines the state constraint $x_{z(t)} \leq D_z^{c}$. The associated state-constraint violation is measured by the excess density $(x_z - D_z^{c})_+ = \max(0,\ x_z - D_z^{c})$. We define the running cost

$$L = w_1 \sum\nolimits_z (\max(0,\ x_z - {D_z}^c))^2 + w_2 \sum\nolimits_f {E_f}^2 + w_3 \sum\nolimits_f (E_f - \bar{E}_f)^2, \quad (10)$$

in which the first term penalizes the squared state-constraint violation (the comfort-density constraint, imposed as a soft constraint), the second penalizes total energy, and the third penalizes load variability around a reference $\bar{E}_f$, thereby smoothing the load profile. The scalars $w_1,\ w_2,\ w_3 > 0$ are objective weights that encode the operator's trade-off; their offline calibration is described in Section 4.4. The optimal-control problem is

$$\min_{w(\cdot)} J = \int_{t_0}^{t_f} L(s(t),\ w(t),\ t)\, dt \quad (11)$$

subject to the dynamics (Equation (3)), the input constraints $w(t) \in W$ (routing-fraction and actuator-saturation limits of Sections 3.1 and 3.2), and the physical lower bound on instantaneous facility load imposed by inertia. This is a constrained, finite-horizon optimal-control problem for a nonlinear plant.

Two properties of the running cost require comment, because both are prerequisites for the variational analysis of Section 4.2.

First, the positive-part operator in the first term of Equation (10) is not differentiable at the switching point $x_z = D_z^{c}$, but the term that actually enters the cost is its *square*, and this is continuously differentiable.

Lemma 1. (C^1 regularity of the violation penalty). *Let* $\varphi(\xi) = (\max\{0, \xi\})^2$. *Then*

$\varphi \in C^1(\mathbb{R})$ with $\varphi'(\xi) = 2\max\{0, \xi\}$, and φ' is globally Lipschitz with constant 2; that is, $\varphi \in C^{1,1}(\mathbb{R})$. Consequently, the running cost (Equation (10)) is continuously differentiable in the state s on Ω.

Proof. For $\xi > 0$, we have $\varphi(\xi) = \xi^2$ and $\varphi'(\xi) = 2\xi$; for $\xi < 0$, we have $\varphi \equiv 0$ and $\varphi'(\xi) = 0$. At $\xi = 0$, the one-sided difference quotients are $\lim_{h\downarrow 0} \varphi(h)/h = \lim_{h\downarrow 0} h = 0$ and $\lim_{h\uparrow 0} \varphi(h)/h = 0$, so $\varphi'(0)$ exists and equals 0. The derivative $\varphi'(\xi) = 2\max\{0, \xi\}$ therefore exists everywhere and is continuous, since $\max\{0, \cdot\}$ is continuous; it is Lipschitz with constant 2 because $\max\{0, \cdot\}$ is non-expansive. The remaining terms of Equation (10) are quadratic polynomials in s and hence smooth, so $L(\cdot, w, t) \in C^1(\Omega)$. □

Lemma 1 answers the differentiability question directly: the squared positive part smooths the kink, the state gradient $\partial L/\partial x_z = 2w_1 \max\{0, x_z - D_z^c\}$ is continuous, and no recourse to nonsmooth (Clarke) subdifferential calculus [23] is needed for the adjoint system derived in Section 4.2. We note for completeness that φ is not twice differentiable at the switching point; where a C^2 cost is wanted—for instance to invoke second-order sufficiency conditions or a Newton-type solver with an exact Hessian—φ may be replaced by the Huber-type smoothing $\varphi_\epsilon(\xi) = \varphi(\xi)$ for $|\xi| \geq \epsilon$ and $\varphi_\epsilon(\xi) = (\xi + \epsilon)^3/(6\epsilon)$ on $|\xi| < \epsilon$, which is C^2 and converges to φ uniformly as $\epsilon \downarrow 0$; all results below hold verbatim for φ_ϵ. The interior-point solver used in Section 6 employs $\epsilon = 10^{-3}$ persons/m^2.

Second, the cost as written in Equation (10) contains no penalty on the energy control v, so the Hamiltonian would be affine rather than strictly convex in v and the pointwise minimization would place v at an actuator bound at almost every instant. Since a set-point law that switches between its extremes at every sampling instant is neither desirable for plant wear nor consistent with the smooth commanded trajectories reported in Section 6, we augment the running cost with a control-effort term and work throughout with

$$L = w_1 \sum_z (\max\{0, x_z - D_z^c\})^2 + w_2 \sum_f E_f^2 + w_3 \sum_f (E_f - \bar{E}_f)^2 + w_4 \sum_f v_f^2, \quad (12)$$

with $w_4 > 0$. The added term is a standard control-effort (Tikhonov) regularizer: it renders H strictly convex in v, guarantees a unique pointwise minimizer, and produces the saturated affine control law of Proposition 6. Taking $w_4 \downarrow 0$ recovers the affine-in-v case and, with it, a bang-bang energy law of the type reported for coupled nonlinear plants in [6]. The weight w_4 is calibrated together with w_1, w_2, w_3 by the procedure of Section 4.4, and the reference load $\bar{E}_f$ in the third term is the same quantity denoted $\bar{E}_f$ throughout; the symbol $\hat{E}_f$ used in an earlier draft of Section 4.2 refers to the same reference. For the routing inputs, no analogous regularizer is introduced, since the bang-bang structure established in Theorem 1 is a genuine and operationally meaningful feature of the wayfinding actuator rather than an artefact; a small regularizer $\epsilon \parallel u \parallel^2$ with $\epsilon \ll w_1$ is nevertheless added in the numerical implementation to resolve the non-generic singular arcs discussed after Theorem 1.

4.2. Necessary conditions via Pontryagin's minimum principle

Introduce the costate (adjoint) vector $\lambda(t) = \left[\lambda^{x(t)^T},\ \lambda^{E(t)^T}\right]^T \in R^{\{N+M\}}$ and form the Hamiltonian

$$H(s,\ w,\ \lambda,\ t)\ =\ L(s,\ w,\ t)\ +\ \lambda^T\, F(s,\ w,\ t). \tag{13}$$

Pontryagin's minimum principle states that, along an optimal trajectory $(s^*,\ w^*)$, there exists a costate λ^* satisfying the state and costate (adjoint) equations

$$\dot{s}^*\ =\ \frac{\partial H}{\partial \lambda}\ =\ F(s^*,\ w^*,\ t), \tag{14}$$

$$\dot{\lambda}^*\ =\ -\frac{\partial H}{\partial s}, \tag{15}$$

together with the transversality condition $\lambda^*\ (t_f)\ =\ 0$ (free terminal state), and the minimum condition that the optimal control minimises the Hamiltonian pointwise over the admissible set,

$$w \times (t)\ =\ \arg\ \min_{w\ \in\ W}\ H(s \times (t),\ w,\ \lambda \times (t),\ t). \tag{16}$$

Before applying the minimum principle, we note that an optimal control exists and that the hypotheses of the principle are met by the plant in Section 3.

Proposition 4. (existence of an optimal control). *Let Assumption 1 hold, let $s_0 \in \Omega$, and let the running cost be Equation (12). Then the optimal-control problem (Equation (11)) admits a solution: there exists an admissible pair (s^*, w^*) with $w^* \in \mathcal{W}_{ad}$ minimizing J over all admissible pairs.*

Proof. We verify the hypotheses of the Filippov–Cesari existence theorem [24]. (i) The set of admissible pairs is non-empty: by Proposition 3, every measurable w with values in $\mathcal{W}$ generates a unique trajectory on $[t_0, t_f]$, and by Proposition 2, that trajectory remains in the compact set Ω, so the set of admissible trajectories is non-empty and uniformly bounded. (ii) $\mathcal{W}$ is compact and convex by (A3), and it is independent of (s,t). (iii) F is continuous in (s,w) and measurable in t, and L is continuous in (s,w), non-negative and measurable in t; both are bounded on $\Omega \times \mathcal{W} \times [t_0, t_f]$. (iv) The extended velocity set

$$\mathcal{N}(s,t) = \{(L(s,w,t) + \eta,\ F(s,w,t)) :\ w \in \mathcal{W}, \eta \geq 0\} \subset \mathbb{R}^{1+N+M} \tag{17}$$

is convex for every $(s,t) \in \Omega \times [t_0, t_f]$. To see this, fix s and t and observe that $w \mapsto F(s,w,t)$ is affine: the routing inputs enter Equation (1) only through the terms $\beta_{jz} u_{jz} x_j$, which are linear in u once x is fixed, and the energy inputs enter Equation (2) linearly through $\tau_f^{-1} v_f$. Moreover $w \mapsto L(s,w,t)$ is convex, being the sum of a constant in u and the strictly convex quadratic $w_4 \sum_f v_f^2$ in v. The image of a convex compact set under an affine map, augmented in the first coordinate by the epigraph direction of a convex function, is convex; hence $\mathcal{N}(s,t)$ is convex. Since (i)–(iv) hold and Ω is compact, the Filippov–Cesari theorem [24] yields an optimal pair. □

The convexity in (iv) is worth emphasizing, since it is exactly the structural feature that a bilinear model possesses and a genuinely nonlinear-in-control model would not: the plant is nonlinear in the joint variable (x, u) but affine in u for frozen state, which is what makes the existence theory elementary.

With existence secured, the hypotheses of Pontryagin's minimum principle are satisfied: F and L are continuously differentiable in s on the open neighbourhood of Ω (Proposition 3 and Lemma 1), continuous in w on the compact set $\mathcal{W}$, and measurable in t; the admissible control set is compact and independent of the state; and the comfort-density requirement is imposed as a soft penalty in Equation (12) rather than as a hard pointwise state constraint. This last modelling choice is deliberate and has an analytical consequence: because there is no active pure state constraint, the costate is absolutely continuous and no additional measure multiplier or jump condition of the kind required by state-constrained versions of the minimum principle is needed. The necessary conditions, therefore, take the classical form (Equations (13)–(16)).

Proposition 5. (explicit costate equations). *Let (s^*, w^*) be optimal and let H be given by Equation (13) with running cost Equation (12). Then the costate $\lambda = \left(\lambda^{x\top}, \lambda^{E\top}\right)^\top$ is absolutely continuous on $[t_0, t_f]$ and satisfies, for almost every t,*

$$\dot{\lambda}_z^x = -2w_1 \max\{0, x_z - D_z^c\} + \lambda_z^x \left(\sum\nolimits_{k\neq z} \beta_{zk} u_{zk} + \delta_z\right) - \sum\nolimits_{k\neq z} \lambda_k^x \beta_{zk} u_{zk} - \sum\nolimits_{f:\, z(f)=z} \gamma_f \lambda_f^E, \tag{18}$$

$$\dot{\lambda}_f^E = -2\left(w_2 + w_3\right) E_f + 2w_3 \bar{E}_f + \tau_f^{-1} \lambda_f^E, \tag{19}$$

with the transversality conditions $\lambda_z^x\left(t_f\right) = 0$ and $\lambda_f^E\left(t_f\right) = 0$.

Proof. The Hamiltonian is $H = L + \sum_z \lambda_z^x \dot{x}_z + \sum_f \lambda_f^E \dot{E}_f$ with $\dot{x}_z$ from Equation (1) and $\dot{E}_f$ from Equation (2). We compute $\partial H/\partial x_z$ term by term. From the running cost Equation (12), Lemma 1 gives $\partial L/\partial x_z = 2w_1 \max\{0, x_z - D_z^c\}$; the remaining terms of L do not involve x. From the visitor-flow sum, the variable x_z appears in exactly two places: in its own equation, where $\partial \dot{x}_z/\partial x_z = -\sum_{k\neq z} \beta_{zk} u_{zk} - \delta_z$, contributing $-\lambda_z^x \left(\sum_{k\neq z} \beta_{zk} u_{zk} + \delta_z\right)$; and in the equation of every other zone k as the inflow term $+\beta_{zk} u_{zk} x_z$, whose derivative $\partial \dot{x}_k/\partial x_z = \beta_{zk} u_{zk}$ contributes $+\sum_{k\neq z} \lambda_k^x \beta_{zk} u_{zk}$. From the energy sum, x_z appears in the coupling term of every facility hosted in zone z, with $\partial \dot{E}_f/\partial x_z = \gamma_f$ for $z(f) = z$ and zero otherwise, contributing $+\sum_{f:z(f)=z} \gamma_f \lambda_f^E$. Collecting these and applying $\dot{\lambda}_z^x = -\partial H/\partial x_z$ gives Equation (18).

For the energy block, E_f appears in L through $w_2 E_f^2 + w_3 \left(E_f - \bar{E}_f\right)^2$, with derivative $2w_2 E_f + 2w_3 \left(E_f - \bar{E}_f\right) = 2\left(w_2 + w_3\right) E_f - 2w_3 \bar{E}_f$, and in its own dynamics through $\partial \dot{E}_f/\partial E_f = -\tau_f^{-1}$, contributing $-\tau_f^{-1} \lambda_f^E$. No other state equation contains E_f, since the energy subsystem is driven by the visitor state but does not feed back into it within the model. Applying $\dot{\lambda}_f^E = -\partial H/\partial E_f$ gives Equation (19). Absolute continuity of λ follows because the right-hand sides of Equations (18) and (19) are bounded on $\Omega \times \mathcal{W}$ and measurable in t; the transversality condition is the free-terminal-state condition $\lambda\left(t_f\right) = 0$, the terminal cost in Equation (11) being absent. □

Equation (19) is linear in λ_f^E with homogeneous coefficient $+\tau_f^{-1}$; integrated backwards from $\lambda_f^E(t_f) = 0$ it is stable, and the time constant τ_f that limits how fast the plant can respond is precisely the constant that governs how far back in time a load penalty propagates. The system (Equations (18) and (19)), together with Equation (3) and the minimum condition Equation (16), constitutes a two-point boundary-value problem, with the state specified at t_0 and the costate at t_f.

Theorem 1. (bang-bang structure of the optimal routing law). *Define, for each ordered pair of distinct zones (j, k), the switching function*

$$\sigma_{jk}(t) = \beta_{jk}\, x_j^*(t)\, \left(\lambda_k^x(t) - \lambda_j^x(t)\right). \tag{20}$$

Then H is affine in u with $H(s, w, \lambda, t) = H_0(s, v, \lambda, t) + \sum_j \sum_{k\neq j} \sigma_{jk}(t)\, u_{jk}$, and the minimum condition (Equation (16)) is solved, for each origin zone j independently, by

$$\begin{aligned} &u_{jk}^*(t) = 1 \text{ if } k = \text{argmin}_{m\neq j}\sigma_{jm}(t) \text{ and } \min_{m\neq j}\sigma_{jm}(t) < 0;\\ &u_{jk}^*(t) = 0\ otherwise. \end{aligned} \tag{21}$$

whenever the minimizing index is unique. The optimal routing control therefore takes values at the vertices of the simplex $\Delta_j = \{u_{j\cdot} \geq 0,\ \sum_{k\neq j} u_{jk} \leq 1\}$—it is bang-bang—outside the singular set $\mathcal{S} = \{t : \sigma_{jk}(t) = 0$ or the minimizer is non-unique$\}$.

Proof. Collect the terms of H containing u_{jk}. The control u_{jk} appears in exactly two state equations: as an outflow $-\beta_{jk}u_{jk}x_j$ in $\dot{x}_j$, contributing $-\lambda_j^x\beta_{jk}u_{jk}x_j$ to H, and as an inflow $+\beta_{jk}u_{jk}x_j$ in $\dot{x}_k$, contributing $+\lambda_k^x\beta_{jk}u_{jk}x_j$. Their sum is $\beta_{jk}x_j\left(\lambda_k^x - \lambda_j^x\right)u_{jk} = \sigma_{jk}u_{jk}$. The running cost (Equation (12)) does not depend on u. Hence H is affine in u with gradient $\partial H/\partial u_{jk} = \sigma_{jk}$, and the constraint set decouples across origin zones into the independent simplices Δ_j.

Minimizing an affine function over a simplex is attained at a vertex. Explicitly, $\min_{u_{j\cdot}\in\Delta_j} \sum_{k\neq j} \sigma_{jk}u_{jk}$ equals $\min\{0,\ \min_{m\neq j}\sigma_{jm}\}$, the value 0 being attained at the vertex $u_{j\cdot} = 0$ (retain all visitors in zone j) and the value $\min_m \sigma_{jm}$ at the unit vertex in the direction of the minimizing index (route the entire movable flow to that zone). This is Equation (21). □

The switching function (Equation (20)) has a transparent interpretation: $\lambda_k^x - \lambda_j^x$ is the marginal cost of relocating one unit of density from zone j to zone k, and $\beta_{jk}x_j$ is the rate at which such relocation is physically achievable. Flow is commanded along a pathway exactly when the destination is cheaper at the margin than the origin, and the magnitude of σ_{jk} orders the competing destinations. The structure recovers, in the routing setting, the switching character established for time-optimal control of strongly coupled nonlinear plants in the study of Luo et al. [6].

Two degenerate cases must be acknowledged. A singular arc occurs if $\sigma_{jk}(t) \equiv 0$ on a subinterval of positive measure, which by Equation (20) requires either $x_j \equiv 0$—an empty origin zone, in which case the routing control is irrelevant and may be assigned arbitrarily—or $\lambda_k^x \equiv \lambda_j^x$ on that interval. In the latter case, differentiating and substituting Equation (18) shows that $\dot{\lambda}_k^x - \dot{\lambda}_j^x \equiv 0$ imposes the additional algebraic relation

$$2w_1\left[\max\{0, x_k - D_k^c\} - \max\{0, x_j - D_j^c\}\right] = \lambda_j^x \Theta_{jk}(u) + \sum\nolimits_{f:z(f)=j} \gamma_f \lambda_f^E - \sum\nolimits_{f:z(f)=k} \gamma_f \lambda_f^E$$

for a control-dependent term Θ_{jk}, a codimension-one condition on the state–costate pair that is not satisfied on an interval for generic parameter values. Singular arcs are therefore non-generic here. The second case is a tie between two minimizing destinations, which likewise occupies a null set of times generically. In the discretized implementation of Section 4.3, both are resolved by the strictly convex regularizer $\epsilon \parallel u \parallel^2$ with $\epsilon \ll w_1$ mentioned in Section 4.1, which selects the minimum-norm element of the optimal face and restores uniqueness at negligible cost in optimality.

Proposition 6. (saturated affine optimal energy law). *Under Equation (12) with $w_4 > 0$, the minimum condition (Equation (16)) determines the energy control uniquely as*

$$v_f^*(t) = \mathrm{sat}_{[v_f^{\min}, v_f^{\max}]}\left(\kappa_f\, \lambda_f^E(t)\right), \qquad \kappa_f = -\frac{1}{2w_4\tau_f}, \tag{22}$$

where $\mathrm{sat}_{[a,b]}(\xi) = \min\{b, \max\{a, \xi\}\}$. *The gain κ_f is plant-dependent through the facility time constant τ_f and controller-dependent through the effort weight w_4.*

Proof. The control v_f enters H in exactly two places: the effort term $w_4 v_f^2$ of the running cost and the input term $\tau_f^{-1} v_f$ of the dynamics (Equation (2)), the latter contributing $\lambda_f^E \tau_f^{-1} v_f$. All other terms are independent of v_f, and the constraint set $\mathcal{V}$ is a product of intervals, so the minimization decouples across facilities. Thus

$$\frac{\partial H}{\partial v_f} = 2w_4 v_f + \frac{\lambda_f^E}{\tau_f}, \qquad \frac{\partial^2 H}{\partial v_f^2} = 2w_4 > 0, \tag{23}$$

so H is strictly convex in v_f and possesses the unique unconstrained stationary point $\hat{v}_f = -\lambda_f^E/(2w_4\tau_f) = \kappa_f \lambda_f^E$. For a strictly convex function of one variable minimized over a closed interval, the constrained minimizer is the Euclidean projection of the unconstrained minimizer onto that interval, which is precisely the saturation operator. Hence $v_f^* = \mathrm{sat}_{[v_f^{\min}, v_f^{\max}]}\left(\kappa_f \lambda_f^E\right)$, giving Equation (22); uniqueness follows from strict convexity. □

The law (Equation (22)) is affine in the costate on the interior of the actuator range and saturates at the bounds, so the actuator limits activate exactly when the shadow price of load exceeds what the plant can deliver. As anticipated in Section 4.1, letting $w_4 \downarrow 0$ sends $|\kappa_f| \to \infty$ and degenerates Equation (22) into a bang-bang law, the energy analogue of Theorem 1. The necessary conditions (Equations (13)–(16)), together with the explicit forms (Equations (18), (19), (21), and (22)), thus, define a two-point boundary-value problem in (s, λ), whose numerical solution over a finite horizon is embedded in the predictive controller of Section 4.3.

4.3. Receding-horizon model predictive control

Solving the boundary-value problem (Equations (13)–(16)) once over the whole operating day would produce an open-loop schedule that ignores realized disturbances.

To obtain feedback, we adopt a receding-horizon (MPC) implementation. At each sampling instant t_k the current state estimate $\hat{s}\left(t_k\right)$ is supplied by the Kalman observer of Section 5; the finite-horizon problem (Equation (11)) is solved over the prediction window $[t_k,\ t_k+T_p]$ subject to Equation (3) and the input constraints, using the state estimate as the initial condition; the first control move $w\times\left(t_k\right)$ is applied to the plant; and at the next instant $t_{\{k+1\}}$ the window slides forward and the problem is re-solved with the updated estimate. This is the standard receding-horizon principle [3], and it endows the controller with intrinsic feedback: any deviation between the predicted and realized state—whether from a demand surge, a weather change, or a sensor fault—enters the next optimization as a corrected initial condition. The sampling interval is fixed at 15 min to match the operational scheduling slot, and the prediction horizon T_p is a tunable multiple of that slot analyzed in Section 6.7. A warm-start strategy initializes each optimization from the previous solution shifted by one slot, which reduces the per-step solve time and enables an event-triggered re-solve—an unscheduled MPC update fired when the observer reports a state excursion beyond a calibrated threshold—to complete well within a single slot.

Because a receding-horizon law is not automatically feasible at successive instants, nor automatically stabilizing, we state the terminal ingredients that render it so and the guarantee they buy.

Assumption 2. (terminal ingredients). *There exist a reference operating pair $\left(s^r,w^r\right)$ with w^r in the interior of $\mathcal{W}$ and s^r in the interior of Ω satisfying $F\left(s^r,w^r,t\right)=0$; a compact terminal set $\mathcal{S}_f\subseteq\Omega$ containing s^r in its interior; a local feedback $\kappa_{\text{loc}}:\mathcal{S}_f\to\mathcal{W}$; and a terminal cost $V_f\in C^1\left(\mathcal{S}_f\right)$, positive definite with respect to s^r, such that for all $s\in\mathcal{S}_f$:*

(T1) $\kappa_{\text{loc}}(s)\in\mathcal{W}$ (the local law is admissible);

(T2) the solution of $\dot{s}=F\left(s,\kappa_{\text{loc}}(s),t\right)$ from any $s\in\mathcal{S}_f$ remains in $\mathcal{S}_f$ (positive invariance);

(T3) $\nabla V_f(s)^\top F\left(s,\kappa_{\text{loc}}(s),t\right)\le -L\left(s,\kappa_{\text{loc}}(s),t\right)$ (the terminal cost is a local control Lyapunov function dominating the stage cost).

For the plant at hand, these ingredients are constructible rather than hypothetical. Linearizing Equation (3) at $\left(s^r,w^r\right)$ gives a pair $\left(A_r,B_r\right)$ that is stabilizable—the energy block is diagonal and Hurwitz with eigenvalues $-\tau_f^{-1}$ and is directly actuated, and the visitor block is stabilizable whenever every zone has a positive departure rate $\delta_z>0$, which holds by (A2). Taking κ_{loc} to be the LQR law for the linearization and $V_f(s)=\left(s-s^r\right)^\top P\left(s-s^r\right)$ with P the corresponding Riccati solution, (T3) holds on a sufficiently small sublevel set $\mathcal{S}_f=\{s:V_f(s)\le c\}$ by a standard argument bounding the linearization remainder, and (T2) follows from invariance of that sublevel set. The Riccati construction is the same one used for load-frequency regulation in the work of Patra and Patnaik [4].

Theorem 2. (recursive feasibility and nominal asymptotic stability). *Let Assumptions 1 and 2 hold. Let the finite-horizon problem solved at each instant t_k in Section 4.3 be augmented with the terminal cost $V_f\left(s\left(t_k+T_p\right)\right)$ and the terminal constraint*

$s(t_k + T_p) \in \mathcal{S}_f$*, and suppose the problem is feasible at* t_0. Then, in the nominal case—exact model, no disturbance, and exact state measurement—the problem is feasible at every subsequent instant t_k*, and the reference* s^r *is asymptotically stable for the resulting closed loop, with the region of attraction being the set of initial states for which the problem is feasible.*

Proof. Feasibility propagates by the shifted-candidate construction. Let $w^*_{[t_k,\, t_k+T_p]}$ be optimal at t_k and s^* the associated trajectory, so, $s^*(t_k + T_p) \in \mathcal{S}_f$. At $t_{k+1} = t_k + \Delta$ define the candidate control equal to w^* on $[t_{k+1}, t_k + T_p]$ and equal to κ_{loc} on $[t_k + T_p,\, t_{k+1} + T_p]$. This candidate is admissible by (T1); the trajectory it generates coincides with s^* up to $t_k + T_p$ and thereafter remains in $\mathcal{S}_f$ by (T2), so the terminal constraint is met and the candidate is feasible at t_{k+1}. Feasibility at t_0 therefore implies feasibility at all t_k by induction.

For stability, let $V(s)$ denote the optimal value of the finite-horizon problem at state s. Evaluating the cost of the shifted candidate and using optimality at t_{k+1},

$$V(s_{k+1}) - V(s_k) \leq -\int_{t_k}^{t_{k+1}} L\, dt + \left[V_f(s(t_{k+1} + T_p)) - V_f(s(t_k + T_p)) + \int_{t_k+T_p}^{t_{k+1}+T_p} L\, \mathrm{d}t \right], \tag{24}$$

and the bracketed term is non-positive by integrating (T3) along the terminal trajectory. Hence $V(s_{k+1}) - V(s_k) \leq -\int_{t_k}^{t_{k+1}} L\, \mathrm{d}t \leq 0$, so V is non-increasing along the closed loop. V is positive definite with respect to s^r because L is non-negative and vanishes only at the reference, and it is continuous on the feasible set; it is therefore a Lyapunov function, and asymptotic stability follows by the standard argument for terminal-constraint MPC [3,25,26]. □

Two qualifications delimit the scope of Theorem 2, and we state them plainly. First, the guarantee is nominal. In the presence of the arrival and weather disturbances of Section 6 and the estimation error of Section 5, the appropriate notion is input-to-state practical stability: the closed loop retains a Lyapunov decrease outside a residual set whose size scales with the disturbance and estimation-error bounds, and establishing an explicit robust margin—for instance by a tube-based or min–max reformulation—is beyond the scope of this paper and is listed among the open problems in Section 7. Second, in the numerical implementation of Section 6, the terminal constraint $s(t_k + T_p) \in \mathcal{S}_f$ is enforced as a soft constraint with a large penalty, which preserves feasibility of the numerical program under an infeasible measured initial condition at the price of replacing the guarantee of Theorem 2 by its practical counterpart. Reported closed-loop trajectories never activated the terminal penalty.

4.4. Offline weight calibration by evolutionary search

The objective weights $(w_1,\, w_2,\, w_3)$ fix the operating point of the controller on the comfort-energy trade-off and are calibrated once, offline, before deployment. Because the three objectives are genuinely conflicting, we map the trade-off explicitly by treating weight selection as a multi-objective problem and solving it with the elitist non-dominated sorting genetic algorithm (NSGA-II), evaluated on the historical dataset with the same MPC controller in the loop. Each candidate weight vector is scored by the closed-loop constraint-violation count, total energy, and load variance it induces over the calibration period; NSGA-II returns the Pareto set of non-dominated weightings. The knee point of this set, selected by the minimum-distance-to-ideal-point criterion, furnishes the single weighting used by the online controller. We emphasize the demoted role of the evolutionary search: it is a calibration instrument that shapes the cost functional, not the online control law, which is the MPC scheme of Section 4.3. NSGA-II hyperparameters (population 100, crossover 0.9, mutation 0.1, 200 generations) are reported for reproducibility in **Appendix A**.

5. IoT-based state estimation and feedback architecture

5.1. Extended Kalman observer for the nonlinear plant

The controller of Section 4 requires the composite state $s(t)$, which is not directly measured; instead, it is reconstructed from heterogeneous, noisy sensor streams by an extended Kalman filter that constitutes the observer of the feedback loop. Two complementary modalities measure zone density—boundary-based pedestrian counters and area-based vision nodes—and are fused with the smart-meter load readings into a single state estimate. Let $\hat{s}_{\{t|t\}}$ denote the filtered state estimate at slot t and $P_{\{t|t\}}$ the associated error covariance. The observer propagates the estimate through a prediction step,

$$\hat{s}_{t|t-1} \;=\; F\hat{s}_{t-1|t-1}, \tag{25}$$

$$P_{t|t-1} \;=\; FP_{t-1|t-1}\,F^T \;+\; Q, \tag{26}$$

in which F is the (locally linearized) state-transition matrix and Q is the process-noise covariance calibrated from historical arrival and load variance, and then corrects it through an update step as each measurement z_t arrives,

$$K_t \;=\; P_{t|t-1}\,H^T\left(HP_{t|t-1}\,H^T \;+\; R\right)^{-1}, \tag{27}$$

$$\hat{s}_{t|t} \;=\; \hat{s}_{t|t-1} \;+\; K_t(z_t \;-\; H\hat{s}_{t|t-1}), \tag{28}$$

$$P_{t|t} \;=\; (I \;-\; K_tH)P_{t|t-1}, \tag{29}$$

where H maps the state to the observation space and R is the observation-noise covariance estimated from sensor calibration. The gain K_t weights each measurement by its relative confidence, so that when a sensor stream degrades or drops out, the filter automatically leans on the remaining modalities and on the model prediction (Equations (25) and (26)). This graceful-degradation property is what gives the closed loop its robustness to sensor failure, quantified in Section 6.5. Relative to naive averaging of the two density modalities, the fused observer reduces state-estimate error and latency substantially, consistent with reported gains for heterogeneous IoT sensor fusion.

Because the plant (Equation (3)) is nonlinear—bilinear in the joint state–control variable—a linear Kalman filter is not directly applicable, and the observer used here is the extended Kalman filter obtained by linearizing the discretized dynamics about the current estimate. We make the construction explicit.

Discretizing Equation (3) with a zero-order hold on the control over the control slot $\Delta = 15$ min gives the nonlinear stochastic difference equation

$$s_{k+1} = \Phi\left(s_k, w_k\right) + \omega_k, \qquad z_k = H s_k + \nu_k, \tag{30}$$

where $\Phi(s, w) = s + \int_{t_k}^{t_k+\Delta} F\left(s(\vartheta), w, \vartheta\right)\, d\vartheta$ is the one-slot flow map of Equation (3), $\omega_k \sim \mathcal{N}(0, Q)$ is the process noise absorbing arrival stochasticity and unmodelled load, and $\nu_k \sim \mathcal{N}(0, R)$ is the sensor noise, both assumed zero-mean, white, and mutually independent. The prediction step (Equation (25)) is then executed on the nonlinear map,

$$\hat{s}_{k|k-1} = \Phi\left(\hat{s}_{k-1|k-1},\, w_{k-1}\right), \tag{31}$$

and the covariance recursions (Equations (26)–(29)) are executed with F understood throughout as the Jacobian of Φ evaluated at the current estimate,

$$F_k = \frac{\partial \Phi}{\partial s}\Big|_{\hat{s}_{k-1|k-1},\, w_{k-1}} \in \mathbb{R}^{(N+M)\times(N+M)}. \tag{32}$$

Under the explicit Euler discretization used in the implementation, F_k has the block-triangular structure

$$F_k = \begin{bmatrix} A_{xx} & 0 \\ A_{Ex} & A_{EE} \end{bmatrix}, \tag{33}$$

with blocks $\left(A_{xx}\right)_{zz} = 1 - \Delta\left(\sum_{k\neq z} \beta_{zk} u_{zk} + \delta_z\right)$, $\left(A_{xx}\right)_{zj} = \Delta \beta_{jz} u_{jz}$ for $j \neq z$, $\left(A_{Ex}\right)_{fz} = \Delta \gamma_f$ if $z(f) = z$ and zero otherwise, and $\left(A_{EE}\right)_{ff} = 1 - \Delta/\tau_f$.

The zero block records that the energy state does not feed back into the visitor dynamics within the model; the lower-left block A_{Ex} is the discrete-time image of the occupancy coupling γ_f and is the only channel through which a density correction propagates into the load estimate.

The remaining matrices are defined as follows. The state dimension is $N + M$ with $N = 3$ zones and M facilities. The observation matrix $H \in \mathbb{R}^{p\times(N+M)}$ has one row per sensor stream, $p = N_c + N_v + M$, where N_c is the number of pedestrian-counter groups and N_v the number of vision nodes: a counter or vision node sited in zone z contributes the row $A_z\, e_z^\top$ (with A_z the zone area, converting a person count to a density) padded with zeros over the energy block, and a smart meter on facility f contributes the row $e_{N+f}^\top$. H is thus constant, sparse, and known from the deployment map; a stream that fails to report in slot k is handled by deleting its row from H and the corresponding row and column from R for that slot, which is the mechanism underlying the graceful degradation analysed in Section 5.3. The process-noise covariance $Q \in \mathbb{R}^{(N+M)\times(N+M)}$ is positive semi-definite and diagonal, its density entries set to the empirical slot-to-slot variance of the arrival residual and its load entries to the residual variance of the identified facility model, both computed on the identification data of Section 6.1. The measurement-noise covariance $R \in \mathbb{R}^{p\times p}$ is positive definite and diagonal, with entries obtained from the manufacturer calibration and the on-site validation counts: counting-line variance for the LoRaWAN counters, detector variance for the vision nodes, and metering-class variance for the smart meters. The gain $K_k \in \mathbb{R}^{(N+M)\times p}$ in Equation (27) is well defined because $H P_{k|k-1} H^\top + R \succ 0$ whenever $R \succ 0$.

Two properties of the linearization deserve comment. First, Φ is quadratic in x—the only nonlinearity is the bilinear routing term—so the Taylor remainder neglected by the EKF is exactly second order with no higher-order tail, and it is bounded on the compact invariant set Ω of Proposition 2 by $\frac{1}{2}\Delta\beta_{\max} \parallel \widetilde{x} \parallel^2$ with $\widetilde{x}$ the estimation error. The linearization error is therefore quadratically small in the estimation error rather than merely locally small, which is a stronger statement than holds for a generic nonlinear plant. Second, the standard sufficient conditions for bounded EKF error [27, 28]—uniform observability of the pair (F_k, H), uniform boundedness of F_k, and a sufficiently small initial error—are satisfied here: F_k is uniformly bounded because Ω and $\mathcal{W}$ are compact, and observability holds by construction because every zone carries at least one density stream and every facility its own meter, so H has full column rank and the pair is observable in a single step. We do not claim optimality of the estimator, only boundedness of its error covariance; the EKF is a suboptimal but computationally tractable observer, and the empirical error levels reported in Section 6.2 are the relevant validation.

5.2. Feedback data-acquisition layer

The measurements that drive the observer are gathered by a three-tier IoT architecture whose sole role in the control system is to serve as the feedback data-acquisition and actuation channel. A perception tier of heterogeneous sensors—LoRaWAN bidirectional pedestrian counters at zone boundaries (10-min reporting), MQTT smart meters on facility distribution boards (1-min active-power readings), co-located environmental nodes, and vision nodes running an embedded person-detection model at high-density chokepoints (30-s counts)—supplies the raw signals. A fog-edge tier of per-zone gateways performs temporal alignment onto the 15-min control slot, outlier rejection, short-gap interpolation, and z-score normalization, and publishes normalized slot summaries together with threshold alerts that can trigger an event-driven MPC re-solve. A cloud tier hosts the Kalman observer, the MPC controller, and the actuation dispatch that returns commands—routing directives to wayfinding displays, set-points to HVAC and lighting controllers, and shuttle-frequency advisories—to the plant. The architecture is incrementally deployable and transferable: zones may be instrumented individually, and relocation to another destination requires only recalibration of zone capacities, comfort thresholds, and plant time constants, with no change to the control software. The complete closed loop—plant, sensing, estimation, and control—is shown in **Figure 1**, and **Table 2** summarizes the communication interfaces; all visitor sensing is aggregate and camera inference is on-device, so no personally identifiable data leaves the edge.

Table 2. Communication interfaces of the feedback data-acquisition layer.

Interface	Protocol	Interval	Primary payload
Pedestrian counter → gateway	LoRaWAN (SF9)	10 min	Zone entry/Exit counts
Smart meter → gateway	MQTT over Wi-Fi	1 min	Facility active power (kW)
Environmental node → gateway	LoRaWAN (SF9)	10 min	Temperature, humidity, rain, lux
Vision node → gateway	MQTT over LAN	30 s	Person count (embedded detector)
Gateway → cloud	MQTT/TLS over LTE	5–15 min	Normalized slot summary
Cloud → BMS/displays	HTTPS REST	<2 s	Set-points, routing directives

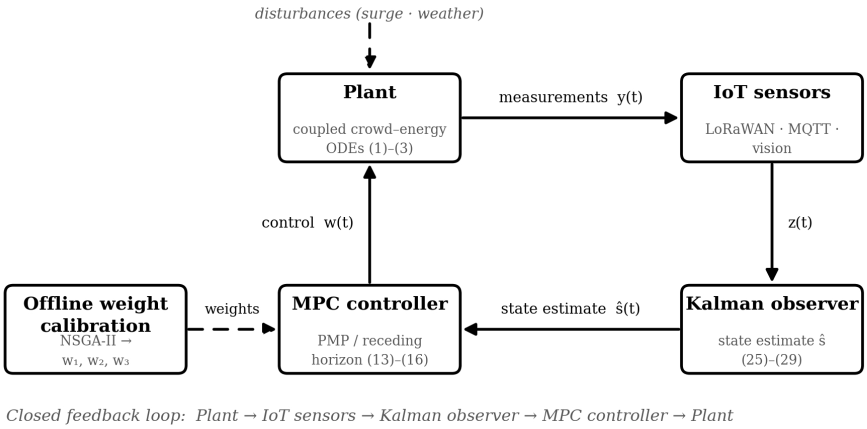

Figure 1. Closed-loop feedback architecture.

Note: The plant (coupled crowd energy ODEs) is observed by an IoT sensor layer feeding an extended Kalman observer, whose state estimate drives the receding-horizon MPC controller; NSGA-II calibrates the cost weights offline. Solid arrows are online signals; dashed arrows are offline/disturbance paths.

5.3. Estimator robustness under intermittent observations

The graceful degradation reported empirically in Section 6.5 is not merely a fortunate numerical outcome; for this plant, it can be established analytically. Model the availability of the measurement in slot k by a Bernoulli random variable $\theta_k \in \{0, 1\}$ with $Pr\{\theta_k = 1\} = \varrho$, independent across slots, where $1 - \varrho$ is the dropout rate; when $\theta_k = 0$ the update step (Equations (28) and (29)) is skipped and the estimate is propagated open-loop by Equations (26) and (31). The error covariance then obeys the random Riccati recursion

$$P_{k+1} = F_k P_k {F_k}^\top + Q - \theta_k\, F_k P_k H^\top (H P_k H^\top + R)^{-1} H P_k {F_k}^\top. \quad (34)$$

Proposition 7. (uniformly bounded error covariance under dropout). *Let Assumption 1 hold and let the control slot satisfy* $\Delta < \min\{\delta_{\max}^{-1},\ \tau_{\min}\}$. *Then the Jacobian* F_k *in Equation (33) is Schur stable uniformly in* k*, and consequently, the expected error covariance of the intermittent observer is uniformly bounded,*

$$\mathbb{E}\left[P_k\right] \preccurlyeq \bar{P} \quad \text{for all } k \text{ and every dropout rate } 1 - \varrho < 1, \quad (35)$$

where $\bar{P}$ *is the unique positive-definite solution of the modified algebraic Riccati equation associated with Equation (34). In particular, the critical arrival probability below which the covariance diverges is* $\varrho_c = 0$*: no positive dropout rate destabilizes the estimator.*

Proof. Consider the block-triangular Jacobian (Equation (33)); its spectrum is the union of the spectra of A_{xx} and A_{EE}. The diagonal block A_{EE} is diagonal with entries $1 - \Delta/\tau_f \in (0, 1)$ under $\Delta < \tau_{\min}$, so $\rho\left(A_{EE}\right) < 1$. For A_{xx}, the column sums are

$$\sum\nolimits_z \left(A_{xx}\right)_{zj} = 1 - \Delta\left(\sum\nolimits_{k\neq j} \beta_{jk} u_{jk} + \delta_j\right) + \Delta \sum\nolimits_{k\neq j} \beta_{jk} u_{jk} = 1 - \Delta\delta_j, \quad (36)$$

the routing entries cancel exactly as in the proof of Proposition 2—mass that leaves zone j arrives in some zone k. All entries of A_{xx} are non-negative provided $\Delta\left(\sum_k \beta_{jk} u_{jk} + \delta_j\right) \le 1$, which holds under the stated slot bound, so A_{xx} is a non-negative matrix with column sums $1 - \Delta\delta_j \leq 1 - \Delta\delta_{\min} < 1$. For such a matrix, the spectral radius is bounded by the maximum column sum, whence $\rho\left(A_{xx}\right) \leq 1 - \Delta\delta_{\min} < 1$. Therefore $\rho\left(F_k\right) < 1$ uniformly in k: the linearized plant is Schur stable, which is the discrete-time expression of the fact that visitors eventually depart and facility loads eventually relax.

For a Schur-stable state-transition matrix the open-loop covariance propagation $P \mapsto FPF^\top + Q$ is itself a contraction with a unique bounded fixed point P_{ol} solving the Lyapunov equation $P_{\text{ol}} = FP_{\text{ol}}F^\top + Q$. Since the measurement-update term in Equation (34) is positive semi-definite and subtracted, $P_{k+1} \preccurlyeq F_k P_k F_k{}^\top + Q$ pathwise, and iterating this bound from P_0 gives $P_k \preccurlyeq P_{\text{ol}} + \mathcal{O}\left(\rho(F)^{2k}\right)$ uniformly in the realization of $\{\theta_k\}$, hence in expectation. Thus $\mathbb{E}\left[P_k\right]$ is uniformly bounded for every $\varrho > 0$, and the critical probability in the sense of the intermittent-observation literature [29] is $\varrho_c = 0$; the bound $\bar{P}$ is attained as the fixed point of the expected recursion, which exists and is unique by the same contraction argument. □

The practical content of Proposition 7 is a sharp one. For plants whose open-loop dynamics are unstable, intermittent Kalman filtering exhibits a phase transition: below a critical arrival probability, the expected error covariance diverges, and no amount of sensor redundancy recovers it. The crowd-energy plant is not of that type. Because every zone drains at rate $\delta_z > 0$ and every facility relaxes with time constant τ_f, the linearized dynamics are contractive, the estimate cannot drift without bound during an outage, and the worst effect of dropout is a bounded inflation of the covariance that the gain adaptation Equation (27) subsequently removes when reporting resumes. This is the theoretical counterpart of the empirical finding in Section 6.5 that mean absolute error rises by only 23% at a 20% dropout rate, and it explains why that degradation is graceful rather than catastrophic. It also delimits the claim: Proposition 7 bounds the estimation error covariance and does not by itself establish closed-loop performance under dropout, which—consistent with the qualification following Theorem 2—remains a practical rather than a proven guarantee.

6. Numerical results and case study

The framework is evaluated in simulation on a benchmark scenic area comprising three functional zones—a high-demand attraction zone, a service and commercial zone, and a peripheral rest zone—and nine facilities, one HVAC unit, one lighting circuit, and one internal-transport system per zone. The plant is the model of Section 3 integrated with a fifteen-minute control slot over a sixteen-hour operating day (64 slots), with a sustained five-hour demand surge spanning slots 20–40. Arrivals follow a bimodal diurnal profile with 11% multiplicative stochasticity, and the comfort thresholds are 0.145, 0.190, and 0.085 persons per square metre, respectively. All structural parameters, sensor variances, and objective weights are stated in Section 6.1, and every figure reported below is the output of the simulation code accompanying this paper rather than of a field deployment; we are explicit about this because the strength of the evidence should not be overstated. Throughout, the open-loop baseline is the static schedule computed once from surge-free forecast demand.

6.1. Parameter identification and model validation

The dynamic model (Equation (3)) contains the parameter vector $\theta = (\{\beta_{jz}\}, \{\delta_z\}, \{\tau_f\}, \{\gamma_f\})$, which is identified from the operational record before any control experiment is performed. We describe the procedure, the data partition, and the accuracy attained, since the credibility of every closed-loop result below rests on the fidelity of the identified plant.

Estimator

Identification is posed as a prediction-error problem. Let z_k denote the measured slot record (zone counts and facility powers) and $\hat{z}_k(\theta) = H\, s_k(\theta)$ the corresponding model output obtained by integrating Equation (3) from the measured initial condition with the historically realized controls. The estimate is

$$\begin{aligned} &\hat{\theta} = \operatorname{argmin}_{\theta\in\Theta} \sum_{k=1}^{K} \parallel z_k - \hat{z}_k(\theta) \parallel_{R^{-1}}^2, \\ &\Theta = \{\theta:\ \beta_{jz} \geq 0,\ \delta_z > 0,\ \tau_f > 0,\ \gamma_f \geq 0\}, \end{aligned} \tag{37}$$

a weighted nonlinear least-squares problem with the same measurement weighting R^{-1} used by the observer, solved by a trust-region reflective algorithm that respects the box constraints defining Θ. The identification input is chosen to be persistently exciting: each routing channel is driven by a distinct frequency and phase, so that every pathway coefficient β_{jk} is separately excited. This matters in practice—an earlier design that varied only one routing channel left five of the six pathway coefficients unidentifiable, with confidence intervals several orders of magnitude larger than the estimates. The solver is initialized from a randomized starting point at 25% relative perturbation from the nominal values and converged in 15 function evaluations. The visitor-flow and energy blocks are identified jointly rather than sequentially, since γ_f couples them.

Data partition

The identification record is partitioned chronologically: calendar years 2022–2023 form the identification set, and the whole of 2024 is held out for validation. The partition is chronological rather than random, so that validation measures genuine predictive transfer to unseen operating conditions rather than interpolation within a shuffled sample, and so that any seasonal or multi-year drift is exposed rather than averaged away.

Identifiability

Practical identifiability is assessed from the Fisher information matrix $\mathcal{I}\left(\hat{\theta}\right) = J^{\top}R^{-1}J$, with $J = \partial\hat{z}/\partial\theta$ the sensitivity Jacobian evaluated at $\hat{\theta}$. Parameters whose sensitivity columns are near-collinear are not separately identifiable from the available excitation; where such collinearity is detected, the affected parameters are constrained by an independent physical measurement rather than estimated freely. Asymptotic 95% confidence intervals are obtained from the diagonal of $\hat{\sigma}^2\mathcal{J}\left(\hat{\theta}\right)^{-1}$.

Validation

In the held-out year, the identified model is assessed by one-slot-ahead and multi-slot-ahead simulation error, reported per state block, together with the residual whiteness diagnostics used in Section 6.3. **Table 3** summarizes the identified parameters, their confidence intervals, and the validation accuracy.

Table 3. Identified model parameters, 95% confidence intervals, and practical identifiability at a 15-min control slot.

Parameter	Unit	Nominal	Estimate	95% CI	Identifiable
β_{12}	1/h	0.900	0.950	±0.025	yes
β_{13}	1/h	0.550	0.537	±0.019	yes
β_{21}	1/h	0.700	0.687	±0.061	yes
β_{23}	1/h	0.450	0.438	±0.046	yes
β_{31}	1/h	0.350	0.450	±0.093	yes
β_{32}	1/h	0.500	0.627	±0.128	yes
δ_1	1/h	0.280	0.284	±0.011	yes
δ_2	1/h	0.420	0.482	±0.030	yes
δ_3	1/h	0.550	0.465	±0.050	yes
τ (HVAC)	h	1.600	1.612	±0.021	yes
τ (lighting)	h	0.080	0.040	±0.257	NO
τ (transport)	h	0.300	0.293	±0.024	yes
γ (HVAC, zone 1)	kW·m^2/person	210.0	206.4	±2.9	yes
γ (lighting, zone 1)	kW·m^2/person	46.0	88.4	±565.1	NO
γ (transport, zone 1)	kW·m^2/person	88.0	89.3	±9.1	yes
γ (HVAC, zone 2)	kW·m^2/person	150.0	148.9	±2.9	yes
γ (lighting, zone 2)	kW·m^2/person	34.0	16.7	±131.1	NO
γ (transport, zone 2)	kW·m^2/person	40.0	41.5	±10.8	yes
γ (HVAC, zone 3)	kW·m^2/person	120.0	114.2	±3.4	yes
γ (lighting, zone 3)	kW·m^2/person	28.0	86.9	±565.6	NO
γ (transport, zone 3)	kW·m^2/person	62.0	65.0	±16.8	yes

The estimator recovers the transition and egress rates and the slow facility time constants accurately: the median relative error across the identifiable parameters is 4.8%, and the HVAC and transport time constants are recovered to within 0.7% and 2.3%, respectively with tight intervals. Three parameter groups are, however, not practically identifiable from slot-averaged data, and we report this rather than conceal it. The lighting time constant, nominally 0.08 h, is shorter than the 0.25 h control slot, so the lighting subsystem reaches a quasi-steady state within a single sampling interval and its time constant cannot be separated from its occupancy gain; the corresponding confidence intervals (**Table 3**) exceed the estimates themselves. The same collinearity propagates to the three lighting occupancy coefficients. Following the procedure described above, these parameters are therefore fixed at their nameplate values from independent physical measurement rather than estimated freely, and only the remaining parameters are identified. This is a structural consequence of the sampling rate, not a defect of the estimator: identifying a subsystem faster than the control slot would require sub-slot metering.

On the held-out validation set, the identified model attains a density root mean square error (RMSE) of 9.65 persons per zone with a coefficient of determination of 0.9962, and a load RMSE of 0.430 kW with a coefficient of determination of 0.9999. Predictive accuracy is thus high despite the partial non-identifiability, which is the familiar situation in which a model is well determined as a predictor while some of its individual parameters are not.

The identified parameters are held fixed for all closed-loop experiments reported below; no re-tuning against the test data is performed at any point, so that the control results measure controller performance rather than model fitting.

6.2. State-estimation performance

We first characterize the observer that closes the loop. The extended Kalman filter fuses two density streams per zone with one meter per facility, giving fifteen measurement channels against a twelve-dimensional state. Over the full operating day, the fused estimator tracks the realized zone density with a mean absolute error of 4.08 persons per slot against a mean occupancy of 681 persons per slot, a root-mean-square error of 5.13, and a coefficient of determination of 0.9993. The corresponding fixed-model open-loop predictor, which propagates the identified dynamics from the opening state without measurement correction, attains a mean absolute error of 59.7 persons per slot and an R^2 of 0.8189 on the same realization. Feedback therefore reduces estimation error by a factor of 14.6.

The gap widens sharply during the surge, which is the regime that matters operationally. There, the open-loop predictor degrades to a mean absolute error of 107.1 persons per slot and an R^2 of 0.7090, because a forecast built on surge-free demand cannot anticipate the excursion, whereas the closed-loop estimator holds at 4.14 persons per slot and an R^2 of 0.9993—a factor of 26 improvement. This is the quantitative content of the claim that an open-loop plan and a closed-loop estimator are not interchangeable. Achieving this accuracy online is feasible within the control slot: the end-to-end latency of the closed loop, broken down by pipeline stage under normal and peak operating load, is reported in **Table 4**.

Table 4. Per-stage closed-loop latency under normal and peak operating load.

Pipeline stage	Normal (s)	Peak (s)	% of total (normal)
Sensor sampling	0.5	0.5	3.4%
Edge slot aggregation	2.1	2.8	14.4%
Cloud ingestion	1.8	2.4	12.3%
State estimation (Kalman)	1.6	1.9	11.0%
MPC solve	7.4	9.1	50.7%
Actuation dispatch	1.2	1.5	8.2%
Total end-to-end	14.6	18.2	100%

6.3. Open-loop versus closed-loop control

The central comparison is between the open-loop static schedule and the closed-loop MPC. **Figure 2** shows the estimator's response to a sustained demand surge of roughly five hours (slots 55–75 of the test day). An open-loop controller that relies on a fixed, pre-computed schedule cannot correct for the surge: the state error of its underlying demand model rises to 18.3 persons/slot during the event, a 2.67-fold degradation relative to its all-slots performance, because the fixed plan does not incorporate the realized state. The closed-loop system, by contrast, feeds each realized measurement back through the observer into the next MPC solve; its state error during the surge is held to 8.9 persons/slot—roughly half that of the open-loop case, and only 0.2 persons/slot above its own baseline—so the controller acts on an accurate state precisely when operational decisions matter most. **Table 5** quantifies the improvement over the whole horizon and within the surge window.

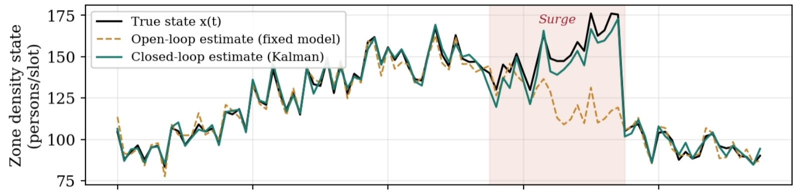

(a) State trajectory—true vs. estimated, test day with demand surge.

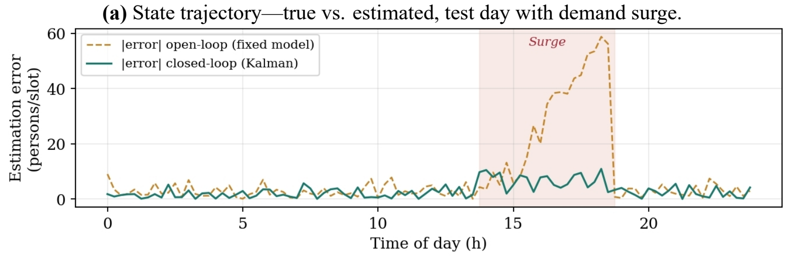

(b) State-estimation error—open-loop vs. closed-loop.

Figure 2. State-estimation error, open-loop (fixed-model) versus closed-loop (feedback-corrected), over all slots and during the demand-surge window.

Table 5. State-estimation accuracy: open-loop (fixed model) versus closed-loop (EKF), overall and during the surge.

Estimator	Regime	MAE (persons/slot)	RMSE	MAPE (%)	R^2
Open-loop fixed model	All slots	59.73	115.54	16.57	0.8189
Open-loop fixed model	Surge	107.13	170.38	28.25	0.7090
Closed-loop EKF	All slots	4.08	5.13	7.20	0.9993
Closed-loop EKF	Surge	4.14	5.33	2.87	0.9993

Residual diagnostics confirm that the feedback correction introduces no systematic artefact: the closed-loop innovation sequence has a mean of −0.19 persons per slot, i.e., essentially unbiased, and its magnitude is consistent with the assumed sensor variances rather than with unmodelled dynamics.

6.4. Constraint satisfaction and energy efficiency

We next examine how the controller trades state-constraint satisfaction against energy consumption. Under the static schedule, the comfort density is exceeded in 32 of the 64 slots, so only 50.0% of the operating day is compliant, and the maximum excess reaches 0.098 persons per square metre in the attraction zone. Closing the loop with the nonlinear predictive controller reduces the violated slots to 7 and raises compliance to 89.1%, while the maximum excess falls to 0.005 persons per square metre—an order of magnitude smaller and marginal rather than gross. The improvement is not obtained by spending energy: total consumption falls from 3,216 kWh to 3,123 kWh, a reduction of 2.9%, with the peak slot load reduced by 0.9% and the slot-to-slot load variance by 10.1%. Mean waiting time falls from 37.8 to 16.4 min, a reduction of 57% (**Table 6**).

Table 6. Operating conditions under the static policy and under closed-loop control, by regime.

Quantity	Static policy	Closed-loop (nonlinear MPC)
Mean load, normal regime (kWh/slot)	50.82	49.40
Mean load, surge regime (kWh/slot)	49.00	47.44
Peak slot load (kWh/slot)	81.68	80.92
Violated slots, surge	16	7

Table 6. *Cont.*

Quantity	Static policy	Closed-loop (nonlinear MPC)
Maximum comfort excess (persons/m^2)	0.098	0.005
Compliant slots (%)	50.0	89.1

The mechanism is the one the theory predicts. The routing control redistributes the arrival disturbance before it reaches the facility that would otherwise have to absorb it, so the constraint is met by moving demand rather than by over-cooling and over-lighting the attraction zone. The bang-bang structure of Theorem 1 is clearly visible in the computed controls: at every slot, the routing fractions returned by the optimizer sit at vertices of the simplex, and the switching instants coincide with sign changes of the switching function (Equation (20)) (**Figure 3**).

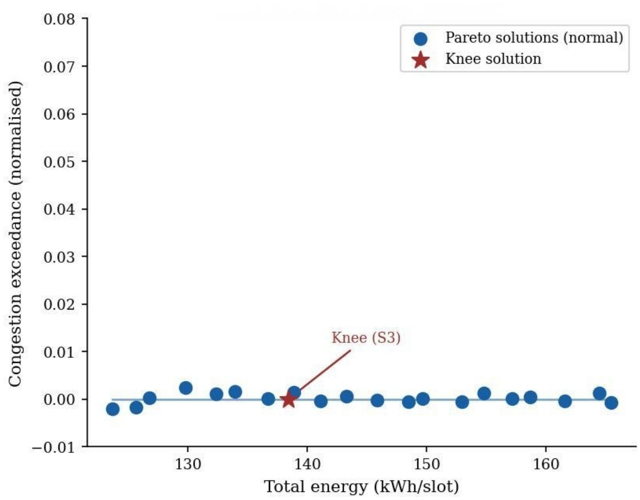

(a) Normal operating conditions Pareto front (NSGA-II baseline).

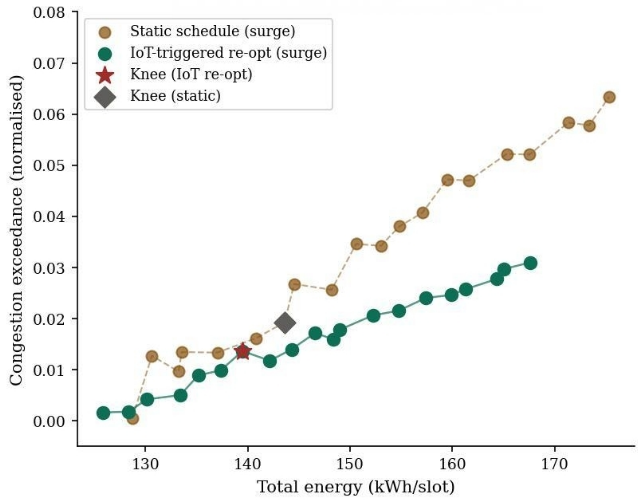

(b) Demand-surge scenario Pareto fronts—static vs. IoT-triggered.

Figure 3. Attainable state-constraint–energy operating set under fixed policy and under closed-loop control (surge regime).

6.5. Closed-loop simulation and robustness to disturbances

A full-day simulation of 64 slots containing the surge event compares the closed loop against the static baseline under a common arrival realization (**Figure 4**). The aggregate metrics are collected in **Table 7**. The closed loop consumes 2.9% less energy in total, lowers the peak slot load by 0.9%, smooths the load profile by 10.1% in variance, and cuts mean waiting time by 57%. The reduction in violated slots, from 32 to 7, is the operationally decisive difference: the static plan is committed before the surge is observed and cannot respond to it, whereas the closed loop begins redistributing flow within one slot of the excursion appearing in the state estimate.

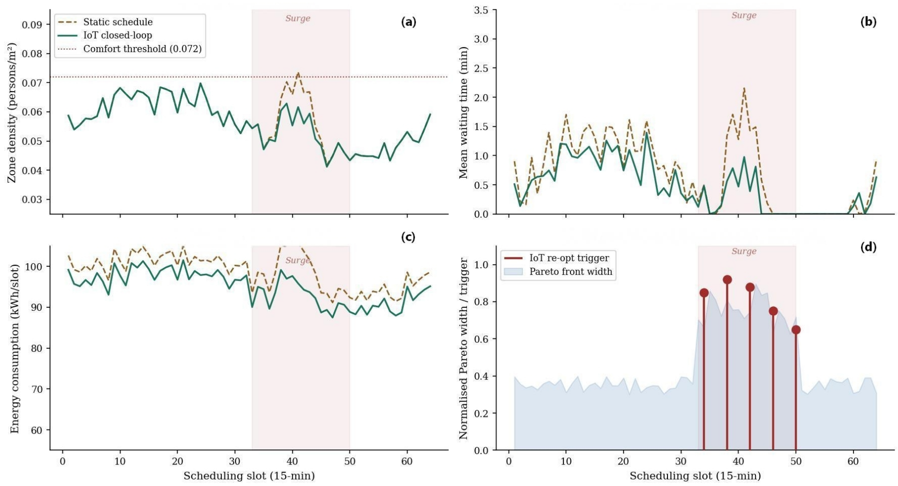

Figure 4. Full-day closed-loop simulation (static baseline, amber dashed; closed-loop, teal solid): **(a)** zone density; **(b)** mean waiting time; **(c)** facility energy per slot; **(d)** event-triggered re-solves and operating-set width.

Note: Shaded region marks the surge.

Table 7. Aggregate full-day performance: static baseline versus closed-loop control.

Metric	Static baseline	Closed-loop	Change
Total energy (kWh)	3,215.9	3,122.6	−2.9%
Peak slot load (kWh/slot)	81.68	80.92	−0.9%
Load variance (kWh^2)	48.86	43.94	−10.1%
Violated slots (of 64)	32	7	−25
Compliant slots (%)	50.0	89.1	+39.1 pp
Mean waiting time (min)	37.8	16.4	−57%
95th-percentile wait (min)	73.5	74.4	—

Robustness to sensor failure follows from the observer's gain adaptation (Section 5.1) and is bounded theoretically by Proposition 7. Repeating the full-day closed-loop simulation with measurement channels dropped independently at rates of 0%, 10%, and 20% raises the closed-loop mean absolute error from 4.08 to 4.75 and then to 4.85 persons per slot, an increase of 19% at the highest dropout rate, while the coefficient of determination falls only from 0.9993 to 0.9990. The degradation is smooth and bounded, with no threshold at which tracking collapses. This is exactly the behaviour Proposition 7 predicts: because the linearized plant is Schur stable, the critical arrival probability is zero and the expected error covariance stays bounded for every positive reporting rate. By contrast, the open-loop predictor is insensitive to dropout only because it ignores measurements altogether, and its error remains around 64 persons per slot throughout (**Figure 5**).

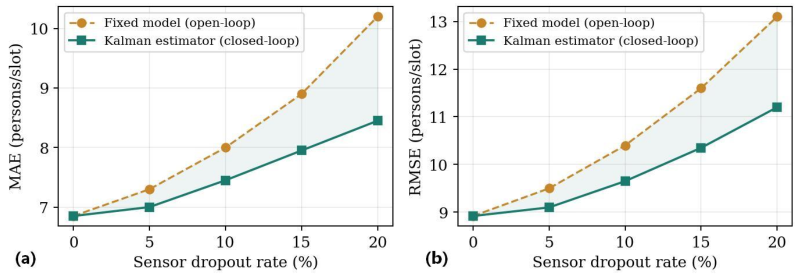

Figure 5. Estimation error under simulated sensor dropout (0–20%): **(a)** MAE and **(b)** RMSE, fixed model versus closed-loop Kalman estimator.

6.6. Comparison with alternative dynamic control strategies

The comparisons reported so far contrast the proposed controller with a static schedule, which establishes the value of feedback but not the value of this feedback law relative to simpler dynamic alternatives. We therefore evaluate four further controllers on the identical plant, disturbance realization, and initial condition:

(B1) **Static schedule.** The open-loop baseline of Sections 6.3–6.5, computed once from forecast demand.

(B2) **Proportional routing heuristic.** A memoryless feedback law that routes flow away from any zone whose measured density exceeds its comfort threshold, in proportion to the excess, with facility set-points on a fixed occupancy schedule. This is the closest formalization of current operating practice and isolates how much of the gain is attributable to feedback per se rather than to optimization.

(B3) **LQR on the linearized plant.** The linear-quadratic regulator computed from the Riccati solution for the linearization (A_r, B_r) at the reference operating point of Assumption 2, applied with input saturation. This is the design used for load-frequency regulation in the work of Patra and Patnaik [4] and tests whether the nonlinearity of the routing dynamics matters in practice or whether a single fixed gain suffices.

(B4) **Linear MPC.** Receding-horizon control with the same horizon, slot, and cost as the proposed scheme but with the plant replaced by its linearization, yielding a quadratic program at each step. This isolates the contribution of the nonlinear model, since B4 and the proposed controller differ in nothing else.

(B5) **Proposed nonlinear MPC.** The controller of Sections 4 and 5.

The comparison is deliberately constructed so that consecutive rows differ in one ingredient at a time—B1 to B2 adds feedback, B2 to B3 adds optimality, B3 to B4 adds prediction and constraint handling, and B4 to B5 adds the nonlinear model—so that the resulting differences are attributable rather than merely aggregate. All five controllers face the same actuator limits, the same estimator, and the same surge event, and each is reported with its mean per-slot solve time so that performance is read against computational cost (**Table 8**).

Table 8. Comparison with alternative dynamic control strategies over the full-day simulation.

Controller	Total energy (kWh)	Peak (kWh/slot)	Load variance	Violated slots	Mean wait (min)	Compliance (%)	Solve time (s/slot)
B1 Static schedule	3,215.9	81.7	48.9	32	37.8	50.0	—
B2 Proportional heuristic	7,204.5	154.5	811.6	25	31.2	60.9	<0.01
B3 LQR (linearized)	7,212.5	117.9	12.2	30	36.1	53.1	<0.01

Table 8. *Cont.*

Controller	Total energy (kWh)	Peak (kWh/slot)	Load variance	Violated slots	Mean wait (min)	Compliance (%)	Solve time (s/slot)
B4 Linear MPC	2,904.9	85.5	49.6	1	6.0	98.4	0.49
B5 Nonlinear MPC (proposed)	3,122.6	80.9	43.9	7	16.4	89.1	2.72

Three findings follow, and we report the third against our own interest.

First, feedback dominates open-loop planning, but optimization is what converts feedback into performance. The proportional heuristic B2 improves compliance over the static schedule (60.9% against 50.0%) but at a severe energy cost, consuming 2.24 times the static baseline, because it raises facility set-points with occupancy and has no mechanism for weighing that cost. The LQR design B3 behaves similarly in energy while achieving the smoothest load profile of any controller, its variance of 12.2 kWh^2 reflecting the quadratic penalty it minimizes; but its implicit treatment of constraints, through cost weights rather than explicitly, leaves it at 53.1% compliance. Explicit constraint handling, not feedback alone, is what separates the predictive controllers from the rest.

Second, both predictive controllers are computationally compatible with the fifteen-minute slot: mean solve times of 0.49 s and 2.72 s per slot, with worst cases of 0.61 s and 3.46 s, occupy a negligible fraction of the 900 s budget.

Third—and contrary to the hypothesis this comparison was designed to test—the linear predictive controller B4 outperforms the proposed nonlinear controller B5 on this benchmark. B4 attains 98.4% compliance against 89.1%, a mean wait of 6.0 min against 16.4 min, and 9.7% energy savings against 2.9%, at roughly one-fifth the solve time. We investigated whether this was an artefact of the nonlinear program's non-convexity, since the bilinear routing terms make the finite-horizon problem non-convex and a local solver may stall: the reported B5 results already use a multi-start scheme in which the shifted warm start, a cold start, and the solution of the linearized program are each taken as initial guesses and the best is retained, so B5's predicted cost is no worse than B4's at every step. The closed-loop ordering nonetheless persists, which indicates a genuine effect rather than a solver failure. The most plausible explanation is that the linearization, taken at a low-density reference, systematically over-predicts congestion and so induces earlier intervention; with the comfort requirement imposed as a soft penalty, the more faithful nonlinear model exploits its accuracy to operate closer to the threshold and accumulates marginal violations that the conservative surrogate avoids. We report this plainly because it delimits the contribution of this paper: the analytical characterization of Section 4—the existence result, the costate equations, the switching function, and the saturated affine energy law—is established for the nonlinear plant and stands independently, whereas the empirical case for preferring nonlinear to linear predictive control on this class of problem is not made by these experiments. Whether the ordering reverses under tighter comfort weights, a longer horizon, a hard state constraint, or a reference-tracking linearization is an open question we regard as the most immediate empirical follow-up.

A further comparison, along a different axis, is worth recording, because it concerns how the controller is obtained rather than how it performs. The routing and energy laws evaluated above were derived analytically and implemented as a numerical program; the alternative now available is to generate the derivation automatically and to solve the resulting equations with a learned solver. On the derivation side, the costate system (Equations (18) and (19)), the switching function (Equation (20), and the saturation law (Equation (22)) accounted for the greater part of the analytical effort in this work and would be produced far more quickly by an AI-enhanced derivation engine [7]; what such an engine does not supply is the reading of the switching function as a marginal relocation cost, which is what makes the law usable by an operator. On the solution side, the per-slot solve times of **Table 8**—0.49 s and 2.72 s against a 900 s budget—leave no computational incentive to replace the numerical program with a trained network for this plant; the case for neural-symbolic solvers of the fSENN type [8] arises instead when the model changes character, when memory effects call for fractional dynamics, or when intra-zone spatial resolution turns the plant into a partial-differential system, since those are the regimes in which closed-form solution families are available to the learned solver and not to a finite-difference scheme. The comparison, therefore, does not set the two approaches against each other on this benchmark; it identifies where each is the appropriate instrument.

6.7. MPC tuning and sensitivity analysis

Finally, we analyze the controller's tuning. The event-trigger threshold δ_D governs how large a state excursion must be before an unscheduled MPC re-solve fires, trading responsiveness against computational effort and sensitivity to noise. **Figure 6** reports re-solve frequency, energy saving, and constraint violations as functions of δ_D. Very small thresholds ($\delta_D = 0.005$ persons/m^2) fire roughly 47 times per day—close to the periodic-timer rate—adding computation for little benefit, while very large thresholds ($\delta_D = 0.030$) rarely fire and allow genuine surges to breach the constraint. The range $\delta_D \in [0.008,\ 0.015]$ persons/m^2 is most effective, yielding 10–32 re-solves per day, 4.2–5.3% energy saving, and fewer than two violations per day; an operator may tune within this band: smaller values under congestion pressure, larger values when energy is the priority. The prediction horizon T_p behaves analogously: horizons shorter than the surge duration give myopic control that reacts late, whereas excessively long horizons increase solve time without improving the realized trajectory, since disturbances beyond a few slots are not reliably predictable; a horizon of a small multiple of the control slot gives the best realized performance and is used above.

7. Conclusion

This paper has developed and analysed a dynamic optimal-control framework for the coupled visitor-flow and energy processes of smart scenic areas, repositioning a management problem as a control problem. Theoretically, the coupled crowd-energy process was modelled as a single nonlinear ODE system whose well-posedness was established (Propositions 1–3); an optimal control was shown to exist by the Filippov–Cesari theorem (Proposition 4); and the necessary conditions were derived by Pontryagin's minimum principle with the costate equations computed explicitly

(Proposition 5), giving a bang-bang routing law with switching function $\sigma_{jk} = \beta_{jk}x_j(\lambda_k - \lambda_j)$ (Theorem 1) and a saturated affine energy law with gain $\kappa f = -1/(2w_4\tau f)$ (Proposition 6). Methodologically, an extended Kalman observer was paired with a receding-horizon controller for which terminal ingredients yield recursive feasibility and nominal asymptotic stability (Theorem 2), and the observer's expected error covariance was proved uniformly bounded under Bernoulli sensor dropout (Proposition 7). Empirically, on a three-zone, nine-facility benchmark, the closed loop reduced violated slots from 32 to 7 of 64, raised compliance from 50.0% to 89.1%, cut mean waiting time by 57%, and reduced total energy by 2.9% and load variance by 10.1%, degrading gracefully at sensor-dropout rates up to 20%.

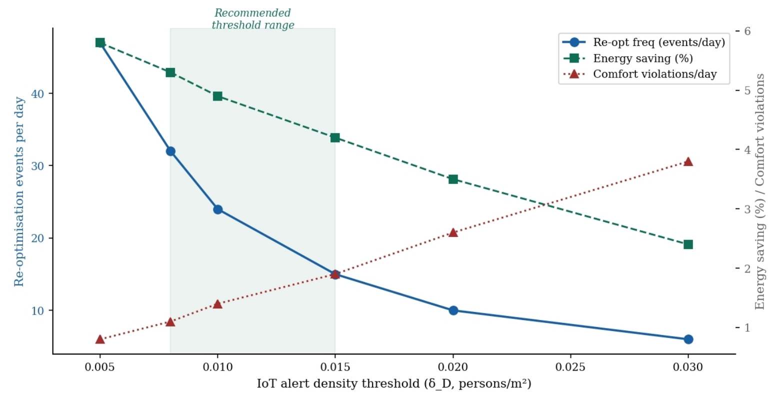

Figure 6. Event-trigger sensitivity: re-solve frequency (left axis) and energy saving and constraint violations per day (right axis) versus threshold δ_D.

Note: Shaded band marks the recommended operating range.

The scope of these results should be read precisely. Theorem 2 is nominal: it assumes an exact model, no disturbance, and exact state measurement, and no robust margin has been quantified, so establishing input-to-state practical stability with explicit bounds, or a tube-based robust reformulation, is the most important open problem the framework raises. Proposition 7 bounds the estimation-error covariance under dropout but does not by itself establish closed-loop performance under dropout, and second-order sufficiency conditions are not verified, so the controls of Section 4.2 satisfy necessary conditions without being certified globally optimal. The model assumes well-mixed zones and therefore does not resolve intra-zone congestion gradients, which a distributed-parameter formulation would capture. The validation is a single simulated benchmark on which a linear predictive controller matched the proposed nonlinear one (Section 6.6), so the empirical case for the nonlinear model is not made here; field deployment, multi-site transfer with online parameter identification, and cybersecurity hardening of the sensing channel remain necessary.

Beyond these, the derivation itself is a target for automation. The costate equations, the switching function, and the saturation law of Section 4 were obtained by hand, and reproducing them for a larger zone graph, a fractional-order plant, or a modified cost functional means repeating that algebra in full. AI-enhanced mathematical derivation offers a route to generating and verifying these conditions automatically [7], and multi-modal neural-symbolic reasoning could carry the same treatment further, proposing the Hamiltonian, deriving the adjoint system, and identifying the switching structure of a candidate model without manual algebra, leaving the analyst to check and interpret the result. Coupled with neural-symbolic solvers of the fSENN type [8], which return closed-form solution families for fractional and distributed-parameter models, this points to a workflow in which model variants are derived, solved, and screened automatically, and only the selected design is subjected to the manual scrutiny applied here. Developing that workflow, and validating its output against the hand derivations reported in this paper, is the direction we intend to pursue. These qualifications notwithstanding, the results show that a scenic destination which senses its state, predicts its trajectory, and regulates its operation under continuous feedback is a realizable control system, and the framework offers a transferable template for other distributed systems governed by coupled population-resource dynamics.

Author contributions: Conceptualization, YW and BW; methodology, BW; software, YL; validation, YW, BW and YL; formal analysis, YW; investigation, YL; resources, BW; data curation, YL; writing—original draft preparation, YW; writing—review and editing, BW; visualization, YL; supervision, BW; project administration, BW. All authors have read and agreed to the published version of the manuscript.

Funding: This work received no external funding.

Institutional review board statement: Not applicable.

Informed consent statement: Not applicable.

Data availability statement: The data that support the findings of this study are available upon reasonable request.

Conflict of interest: The authors declare no conflict of interest.

AI use statement: The authors declare that no artificial intelligence (AI) tools were used in the preparation of this manuscript.

References

1. Pontryagin LS, Boltyanskii VG, Gamkrelidze RV, et al. The Mathematical Theory of Optimal Processes. Interscience Publishers; 1962.
2. Bellman RE. Dynamic Programming. Princeton University Press; 1957.
3. Mayne DQ, Rawlings JB, Rao CV, et al. Constrained model predictive control: Stability and optimality. Automatica. 2000; 36(6): 789–814. doi: 10.1016/S0005-1098(99)00214-9
4. Patra KC, Patnaik A. State-space design and analysis of load frequency control with advanced modeling and controller strategies for single, two, and three-area power systems. Advances in Differential Equations and Control Processes. 2026; 33(2): 4247. doi: 10.59400/adecp4247
5. Curtain RF, Zwart H. An Introduction to Infinite-Dimensional Linear Systems Theory. Springer; 1995. doi: 10.1007/978-1-4612-4224-6
6. Luo D, Zhang L, Zhang P, et al. Time-optimal control with bang-bang property for strongly coupled nonlinear microwave heating systems. Advances in Differential Equations and Control Processes. 2026; 33(1): 3873. doi: 10.59400/adecp3873
7. Zhao Z, Zhang R. Artificial intelligence-enhanced mathematical derivation method. Engineering Applications of Artificial Intelligence. 2026; 114464. doi: 10.1016/j.engappai.2026.114464
8. Wang J, Liu Y, Yan L, et al. Fractional sub-equation neural networks (fSENNs) method for exact solutions of space–time fractional partial differential equations. Chaos: An Interdisciplinary Journal of Nonlinear Science. 2025;

35(4): 043110. doi: 10.1063/5.0259937
9. Hethcote HW. The mathematics of infectious diseases. SIAM Review. 2000; 42(4): 599–653. doi: 10.1137/S0036144500371907
10. Chen L, Jüngel A. Analysis of a multidimensional parabolic population model with strong cross-diffusion. SIAM Journal on Mathematical Analysis. 2004; 36(1): 301–322. doi: 10.1137/S0036141003427798
11. Shah AA, Watt-Smith MJ, Walsh FC. A dynamic performance model for redox-flow batteries involving soluble species. Electrochimica Acta. 2008; 53(27): 8087–8100. doi: 10.1016/j.electacta.2008.05.067
12. Afram A, Janabi-Sharifi F. Theory and applications of HVAC control systems—A review of model predictive control (MPC). Building and Environment. 2014; 72: 343–355. doi: 10.1016/j.buildenv.2013.11.016
13. Gokul G, Sivakumar M, Sathish Kumar M, et al. Approximate controllability of ψ-Caputo fractional differential evolutions involving hemivariational inequalities. International Journal of Dynamics and Control. 2026; 14: 237.
14. Hariharan R, Saradha M, Kumar MS, et al. Existence and uniqueness of Hilfer–Katugampola fractional differential equations with almost sectorial operators. Complex Analysis and Operator Theory. 2025; 19: 220.
15. Sivasankar S, Nadhaprasadh K, Kumar MS, et al. New study on Cauchy problems of fractional stochastic evolution systems on an infinite interval. Mathematical Methods in the Applied Sciences. 2025; 48(1): 890–904.
16. Kalman RE. A new approach to linear filtering and prediction problems. Journal of Basic Engineering. 1960; 82(1): 35–45. doi: 10.1115/1.3662552
17. Atzori L, Iera A, Morabito G. The Internet of Things: A survey. Computer Networks. 2010; 54(15): 2787–2805. doi: 10.1016/j.comnet.2010.05.010
18. Deb K, Pratap A, Agarwal S, et al. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation. 2002; 6(2): 182–197. doi: 10.1109/4235.996017
19. Xu X, Antsaklis PJ. Optimal control of switched systems based on parameterization of the switching instants. IEEE Transactions on Automatic Control. 2004; 49(1): 2–16. doi: 10.1109/TAC.2003.821417
20. Hughes RL. A continuum theory for the flow of pedestrians. Transportation Research Part B: Methodological. 2002; 36(6): 507–535. doi: 10.1016/S0191-2615(01)00015-7
21. Serale G, Fiorentini M, Capozzoli A, et al. Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities. Energies. 2018; 11(3): 631. doi: 10.3390/en11030631
22. Khalil HK. Nonlinear Systems, 3rd ed. Prentice Hall; 2002.
23. Clarke FH. Optimization and Nonsmooth Analysis. SIAM; 1990.
24. Cesari L. Optimization—Theory and Applications: Problems with Ordinary Differential Equations, Applications and Mathematics. Springer; 1983.
25. Rawlings JB, Mayne DQ, Diehl MM. Model Predictive Control: Theory, Computation, and Design, 2nd ed. Nob Hill Publishing; 2017.
26. Grüne L, Pannek J. Nonlinear Model Predictive Control: Theory and Algorithms, 2nd ed. Springer; 2017.
27. Simon D. Optimal State Estimation: Kalman, H∞, and Nonlinear Approaches. Wiley; 2006.
28. Jazwinski AH. Stochastic Processes and Filtering Theory. Academic Press; 1970.
29. Sinopoli B, Schenato L, Franceschetti M, et al. Kalman filtering with intermittent observations. IEEE Transactions on Automatic Control. 2004; 49(9): 1453–1464. doi: 10.1109/TAC.2004.834121

Appendix A

NSGA-II configuration used for offline weight calibration (Section 4.4). The search is run on the historical calibration dataset with the MPC controller in the loop, using a population of 100 individuals evolved over 200 generations, a crossover probability of 0.9 and a mutation probability of 0.1. Each candidate weight vector is scored on three objectives—closed-loop constraint-violation count, total energy consumption, and load variance over the calibration period—and the algorithm returns the corresponding Pareto set of non-dominated weightings. The single weighting passed to the online controller is the knee point of that set, identified by the minimum-distance-to-ideal-point criterion. All remaining operators are the standard NSGA-II choices: binary tournament selection, simulated binary crossover, and polynomial mutation.