

Article

Communication-Efficient Multi-Objective Edge Energy Management for a Grid-Connected Microgrid Using a Day-Ahead Strategy Library

Hanyu Dong ^{1,2}, Jun Lai ¹, Kaiyun Zhou ¹, Changsheng Liu ^{1,2}, Yuming Liao ^{3,*}  and Heng Nian ³ 

¹ Huzhou Power Supply Company, State Grid Zhejiang Electric Power Co., Ltd., Huzhou 313000, China

² Zhejiang Talent Electric Group Co., Ltd., Huzhou 313000, China

³ College of Electrical Engineering, Zhejiang University, Hangzhou 310027, China; nianheng@zju.edu.cn

* Correspondence: yumingliao@zju.edu.cn

Abstract

To reduce the communication and centralized-computation burden associated with frequent intraday upper-level updates under renewable-energy uncertainty, this paper proposes an edge-autonomous multi-objective energy management strategy with an offline–online architecture. In the day-ahead stage, Gaussian Copula modeling and trajectory screening generate temporally correlated photovoltaic (PV) scenarios. The conventional Strength Pareto Evolutionary Algorithm 2 (SPEA2) first constructs a base library C0. A dual-space coordinated selection mechanism then retains every C0 strategy exactly and adds complementary trajectories according to their objective responses and differences in energy storage system power trajectories. In the intraday stage, the edge controller combines current local measurements with preloaded PV, load, and price profiles, propagates every stored candidate over a receding horizon, and applies the first action of the strict minimum-composite-cost candidate. The case-study results show that, without relying on intraday upper-level communication, the proposed method achieves a composite-objective value close to that of the perfect-information centralized reference.

Keywords: microgrid; multi-objective energy management; communication-efficient operation; edge autonomy; strategy library; receding-horizon candidate evaluation; offline–online hierarchical architecture

1. Introduction

Microgrids provide a practical framework for integrating distributed renewable energy and flexible demand at the local level. However, increasing photovoltaic (PV) penetration also introduces stronger variability and forecasting uncertainty, which complicate the coordination of operating cost, load service, and renewable-energy utilization [1,2]. These challenges are particularly relevant when the energy management system must operate with limited communication resources and modest edge-computing capability.

Centralized microgrid energy management generally relies on frequent upper-level acquisition of source–load states and repeated delivery of updated control instructions, while a fixed day-ahead schedule cannot continuously accommodate intraday deviations. Repeatedly solving a multi-objective optimization problem at short control intervals may also exceed the computational capability of an edge controller. The present study therefore investigates whether a controller based on local measurements and actuation can reduce the need for intraday upper-level updates after successful delivery of a day-ahead strategy



Received: 23 July 2026

Revised: 6 September 2026

Accepted: 8 September 2026

Published: 11 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

package, while maintaining economic operation, flexible-load service, PV utilization, and energy-storage feasibility.

Existing studies reduce communication dependence through distributed or decentralized coordination, event- or self-triggered updates, and compensation for communication delay and packet loss [3–18]. These approaches primarily address information exchange and dynamic regulation at the control layer. Communication-aware scheduling and edge-cloud architectures further consider uncertainty management and the coordination of computation and communication resources [19–22]. Table 1 summarizes their functional positioning relative to the energy-management method developed in this study.

Table 1. Functional positioning of representative low-communication approaches and the proposed method.

Approach	Main Mechanism	Primary Function	Relevance and Remaining Issue for Multi-Objective Energy Management
Distributed or decentralized coordination [3–6]	Reference sharing, local observations, and consensus-based coordination	Reduce centralized information dependence for voltage, frequency, current, or power regulation	Mainly addresses control-layer coordination; day-ahead uncertainty management and intraday multi-objective scheduling are generally not integrated.
Event-, self-, or hybrid-triggered control [7–11]	Communication and control updates are activated by state- or error-based conditions	Reduce sampling and communication frequency	Communication is still required when triggering occurs, and prolonged outages or strategy diversity for local energy management are not the main focus.
Delay- and packet-loss-compensated control [12–18]	Predictive compensation, buffering, filtering, distributed control, or gain scheduling	Maintain dynamic regulation under delayed or missing information	Improves tolerance to impaired communication, but usually retains online information exchange and focuses on dynamic control rather than multi-objective dispatch.
Communication-aware scheduling and architecture [19–22]	Distributed robust optimization, edge-cloud coordination, software-defined networking, or local reference compensation	Coordinate scheduling and communication reliability	Considers uncertainty or communication architecture, but finite-candidate local execution with a diverse precomputed strategy set is not jointly addressed.
Proposed method	Day-ahead scenario-based strategy-library construction and intraday receding-horizon candidate evaluation	Separate complex offline optimization from minute-level edge execution	Targets multi-objective energy management under limited communication; applicability is bounded by the strategy-library validity period, local sensing, and the modeled physical constraints.

The comparison concerns functional positioning rather than a direct performance ranking because the referenced methods address different control layers, system models, and test conditions.

Despite the progress above, three gaps remain relevant to communication-constrained microgrid energy management. First, a static day-ahead schedule cannot continuously adapt to intraday source–load deviations. Second, selecting candidate solutions only according to objective-space performance may retain storage trajectories with similar physical actions and thereby restrict the choices available to a local controller. Third, repeated online solution of a complex multi-objective model is difficult to implement on a resource-limited edge device. These gaps motivate a layered mechanism that connects day-ahead multi-objective search, strategy diversity, and physically feasible local execution.

To address these issues, this paper proposes a communication-efficient edge-autonomous multi-objective energy management strategy. The method adopts a spatio-

temporally decoupled architecture that combines day-ahead offline optimization with intraday local execution. Temporally correlated PV scenarios are generated using a Gaussian Copula model and screened using a max–min trajectory-distance criterion. The conventional Strength Pareto Evolutionary Algorithm 2 (SPEA2) first forms a base library. The proposed dual-space coordinated selection mechanism then retains the complete base library and adds complementary candidates by jointly considering their objective contributions, decision-space separation, and scenario-response diversity. The resulting variable-size strategy library, together with the required day-ahead PV, load, and price profiles, is delivered to the local controller before intraday operation. During execution, the library is treated as a finite set of candidate energy storage system (ESS) power sequences. At each interval, current PV output, load demand, and SOC are obtained from local measurements, while the preloaded profiles provide the PV, load, and price values for subsequent steps in the prediction horizon.

The main contributions are summarized as follows:

1. An uncertainty-aware day-ahead strategy-library construction method is developed by combining temporally correlated PV scenario generation, trajectory-based scenario screening, and dual-space coordinated selection. The mechanism retains all conventional SPEA2 base strategies and adds complementary candidates to broaden coverage in both objective and decision spaces.
2. A receding-horizon edge execution mechanism is formulated to evaluate precomputed candidate trajectories using current local PV, load, and state-of-charge measurements together with preloaded forecast and price profiles. Grid-interaction cost, unserved flexible-load demand, PV curtailment, the in-horizon terminal-SOC condition, and dynamic storage constraints are considered without repeatedly solving the full day-ahead multi-objective problem online.
3. A comparative evaluation framework is established that includes a conventional strategy library, a static open-loop scheme, a perfect-information centralized reference, deterministic online model predictive control (MPC), scenario-based stochastic MPC, periodic and event-triggered communication, communication outages, seven constructed day types, persistent multi-day forecast-error sequences, point of common coupling (PCC) limit stress cases, and computational profiling. Communication counts and centralized optimization counts are reported separately to distinguish information exchange from optimization effort.

2. Edge-Autonomous Energy Management Framework Under Upper-Level Communication Constraints

2.1. Strategy Construction

The proposed framework separates computationally intensive day-ahead optimization from minute-level intraday execution, as illustrated in Figure 1. In the day-ahead stage, a set of screened PV scenarios is embedded in a multi-objective SPEA2 optimization. Each candidate ESS trajectory is evaluated under all retained scenarios, and the mean objective values are used for multi-scenario fitness assessment. Environmental selection maintains an external archive of nondominated solutions.

After convergence, the conventional archive forms the base library C0. The proposed dual-space coordinated selection mechanism enriches this base with complementary trajectories without deleting or modifying any C0 trajectory. The resulting daily library C is transmitted to the local controller before intraday operation.

At each intraday control interval, the local controller acquires the current PV output, base-load demand, flexible-load demand, and ESS state of charge (SOC) through local measurements. The measured SOC initializes the state prediction, while the measured PV

and load values are used in the first-step power allocation. For subsequent steps in the prediction horizon, the controller uses the PV, load, and price profiles preloaded before operation. The candidate strategies are then evaluated using the local power-allocation model and matching cost defined below, and the first action of the selected candidate is executed as described in Section 4.




Figure 1. Schematic diagram of the proposed edge-autonomous multi-objective energy management framework. PV, photovoltaic; SPEA2, Strength Pareto Evolutionary Algorithm 2; ESS, energy storage system; SOC, state of charge; DC/AC, direct current/alternating current.

2.2. Matching Cost

For candidate strategy k at time t , the internal supply, base-load deficit, and local surplus before grid exchange and PV curtailment are

$$\begin{cases} P_{\text{int},k}(t) = P_{\text{pv,act}}(t) + P_{\text{ess},k}(t), \\ D_{\text{base},k}(t) = \max\{P_{\text{base,act}}(t) - P_{\text{int},k}(t), 0\}, \\ S_k(t) = \max\{P_{\text{int},k}(t) - P_{\text{base,act}}(t), 0\}. \end{cases} \quad (1)$$

Here, $P_{\text{ess},k}(t) > 0$ denotes ESS discharging and $P_{\text{ess},k}(t) < 0$ denotes charging. The PCC import and export capacities are denoted by $\bar{P}_{\text{grid,buy}}$ and $\bar{P}_{\text{grid,sell}}$, respectively.

The grid first supplies any base-load deficit, after which its remaining import capacity may supply flexible load during valley-price periods:

$$\begin{cases} P_{\text{grid,buy},k}^{\text{base}}(t) = \min[D_{\text{base},k}(t), \bar{P}_{\text{grid,buy}}], \\ P_{\text{base,unm},k}(t) = \max\{D_{\text{base},k}(t) - P_{\text{grid,buy},k}^{\text{base}}(t), 0\}, \\ P_{\text{flex,loc},k}(t) = \min[S_k(t), P_{\text{flex,dem}}(t)], \\ P_{\text{grid,buy},k}^{\text{flex}}(t) = I_{\text{val}}(t) \min\{P_{\text{flex,dem}}(t) - P_{\text{flex,loc},k}(t), \max\{\bar{P}_{\text{grid,buy}} - P_{\text{grid,buy},k}^{\text{base}}(t), 0\}\}, \\ P_{\text{grid,buy},k}(t) = P_{\text{grid,buy},k}^{\text{base}}(t) + P_{\text{grid,buy},k}^{\text{flex}}(t), \\ P_{\text{unm},k}(t) = P_{\text{flex,dem}}(t) - P_{\text{flex,loc},k}(t) - P_{\text{grid,buy},k}^{\text{flex}}(t). \end{cases} \quad (2)$$

where $I_{\text{val}}(t) = 1$ during valley-price periods and $I_{\text{val}}(t) = 0$ otherwise.

After local flexible-load service, the remaining surplus is exported up to the PCC limit. Any residual PV power is curtailed:

$$\begin{cases} R_k(t) = \max\{S_k(t) - P_{\text{flex,loc},k}(t), 0\}, \\ P_{\text{grid,sell},k}(t) = \min[R_k(t), \bar{P}_{\text{grid,sell}}], \\ P_{\text{pv,cur},k}(t) = \min\{\max\{R_k(t) - P_{\text{grid,sell},k}(t), 0\}, P_{\text{pv,act}}(t)\}, \\ P_{\text{spill},k}(t) = \max\{R_k(t) - P_{\text{grid,sell},k}(t) - P_{\text{pv,cur},k}(t), 0\}. \end{cases} \quad (3)$$

The allocation is feasible only if $P_{\text{base,unm},k}(t) = 0$ and $P_{\text{spill},k}(t) = 0$. Thus, a candidate is rejected when its base-load deficit exceeds the combined local supply and PCC import capability, or when ESS discharge creates a surplus that cannot be absorbed, exported, or curtailed as PV. For a feasible candidate, the minute-level power balance is

$$\begin{aligned} P_{\text{pv,act}}(t) - P_{\text{pv,cur},k}(t) + P_{\text{ess},k}(t) + P_{\text{grid,buy},k}(t) \\ = P_{\text{base,act}}(t) + P_{\text{flex,dem}}(t) - P_{\text{unm},k}(t) + P_{\text{grid,sell},k}(t). \end{aligned} \quad (4)$$

The one-step matching cost is defined as follows:

$$\begin{aligned} j_k(t) = [c_{\text{buy}}(t)P_{\text{grid,buy},k}(t) - c_{\text{sell}}(t)P_{\text{grid,sell},k}(t) \\ + c_{\text{cyc}}|P_{\text{ess},k}(t)| + \lambda_{\text{flex}}P_{\text{unm},k}(t) + \lambda_{\text{curt}}P_{\text{pv,cur},k}(t)]\Delta t, \end{aligned} \quad (5)$$

where $c_{\text{buy}}(t)$ and $c_{\text{sell}}(t)$ are the electricity purchase and selling prices, respectively; c_{cyc} is the ESS cycling-degradation coefficient; and $P_{\text{pv,cur},k}(t)$ is curtailed PV power. The coefficients λ_{flex} and λ_{curt} are the penalty coefficients for unserved flexible-load energy and curtailed PV energy, respectively, with units of CNY/kWh. Their sensitivity is evaluated in Section 5.4.5.

2.3. Objective Functions

The day-ahead strategy library is generated using three objectives: economic operating cost, unserved flexible-load ratio, and PV-curtailment ratio. The first objective is defined as follows:

$$F_1 = \sum_{t=1}^T [c_{\text{buy}}(t)P_{\text{grid,buy}}(t) - c_{\text{sell}}(t)P_{\text{grid,sell}}(t) + c_{\text{cyc}}|P_{\text{ess}}(t)|]\Delta t, \quad (6)$$

The terminal SOC is imposed within the same daily horizon through the constraints in Section 2.4. Consequently, all grid exchange and ESS cycling associated with satisfying the terminal condition are already included in Equation (6).

The unserved flexible-load ratio is defined as follows:

$$F_2 = \frac{\sum_{t=1}^T P_{\text{unm}}(t)\Delta t}{\sum_{t=1}^T P_{\text{flex,dem}}(t)\Delta t}. \quad (7)$$

The PV-curtailment ratio is defined as follows:

$$F_3 = \frac{\sum_{t=1}^T P_{\text{pv,cur}}(t)\Delta t}{\sum_{t=1}^T P_{\text{pv,act}}(t)\Delta t}. \quad (8)$$

Accordingly, the offline multi-objective optimization is formulated as follows:

$$\min_x F(x), \quad F(x) = [F_1(x), F_2(x), F_3(x)]^T. \quad (9)$$

Here, α denotes the ESS strategy vector, and the constraints are given in Section 2.4.

The three-objective formulation is retained in the SPEA2 search to preserve Pareto trade-offs. A weighted full-day indicator is used for consistent case-study comparison:

$$J_{\text{day}} = C_{\text{eco}} + \lambda_{\text{flex}} E_{\text{flex,unm}} + \lambda_{\text{curt}} E_{\text{curt}}, \quad (10)$$

where $C_{\text{eco}} = F_L$, $E_{\text{flex,unm}} = \sum_t P_{\text{unm}}(t) \Delta t$, and $E_{\text{curt}} = \sum_t P_{\text{pv,cur}}(t) \Delta t$. The same economic and penalty terms are used for the additive library selection in Section 3.2 and the rolling candidate evaluation in Section 4, while J_{day} is used for complete-day reporting.

2.4. Constraints

The local feasibility constraints are independent of remote commands and are applied to every candidate trajectory. The ESS SOC constraint is written as follows:

$$S_{\text{SOC,min}} \leq S_{\text{SOC}}(t) \leq S_{\text{SOC,max}}. \quad (11)$$

The ESS power constraint is written as follows:

$$P_{\text{min}}(t) \leq P_{\text{ess}}(t) \leq P_{\text{max}}(t), \quad (12)$$

where the state-dependent limits are defined as follows:

$$\begin{cases} P_{\text{max}}(t) = \min \left[P_{\text{ess, rat}}, \frac{(S_{\text{SOC}}(t) - S_{\text{SOC,min}}) E_{\text{ess, cap}} \eta}{\Delta t} \right] \\ P_{\text{min}}(t) = -\min \left[P_{\text{ess, rat}}, \frac{(S_{\text{SOC,max}} - S_{\text{SOC}}(t)) E_{\text{ess, cap}}}{\eta \Delta t} \right] \end{cases} \quad (13)$$

The SOC-based limits are combined with the PCC/base-load conditions used by the physical allocator:

$$\begin{cases} P_{\text{min}}^f(t) = \max \left[P_{\text{min}}(t), P_{\text{base, act}}(t) - P_{\text{pv, act}}(t) - \bar{P}_{\text{grid, buy}} \right], \\ P_{\text{max}}^f(t) = \min \left[P_{\text{max}}(t), P_{\text{base, act}}(t) + P_{\text{flex, dem}}(t) + \bar{P}_{\text{grid, sell}} \right]. \end{cases} \quad (14)$$

The lower limit prevents PCC import capacity from being exceeded when supplying the base load. The upper limit prevents ESS discharge from exceeding the power that can be absorbed by local loads and permitted export even when all PV can be curtailed. If $P_{\text{min}}^f(t) > P_{\text{max}}^f(t)$, no feasible ESS action exists for that candidate and operating condition.

The SOC transition is defined as follows:

$$S_{\text{SOC}}(t+1) = S_{\text{SOC}}(t) - \lambda(t) \frac{P_{\text{ess}}(t) \Delta t}{E_{\text{ess, cap}}}, \quad (15)$$

where the efficiency factor is defined as follows:

$$\lambda(t) = \begin{cases} 1/\eta, & P_{\text{ess}}(t) \geq 0, \\ \eta, & P_{\text{ess}}(t) < 0. \end{cases} \quad (16)$$

The terminal SOC is constrained at the end of the reported daily horizon:

$$S_{\text{SOC}}(T+1) = S_{\text{SOC, ref}}. \quad (17)$$

During the final N_{cl} control intervals, let $r_t = T - t$ denote the number of intervals remaining after the current action. Based on the remaining ESS charging/discharging capability, the next SOC is restricted to

$$\begin{aligned} L_t &= \max \left[S_{SOC,min}, S_{SOC,ref} - \frac{\eta P_{ess, rat} r_t \Delta t}{E_{ess, cap}} \right], \\ U_t &= \min \left[S_{SOC, max}, S_{SOC, ref} + \frac{P_{ess, rat} r_t \Delta t}{\eta E_{ess, cap}} \right], \\ L_t &\leq S_{SOC}(t+1) \leq U_t. \end{aligned} \quad (18)$$

The original action is retained when it satisfies the power limits in Equation (14) and the next-SOC condition in Equation (18); otherwise, it is minimally corrected subject to both. At $t = T$, $r_t = 0$ and Equation (18) reduces to Equation (17). A candidate is infeasible if these conditions cannot be jointly satisfied.

The PCC and PV-curtailment constraints are

$$\begin{cases} 0 \leq P_{grid, buy}(t) \leq \bar{P}_{grid, buy}, \\ 0 \leq P_{grid, sell}(t) \leq \bar{P}_{grid, sell}, \\ 0 \leq P_{pv, cur}(t) \leq P_{pv, act}(t). \end{cases} \quad (19)$$

Together with Equations (2)–(4), these constraints enforce base-load priority, prevent simultaneous grid import and export, and maintain minute-level power balance. The effects of finite PCC limits are evaluated in the stress cases.

3. Day-Ahead Construction of a Diverse Strategy Library via Dual-Space Coordinated Selection

3.1. Temporally Correlated PV Scenario Generation and Screening

The ability of the day-ahead strategy library to accommodate renewable-energy uncertainty depends on how well the retained PV scenarios represent temporal variation. Because PV forecast errors at adjacent time instants are generally correlated, independently sampled errors may underrepresent continuous ramps and boundary changes. A Gaussian Copula model is therefore used to describe the temporal dependence of PV forecast errors, followed by trajectory-based scenario screening [23].

For a scheduling horizon containing T time steps, the PV forecast-error vector is defined as follows:

$$\boldsymbol{\varepsilon}_{pv} = [\varepsilon_{pv}(1), \varepsilon_{pv}(2), \dots, \varepsilon_{pv}(T)]^T. \quad (20)$$

Its temporal correlation matrix is defined using an exponential-decay model:

$$[\mathbf{R}_{corr}]_{i,j} = \exp\left(-\frac{|i-j|}{\tau_{corr}}\right), \quad (21)$$

where τ_{corr} controls the decay rate of temporal dependence.

A Cholesky decomposition is applied:

$$\mathbf{R}_{corr} = \mathbf{L}\mathbf{L}^T. \quad (22)$$

For a standard Gaussian vector $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the correlated latent vector is defined as follows:

$$\mathbf{Z} = \mathbf{L}\mathbf{v}. \quad (23)$$

Each component is mapped to the unit interval by the standard normal cumulative distribution function:

$$u_{pv,t} = \Phi(Z_t) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{Z_t}{\sqrt{2}} \right) \right]. \quad (24)$$

For scenario n , the sampled probability $u_{pv,t}^{(n)}$ is mapped through the inverse marginal distribution of the forecast error at time t . The resulting PV output is clipped to its physical range:

$$P_{pv,t}^{(n)} = \min \left\{ P_{pv, \text{rat}}, \max \left[0, P_{pv, \text{for}}(t) + F_{\varepsilon,t}^{-1} \left(u_{pv,t}^{(n)} \right) \right] \right\}, \quad (25)$$

where $P_{pv, \text{for}}(t)$ is the day-ahead PV forecast, $F_{\varepsilon,t}^{-1}(\cdot)$ is the inverse marginal distribution of the forecast error, and $P_{pv, \text{rat}}$ is the rated PV power.

Directly evaluating a large Monte Carlo scenario set would substantially increase the computational burden of the offline optimization. The scenario set is therefore screened using the Euclidean distance between full-day PV trajectories as the trajectory-distance measure [24]:

$$d(s_i, s_j) = \left[\sum_{t=1}^T (P_{pv,i}(t) - P_{pv,j}(t))^2 \right]^{1/2}. \quad (26)$$

The scenario closest to the mean trajectory is selected first. Thereafter, the k th retained scenario is chosen according to the following max–min criterion:

$$s^{(k)} = \arg \max_{s \in S_{\text{unsel}}^{(k-1)}} \min_{\hat{s} \in S_{\text{sel}}^{(k-1)}} d(s, \hat{s}). \quad (27)$$

This procedure preferentially retains trajectories that are separated from the current selected set, thereby reducing the number of scenarios while maintaining geometric diversity among the retained PV trajectories. Figure 2 compares the PV power scenarios generated by independent sampling with the screened temporally correlated PV scenario set.




Figure 2. Comparison of (a) PV power scenarios generated by independent sampling and (b) the screened temporally correlated PV scenario set.

3.2. Dual-Space Coordinated Selection for Additive Library Construction

Conventional SPEA2 produces the base archive $\mathcal{L}_0 = \{(x_i, f_i)\}_{i=1}^{K_0}$. Strategy library C retains all rows of $\mathcal{L}_0$ in their original order and adds complementary trajectories through two successive augmentation steps.

For a candidate with objective vector $f = [F_1, F_2, F_3]$, the composite score used for this comparison is

$$G(f) = F_1 + \lambda_{\text{flex}} E_{\text{flex}} F_2 + \lambda_{\text{curt}} E_{\text{pv}} F_3, \quad (28)$$

where E_{flex} and E_{pv} are the forecast flexible-load and PV energies used in the corresponding offline evaluation. The scalar $G(f)$ is used only as a consistent offline score for candidate comparison, admission, and validation; it does not replace the three-objective formulation in Equation (9) or represent the realized multi-scenario operating cost. Objective cells are constructed by normalizing each objective using its minimum and range in $\mathcal{L}_0$. The grid resolution is determined by the base-archive size and the number of objectives. Candidates outside the reference objective ranges are not clipped to the boundary cells. A candidate retained from the same SPEA2 search history is eligible for addition only when it lies in the same objective cell as a non-locked base representative and satisfies $G_{\text{cand}} \leq G_{\text{base}} + \delta_{\text{num}}(G_{\text{base}})$, where $\delta_{\text{num}}(G_{\text{base}}) = \varepsilon_{\text{num}} \max\{|G_{\text{base}}|, 1\}$ and $\varepsilon_{\text{num}} = 10^{-9}$. The strategies minimizing the individual objectives, the minimum- G strategy, and the minimum- G representative in each occupied cell are locked and excluded as base representatives for this addition test.

For each addition-eligible trajectory, decision-space separation is evaluated using both action magnitude and trajectory shape:

$$d_{\text{dec}}(i, j) = \left[\frac{\|x_i - x_j\|_2}{2P_{\text{ess, rat}} \sqrt{N_b}} \cdot \frac{1 - \rho(x_i, x_j)}{2} \right]^{1/2}, \quad (29)$$

where N_b is the number of hourly ESS-power blocks in each strategy, and $\rho(x_i, x_j)$ is the correlation coefficient between the two block vectors. Both factors are clipped to $[0, 1]$. In the search-history augmentation, an eligible candidate is appended only if the mean decision-space separation increases.

The search-history augmentation is followed by an augmentation using model-consistent trajectories generated by linear programming (LP) for training PV scenarios and prescribed load multipliers. All previously retained strategies are preserved. Candidate generation uses the same power-balance, ESS, load-service, and in-horizon terminal-SOC constraints as the main framework.

Each training context $q \in \{1, \dots, Q\}$ combines one training PV scenario with a prescribed multiplier applied to both the base- and flexible-load forecasts. Let $G_q(x)$ denote the offline composite score of strategy x under context q . The retained library $\mathcal{L}$ is initialized with the complete output of the search-history augmentation, and its best score in context q is

$$b_q(\mathcal{L}) = \min_{x \in \mathcal{L}} G_q(x). \quad (30)$$

The marginal objective contribution of candidate y is

$$\Delta(y | \mathcal{L}) = \sum_{q=1}^Q \max\{b_q(\mathcal{L}) - G_q(y), 0\}. \quad (31)$$

In this subsequent augmentation, an admissible candidate must strictly improve at least one training context relative to the current retained library.

To define the scenario-response distance, let $\mathcal{R}$ contain the library at the start of this augmentation and the candidates passing the initial context-improvement screen. Let $G_q^{\min}$ and $G_q^{\max}$ denote the minimum and maximum scores over $\mathcal{R}$ in context q . The normalized response is

$$r_q(\mathbf{x}) = \frac{G_q(\mathbf{x}) - G_q^{\min}}{s_q}, \quad s_q = \begin{cases} G_q^{\max} - G_q^{\min}, & G_q^{\max} - G_q^{\min} > 10^{-16}, \\ 1, & \text{otherwise.} \end{cases}$$

The reference set and normalization scales remain fixed throughout this augmentation. With $\mathbf{r}(\mathbf{x}) = [r_1(\mathbf{x}), \dots, r_Q(\mathbf{x})]^T$, the scenario-response distance is

$$d_{\text{resp}}(\mathbf{x}, \mathbf{y}) = \frac{\|\mathbf{r}(\mathbf{x}) - \mathbf{r}(\mathbf{y})\|_2}{\sqrt{Q}}.$$

At each addition, candidates whose marginal contribution is at least 98% of the best available contribution form the contribution shortlist. Within this shortlist, the candidate maximizing the geometric mean of its minimum decision-trajectory distance and minimum scenario-response distance to the current retained library is selected.

The resulting library is therefore

$$\mathcal{L}_C = \mathcal{L}_0 \cup \mathcal{Q}_+, \quad \mathcal{L}_0 \subseteq \mathcal{L}_C, \quad K_0 \leq K_C \leq K_{\max}, \quad (32)$$

where $\mathcal{Q}_+$ contains the candidates admitted in either of the two augmentation steps. The actual K_C is variable, and $K_{\max}$ is the computational cap. Figure 3 illustrates the resulting ESS trajectory coverage.




Figure 3. ESS power-trajectory coverage in the representative library: (a) conventional base library C0; (b) strategy library C obtained through dual-space coordinated selection, with C0 retained as an exact subset.

3.3. Strategy Library Generation, Delivery, and Validity Period

At the end of the offline optimization and dual-space coordinated selection, the retained solutions form the day-ahead strategy library:

$$\mathcal{L} = \{(x_k, f_k)\}_{k=1}^K, \quad (33)$$

where $x_k = [x_{k,1}, \dots, x_{k,N_b}]^T$ contains the $N_b = 24$ hourly ESS-power blocks of a daily strategy, and f_k contains its multi-scenario objective values. Each block provides the raw ESS-power value for its corresponding hour and is expanded to the one-minute candidate sequence used in Section 4. Only solutions satisfying the physical constraints in Section 2.4 are retained. For C0, $K = K_0$; for C, K is the number admitted by Equation (32) and may differ between daily libraries.

The strategy library and the required day-ahead PV, load, and price profiles are delivered to the edge controller before intraday operation through a low-frequency communication process. During the predefined one-day validity period, the stored trajectories form the finite set of candidate ESS power sequences used for rolling matching. No continuous upper-level state acquisition or repeated centralized dispatch update is required, although current local PV, load, and SOC measurements remain necessary.

When forecasts cover multiple days, a separate strategy library can be constructed for each day and delivered in advance together with the corresponding PV, load, and price profiles. Each library is activated on its designated operating day and renewed when the forecasts, operating parameters, or physical limits used in its construction are updated. Thus, the proposed method reduces intraday communication dependence but does not remove the need for advance library delivery, local sensing, or periodic library renewal.

4. Intraday Edge-Autonomous Receding-Horizon Strategy Evaluation

The day-ahead strategy library supplies a finite set of candidate ESS power sequences to the intraday edge controller. At each control interval, the measured PV output, load demand, and SOC are used to initialize the rolling prediction, while the preloaded PV, load, and price profiles provide the inputs for subsequent prediction steps. Each stored sequence is propagated over the prediction horizon, corrected before state propagation to satisfy the dynamic physical limits, and evaluated using the cumulative cost. Only the first action of the selected sequence is applied.

The online decision variable is the discrete library index $k \in \{1, \dots, K\}$. Its feasible set and minimum-cost selection are defined in Equations (41) and (43), respectively.

4.1. Online Prediction and Strategy-Library Evaluation Model

At each control interval, a fixed prediction horizon of $N_p = 1440$ steps is used. The full-day horizon retains the intertemporal coupling among time-of-use prices, ESS energy shifting, flexible-load service, and the in-horizon terminal-SOC condition. In the evaluated daily-library implementation, the minute-level candidate trajectories and preloaded PV, load, and price profiles are periodically extended beyond the current operating day to provide a prediction-tail approximation. Physical execution on the next day uses that day's designated library. For the one-day index range $T = 1440$, the cyclic mapping is

$$\chi(t + \tau) = 1 + \text{mod}(t + \tau - 1, T), \quad \tau = 0, 1, \dots, N_p - 1. \quad (34)$$

The K candidate ESS power sequences available for rolling evaluation are collected in

$$\mathcal{U}(t) = \{u_1(t), u_2(t), \dots, u_K(t)\}. \quad (35)$$

For candidate k , the hourly block vector defined in Section 3.3 is expanded into the raw minute-level ESS-power profile:

$$P_{\text{ess},k}(j) = x_{k,b(j)}, \quad b(j) = 1 + \left\lfloor \frac{j-1}{T/N_b} \right\rfloor, \quad j = 1, \dots, T.$$

With $N_b = 24$ and $T = 1440$, each hourly value is repeated over 60 one-minute intervals. The raw candidate sequence extracted at time t is then defined as follows:

$$\mathbf{u}_k(t) = [P_{\text{ess},k}(\chi(t)), P_{\text{ess},k}(\chi(t+1)), \dots, P_{\text{ess},k}(\chi(t+N_p-1))]. \quad (36)$$

For notational convenience, $P_{\text{ess},k}(t+\tau|t)$ denotes the corresponding element of this cyclically indexed candidate sequence. The same index mapping is applied to the preloaded PV, load, and price profiles for future prediction steps. At $\tau = 0$, the preloaded PV and load values are replaced by the current local measurements.

At every prediction step, the raw action is projected onto the admissible ESS-power interval. The SOC- and PCC-dependent bounds are first calculated from the predicted SOC and the corresponding PV and load values using Equations (13) and (14). During a terminal-closing step, this interval is further restricted by the next-SOC condition in Equation (18). Denoting the resulting lower and upper bounds by $\underline{P}_k$ and $\bar{P}_k$, respectively, the corrected action is

$$\tilde{P}_{\text{ess},k}(t+\tau|t) = \min\{\bar{P}_k(t+\tau|t), \max[\underline{P}_k(t+\tau|t), P_{\text{ess},k}(t+\tau|t)]\}. \quad (37)$$

A candidate is infeasible if the resulting interval is empty. The corrected action, rather than the raw library value, is used consistently in the subsequent SOC propagation, local power allocation, and cost evaluation.

Using the measured SOC as the initial state, the predicted SOC of candidate k is updated as follows:

$$S_{\text{SOC},k}(t+\tau+1|t) = S_{\text{SOC},k}(t+\tau|t) - \mu_k(t+\tau|t) \frac{\tilde{P}_{\text{ess},k}(t+\tau|t)\Delta t}{E_{\text{ess},\text{cap}}}, \quad (38)$$

where $\tau = 0, 1, \dots, N_p - 1$. The charging/discharging efficiency factor is defined as follows:

$$\mu_k(t+\tau|t) = \begin{cases} 1/\eta, & \tilde{P}_{\text{ess},k}(t+\tau|t) \geq 0, \\ \eta, & \tilde{P}_{\text{ess},k}(t+\tau|t) < 0. \end{cases} \quad (39)$$

The sign convention is consistent with Section 2.4: positive ESS power denotes discharging and negative power denotes charging.

The cumulative prediction cost is evaluated using the same economic and penalty terms as the one-step matching cost:

$$\begin{aligned} J_k(t) = \sum_{\tau=0}^{N_p-1} & \left[c_{\text{buy}}(t+\tau)P_{\text{grid,buy},k}(t+\tau|t) - c_{\text{sell}}(t+\tau)P_{\text{grid,sell},k}(t+\tau|t) \right. \\ & + c_{\text{cyc}}|\tilde{P}_{\text{ess},k}(t+\tau|t)| + \lambda_{\text{flex}}P_{\text{unm},k}(t+\tau|t) \\ & \left. + \lambda_{\text{curt}}P_{\text{pv,cur},k}(t+\tau|t) \right] \Delta t. \end{aligned} \quad (40)$$

The grid-interaction, unserved-load, and PV-curtailment terms are calculated from $\tilde{P}_{\text{ess},k}$ using the local power-allocation rules in Section 2.2. For a prediction step with day-relative index $j = \chi(t+\tau)$, Equation (18) is applied when $T - N_{\text{cl}} + 1 \leq j \leq T$, with

$r_j = T - j$. At $j = T$, the next predicted SOC satisfies Equation (17). SOC is propagated continuously using Equation (38) and is not reset by the cyclic index. The corresponding ESS and grid actions are included directly in Equation (40).

4.2. Rolling Selection and Physical Constraints

Forecast errors may cause the realized state to differ from the day-ahead assumptions. After the corrected propagation in Equation (37), define the feasible candidate-index set as

$$\mathcal{K}_f(t) = \{k \in \{1, \dots, K\} : \text{Feas}_k(t) = 1\}. \quad (41)$$

Here, $\text{Feas}_k(t) = 1$ only if the corrected prediction satisfies the SOC and ESS-power limits, power balance, PCC/base-load conditions, and the in-horizon terminal-SOC condition. Only indices in $\mathcal{K}_f(t)$ participate in candidate ranking.

For $\mathcal{K}_f(t) \neq \emptyset$, the minimum candidate cost is

$$J_{\min}(t) = \min_{k \in \mathcal{K}_f(t)} J_k(t). \quad (42)$$

In exact arithmetic, the selected index satisfies

$$k^*(t) \in \arg \min_{k \in \mathcal{K}_f(t)} J_k(t). \quad (43)$$

In the implementation, candidates satisfying $0 \leq J_k(t) - J_{\min}(t) \leq \delta_{\text{tie}}$, with an absolute tolerance of $\delta_{\text{tie}} = 10^{-9}$ CNY, are treated as numerical ties. These ties are resolved lexicographically by greater predicted flexible-load supply, smaller ESS-action variation, retention of the previously selected strategy, and lower strategy index.

If $\mathcal{K}_f(t) = \emptyset$, a zero-ESS-power fallback is evaluated through the same physical allocation model. Its activation and any resulting constraint violations are recorded; full-horizon feasibility is not assumed for the fallback.

For a nonempty feasible set, only the first corrected action of the selected sequence is applied:

$$P_{\text{ess}}(t) = \tilde{P}_{\text{ess},k^*}(t|t). \quad (44)$$

After execution, the horizon is shifted forward by one interval, the initial state is refreshed using the latest local measurements, and the evaluation is repeated. The on-line layer does not solve a new continuous trajectory-optimization problem. Instead, it projects the stored candidate sequences onto the current feasible power limits, evaluates the resulting corrected sequences, and selects one of them for receding-horizon execution.

4.3. Feasibility and Computational Characteristics

The following feasibility argument uses the SOC dynamics in Equations (15) and (16) and assumes that the corrected ESS commands are executed as calculated. The online correction in Equation (37) enforces the state-dependent admissible ESS-power interval at every prediction step. Suppose that $S_{\text{SOC},k}(t + \tau|t) \in [S_{\text{SOC},\min}, S_{\text{SOC},\max}]$. For a nonnegative corrected action, Equations (13) and (37) give

$$\begin{aligned} S_{\text{SOC},k}^+ &= S_{\text{SOC},k} - \frac{\tilde{P}_{\text{ess},k} \Delta t}{\eta E_{\text{ess},\text{cap}}} \\ &\geq S_{\text{SOC},k} - \frac{(S_{\text{SOC},k} - S_{\text{SOC},\min}) E_{\text{ess},\text{cap}} \eta}{\Delta t} \frac{\Delta t}{\eta E_{\text{ess},\text{cap}}} = S_{\text{SOC},\min}. \end{aligned} \quad (45)$$

For a negative corrected action, the corresponding upper-bound relation is

$$\begin{aligned} S_{\text{SOC},k}^+ &= S_{\text{SOC},k} + \frac{\eta |\tilde{P}_{\text{ess},k}| \Delta t}{E_{\text{ess,cap}}} \\ &\leq S_{\text{SOC},k} + \frac{\eta \Delta t}{E_{\text{ess,cap}}} \frac{(S_{\text{SOC,max}} - S_{\text{SOC},k}) E_{\text{ess,cap}}}{\eta \Delta t} = S_{\text{SOC,max}}, \end{aligned} \quad (46)$$

Equations (45) and (46) show that a corrected discharging action preserves the lower SOC bound and a corrected charging action preserves the upper bound. The opposite bounds also hold because discharging cannot increase SOC and charging cannot decrease it. Starting from a feasible measured SOC, repeating this correction over the prediction and execution steps establishes SOC forward invariance under the modeled dynamics, provided that the corrected commands are executed as calculated and $\mathcal{K}_f(t)$ remains nonempty at every update.

During the final closing window, Equation (18) additionally restricts the next SOC to $[L_t, U_t]$ jointly with the power limits in Equation (14). At $t = T$, this interval reduces to the single value $S_{\text{SOC,ref}}$. Under the same feasibility and execution assumptions, the implemented sequence therefore satisfies $S_{\text{SOC}}(T+1) = S_{\text{SOC,ref}}$. This establishes conditional satisfaction of the terminal-SOC constraint; nonemptiness of the feasible candidate set at future updates remains an assumption rather than a result of the SOC-bound argument.

The additive library relation also provides a common-state cost ordering. At the same measured state and with the same preloaded PV, load, and price data, every C0 candidate has the same projected trajectory and cost when evaluated inside C. Hence $\mathcal{K}_{f,C0}(t) \subseteq \mathcal{K}_{f,C}(t)$ and, whenever the C0 feasible set is nonempty,

$$J_{\min,C}(t) = \min_{k \in \mathcal{K}_{f,C}(t)} J_k(t) \leq \min_{k \in \mathcal{K}_{f,C0}(t)} J_k(t) = J_{\min,C0}(t). \quad (47)$$

Equation (47) compares the minimum candidate costs evaluated from the same current state and with the same input data. It does not establish an ordering of the realized full-day closed-loop objectives, because different selected actions may lead to different subsequent SOC states.

For finite K and N_p , the online arithmetic complexity is $\mathcal{O}(KN_p)$ per control interval. The stored hourly ESS block array for one daily library requires $\mathcal{O}(KN_b)$ space, with $N_b = 24$. An explicitly expanded minute-level representation has size $\mathcal{O}(KT)$ and is distinct from the block-level library storage. Retaining all candidate prediction trajectories requires $\mathcal{O}(KN_p)$ workspace, whereas retaining one candidate prediction at a time requires $\mathcal{O}(N_p)$ workspace. For M independent microgrids, the serial computation is $\mathcal{O}(\sum_{i=1}^M K_i N_{p,i})$ and can be distributed among their edge controllers. Measured computation times for the tested K and N_p ranges are reported in Section 5.4.7.

4.4. Operation Under Communication-Constrained Conditions

The resulting execution process is illustrated in Figure 4. Before intraday operation, the upper-level controller delivers the strategy library through a low-frequency communication process. During the library validity period, the edge controller performs the following local loop:

1. Acquire the current local PV, base-load, flexible-load, and SOC measurements;
2. Initialize the first prediction step using the current measurements and obtain subsequent PV, load, and price values from the preloaded profiles;
3. Update candidate-state predictions, dynamic physical limits, and terminal-SOC reachability;

4. Evaluate the corrected candidate sequences over the prediction horizon;
5. Select the minimum-cost feasible candidate and apply its first corrected control action, or use the zero-ESS-power fallback if the feasible candidate set is empty; and
6. Repeat the procedure at the next control interval.




Figure 4. Flowchart of the intraday receding-horizon strategy-library evaluation and edge execution process.

Because the local loop uses the preloaded strategy library and forecast profiles together with current local measurements, it does not require continuous upper-level state acquisition or repeated centralized dispatch updates. The method therefore continues the predefined one-day local execution process during an upper-level communication outage, provided that the strategy library remains valid and the local sensing and actuation channels operate normally.

For a library of K candidates and a prediction horizon of N_p , the principal online computation consists of finite trajectory evaluation with complexity proportional to KN_p . The method reduces online optimization effort by moving the multi-objective evolutionary search to the day-ahead stage, rather than by removing local prediction, feasibility checks, or action selection.

5. Case Study

5.1. Case Study Setup and Evaluation Methodology

5.1.1. System Parameters and Input Data

The proposed strategy is evaluated in a grid-connected microgrid comprising PV generation, an ESS, base loads, flexible loads, and the utility grid. Each operating day comprises $T = 1440$ control intervals with a one-minute execution step. For Schemes C0 and C, the nominal candidate-evaluation horizon is $N_p = 1440$ steps, corresponding to 24 h. Daily performance indicators include only the realized operation within the corresponding day, while the seven-day indicators aggregate the consecutive daily results.

For the baseline comparison, surplus PV remaining after local load supply and ESS charging is exported to the grid. Separate PCC-limited and 20-times-PV/zero-export cases are defined in Section 5.4.6. The main system and algorithm parameters are listed in Table 2.

Table 2. System and algorithm parameters of the microgrid.

Parameter	Symbol or Setting	Value
Daily execution intervals	T	1440 steps (24 h)
Hourly ESS-power blocks per strategy	N_b	24
Execution time step	Δt	1 min
Rated PV power	$P_{pv, rat}$	23.4 kW
Rated ESS power	$P_{ess, rat}$	500 kW
Rated ESS capacity	$E_{ess, cap}$	500 kWh
Initial state of charge	$S_{SOC, 0}$	0.4
SOC range	$[S_{SOC, min}, S_{SOC, max}]$	[0.1, 0.9]
ESS charging and discharging efficiency	η	0.95
ESS cycling degradation coefficient	c_{cyc}	0.05 CNY/kWh
Fixed rolling candidate-evaluation horizon	N_p	1440 steps (24 h)
Unserved flexible-load penalty coefficient	λ_{flex}	0.9 CNY/kWh
PV-curtailment penalty coefficient	λ_{curt}	0.9 CNY/kWh
Time-of-use electricity purchase price	$c_{buy}(t)$	0.316465, 0.703256, and 1.054884 CNY/kWh
Electricity selling price	$c_{sell}(t)$	0.4 CNY/kWh
Grid-import power limit	—	Unlimited
Grid-export power limit	—	Unlimited (nominal case)
Terminal SOC reference and closing-window length	$S_{SOC, ref}, N_{cl}$	0.4 and 60 steps (60 min)
Online absolute tie tolerance	δ_{tie}	10^{-9} CNY
Adopted relative near-optimal threshold	ε	0
SPEA2 iterations and population	L, N_{pop}	200 and 200
Base archive size and computational cap	K_0, K_{max}	200 and 400

In the case study, λ_{flex} is set empirically to 0.9 CNY/kWh, between the flat- and peak-period purchase prices, to specify the cost-service trade-off. The same tariff-referenced monetary scale is assigned to λ_{curt} to penalize curtailed PV energy and represent the case-study preference for PV utilization.

For the seven day-specific libraries, the base archive contains $K_0 = 200$ strategies with three objectives, and the objective grid uses six divisions per axis. Each library uses 40 training PV scenarios combined with load multipliers of 0.90, 1.00, and 1.10, applied to both the base and flexible loads, yielding $Q = 120$ training contexts.

The source-side inputs include the day-ahead PV forecast and the actual intraday PV output. The day-ahead PV scenarios are used for offline strategy-library construction, while the preloaded PV forecast profile provides the future inputs for the rolling candidate evaluation described in Section 4.1. During intraday execution, the measured actual PV output replaces the forecast value at the current control interval, thereby allowing the controller response to forecasting errors to be evaluated. The day-ahead PV forecast and actual output are shown in Figure 5, together with samples from the PV forecast-error distributions modeled in Section 3.1.

The demand side consists of base loads and interruptible flexible loads. Base loads represent non-dispatchable demand that must be fully supplied. Flexible loads are modeled as interruptible demand whose supplied power may be reduced according to PV output, the ESS state, and time-of-use electricity prices, without an intertemporal energy-recovery requirement in the present case study. The forecasted and actual load profiles are shown in Figure 6.




Figure 5. Day-ahead forecasted and actual PV power with forecasting errors.




Figure 6. Forecasted and actual power profiles of base and flexible loads.

5.1.2. Reproducible Construction of the Seven Day Types

Six deterministic source–load transformations are applied to the baseline day to construct the additional reproducible day types. For a profile with N_x samples, define $\theta_{x,n} = 2\pi(n + 0.5)/N_x$. The day- d PV and load modulation functions are

$$\begin{aligned} m_{pv,d}(\theta) &= \max\{0.20, 1 + a_{pv,d} \sin(3\theta + \phi_{pv,d}) + 0.5a_{pv,d} \sin(7\theta + 0.3\phi_{pv,d})\}, \\ m_{load,d}(\theta) &= \max\{0.50, 1 + a_{load,d} \sin(2\theta + \phi_{load,d})\}. \end{aligned} \quad (48)$$

For each available baseline forecast, retained scenario, and actual trajectory x , the transformed profiles are

$$\begin{aligned} P_{pv,d}^x(n) &= \text{clip}\left[\alpha_{pv,d} m_{pv,d}(\theta_{x,n}) P_{pv,1}^x(n), 0, P_{pv, \text{rat}}\right], \\ P_{load,d}^x(n) &= \max\left[0, \alpha_{load,d} m_{load,d}(\theta_{x,n}) P_{load,1}^x(n)\right]. \end{aligned} \quad (49)$$

The same load transformation is applied to the base- and flexible-load forecasts. Forecasts, scenarios, and actual profiles are transformed synchronously so that each day retains the baseline forecast-error structure while changing energy level and intraday shape. Electricity prices are unchanged. Table 3 lists all fixed parameters; phases are in radians and Day 1 is unchanged.

These deterministic day-type transformations represent changes in clear-sky level, cloud-induced intraday shape, and load level. A separate uncertainty test fixes each realized trajectory and constructs the controller-visible forecast in the form $\hat{y}_\alpha = y_{\text{act}} - \alpha(e_0 + \xi)$, where $e_0 = y_{\text{act}} - \hat{y}_0$ and ξ is a zero-mean seeded direction whose root-mean-square (RMS) value equals 12% of the baseline PV-error RMS or 10% of the baseline load-error RMS. The amplification factors are $\alpha = 0.5, 1.0, 1.5, 2.0, 2.5$, with ten seeded trajectories at each level. To represent a prolonged rather than independently redrawn daily error,

an additional stress test concatenates all seven day types and continues one stationary first-order autoregressive (AR(1)) perturbation over $q = 1, \dots, 10,080$ min:

$$\rho_{\min} = r_d^{1/1440}, \quad \xi_q = \rho_{\min} \xi_{q-1} + \sigma \sqrt{1 - \rho_{\min}^2} z_q, \quad \hat{y}_q = y_q^{\text{act}} - 2.5(e_{0,q} + \xi_q). \quad (50)$$

Here $z_q \sim \mathcal{N}(0, 1)$, σ follows the same PV/load RMS rule above, and the theoretical one-day-lag persistence is set to $r_d = 0.90$ or 0.97 , corresponding to $\rho_{\min} = 0.99992684$ or 0.99997885 . Three seeds are fixed before execution for each setting. The perturbation is not reset at a day boundary; actual trajectories, objective weights, selectors, and the seven precomputed day-specific libraries are unchanged. The 20-times-PV/zero-export case in Section 5.4.6 separately represents an extreme surplus event.

Table 3. Fixed parameters used to construct the seven reproducible source–load day types.

Day Type	α_{pv}	α_{load}	a_{pv}	ϕ_{pv}	a_{load}	ϕ_{load}
Baseline	1.00	1.00	0.00	0.00	0.00	0.00
Clear/low load	0.97	0.90	0.04	0.35	0.03	0.60
Cloudy/normal load	0.78	1.00	0.15	1.10	0.02	1.40
Overcast/high load	0.52	1.12	0.10	2.00	0.04	2.20
Mixed sky/high load	0.88	1.08	0.18	2.70	0.03	0.20
Low PV/low load	0.62	0.92	0.12	3.50	0.04	3.30
Variable PV/normal load	0.82	1.02	0.22	4.40	0.03	4.50

5.1.3. Comparison Schemes and Extended Cases

Schemes A, B, C0, C, D, E, and F are considered to compare overall operating performance, strategy-library construction, communication dependence, and online computational burden. Their definitions are provided in Table 4. Except for the scheme-specific information availability, communication/update mechanism, and control or optimization procedure defined in Table 4, the comparison uses the same realized source–load conditions, system parameters, electricity prices, physical constraints, and objective weights.

Schemes C0 and C use the same intraday receding-horizon evaluation, physical constraints, in-horizon terminal-SOC condition, and strict composite-cost selector. Their only structural difference is the day-ahead library: C0 contains the conventional $K_0 = 200$ archive, whereas C retains C0 as an exact prefix and adds complementary trajectories through the two successive augmentation steps in Section 3.2. The comparison therefore evaluates the complete additive construction under a common online controller, rather than the isolated contribution of either augmentation step.

Scheme B is a 1 min perfect-information reference. Its optimization objective and minute-level physical execution differ by less than the numerical consistency tolerance in every tested day. The offline computation used to construct its reference trajectory is excluded from the count of online optimizations. For communication evaluation, the centralized execution layer is assumed to exchange operating-state information and control instructions once per minute, resulting in 1440 equivalent intraday communications. This communication count represents the assumed high-frequency centralized execution architecture rather than the offline procedure used to generate the reference result. Scheme D1 is the conventional deterministic online-MPC benchmark: unlike C, it solves a new finite-horizon LP at every minute from the current measured state and corrected forecast. D5–D60 use the same model but reduce update frequency and retain the latest valid instruction between updates. Robust MPC, reinforcement learning, and approximate dynamic programming are not included because controlled implementations would require additional uncertainty-set or training assumptions that are not shared by the present model. Their systematic comparison is left for future work.

Table 4. Control characteristics of the comparison schemes.

Scheme	Control Method	Optimization or Strategy-Library Setting	Intraday Communication and Computation Characteristics
A	Day-ahead static open-loop control	A full-day ESS schedule is generated from day-ahead forecasts and executed at 1 min intervals without intraday rematching or reoptimization.	No intraday upper-level communication or centralized online optimization.
B	Perfect-information centralized operating reference	A full-day 1 min mixed-integer linear programming (MILP) model uses actual PV and load data. Optimization and physical execution use the same temporal resolution, and objective–execution consistency is checked.	For communication evaluation, the assumed centralized execution architecture involves one upper-level state-and-command exchange per minute, resulting in 1440 intraday communications.
C0	Local evaluation with a conventional strategy library	A conventional SPEA2 archive with $K_0 = 200$ is followed by local receding-horizon candidate propagation and strict composite-cost selection.	No intraday dependence on upper-level communication or centralized online optimization.
C	Edge-autonomous control with a strategy library constructed by dual-space coordinated selection	The complete C0 library is retained exactly, and complementary trajectories are added according to objective contribution and dual-space separation. The actual K is variable and $K_{\max} = 400$ is only a cap.	The same local selector as C0 is used, with no intraday upper-level communication or centralized online optimization.
D	Deterministic online MPC and periodic-update variants	At each update instant, the upper-level controller solves a new deterministic finite-horizon LP from the measured state and updated forecast. D1 repeats this optimization every minute; the coarser variants retain the latest valid first-step instruction between updates.	Communication and online-optimization periods of 1, 5, 15, 30, and 60 min.
E	Event-triggered rolling optimization	Communication and control updates are triggered when source–load deviations exceed a prescribed threshold. Otherwise, the latest valid instruction is retained.	Event-trigger thresholds of 2%, 4%, and 6%.
F	Scenario-based stochastic MPC	A common ESS control sequence is optimized over five PV scenarios using a 15 min internal control block and 1 min rolling updates.	One upper-level update and centralized online optimization per minute.

Schemes C0 and C use the same rolling candidate-selection procedure defined in Section 4.2. Equation (43) gives priority to the strict minimum composite cost; flexible-load supply and ESS-action variation are consulted only for numerical ties.

In addition to the basic scheme comparison, extended cases are designed for statistical strategy-library validation, periodic communication, event-triggered communication, communication outages, and forecasting-error sensitivity. The settings are listed in Table 5.

The periodic communication cases examine the trade-off between communication frequency and operating performance. The event-triggered cases assess communication reduction through on-demand updates. The communication-outage cases evaluate continued operation when no new upper-level instruction is available. The forecasting-error cases examine the response of the strategy library and local matching mechanism to uncertainties in source and load forecasts. Parameter, stress, and profiling cases use prespecified settings and are reported separately from the primary comparison.

Table 5. Settings of the extended simulation cases.

Test	Case Setting
Paired strategy-library validation	Fifty independent common SPEA2 searches are performed. In every pair, C is formed by exact additive enhancement of the corresponding C0 archive; no C0 strategy is deleted.
Seven-day online comparison	Seven distinct constructed day types; C0 and C are rebuilt separately for each day from that day's forecasts and scenarios, using five common-search runs per day.
Periodic communication test	Communication periods of 1, 5, 15, 30, and 60 min.
Event-triggered communication test	Trigger thresholds of 2%, 4%, and 6%.
Communication-outage starting times	07:00, 12:00, and 18:00.
Communication-outage durations	3, 6, and 12 h, together with a full-day 24 h outage starting at 00:00.
Forecasting-error amplification factors	0.5, 1.0, 1.5, 2.0, and 2.5.
Random scenarios at each error level	10.
Principle of the forecasting-error test	Actual PV and load trajectories are fixed, while only the forecasts available to the controller are varied.
Persistent seven-day forecasting-error stress	One AR(1) perturbation is continued without reset for 10,080 min at $\alpha = 2.5$; theoretical one-day-lag persistence $r_d = 0.90$ and 0.97 , with three fixed seeds per setting.
Treatment of B in the forecasting-error test	The same perfect-information optimization result is used as a fixed reference and does not vary with forecasting error.
Online-parameter sensitivity	$N_p = 60, 240, 720, 1440$ min; $\varepsilon = 0, 1, 2, 3, 5\%$; $\lambda = 0.5, 0.9, 1.3$; and nested $K = 50, 100, 150, 200, K_{full}$.
Offline/system sensitivity and stress	10, 20, and 40 retained PV scenarios; ESS capacity factors 0.5, 1.0, and 1.5; three PCC-limit combinations; and a 20-times-PV zero-export case.
Computational profiling	$K = 100, 200, 300, 400$ at $N_p = 1440$, and $N_p = 60, 240, 720, 1440$ at $K = 237$, using an independent candidate pool for controlled scaling.

5.1.4. Evaluation Indicators and Statistical Methods

The economic operating cost, flexible-load unserved ratio, and PV-curtailment ratio are calculated according to the definitions in Section 2.3. The full-day composite performance indicator is represented by J_{day} , as defined in the methodological section. A smaller J_{day} indicates a more favorable overall balance among operating cost, flexible-load supply, and PV utilization.

The flexible-load satisfaction ratio is calculated from the flexible-load unserved ratio F_2 as follows:

$$\rho_{flex} = 1 - F_2. \quad (51)$$

All power quantities are evaluated at a 1 min resolution, and the corresponding energy is calculated as follows:

$$E = \sum_{t=1}^{1440} P(t)\Delta t, \quad \Delta t = \frac{1}{60} \text{ h}. \quad (52)$$

In addition to the composite objective, the schemes are evaluated in terms of load service, renewable-energy utilization, ESS operation, physical feasibility, communication

and computational burdens, strategy-library diversity, and robustness. The main indicators are summarized in Table 6.

Table 6. Main evaluation indicators for the comparison schemes.

Indicator Category	Main Indicators	Evaluation Purpose
Overall performance	Economic operating cost and composite objective value J_{day}	Evaluate economic performance and the overall multi-objective trade-off.
Load service	Base-load satisfaction ratio, flexible-load satisfaction ratio ρ_{flex} , and unserved flexible-load energy	Evaluate base-load supply security and flexible-load service.
Renewable-energy utilization	PV-curtailment energy, PV-curtailment ratio F_3 , and grid-export energy	Evaluate PV utilization and surplus-power export.
ESS operation	Minimum, maximum, and terminal SOC; equivalent cycle count; and total variation in ESS actions	Evaluate ESS utilization, operating safety, and action smoothness.
Physical feasibility	Minute-level power-balance residual, constraint-violation counts, candidate-set availability, and fallback activation when applicable	Verify compliance with system physical constraints.
Communication burden	Number of intraday upper-level communications, average communication interval, and maximum duration without upper-level communication	Measure dependence on the upper-level communication link.
Computational burden	Number of centralized online optimizations, mean and percentile computation times per step, and cumulative intraday computation time	Measure online computational requirements.
Strategy-library characteristics	Strategy-library size, number of strategies invoked, and number of strategy switches	Evaluate candidate coverage and intraday utilization.
Strategy-library diversity	Mean Euclidean distance and mean correlation distance	Evaluate differences in the magnitude and shape of ESS trajectories.
Robustness	Changes in the composite objective and constraint satisfaction under communication outages and forecasting errors	Evaluate continued operation under uncertain conditions.

To reduce the influence of random optimization outcomes, 50 independent common SPEA2 searches are performed. Each search produces C0 and its exact additive enhancement C, so pairing does not require two unrelated optimizer outcomes. The means and standard deviations of their holdout coverage scores are reported, together with the mean, standard deviation, and 95% confidence interval of the paired differences. Paired *t*-tests and Wilcoxon signed-rank tests are also reported.

The subsequent case-study analysis is organized as follows: strategy-library diversity and statistical validation, operating performance of the basic schemes, comparison of communication mechanisms, adaptation to communication outages, sensitivity to forecasting errors, parameter and system sensitivity, PCC stress tests, and computational scaling.

5.2. Offline Strategy Library Diversity and Statistical Validation

Before comparing intraday operation, the effect of the day-ahead library construction is examined separately. Every paired run uses one common SPEA2 search. C0 is the conventional archive and C is its exact additive enhancement; both are evaluated with the same holdout scenarios, objectives, and physical model.

5.2.1. Strategy-Library Diversity

In every paired run, the first 200 rows of C were numerically identical to C0 and retained in the same order. The seven-day experiment rebuilt a separate library for each day, yielding daily C libraries with 221–265 trajectories; the 50-run library statistics and day-specific libraries are therefore reported separately.

Across the 50 common-search runs, Table 7 shows that the mean library size increased from 200.00 to 237.38. The mean Euclidean and correlation distances increased by 21.83% and 43.88%, respectively, indicating broader magnitude and shape coverage without deleting base strategies. Figure 7 presents the corresponding distributions.

Table 7. Strategy-library characteristics over 50 paired independent runs.

Indicator	C0	C	Relative Change
Strategy-library size	200.00 ± 0.00	237.38 ± 18.74	+18.69%
Mean Euclidean distance	133.72 ± 27.19	162.91 ± 33.58	+21.83%
Mean correlation distance	0.0554 ± 0.0195	0.0797 ± 0.0258	+43.88%

Values are reported as mean ± standard deviation over 50 paired independent runs.




Figure 7. Distributions of strategy-library size and diversity indicators for Schemes C0 and C across 50 paired independent optimization runs: (a) Strategy-Library Size; (b) Mean Euclidean Distance; (c) Mean Correlation Distance.

These results establish the structural effect of the additive mechanism: C is larger and more diverse while preserving C0. They do not by themselves establish an online economic improvement, because the additional trajectories still have to be selected under realized minute-level conditions.

5.2.2. Paired Validation Performance

The 50 PV scenarios are partitioned by a seeded random permutation into 40 training scenarios and 10 validation scenarios. The validation score of a library is

$$J^{\text{val}}(\mathcal{L}) = \frac{1}{N_{\text{val}}} \sum_{s=1}^{N_{\text{val}}} \min_{x \in \mathcal{L}} G_s(x), \quad N_{\text{val}} = 10,$$

where $G_s(x)$ is the offline composite score under validation scenario s . The minimum library score is calculated separately for each scenario and then averaged to characterize scenario coverage.

For the r th paired run, the validation improvement of C relative to C0 is defined as follows:

$$\Delta J_r = J_{C0,r}^{\text{val}} - J_{C,r}^{\text{val}}, \quad (53)$$

where a positive value indicates a lower validation objective for C.

The paired holdout results are summarized in Table 8. Because C retains C0 as an exact subset, the minimum-score definition of J^{val} implies $J^{\text{val}}(\mathcal{L}_C) \leq J^{\text{val}}(\mathcal{L}_0)$ on the same holdout scenarios. The observed result is stronger than this non-deterioration property: all 50 paired runs show a strict reduction, with the mean score decreasing from 2006.7359 for C0 to 1997.7417 for C. The mean improvement is 8.9942 (0.4482%), with a 95% confidence interval of 7.2673–10.7210. Although the generated C libraries differ in size and composition across runs, their validation scores are identical to the reported precision. This is consistent

with the definition of J^{val} , which depends only on the scenario-wise minimum scores rather than on all retained candidates. Figure 8 presents the paired scores and improvements. These offline scores characterize library coverage and are not used to claim that online operation must improve.

Table 8. Paired validation results for the conventional and dual-space strategy libraries.

Statistic	C0	C	Paired Result, ΔJ_r
Mean $\pm$ standard deviation	2006.7359 $\pm$ 6.0762	1997.7417 $\pm$ 0.0000	8.9942 $\pm$ 6.0762
Median	2005.6711	1997.7417	7.9294
Range	2000.5070–2034.5950	1997.7417–1997.7417	2.7653–36.8533
Mean relative improvement	—	—	0.4482%
95% confidence interval of the mean improvement	—	—	7.2673–10.7210
Improved/tied/deteriorated runs	—	—	50/0/0
Paired t -test	—	—	$t = 10.467$, $p = 4.33 \times 10^{-14}$
Wilcoxon signed-rank test	—	—	$W = 0$, $p = 1.78 \times 10^{-15}$




Figure 8. Paired validation results for Schemes C0 and C across 50 independent optimization runs: (a) paired validation objective values; (b) paired improvement $\Delta J_r = J_{C0,r}^{\text{val}} - J_{C,r}^{\text{val}}$.

Both paired tests reject a zero difference at the 5% level, and every paired improvement is positive. The result supports improved coverage on the evaluated holdout set but does not constitute an offline guarantee of superior realized operation. The minute-level and seven-day results provide the operational test.

5.3. Operational Comparison of the Basic Schemes

This subsection compares the minute-level operation of Schemes A, B, C0, and C on Day 1. A and B use their Day-1 results, while C0 and C use the Day-1 libraries rebuilt for the seven-day experiment. Scheme B is the 1 min perfect-information centralized reference.

5.3.1. Power and SOC Trajectories

Figure 9 presents the hourly ESS net-power planning profiles, whereas the operating indicators are calculated from minute-level execution. During intraday operation, Scheme A follows a fixed day-ahead schedule, whereas C0 and C select candidate trajectories using local measurements and rolling evaluation. The Day-1 C library contains 237 strategies and records 25 strategy switches, compared with 200 strategies and 103 switches for C0. Thus, admitting additional candidates reduced rather than increased strategy switching in this case.

Figure 10 presents the hourly grid net-power planning profiles. The trajectory deviations in Table 9 are calculated from minute-level execution rather than from these hourly planning profiles. C has a larger trajectory root-mean-square error (RMSE) relative to B than A or C0, even though its full-day composite objective is close to that of B. This distinction is important: physically different ESS trajectories can realize similar total cost, so trajectory similarity is not used as a performance objective.




Figure 9. Hourly ESS net-power planning profiles of Schemes A, B, C0, and C.




Figure 10. Hourly grid net-power planning profiles of Schemes A, B, C0, and C.

Table 9. Trajectory deviations from the perfect-information centralized reference B.

Scheme	ESS-Power RMSE (kW)	Grid-Power RMSE (kW)	SOC RMSE
A	127.765	127.761	0.1066
C0	126.039	125.784	0.1060
C	149.154	149.154	0.1741

The SOC trajectories in Figure 11 remain within the prescribed range of 0.1–0.9 for all schemes. Scheme B exhibits a substantially larger total variation in ESS power than the other schemes, indicating frequent changes in the reference dispatch. It is therefore used mainly as a performance reference and not as a benchmark for action smoothness.




Figure 11. Minute-level SOC trajectories of Schemes A, B, C0, and C.

5.3.2. Load Service and Overall Operating Performance

Figure 12 compares supplied flexible-load power. All schemes fully satisfy the base load. C0 supplies more flexible load than the other three schemes, but its additional service is accompanied by higher economic cost. Since C0 and C use the same minimum-cost selector, their different cost–service outcomes reflect the different candidate libraries and resulting closed-loop trajectories.




Figure 12. Minute-level supplied flexible-load power under Schemes A, B, C0, and C.

The principal operating indicators are listed in Table 10. Relative to C0, C lowers the Day-1 composite objective by 7.266 and the economic cost by CNY 21.502, while flexible-load satisfaction decreases by 7.609 percentage points. The result reflects the weighted balance between economic cost and flexible-load service rather than simultaneous improvement of every component. Relative to B, C differs by only 0.142 in the composite objective and has the same flexible-load satisfaction ratio. Relative to A, C reduces the composite objective by 8.008 and the economic cost by 8.269 CNY.

Table 10. Operating performance of the basic comparison schemes.

Indicator	A	B	C0	C
Composite objective J_{day}	1938.303	1930.153	1937.561	1930.295
Economic operating cost (CNY)	1805.588	1797.178	1818.821	1797.319
Flexible-load satisfaction ratio (%)	29.06	28.92	36.53	28.92
Unserved flexible-load energy (kWh)	147.461	147.751	131.934	147.751
Grid-purchased energy (kWh)	3059.464	3053.945	3075.985	3054.157
Grid-exported energy (kWh)	1.180	0.000	0.107	0.212
ESS equivalent full cycles	1.681	1.602	1.721	1.602
Total variation in ESS power (kW)	1695.65	121,642.62	2208.80	1775.43
End-of-day SOC	0.4000	0.4000	0.4000	0.4000
Strategy-library size	—	—	200	237
Number of strategy switches	—	—	103	25

For all four schemes, base-load satisfaction is 100%, and no PV curtailment, SOC violation, or ESS-power violation is recorded. The maximum minute-level power-balance residual is below 9.60×10^{-14} kW. For every scheme, the grid exchange and ESS cycling associated with the in-horizon terminal-SOC correction are included in the reported economic cost and composite objective.

5.3.3. Seven-Day Comparison with Daily Library Rebuilding

To test whether the Day-1 result is exceptional, a distinct C0/C library is rebuilt for each of the seven deterministic day types generated by Equations (48) and (49) and Table 3. Each day uses its own forecasts and PV scenarios and five common-search runs with the same training-validation split. The representative pair is selected from the run with the lowest C0 validation score, and C is constructed from the corresponding common-search result. The actual minute-level trajectory is not used for this selection. Table 11 reports the unified full-day metric.

C is lower than C0 on all seven days (7/0/0 improved/tied/deteriorated). Its mean composite objective is 0.298% lower and its mean economic cost is 0.547% lower. The mean flexible-load satisfaction is 2.174 percentage points lower than C0 because C0 serves extra flexible demand at a larger economic cost on some days. C remains close to B, with a mean difference of 0.123. The Day-6 value is slightly lower than B; therefore, B is used as an operating reference rather than a certified mathematical lower bound. Every day passes the SOC-continuity, terminal-recovery, power-balance, and physical-limit checks, and both C0 and C begin the next day from the same recovered SOC of 0.4.

Table 11. Seven-day results with a separately rebuilt strategy library for each day.

Constructed Day Type	J_B	J_{C0}	J_C	$J_{C0} - J_C$	K_C
Day 1: baseline	1930.153	1937.561	1930.295	7.266	237
Day 2: clear/low load	1688.448	1696.807	1688.577	8.229	265
Day 3: cloudy/normal load	1958.015	1964.322	1958.429	5.893	225
Day 4: overcast/high load	2285.441	2286.882	2285.600	1.283	247
Day 5: mixed sky/high load	2146.374	2154.260	2146.374	7.887	246
Day 6: low PV/low load	1768.518	1773.227	1768.506	4.720	249
Day 7: variable PV/normal load	1996.788	2002.644	1996.814	5.830	221
Mean	1967.677	1973.672	1967.799	5.873	241.4

Across the representative and constructed-day cases, C reduces the complete objective relative to C0 while satisfying the same physical and terminal-SOC criteria.

5.4. Communication–Performance Trade-Off and Robustness Evaluation

5.4.1. Communication and Computational Burdens

Table 12 compares the representative C library with the centralized reference B, periodic cases D, event-triggered cases E, and stochastic MPC case F. For D, increasing the update interval from 1 to 60 min reduces communications and online optimizations from 1440 to 24 while increasing the composite objective from 1937.688 to 2012.917. The event-triggered results are non-monotonic: the 4% case outperforms the 2% case even though it communicates less, whereas the 6% case deteriorates markedly. Thus, the event threshold interacts with the timing of realized deviations and does not have a monotonic effect on performance.

Scheme C requires no intraday upper-level communication or centralized online optimization. In this fixed-library comparison, only B has a lower objective; C is lower than D1/D5/D15/D30/D60, all E cases, and F. Its measured Python implementation takes 0.149 s per control interval on average and 214.994 s cumulatively. The cumulative time exceeds that of D1 and F in the reported comparison, so the benefit is the removal of repeated upper-level updates and centralized solves rather than uniformly lower online computation time. The average evaluation time is well below the one-minute interval on the tested computer; performance on low-cost edge hardware has not been measured.

Table 12. Communication burden, computation time, and operating performance of the communication-related schemes.

Scheme	Intraday Upper-Level Communications	Centralized Online Optimizations	Longest No-Update Run (Steps)	Cumulative Computation Time (s)	J_{day}	Difference from B
B	1440	0	0	0.000	1930.153	0
C	0	0	1440	214.994	1930.295	0.142
D, 1 min	1440	1440	0	145.818	1937.688	7.535
D, 5 min	288	288	4	37.095	1938.131	7.978
D, 15 min	96	96	14	13.539	1951.554	21.401
D, 30 min	48	48	29	7.116	1981.760	51.606
D, 60 min	24	24	59	3.635	2012.917	82.764
E, 2%	500	500	301	44.900	2007.264	77.111
E, 4%	248	248	428	23.186	1939.007	8.854
E, 6%	72	72	674	5.376	2160.512	230.359
F	1440	1440	0	35.592	1955.017	24.863

For Scheme B, the single offline MILP solve is excluded from the online-optimization count and cumulative intraday computation time. Its 1440 communications represent assumed minute-level state-and-command exchanges. For C, the table reports local candidate-evaluation time, not centralized optimization time. The difference from B is $J_{\text{day}} - J_{\text{day},B}$. The longest no-update run counts consecutive one-minute control intervals without an upper-level update; intervals containing an update are excluded.

5.4.2. Operation Under Communication Outages

A communication outage is applied to D1 rather than to B, because D1 is an actual rolling controller that depends on minute-level upper-level state acquisition and instruction updates, whereas B is an offline perfect-information reference. During a blocked interval, D1 retains its most recently received instruction. Scheme C does not require an intraday upper-level link and therefore follows the same local process in every outage case.

Table 13 reports all effective outage windows. A full-day D1 outage increases the objective by 447.575. Outages beginning at 07:00 have a larger effect than late outages, while the 18:00 cases change the result very little because fewer active scheduling intervals remain. C maintains the same objective in every case. All listed cases complete terminal recovery and record no SOC or ESS-power violation.

Table 13. Representative communication-outage results.

Outage Case	$J_{\text{day},D1}$	Change from Normal D1	$J_{\text{day},C}$	Change from Normal C
Normal operation	1937.688	0	1930.295	0
07:00, 3 h	2158.444	220.756	1930.295	0
07:00, 6 h	2158.852	221.164	1930.295	0
07:00, 12 h	2349.391	411.703	1930.295	0
12:00, 3 h	2000.016	62.328	1930.295	0
12:00, 6 h	2097.017	159.329	1930.295	0
12:00, 12 h	2173.430	235.742	1930.295	0
18:00, 3 h	1937.688	0.0001	1930.295	0
18:00, 6 h	1937.867	0.179	1930.295	0
00:00, 24 h	2385.263	447.575	1930.295	0

The unchanged C results follow directly from its local execution design; this robustness statement assumes that local measurement and actuation remain available and that the preloaded daily library is still valid.

5.4.3. Sensitivity to Forecasting Errors

In the forecasting-error test, the actual PV and load trajectories remain fixed, while only the forecasts available to the fixed representative C0 and C libraries are perturbed. The perfect-information reference B is unchanged. As shown in Table 14, mean PV RMSE increases from 0.508 to 2.431 kW and mean load RMSE increases from 2.842 to 14.229 kW as the amplification factor increases from 0.5 to 2.5.

C has a lower composite objective than C0 in all 50 paired perturbations. The mean paired reduction is 7.356, with individual reductions ranging from 7.085 to 8.325. Across the five amplification levels, the mean objective of C ranges from 1930.153 to 1930.383, whereas that of C0 ranges from 1937.486 to 1937.977. The consistent paired reductions show that the additional candidates remain effective across the five tested forecast-error levels.

Table 14. Sensitivity of C0 and C to forecasting-error amplification.

Error Factor	PV RMSE (kW)	Load RMSE (kW)	$J_{\text{day},C0}$	$J_{\text{day},C}$	C–B Difference
0.5	0.508	2.842	1937.535 ± 0.044	1930.383 ± 0.040	0.230
1.0	1.014	5.695	1937.574 ± 0.015	1930.295 ± 0.001	0.142
1.5	1.506	8.546	1937.531 ± 0.020	1930.249 ± 0.080	0.096
2.0	1.973	11.388	1937.486 ± 0.042	1930.244 ± 0.085	0.091
2.5	2.431	14.229	1937.977 ± 0.397	1930.153 ± 0.000	0.000

Composite objective values are mean ± standard deviation over 10 random scenarios per error factor. The C–B difference uses $J_{\text{day},B} = 1930.153$.

All 100 C0/C executions pass the SOC, ESS-power, PCC/base-load, power-balance, and terminal-recovery checks.

5.4.4. Persistent Seven-Day Correlated Forecast Errors

The prolonged-error test applies Equation (50) at the largest examined amplification factor, $\alpha = 2.5$, while retaining the seven precomputed day-specific C0/C libraries. Each scheme inherits its own recovered SOC from the preceding day. Across the six tested sequences, the realized PV-error RMSE is 2.273–2.337 kW and the load-error RMSE is 14.468–14.704 kW. The largest AR(1)-perturbation change across any midnight boundary is only 0.0043 kW for PV and 0.0137 kW for load, confirming that the perturbation is continued rather than redrawn at midnight. Table 15 reports means over the three seeds at each theoretical one-day-lag persistence setting.

Table 15. C0/C results under one correlated forecasting-error process continued over seven days.

r_d	$\rho_{\min}$	$J_{7d,C0}$	$J_{7d,C}$	$J_{C0} - J_C$	Economic Saving (CNY)	ΔFlex (p.p.)	Feasible Sequences
0.90	0.99992684	13,817.512	13,775.634	41.878	54.974	−0.994	3/3
0.97	0.99997885	13,818.157	13,775.617	42.540	55.388	−0.975	3/3
Overall	—	13,817.835	13,775.626	42.209	55.181	−0.984	6/6

r_d is the theoretical one-day-lag correlation implied by the stationary minute-level AR(1) model. ΔFlex is C minus C0 in flexible-load satisfaction percentage points. Economic cost includes terminal-SOC recovery.

C has a lower seven-day composite objective in all six sequences and on all 42 paired scheme–days. Averaged over the six sequences, the reduction is 42.209 (0.3055%), and the economic saving is 55.181 CNY (0.4279%). Flexible-load satisfaction is 0.984 percentage points lower, so the result is a weighted-objective trade-off rather than simultaneous dominance of every component. Because B uses the unchanged actual trajectories and perfect information, its seven-day total remains $J_{7d,B} = 13,773.737$; mean C is 1.889 (0.0137%) above this reference. All 84 scheme–day executions have a nonempty feasible candidate set at every control interval and require no zero-ESS-power fallback. No SOC, ESS-power, PCC/base-load, or terminal-recovery violation is recorded, and the maximum power-balance residual is 5.68×10^{-14} kW.

5.4.5. Parameter Sensitivity

Table 16 separates the baseline setting from sensitivity cases. For the threshold sensitivity, candidates satisfying $J_k(t) \leq (1 + \varepsilon)J_{\min}(t)$ enter the secondary ranking, with $\varepsilon = 0$ adopted in the proposed controller. A nonzero threshold raises flexible-load satisfaction but also raises the complete objective; for example, $\varepsilon = 2\%$ changes J_{day} from 1930.295 to 1984.867. Increasing the horizon improves performance at additional computational cost. For the library-size sensitivity test, reduced libraries are constructed from the full C library by max–min distance sampling of the coordinate-wise normalized ESS-power trajectories. The trajectory farthest from the candidate-set centre is selected first, followed by the trajectory maximizing its minimum distance to the selected set. Different library sizes use the same deterministic selection order and therefore form nested subsets.

For this nested library-size test, the objective varies by 0.374 over $K = 50$ –237 and is unchanged to three decimals for $K = 150$, 200, and 237, indicating saturation in this case. The penalty coefficients are varied jointly as $\lambda_{\text{flex}} = \lambda_{\text{curt}} = \lambda$. The settings $\lambda = 0.5$ and 0.9 produce the same dispatch, economic cost, and flexible-load satisfaction. Increasing λ to 1.3 raises flexible-load satisfaction from 28.92% to 40.19%, at an additional economic cost of 30.068 CNY, illustrating the trade-off between operating cost and flexible-load service.

Table 16. Sensitivity of the fixed representative C library to online parameters.

Parameter	Setting	J_{day}	Economic Cost (CNY)	Flexible Satisfaction (%)	Cumulative Time (s)
λ	0.5	1871.195	1797.319	28.92	318.3
λ (baseline)	0.9	1930.295	1797.319	28.92	317.1
λ	1.3	1989.026	1827.387	40.19	317.0
ε (adopted)	0%	1930.295	1797.319	28.92	317.1
ε	1%	1957.542	1858.467	47.04	319.6
ε	2%	1984.867	1898.028	53.58	307.8
ε	3%	1988.204	1922.156	64.70	317.2
ε	5%	2024.517	1969.365	70.52	305.4
N_p	60 min	2078.865	1945.889	28.92	11.7
N_p	240 min	1960.551	1827.575	28.92	50.0
N_p	720 min	1933.097	1808.129	33.20	159.9
N_p (adopted)	1440 min	1930.295	1797.319	28.92	317.1
Nested K	50	1930.449	1797.473	28.92	291.1
Nested K	100	1930.669	1797.693	28.92	285.4
Nested K	150	1930.295	1797.319	28.92	257.8
Nested K	200	1930.295	1797.319	28.92	205.8
Full K	237	1930.295	1797.319	28.92	317.1

For Scheme C, the cumulative timings in Tables 12 and 16 sum the wall-clock durations of minute-level candidate evaluation and strategy selection. The two tables use the same timing definition but report separate execution batches; other simulation operations are excluded. Values of J across different λ rows are not directly comparable because the objective weighting changes. All other groups use the baseline coefficients.

5.4.6. Offline/System Sensitivity and PCC Stress Tests

Each sensitivity setting in Table 17 is evaluated with a separately reconstructed strategy library rather than by reusing the nominal representative library, including the reference-level settings that numerically coincide with the nominal parameters. Table 17 shows that all scenario-count and capacity cases complete without a zero-ESS-power fallback and recover to SOC 0.4. Larger ESS capacity lowers the objective in this case, while the retained-scenario results are not monotonic, reflecting the variability of the stochastic offline optimization.

The three PCC tests set the import/export limits to 492.63/1.00, 394.11/0.50, and 295.58/0.00 kW, respectively, and rebuild the strategy library for each pair. These cases use the common zero-ESS-power fallback for 307–452 min, with candidate-set availability of 68.61–78.68%, and complete terminal recovery. The 20-times-PV zero-export case results in 561.206 kWh (20.00%) of PV curtailment, candidate-set availability of 51.04%, 705 min of fallback, and a terminal SOC of 0.724. The recorded PCC/base-load violation counter remains zero in all four cases.

Table 17. Offline/system sensitivity and PCC/high-PV operating-envelope tests.

Test Group	Setting	J_{day}	Curtailment (kWh)	Candidate-Set Availability (%)	Zero-ESS Fallback (min)
Retained PV scenarios	10	1930.749	0	100.00	0
Retained PV scenarios	20	1930.174	0	100.00	0
Retained PV scenarios	40	1930.297	0	100.00	0
ESS capacity factor	0.5 (250 kWh)	2157.708	0	100.00	0
ESS capacity factor	1.0 (500 kWh)	1930.293	0	100.00	0
ESS capacity factor	1.5 (750 kWh)	1751.435	0	100.00	0
PCC import/export limits	492.63/1.00 kW	2195.106	0.294	78.68	307
PCC import/export limits	394.11/0.50 kW	2218.811	0	76.11	344
PCC import/export limits	295.58/0.00 kW	2231.906	0	68.61	452
High-PV stress	20× PV, zero export	—	561.206	51.04	705

Changed physical/input cases are not directly comparable through the absolute objective value. Candidate-set availability is the percentage of control intervals for which $K_f(t) \neq \emptyset$ before fallback. “—” denotes a stress case excluded from objective comparisons requiring the same terminal SOC.

5.4.7. Computational Scaling and Storage

The computational measurements were obtained on a computer equipped with an Intel Core i9-14900HX processor and 32 GB of RAM, running 64-bit Windows 11. The implementation uses Python 3.13.6, NumPy 2.5.1, and SciPy 1.18.0, with HiGHS accessed through SciPy for linear and mixed-integer optimization. Computation times are reported as wall-clock durations over the measurement intervals specified in the corresponding tables.

The principal online workload is $\mathcal{O}(KN_p)$ candidate-step evaluations. Table 18 shows a modest increase in measured time over the tested K range and approximately proportional growth with N_p ; these wall-clock measurements are distinct from the arithmetic complexity. Each strategy stores $N_b = 24$ hourly ESS-power values in double precision, requiring $8KN_b$ bytes for the raw ESS block array. This corresponds to 0.0366 MiB for C0 ($K = 200$) and 0.0434 MiB for C ($K = 237$), where 1 MiB = 2^{20} bytes. These values describe the ESS block arrays rather than the complete delivery package, which also contains the required profiles and metadata. These values show that storage is small, whereas full-horizon propagation dominates computation.

Table 18. Timing and peak Python allocation of the candidate-evaluation function.

K	N_p (min)	Candidate Steps	Mean Time (s)	95th Percentile (s)	Python Peak Allocation (MiB)
100	1440	144,000	0.1252	0.1270	0.122
200	1440	288,000	0.1302	0.1316	0.127
300	1440	432,000	0.1357	0.1374	0.159
400	1440	576,000	0.1371	0.1394	0.191
237	60	14,220	0.0056	0.0057	0.086
237	240	56,880	0.0218	0.0224	0.093
237	720	170,640	0.0665	0.0673	0.111
237	1440	341,280	0.1323	0.1348	0.138

Measurements use a fixed initial time and SOC, with two warm-up calls followed by five timed calls. Input preparation, outer strategy selection, and physical execution are excluded. The reported 95th percentile is calculated from the five timed calls. Peak allocation is measured using Python `tracemalloc` during the evaluation call and reported in MiB; it is not the total process working set. Profiling uses an independent candidate pool for controlled scaling.

The largest profiled configuration, $K = 400$ and $N_p = 1440$, has a 95th-percentile evaluation time of 0.1394 s, corresponding to 0.23% of the 60 s control interval. The ratio $60/0.1394 \approx 430$ describes the measured percentile margin on the tested platform, not a worst-case execution-time bound. This result is not presented as a low-cost-hardware benchmark: before deployment, the same complete-call test must be repeated on the target processor. If its worst-case time approaches the control interval, the explicit $K_{\max}$ cap, a shorter N_p , sequential/streaming evaluation, or parallel independent edge controllers provide direct ways to bound computation; network-coupled microgrids require an additional network model.

For the offline stage, multi-scenario objective evaluation contributes approximately $\mathcal{O}(N_{\text{eval}}N_sT)$ model-step evaluations, where N_{eval} is the number of candidate objective evaluations and N_s is the number of scenarios used in each evaluation. Pairwise objective-space calculations contribute $\mathcal{O}(L(N_{\text{pop}} + K_0)^2m)$ operations, where L is the number of SPEA2 iterations and $m = 3$. The stage-specific offline timing and process working-set measurements are summarized in Table 19.

The five independent SPEA2 searches and associated search-history augmentation sum to 12,617.404 s (3.505 h) of per-process wall-clock time; the independent runs can be parallelized. In the selected-seed SPEA2 and search-history augmentation measurement, the baseline and peak working sets were 136.316 and 690.656 MiB, respectively, a 554.340 MiB increment. Candidate generation and scoring for training-context augmentation, together with holdout validation, required 70.983 s in the repeated postprocessing measurement,

whereas input loading and day construction required 0.222 s. Thus, evolutionary search dominates offline wall-clock time, while measured memory remains approximately 691 MiB for this single-microgrid configuration. For M independent microgrids with the same K and N_p , serial online computation scales as $\mathcal{O}(MKN_p)$; network-coupled power-flow constraints require an additional model.

Table 19. Day-1 offline timing and process working set.

Offline Stage	Timing Basis	Wall-Clock Time (s)	Peak Working Set (MiB)
Gaussian-copula generation and max–min screening	Three repeated measurements	0.284	133.470
SPEA2 base-archive search	Five independent runs	2522.393 ± 15.707	—
Search-history augmentation	Five independent runs	1.088 ± 0.201	Included in complete run
Training-context augmentation and holdout validation	Three repeated measurements	70.983	127.979
Selected-seed SPEA2 search and search-history augmentation	One measurement	2441.746	690.656

Independent-run entries are per-process wall times. The selected-seed measurement covers SPEA2 and search-history augmentation and reproduces the corresponding stored arrays; training-context augmentation and holdout validation are measured separately. The selected-seed measurement is reported for reproducibility and peak memory rather than added to the five-run mean. The sum of the recorded per-process wall-clock durations describes the accumulated durations of the independent runs, not the elapsed completion time of a parallel batch. The 70.983 s entry covers the combined post-search construction and validation replay.

6. Discussion

The proposed library is an exact additive extension of the conventional library, $\mathcal{L}_C = \mathcal{L}_0 \cup \mathcal{Q}_+$. Every C0 candidate is retained in its original order and with identical decision values. Search-history augmentation is followed by training-context augmentation, with all previously retained strategies preserved at each step. The C0–C comparison therefore evaluates the complete additive construction under a common online controller. The validation score characterizes scenario coverage, whereas the closed-loop comparisons assess how the augmented candidate set is used during minute-level operation.

The adopted selector gives priority to the minimum predicted composite cost, with flexible-load supply and ESS-action variation used only to resolve numerical ties. Since unserved flexible-load energy is already penalized in this cost, a positive near-optimal threshold allows an additional service preference through secondary ranking. Both C0 and C use the same selector, so their comparison concerns the effect of the candidate libraries on the resulting cost–service trade-off. The in-horizon terminal-SOC constraint defines a common daily endpoint for ESS energy, with the associated grid exchange and cycling included in the same operating horizon. Scheme B serves as the perfect-information centralized reference defined in Section 5.1.3. The comparison with B concerns the composite objective, while ESS-power, grid-power, and SOC deviations describe differences in operating trajectories. These are distinct aspects of performance. Scheme C selects among precomputed trajectories using local measurements and preloaded forecasts; tracking the reference trajectory is not part of its selection objective.

The communication benefit follows from the offline–online separation. Once the daily library and required profiles have been delivered, C performs local rolling evaluation without intraday upper-level updates during the library validity period. Upper-level interruptions do not alter this local information path, provided that local sensing and actuation remain operational. The reduction concerns repeated upper-level exchanges and centralized solves, while local prediction, constraint correction, and strategy selection remain necessary. Block-level library storage scales as $\mathcal{O}(KN_b)$, whereas minute-level candidate

propagation requires $\mathcal{O}(KN_p)$ arithmetic work per control interval. The communication and computational indicators are therefore assessed separately.

The single-day forecast perturbations and persistent multi-day error sequences examine different temporal patterns of forecasting uncertainty. Forecast changes can alter the predicted ranking of stored candidates, while current local measurements refresh the operating state at each update. The library-size sensitivity test in Section 5.4.5 uses nested subsets sampled from the full C library to examine the effect of candidate-set size under the same online controller. Under restricted grid exchange, candidate-set availability and zero-ESS-power fallback duration describe how often the library supplies a feasible candidate. PV curtailment and terminal SOC relative to the common reference are reported alongside these indicators to characterize the operating conditions reached in the stress cases.

Several limitations remain. The seven profiles are constructed representative day types rather than a measured seasonal data set, and each day uses its own day-ahead library; the results therefore do not establish year-round performance or a single library valid across multiple days. The persistent-error evidence comprises six synthetic AR(1) sequences at the maximum examined amplification and two theoretical persistence levels; it is not a substitute for measured multi-season residuals or an exhaustive bound on error duration and distribution. Computational measurements are platform dependent. Deployment assessment should measure the complete local control cycle on the target processor, including input preparation, candidate evaluation, strategy selection, and command execution. The communication-outage model assumes that local sensing and actuation remain operational. The single-bus formulation represents aggregate PCC limits but not feeder-level power-flow or voltage constraints. Future work should use measured multi-season data and forecast residuals, adaptive library renewal, explicit sensor/actuator faults and packet delays, feeder-aware boundary modeling, and hardware-in-the-loop experiments.

7. Conclusions

This study proposed an edge-autonomous multi-objective energy management strategy for microgrids with limited upper-level communication. The method separates computationally intensive day-ahead optimization from minute-level local execution. Temporally correlated scenarios, a strategy library constructed through dual-space coordinated selection while retaining all C0 strategies, strict receding-horizon candidate selection, and an in-horizon terminal-SOC condition are integrated in a single offline–online framework.

The main conclusions are as follows:

1. The proposed construction preserves all C0 strategies exactly and adds complementary trajectories through successive search-history and training-context augmentation. Across 50 common-search runs, it increased mean library size by 18.69%, Euclidean trajectory dispersion by 21.83%, and correlation-distance dispersion by 43.88%. All 50 paired validation comparisons improved, with a mean validation-score reduction of 8.9942 and a 95% confidence interval of [7.2673, 10.7210].
2. In the representative case, C reduced the daily composite objective by 7.266 and economic operating cost by 21.502 CNY relative to C0; relative to the static day-ahead Scheme A, the corresponding reductions were CNY 8.008 and 8.269. C also achieved a lower daily composite objective than C0 in all seven constructed day types and all 42 paired daily comparisons under the six persistent correlated-error sequences, with the recorded physical checks satisfied and the end-of-day SOC restored to 0.4.
3. After daily library and profile delivery, C requires neither intraday upper-level communication nor repeated centralized optimization. With local sensing and actuation available, it remained feasible during the tested upper-level outages. In the representative comparison, C achieved a lower daily composite objective than all tested

periodic-update, event-triggered, and stochastic-MPC schemes. Its composite objective was only 0.0074% above that of the perfect-information centralized reference B.

These results support the proposed architecture for the examined daily operating and communication conditions. Future work will extend the evaluation using measured seasonal data and low-cost edge hardware, enrich strategy libraries near constrained operating boundaries, and investigate local-device faults through hardware-in-the-loop experiments.

Author Contributions: Conceptualization, H.D., Y.L. and H.N.; methodology, H.D. and Y.L.; software, H.D. and J.L.; validation, J.L., K.Z. and C.L.; formal analysis, H.D. and J.L.; investigation, H.D., J.L. and K.Z.; resources, C.L. and Y.L.; data curation, J.L. and K.Z.; writing—original draft preparation, H.D.; writing—review and editing, Y.L. and H.N.; visualization, H.D. and J.L.; supervision, Y.L. and H.N.; project administration, C.L. and Y.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Science and Technology Project of Former Collective Enterprises of State Grid Zhejiang Electric Power Co., Ltd., grant number 521148250004.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (GPT-5.6 Sol, OpenAI, last accessed in September 2026) for English-language editing and LaTeX formatting assistance. The authors reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: Authors Hanyu Dong, Jun Lai, Kaiyun Zhou and Changsheng Liu were employed by the company Huzhou Power Supply Company, State Grid Zhejiang Electric Power Co., Ltd. Authors Hanyu Dong and Changsheng Liu were employed by the company Zhejiang Talent Electric Group Co., Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The authors declare that this study received funding from State Grid Zhejiang Electric Power Co., Ltd. The funder was not involved in the study design, collection, analysis, interpretation of data, the writing of this article or the decision to submit it for publication.

Abbreviations

The following abbreviations are used in this manuscript:

CNY	Chinese yuan
ESS	Energy storage system
MPC	Model predictive control
PCC	Point of common coupling
LP	Linear programming
MILP	Mixed-integer linear programming
PV	Photovoltaic
RMSE	Root-mean-square error
SOC	State of charge
SPEA2	Strength Pareto Evolutionary Algorithm 2

References

1. Qu, K.; Si, G.Q.; Wang, Q.Y.; Xu, M.L.; Shan, Z.H. Improving economic operation of a microgrid through expert behaviors and prediction intervals. *Appl. Energy* **2025**, *383*, 125391. [[CrossRef](#)]
2. Dong, W.; Zhang, F.; Li, M.; Fang, X.L.; Yang, Q. Imitation learning based real-time decision-making of microgrid economic dispatch under multiple uncertainties. *J. Mod. Power Syst. Clean Energy* **2024**, *12*, 1183–1193.

3. Tian, J.; Shangguan, X.C.; Zhang, C.K.; He, Y.; Yao, W. Adjustable gain and reference signal sharing-based distributed secondary control for voltage restoration and current equalization in DC microgrid with time delay. *IEEE Trans. Ind. Electron.* **2026**, *73*, 7841–7852. [\[CrossRef\]](#)
4. Lian, Z.J.; Wen, C.Y.; Guo, F.H.; Lin, P.; Wu, Q. Decentralized secondary control for frequency restoration and power allocation in islanded AC microgrids. *Int. J. Electr. Power Energy Syst.* **2023**, *148*, 108927. [\[CrossRef\]](#)
5. Su, H.Y.; Lai, J.G. Distributed robust iterative learning control for AC microgrids with external disturbances. *Sci. China Inf. Sci.* **2025**, *68*, 202201. [\[CrossRef\]](#)
6. Li, L.; Luo, Z.X.; Tian, P.; Zhou, K.; Shen, S.; Liu, Y.; Zhang, C. A distributed decentralized cooperative control scheme for microgrids. *IET Gener. Transm. Distrib.* **2023**, *17*, 1059–1069. [\[CrossRef\]](#)
7. Duan, Q.S.; Tang, Z.; Ding, D. Distributed adaptive optimal secondary control for AC islanded microgrid under multiple event-triggered mechanisms. *ISA Trans.* **2025**, *163*, 65–75. [\[CrossRef\]](#)
8. Shahabadi, M.Z.; Arianfar, H.; Gharehpetian, G.B.; Yazdani, A. Distributed self-triggered control of islanded microgrids under switching communication topologies. *IEEE Trans. Ind. Inform.* **2024**, *20*, 13904–13915. [\[CrossRef\]](#)
9. Zhao, G.L.; Jin, L.Q.; Wang, Y.F. Distributed event-triggered secondary control for islanded microgrids with disturbances: A hybrid systems approach. *IEEE Trans. Power Syst.* **2023**, *38*, 1420–1431. [\[CrossRef\]](#)
10. Etivand, K.; Kalantar, M. Dynamic event-triggered fixed-time distributed secondary control for DC microgrids under control signal uncertainty. *Int. J. Electr. Power Energy Syst.* **2026**, *174*, 111449. [\[CrossRef\]](#)
11. Calpinici, A.; Irmak, E.; Kabalci, E. Design and implementation of an energy management system with event-triggered distributed secondary control in DC microgrids. *Energies* **2024**, *17*, 662. [\[CrossRef\]](#)
12. Babayomi, O.O.; Igboanusi, I.S.; Ahakonye, L.A.C.; Kim, D.S. Integrated blockchain and federated learning for the cybersecurity of distributed energy resources. *Int. J. Electr. Power Energy Syst.* **2025**, *173*, 111286. [\[CrossRef\]](#)
13. Mohan, A.M.; Meskin, N. Observer-based false data injection attack resilient event-triggered control of microgrid load frequency control system. *ISA Trans.* **2025**, *164*, 46–60. [\[CrossRef\]](#)
14. Yu, Y.; Liu, G.P.; Hu, W.S. Coordinated distributed predictive control for voltage regulation of DC microgrids with communication delays and data loss. *IEEE Trans. Smart Grid* **2023**, *14*, 1708–1722. [\[CrossRef\]](#)
15. Dai, X.R.; Liu, G.P.; Hu, W.S.; Deng, Q. Consensus-based secondary control for DC microgrids with communication delays via a networked predictive control strategy. *ISA Trans.* **2024**, *148*, 387–396. [\[CrossRef\]](#)
16. Yan, C.F.; Li, Q.R.; Zhang, X.M.; Han, Y.; Yang, P.; Zalhaf, A.S. A distributed coordinated control strategy for isolated AC microgrids based on consensus algorithm considering communication delay. *Int. J. Circuit Theory Appl.* **2025**, *53*, 2318–2339. [\[CrossRef\]](#)
17. Dai, X.R.; Liu, G.P.; Hu, W.S.; Deng, Q.; Lei, Z. Distributed predictive control for networked DC microgrids with communication delays and packet dropouts. *IEEE Trans. Ind. Inform.* **2025**, *21*, 2294–2303. [\[CrossRef\]](#)
18. Mallick, A.; Pradhan, J.K.; Pradhan, R.; Achary, K.S.K.; Rath, S.; Guha, A. A gain scheduling control framework for mitigation of time-varying network latency in autonomous AC microgrids. *IEEE Trans. Ind. Inform.* **2025**, *21*, 347–356. [\[CrossRef\]](#)
19. Wang, J.J.; Yao, L.Z.; Liang, J.; Wang, J.; Cheng, F. Distributed optimization strategy for networked microgrids based on network partitioning. *Appl. Energy* **2025**, *378*, 124834. [\[CrossRef\]](#)
20. Lee, C.D.; Kwon, H.; Lee, Y.I. OPC UA-based three-layer architecture for aggregated microgrids integrating edge cloud computing and IEC 62264. *J. Ind. Inf. Integr.* **2025**, *48*, 100965. [\[CrossRef\]](#)
21. Perez, R.; Rivera, M.; Araya, B.; Gomez, J.S.; Salgueiro, Y.; Restrepo, C.; Wheeler, P.; You, M.; Sumner, M. Enhanced microgrid reliability through software-defined networking and extended horizon predictive control. *Sustain. Energy Grids Netw.* **2025**, *42*, 101635. [\[CrossRef\]](#)
22. Kim, J.H.; Chang, J.W.; Kim, Y.S. A novel frequency reference compensation method in decentralized secondary frequency control for parallel operation of multiple generators in an islanded microgrid. *IEEE Access* **2024**, *12*, 127993–128006. [\[CrossRef\]](#)
23. Zhang, H.Y.; Zandehshahvar, R.; Tanneau, M.; Van Hentenryck, P. Weather-informed probabilistic forecasting and scenario generation in power systems. *Appl. Energy* **2025**, *384*, 125369. [\[CrossRef\]](#)
24. Bertsimas, D.; Mundru, N. Optimization-based scenario reduction for data-driven two-stage stochastic optimization. *Oper. Res.* **2023**, *71*, 1343–1361. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.