symmetry

MDPI

Article

Sensitivity-Guided BESS Siting and Sizing with Uncertainty-Aware Scheduling in Renewable-Rich Distribution Networks

Jun Ma *, Jishen Peng, Haotong Han, Liye Song and Hao Liu

Faculty of Electrical and Control Engineering, Liaoning Technical University, Huludao 125105, China; pengjishen@163.com (J.P.); hht4work@163.com (H.H.); songliye@lntu.edu.cn (L.S.); l1405110210@163.com (H.L.)
* Correspondence: thisismajunsms@163.com

Abstract

High penetrations of wind and photovoltaic generation create simultaneous challenges for battery energy storage system (BESS) planning in distribution networks, including differences in nodal regulation value, power-energy configuration, and day-ahead operation under forecast uncertainty. This study develops a sequential planning-to-operation workflow comprising candidate-bus generation, siting and sizing within the candidate set, and finite-scenario day-ahead scheduling for a fixed configuration. First, nodal net-injection sensitivities, Jacobian-assisted pre-screening, and deterministic topology/support safeguards are used to generate the main candidate set, and alternating-current (AC) finite-difference refinement is performed only for the sensitivity-led fast set; in the IEEE-33 system, this refinement reduces the number of AC power-flow calls from 65 for full-node analysis to 17. Next, Sensitivity-Guided Envelope-Based Nonanticipative Adjustable Recourse Optimal Power Flow (SG-ENAR-OPF) is solved separately for each bus in the main candidate set; the BESS location and power/energy capacities are jointly determined subject to the P-Q LinDistFlow model, BESS duration constraints, a shared affine response, and finite-scenario constraints. The full-node audit serves only as an independent paper-level validation benchmark and is not part of the deployable workflow. After the configuration is fixed, interval forecasts for load, photovoltaic (PV) output, and wind-turbine (WT) output at $q^{05}/q^{50}/q^{95}$ are used to construct 25 static load-renewable disturbance points and seven temporal stress paths, over which a shared finite-scenario day-ahead policy is optimized. The IEEE-33 MAIN case selects Bus 30, with BESS capacities of approximately 10.66 MW/10.66 MWh. Using scaled public time-series data, the final policy is replayed over 46 consecutive 24 h execution windows, comprising 1104 h actual trajectories; under the 0.002 MW/MWh storage-engineering criterion, all 46/46 windows pass storage engineering validation, and energy continuity is maintained across all 45/45 interday boundaries. Further nonlinear AC post-validation converges at all 1104/1104 operating points, of which 1058/1104 satisfy the complete voltage and branch-capacity constraints. Supplementary results for IEEE-69 show that the workflow can be executed on a second radial test feeder. However, the conclusions are strictly limited to the tested feeders, finite scenarios, scaled public-data settings, and stated engineering tolerances and do not constitute a formal guarantee over a continuous uncertainty domain or of general cross-system applicability.

Keywords: battery energy storage system; siting and sizing; sensitivity analysis; robust optimal power flow; scenario-based day-ahead scheduling; interval forecasting; scenario-based uncertainty

Academic Editor: Vasilis K. Oikonomou

Received: 29 July 2026
Revised: 1 September 2026
Accepted: 3 September 2026
Published: 4 September 2026

1. Introduction

As energy systems continue to decarbonize and distributed energy resources develop rapidly, the penetration of wind and photovoltaic generation in distribution networks is increasing [1]. The superposition of renewable-generation fluctuations and load variations is shifting distribution networks from conventional unidirectional, deterministic operation toward active operation with bidirectional power flows and multiperiod uncertainty, potentially causing voltage violations, branch congestion, increased network losses, and renewable-energy curtailment [2]. Against this background, enhancing the ability of distribution networks to regulate renewable fluctuations and forecast deviations has become important for maintaining secure operation and improving renewable-energy utilization [3].

A battery energy storage system (BESS) provides both rapid power regulation and intertemporal energy-shifting capabilities and can be used for peak shaving and valley filling, voltage support, power-flow regulation, and improved renewable-energy utilization [4]. However, the effectiveness of storage depends not only on its capacity but also on its connection location, network structure, and operating strategy [5]. The connection location determines the range over which the BESS affects nodal voltages and branch flows, the power capacity governs instantaneous regulation capability, and the energy capacity determines the duration of sustained support. Meanwhile, day-ahead scheduling must address forecast errors in load, photovoltaic generation, and wind generation [6]. Consequently, storage planning and operation are not independent; together, they involve bus screening, siting and sizing, and scheduling under uncertainty.

With high renewable penetration, radial topology, electrical distance, and the locations of loads and renewable resources cause the same power adjustment to produce markedly different voltage, network-loss, and branch-flow responses at different buses [7]. This network locational heterogeneity is addressed through nodal sensitivity analysis and the subsequent optimal power flow (OPF) model. Day-ahead uncertainty, by contrast, is represented using system-level load and renewable forecasts: finite scenarios are constructed in a two-dimensional standardized load-renewable disturbance plane, and temporal stress paths describe the cumulative effects of different disturbance magnitudes, durations, and occurrence sequences on the state of charge (SOC) [8]. Thus, this study treats network locational heterogeneity and the temporal evolution of disturbances separately rather than grouping them under physical spatial uncertainty.

1.1. Literature Review

Existing studies have first focused on energy-storage siting and sizing. Mathematical programming, intelligent optimization, and hybrid algorithms have been used to jointly determine storage locations and capacities while considering investment cost, electricity-purchase cost, network losses, voltage deviation, and renewable-energy curtailment [9]. Das et al. [10] optimized the locations and capacities of distributed storage in the IEEE-33-bus distribution network with high wind and photovoltaic penetration, effectively reducing voltage deviations, line loading, and power losses. Wichitkrailat et al. [11] used the crayfish optimization algorithm to determine the locations and capacities of multiple BESS units, thereby improving voltage levels and reducing peak load and system losses. However, including all candidate buses, capacity variables, and multiscenario constraints simultaneously can substantially increase problem size and solution time.

To improve the siting-and-sizing process, Li et al. [12] developed a bilevel joint planning model for distributed generation and storage, in which the upper level determines locations and capacities and the lower level optimizes the storage operating strategy, thereby reducing planning deviations caused by renewable-energy uncertainty. Mumtahina et al. [13] adopted a two-stage method: a storage configuration is first obtained through mixed-integer linear programming (MILP), after which the mountain-gazelle optimization algorithm evaluates its voltage-regulation and loss-reduction performance. Liu et al. [14] partitioned clusters according to network structure, load characteristics, and electrical coupling and combined hybrid particle swarm optimization with power-flow calculations to coordinate photovoltaic generation and storage. Al-Mahammedi et al. [15] further compared several load models and metaheuristic algorithms, confirming that appropriate storage configuration improves voltage stability and reduces active-power losses. These studies improved the configuration process through hierarchical planning, hybrid solution methods, and network partitioning, but did not clearly distinguish the functions of physically informed candidate-bus screening and formal siting-and-sizing optimization.

Robust optimization and flexibility constraints have been applied to power-system operation and energy-storage planning to address renewable-energy uncertainty and system flexibility. In a representative study of adaptive robust optimization, Bertsimas et al. [16] represented nodal net-injection deviations using a deterministic uncertainty set, balancing system security and operating economics without requiring a probability distribution in advance. Jing et al. [17] introduced system-flexibility constraints into storage siting and sizing while considering storage investment, network losses, electricity purchases, photovoltaic curtailment, and scheduling costs. Such methods improve the ability of a solution to accommodate uncertainty, but their effectiveness still depends on the construction of uncertainty sets or operating scenarios; a small number of representative scenarios may not adequately reflect combinations of multisource disturbances and their sustained effects [18].

Probabilistic and interval forecasts provide predictive distributions, quantiles, and lower and upper bounds, thereby supplying a data basis for constructing day-ahead uncertainty scenarios [19]. Telle et al. [20] developed a probabilistic net-load forecasting framework for distributed integrated renewable-energy systems; Dumas et al. [21] used normalizing flows to generate multistep probabilistic scenarios for wind power, photovoltaic generation, and load; Ma et al. [22] coordinated photovoltaic inverters and storage according to photovoltaic forecast-error scenarios to provide voltage control under uncertainty. Van der Meer et al. [23] noted that no single probabilistic forecasting model is optimal under all application conditions.

Recent studies have further advanced BESS planning and uncertainty-aware operation from several perspectives. Li et al. [24] combined voltage-sensitivity-based clustering and partitioning with multilevel optimization to improve the siting and sizing of distributed storage. Xu et al. [25] developed a robust configuration model for distributed modular storage using a correlation-aware polyhedral uncertainty set and a coordinated voltage-control design. From an operational perspective, Zhao et al. [26] proposed a multistage robust BESS scheduling method based on uncertainty-set decomposition and adaptive rolling decisions, whereas Yu et al. [27] explicitly modeled the dynamic spatiotemporal correlation of renewable-energy uncertainty within a distributionally robust scheduling framework. Hosseini et al. [28] further considered BESS siting, sizing, multistage operation, and network reconfiguration in a unified optimization model. These studies indicate that sensitivity-based planning, robust uncertainty modeling, adaptive scheduling, and planning-operation coordination are themselves relatively mature research directions. BESS planning can also be classified as centralized or distributed according to configuration form and as short-duration or longer-duration according to rated duration. The former emphasizes unified siting and sizing and system-level operation, whereas the latter places greater emphasis on multibus coordination or intertemporal energy support. This study focuses on centralized single-BESS planning and jointly considers power and energy capacities through explicit duration constraints.

In addition, for distribution systems with high penetrations of inverter-based resources (IBRs), frequency-domain stability caused by converter-network interaction warrants attention alongside steady-state voltage deviation, network losses, and branch-flow metrics. Xu et al. [29] proposed the frequency-spectra impedance margin ratio (fIMR), which constructs a unified stability-margin metric from impedance spectra. It can evaluate the relative stability margins of different point of interconnection (POI) locations and identify potentially weak locations, providing another quantitative tool for impedance-stability analysis in systems with high IBR penetration. This frequency-domain impedance criterion addresses a different physical dimension from the steady-state voltage/loss/branch-flow sensitivity metrics used in this study and is therefore complementary. The present study remains limited to BESS candidate-bus screening, siting and sizing, and steady-state network feasibility in day-ahead operation; it does not explicitly incorporate the impedance-based small-signal stability margin into the proposed formulation. Further coupling frequency-domain stability metrics such as fIMR with candidate-bus evaluation and the planning model is a worthwhile direction for future research.

Substantial progress has been made separately in sensitivity analysis, BESS planning, scheduling under uncertainty, and operational validation. This study further examines how these stages connect within a unified workflow. Specifically, sensitivity analysis is used only to reduce the candidate-bus range; the final BESS location and capacity are still determined by the OPF model, and a full-node audit tests whether candidate screening retains the preferred buses. During day-ahead operation, system-level load, photovoltaic (PV), and wind-turbine (WT) interval forecasts are used to construct a two-dimensional standardized disturbance space in which finite combination scenarios are generated. Temporal stress paths then describe the cumulative effects of disturbance duration and occurrence sequence on SOC. This disturbance space is only a mathematical coordinate representation for the scheduling policy; it does not correspond to node-level or multisite geographic forecast errors. The resulting affine scheduling policy is frozen for execution on actual trajectories: BESS active power is not re-optimized, and subsequent network audits evaluate operating feasibility and performance.

1.2. Aim of the Work

For centralized single-BESS planning and day-ahead operation under uncertainty, this study develops a sequential planning-to-operation workflow comprising candidate-bus generation, siting and sizing within the candidate set, finite-scenario day-ahead scheduling for a fixed configuration, and actual fixed-policy validation. The contribution is not the isolated introduction of a new sensitivity analysis, affine policy, or finite-scenario technique, but the explicit assignment of decision responsibilities to each stage and the connection of candidate reduction, unified planning, shared-policy scheduling, and post hoc validation of a frozen policy into an auditable sequential process. The main contributions are as follows.

1. A sensitivity-guided candidate-bus generation method with clearly defined responsibilities is proposed. Multimetric nodal sensitivities identify system-level regulation buses and buses that support the target weak area, while Jacobian-assisted pre-screening reduces the number of buses requiring nodal alternating-current (AC) finite-difference refinement. Structurally safeguarded candidate buses are then added using measurable indicators, including feeder depth, downstream load, target-area support, and the relief of branch/renewable-integration stress. All candidate-generation rules are fixed before OPF is solved; sensitivity analysis performs only candidate prioritization and candidate-space reduction, and does not directly determine the final BESS installation location.

2. A candidate-set BESS siting-and-sizing procedure with technical duration constraints is established. For each retained candidate bus, Sensitivity-Guided Envelope-Based

Nonanticipative Adjustable Recourse Optimal Power Flow (SG-ENAR-OPF) is solved, and power-energy duration constraints restrict the BESS power and energy capacities to a predefined technically feasible domain. The final BESS location and capacities are determined by the unified candidate-wise optimization results rather than directly by sensitivity rankings. The full-node audit serves only as a paper-level benchmark, used to determine whether candidate-space reduction omits a preferred location found by the complete bus search; it is not a necessary step in the deployable workflow.

3. An interval-forecast-driven finite-scenario shared-policy scheduling method is proposed. Load, PV, and WT q05/q50/q95 interval forecasts are mapped into a two-dimensional standardized load-renewable disturbance space, from which 25 static disturbance points are generated. Seven temporal stress paths are further introduced to represent the effects of disturbance duration, occurrence sequence, and switching pattern on SOC accumulation. All modeled scenarios share the same day-ahead baseline plan, affine-response coefficients, and operating-mode variables, forming a finite-scenario nonanticipative shared policy. The robustness claim is strictly limited to the modeled finite scenario set and is not extended to a formal feasibility guarantee over a continuous uncertainty domain.

4. A validation procedure combining actual trajectory and nonlinear AC validation for a frozen day-ahead policy is established. Once set, the day-ahead policy remains fixed; actual load, PV, and WT trajectories are used only for post hoc execution. BESS active power is not re-optimized, and the final charging/discharging commands are not subjected to clipping, projection, or corrective re-optimization. Storage-engineering residual checks and a nonlinear AC power-flow audit then independently provide execution evidence for the fixed policy on the tested actual trajectories. The IEEE-33 system is used for the primary rolling-window validation; the IEEE-69 system serves as a second radial test feeder for supplementary examination of execution performance under changes in network size, topology, and renewable-resource locations, without claiming general cross-system applicability.

2. Materials and Methods

2.1. Problem Analysis and Overall Framework

To balance candidate-bus screening efficiency, storage-configuration precision, and the validation of fixed-policy day-ahead operation, the overall workflow is divided into three sequential stages: candidate-bus generation, siting and sizing within the candidate set, and day-ahead scenario-based scheduling for a fixed configuration. The full-node audit is not part of the operational workflow; it serves only as an independent validation benchmark for evaluating the effectiveness of candidate-space reduction.

Figure 1 presents the overall workflow. Stage I uses nodal net-injection sensitivities and Jacobian-assisted pre-screening to identify buses with high regulation value and applies predefined candidate-generation rules to form the main candidate set. Stage II solves the Sensitivity-Guided Envelope-Based Nonanticipative Adjustable Recourse Optimal Power Flow (SG-ENAR-OPF) sequentially only for buses in the main candidate set and determines the BESS location and capacities using uniform evaluation rules. Stage III fixes the selected BESS configuration, constructs finite day-ahead scenarios from interval forecasts, and optimizes an affine scheduling policy, followed by post hoc validation using actual trajectories and the network model. To examine whether candidate-space reduction omits preferred buses, the same SG-ENAR-OPF is additionally solved for all eligible nonslack buses to conduct a one-time full-node audit. This audit neither participates in candidate-set construction nor constitutes a necessary step in practical application.

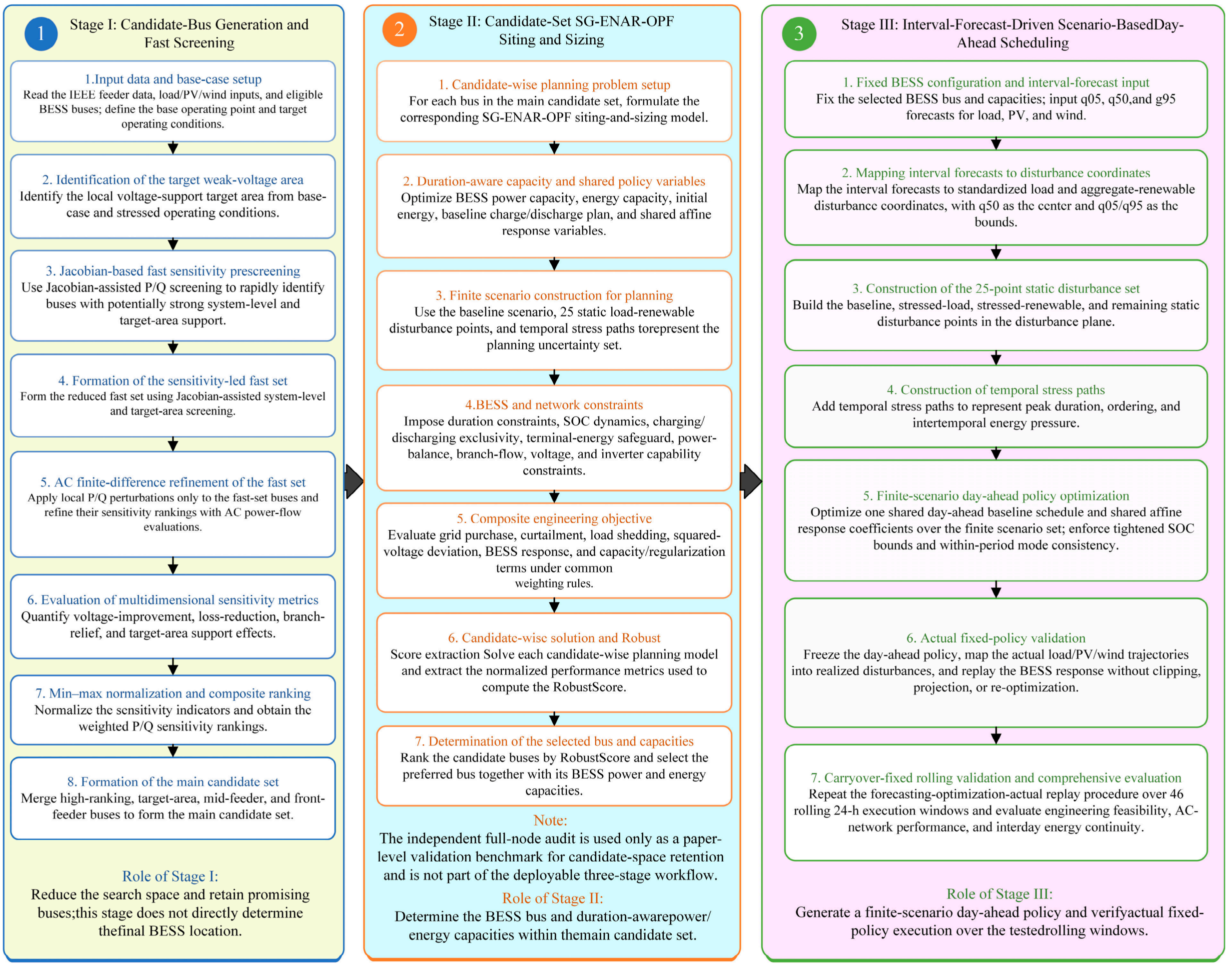

Figure 1. Overall framework of sensitivity-guided BESS siting and sizing and uncertainty-aware day-ahead scheduling.

2.2. *Sensitivity-Guided Generation of BESS Candidate Buses*

Energy storage, distributed photovoltaic generation, wind power, and inverter reactive-power support can all affect distribution-network operation by changing active or reactive net injections at buses. Although these flexibility resources differ in physical characteristics and operating constraints, their local regulation effects can be represented uniformly as changes in nodal net injections. Accordingly, nodal net-injection sensitivities are used to identify candidate buses with high regulation value, providing a reduced candidate space for subsequent BESS siting and sizing.

2.2.1. AC Finite-Difference Sensitivity Based on Nodal Net-Injection Perturbations

Let the set of distribution-network buses be $\mathcal{N} = \{1, 2, \cdots, N\}$, the set of branches be $\mathcal{E}$, and the set of slack buses be $\mathcal{N}_{slack}$. Define the set of all eligible nonslack installation buses as $\mathcal{N}_{elig} = \mathcal{N} \backslash \mathcal{N}_{slack}$, hereafter referred to as the eligible/installable bus set, and denote the target weak-voltage area by $\mathcal{N}_{tar} \subseteq \mathcal{N}$. The set $\mathcal{N}_{tar}$ can be determined from low-voltage buses under the base case or a typical heavy-load case, or it can be specified directly according to engineering mitigation needs. This area is used to evaluate local voltage-support capability, whereas the improvement in the minimum system-wide voltage is represented by a separate metric.

Let the AC power-flow solution process be represented by the nonlinear mapping $x = \phi(p, q)$, where p and q are the nodal net active- and reactive-power injection vectors, respectively, and x contains operating states such as bus voltages, branch flows, and system losses. The base operating point is $x^0 = \phi(p^0, q^0)$. For any $i \in \mathcal{N}_{elig}$, apply an active-power net-injection perturbation ΔP or a reactive-power net-injection perturbation ΔQ. For any operating metric y, its finite-difference sensitivity to active-power injection at bus i is:

$$\frac{\partial y}{\partial P_i} \approx \frac{y(p^0 + \Delta P e_i, q^0) - y(p^0, q^0)}{\Delta P} \tag{1}$$

and the corresponding reactive-power sensitivity is

$$\frac{\partial y}{\partial Q_i} \approx \frac{y(p^0, q^0 + \Delta Q e_i) - y(p^0, q^0)}{\Delta Q} \tag{2}$$

where e_i is the unit vector whose ith element equals 1. Let the number of buses in the eligible/installable bus set be $N_c = \left|\mathcal{N}_{elig}\right|$. A complete active- and reactive-power finite-difference analysis therefore requires $2N_c + 1$ power-flow calculations: 1 base power flow and $2N_c$ nodal-perturbation power flows.

Forward finite differences are used to calculate nodal sensitivities, balancing computational efficiency and ranking requirements during candidate-bus screening. Compared with central finite differences, forward finite differences require only one-sided perturbation power-flow calculations and therefore avoid the additional reverse-perturbation calculations. Because the sensitivity analysis compares the relative regulation capabilities of different buses rather than obtaining high-precision gradients for continuous optimization, this method is adequate for candidate generation. Section 3.3.3 further compares the ranking consistency of forward and central finite differences at different perturbation step sizes to examine the effects of the finite-difference form and perturbation magnitude.

To evaluate nodal regulation value comprehensively, four metrics are constructed: improvement in the minimum voltage, reduction in system active-power losses, relief of the maximum branch flow, and voltage support in the target area. Taking an active-power perturbation as an example, the four metrics for candidate bus i are defined as follows:

$$a^P_{1,i} = \frac{\partial V_{\min}}{\partial P_i} \tag{3}$$

$$a^P_{2,i} = -\frac{\partial P_{loss}}{\partial P_i} \tag{4}$$

$$a^P_{3,i} = -\frac{\partial S_{\max}}{\partial P_i} \tag{5}$$

$$a^P_{4,i} = \frac{1}{|\mathcal{N}_{tar}|} \sum_{j \in \mathcal{N}_{tar}} \frac{\partial V_j}{\partial P_i} \tag{6}$$

where $V_{\min}$ is the minimum system-wide voltage, P_{loss} is the system active-power loss, $S_{\max}$ is the maximum branch apparent power, and $\mathcal{N}_{tar}$ is the target weak-voltage area. Because lower network losses and maximum branch flows are preferred, the signs of the corresponding sensitivities are reversed so that a larger value indicates greater nodal regulation value for every metric. The reactive-power metrics $a^Q_{1,i} = \frac{\partial V_{\min}}{\partial Q_i}$, $a^Q_{2,i} = -\frac{\partial P_{loss}}{\partial Q_i}$, $a^Q_{3,i} = -\frac{\partial S_{\max}}{\partial Q_i}$, and $a^Q_{4,i} = \frac{1}{|\mathcal{N}_{tar}|} \sum_{j \in \mathcal{N}_{tar}} \frac{\partial V_j}{\partial Q_i}$ are calculated analogously. To eliminate dimensional differences, each metric is normalized using min–max normalization:

$$\hat{a}^P_{m,i} = \frac{a^P_{m,i} - \min_{j \in \mathcal{N}_{elig}} a^P_{m,j}}{\max[\max_{j \in \mathcal{N}_{elig}} a^P_{m,j} - \min_{j \in \mathcal{N}_{elig}} a^P_{m,j}, \varepsilon_{norm}]} \tag{7}$$

$$\hat{a}^{Q}_{m,i} = \frac{a^{Q}_{m,i} - \min\limits_{j\in\mathcal{N}_{elig}} a^{Q}_{m,j}}{\max[\max\limits_{j\in\mathcal{N}_{elig}} a^{Q}_{m,j} - \min\limits_{j\in\mathcal{N}_{elig}} a^{Q}_{m,j}, \varepsilon_{norm}]} \tag{8}$$

where $m \in \{1,2,3,4\}$ is the sensitivity-metric index, $j \in \mathcal{N}_{elig}$ denotes the eligible/installable buses used to determine the normalization extrema, $\hat{a}^{P}_{m,i}$ and $\hat{a}^{Q}_{m,i}$ are the mth normalized active- and reactive-power sensitivity metrics, respectively, and $\varepsilon_{norm} = 10^{-12}$ is a numerical floor that prevents degeneration of the normalization denominator. If the max–min range of a metric over the candidate buses is zero, all of its normalized values are set to zero. If the range is smaller than ε_{norm}, ε_{norm} is used as the denominator floor, preventing differences near numerical precision from being amplified into substantial ranking differences.

Because $V_{\min} = \min\limits_{j\in\mathcal{N}}|V_j|$ and $S_{\max} = \max\limits_{l\in\mathcal{E}}|S_l|$ are determined by the minimum system-wide voltage and the maximum branch apparent power, respectively, the corresponding limiting bus or branch may change with the operating point and perturbation. Accordingly, the finite-difference calculation records both the extrema and the limiting buses/branches under the base and perturbed conditions. If the limiting element changes, the case is marked as an extremum switch and the corresponding result is treated as a potentially nonsmooth response. If extrema are tied within numerical precision, the associated buses or branches are jointly treated as the limiting set.

The metrics $V_{\min}$ and $S_{\max}$ are retained rather than replaced by soft-min, soft-max, or other smooth aggregate surrogate because they directly represent the weakest voltage and the most heavily loaded branch and therefore have clear network-security interpretations. Section 3.3.3 further examines their stability across operating points, perturbation step sizes, and finite-difference forms.

On this basis, the composite active- and reactive-power scores of candidate bus i are defined as $Score^{P}_{i} = \sum\limits_{m=1}^{4} \omega_m \hat{a}^{P}_{m,i}$ and $Score^{Q}_{i} = \sum\limits_{m=1}^{4} \omega_m \hat{a}^{Q}_{m,i}$, respectively, where the weights satisfy $\omega_m \geq 0$ and $\sum\limits_{m=1}^{4} \omega_m = 1$ and represent the relative engineering preference between system-wide mitigation and support for the target weak area.

The default weight vector is [0.35, 0.25, 0.20, 0.20], corresponding to minimum-voltage improvement, network-loss reduction, branch-flow relief, and target-area voltage support, respectively. These weights represent the relative preferences among evaluation metrics during candidate screening; their effects are examined in Section 3.3.2 using multiple fixed and random weight vectors.

Nodal finite-difference sensitivities may also depend on the operating point and perturbation step size. The main screening setup uses the IEEE-33 base operating condition with $\Delta P = 0.01$ MW and $\Delta Q = 0.01$ MVAr. Section 3.3.3 further examines result stability across operating conditions, perturbation step sizes, and forward/central finite differences while recording switches in the limiting bus or branch associated with each extremum metric.

2.2.2. Jacobian-Assisted Two-Stage Fast Candidate Screening

To reduce repeated power-flow calculations for finite differences over all candidates, the AC power-flow equations are linearized to first order near the base operating point:

$$\Delta x = J^{-1}_{red} \Delta s \tag{9}$$

where J^{-1}_{red} is the reduced Jacobian matrix, Δs is the nodal power-perturbation vector, and Δx is the vector of voltage-angle and voltage-magnitude changes. The Jacobian describes changes in system states caused by nodal-injection perturbations rather than changes in evaluation metrics directly. The voltage-magnitude and angle states are then updated using Δx, and the effect of a perturbation on voltage, network losses, branch flows, and other metrics can be approximated through the first-order expansion of operating metric $I(x)$:

$$\Delta I \approx \frac{\partial I}{\partial x}\Delta x \tag{10}$$

For the voltage metric, $I(x)$ corresponds to the nodal-voltage state; for the network-loss metric, $I(x)$ corresponds to the total active-power loss calculated from nodal voltages and angles; and for the branch-flow metric, $I(x)$ corresponds to the maximum branch apparent power calculated from the power-flow equations.

Using this linear mapping, the system-level active-power candidate set $\mathcal{C}_{p,k}$, system-level reactive-power candidate set $\mathcal{C}_{Q,k}$, target-area active-power support candidate set $\mathcal{C}^{tar}_{P,k_t}$, and target-area reactive-power support candidate set $\mathcal{C}^{tar}_{Q,k_t}$ are constructed separately, and their union forms the fast pre-screening set:

$$\mathcal{N}_{fast} = \mathcal{C}_{p,k} \cup \mathcal{C}_{Q,k} \cup \mathcal{C}^{tar}_{P,k_t} \cup \mathcal{C}^{tar}_{Q,k_t} \tag{11}$$

where k is the truncation depth for the system-level P/Q fast composite scores and k_t is the independent truncation depth for the target-area P/Q support rankings. The IEEE-33 main workflow uses $k = 8$ and $k_t = 3$. The resulting fast set is:

$$\mathcal{N}_{fast} = \{18, 17, 33, 32, 31, 16, 15, 14\} \tag{12}$$

$k = 8$ is the minimum truncation depth that simultaneously retains the high-ranking system-level P/Q buses and the four types of key buses under the tested settings, whereas $k_t = 3$ retains the main support buses for the target weak area. $\mathcal{N}_{fast}$ determines only which buses require the costly nodal AC finite-difference evaluation and is not equivalent to the final $\mathcal{N}_{cand}$.

Active- and reactive-power AC finite-difference checks are performed for the buses in $\mathcal{N}_{fast}$. The resulting number of AC power-flow calls is $N^{FD}_{PF} = 1 + 2\left|\mathcal{N}_{fast}\right|$, where 1 denotes the shared base AC power flow. Relative to complete P/Q finite differences for all N_c candidate buses, the reduction in computational effort is:

$$\eta_{FD} = 1 - \frac{2\left|\mathcal{N}_{fast}\right| + 1}{2N_c + 1} \tag{13}$$

For the IEEE-33 system, a complete 32-bus analysis requires 65 AC power-flow calculations, whereas $k = 8$ requires only 17, reducing finite-difference power-flow calls by 73.85%. The Jacobian stage provides low-cost pre-screening, and $\mathcal{N}_{fast}$ determines only which buses require AC finite-difference checking; it is not equivalent to the final main candidate set. Neither stage directly determines the final BESS location. Section 2.2.3 further adds candidate buses using predefined topology and support indicators and then applies the subsequent SG-ENAR-OPF evaluation uniformly to all retained buses.

2.2.3. Construction of the Main Candidate Set

To reduce the risk that a single Top-K sensitivity cutoff omits buses at different network locations, the main candidate set combines buses with high system-level sensitivities, buses supporting the target weak area, mid-feeder buffer buses, and upstream buffer buses for the target weak area:

$$\mathcal{N}_{cand} = \mathcal{N}_{sys} \cup \mathcal{N}_{tar} \cup \mathcal{N}_{mid} \cup \mathcal{N}_{front} \tag{14}$$

The four bus categories are selected in a fixed order with mutually exclusive membership: once a bus enters an earlier category, it is not selected again in a later category. The corresponding truncation depths are denoted by K_{sys}, K_{tar}, K_{mid}, and K_{front}, respectively. For the IEEE-33 case, they are set to 5, 3, 5, and 2, respectively, and all are fixed before the subsequent OPF solution.

Here, $\mathcal{N}_{sys}$ retains the top K_{sys} buses in the system-level composite sensitivity rankings from Sections 2.2.1 and 2.2.2, whereas $\mathcal{N}_{tar}$ retains the top K_{tar} buses ranked by support capability for the target weak-voltage area. For candidate buses not yet selected, the following mid-feeder buffer score is further constructed:

$$G_i^{mid} = 0.25R_i^{depth} + 0.25R_i^{load} + 0.25R_i^{br} + 0.25R_i^{cross} \tag{15}$$

where R_i^{depth} represents the mid-feeder location characteristic, R_i^{load} the downstream load level, R_i^{br} the branch-relief capability, and R_i^{cross} the cross-area support capability; all are normalized in advance to [0, 1].

(1) Mid-feeder location score:

First, for bus i, the cumulative impedance $Z_i^{path} = \sum\limits_{e \in \mathcal{P}_i} \sqrt{r_e^2 + x_e^2}$ is calculated along the unique radial path from the slack bus, where $\mathcal{P}_i$ is the set of branches on the path from bus i to the slack bus. The feeder electrical depth $\delta_i^{depth} = \frac{Z_i^{path}}{\max\limits_{j \in \mathcal{N}} Z_j^{path}}$ is then obtained by normalization, on the basis of which the mid-feeder location function is defined:

$$\widetilde{R}_i^{depth} = exp\left[-(\frac{\delta_i^{depth}{}_i - 0.52}{0.18})^2 \right] \tag{16}$$

where $\delta_i^{depth} \in [0, 1]$ is the cumulative impedance along the unique radial path from the slack bus to bus i. Min–max normalization of $\widetilde{R}_i^{depth}$ yields the final $R_i^{depth} = Norm(\widetilde{R}_i^{depth})$, which describes the feeder-location characteristic of a bus. Buses with an electrical depth close to 0.52 receive higher scores, and the score decreases as a bus approaches either the feeder head or end.

(2) Downstream-load score:

Let $\mathcal{D}_i$ be the set of downstream buses rooted at bus i. First, the 24 h baseline active loads at all buses are summed to obtain $P_{L,j}^0 = \sum\limits_t P_{L,j,t}^0$. The downstream load covered by bus i is $L_i^{down} = \sum\limits_{j \in \mathcal{D}_i} P_{L,j}^0$, its share of the network-wide baseline load is $\widetilde{R}_i^{load} = \frac{L_i^{down}}{\sum\limits_{j \in \mathcal{N}} P_{L,j}^0}$, and normalization gives $R_i^{load} = Norm\left(\widetilde{R}_i^{load}\right)$.

(3) Upstream branch-stress relief score:

For each branch e, the maximum and mean absolute active-power flows under the scenarios are first calculated:

$$P_{e,\max}^{abs} = \max_{s,t} |P_{e,t}^s| \tag{17}$$

$$P_{e,ave}^{abs} = \frac{1}{|\Omega|} \sum_s \frac{1}{T} \sum_t |P_{e,t}^s| \tag{18}$$

The branch-stress indicator $\widetilde{B}_e = 0.55P_{e,\max}^{abs} + 0.45P_{e,avg}^{abs}$ is then constructed and min–max normalized to obtain the branch-stress score $B_e = Norm\left(\widetilde{B}_e\right)$. Because a power injection at bus i primarily affects the branches on its path to the slack bus, the branch-stress scores along this path are summed to obtain $\widetilde{R}_i^{br} = \sum\limits_{e \in P_i} B_e$, which is normalized as $R_i^{br} = Norm\left(\widetilde{R}_i^{br}\right)$. A larger value indicates that bus i is downstream of more highly stressed upstream branches and that regulating its injection has greater potential to relieve those branches.

(4) Cross-area support score for the target area:

Using LinDistFlow sensitivities, the mean voltage responses in the target weak-voltage area to active- and reactive-power injections at bus i are first calculated and normalized:

$$R_i^{P,tar} = Norm\left[\frac{1}{|\mathcal{N}_{tar}|}\sum_{j\in\mathcal{N}_{tar}} S_{V,P}(j,i)\right] \tag{19}$$

$$R_i^{Q,tar} = Norm\left[\frac{1}{|\mathcal{N}_{tar}|}\sum_{j\in\mathcal{N}_{tar}} S_{V,Q}(j,i)\right] \tag{20}$$

The cross-area support score in the mid-feeder buffer metric is defined as:

$$R_i^{cross} = Norm\left(0.50R_i^{P,tar} + 0.50R_i^{Q,tar}\right) \tag{21}$$

The four metrics describe feeder location, downstream load scale, upstream branch stress, and target-area support capability. Because they retain complementary network information of different types, they are combined with equal weights. After excluding buses already included in $\mathcal{N}_{sys}$ and $\mathcal{N}_{tar}$, buses are ranked in descending order of G_i^{mid} , and the top K_{mid} form $\mathcal{N}_{mid}$.

To further retain buses with a buffering role on the upstream radial paths of the target weak area, the union of all unique paths from target-area buses to the slack bus is first determined. The slack bus, target-area buses, and previously selected buses are excluded. The following upstream-buffer score is constructed for the remaining buses:

$$G_i^{front} = 0.40R_i^{pos} + 0.25R_i^{tar} + 0.25R_i^{ren} + 0.10R_i^{load} \tag{22}$$

(5) Target-area upstream-path location score:

Let the union of all unique paths from the target area to the slack bus be $P_{tar} = \bigcup_{j\in\mathcal{N}_{tar}} P_j$. For buses on these paths but outside the target area, the normalized electrical-depth interval $0.55 \le \delta_i^{depth} \le 0.93$ is used to construct the location score $\widetilde{R}_i^{pos} = \frac{\delta_i^{depth}-0.55}{0.93-0.55}$. For buses outside this interval, $\widetilde{R}_i^{pos} = 0$, and final normalization yields $R_i^{pos} = Norm(\widetilde{R}_i^{pos})$. Thus, within the predefined upstream-path interval, buses closer to the target area receive higher location scores.

(6) Composite target-area support score:

The composite target-area support score is $R_i^{tar} = 0.40R_i^{P,tar} + 0.60R_i^{Q,tar}$. Compared with R_i^{cross}, the target-area support evaluation for upstream-buffer buses places greater emphasis on reactive-power voltage support.

(7) Renewable-energy branch-stress relief score:

For high-renewable, RES, or low-net-load scenarios, the mean absolute active-power flow of each branch is calculated to form the renewable-energy branch stress $\widetilde{B}_e^{ren} = \sum_{s\in\Omega_{ren}} \frac{1}{T}\sum_t |P_{e,t}^s|$, which is normalized to obtain $B_e^{ren} = Norm(\widetilde{B}_e^{ren})$. The renewable-stress relief indicator of bus i is obtained by summing the relevant branch stresses along its upstream path to give $\widetilde{R}_i^{ren} = \sum_{e\in P_i} B_e^{ren}$, which is further normalized as $R_i^{ren} = Norm(\widetilde{R}_i^{ren})$. This indicator measures the potential of a bus to relieve upstream branch stress under high renewable output.

The upstream-buffer category is intended mainly to retain buses with a buffering role on the upstream paths of the target area; the location score is therefore assigned a relatively high weight. Target-area support and renewable-stress relief represent voltage-support and power-flow-relief capabilities, respectively, while downstream load scale is an auxiliary indicator. Buses are ranked in descending order of G_i^{front}, and the top K^{front} form $\mathcal{N}_{front}$. All candidate-generation metrics above are calculated from network topology, bus and branch parameters, baseline loads, predefined operating scenarios, and sensitivity information available before OPF solution. They do not use subsequent SG-ENAR-OPF results or full-node audit results.

For the IEEE-33 case, these rules produce, in sequence:

$$\mathcal{N}_{sys} = \{18, 17, 16, 15, 14\} \tag{23}$$

$$\mathcal{N}_{tar} = \{33, 32, 31\} \tag{24}$$

$$\mathcal{N}_{mid} = \{12, 30, 11, 13, 10\} \tag{25}$$

$$\mathcal{N}_{front} = \{29, 28\} \tag{26}$$

$$\mathcal{N}_{cand} = \{18, 17, 16, 15, 14, 33, 32, 31, 12, 30, 11, 13, 10, 29, 28\} \tag{27}$$

The first 8 buses are from the sensitivity-based pre-screening set in Section 2.2.2 and have undergone AC finite-difference checking; the remaining 7 buses are added by the topology and support rules above. All candidate buses are ultimately evaluated using the same SG-ENAR-OPF. Thus, screening and supplementation only reduce the search space and do not directly determine the final BESS installation location. The full-node audit serves as an independent paper-level validation benchmark, as detailed in Section 2.4.1.

2.3. Sensitivity-Guided Envelope-Based Nonanticipative Adjustable Recourse Optimal Power Flow (SG-ENAR-OPF) Siting and Sizing Model

This study uses the Sensitivity-Guided Envelope-Based Nonanticipative Adjustable Recourse Optimal Power Flow (SG-ENAR-OPF) for BESS siting and sizing. In its name, Sensitivity-Guided indicates that the candidate buses are determined by Section 2.2 before OPF is solved; Envelope-Based indicates that operating constraints are imposed simultaneously on a predefined finite scenario set; Nonanticipative Adjustable Recourse indicates that the scenarios share a day-ahead baseline plan and affine-response coefficients and produce corresponding operating responses to disturbances. The network is represented by a multiperiod P–Q LinDistFlow MILP suitable for radial distribution networks.

For subsequent model description, Table 1 summarizes the main sets, indices, capacity variables, affine-policy variables, network operating variables, and auxiliary variables. Unless otherwise specified, s denotes a finite scenario, t denotes a scheduling period, and b denotes the BESS connection bus fixed in advance for a single candidate-bus optimization.

Table 1. Compact nomenclature for the SG-ENAR-OPF formulation.

Symbol	Definition	Unit/Type
$\mathcal{N}, \mathcal{E}, \mathcal{T}$	Bus, branch, and scheduling-period sets	set
$\Omega, \Omega_{size}, \Omega_{DA}$	Generic, sizing, and day-ahead finite scenario sets	set
$\mathcal{N}_{elig}$	Eligible/installable nonslack bus set before candidate screening	set
$\mathcal{N}_{cand}$	Screened main BESS candidate-bus set determined before OPF	set
b	Fixed BESS installation bus in one candidate-wise MILP	bus index
$s, t, i, (i, j)$	Scenario, time, bus, and branch indices	index
$\xi^s_{L,t}, \xi^s_{R,t}$	Standardized load and aggregate renewable policy coordinates	dimensionless
$\rho^s_{L,t}$	System-level load multiplier mapped from $\xi^s_{L,t}$	dimensionless
P^{cap}, E^{cap}	BESS rated power and energy capacity	MW, MWh
E^{init}	Initial BESS stored energy	MWh

Table 1. *Cont.*

Symbol	Definition	Unit/Type
$P^0_{ch,t}, P^0_{dis,t}$	Shared day-ahead baseline charging/discharging powers	MW
$\beta_{\sigma,\mathcal{V},t}$	Shared affine response coefficient	MW
$P^s_{ch,t}, P^s_{dis,t}, P^s_{bat,t}$	Scenario charging, discharging, and net BESS active power	MW
$Q^s_{bat,t}$	BESS reactive-power injection	MVAr
E^s_t, u^s_t	Stored energy and charge/discharge mode	MWh, binary
$P^s_{ij,t}, Q^s_{ij,t}$	Active/reactive branch power flow	MW, MVAr
$v^s_{i,t}$	Squared voltage magnitude	p.u.^{2}
$P^s_{G,t}, Q^s_{G,t}$	Active/reactive exchange with upstream grid	MW, MVAr
$P^{ava,s}_{g,t}, P^s_{g,t}, Q^s_{g,t}$	Resource-level available, dispatched active, and reactive power of $g \in \{PV, WT\}$	MW, MW, MVAr
g, i_g	Renewable-resource index and its fixed connection bus, $g \in \{PV, WT\}$	index, bus index
$P^s_{g,i,t}, Q^s_{g,i,t}$	Node-level active/reactive injections mapped from renewable resource g	MW, MVAr
$P^s_{shed,i,t}, Q^s_{shed,i,t}$	Active/reactive load-loss variables	MW, MVAr
S_{base}	System apparent-power base	MVA
$d^s_{i,t}$	Absolute squared-voltage-deviation auxiliary variable	p.u.^{2}
ζ^s_t	Absolute BESS response-power auxiliary variable	MW
$z^{\beta}_{\sigma,\mathcal{V},t}$	Absolute-value auxiliary variable for affine coefficient	MW
π_s	Deterministic objective weight of scenario s	dimensionless
η_{ch}, η_{dis}	Charging/discharging efficiencies	dimensionless
ε_{rank}	Numerical floor for RobustScore metric-range normalization	dimensionless
$\tau_{\min}, \tau_{\max}$	Lower/upper rated-duration bounds	h
Δt	Scheduling time step	h

2.3.1. Sets, Variables, and Uncertainty Description

Let the time set be $\mathcal{T} = \{1, 2, \cdots, T\}$, the main BESS candidate-bus set be $\mathcal{N}_{cand}$, the finite scenario set used in the siting-and-sizing stage be Ω_{size}, and the day-ahead scheduling scenario set after the BESS configuration is fixed be Ω_{DA}. When a constraint applies to both stages, Ω denotes the finite scenario set used in the corresponding stage.

For each candidate bus $b \in \mathcal{N}_{cand}$, a separate multiscenario MILP is formulated to jointly determine the BESS power capacity P^{cap}, energy capacity E^{cap}, baseline charging/discharging schedule, and affine-response coefficients. Candidate bus b is a fixed parameter within each MILP, and all candidate buses are solved using the same model and parameter system.

Two standardized disturbance variables describe the inputs to the BESS policy:

$$\xi^s_t = \left[\xi^s_{L,t}, \xi^s_{R,t}\right], s \in \Omega, t \in \mathcal{T} \tag{28}$$

where $\xi^s_{L,t}$ is the shared load-disturbance coordinate and $\xi^s_{R,t}$ is the aggregate renewable-disturbance coordinate. On the load side, standardized disturbance coordinate $\xi^s_{L,t}$ is first converted through the quantile mapping defined in Section 2.5.1 into the corresponding system-level load multiplier $\rho^s_{L,t}$, which is then applied uniformly to the baseline load at each bus. Because active and reactive loads vary with the same system-level load multiplier:

$$P^s_{L,i,t} = P^0_{L,i}\, \rho^s_{L,t} \tag{29}$$

$$Q^s_{L,i,t} = Q^0_{L,i}\, \rho^s_{L,t} \tag{30}$$

Consequently, active and reactive loads have the same relative disturbance and are both driven by ξ_L after mapping through $\rho^s_{L,t}$. This treatment preserves each bus's baseline power factor while avoiding the introduction of fully collinear independent load-disturbance coordinates.

PV and WT retain separate available active-power outputs, utilized powers, and reactive-power variables in the network model; disturbance information is conveyed through the aggregate renewable coordinate ξ_R only at the BESS affine-policy interface. Therefore, ξ_R does not indicate that PV and WT are physically merged. The corresponding interval-forecast mapping and aggregation are described in Section 2.5.1.

For any scenario s and period t, the BESS charging and discharging powers are generated by the shared baseline schedule and affine coefficients:

$$P^s_{ch,t} = P^0_{ch,t} + \beta_{ch,L,t}\xi^s_{L,t} + \beta_{ch,R,t}\xi^s_{R,t} \tag{31}$$

$$P^s_{dis,t} = P^0_{dis,t} + \beta_{dis,L,t}\xi^s_{L,t} + \beta_{dis,R,t}\xi^s_{R,t} \tag{32}$$

where $P^0_{ch,t}$ and $P^0_{dis,t}$ are the day-ahead baseline charging and discharging schedules; $\beta_{ch,L,t}$ and $\beta_{dis,L,t}$ are the charging- and discharging-power response coefficients for the load disturbance, respectively; and $\beta_{ch,R,t}$ and $\beta_{dis,R,t}$ are the corresponding response coefficients for the renewable disturbance. The baseline and scenario-dependent BESS net active powers are defined, respectively, as:

$$P^0_{bat,t} = P^0_{dis,t} - P^0_{ch,t} \tag{33}$$

$$P^s_{bat,t} = P^s_{dis,t} - P^s_{ch,t} \tag{34}$$

where $P^s_{bat,t} > 0$ indicates BESS discharge to the distribution network and $P^s_{bat,t} < 0$ indicates BESS charging from the distribution network. The affine-response coefficients satisfy the prescribed bounds:

$$-\overline{\beta}_{\sigma,\mathcal{V},t} \le \beta_{\sigma,\mathcal{V},t} \le \overline{\beta}_{\sigma,\mathcal{V},t},\ \sigma \in \{ch, dis\},\ \mathcal{V} \in \{L, R\} \tag{35}$$

where $\overline{\beta}_{\sigma,\mathcal{V},t} = P^{cap}$ is the permitted upper bound of the corresponding affine-response coefficient. Thus, $-P^{cap} \le \overline{\beta}_{\sigma,\mathcal{V},t} \le P^{cap}$. This constraint limits the response gain per unit standardized disturbance and keeps the scale of the affine policy consistent with the BESS power capacity.

This parameterization makes all scenarios share the same day-ahead baseline schedule and affine-response coefficients, with corresponding operating responses produced only through (ξ_L, ξ_R), thereby forming a nonanticipative adjustable-recourse policy. Robust feasibility here refers only to the simultaneous satisfaction of constraints over the specified finite scenario set; it is not extended to a formal feasibility guarantee for unmodeled trajectories over a continuous disturbance domain.

2.3.2. BESS and Distribution-Network Operating Constraints

For any scenario $s \in \Omega$ and period $t \in \mathcal{T}$, the BESS power capacity P^{cap}, energy capacity E^{cap}, and initial stored energy E^{init} satisfy the capacity and SOC constraints. All scenarios share the same initial stored energy:

$$SOC^{\min}_{opt}E^{cap} \le E^{init} \le SOC^{\max}_{opt}E^{cap} \tag{36}$$

where $E^s_1 = E^{init}$ and $\forall s \in \Omega$. The BESS power and energy capacities satisfy, respectively:

$$0 \le P^{cap} \le \overline{P}^{cap} \tag{37}$$

$$0 \le E^{cap} \le \overline{E}^{cap} \tag{38}$$

where $\overline{P}^{cap}$ and $\overline{E}^{cap}$ are the maximum permitted power and energy capacities, respectively. To prevent independent optimization of power and energy capacities from producing a technically unreasonable, extremely short-duration configuration, the following rated-duration constraint is further imposed:

$$\tau_{\min} P_{cap} \;\leq\; E_{cap} \;\leq\; \tau_{\max} P_{cap} \tag{39}$$

where $\tau_{\min}$ and $\tau_{\max}$ are the minimum and maximum permitted rated durations, respectively. The main case uses 1–4 h, and the technical sensitivity case uses 2–4 h, corresponding to rated power-energy ranges of approximately 0.25C–1C and 0.25C–0.5C for P_{cap}/E_{cap}, respectively.

For each scenario, the BESS charging and discharging powers satisfy

$$0 \leq P^{s}_{ch,t} \leq P^{cap} \tag{40}$$

$$0 \leq P^{s}_{dis,t} \leq P^{cap} \tag{41}$$

The stored-energy state evolves recursively as follows:

$$E^{s}_{t+1} = E^{s}_{t} + \eta_{ch} P^{s}_{ch,t} \Delta t - \frac{1}{\eta_{dis}} P^{s}_{dis,t} \Delta t \tag{42}$$

where η_{ch} and η_{dis} are the charging and discharging efficiencies, respectively, and Δt is the interval between adjacent scheduling periods. The stored-energy state in every scenario and period satisfies:

$$SOC^{\min}_{opt} E^{cap} \leq E^{s}_{t} \leq SOC^{\max}_{opt} E^{cap} \tag{43}$$

The physical BESS SOC range is 0.10–0.90, whereas day-ahead optimization uses the tightened range 0.12–0.88; thus, $0.12E^{cap} \leq E^{s}_{t} \leq 0.88E^{cap}$ is used in day-ahead optimization, while actual-trajectory validation continues to use the physical bounds $0.10E^{cap} \leq E^{act}_{t} \leq 0.90E^{cap}$ to determine whether the physical SOC is violated. Under the single-day closure formulation, every modeled scenario also satisfies the terminal stored-energy closure constraint $E^{s}_{T+1} = E_{init}, \forall s \in \Omega$. This constraint provides terminal energy closure for the single-day day-ahead model. Continuous multiday rolling validation does not reinitialize the storage state in each window; instead, the actual terminal stored energy from the preceding day is passed to the next day under the interday state protocol described in Section 2.5.3.

To enforce mutually exclusive charging and discharging, the binary operating-state variable u^{s}_{t} is introduced. When $u^{s}_{t} = 1$, the storage unit can only charge; when $u^{s}_{t} = 0$, it can only discharge. The charging and discharging powers satisfy:

$$0 \leq P^{s}_{ch,t} \leq \overline{P}^{cap} u^{s}_{t} \tag{44}$$

$$0 \leq P^{s}_{dis,t} \leq \overline{P}^{cap} (1 - u^{s}_{t}) \tag{45}$$

$$P^{s}_{ch,t} + P^{s}_{dis,t} \leq P^{cap} \tag{46}$$

The final inequality tightens the continuous relaxation of the MILP without changing its integer-feasible set. Day-ahead scheduling with a fixed BESS configuration also requires different scenarios to share the operating mode in each period, as detailed in Section 2.5.2. The distribution network is represented by the P-Q LinDistFlow model suitable for radial structures. Let i_0 be the unique point of connection to the upstream grid, namely the slack/substation bus; for the IEEE-33 system in this study, $i_0 = 1$. To express nodal power balance uniformly for the slack and nonslack buses, define δ_{i,i_0}: when $i = i_0$, $\delta_{i,i_0} = 1$, and when $i \neq i_0$, $\delta_{i,i_0} = 0$. For any bus $i \in \mathcal{N}$, the active-power balance is:

$$\sum_{k:(k,i)\in\mathcal{E}} P^s_{ki,t} - \sum_{j:(i,j)\in\mathcal{E}} P^s_{ij,t} + \delta_{i,i_0} P^s_{G,t} + P^s_{PV,i,t} + P^s_{WT,i,t} + P^s_{bat,i,t} + P^s_{shed,i,t} = P^s_{L,i,t} \quad (47)$$

The reactive-power balance is:

$$\sum_{k:(k,i)\in\mathcal{E}} Q^s_{ki,t} - \sum_{j:(i,j)\in\mathcal{E}} Q^s_{ij,t} + \delta_{i,i_0} Q^s_{G,t} + Q^s_{PV,i,t} + Q^s_{WT,i,t} + Q^s_{bat,i,t} + Q^s_{shed,i,t} = Q^s_{L,i,t} \quad (48)$$

where $P^s_{G,t}$ and $Q^s_{G,t}$ are the active- and reactive-power exchanges with the upstream grid at the slack/substation bus, respectively. For renewable resources, let $g \in \{PV, WT\}$, and let i_g denote the fixed connection bus of resource g. Resource-level variables $P^s_{g,t}$ and $Q^s_{g,t}$ represent the utilized active power and reactive-power output of resource g in scenario s and period t, respectively. Their nodal injections are mapped through the fixed connection locations as:

$$P^s_{g,i,t} = \begin{cases} P^s_{g,t}, & i = i_g \\ 0, & i \neq i_g \end{cases} \quad (49)$$

$$Q^s_{g,i,t} = \begin{cases} Q^s_{g,t}, & i = i_g \\ 0, & i \neq i_g \end{cases} \quad (50)$$

Therefore, $P^s_{PV,i,t}$, $Q^s_{PV,i,t}$, $P^s_{WT,i,t}$, and $Q^s_{WT,i,t}$ in Equations (47) and (48) are all nodal injections mapped from the corresponding resource-level variables rather than additional independent renewable-energy decision variables. $P^s_{bat,i,t}$ and $Q^s_{bat,i,t}$ are the storage active- and reactive-power injections, and $P^s_{shed,i,t}$ and $Q^s_{shed,i,t}$ are the active- and reactive-load shedding variables.

In the current single-BESS model, only the fixed candidate bus b has a storage injection $P^s_{bat,b,t} = P^s_{dis,t} - P^s_{ch,t}$, whereas when $i \neq b$, $P^s_{bat,i,t} = 0$ and $Q^s_{bat,i,t} = 0$.

Define the squared voltage magnitude at a bus as $v^s_{i,t} = \left| V^s_{i,t} \right|^2$. When branch resistance and reactance are expressed in p.u. and branch active power $P^s_{ij,t}$ and reactive power $Q^s_{ij,t}$ are expressed in MW and MVAr, the voltage drop across branch (i, j) is written as:

$$v^s_{j,t} = v^s_{i,t} - 2\left(r_{ij} \frac{P^s_{ij,t}}{S_{base}} + x_{ij} \frac{Q^s_{ij,t}}{S_{base}}\right) \quad (51)$$

where r_{ij} and x_{ij} are the resistance and reactance of branch (i, j), respectively. The corresponding nodal-voltage constraint is $(V^{min})^2 \leq v^s_{i,t} \leq (V^{max})^2$. With $V^{\min} = 0.90$ p.u. and $V^{\max} = 1.10$ p.u. used in this study, the squared-voltage range in the optimization model is $0.81 \leq v^s_{i,t} \leq 1.21$. The nodal-voltage magnitudes reported in Section 3 are recovered from $\left| V^s_{i,t} \right| = \sqrt{v^s_{i,t}}$ rather than interpreting v directly as a voltage magnitude.

To preserve the MILP structure, component-wise limits are imposed separately on branch active- and reactive-power flows during optimization:

$$-\overline{P}_{ij} \leq P^s_{ij,t} \leq \overline{P}_{ij} \quad (52)$$

$$-\overline{Q}_{ij} \leq Q^s_{ij,t} \leq \overline{Q}_{ij} \quad (53)$$

The rectangular feasible region above is not mathematically equivalent to the physical branch apparent-power constraint $(P^s_{ij,t})^2 + (Q^s_{ij,t})^2 \leq (S^{rated}_{ij})^2$. When $\overline{P}_{ij} = \overline{Q}_{ij} = S^{rated}_{ij}$, the rectangle is an outer relaxation of the circular MVA feasible region rather than a conservative inner approximation; the apparent power permitted at a rectangle corner can reach $\sqrt{2}S_{ij}^{rated}$. Consequently, the component-wise branch limits are not described as an exact or conservative apparent-power constraint.

Consistent with the branch constraints, active-power exchange with the upstream grid satisfies $0 \leq P_{G,t}^{s} \leq \overline{P}_{G}$, whereas reactive-power exchange can vary bidirectionally according to $-\overline{Q}_{G} \leq Q_{G,t}^{s} \leq \overline{Q}_{G}$. Thus, the model does not permit active-power sales to the upstream grid through $P_G < 0$, and the day-ahead price signal applies only to active-power purchases. The specific exchange limits and price parameters are given in Section 3.1.

For PV and WT, the available active power satisfies:

$$0 \leq P_{PV,t}^{s} \leq P_{PV,t}^{ava,s} \tag{54}$$

$$0 \leq P_{WT,t}^{s} \leq P_{WT,t}^{ava,s} \tag{55}$$

The corresponding renewable-energy curtailments are defined, respectively, as:

$$P_{PV,t}^{curt,s} = \overline{P}_{PV,t}^{s} - P_{PV,t}^{s} \tag{56}$$

$$P_{WT,t}^{curt,s} = \overline{P}_{WT,t}^{s} - P_{WT,t}^{s} \tag{57}$$

The load-shedding variables satisfy:

$$0 \leq P_{shed,i,t}^{s} \leq P_{L,i,t}^{s}, \forall i \in \mathcal{N}, t \in \mathcal{T}, s \in \Omega \tag{58}$$

$$0 \leq Q_{shed,i,t}^{s} \leq Q_{L,i,t}^{s}, \forall i \in \mathcal{N}, t \in \mathcal{T}, s \in \Omega \tag{59}$$

With high penalty coefficients, these variables serve only as feasibility relaxations under extreme conditions.

PV, WT, and BESS inverters can all provide reactive-power support. To retain the MILP structure, linear P-Q capability bounds determined by rated active power and minimum power factor are used.

For $g \in \{PV, WT\}$, the active-power output satisfies:

$$0 \leq P_{g,t}^{s} \leq P_{g,t}^{ava,s} \leq P_{g}^{rated} \tag{60}$$

whereas the reactive-power output satisfies $-\overline{Q}_{g} \leq Q_{g,t}^{s} \leq \overline{Q}_{g}$, where $\overline{Q}_{g} = tan\left[arccos\left(pf_{g}^{\min}\right)\right]P_{g}^{rated}$, with an equivalent rated apparent power of $S_{g}^{eq} = \frac{P_{g}^{rated}}{pf_{g}^{min}}$. For the BESS, the reactive-power output satisfies:

$$-\overline{Q}_{bat} \leq Q_{bat,t}^{s} \leq \overline{Q}_{bat} \tag{61}$$

$$\overline{Q}_{bat} = P^{cap}\tan\left[\arccos\left(pf_{bat}^{min}\right)\right] \tag{62}$$

The corresponding equivalent $S_{bat}^{eq} = \frac{P^{cap}}{pf_{bat}^{min}}$, and $pf_{PV}^{min} = pf_{WT}^{min} = pf_{bat}^{min} = 0.95$ is used uniformly in this study. These linear rectangular P-Q bounds can include corner regions outside the circular inverter-capability constraint $P^2 + Q^2 \leq S^2$; they are therefore an engineering approximation used to preserve the MILP structure and are not equivalent to the complete inverter-capability curve.

Finally, the BESS, distribution-network, renewable-energy, and grid-exchange constraints above are imposed simultaneously for every scenario in the finite scenario set. Let $\boldsymbol{y}$ denote the optimization decision vector for the corresponding stage; the constraints can then be written uniformly as $g^{s}(\boldsymbol{y}) \leq 0,\ h^{s}(\boldsymbol{y}) = 0,\ \forall s \in \Omega$. Accordingly, finite-scenario robust feasibility means that the same policy simultaneously satisfies the stated physical and network constraints over the specified finite scenario set. Scenario weights enter only the objective function, as defined in Section 2.3.3.

2.3.3. Objective Function and Comprehensive Candidate-Bus Evaluation

For a given BESS connection bus b, SG-ENAR-OPF uses a composite engineering objective comprising network operation, renewable-energy utilization, storage configuration, and response-policy terms:

$$\min J_b = C_{grid} + C_{curt} + C_{shed} + C_{vdev} + C_{size} + C_{through} + C_{rec} + C_{\beta} \tag{63}$$

where $C_{grid} = \sum\limits_{s\in\Omega} \pi_s \sum\limits_{t\in\mathcal{T}} c_{G,t} P^s_{G,t} \Delta t$ is the grid-purchase-related term and $c_{G,t}$ is the grid-purchase coefficient in period t. $C_{curt} = \sum\limits_{s\in\Omega} \pi_s \sum\limits_{t\in\mathcal{T}} c_{curt} (P^{curt,s}_{PV,t} + P^{curt,s}_{WT,t}) \Delta t$ is the renewable-energy curtailment penalty, and c_{curt} is its coefficient. $C_{shed} = \sum\limits_{s\in\Omega} \sum\limits_{t\in\mathcal{T}} \sum\limits_{i\in\mathcal{N}} \pi_s (c^P_{shed} P^s_{shed,i,t} + c^Q_{shed} Q^s_{shed,i,t}) \Delta t$ is the load-shedding penalty; c^{shed}_P and c^{shed}_Q are both set substantially higher than regular operating terms so that load shedding serves only as a feasibility slack under extreme conditions.

Voltage deviation is linearized using the auxiliary variable $d^s_{i,t}$, which satisfies $d^s_{i,t} \geq v^s_{i,t} - 1$ and $d^s_{i,t} \geq 1 - v^s_{i,t}$, $d^s_{i,t} \geq 0$. The corresponding voltage-deviation penalty is $C_{vdev} = c_v \sum\limits_{s\in\Omega} \pi_s \sum\limits_{t\in\mathcal{T}} \sum\limits_{i\in\mathcal{N}} d^s_{i,t}$. Thus, when $c_v > 0$, $d^s_{i,t} = \left| v^s_{i,t} - 1 \right|$. Because $v^s_{i,t} = |V^s_{i,t}|^2$, this term actually penalizes $\left| |V^s_{i,t}|^2 - 1 \right|$ rather than the voltage-magnitude deviation $\left| \left| V^s_{i,t} \right| - 1 \right|$. Voltage deviation in the subsequent AC power-flow validation is calculated directly from voltage magnitudes.

The BESS capacity-size regularization term is defined as:

$$C_{size} = c_P P^{cap} + c_E E^{cap} \tag{64}$$

To suppress unnecessary BESS charging/discharging throughput, define:

$$C_{through} = c_{through} \sum_{s\in\Omega} \pi_s \sum_{t\in\mathcal{T}} \left(P^s_{ch,t} + P^s_{dis,t} \right) \Delta t \tag{65}$$

$C_{through}$ constrains unnecessary charging/discharging energy exchange during operation and is not interpreted as a battery degradation or life-loss model. To measure the magnitude of each scenario's response relative to the day-ahead baseline schedule, the nonnegative auxiliary variable ζ^s_t is introduced and satisfies $\zeta^s_t \geq P^s_{bat,t} - P^0_{bat,t}$, $\zeta^s_t \geq -\left(P^s_{bat,t} - P^0_{bat,t} \right)$, and $\zeta^s_t \geq 0$. The corresponding response power is:

$$C_{rec} = c_{rec} \sum_{s\in\Omega} \pi_s \sum_{t\in\mathcal{T}} \zeta^s_t \Delta t \tag{66}$$

When $c_{rec} > 0$, $\zeta^s_t = \left| P^s_{bat,t} - P^0_{bat,t} \right|$; therefore, C_{rec} directly measures the magnitude of the scenario response relative to the shared baseline schedule.

The affine-coefficient regularization term is defined as:

$$C_{\beta} = c_{\beta} \sum_{t} \left(\left| \beta_{ch,L,t} \right| + \left| \beta_{ch,R,t} \right| + \left| \beta_{dis,L,t} \right| + \left| \beta_{dis,R,t} \right| \right) \tag{67}$$

To preserve the MILP form, a nonnegative auxiliary variable $z^{\beta}_{\sigma,\mathcal{V},t}$ is introduced for each shared affine-response coefficient and satisfies $z^{\beta}_{\sigma,\mathcal{V},t} \geq \beta_{\sigma,\mathcal{V},t}$, $z^{\beta}_{\sigma,\mathcal{V},t} \geq -\beta_{\sigma,\mathcal{V},t} z^{\beta}_{\sigma,\mathcal{V},t} \geq 0$, where $\sigma \in \{ch, dis\}$ and $\mathcal{V} \in \{L, R\}$. The affine-coefficient regularization term is then defined as:

$$C_\beta = c_\beta \sum_{t\in\mathcal{T}} \sum_{\sigma\in\{ch,dis\}} \sum_{\mathcal{V}\in\{L,R\}} z^{\beta}_{\sigma,\mathcal{V},t} \tag{68}$$

When $c_\beta > 0$, $z^{\beta}_{\sigma,\mathcal{V},t} = \left|\beta_{\sigma,\mathcal{V},t}\right|$, with absolute values linearized using nonnegative auxiliary variables. This term suppresses excessively large affine-response gains and provides L_1-type regularization for the four sets of shared affine-response coefficients.

$C_{through}$, C_{rec}, and C_β act on total storage charging/discharging throughput, the magnitude of scenario responses relative to the baseline schedule, and the scale of affine-response coefficients, respectively; the three terms therefore regulate different quantities.

The objective above is a composite engineering objective whose terms have different physical dimensions. In C_{size}, the capacity coefficients c_p and c_E balance configuration size and operating performance, while $C_{through}$ suppresses unnecessary energy throughput; they are not equivalent to actual capital cost and battery degradation cost, respectively. Therefore, J_b is not interpreted as lifecycle economic cost in a unified monetary unit. The coefficients, units, and engineering meanings of the objective terms are provided in Section 3.1, while actual equipment investment and annualized ownership cost are evaluated separately in a subsequent post hoc economic audit.

For the finite scenario set Ω_{size} used in the siting-and-sizing stage, if the original nonnegative weights produced during scenario construction are $\widetilde{\pi}_s^{size}$, they are normalized uniformly as:

$$\pi_s^{size} = \frac{\widetilde{\pi}_s^{size}}{\sum\limits_{q\in\Omega_{\mathrm{size}}} \widetilde{\pi}_q^{size}} \tag{69}$$

so that $\pi_s^{size} \geq 0$ and $\sum\limits_{s\in\Omega_{\mathrm{size}}} \pi_s^{size} = 1$. Day-ahead scheduling with a fixed BESS configuration uses the 32 finite scenarios defined in Section 2.5. The central scenario $(\xi_L, \xi_R) = (0, 0)$ has weight $\pi^{DA}_{center} = 0.25$, while the other 24 noncentral disturbance-space scenarios share a total weight of 0.50. Thus:

$$\pi_s^{DA} = \frac{0.50}{24} = 0.0208333,\ s \in \Omega_{grid} \backslash center \tag{70}$$

The 7 temporal stress paths share a total weight of 0.25; therefore, the weight of each path is:

$$\pi_s^{DA} = \frac{0.25}{7} = 0.0357143,\ s \in \Omega_{path} \tag{71}$$

Scenario weights only adjust the relative contributions of different scenarios to the composite engineering objective and are not interpreted as probabilities estimated from historical frequencies. Finite-scenario robust feasibility results from the simultaneous satisfaction of operating constraints by all modeled scenarios in Section 2.3.2, rather than from the scenario weights themselves.

For candidate bus b, let the optimal SG-ENAR-OPF objective value be $J_b^* = \min\limits_{\mathbf{y}_b\in\mathcal{X}_b} J_b$, where $\mathcal{X}_b$ denotes the feasible region associated with candidate bus b, jointly defined by the BESS capacities, shared affine policy, network flows, stored-energy states, reactive-power variables, auxiliary variables, and binary operating variables. J_b^* evaluates the internal optimization objective for a given connection bus. To further compare the configuration scale, operating performance, and network-security characteristics of different buses, 8 evaluation metrics are extracted from the optimization results for each candidate bus:

$$m_b = [m_{b,1}, \cdots, m_{b,8}] \tag{72}$$

The 8 metrics are GridEnergy, Curtailment, VoltageDeviation, WorstVoltage, RecoursePower, LoadLoss, StoragePower, and StorageEnergy, in that order. WorstVoltage is treated as a benefit-type metric, whereas the remaining metrics are treated as cost-type metrics.

To prevent the normalization scale of RobustScore from changing with the set of buses being compared, for each test feeder and technical duration case, define $\mathcal{N}_{full} = \mathcal{N}_{elig}$ as the complete installable-bus set used in the full-node audit. The common normalization bounds for the ℓth metric are determined uniformly over this set:

$$m_\ell^{\min} = \min_{b \in \mathcal{N}_{full}} m_{b,\ell} \tag{73}$$

$$m_\ell^{\max} = \max_{b \in \mathcal{N}_{full}} m_{b,\ell} \tag{74}$$

To prevent numerical amplification when a metric has a very small value range, define the RobustScore normalization tolerance $\varepsilon_{rank} = 10^{-2}$. When $m_\ell^{\max} - m_\ell^{\min} \le \varepsilon_{rank}$, the ℓth metric is regarded as having no effective discriminatory power over the current candidate set and its normalized contribution is set to zero; otherwise, it is normalized using the corresponding min–max rule.

For a cost-type metric, define:

$$\hat{m}_{b,\ell} = \frac{m_{b,\ell} - m_\ell^{\min}}{m_\ell^{\max} - m_\ell^{\min}} \tag{75}$$

For the benefit-type metric WorstVoltage, define:

$$\hat{m}_{b,\ell} = \frac{m_\ell^{\max} - m_{b,\ell}}{m_\ell^{\max} - m_\ell^{\min}} \tag{76}$$

On the basis of uniform normalization, define:

$$RobustScore_b = \sum_{\ell=1}^{8} \lambda_\ell \hat{m}_{b,\ell} \tag{77}$$

where $\lambda_\ell \ge 0$ and $\sum_{\ell=1}^{8} \lambda_\ell = 1$. The units and weights ultimately used in this study are:

$$\lambda = [0.1818, 0.1818, 0.2273, 0.1364, 0.1364, 0.04545, 0.04545, 0.04545] \tag{78}$$

corresponding to GridEnergy, Curtailment, VoltageDeviation, WorstVoltage, RecoursePower, LoadLoss, StoragePower, and StorageEnergy, respectively.

VoltageDeviation is assigned a relatively high weight to emphasize distribution-network voltage quality; GridEnergy and Curtailment represent system power demand and renewable-energy utilization; WorstVoltage and RecoursePower represent the security boundary and disturbance-response magnitude, respectively; and the other three metrics supplement the evaluation of load reliability and storage-configuration scale. These unit-sum weights represent only the engineering evaluation preferences adopted in this study and are not unique or universally optimal evaluation coefficients. Their sensitivity is further examined in Section 3.3.5.

Finally, the preferred bus within the main candidate set is defined as:

$$b^* = \arg\min_{b \in \mathcal{N}_{cand}} RobustScore_b \tag{79}$$

where J_b^* is the optimal value of the internal SG-ENAR-OPF composite engineering objective at a given bus, whereas $RobustScore_b$ is used for composite comparison among candidate buses; the two quantities therefore serve different levels of evaluation.

The preferred bus in this study is the best audited bus under the specified model, common normalization scale, and explicit evaluation weights; it is not interpreted as a universal global optimum independent of model settings and evaluation preferences.

2.4. Solution, Audit, and Fixed-Policy Validation

2.4.1. Candidate-Set Solution and Independent Full-Node Benchmark

After the main candidate set is obtained in Section 2.2, every bus in the set is solved candidate-wise using the same SG-ENAR-OPF model and uniform solver settings. For each fixed candidate bus $b \in \mathcal{N}_{cand}$, the multiscenario MILP defined in Section 2.3 is solved independently to obtain the corresponding BESS power capacity, energy capacity, day-ahead baseline schedule, shared affine-response coefficients, and candidate-evaluation metrics.

To improve MILP numerical stability and solution efficiency, the model uses the valid inequalities in Section 2.3.2 to tighten the continuous relaxation and sets upper bounds on capacity variables from known physical ranges. The charging/discharging mode logic does not use an artificially enlarged uniform Big-M constant; instead, BESS power capacity P_{cap} directly provides a tight power upper bound, whose maximum effective value is further limited by $P^{cap} \le \overline{P}^{cap}$. The MIP gap, integer-feasibility tolerance, and time limit for each candidate solve are given in Section 3.1.

After all candidate buses are solved, they are ranked uniformly using the common normalization scale and RobustScore defined in Section 2.3.3, and the preferred bus is selected within the main candidate set. A one-time full-node audit is additionally conducted as an independent test of whether candidate-space reduction omits high-ranking buses. Using Section 2.2.1, all eligible/installable buses, namely $\mathcal{N}_{full} = \mathcal{N}_{elig} = \mathcal{N} \backslash \mathcal{N}_{slack}$, are included; every bus in $b \in \mathcal{N}_{full}$ is solved and ranked independently using exactly the same SG-ENAR-OPF model, solver settings, common normalization scale, and composite evaluation weights as the main candidate set. Let the top-K set in the complete-bus ranking be:

$$\mathcal{N}_K^{(full)} = TopK(RobustScore, \mathcal{N}_{full}) \tag{80}$$

The coverage of the complete-bus top K by the main candidate set is then defined as:

$$CR_K = \frac{\left| \mathcal{N}_K^{(full)} \cap \mathcal{N}_{cand} \right|}{K} \tag{81}$$

where $CR_1 = 1$ indicates that the main candidate set retains the best audited location in the complete-bus analysis under the current model, common normalization scale, and explicit evaluation weights.

The full-node audit is used only to evaluate result retention after candidate-space reduction. It neither participates in constructing the main candidate set nor feeds back to modify candidate buses. The audit is therefore an independent validation benchmark for this paper rather than a necessary step in the practical candidate-screening workflow.

2.4.2. Out-of-Sample Fixed-Policy Validation and AC Power-Flow Validation

After siting and sizing, denote the fixed policy by $\mathcal{F}^{fix} = (b^*, P^{cap*}, E^{cap*}, E^{init*}, \Pi^*)$, where Π^* contains the optimized baseline charging/discharging schedule and shared affine-response coefficients. During post hoc validation, this fixed policy generates BESS operating commands, and the stored-energy state is propagated using the energy-state equation in Section 2.3.2. The mapping of actual disturbances to the fixed policy is described in Section 2.5.3.

For any validation scenario s, storage-constraint residuals during fixed-policy execution are first defined. The maximum power-limit violation is:

$$\Delta P^s_{viol} = \max_t \left\{ \begin{array}{l} \max(0, -P^s_{ch,t}), \\ \max(0, -P^s_{dis,t}), \\ \max(0, P^s_{ch,t} - P^{cap}), \\ \max(0, P^s_{dis,t} - P^{cap}) \end{array} \right\} \tag{82}$$

This metric checks both the nonnegative lower bounds and rated upper bounds of charging and discharging power. The maximum stored-energy limit violation is defined as:

$$\Delta E^s_{viol} = \max_t \left\{ \max(0, SOC^{\min}_{phy} E^{cap} - E^s_t), \max(0, E^s_t - SOC^{\max}_{phy} E^{cap}) \right\} \tag{83}$$

The physical BESS SOC bounds are used in post hoc actual-operation validation. For validation cases requiring single-day terminal energy closure, the terminal stored-energy deviation is defined as:

$$\Delta E^s_{end} = \left| E^s_{T+1} - E^{init} \right| \tag{84}$$

The simultaneous charging/discharging residual is defined as:

$$\Delta CD^s = \max_t \min(P^s_{ch,t}, P^s_{dis,t}) \tag{85}$$

where ΔP^s_{viol}, ΔE^s_{viol}, ΔE^s_{end}, and ΔCD^s are the maximum power-limit violation, maximum stored-energy limit violation, terminal stored-energy deviation, and maximum simultaneous charging/discharging residual, respectively. The corresponding engineering-feasibility tolerances are given in Section 3.1. Note that ΔE^s_{end} is used for validation cases requiring single-day terminal closure. In carryover-fixed continuous multiday validation, the interday storage state is not required to return to the same initial value every day; it is evaluated according to the state-carryover rule and interday continuity error defined in Section 2.5.3.

Because the main optimization model uses the LinDistFlow approximation, fixed operating points are further subjected to independent network checks using nonlinear AC power flow. For any validation scenario s and period t, nodal loads, nodal renewable outputs mapped through fixed connection locations, BESS injections, and load shedding are converted into equivalent nodal active and reactive loads in the AC power-flow model:

$$P^{AC}_{D,i,t} = P^s_{L,i,t} - P^s_{PV,i,t} - P^s_{WT,i,t} - P^s_{bat,i,t} - P^s_{shed,i,t} \tag{86}$$

$$Q^{AC}_{D,i,t} = Q^s_{L,i,t} - Q^s_{PV,i,t} - Q^s_{WT,i,t} - Q^s_{bat,i,t} - Q^s_{shed,i,t} \tag{87}$$

Substituting these nodal injections into the standard nonlinear AC power-flow model gives the following nodal active- and reactive-power balances:

$$P_i = V_i \sum_{j \in \mathcal{N}} V_j (G_{ij} \cos\theta_{ij} + B_{ij} \sin\theta_{ij}) \tag{88}$$

$$Q_i = V_i \sum_{j \in \mathcal{N}} V_j (G_{ij} \sin\theta_{ij} - B_{ij} \cos\theta_{ij}) \tag{89}$$

where P_i and Q_i are the net active- and reactive-power injections at bus i, respectively; G_{ij} and B_{ij} are the conductance and susceptance of element (i, j) of the nodal admittance matrix, respectively; and $V_i = |V_i| e^{\mathrm{j}\theta_i}$, $\theta_{ij} = \theta_i - \theta_j$, where $\mathrm{j} = \sqrt{-1}$ is the imaginary unit.

Each validation operating point is mapped to MATPOWER 8.0, and runpf is called to execute a Newton–Raphson AC power flow. The power-mismatch termination tolerance is set to 10^{-8} p.u., and the maximum number of Newton iterations is 10. In the IEEE-33 system, Bus 1 is the reference/slack bus. Network topology and transformer/tap settings remain fixed throughout the AC checks; no network reconfiguration or tap optimization is performed. This study sets pf.enforce_q_lims = 0; therefore, the AC power-flow stage does not redispatch the reactive-power outputs of PV, WT, or BESS but retains the active- and reactive-power injections associated with the fixed policy. The slack bus supplies the remaining system power imbalance according to the standard AC power-flow definition.

For branch (i,j), the apparent powers at both ends are calculated from the from-end and to-end powers obtained by the AC power flow, and the larger value is used:

$$S_{ij,t}^{AC,s} = \max\left\{\sqrt{\left(P_{ij,t}^{\mathrm{fr},s}\right)^2 + \left(Q_{ij,t}^{\mathrm{fr},s}\right)^2}, \sqrt{\left(P_{ij,t}^{\mathrm{to},s}\right)^2 + \left(Q_{ij,t}^{\mathrm{to},s}\right)^2}\right\} \tag{90}$$

If the effective rating of the branch is S_{ij}^{rated}, the branch-feasible condition is $S_{ij,t}^{AC,s} \le S_{ij}^{rated}$. Section 3.1 specifies branch ratings and the default used when an effective rating is unavailable. The relationship between the component-wise P-Q box and the circular MVA constraint in Section 2.3.2 is not repeated here; the AC check directly compares the actual apparent power at both branch ends with the physical rating. As specified above, RATE_A is used when a valid value is available; otherwise, the default rating is 10 MVA.

For a given fixed policy, define the AC-feasibility indicator in scenario s and period t as:

$$\Phi_{AC}(s,t) = \begin{cases} & \text{\textit{AC power flow converges}} \\ 1, & V^{min} \le \left|V_{i,t}^s\right| \le V^{max},\ \forall i \in \mathcal{N} \\ & S_{ij,t}^{AC,s} \le S_{ij}^{rated},\ \forall (i,j) \in \mathcal{E} \\ 0, & \text{\textit{otherwise}} \end{cases} \tag{91}$$

Thus, an operating point is classified as AC feasible only if the AC power flow converges and all nodal-voltage magnitude and branch-rating constraints are satisfied. If the AC power flow does not converge, $\Phi_{AC}(s,t) = 0$ is set directly. The corresponding AC-feasibility rate is defined as:

$$FR_{AC} = \frac{1}{|\Omega_{val}|T} \sum_{s \in \Omega_{val}} \sum_{t \in \mathcal{T}} \Phi_{AC}(s,t) \tag{92}$$

where Ω_{val} denotes the corresponding set of post hoc validation samples.

AC nodal-voltage and branch-rating feasibility are determined directly using the physical bounds above, without additional engineering tolerances. In Section 3.1, the 0.002 MW/MWh threshold is used only for power-limit violations, energy-limit violations, terminal-energy deviations, and charging/discharging overlap residuals in the BESS fixed-policy replay; it is not used to relax AC voltage or branch-rating feasibility.

If a fixed-policy operating point fails the nonlinear AC check, the failure is recorded directly; the BESS location, capacities, and prescribed operating injections are not modified during the AC audit. Nonlinear AC power flow therefore serves as an independent post hoc network audit of the LinDistFlow planning and scheduling results, rather than as a feedback stage that corrects the original optimization results.

2.5. Interval-Forecast-Driven Scenario-Based Robust Day-Ahead Scheduling

The preceding siting-and-sizing stage determines the BESS connection location and its power and energy capacities. After the BESS configuration is fixed, the day-ahead operating stage constructs finite uncertainty scenarios using load, PV, and WT interval forecasts and re-optimizes the day-ahead baseline schedule and shared affine-response coefficients under the specified network model and BESS capacity constraints. This stage does not change the BESS location or capacities; its purpose is to obtain a shared day-ahead operating policy for subsequent validation on actual trajectories under forecast uncertainty.

Load and renewable quantile forecasts are first mapped into a two-dimensional standardized disturbance space, on which 25 static disturbance combinations are constructed. Temporal stress paths then incorporate disturbance duration, occurrence sequence, and switching pattern. Actual trajectories are used for post hoc execution validation after the day-ahead policy is frozen, as described in Section 2.5.3.

2.5.1. Mapping Interval Forecasts to a Standardized Disturbance Space

A lightweight quantile temporal convolutional network (Quantile-TCN) is used to generate conditional quantile forecasts separately for system load, PV, and onshore wind. Independent models are built for the three forecast targets, using the same data-processing procedure and network architecture. Each sample uses the preceding 168 h of historical information to predict q_t^{05}, q_t^{50}, and q_t^{95} for the next period, where q^{05}–q^{95} form the nominal 90% prediction interval and q^{50} gives the predictive median.

The Quantile-TCN uses three 64-channel temporal convolutional layers, a kernel size of 3, dropout of 0.10, and training with the pinball loss. The output layer constructs nonnegative lower and upper interval widths around q^{50}, thereby ensuring $q_t^{05} \leq q_t^{50} \leq q_t^{95}$.

The model inputs include historical load actuals, available load forecasts, solar actuals, onshore wind, total wind, offshore wind, and day-ahead prices, together with sine/cosine cyclical features for hour of day, month, and day of year.

The 8760 UTC hours in 2019 are divided chronologically into 70% training, 15% validation, and 15% test subsets without random shuffling on the original time axis. Standardization statistics for continuous features are computed only from the training subset and then fixed for the validation and test subsets. The validation subset is used for validation-loss evaluation, early stopping, and best-model selection, whereas the test subset is not used for parameter training or model selection. Because every test sample requires 168 h of historical input, the final dataset contains 1146 consecutive out-of-sample test forecasts.

To incorporate forecast intervals for different physical quantities into a unified day-ahead policy, the load, PV, and WT quantile forecasts are mapped separately to standardized disturbance coordinates. For any forecast variable $\chi \in \{L, PV, WT\}$, define the standardized disturbance coordinate $\xi_{\chi,t} \in [-1, 1]$ such that:

$$\begin{gathered} \xi_{\chi,t} = -1 \Leftrightarrow x_t = q_{\chi,t}^{05} \\ \xi_{\chi,t} = 0 \Leftrightarrow x_t = q_{\chi,t}^{50} \\ \xi_{\chi,t} = +1 \Leftrightarrow x_t = q_{\chi,t}^{95} \end{gathered} \tag{93}$$

For a given standardized coordinate $\xi_{\chi,t}$, the corresponding physical forecast value is obtained through a piecewise-linear mapping centered on q_{50}:

$$\hat{x}_t(\xi_{\chi,t}) = \begin{cases} q_{\chi,t}^{50} + \xi_{\chi,t}\left(q_{\chi,t}^{95} - q_{\chi,t}^{50}\right), & 0 \leq \xi_{\chi,t} \leq 1 \\ q_{\chi,t}^{50} + \xi_{\chi,t}\left(q_{\chi,t}^{50} - q_{\chi,t}^{05}\right), & -1 \leq \xi_{\chi,t} < 0 \end{cases} \tag{94}$$

This mapping separately retains the different interval widths on the q^{05}–q^{50} and q^{50}–q^{95} sides and therefore does not require the prediction interval to have the same physical scale on both sides of the median.

(1) Load-disturbance coordinate:

The day-ahead policy uses a shared load coordinate $\xi_{L,t} \in [-1,1]$ to describe the position of system-level load relative to the predictive median. For the load side, $\rho^{05}_{L,t}$, $\rho^{50}_{L,t}$, and $\rho^{95}_{L,t}$ are the system-load forecast multipliers in period t corresponding to q^{05}, q^{50}, and q^{95}, respectively. The standardized load-disturbance coordinate $\xi^s_{L,t}$ is mapped to the system-level load multiplier through q^{05}–q^{50}–q^{95} piecewise-linear interpolation:

$$\rho^s_{L,t} = \begin{cases} \rho^{50}_{L,t} + \xi^s_{L,t}\left(\rho^{95}_{L,t} - \rho^{50}_{L,t}\right), & \xi^s_{L,t} \geq 0 \\ \rho^{50}_{L,t} + \xi^s_{L,t}\left(\rho^{50}_{L,t} - \rho^{05}_{L,t}\right), & \xi^s_{L,t} < 0 \end{cases} \tag{95}$$

Thus, when $\xi^s_{L,t} = -1$, $\rho^s_{L,t} = \rho^{05}_{L,t}$; when $\xi^s_{L,t} = 0$, $\rho^s_{L,t} = \rho^{50}_{L,t}$; and when $\xi^s_{L,t} = 1$, $\rho^s_{L,t} = \rho^{95}_{L,t}$. The resulting system-level load multiplier is then applied uniformly to the baseline active and reactive loads at every bus: $P_{L,i,t} = P^0_{L,i}\,\rho_{L,t}$ and $Q_{L,i,t} = Q^0_{L,i}\,\rho_{L,t}$. Consequently, $\xi^s_{L,t}$ is only a standardized policy-disturbance coordinate, whereas $\rho^s_{L,t}$ is the system-level load multiplier that actually enters the distribution-network load model. The current model does not define separate, fully collinear active-load and reactive-load policy-disturbance coordinates. Instead, the single coordinate $\xi^s_{L,t}$, after mapping through $\rho^s_{L,t}$, simultaneously drives the nodal P/Q load changes while preserving each bus's baseline power-factor relationship.

(2) Renewable-disturbance coordinate:

The physical forecasts and network injections of PV and WT remain separate, but the BESS affine policy uses the shared renewable-policy coordinate $\xi_{R,t} \in [-1,1]$. During day-ahead scenario construction, the same $\xi_{R,t}$ is applied separately to the respective PV and WT quantile forecast intervals. For $g \in \{PV, WT\}$, the available output under a scenario is defined as:

$$P^{ava}_{g,t}(\xi_{R,t}) = \begin{cases} q^{50}_{g,t} + \xi_{R,t}\left(q^{95}_{g,t} - q^{50}_{g,t}\right), & 0 \leq \xi_{R,t} \leq 1 \\ q^{50}_{g,t} + \xi_{R,t}\left(q^{50}_{g,t} - q^{05}_{g,t}\right), & -1 \leq \xi_{R,t} < 0 \end{cases} \tag{96}$$

Thus, $\xi_R = +1$ indicates that PV and WT are each on their respective q^{95} side, $\xi_R = 0$ indicates that they are each on their respective q^{50} side, and $\xi_R = -1$ indicates that they are each on their respective q^{05} side. It is important that ξ_R is only the shared renewable-disturbance coordinate at the BESS policy layer. PV and WT retain separate $P^{ava}_{PV,t}$ and $P^{ava}_{WT,t}$, respectively, and enter network power balance at their own connection buses; they are not combined into one renewable unit in the physical network model.

On this basis, the day-ahead affine policy uses the two-dimensional standardized disturbance vector $\xi_t = \begin{bmatrix} \xi_{L,t} \\ \xi_{R,t} \end{bmatrix} \in [-1,1]^2$. Define the discrete level set for a single coordinate as $\Xi = \{-1, -0.5, 0, 0.5, 1\}$ and the two-dimensional static disturbance scenario set as $\Omega_{grid} = \Xi \times \Xi$, yielding $\left|\Omega_{grid}\right| = 5 \times 5 = 25$ finite disturbance combinations.

For example, $(\xi_L, \xi_R) = (+1, -1)$ represents high load and low renewable output, $(\xi_L, \xi_R) = (-1, +1)$ represents low load and high renewable output, and $(\xi_L, \xi_R) = (0, 0)$ represents the predictive-median state. The remaining grid points supplement different directions and intermediate magnitudes of load-renewable deviations. If $(\xi_L, \xi_R) \in \Omega_{grid}$, then $(-\xi_L, -\xi_R) \in \Omega_{grid}$. The grid therefore has a centrally paired geometric organization in the standardized disturbance-coordinate plane. This symmetry describes only the geometric organization of the selected discrete disturbance coordinates. Because the q^{05}–q^{50}

and q^{50}–q^{95} interval widths can differ for load, PV, and WT, the corresponding physical predictive distributions are not required to be statistically symmetric.

The disturbance space in this study is the mathematical policy-coordinate space formed by the shared load coordinate ξ_L and aggregate renewable policy coordinate ξ_R, rather than physical geographic space in the distribution network. The current model uses a system-level load multiplier and single-site PV/WT resources at fixed connection locations; it therefore does not explicitly describe spatially heterogeneous forecast errors or their statistical correlations among load buses or multiple renewable sites. The 25-point grid constructs finite combinations of load-renewable disturbance directions; time-dependent scenario construction is defined further in Section 2.5.2.

2.5.2. Temporal-Path Envelope and Terminal-SOC Safeguard

The 25 disturbance-space grid scenarios in Section 2.5.1 cover different disturbance directions in the two-dimensional load-renewable policy space, but ξ_L and ξ_R remain constant throughout the scheduling horizon in every static scenario. Because the BESS energy state evolves recursively across periods, disturbance duration, occurrence sequence, and switching pattern can produce different SOC accumulation even when instantaneous magnitudes are identical. Therefore, on the basis of the static disturbance grid Ω_{grid}, a temporal stress path set $\Omega_{path} = \{w_1, w_2, \cdots, w_7\}$ is further constructed to provide finite stress testing of intertemporal energy accumulation.

Define the high-net-load state as $(\xi_L, \xi_R) = (+1, -1)$ and the high-renewable state as $(\xi_L, \xi_R) = (-1, +1)$. The 7 baseline paths are listed in Table 2.

Table 2. Seven types of temporal stress-path scenarios.

Path	Path Type	ξ_L	ξ_R	Primary Function
w_1	High-net-load path during load-peak periods	Set to +1 in the six periods with the highest forecast-load q^{50}; 0 otherwise	Set to −1 in the corresponding six periods; 0 otherwise	Evaluate discharge and SOC drop under sustained high-load, low-renewable-energy conditions
w_2	High-renewable path during renewable-peak periods	Set to −1 in the six periods with the highest forecast PV + WT q^{50}; 0 otherwise	Set to +1 in the corresponding six periods; 0 otherwise	Evaluate charge and SOC rise under sustained high-renewable-energy, low-load conditions
w_3	High net load followed by high renewable generation	+1 in the first 12 h and −1 in the last 12 h	−1 in the first 12 h and +1 in the last 12 h	Evaluate sequence effects under conditions of discharge followed by charge
w_4	High renewable generation followed by high net load	−1 in the first 12 h and +1 in the last 12 h	+1 in the first 12 h and −1 in the last 12 h	Compare with w_3 to analyze the impact of different disturbance sequences on SOC accumulation
w_5	Transition from high net load to high renewable generation	Linearly decreases from +1 to −1	Linearly increases from −1 to +1	Evaluate energy response under continuous state transitions
w_6	Transition from high renewable generation to high net load	Linearly increases from −1 to +1	Linearly decreases from +1 to −1	Evaluate energy response under continuous reverse transitions

Table 2. *Cont.*

Path	Path Type	ξ_L	ξ_R	Primary Function
w_7	Hourly alternation between high net load and high renewable generation	+1 in odd periods and −1 in even periods	−1 in odd periods and +1 in even periods	Evaluate characteristics of high-frequency switching between positive and negative disturbances and energy cancellation

The linear transitions in paths w_5 and w_6 are further defined as:

$$\xi_{L,t}^{(5)} = 1 - 2\frac{t-1}{T-1},\ \xi_{R,t}^{(5)} = -1 + 2\frac{t-1}{T-1} \tag{97}$$

$$\xi_{L,t}^{(6)} = -1 + 2\frac{t-1}{T-1},\ \xi_{R,t}^{(6)} = 1 - 2\frac{t-1}{T-1} \tag{98}$$

Paths w_3 and w_4 contain the same two extreme states but reverse their occurrence order and therefore explicitly test the dependence of SOC recursion on disturbance timing. w_5 and w_6 describe continuous transitions between the two extreme states, whereas w_7 examines charging/discharging responses to rapid switching between positive and negative disturbances. These temporal paths do not add a new uncertainty dimension; they operate within the two-dimensional policy space (ξ_L, ξ_R) defined in Section 2.5.1.

They introduce temporal structure into that space. The final finite day-ahead scenario set is defined as:

$$\Omega_{DA} = \Omega_{grid} \cup \Omega_{path} \tag{99}$$

where $\left|\Omega_{grid}\right| = 25$ and $\left|\Omega_{path}\right| = 7$. The 25 static disturbance-space scenarios mainly supplement combinations of different load-renewable disturbance directions and magnitudes, whereas the 7 temporal stress paths further represent the effects of disturbance persistence, occurrence sequence, continuous transition, and frequent switching on the intertemporal BESS energy state.

To preserve the nonanticipativity of the day-ahead policy, all scenarios share the same baseline day-ahead charging/discharging schedule and affine-response coefficients. The charging/discharging operating mode is also shared across scenarios within each period and can be expressed as:

$$u_t^s = u_t^{DA}, \forall s \in \Omega_{DA} \tag{100}$$

Different scenarios can therefore produce different BESS power responses and SOC trajectories, but an independent day-ahead operating policy cannot be reselected separately for any individual scenario.

(1) Tightened SOC interval and terminal-energy response:

To retain an energy margin before post hoc fixed-policy execution, day-ahead optimization uses an SOC interval strictly inside the physical SOC range:

$$SOC_{phy}^{\min} < SOC_{opt}^{\min} < SOC_{opt}^{\max} < SOC_{phy}^{\max} \tag{101}$$

In the final model, day-ahead optimization and physical validation use SOC ranges of 0.12–0.88 and 0.10–0.90, respectively; the parameter settings are summarized in Section 3.1. These two SOC levels are also used uniformly in the final scenario-construction ablation. Further, using the shared affine policy in Section 2.3.1 and the stored-energy recursion in

Section 2.3.2, define the response coefficients of the terminal-energy change to the shared load and renewable disturbances in period t:

$$\kappa_{L,t} = \Delta t(\eta_{ch}\beta_{ch,L,t} - \frac{\beta_{dis,L,t}}{\eta_{dis}}) \tag{102}$$

$$\kappa_{R,t} = \Delta t(\eta_{ch}\beta_{ch,R,t} - \frac{\beta_{dis,R,t}}{\eta_{dis}}) \tag{103}$$

For any temporal disturbance trajectory, its terminal-energy response relative to the baseline schedule can be written as $\Delta E_T = \sum_{t=1}^{T} \kappa_{L,t}\xi_{L,t} + \sum_{t=1}^{T} \kappa_{R,t}\xi_{R,t}$. Because the policy coordinates satisfy $|\xi_{L,t}| \leq 1$ and $|\xi_{R,t}| \leq 1$, it follows that $|\Delta E_T| \leq \sum_{t=1}^{T} (|\kappa_{L,t}| + |\kappa_{R,t}|)$. This relationship shows that even when different temporal paths are formed from the same two-dimensional disturbance bounds, their durations and temporal structures produce different cumulative energy responses through SOC recursion. Beyond the static 25-point disturbance space, temporal stress paths are therefore added.

To limit this terminal-energy response in the linear model, introduce the nonnegative auxiliary variables $z^E_{L,t} \geq \kappa_{L,t}$, $z^E_{L,t} \geq -\kappa_{L,t}$, $z^E_{R,t} \geq \kappa_{R,t}$, and $z^E_{R,t} \geq -\kappa_{R,t}$, such that $z^E_{L,t} \geq |\kappa_{L,t}|$ and $z^E_{R,t} \geq |\kappa_{R,t}|$. The final temporal-path formulation enables the terminal-energy safeguard:

$$\sum_{t=1}^{T} (z^E_{L,t} + z^E_{R,t}) \leq \varepsilon_E \tag{104}$$

where this study uses $\varepsilon_E = 10^{-6}$ MWh. This constraint suppresses excessive terminal-energy drift of the shared affine policy under bounded temporal disturbances and, together with the temporal stress paths, forms the intertemporal energy safeguard in the final day-ahead model. The final ablation study also keeps the same SOC ranges in all three configurations and enables the terminal-energy guard only in the complete 25-point grid + 7 temporal paths configuration.

The daily-closed and carryover-fixed multiday operating protocols are not repeated here. The single-day state constraint for the former is defined in Section 2.3.2, while the interday carryover rule for actual terminal energy under the latter is described uniformly in Section 2.5.3.

(2) Scope of finite-scenario robustness:

The scenario-based robust scheduling in this study requires the same day-ahead baseline schedule, shared affine-response coefficients, and operating modes to satisfy the operating constraints of every modeled scenario simultaneously, namely:

$$g^s(\boldsymbol{y}) \leq 0,\ h^s(\boldsymbol{y}) = 0,\ \forall s \in \Omega_{DA} \tag{105}$$

Accordingly, the BESS power, SOC, terminal energy, nodal power balance, voltage, and branch operating constraints are imposed separately on all 32 scenarios. The scenario weights defined in Section 2.3.3, π_s, only adjust the relative contribution of each scenario to the composite engineering objective and do not alter the requirement that every scenario satisfy the operating constraints. In other words, robustness here results from joint feasibility over the modeled finite scenario set, rather than from scenario probability weights or a worst-case objective.

It must also be made clear that Ω_{DA} is only a finite discrete representation of the continuous standardized disturbance space and possible temporal trajectories. This study does not formulate an $\max(\xi \in U)$-type continuous-domain worst-case robust counterpart, nor does it use adversarial separation or column-and-constraint generation to verify every unsampled disturbance. Therefore, satisfaction of the constraints by $\forall s \in \Omega_{DA}$ does not imply a formal feasibility guarantee for $\forall(\xi_L, \xi_R) \in [-1,1]^2$ or for all possible 24 h disturbance sequences. The term 'robust' in this study is strictly limited to joint constraint satisfaction over the predefined finite scenario set.

Finally, the 7 temporal stress paths above are predefined engineering stress templates rather than probabilistic scenarios fitted directly from historical forecast-error trajectories. The 6 h peak duration, 12/12 h segmentation, full-horizon linear transition, and 1 h alternation frequency are engineering-interpretable but not theoretically unique design parameters. To evaluate dependence on these manual parameters, Section 3.3.8 further varies peak duration, segment location, alternation frequency, number of paths, and transition shape in a path-design sensitivity analysis. Following this logic, the revised Section 3 reports variants including PEAK_H4/H8, SPLIT_H8/H16, ALT_2H, COUNT_3/5, and RAMP_12H.

2.5.3. Actual Fixed-Policy and Multiday Rolling Validation

After day-ahead scheduling, the BESS location and capacities and the resulting baseline charging/discharging schedule and shared affine-response coefficients are fixed. Actual observations do not participate in day-ahead scenario construction or policy optimization and are used only for fixed-policy execution replay after the policy has been determined. To map actual load, PV, and WT observations to the two-dimensional policy interface defined in Section 2.3.1, the quantile mapping in Section 2.5.1 is first used to calculate the raw standardized deviation of each physical quantity.

For $\chi \in \{L, PV, WT\}$, denote the actual observation by $x_{\chi,t}^{act}$. Its raw standardized disturbance coordinate is defined as:

$$\xi_{\chi,t}^{raw} = \begin{cases} \frac{x_{\chi,t}^{act} - q_{\chi,t}^{50}}{q_{\chi,t}^{95} - q_{\chi,t}^{50} + \varepsilon_\xi}, x_{\chi,t}^{act} \geq q_{\chi,t}^{50} \\ \frac{x_{\chi,t}^{act} - q_{\chi,t}^{50}}{q_{\chi,t}^{50} - q_{\chi,t}^{05} + \varepsilon_\xi}, x_{\chi,t}^{act} \leq q_{\chi,t}^{50} \end{cases} \tag{106}$$

Thus, when the actual value lies inside the prediction interval, $\xi_{\chi,t}^{raw} \in [-1,1]$, whereas when it exceeds q^{95} or falls below q^{05}, $\left|\xi_{\chi,t}^{raw}\right| > 1$. The raw coordinate only represents the degree to which an actual observation departs from the prediction interval and is used for out-of-interval disturbance diagnostics in Section 3.3.7. For PV and WT, over periods in which both raw coordinates are finite, the following is the capacity-weighted aggregate renewable raw diagnostic coordinate:

$$\xi_{R,t}^{(raw,agg)} = \frac{P_{PV}^{rated}\xi_{PV,t}^{raw} + P_{WT}^{rated}\xi_{WT,t}^{raw}}{P_{PV}^{rated} + P_{WT}^{rated}},\ t \in \mathcal{T}_{raw}^{valid} \tag{107}$$

where $\mathcal{T}_{raw}^{valid} = \left\{t : \xi_{PV,t}^{raw} and \xi_{WT,t}^{raw} are\ both\ finite\right\}$. This $\xi_{R,t}^{(raw,agg)}$ is used only for post hoc raw-coordinate distribution diagnostics and is not an input to the frozen affine policy. For degenerate periods in which the prediction-interval width is near zero, a small numerical floor $\varepsilon_\xi = 10^{-8}$ is introduced to prevent denominator degeneration.

In particular, nighttime PV can exhibit $q_{PV,t}^{05} \approx q_{PV,t}^{50} \approx q_{PV,t}^{95} \approx 0$. When the corresponding quantile-interval width does not exceed ε_ξ, an amplified standardized deviation is not calculated from the near-zero width; instead, the PV policy coordinate for that period is set to the neutral value $\xi_{PV,t} = 0$. This treatment affects only the numerical mapping of the policy input and does not change the actual physical PV output. A PV raw coordinate that cannot be stably defined is excluded only from the raw-coordinate distribution statistics.

(1) Bounded policy coordinates:

Because the day-ahead affine rule is constructed and constrained over the policy domain $[-1,1]^2$, actual execution does not extrapolate the unbounded ξ^{raw} directly through the affine policy. Instead, the policy-input coordinates are first saturated:

$$\overline{\xi}_{\chi,t} = sat\left(\xi_{\chi,t}^{raw}, -1, 1\right) = \min\left\{1, \max\left(-1, \xi_{\chi,t}^{raw}\right)\right\} \tag{108}$$

For load, $\xi_{L,t}^{act} = \overline{\xi}_{L,t}$. It is important that this saturation occurs only for the BESS affine-policy input coordinate. If actual load exceeds q^{95} or falls below q^{05}, the true physical load still enters subsequent network validation as the actual observation and is not truncated to the prediction interval.

(2) Actual PV/WT coordinates and the aggregate renewable policy coordinate:

The physical outputs of PV and WT are modeled separately in the day-ahead model, whereas the BESS policy interface uses the shared renewable coordinate ξ_R. During actual execution, $\xi_{PV,t}^{raw}$ and $\xi_{WT,t}^{raw}$ are therefore calculated separately, yielding the following bounded coordinates:

$$\overline{\xi}_{PV,t} = sat(\xi_{PV,t}^{raw}, -1, 1) \tag{109}$$

$$\overline{\xi}_{WT,t} = sat(\xi_{WT,t}^{raw}, -1, 1) \tag{110}$$

The actual aggregate renewable policy coordinate is then constructed using the rated active-power capacities of PV and WT:

$$\xi_{R,t}^{act} = \left[\frac{P_{PV}^{rated}\overline{\xi}_{PV,t} + P_{WT}^{rated}\overline{\xi}_{WT,t}}{P_{PV}^{rated} + P_{WT}^{rated}}\right] \in [-1,1] \tag{111}$$

This capacity-weighted mapping allows actual PV and WT forecast errors to have opposite directions: if one is above its median and the other below, their standardized responses partially cancel in the aggregate policy coordinate. This aggregation is a policy-dimension-reduction treatment and does not imply that a single ξ_R retains every independent joint PV-WT error combination. The corresponding limitation is discussed further in the Discussion section.

(3) Execution of the frozen affine policy:

After the actual policy coordinates are obtained, $(\xi_{L,t}^{act}, \xi_{R,t}^{act})$ is substituted into the shared affine policy already determined by the day-ahead problem. The actual charging and discharging commands are, respectively:

$$P_{ch,t}^{act} = P_{ch,t}^{0} + \beta_{ch,L,t}\xi_{L,t}^{act} + \beta_{ch,R,t}\xi_{R,t}^{act} \tag{112}$$

$$P_{dis,t}^{act} = P_{dis,t}^{0} + \beta_{dis,L,t}\xi_{L,t}^{act} + \beta_{dis,R,t}\xi_{R,t}^{act} \tag{113}$$

The corresponding net BESS active-power injection is $P_{bat,t}^{act} = P_{dis,t}^{act} - P_{ch,t}^{act}$, and the actual stored energy is propagated hourly using the same physical state equation:

$$E_{t+1}^{act} = E_{t}^{act} + \eta_{ch}P_{ch,t}^{act}\Delta t - \frac{P_{dis,t}^{act}}{\eta_{dis}}\Delta t \tag{114}$$

The policy coordinate is bounded; the final BESS command is not. During actual validation, $P_{ch,t}^{act}$ and $P_{dis,t}^{act}$ are not re-optimized and are not subjected to clipping, projection, or other post hoc correction. Consequently, if the frozen affine policy produces negative charging/discharging power, power exceeding P^{cap}, an SOC violation, a terminal-energy deviation, or simultaneous charging and discharging, the issue is retained as an actual fixed-policy execution residual and evaluated using ΔP_{viol}, ΔE_{viol}, ΔE_{end}, and ΔCD as defined in Section 2.4.2.

Meanwhile, the actual physical load, PV, and WT retain their observed values and enter the network-operation and nonlinear AC audits defined in Section 2.4.2 together with the fixed BESS injection. Physical disturbances outside the prediction intervals are therefore not removed from the validation data.

(4) Multiday rolling execution replay:

To avoid evaluating fixed-policy execution performance from a single day alone, a multiday rolling execution replay is conducted on chronologically ordered test-period data. Each rolling window contains 48 h of data, comprising an initial 24 h = actual execution horizon, followed by a 24 h = planning look-ahead horizon. Adjacent windows advance in 24 h steps. Thus, adjacent 48 h planning horizons overlap by 24 h, whereas their actual execution horizons do not overlap. The 46 consecutive rolling windows yield $46 \times 24 = 1104$ consecutive, nonrepeated actual execution hours. The 46 windows are therefore not interpreted as 46 statistically independent samples but as a continuous energy-management-system execution replay over one time series.

Each window follows the sequence of forecast/scenario construction, day-ahead optimization, policy freezing, actual replay, and storage/network validation. Within the actual execution horizon, observations are used only for post hoc validation after the day-ahead policy has been determined; they do not participate in retraining the day-ahead problem or feeding scenario corrections back into it.

(5) Daily-closed and carryover-fixed state protocols:

For interday SOC coupling, two validation protocols are distinguished: daily-closed and carryover-fixed.

In the daily-closed formulation, the initial stored energy E_{init} of each day-ahead problem is determined within the range defined in Section 2.3.2, namely the optimization SOC range, and satisfies single-day terminal closure $E^s_{T+1} = E_{init}$, $s \in \Omega_{DA}$ in the modeled scenarios. This formulation is used for single-day day-ahead scheduling and the corresponding fixed-policy validation, but each rolling day can independently establish its day-ahead initial state; therefore, the preceding day's actual terminal energy is not explicitly passed forward.

Under the carryover-fixed protocol, interday SOC coupling is retained explicitly by applying a continuous multiday carryover-fixed validation to the final configuration. Day 1 uses the initial energy $E^{init}_{1,1}$ determined by its day-ahead problem. Starting from day 2, the terminal energy at the end of the preceding day's actual execution is fixed, without modification, as the initial state for the next day's day-ahead problem:

$$E^{init}_{d+1,1} = E^{act}_{d,T+1},\ d = 1, \ldots, D-1 \tag{115}$$

Given this inherited state, for day d + 1 the day-ahead problem is re-solved to obtain the new baseline schedule and affine-response coefficients for that day. The day-ahead policy is then frozen again, and the day's actual trajectory is used for fixed-policy replay. Accordingly, carryover-fixed does not simply concatenate 46 independent daily-closed frozen policies; it proceeds as $E^{act}_{d,T+1} \rightarrow E^{init}_{d+1,1} \rightarrow$ re-solve day-ahead $\rightarrow$ freeze new policy $\rightarrow$ actual replay.

The interday state transfer itself applies no SOC clipping; $E^{act}_{d,T+1}$ is not artificially projected back to a predefined initial SOC but is used directly as the fixed initial state of the next day's day-ahead problem. The actual fixed-policy execution stage likewise applies no clipping, projection, or secondary optimization to the BESS charging/discharging commands.

The interday continuity error between adjacent days is defined as $\Delta E_{bdry,d} = \left| E^{init}_{d+1,1} - E^{act}_{d,T+1} \right|$, with $\Delta E_{bdry,d} = 0$ under ideal continuity. This metric evaluates whether SOC-boundary transfer is continuous under the carryover-fixed protocol;

actual results are reported in Section 3.3.11. Finally, carryover-fixed validation shows that the fixed-policy execution can operate with explicit interday SOC transfer over the tested consecutive windows, but it does not establish a theoretical guarantee of recursive-feasibility for arbitrary indefinite actual-disturbance sequences.

3. Results

3.1. Test System and Unified Parameter Settings

3.1.1. Test System and Basic Settings

Validation is performed primarily on the IEEE-33-bus radial distribution system, which contains 32 analyzable buses excluding the slack bus. Based on the low-voltage characteristics at the feeder end under the base operating condition, Buses 31, 32, and 33 are designated as the target weak-voltage area. PV is connected at Bus 18 and WT at Bus 25; therefore, the fixed renewable connection buses defined in Section 2.3.2 are $i_{PV} = 18$ and $i_{WT} = 25$, respectively. The rated active-power capacities of PV and WT are 2.5 MW and 2.0 MW, respectively. The system-level load multiplier is mapped to the network according to each bus's share of baseline load, and the PV and WT forecast profiles are multiplied by their installed capacities to obtain available output.

The main candidate-screening procedure uses the original IEEE-33 base operating condition as the reference point and applies net-injection perturbations of $\Delta P = 0.01$ MW and $\Delta Q = 0.01$ MVAr separately at each candidate bus. AC finite differences are used to calculate nodal sensitivities for minimum voltage, system active-power loss, maximum branch apparent power, and target-area voltage support. This reference setting is used only to obtain the main candidate-screening result; its dependence on operating condition, perturbation magnitude, and finite-difference form is audited separately in Section 3.3.3.

To examine the stability of sensitivity rankings with respect to operating condition and finite-difference setting, five representative operating points are selected from the 46 official rolling windows, which contain 1104 actual execution points: median actual, peak load, high PV, high wind, and minimum grid import. A synthetic reverse-flow stress point is also added. At every operating point, three perturbation magnitudes, 0.005, 0.010, and 0.020 MW/MVAr, are applied, and both forward and central finite differences are calculated.

An AC power-flow scan of all 1104 actual time points found no actual reverse-flow state with slack active power below zero. Therefore, the reverse-flow case is not selected directly from an actual sample; instead, starting from the minimum-import actual point, renewable penetration is increased to construct a synthetic stress case for sensitivity-stability testing. This operating condition is used only in Section 3.3.3 for the sensitivity stress audit and is not treated as an actual operating sample.

Figure 2 shows the IEEE-33 system topology, main candidate-bus categories, and the PV, WT, and final BESS connection locations. The modified feeder used here retains the original MATPOWER network topology and electrical parameters; renewable connection locations and storage-related configurations are introduced as additional simulation settings.

All experiments are conducted using an Intel Core i5-14600K, 32 GB RAM, Windows 11, and MATLAB R2023a environment. In the siting-and-sizing stage, MATLAB intlinprog solves the candidate-wise MILP; in day-ahead scheduling with a fixed configuration, YALMIP invokes CPLEX 12.0. Table 3 lists the uniform siting-and-sizing solver settings used for every candidate bus:

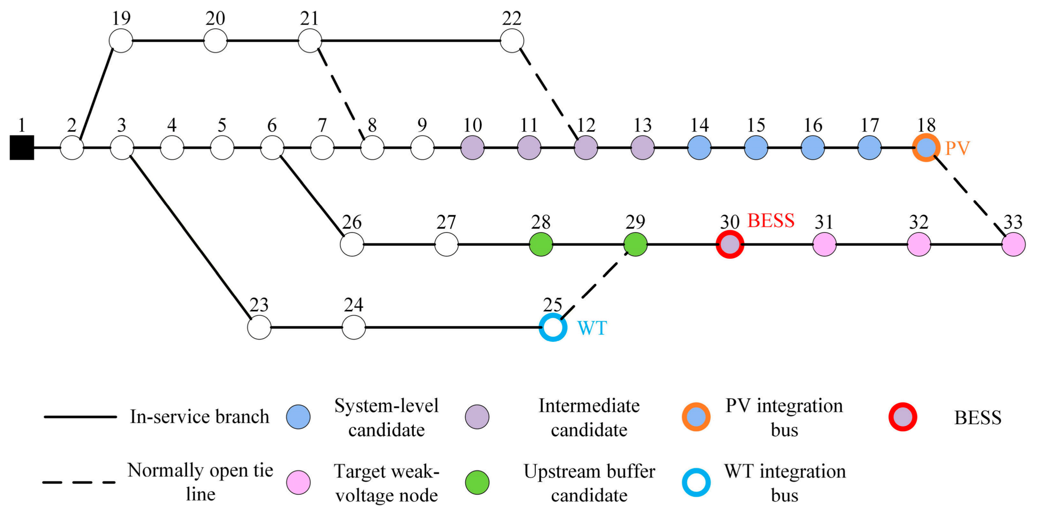

Figure 2. Topology, candidate-bus classification, and PV–wind–BESS integration locations in the IEEE-33-bus distribution system.

Table 3. Candidate-wise planning and solver settings.

Item	Setting
Main candidate number	15
Planning solver	MATLAB intlinprog
Relative MIP gap	10^{-4}
Absolute MIP gap	10^{-6}
Integer tolerance	10^{-6}
Time limit per candidate	600 s

The MIP gap for day-ahead scheduling is set to 10^{-3}, the maximum solution time for a single rolling window is 1800 s, and the maximum solution time for actual fixed-policy validation is 600 s. Apart from these settings, CPLEX's internal FeasibilityTol, OptimalityTol, and IntFeasTol retain their default values.

3.1.2. BESS, Scenario, and Validation Parameters

The day-ahead scheduling time step is $\Delta t = 1$ h, and both BESS charging and discharging efficiencies are set to $\eta_{ch} = \eta_{dis} = 0.95$. In the final duration-aware siting-and-sizing model, the upper bounds on BESS power and energy capacities are respectively $\overline{P}^{cap} = 20$ MW, $\overline{E}^{cap} = 80$ MWh. To examine the effect of the rated-duration constraint on the configuration result, two technical cases are specified:

$$MAIN : 1\,\mathrm{h} \leq \frac{E^{cap}}{P^{cap}} \leq 4\,\mathrm{h} \tag{116}$$

$$STRICT : 2\,\mathrm{h} \leq \frac{E^{cap}}{P^{cap}} \leq 4\,\mathrm{h} \tag{117}$$

The two cases differ only in the minimum permitted rated duration; their final siting-and-sizing results are compared in Section 3.3.5. Day-ahead scheduling with a fixed BESS configuration uses the tightened SOC range $0.12 \leq SOC_{opt} \leq 0.88$, whereas post hoc actual fixed-policy validation evaluates stored-energy feasibility using the physical SOC range $0.10 \leq SOC_{phy} \leq 0.90$. The daily-closed and carryover-fixed state protocols are defined in Section 2.5.3 and are not repeated here.

The minimum power factors of PV, WT, and the BESS inverter are all set to $pf_{PV}^{\min} = pf_{WT}^{\min} = pf_{bat}^{\min} = 0.95$. In the IEEE-33 system, the rated active-power capacities of PV and WT are 2.5 MW and 2.0 MW, respectively. According to Section 2.3.2, using the PF-derived capability approximation, their reactive-power limits are:

$$\overline{Q}_{PV} = 2.5\tan[\arccos(0.95)] = 0.8217 \text{ MVAr} \tag{118}$$

$$\overline{Q}_{WT} = 2.0\tan[\arccos(0.95)] = 0.6574 \text{ MVAr} \tag{119}$$

The corresponding equivalent nameplate apparent-power ratings are, respectively:

$$S_{PV}^{eq} = \frac{2.5}{0.95} = 2.6316 \text{ MVA} \tag{120}$$

$$S_{WT}^{eq} = \frac{2.0}{0.95} = 2.1053 \text{ MVA} \tag{121}$$

The BESS reactive-power limit and equivalent apparent-power rating are calculated from the optimized P^{cap} using the same relationship. The final MAIN and STRICT values are reported in Section 3.3.5. This section lists only the parameter settings and does not repeat the geometric difference between the linear P-Q capability approximation and the circular inverter capability described in Section 2.3.2.

The physical nodal-voltage magnitude range is $0.90 \leq \left|v_{i,t}^s\right| \leq 1.10$ p.u., whereas the LinDistFlow optimization model uses the squared-voltage state $v_{i,t}^s = \left|V_{i,t}^s\right|^2$; thus, the squared-voltage state in LinDistFlow satisfies $0.81 \leq v_{i,t}^s \leq 1.21$. Subsequent nonlinear AC validation directly uses voltage magnitude $\left|V_{i,t}^s\right|$ and checks the physical voltage bounds of 0.90–1.10 p.u. In the siting-and-sizing stage, the active- and reactive-power component limits of every branch are respectively $\left|P_{ij,t}^s\right| \leq 30$ MW, $\left|Q_{ij,t}^s\right| \leq 30$ MVAr; the corresponding upstream-grid exchange ranges are $0 \leq P_{G,t}^s \leq 30$ MW and $\left|Q_{G,t}^s\right| \leq 30$ MVAr, whereas in day-ahead scheduling with a fixed configuration, the upstream-grid exchange ranges are tightened to $0 \leq P_{G,t}^s \leq 20$ MW and $\left|Q_{G,t}^s\right| \leq 20$ MVAr. In addition, the system base capacity used by the LinDistFlow model is $S_{base} = 10$ MVA, consistent with the MATPOWER baseMVA setting of the IEEE-33 case33bw test system.

For branch ratings, when the original MATPOWER feeder data provide a valid RATE_A, that value is used; branches without a valid rating use $S_{ij}^{rated} = 10$ MVA uniformly. The distinction between branch component-wise P–Q bounds and the physical MVA constraint is explained in Section 2.3.2, and apparent-power evaluation in the nonlinear AC check is described in Section 2.4.2, with only the specific parameters retained here.

The day-ahead policy-disturbance coordinates satisfy $\xi_L, \xi_R \in [-1, 1]$, and the static disturbance grid uses $\{-1, -0.5, 0, 0.5, 1\}^2$. The final day-ahead model uses the scenarios defined in Section 2.5: 25 static disturbance-space scenarios and seven temporal stress paths. The terminal-energy safeguard is set to $\varepsilon_E = 10^{-6}$ MWh. Objective-aggregation weights for all scenarios are defined in Section 2.3.3; therefore, the specific allocation of 0.25, 0.50/24, and 0.25/7 is not repeated here. For fixed-policy storage replay, the engineering-feasibility thresholds are $\varepsilon_P^{eng} = 0.002$ MW and $\varepsilon_E^{eng} = 0.002$ MWh. These thresholds are used to determine engineering feasibility for the power-limit violation, stored-energy limit violation, terminal-energy deviation, and charging/discharging overlap residual defined in Section 2.4.2; they are not used to relax AC nodal-voltage or branch-rating constraints.

3.1.3. Solution and Objective-Function Settings

Both the siting-and-sizing and day-ahead scheduling stages use the Section 2.3.3-defined composite engineering objective. Because the physical quantities in the objective terms include energy, power, squared-voltage deviation, affine-response magnitude, and BESS power/energy capacities, their coefficients do not share a common monetary unit. For reproducibility, Table 4 summarizes the numerical coefficient, weighted physical quantity, unit, and engineering role of each objective term.

Table 4. Engineering objective coefficients for the siting-and-sizing and day-ahead scheduling stages.

Stage	Objective Term	Coefficient/ Setting	Weighted Quantity	Quantity Unit	Coefficient Unit/Scaling Role	Engineering Interpretation
Siting/sizing	Grid-purchase-related term	1.0	Grid-import energy	MWh	o.u./MWh	Controls the contribution of upstream-grid import in the engineering objective
	Renewable-curtailment penalty	300	Curtailed renewable energy	MWh	o.u./MWh	Discourages renewable curtailment
	Load-shedding penalty	10^7	Active/reactive load shedding	MWh/MVArh	o.u./MWh; o.u./MVArh	High feasibility penalty used only under extreme conditions
	Voltage-deviation penalty	100	Aggregated $\left\|v_{i,t}^{s}-1\right\|$	p.u.^{2}	o.u./p.u.^{2}; scales voltage-quality contribution	Penalizes deviation from nominal voltage state
	BESS-throughput regularization	1.0	Charge/discharge energy throughput	MWh	o.u./MWh	Suppresses unnecessary energy cycling
	Recourse regularization	20	Time-integrated absolute BESS response deviation	MWh	o.u./MWh	Suppresses excessive scenario-dependent response
	Affine-coefficient regularization	0.10	$\|\beta_{\sigma,\mathcal{V},t}\|$	MW	o.u./MW; scales affine-response magnitude	Limits excessive affine-response gain
	Power-capacity regularization	100	P^{cap}	MW	o.u./MW	Controls BESS power-size contribution
	Energy-capacity regularization	20	E^{cap}	MWh	o.u./MWh	Controls BESS energy-size contribution
Day-ahead	Grid-purchase term	price profile, scale 1.0	Grid-import energy	MWh	dataset-native price unit/MWh	Uses the day-ahead electricity-price signal
	Renewable-curtailment penalty	100	Curtailed renewable energy	MWh	o.u./MWh	Discourages renewable curtailment
	Active-load-shedding penalty	10^5	Active-load shedding	MWh	o.u./MWh	Emergency active-power feasibility penalty
	Reactive-load-shedding penalty	0.1 × active penalty	Reactive-load shedding	MVArh	o.u./MVArh	Emergency reactive-power feasibility penalty
	Voltage-deviation penalty	10	Aggregated $\left\|v_{i,t}^{s}-1\right\|$	p.u.^{2}	o.u./p.u.^{2}; scales voltage-quality contribution	Maintains voltage quality under uncertain operation
	BESS-throughput regularization	0.05	Charge/discharge energy throughput	MWh	o.u./MWh	Mildly suppresses unnecessary throughput

In Table 4, o.u. denotes objective unit and is used only to describe the numerical scaling among different physical quantities in the composite scalarization; it does not correspond to a monetary unit such as USD or EUR. The day-ahead grid-purchase term uses the dataset-provided day-ahead import-price profile, scaling factor of 1.0. Because the model imposes $P_G \geq 0$, active-power export is excluded, and there is no feed-in tariff or export-revenue term.

It should be emphasized that $100P^{cap}$ and $20E^{cap}$ in the capacity term are engineering capacity regularization coefficients rather than actual $/MW or $/MWh investment costs; BESS-throughput is a term that only suppresses operating throughput and does not constitute a battery degradation, capacity-fade, or replacement-cost model. Actual BESS capital cost and annualized ownership cost are evaluated separately through the post hoc economic audit in Section 3.3.13 and do not participate in SG-ENAR-OPF siting, sizing, or candidate ranking.

3.2. Experimental Design and Validation Logic

The experiments follow the methodological workflow proposed in Section 2. First, under the base operating condition of the IEEE-33-bus system, candidate generation is conducted, and SG-ENAR-OPF is solved within the resulting main candidate set to determine the BESS installation location and capacities.

The BESS configuration is then fixed, and a day-ahead operating policy is generated under the Section 2.5-defined finite-scenario scheduling framework. The policy contains a baseline schedule and shared affine-response coefficients and is subjected to post hoc execution validation on the actual trajectory.

The following experiments are designed to evaluate the proposed method comprehensively:

(1) Candidate-generation validation: The sensitivity-guided candidate set is compared with the full-node audit and other candidate-selection benchmarks to evaluate result retention after candidate-space reduction.

(2) BESS planning validation: Changes in storage configuration under different duration constraints and evaluation settings are analyzed, and the stability of the final BESS location and capacity selection is tested.

(3) Finite-scenario scheduling validation: The execution performance of the fixed day-ahead policy under different uncertainty disturbances is evaluated through finite-scenario ablation, temporal stress-path comparison, and actual rolling replay.

(4) Second-feeder supplementary validation: The complete workflow is repeated on the IEEE-69-bus radial feeder, and an independent economic audit is conducted to evaluate execution performance under network changes and the economic scale of the configuration.

3.3. Experimental Results and Analysis

3.3.1. Full-Node Sensitivity Analysis and Main-Candidate Construction

To validate the Section 2.2-defined sensitivity-guided candidate generation method, nodal sensitivity analysis is first conducted for the IEEE-33-bus system under the base operating condition. Comparing nodal rankings under different sensitivity metrics reveals how system-level regulation capability, target-area support capability, and network-structure factors affect candidate-bus identification. Table 5 lists the main high-ranking buses under different sensitivity metrics and their coverage by the main candidate set.

Table 5. Full-node sensitivity rankings and coverage by the main candidate set.

Ranking Type	Top 6 Buses	Best Bus and Value	Top 6 Coverage	Top 10 Coverage	Uncovered Top 10 Buses
Comprehensive P-side sensitivity	Buses 18, 17, 16, 15, 14, 13	Bus 18: 0.8722	100%	100%	None
Comprehensive Q-side sensitivity	Buses 18, 17, 33, 32, 31, 16	Bus 18: 0.7809	100%	100%	None
Target-area P support	Buses 33, 32, 31, 30, 29, 28	Bus 33: 0.04543	100%	80%	Buses 27 and 26
Target-area Q support	Buses 33, 32, 31, 30, 29, 28	Bus 33: 0.03563	100%	80%	Buses 27 and 26

Table 5 shows that different sensitivity metrics favor different buses. The system-level composite sensitivity mainly identifies buses with potential for system-wide voltage improvement and loss reduction, whereas target-area sensitivity is more concentrated in the weak-voltage area near Bus 31, Bus 32, and Bus 33. Because a single metric reflects only one aspect of network operation, a multi-role candidate generation strategy is used to form the final candidate set. The final main candidate set contains 15 buses:

$$\mathcal{N}_{cand} = \{18, 17, 16, 15, 14, 33, 32, 31, 12, 30, 11, 13, 10, 29, 28\} \tag{122}$$

Of these, eight buses are sensitivity-led candidates and seven are deterministic safeguard candidates. The former use Jacobian-assisted pre-screening to reduce the number of buses requiring AC finite-difference refinement, while the latter supplement buses with potential value in feeder structure, weak-area support, and network-stress relief. It is important that the candidate generation stage only reduces the search range for subsequent SG-ENAR-OPF solution and does not directly determine the final BESS installation location. All candidate buses are optimized using the same siting-and-sizing model, and the final BESS configuration is determined by the SG-ENAR-OPF defined in Section 2.3.

To further evaluate retention after candidate-set reduction, the proposed method is compared with direct full-node search, minimum-voltage sensitivity Top 15, loss sensitivity Top 15, and the electrical-distance-based heuristic. Each method generates its candidate set independently, without adjustment using the final planning results.

Table 6 shows that the different candidate-selection rules have distinctly different bus-retention characteristics. The minimum-voltage sensitivity Top 15 retains system-level voltage-related buses relatively well, but its target-area Top 5 coverage is low and, in the full-node planning benchmark, it does not include the final preferred bus Bus 30. In contrast, the loss-sensitivity Top 15 performs strongly under the current base operating point and completely covers the final planning Top 5, indicating that a single metric can still yield an effective candidate set under particular operating conditions. This study therefore does not claim that the proposed method is strictly superior to every other rule in all single-metric comparisons.

The electrical-distance-based heuristic further illustrates the bias of a single topology rule. Electrical-near-target covers target-area buses relatively well but lacks system-level regulation capability, whereas electrical-far-slack retains some system-level sensitivity buses but provides insufficient target-area support. In comparison, the main candidate set simultaneously maintains complete coverage of the system-level P/Q composite-sensitivity Top 5 and the target-area P/Q Top 5 and, in the full-node planning benchmark, retains the final Top 5 buses.

Regarding computational burden, because all baselines are post-processed from the same full-node sensitivity matrix or fixed topology information, their wall-clock time is not compared as an independent end-to-end runtime. Instead, AC power-flow call count and candidate-set size are used as the primary computational indicators. The measured time for the complete 32-bus AC-FD reference sensitivity calculation is approximately 2.408 s, and the total time including construction and comparison of all equal-budget baselines in the benchmark suite is approximately 2.830 s.

In terms of computational scale, the proposed method performs AC finite-difference refinement only for eight sensitivity-led core buses and adds seven buses requiring no additional AC-FD calculations, namely the deterministic safeguard buses. Consequently, the number of nodal AC-FD calls decreases from a complete 32-bus analysis with 65 calls to 17, a reduction of 73.85%. Because the other baselines are post-processed from the same full-node sensitivity matrix, AC power-flow call count and candidate-set size are used as the main computational-burden indicators rather than separating shared post-processing time into independent wall-clock comparisons.

Table 6. Comparison of the proposed candidate-generation method with conventional sensitivity and topology-based screening baselines in the IEEE-33 system.

Method	Candidate-Bus Set	Number of Candidates	Nodal AC-FD Call Count	Reduction Relative to Complete AC-FD	Retention of the Best Buses Under the Four Sensitivity Categories	P Top 5 Coverage	Q Top 5 Coverage	Target-P Top 5 Coverage	Target-Q Top 5 Coverage	Retention of Final Full-Node Planning Top 1 (Bus 30)	Final Planning Top 5 Coverage	Stability Across Operating Points
Direct full-node benchmark	Buses 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33	32	65	0%	Yes	1.00	1.00	1.00	1.00	Yes	1.00	Reference
Minimum-voltage-sensitivity Top 15	18, 17, 16, 15, 14, 13, 12, 11, 10, 9, 8, 7, 33, 32, 31	15	65	0%	Yes	1.00	1.00	0.60	0.60	No	0.40	Not audited separately
Loss-sensitivity Top 15	33, 32, 31, 18, 17, 16, 15, 30, 14, 13, 29, 12, 11, 10, 28	15	65	0%	Yes	1.00	1.00	1.00	1.00	Yes	1.00	Not audited separately
Electrical-near-target heuristic	31, 32, 33, 30, 29, 28, 27, 26	8	0	100%	No	0.00	0.60	1.00	1.00	Yes	1.00	Not audited separately
Electrical-far-slack heuristic	18, 17, 16, 15, 14, 13, 33, 32	8	0	100%	Yes	1.00	0.80	0.40	0.40	No	0.20	Not audited separately
Proposed main candidate set	18, 17, 16, 15, 14, 33, 32, 31, 12, 30, 11, 13, 10, 29, 28	15	17	73.85%	Yes	1.00	1.00	1.00	1.00	Yes	1.00	14/14 Top 5 coverage

3.3.2. Sensitivity-Screening Stability Under Weight, Operating-Point, and Finite-Difference Variations

To evaluate the stability of sensitivity-guided candidate generation under different evaluation preferences, operating conditions, and finite-difference settings, a robustness audit of ranking stability is further conducted. This section examines the ability of the candidate set to retain high-sensitivity buses; it does not require a single base-point sensitivity ranking to remain identical under every operating condition.

First, for weight selection during candidate-screening, six fixed alternative weight vectors are specified in addition to the default weights, including equal weights, single-metric-priority settings, and target-area-enhanced settings. The resulting bus rankings on the P/Q side for the Top 5 are shown in Table 7.

Table 7. Bus-ranking results under fixed-weight schemes.

Weight Scheme	ω_1	ω_2	ω_3	ω_4	P-Side Top 5	Q-Side Top 5
Default weights	0.35	0.25	0.20	0.20	18, 17, 16, 15, 14	18, 17, 33, 32, 31
Equal weights	0.25	0.25	0.25	0.25	18, 17, 16, 15, 14	33, 32, 31, 18, 17
Voltage priority	0.55	0.15	0.15	0.15	18, 17, 16, 15, 14	18, 17, 16, 15, 14
Loss priority	0.15	0.55	0.15	0.15	18, 17, 16, 15, 14	33, 32, 31, 30, 18
Branch-flow priority	0.15	0.15	0.55	0.15	18, 17, 16, 15, 14	33, 32, 31, 30, 18
Target-area priority	0.15	0.15	0.15	0.55	33, 32, 31, 30, 18	33, 32, 31, 30, 29
Strong target-area priority	0.10	0.10	0.10	0.70	33, 32, 31, 30, 29	33, 32, 31, 30, 29

As shown in Table 7, weight changes mainly affect the ranking relationship between system-level regulation buses and target-area support buses. For the P-side metrics, the regions near Bus 18, Bus 17, Bus 16, Bus 15, and Bus 14 remain highly ranked under most weight combinations. When the target-area weight is increased, the ranking gradually shifts toward the regions near Bus 33, Bus 32, Bus 31, Bus 30, and Bus 29. The Q-side ranking is more sensitive to weight changes but remains concentrated primarily in the two key groups of system-level and target-area buses.

To further reduce dependence on fixed weight combinations, 1000 random nonnegative normalized weight vectors are generated with the fixed random seed 20260714, and the frequencies with which each bus enters the Top 1, Top 3, Top 5, and Top 10 are calculated.

As shown in Table 8, key buses continue to have high occurrence frequencies under random weights. For example, the frequencies with which Bus 18 enters the Top 1, Top 5, and Top 10 on the P side are 0.743, 0.973, and 1.000, respectively; the frequencies with which Bus 33 enters the Top 1 and Top 5 on the Q side are 0.640 and 0.882, respectively. These results indicate that the key buses identified by front-end sensitivity screening do not arise solely from the default weights but remain stable over a broad range of metric preferences. The coverage of high-ranking buses under random weights by the main candidate set is further evaluated.

Table 8. Stability of sensitivity-screening buses under random four-index weights.

Type	Bus	Top 1 Frequency	Top 3 Frequency	Top 5 Frequency	Top 10 Frequency
P-side system-level bus	Bus 18	0.743	0.793	0.973	1.000
P-side system-level bus	Bus 17	0	0.733	0.915	1.000
P-side target-area bus	Bus 33	0.257	0.376	0.496	0.988
Q-side target-area bus	Bus 33	0.640	0.795	0.882	0.997
Q-side target-area bus	Bus 32	0	0.666	0.832	0.989
Q-side system-level bus	Bus 18	0.360	0.427	0.703	1.000

The results in Table 9 show that, for the main candidate set on both the P/Q sides, the Top 1 through Top 5 buses achieve 100% coverage, while the average Top 10 coverage rates reach 99.88% and 99.96%, respectively. Therefore, the proposed candidate generation method retains key buses under different sensitivity-evaluation preferences.

Table 9. Main-candidate coverage under random four-index sensitivity weights.

Type	Top-K	Fraction for Which All Top-K Buses Are in the Main Candidate Set	Average Coverage
P side	Top 1	1.000	1.000
P side	Top 3	1.000	1.000
P side	Top 5	1.000	1.000
P side	Top 10	0.988	0.9988
Q side	Top 1	1.000	1.000
Q side	Top 3	1.000	1.000
Q side	Top 5	1.000	1.000
Q side	Top 10	0.996	0.9996

It should be noted that the preceding experiments assess, in the candidate-screening stage, the effects of the four sensitivity-metric weights ω_1–ω_4 and are not equivalent to the sensitivity analysis of the RobustScore weights λ_1–λ_8 in the final BESS-siting stage. The former determine the candidate space, whereas the latter affect the final planning ranking within that space; their roles are distinct.

3.3.3. Robustness Audit of Sensitivity-Based Candidate Screening

Although the sensitivity-guided candidate-generation method is constructed from a reference operating state, changes in operating states and numerical settings may occur during actual distribution-system operation. Therefore, this section further analyzes the stability of the sensitivity-screening process with respect to changes in operating point, finite-difference form, and perturbation magnitude.

It should be emphasized that this section does not require the bus-sensitivity rankings to remain completely identical under different operating conditions. Instead, it verifies whether the pregenerated candidate set consistently retains key buses with regulation value under different operating states and numerical settings.

(1) Effect of operating-state changes on sensitivity ranking:

The effect of operating-state changes on the bus-sensitivity ranking is first analyzed. Seven representative operating states are selected: the original base state (Original base), actual median state (Median actual), peak-load state (Peak load), high-photovoltaic state (High PV), high-wind state (High wind), minimum-import state (Minimum import), and synthetic reverse-flow stress state (Synthetic reverse-flow).

With the default sensitivity weights and the 0.010 MW/MVAr forward finite difference setting retained, the sensitivity rankings under different operating points are compared, and the results are presented in Table 10.

As shown in Table 10, the sensitivity ranking exhibits a degree of operating-state dependence. Under the original base, actual median, peak-load, high-wind, and minimum-import states, the P-side Top 1 bus remains Bus 18, and the key high-ranking buses are still concentrated mainly in the original critical areas. Under the high-photovoltaic-output state, however, the change in renewable-energy output modifies the system power-flow distribution, and the P- and Q-side Top 1 buses both shift to Bus 33, indicating that bus-sensitivity rankings can be affected by changes in operating state. Accordingly, the sensitivity ranking obtained at the base operating point is not interpreted as a fixed bus-priority order that holds under all operating conditions. The focus is instead on whether the candidate set covers buses with high regulation value under different operating states.

Table 10. Operating-point dependence of sensitivity ranking.

Operating Point	P-Side Top 1	Q-Side Top 1	P-Side Rank Correlation	Q-Side Rank Correlation	Remark
Original base	Bus 18	Bus 18	1.0000	0.99963	Reference
Median actual	Bus 18	Bus 18	1.0000	0.99963	Stable
Peak load	Bus 18	Bus 18	1.0000	1.0000	Stable
High PV	Bus 33	Bus 33	1.0000	1.0000	Ranking rotation
High wind	Bus 18	Bus 18	1.0000	1.0000	Stable
Minimum import	Bus 18	Bus 18	1.0000	0.99963	Stable
Synthetic reverse-flow	Bus 18	Bus 18	1.0000	1.0000	Stable

To further examine the potential nonsmoothness of the extremum-based metrics $V_{\min}$ and $S_{\max}$, the limiting-bus and limiting-branch identities are also recorded under each representative operating point. The results show that the active-extremum identity changes with the operating condition: the minimum-voltage limiting bus shifts from Bus 18 under the base condition to Bus 33 under the high-PV condition, and the maximum-apparent-flow branch shifts from Branch 1 (1–2) under most operating states to Branch 24 (24–25) under the synthetic reverse-flow stress condition. These observations indicate that the extremum identity does not remain fixed across operating states. Therefore, the sensitivities corresponding to $V_{\min}$ and $S_{\max}$ are not interpreted as globally smooth derivatives over the entire operating domain.

(2) Consistency analysis of forward and central finite differences:

Because the four sensitivity metrics used in this study are calculated through AC finite differences, the effect of different difference forms on bus ranking is further examined. With the perturbation magnitudes fixed at $\Delta P = 0.010$ MW and $\Delta Q = 0.010$ MVAr, the rankings obtained using forward finite difference and central finite difference are compared.

As shown in Table 11, forward and central finite differences yield highly consistent sensitivity rankings. Under all tested operating states, the Top 1 buses obtained using the two difference forms are identical. The minimum rank-correlation coefficient is 0.99963, and the minimum Top 5 overlap is 0.80, indicating that forward finite differences closely approximate the ranking obtained by central finite differences over the tested range while avoiding the calculation of reverse perturbations. Therefore, the forward finite difference used in this study provides a suitable balance between computational efficiency and numerical reliability for candidate-bus screening.

Table 11. Forward and central finite-difference comparison.

Operating Condition	Direction	Delta (MW/MVAr)	Top 1 Match	Top 5 Overlap	Rank Correlation
Original base	P	0.010	1	1.00	1.00000
Original base	Q	0.010	1	0.80	0.99963
Median actual	P	0.010	1	1.00	1.00000
Median actual	Q	0.010	1	0.80	0.99963
High PV	P	0.010	1	1.00	1.00000
High PV	Q	0.010	1	1.00	1.00000
Synthetic reverse-flow	P	0.010	1	1.00	1.00000
Synthetic reverse-flow	Q	0.010	1	1.00	1.00000

(3) Analysis of the effect of the finite-difference perturbation magnitude:

The effect of changes in the finite-difference perturbation magnitude on sensitivity-ranking stability is further analyzed. Perturbation magnitudes of $\Delta P, \Delta Q = 0.005$

MW/MVAr, $\Delta P, \Delta Q = 0.010$ MW/MVAr, and $\Delta P, \Delta Q = 0.020$ MW/MVAr are tested, and changes in the bus rankings under different settings are quantified.

As shown in Table 12, changes in the finite-difference step over the perturbation range of 0.005–0.020 MW/MVAr do not materially affect sensitivity-ranking stability. Across the P/Q rankings for all 14 operating-state cases, no change in the Top 1 bus is observed, and all rank-correlation coefficients remain above 0.99963. Thus, within the tested range, the selected finite-difference perturbation scale does not materially affect identification of the key buses.

Table 12. Sensitivity ranking robustness under different finite-difference step sizes.

Perturbation Step ΔP, ΔQ (MW/MVAr)	Number of Tested Conditions	Top 1 Mismatch Count	Minimum Rank Correlation	Minimum Top 5 Overlap	Overall Assessment
0.005	14	0	1.00000	1.00	Stable
0.010	14	0	1.00000	1.00	Reference setting
0.020	14	0	0.99963	0.80	Stable

In addition, whether the local finite-difference calculation triggers limiting-element switching is examined for different perturbation magnitudes. Over the tested perturbation range of 0.005, 0.010, and 0.020 MW/MVAr, no perturbation-induced switching of the limiting bus or limiting branch is recorded. Thus, within the present test range, the local AC finite-difference ranking is not materially affected by active-extremum switching. This conclusion is limited to the operating points and perturbation range tested in this study and does not constitute a guarantee that these extremum-based metrics are globally smooth.

In summary, although the sensitivity-guided candidate-screening results are affected by changes in operating state and evaluation-parameter preferences, their primary purpose is not to guarantee that a single sensitivity ranking remains identical under all conditions, but to retain coverage of key buses under different operating states and numerical settings. The preceding ranking-stability audit shows that the generated candidate set can accommodate changes in operating point and finite-difference settings within the tested ranges, thereby providing a reliable candidate space for subsequent SG-ENAR-OPF optimization. The final BESS siting and sizing results are determined jointly by the subsequent optimization model rather than directly by a single sensitivity ranking.

3.3.4. Jacobian-Assisted AC Finite-Difference Refinement

To determine an appropriate scale for the Jacobian-assisted AC finite-difference refinement, the effects of different system-level sample sizes on k, computational cost, and key-bus coverage are analyzed. The target-area refinement depth $k_t = 3$ is fixed to avoid missing local support buses near the undervoltage-sensitive region because of insufficient system-level sampling. Therefore, the analysis focuses on selecting the system-level refinement k, rather than redesigning the target-area refinement rule.

As shown in Table 13, as k increases, the number of AC power-flow evaluations increases linearly, while key-operating-bus coverage gradually improves. When $k = 3$ or $k = 5$, the computational cost remains low, but insufficient sampling may prevent high-value buses from entering the candidate set.

Table 13. AC finite-difference cost and critical-bus coverage of Jacobian-assisted refinement under different truncation depths.

k	AC-FD Refined Buses	AC PF Runs for FD Refinement	Reduction vs. Full AC-FD	Hit All Four Key Sensitivity Categories	P Top 5 Coverage	Q Top 5 Coverage	Target P Top 5 Coverage	Target Q Top 5 Coverage
3	5	11	83.08%	Yes	0.40	1.00	0.60	0.60
5	6	13	80.00%	Yes	0.60	1.00	0.60	0.60
8	8	17	73.85%	Yes	1.00	1.00	0.60	0.60
10	10	21	67.69%	Yes	1.00	1.00	0.80	0.80
15	15	31	52.31%	Yes	1.00	1.00	1.00	1.00
20	20	41	36.92%	Yes	1.00	1.00	1.00	1.00

When k increases to eight, only 17 AC power-flow evaluations are required. At this setting, all four key sensitivity categories are represented, the system-level P- and Q-side Top 5 coverage both reach 1.00, and the target-area P- and Q-side Top 5 coverage both reach 0.60. Compared with AC finite-difference evaluation of all 32 nonslack buses, which requires 65 evaluations, the computational burden is reduced by 73.85%.

When k is further increased to 10, 15, or 20, the computational cost continues to rise, while the additional improvement occurs mainly in target-area Top 5 coverage, which increases from 0.60 at $k = 8$ to 0.80 at $k = 10$ and 1.00 at $k \geq 15$. Because the target-area refinement depth $k_t = 3$ is handled independently to retain the main local support buses, $k = 8$ is adopted as a balance between AC finite-difference cost, complete system-level P/Q Top 5 retention, and representation of the four key sensitivity categories. At this setting, $\mathcal{N}_{fast} = \{18, 17, 33, 32, 31, 16, 15, 14\}$. Figure 3 further illustrates the tradeoff between computational cost and coverage under different k.

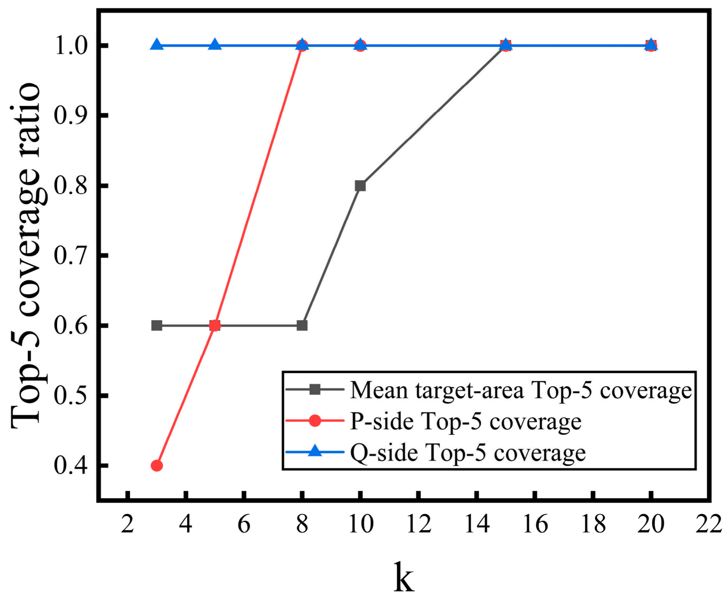

Figure 3. AC finite-difference power-flow cost and critical-bus coverage of Jacobian-assisted refinement under different truncation depths.

3.3.5. Main-Candidate OPF Siting and Sizing and Full-Node Audit

To further examine the effect of candidate-space compression on the final BESS-planning results, the eligible-bus space is subjected to a full-node audit, and the effects of different duration constraints and RobustScore evaluation preferences on the final planning results are analyzed. It should be emphasized that sensitivity screening is used to construct a representative candidate space, whereas the final BESS bus is determined comprehensively by the SG-ENAR-OPF optimization model. Accordingly, the analysis focuses on whether the candidate set retains the high-value buses identified in the final planning results.

(1) Final BESS configurations under different duration constraints:

First, a complete search is performed over the eligible bus set under the MAIN and STRICT duration-constraint settings. The final BESS siting and sizing results are presented in Table 14.

Table 14. Final duration-aware BESS configurations for the IEEE-33-bus system.

Case	Duration Range	Selected Bus	P^{cap}/MW	E^{cap}/MWh	E^{cap}/P^{cap}/h
MAIN	1–4 h	Bus 30	10.6613	10.6613	1.000
STRICT	2–4 h	Bus 30	10.5782	21.1565	2.000

As shown in Table 14, the final siting result is Bus 30 under both the MAIN and STRICT duration constraints, indicating that this bus has high composite planning value under different energy-duration requirements. At the same time, the STRICT constraint mainly increases the energy-capacity requirement, whereas the power-rating change is small, consistent with the effect of a longer duration constraint on the energy-allocation requirement.

(2) Analysis of composite metrics for the top candidate buses in the full-node audit:

To further explain the final bus selection, the original engineering metrics of the Top5 buses in the full-node audit—Bus 30, Bus 28, Bus 29, Bus 31, and Bus 32—are compiled in Table 15, which presents, under the MAIN and STRICT duration constraints, the eight ranking metrics for each candidate bus.

Table 15. Raw values of all eight ranking indices for the top five IEEE-33 buses in the final full-node audit.

Case	Rank	Bus	Weighted Grid Energy/MWh	Weighted Curtailment/MWh	Weighted Voltage-Deviation Index	Worst Minimum Voltage/p.u.	Weighted Recourse-Power Index	Weighted Load Loss/MWh	P^{cap}/MW	E^{cap}/MWh
MAIN	1	30	74.069068	0	14.737337	0.966169	0	0	10.661342	10.661342
MAIN	2	28	74.199935	0	8.352843	0.978355	0	0	14.120905	14.120905
MAIN	3	29	74.119007	0	13.321051	0.968826	0	0	11.360688	11.360688
MAIN	4	31	74.035416	0	18.255520	0.963251	0	0	9.819268	9.819268
MAIN	5	32	74.035466	0	18.969279	0.963230	0	0	9.804441	9.804441
STRICT	1	30	74.069643	0	14.843782	0.965909	0	0	10.578232	21.156463
STRICT	2	28	74.195013	0	8.511993	0.977991	0	0	13.999700	27.999400
STRICT	3	29	74.116292	0	13.403021	0.968727	0	0	11.298370	22.596740
STRICT	4	31	74.037006	0	18.380100	0.963169	0	0	9.725349	19.450698
STRICT	5	32	74.032823	0	19.129773	0.962911	0	0	9.682118	19.364236

As shown in Table 15, Bus 30 ranks first under both the MAIN and STRICT constraints. It should be noted, however, that Bus 30 is not superior in every individual metric. For example, Bus 28 has a weighted voltage deviation of 8.352843 and a minimum voltage of 0.978355 p.u., both better than those of Bus 30, indicating stronger local voltage-support capability. However, Bus 28 requires a BESS rated at 14.120905 MW/14.120905 MWh, whereas Bus 30 requires 10.661342 MW/10.661342 MWh; Bus 28's power and energy capacities are therefore approximately 24.50% larger. Bus 28's weighted losses are also slightly higher than those of Bus 30. In contrast, although Bus 31 and Bus 32 require smaller storage capacities than Bus 30, their weighted voltage deviations increase to 18.255520 and 18.969279, respectively, with corresponding reductions in minimum voltage. Thus, no single bus is superior across all engineering metrics.

Accordingly, the final ranking is not determined by a single voltage metric but is obtained through a composite tradeoff among voltage improvement, network losses, and the required storage scale.

(3) Sensitivity analysis of RobustScore evaluation preferences:

Because the final bus ranking is determined jointly by eight engineering metrics, the effect of different evaluation preferences on the RobustScore ranking is further analyzed. The weight-sensitivity experiment follows Section 2.3.3 and uses the full-node common-normalization basis to ensure that all metrics for the same bus set use the same normalization scale under different evaluation preferences. To ensure reproducibility, random weight sampling uses the fixed random seed 20260815 to generate 5000 nonnegative weights whose sum is 1. The following weight-change schemes are considered: equal-weight, random weight sampling, exact-zero drop-one, and one-at-a-time (OAT). The results are presented in Table 16.

Table 16. Sensitivity of the final RobustScore ranking to the eight evaluation weights.

Case	Adopted Bus	Equal-Weight Top 1	Adopted-Bus Rank Under Equal Weights	Random Top 1	Random Top 3	Random Top 5	Mean Rank
IEEE-33 MAIN	Bus 30	Bus 13	14	0.1124	0.3668	0.4502	13.0826
IEEE-33 STRICT	Bus 30	Bus 4	14	0.1110	0.3416	0.4286	13.7342
IEEE-69 MAIN	Bus 19	Bus 3	37	0.0566	0.0846	0.1274	28.2434
IEEE-69 STRICT	Bus 18	Bus 3	38	0.0218	0.0728	0.1232	29.9120

Table 16 summarizes the adopted-bus rankings under different weight settings for the final planning cases. The results indicate that the RobustScore ranking is clearly dependent on the evaluation preference. For IEEE-33 MAIN, Bus 30 ranks first under the preset weights used in this study; across 5000 random-weight samples, however, its Top 1, Top 3, and Top 5 frequencies are 0.1124, 0.3668, and 0.4502, respectively, and its mean rank is 13.0826. Under completely equal weights, Bus 13 becomes the Top 1 bus, whereas Bus 30 falls to 14th. IEEE-33 STRICT exhibits a similar trend, with Bus 4 becoming the Top 1 bus under equal weights.

The exact-zero drop-one and one-at-a-time (OAT) weight-perturbation results further illustrate the effects of individual metrics on the final ranking. When any one engineering metric is completely removed, the adopted bus remains Top 1 under some, but not all, perturbations. When each weight is multiplied by 0.5 or 1.5 and then renormalized, the adopted buses in the four cases remain Top 1 under 10, 10, nine, and eight of the 16 perturbations, respectively.

These perturbation experiments indicate that RobustScore is not a fixed physical ranking independent of evaluation preferences; rather, it is a composite evaluation tool that explicitly reflects the decision-maker's tradeoffs among multiple engineering objectives. Therefore, the final BESS-siting results should be interpreted under clearly specified engineering objectives and weight settings and should not be treated as a unique optimum that remains invariant over all possible weight combinations.

3.3.6. Nonlinear AC Post-Validation of the Final Fixed Policies

To assess the actual operating safety of the pre-dispatch policy under a nonlinear power-flow model rather than only under LinDistFlow, AC post-validation is further conducted. Specifically, the IEEE-33 MAIN case, its Bus 30 BESS configuration, and 46 complete dispatch periods are considered. The fixed BESS charging/discharging injections at the

1104 actual execution-time points are mapped to the MATPOWER AC nonlinear model, and nonlinear AC power-flow validation is then performed directly.

It should be emphasized that the AC post-validation stage does not re-optimize the dispatch, does not clip the active/reactive BESS outputs, and directly uses the fixed-policy injection determined in the preceding stage. Therefore, this process evaluates the actual performance of the pre-dispatch policy on the tested actual operating trajectories rather than an outcome corrected after re-optimization.

For a fair comparison, the same Bus 30 BESS configuration is also evaluated. The q^{50} forecast median and deterministic dispatch are used to construct the baseline. The two policies use the same BESS installation location, capacity, and AC operating conditions and differ only in their treatment of uncertainty in the preceding stage.

As shown in Table 17, across the 1104 actual operating points, both fixed policies achieve 100% AC power-flow convergence, and all operating points satisfy the branch-flow constraints. Their maximum branch-loading ratios are 0.843779 and 0.835306, respectively, both below the rated limit of 1.0, and no branch overload occurs. In terms of voltage constraints, both policies still have operating points with low-voltage violations, but the uncertainty-aware policy performs slightly better. It has 1058/1104 AC-feasible operating points, corresponding to 95.833%, compared with q50 baseline at 1057/1104 and 95.743%. The infeasible points under both policies arise from voltage deviations rather than branch-flow limits. The minimum voltage under the uncertainty-aware policy is 0.867314 p.u., compared with q^{50} baseline at 0.865004 p.u., and the corresponding maximum voltage violation decreases from 0.034996 p.u. to 0.032686 p.u.

Table 17. Nonlinear AC post-validation of the q^{50} baseline and uncertainty-aware fixed policy under the final IEEE-33 MAIN Bus 30 BESS configuration.

Index	Uncertainty-Aware Policy	Robust Policy
AC power-flow convergence	1104/1104	1104/1104
AC feasible operating points	1057/1104	1058/1104
AC feasibility rate	95.743%	95.833%
Worst minimum voltage/p.u.	0.865004	0.867314
Worst maximum voltage/p.u.	1.034448	1.033676
Maximum voltage violation/p.u.	0.034996	0.032686
Maximum AC branch loading/p.u. of rating	0.843779	0.835306
Branch-feasible operating points	1104/1104	1104/1104
Maximum branch overload/p.u.	0	0
Mean AC voltage deviation/p.u.	0.015945	0.015877
Total AC active-loss energy/MWh	451.002698	444.960142
Active load shedding/MWh	0	0
Reactive load shedding/MVArh	0.683825	0
Renewable curtailment/MWh	0.554326	0
Common actual composite objective	3322.156435	3172.520630

A further comparison of the engineering metrics under the common actual trajectories shows that the uncertainty-aware policy provides better overall operating performance. Its cumulative AC active-power loss is 444.960142 MWh, compared with q^{50} baseline at 451.002698 MWh, a reduction of approximately 1.34%. Under the same engineering-evaluation basis, the common actual composite objective decreases from 3322.156435 to 3172.520630, a reduction of approximately 4.50%. In addition, the uncertainty-aware policy incurs no active/reactive load shedding or renewable-energy curtailment over all 1104 validation periods, whereas the q^{50} baseline incurs 0.683825 MVArh of reactive load shedding and 0.554326 MWh of renewable-energy curtailment.

AC post-validation further shows that day-ahead feasibility in the LinDistFlow model is not equivalent to hourly feasibility in the nonlinear AC network. Although the day-ahead policy remains feasible under the LinDistFlow constraints, after mapping to the MATPOWER AC model, low-voltage deviations caused by linear-network approximation errors remain in some periods. Accordingly, AC power-flow calculation is treated as an independent nonlinear-network audit to identify the applicability boundary between the linear optimization model and the actual AC network. Meanwhile, the baseline comparison in this section uses the same Bus 30 BESS location and capacity and compares only the q^{50} deterministic policy with the uncertainty-aware shared policy, to isolate the effects of uncertainty treatment and policy design on actual operation. Therefore, in Table 17, the common actual composite objective, AC active-loss energy, renewable curtailment, and load-shedding differences should all be interpreted as operating-performance differences under fixed hardware conditions, rather than the overall investment return of installing a BESS relative to having no BESS or using other lower-cost regulation resources. This study does not place OLTC/tap adjustment, independent VAR compensation, targeted curtailment-only control, or smaller-BESS re-optimization within a unified lifecycle-cost framework. It therefore does not claim that the current BESS configuration has a general cost-benefit advantage over all alternatives; the approximate hardware-investment scale is discussed separately in Section 3.3.13 within the post hoc ownership-cost audit.

Overall, the AC post-validation results indicate that, with the BESS configuration fixed, the uncertainty-aware policy operating trends consistent with the day-ahead optimization results in the nonlinear network and achieves lower network losses, fewer constraint violations, and better overall operating performance than the q^{50} baseline. At the same time, this validation clearly identifies the applicability boundary between the linear network approximation and the actual AC network, providing a basis for future research on AC-aware closed-loop correction strategies.

3.3.7. Interval-Forecast Coverage and Standardized-Disturbance Diagnostics

To examine the relationship between the prediction intervals used in Section 2.5 and actual operating deviations, prediction-interval coverage and standardized-disturbance diagnostics are conducted over 46 rolling windows comprising 1104 nonoverlapping actual execution hours. All day-ahead disturbance-space scenarios and temporal stress paths for the corresponding test-period are constructed solely from the $q^{05}/q^{50}/q^{95}$ forecasts; the actual observations are not used in scenario construction or day-ahead policy optimization and are used only for post hoc statistics and fixed-policy validation after the policy is frozen. Therefore, this section evaluates chronological out-of-sample performance over the forecast test period and does not interpret this sequence as an independent later-period holdout.

First, the proportions of the actual load, PV, WT, and their combined quantities that fall within the q^{05}–q^{95} prediction intervals are calculated. PV–wind joint means that PV and WT are both within their respective prediction intervals at the same time, whereas All-component joint means that load, PV, and WT are all within their respective intervals simultaneously. The results are presented in Table 18.

Table 18 reports the empirical coverage of the actual trajectories relative to the q^{05}–q^{95} prediction intervals. It should first be clarified that the q^{05}–q^{95} values used in this study are the lower and upper quantile boundaries output by Quantile-TCN. Their purpose is to construct the day-ahead uncertainty description and scheduling scenarios, not to presume that the test-period actual samples must achieve an empirical hit rate of 90%. Therefore, the overall coverage reported in the table measures the actual agreement between the prediction intervals and the actual operating trajectories and should not be interpreted directly as "forecast accuracy."

Table 18. Prediction-interval coverage statistics for actual trajectories.

Index	Covered Points	Out-of-Interval Points	Overall Coverage	Worst Window	Worst-Window Coverage
Load	928	176	84.06%	WIN_003	50.00%
PV	909	195	82.34%	WIN_018	50.00%
WT	846	258	76.63%	WIN_004	50.00%
Net load	1026	78	92.93%	WIN_042	70.83%
PV–wind joint	708	396	64.13%	WIN_009	37.50%
All-component joint	604	500	54.71%	WIN_003	29.17%

For the individual variables, the empirical coverage rates of Load, PV, and WT are 84.06%, 82.34%, and 76.63%, respectively, all below the nominal 90% interval level, with the largest deviation observed for WT. Further inspection shows that, for Load and PV, the miscoverage occurs mainly when the actual value exceeds q^{95}, i.e., as an upper-tail exceedance, whereas for WT it occurs mainly when the actual value falls below q^{05}, i.e., as a lower-tail exceedance, indicating that interval mismatch has different directions for different forecast variables. By contrast, Net load coverage reaches 92.93%, exceeding the nominal 90% interval level. This indicates a degree of cancellation after load and renewable forecast errors are aggregated through power balance, so the final net-load trajectory agrees with its prediction interval more closely than do the individual components. Because BESS dispatch responds directly to the system's net power imbalance, this result is more representative of the combined disturbance faced in actual dispatch than the marginal coverage of any single source. It should be emphasized that the test-period actual results are not used for post hoc recalibration, nor are these coverage statistics used to modify the $q^{05}/q^{50}/q^{95}$ forecast boundaries retrospectively. Accordingly, the q^{05}–q^{95} intervals serve as finite-scenario construction empirical forecast inputs and are not treated as providing a finite-sample coverage guarantee for a calibrated uncertainty set.

PV–wind joint and All-component joint have coverage rates of 64.13% and 54.71%, respectively. Here, joint coverage requires multiple components to fall within their respective prediction intervals simultaneously, a criterion substantially stricter than univariate marginal coverage. It therefore cannot be compared directly with an individual q^{05}–q^{95} interval, whose nominal level is 90%. The reduction in joint coverage instead indicates that relying only on marginal quantile intervals is insufficient to fully characterize joint realizations of multisource forecast errors. This motivates the use of multiscenario scheduling and subsequent actual fixed-policy validation in this study.

Table 18 also reports worst-window coverage to examine performance in locally difficult periods rather than only on average. In the worst windows, the coverage rates of Load, PV, and WT all decrease to 50%, while All-component joint reaches a minimum of 29.17%. These extreme windows are not representative of the typical level across all 46 rolling windows; they deliberately expose the operating segments with the most pronounced mismatch between the prediction intervals and actual trajectories, thereby providing stricter stress cases for subsequent fixed-policy validation. In other words, low-coverage windows are retained rather than excluded to improve the statistics, and the dispatch policy continues to be tested on these difficult samples.

To further quantify the magnitude by which actual values deviate from the prediction intervals, Table 19 presents the distributions of the unclipped standardized-disturbance coordinates and actual policy inputs. Here, ξ^{raw} is used only to diagnose interval exceedances; the coordinates actually supplied to the frozen affine policy are limited to $[-1, 1]$ in accordance with Section 2.5.3.

Table 19. Distribution of standardized actual-disturbance magnitudes.

Disturbance Coordinate	N	P50 of $\lvert\xi\rvert$	P90 of $\lvert\xi\rvert$	P95 of $\lvert\xi\rvert$	P99 of $\lvert\xi\rvert$	Maximum $\lvert\xi\rvert$	Fraction with $\lvert\xi\rvert$>1
Load, ξ_L^{raw}	1104	0.4213	1.2066	1.5149	2.3010	4.6391	15.94%
PV, ξ_{PV}^{raw}	1047	0.0000	1.1548	1.5634	2.4481	6.6949	13.18%
WT, ξ_{WT}^{raw}	1104	0.5448	1.6314	2.0801	3.0431	6.7175	23.37%
Renewable aggregate, raw diagnostic coordinate	1047	0.3151	0.9440	1.2324	1.7140	4.4021	8.60%
Renewable aggregate, bounded policy input	1104	0.2739	0.7727	0.9340	1.0000	1.0000	0

Table 19 further quantifies the positions of the actual trajectories relative to the prediction intervals in terms of standardized-disturbance magnitude. For Load, PV, and WT, 15.94%, 13.18%, and 23.37% of valid samples, respectively, satisfy $|\xi^{raw}| > 1$, indicating that some actual realization values do exceed the day-ahead disturbance range defined by q^{05}–q^{95}. Exceedance of ξ^{raw} does not indicate failure of the dispatch policy itself; it identifies the intensity of actual disturbances outside the prediction interval and provides the subsequent fixed-policy replay with additional out-of-interval stress tests.

It is worth noting that the capacity-weighted renewable aggregate raw diagnostic coordinate $\xi_{R,t}^{(raw,agg)}$ is calculated only for periods in which both PV and WT raw coordinates are finite, yielding a valid-sample count of 1047; by contrast, after component-level regularization and saturation, the resulting bounded renewable policy input $\xi_{R,t}^{act}$ covers all 1104 actual execution hours. Among valid raw samples, only 8.60% satisfy $\left|\xi_{R,t}^{(raw,agg)}\right| > 1$, which is substantially lower than the individual exceedance rates of PV and WT, and its P90 decreases to 0.9440. This indicates that, at the aggregated renewable-power level corresponding to the actual BESS response, errors from different renewable sources partially cancel, so the combined disturbance faced by the policy is less extreme than that indicated when PV or WT are examined separately. This observation is consistent with the 92.93% net-load coverage reported in Table 18: local departures of individual-source prediction intervals do not propagate proportionally into dispatch deviations at the system net-power level.

On the other hand, the bounded policy coordinate actually supplied to the frozen affine policy always satisfies $|\xi| \leq 1$, and its P99 and maximum value are both 1.0. This bound results from the policy-input mapping defined in Section 2.5.3. Its purpose is not to modify the actual physical trajectories but to ensure that the fixed affine control law is always invoked within its predefined input domain. The actual load, PV, and WT values still enter network operation and post-validation without modification; thus, this treatment does not artificially improve operating feasibility by clipping the physical data.

Taken together, Tables 18 and 19 show that the q^{05}–q^{95} prediction intervals do not completely envelop all actual realizations, but still provide a reference range with a clear quantile interpretation for day-ahead uncertainty modeling. At the same time, net-load and aggregated-renewable coordinates exhibit substantially closer agreement than the individual-source components. More importantly, actual samples outside the intervals are not excluded but are retained for subsequent fixed-policy validation. Accordingly, the role of the prediction-interval coverage statistics in this study is not to demonstrate "90% forecast accuracy," but to identify the gap between the day-ahead uncertainty description and actual operating disturbances and to examine the practical applicability of the proposed policy in the presence of that gap.

Further examination of all 46 d and 1104 h of actual trajectories for the final IEEE-33 MAIN configuration at Bus 30 (10.6613 MW/10.6613 MWh) shows that no BESS power clipping, projection, or re-optimization is used during frozen-policy execution and that all 46 windows pass the storage engineering-feasibility check. The maximum charging and discharging powers are 3.8903 MW and 4.7820 MW, only 36.49% and 44.85% of rated power, respectively; the continuous SOC range is 0.119999889–0.880000100, always within the physical bounds of 0.10–0.90. Thus, even though some raw standardized deviations clearly exceed the q^{05}–q^{95} interval, no negative charging/discharging power or power-capacity exceedance caused by out-of-interval disturbances is observed in the tested 1104 h actual trajectories. This conclusion constitutes only post hoc validation evidence for the current actual trajectory bank and does not imply a continuous-domain feasibility guarantee for arbitrary out-of-interval disturbances.

3.3.8. Sensitivity to Temporal-Path Design

To examine whether the seven temporal stress paths depend strongly on specific artificial duration and path-shape settings, a path-design sensitivity audit is further conducted for the final dispatch model. Three representative windows with different forecast characteristics are selected: WIN008 represents typical operating conditions, WIN019 corresponds to forecast peak-load conditions, and WIN024 corresponds to high forecast renewable output. In addition to the seven-path design defined in Section 2.5, BASELINE7, eight variants are constructed: PEAK_H4, PEAK_H8, SPLIT_H8, SPLIT_H16, ALT_2H, COUNT_3, COUNT_5, and RAMP_12H. These variants perturb the duration of peak stress, sequential switching position, alternating period, number of paths, and ramp shape, respectively. Except for the temporal-path design, all model parameters and operating constraints are unchanged. Day-ahead dispatch is re-solved for each path design over the three representative windows, followed by fixed-policy replay for the corresponding windows, to evaluate both changes in objective and engineering replay performance.

Table 20 shows that the final policy is generally weakly sensitive to the specific path parameters among the tested temporal-path variants. When the peak duration, alternating period, or ramp form is changed, the maximum change in objective does not exceed 0.112%, and the average change does not exceed 0.041%. When the 12/12 h sequential switching pattern is changed to 8/16 h or 16/8 h, the change in objective is only on the order of 10^{-7}%, indicating that this switching-position change has almost no visible effect on the final result. The engineering replay pass rate is 100% for all variants, showing that these path-design changes do not impair engineering replay performance over the tested windows.

Table 20. Sensitivity of the final policy to alternative temporal-path designs.

Variant	Path-Design Modification Relative to BASELINE7	Path Count	Max. Absolute Objective Change/%	Mean Absolute Objective Change/%	Engineering Replay Pass Rate
BASELINE7	6 h peak stress; 12/12 h split; 24 h linear transition; 1 h alternation	7	0	0	100%
PEAK_H4	Peak duration: 6 h → 4 h	7	0.111990	0.040193	100%
PEAK_H8	Peak duration: 6 h → 8 h	7	0.034941	0.014299	100%
SPLIT_H8	Order-switch point: 12 h → 8 h; 8/16 h split	7	3.78×10^{-7}	1.33×10^{-7}	100%
SPLIT_H16	Order-switch point: 12 h → 16 h; 16/8 h split	7	3.17×10^{-7}	1.17×10^{-7}	100%
ALT_2H	Alternation interval: 1 h → 2 h	7	0.041456	0.015510	100%
COUNT_3	Retain paths 1, 2, and 7 only	3	0.411482	0.155640	100%
COUNT_5	Retain paths 1, 2, 3, 4, and 7	5	0.069225	0.028476	100%
RAMP_12H	6 h plateau + 12 h ramp + 6 h opposite plateau	7	0.038224	0.016706	100%

Reducing the number of temporal paths has a more pronounced effect than changing path-shape parameters. When only three paths are retained, the maximum and average changes in objective increase to 0.411482% and 0.155640%, respectively, the largest among all tests. When five paths are retained, the two values decrease to 0.069225% and 0.028476%. This indicates that adding temporal paths improves coverage of different intraday error-evolution patterns, but the marginal effect decreases markedly when the number increases from five to the baseline seven. Therefore, BASELINE7 is not the only feasible path design, but a relatively conservative setting adopted to balance scenario-coverage completeness and computational scale.

Overall, over the ranges of path duration, switching position, alternating period, ramp form, and number of paths tested in this study, the change in the final policy objective remains below 0.5%, and all schemes pass the engineering replay. These results indicate that the principal conclusions of this study are not incidentally driven by one specific set of temporal-path parameters. At the same time, the relatively larger change for COUNT_3 also shows that temporal coverage itself remains important, so path design should not be regarded as a completely irrelevant modeling factor.

3.3.9. Ablation of Disturbance-Space and Temporal-Path Scenario Construction

To further examine the effects of disturbance-space coverage and temporal-path protection on the final dispatch policy, under the final IEEE-33 MAIN Bus 30 configuration, a controlled ablation is conducted. Three representative windows with different operating characteristics are again selected, while the BESS configuration, network model, objective function, SOC ranges, solver settings, and actual-validation protocol are held fixed. The three cases progressively modify the uncertainty-support structure; the final 25-point-grid + seven-temporal-path formulation additionally activates the terminal-energy safeguard associated with the temporal-path design, whereas the three-scenario and 25-point static-grid formulations do not. Three schemes are compared sequentially: three representative scenario points only, a 25-point disturbance-space grid, and a 25-point disturbance-space grid supplemented with seven temporal stress paths and the corresponding terminal-energy safeguard. For each scheme, day-ahead dispatch is re-solved for the corresponding windows, and the resulting policy is frozen and subjected to fixed-policy validation on the actual trajectory from the same window, allowing day-ahead optimization performance and post hoc fixed-policy execution performance to be evaluated separately.

All three ablation cases use normalized weights within their respective scenario set, but their per-scenario weights differ because their scenario structures differ. The three-scenario case contains the q^{50} center, high-net-load, and high-renewable representative scenarios, with weights of 0.50, 0.25, and 0.25, respectively. In the 25-point static-grid case, the center has a weight of 0.30, while the remaining 24 non-center points jointly receive 0.70, or 0.0291667 each. The final 25-point grid + seven temporal paths case uses the weights defined in Section 2.3.3: the center receives 0.25, the remaining 24 grid points jointly receive 0.50, and the seven temporal paths jointly receive 0.25. These weights are used only for composite objective aggregation.

Table 21 shows that the day-ahead optimization results themselves change only slightly as scenario coverage increases. The mean composite objective values for the three schemes are similar, the worst modeled minimum voltage is 0.9000 p.u. in all cases, and neither weighted grid energy nor curtailment deteriorates materially. Meanwhile, as the number of scenarios increases from three to 25 and 32, the mean solution time increases from 0.739 s to 6.350 s and 11.404 s, respectively, reflecting the additional computational burden of a more complete uncertainty description. Based only on the day-ahead objective and modeled constraint results, the differences among the three scenario construction schemes are not substantial. Therefore, the day-ahead optimal value alone is insufficient to assess their effectiveness in actual operation.

Table 21. Day-ahead optimization results of the final Bus 30 scenario-construction ablation over three representative windows.

Method	Scenarios	Mean Composite Objective	Mean Runtime/s	Worst Modeled Minimum Voltage/p.u.	Mean Weighted Grid Energy/MWh	Mean Weighted Curtailment/MWh
Three-scenario representative points	3	2740.5824	0.739	0.9000	79.4005	0
25-point disturbance-space grid	25	2740.9992	6.350	0.9000	79.4348	5.52×10^{-5}
Disturbance-space grid + temporal paths	32	2744.2327	11.404	0.9000	79.4215	0.005567

Each day-ahead policy is then frozen and, on the corresponding actual trajectory, subjected to engineering replay; the results are presented in Table 22.

Table 22 reveals differences that are substantially more pronounced than those in the day-ahead results. It should first be clarified that the power/energy violations in the table are the amounts by which the frozen policy output exceeds the corresponding BESS constraint bounds, not the BESS nameplate power or energy capacity itself; their values can therefore exceed P^{cap} or E^{cap}. With the Bus 30 10.6613 MW/10.6613 MWh hardware configuration fixed, when only three representative scenario points are used, none of the three windows passes the engineering replay. The results include 18.5454 MW for the worst lower-bound power violation, 5.6642 MW for the worst upper-limit power violation, 6.6629 MWh for the worst energy violation, 9.0650 MWh for the worst terminal-energy deviation, and 13.6418 MW for the worst charging/discharging overlap, accompanied by reactive load shedding and renewable-energy curtailment. These large residuals are not caused directly by $|\xi| > 1$ with unbounded affine extrapolation, because the final policy coordinates are restricted to $[-1, 1]$. More precisely, the sparse three-scenario support does not sufficiently constrain the response of the shared affine policy to the actual time-varying disturbance sequence. Thus, although a small set of representative scenarios can yield a day-ahead objective close to that obtained with a more complete scenario set, it is insufficient to ensure engineering executability of the frozen policy in actual replay.

Table 22. Actual fixed-policy validation of the final Bus 30 scenario-construction ablation over three representative windows.

Method	Worst Lower-Bound Power Violation/MW	Worst Upper-Limit Exceedance/MW	Worst Energy Violation/MWh	Worst Terminal-Energy Deviation/MWh	Worst C/D Overlap/MW	Actual Voltage Range/p.u.	Reactive Shedding/MVArh	Curtailment/MWh	Engineering Pass
Three-scenario representative points	18.5454	5.6642	6.6629	9.0650	13.6418	0.9000–1.0357	0.036218	0.192688	0/3
25-point disturbance-space grid	0	0	0.033414	0.500163	0	0.9000–1.0348	0.006635	0	0/3
Disturbance-space grid + temporal paths	4.84×10^{-7}	0	0	9.31×10^{-7}	1.63×10^{-6}	0.901523–1.034429	0	0	3/3

After expansion to the 25-point disturbance-space grid, disturbance-magnitude coverage improves substantially: lower- and upper-bound power violations and charging/discharging overlap all decrease to 0, and curtailment also decreases to 0. However, there remains a maximum 0.033414 MWh energy violation, a 0.500163 MWh terminal-energy deviation, and 0.006635 MVArh reactive shedding; consequently, none of the three windows passes the complete engineering criteria. This result indicates that increasing static disturbance-space sampling provides better coverage of disturbance magnitudes but does not automatically capture the sensitivity of temporally coupled states such as SOC to the order and duration of error evolution.

After the 25-point grid is supplemented with seven temporal stress paths, all three representative windows pass the fixed-policy engineering validation. The energy violation decreases to 0, the terminal-energy deviation is only 9.31×10^{-7} MWh, and the lower-bound power violation and C/D overlap retain only numerical residuals on the order of 10^{-6}. No load shedding or renewable-energy curtailment occurs, and the actual voltage remains within 0.901523–1.034429 p.u. Thus, the primary contribution of the temporal stress paths is not a material improvement in the day-ahead objective, but the addition of intraday temporal stresses that cannot be represented by the static disturbance-space grid, enabling the frozen policy to maintain engineering consistency of the storage state and network operation during actual replay.

In summary, this ablation shows that the three types of scenario construction differ only slightly in their day-ahead optimization metrics but differ substantially in fixed-policy execution performance on the tested actual trajectories. The sparse three-point scenario support does not sufficiently constrain the response of the shared affine policy to actual time-varying disturbances. Expansion to the 25-point grid markedly improves static disturbance-space coverage, but still does not sufficiently constrain the response of intertemporal states such as SOC and terminal energy to error duration and order. After the seven temporal stress paths and their associated terminal-energy safeguard are added in the final formulation, all three representative windows pass the complete engineering validation. Therefore, the comparison indicates that the static disturbance-space grid and the combined temporal-path/terminal-energy protection play complementary roles in the tested fixed-policy execution results. Because the final ablation step introduces both temporal stress paths and the associated terminal-energy safeguard, the present experiment does not isolate their individual causal contributions.

3.3.10. Actual Fixed-Policy Operation on a Representative Day

The first rolling window (WIN001) is selected for the representative-day visualization to illustrate the execution of the frozen day-ahead policy on an actual trajectory. This window serves only as a temporal-visualization example and is not subjected to additional screening for typicality or favorability. Figures 4–6 sequentially present the system net-load prediction interval and actual trajectory, the BESS charging/discharging power response, and the corresponding SOC evolution.

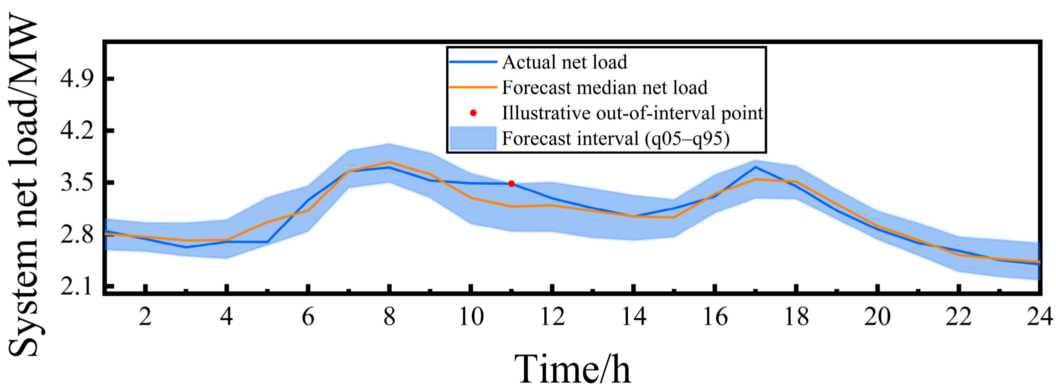

Figure 4. System net-load forecast envelope and actual trajectory on a representative day.

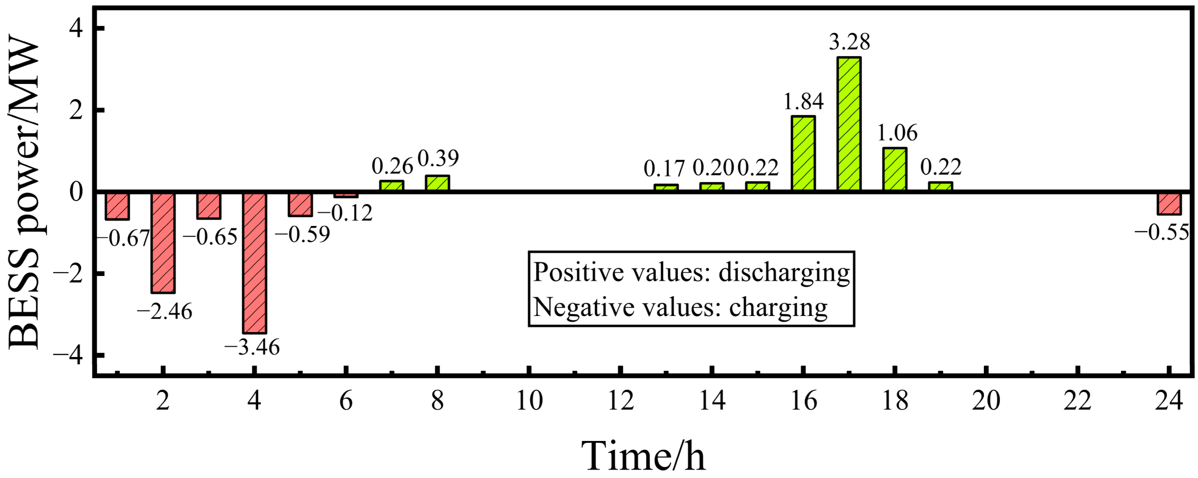

Figure 5. BESS charging/discharging power under the actual trajectory on a representative day.

Figure 4 presents, for WIN001, the system net-load forecast median, q^{05}–q^{95} prediction interval, and actual net-load trajectory. The blue solid line represents the actual net load, the orange solid line represents the forecast median net load, and the blue shading represents the q^{05}–q^{95} forecast interval. A red circle marks an illustrative out-of-interval point to illustrate that an actual realization can fall outside the day-ahead prediction interval.

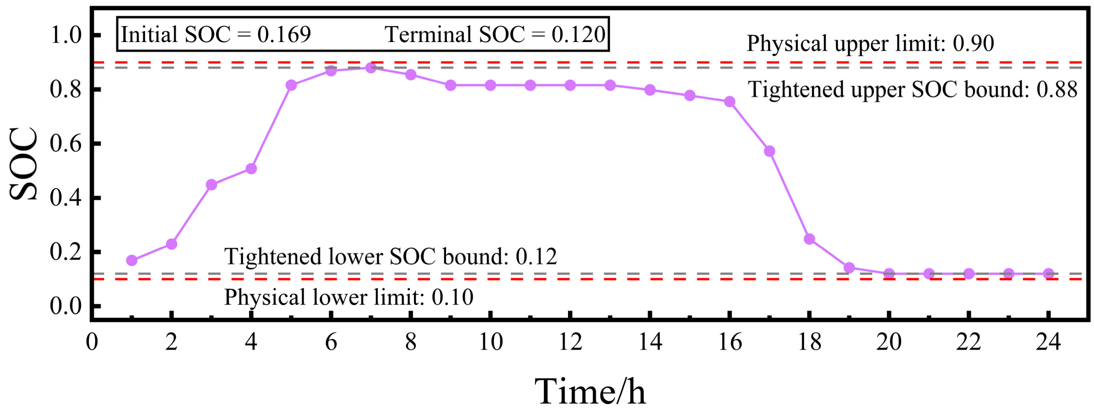

Figure 6. SOC trajectory under the actual trajectory on a representative day.

It should be emphasized that an actual trajectory outside the q^{05}–q^{95} interval does not imply that the actual physical input is clipped back to the prediction interval. In accordance with Section 2.5.3, the fixed-policy execution rules retain the actual physical load and renewable outputs without modification for subsequent network validation; only the policy coordinates used to invoke the frozen affine BESS policy are restricted to the predefined $[-1, 1]$ range.

Figure 5 presents, on the same representative day, the BESS hourly power response produced by the frozen day-ahead policy. In WIN001, at approximately 2 h and 4 h, the charging powers reach approximately 2.46 MW and 3.46 MW, respectively. The BESS then switches to discharging during the higher-net-load period, with the main discharge interval occurring from 16–18 h and, at approximately 17 h, reaching approximately 3.28 MW. In Figure 4, the net-load trajectory changes consistently with this temporal response, demonstrating the actual response of the frozen policy to the actual operating conditions.

Figure 6 presents the result corresponding to Figure 5, namely the SOC trajectory. For WIN001, the initial SOC is 0.16923, and its terminal SOC is 0.12000. As charging proceeds during the first half of the day, SOC gradually increases from its initial value and, at approximately 6–7 h, approaches the tightened upper SOC bound. It then decreases during the middle and latter portions of the day under discharging and reaches 0.12000 at the end of the day.

Throughout the 24 h period, SOC remains within the tightened bounds 0.12–0.88 used in day-ahead optimization and also within the physical limits 0.10–0.90. The red dashed lines indicate the physical SOC limits, and the gray dashed lines indicate the tightened SOC bounds. The result shows that, on this representative actual trajectory, the frozen day-ahead policy maintains the BESS energy state within the prescribed engineering bounds without post hoc re-optimization.

Figures 4–6 jointly show that operating feasibility is not obtained by hourly re-optimization. Instead, when actual net load deviates from the forecast median, the frozen BESS policy shifts charging and discharging over time while maintaining the SOC constraints. This representative-day result provides an intuitive temporal example for the subsequent multiwindow statistical validation, but reflects only the typical behavior of one operating day. Overall engineering feasibility should still be assessed using the fixed-policy validation results for all rolling windows.

3.3.11. Multiday Rolling-Window Actual Validation

To evaluate the fixed-policy execution performance of the final BESS configuration over the tested continuous multiday sequence, for the IEEE-33 MAIN scheme, 46 consecutive rolling windows are subjected to carryover-fixed validation. Throughout the validation, Bus 30 and its final capacity configuration of 10.6613 MW/10.6613 MWh remain fixed; the site is not changed, the capacity is not expanded, and the rating is not adjusted for interday operation. The first window uses the prescribed initial energy of 1.804214 MWh. Thereafter, the BESS energy at the end of the actual execution of each window is used as the initial state of the next, forming 45 continuous interday state boundaries. Day-ahead optimization and the corresponding actual fixed-policy replay are completed for each window.

Table 23 shows that all 46 day-ahead optimization problems and 46 actual validations are completed successfully, with both BESS and network engineering-feasibility rates reaching 100%. More importantly, the state boundaries of all 45 adjacent windows are continuous, and the maximum interday continuity error is 0 MWh, indicating that the carryover-fixed protocol maintains interday state transfer without resetting SOC. Over the entire 46 d period, the actual SOC range is 0.1199999–0.8800001 p.u., differing from the 0.12–0.88 tightened bounds used in optimization only by numerical deviations on the order of 10^{-7}, and always remaining within the physical bounds of 0.10–0.90. The final cumulative energy drift over 46 d is only 7.67×10^{-7} MWh, and no systematic accumulation of SOC drift caused by continuous rolling execution is observed.

In Table 23, the exact-zero numerical diagnostic should be clearly distinguished from engineering feasibility. As an auxiliary numerical diagnostic, if the relevant residuals are checked point by point to determine whether they are exactly 0 at the output precision of the solver, 43/46 windows satisfy the exact-zero condition, while the remaining three windows have only extremely small nonzero numerical residuals. This exact-zero diagnostic is used only to show the actual scale of the numerical residuals and is not an engineering-feasibility criterion. To quantify their scale further, Table 24 reports the most adverse constraint residuals and interday-continuity results over all 46 windows.

As shown in Table 24, all residuals that prevent the exact-zero numerical diagnostic from reaching 46/46 are on the order of 10^{-6} or smaller. The maximum BESS power violation is only 7.8930×10^{-7} MW, the maximum terminal-energy deviation is 1.7193×10^{-6} MWh, and the maximum simultaneous charging/discharging residual is 6.4661×10^{-7} MW; both the energy-limit violation and interday continuity error are 0. Thus, these three windows reflect numerical residuals at the scale of floating-point computation and solver tolerances rather than engineering infeasibility of power, energy, or interday states. Under the unified engineering-feasibility criteria adopted in this study, all 46 windows pass both BESS and network validation.

Table 23. Overall validation results over 46 rolling windows.

Category	Index	Result
Fixed configuration	BESS installation bus P^{cap}/E^{cap}	Bus 30 10.6613 MW/10.6613 MWh
Rolling protocol	Continuous windows Interday state boundaries Day-1 initial energy	46 45 1.804214 MWh
Solution stability	Day-ahead problems completed Actual validations completed BESS engineering feasibility Network engineering feasibility Strict numerical BESS feasibility	46/46 46/46 46/46 46/46 43/46 (93.478%)
SOC continuity	Boundary continuity Maximum boundary continuity error Minimum actual SOC Maximum actual SOC Final 46 d energy drift	45/45 0 MWh 0.1199999 p.u. 0.8800001 p.u. 7.67×10^{-7} MWh
Operating objective	Mean day-ahead composite objective Sum over 46 windows	3143.168368 o.u. 144,585.744934 o.u.
Capacity implication	Additional capacity introduced for carryover validation	None

Table 24. Constraint-residual and interday-continuity audit under carryover-fixed rolling validation.

Index	Worst Value	Engineering Interpretation
Maximum BESS power violation	7.8930×10^{-7} MW	Pass
Maximum BESS energy-limit violation	0 MWh	Pass
Maximum terminal-energy deviation	1.7193×10^{-6} MWh	Pass
Maximum simultaneous charge/discharge residual	6.4661×10^{-7} MW	Pass
Maximum interday continuity error	0 MWh	Pass
Boundary continuity pass rate	45/45	100%
Storage engineering pass rate	46/46	100%
Network engineering pass rate	46/46	100%

In addition, the entire carryover-fixed validation uses the originally determined Bus 30 capacity configuration, without any additional BESS power or energy capacity to accommodate continuous multiday operation. This shows that the preceding single-day siting and sizing result can be applied to the tested 46 d continuous actual operating sequence without post hoc capacity expansion during interday execution. The mean day-ahead composite objective is 3143.168368 o.u., and the cumulative value over 46 windows is 144,585.744934 o.u. These values characterize the unified engineering evaluation during continuous operation rather than lifecycle economic cost.

In summary, the 46-window carryover-fixed validation shows that the final Bus 30 configuration maintains complete interday SOC continuity and an engineering-feasibility rate of 100% with capacity fixed, without cumulative energy drift or additional capacity requirements caused by continuous rolling execution. This result is multiday post hoc validation evidence for the tested 46 d actual trajectory bank and should not be extended to imply an infinite-horizon feasibility guarantee for arbitrary long-term disturbance sequences.

3.3.12. Supplementary Validation on the IEEE-69-Bus Test Feeder

To further examine the execution of the revised method on a second radial test feeder, supplementary validation is conducted on the IEEE-69-bus system. Bus 1 is the slack bus, and PV and WT are connected to Bus 65 and Bus 64, respectively, i.e., $i_{PV} = 65$ and $i_{WT} = 64$, with installed capacities of 2.5 MW and 2.0 MW. Compared with the IEEE-33 system, this experiment changes the network scale, radial topology, and renewable-energy connection locations, while retaining the same system-level forecast inputs, scenario-construction method, and overall planning and rolling-validation procedure. Accordingly, the results in this section provide supplementary feasibility evidence on a second test feeder and are not interpreted as generalized cross-system evidence across different load levels, renewable penetration levels, or uncertainty distributions.

After adding duration-aware constraints, for the IEEE-69 system, planning optimization is repeated over all 68 eligible installation buses. The MAIN scheme uses a 1–4 h duration constraint and ultimately selects Bus 19, yielding $P^{cap} = 7.4552$ MW and $E^{cap} = 7.4552$ MWh, corresponding to a rated duration of 1 h. The STRICT scheme uses a 2–4 h constraint and ultimately selects Bus 18, yielding $P^{cap} = 7.6984$ MW and $E^{cap} = 15.3968$ MWh, corresponding to a rated duration of 2 h. The original Bus 68 10 MW/0.4496 MWh result is no longer treated as a valid configuration under the revised model. To check whether the final configuration is constrained by an artificial capacity upper bound, the power and energy upper bounds in the capacity audit are further relaxed from $P^{cap} \leq 20$ MW, $E^{cap} \leq 80$ MWh to $P^{cap} \leq 30$ MW, $E^{cap} \leq 120$ MWh.

Across 136 planning problems comprising 68 eligible installation buses and two duration cases, the maximum power and energy capacities among the relaxed-bound solutions are both 15.428 MW for MAIN; the corresponding maxima for STRICT are 14.009 MW and 28.017 MWh. All relaxed-bound solutions remain strictly below the original upper bounds of 20 MW/80 MWh. Therefore, the original capacity limits are not binding constraints on the revised IEEE-69 planning results. This audit only shows that the capacity upper bounds do not truncate the optimum within the relaxed search space; it does not imply that all original operating metrics are identical under different upper-bound settings.

Table 25 shows that, in the IEEE-69 system, the MAIN and STRICT schemes select Bus 19 and Bus 18, respectively, indicating that a stricter duration requirement not only increases the energy-capacity requirement but may also change the preferred installation location. The final MAIN configuration is 7.4552 MW/7.4552 MWh, corresponding to a 1 h duration, whereas the STRICT configuration is 7.6984 MW/15.3968 MWh, corresponding to a 2 h duration, consistent with its stricter energy-duration constraint. After the capacity-search upper bounds are further relaxed, the allowable maximum powers for MAIN and STRICT increase to 15.428 MW and 14.009 MW, respectively, while the allowable maximum energies increase to 15.428 MWh and 28.017 MWh, all substantially exceeding the final planned capacities. Thus, the current selected configurations do not result from truncation by the original capacity upper bounds and remain in the interior feasible region after the search space is relaxed.

Table 25. Duration-aware BESS planning and relaxed-capacity-bound audit for the IEEE-69-bus system.

Case	Duration Range	Selected Bus	P^{cap}/MW	E^{cap}/MWh	Relaxed Max P/MW	Relaxed Max E/MW
MAIN	1–4 h	Bus 19	7.4552	7.4552	15.428	15.428
STRICT	2–4 h	Bus 18	7.6984	15.3968	14.009	28.017

On this basis, for both final IEEE-69 configurations, complete 46-window carryover rolling validation is conducted, and the results are presented in Table 26.

Table 26. Final rolling operating validation on the IEEE-69-bus test feeder.

Index	MAIN	STRICT
Selected bus	Bus 19	Bus 18
P^{cap}/MW	7.4552	7.6984
E^{cap}/MWh	7.4552	15.3968
Completed windows	46/46	46/46
Carryover boundaries	45/45	45/45
Storage engineering feasibility	46/46	46/46
Network engineering feasibility	45/46	45/46
Active-load shedding/MWh	0	0
Reactive-load shedding/MVArh	0.010201	0.010201
Renewable curtailment/MWh	0	0
Failed network window	WIN034	WIN034

Table 26 shows that all 46 day-ahead and actual rolling windows are completed for both configurations, all 45 interday carryover boundaries are continuous, and storage engineering feasibility reaches 46/46 in both cases. This indicates that the final BESS configurations maintain complete interday storage-state continuity in the 69-bus system, without engineering failure of BESS power or energy caused by the larger network scale.

On the network side, 45/46 windows pass the complete engineering-feasibility check, with MAIN and STRICT both having their only network-failure window at WIN034. Neither case involves active-load shedding or renewable curtailment, and only 0.010201 MVArh of cumulative reactive-load shedding occurs. Because the failures are concentrated in the same WIN034 window and the two different BESS configurations behave consistently, this result is more likely to reflect local reactive-power support stress in the network operating state corresponding to WIN034 than a failure unique to one siting scheme. This should, however, be interpreted as an experimental observation rather than a rigorous causal demonstration of the failure mechanism.

Overall, the IEEE-69 results show that the proposed sequential workflow can complete duration-aware siting and sizing and continuous rolling execution on a larger feeder, and that the final capacity is not artificially limited by preset capacity bounds. At the same time, the 45/46 network engineering feasibility indicates high but not absolute network feasibility on this system. The local reactive-power issue in WIN034 further shows that extreme operating points in a larger network may still require more refined AC-aware reactive-power support or network-correction mechanisms.

3.3.13. Post Hoc Annualized BESS Ownership-Cost Audit

The preceding SG-ENAR-OPF uses a composite engineering objective, whose capacity regularization term provides a numerical tradeoff in technical planning and is not equivalent to a complete lifecycle monetary cost. To further quantify the realistic investment scale of the final BESS configurations, an independent post hoc annualized ownership-cost audit is conducted for the MAIN and STRICT configurations of IEEE-33 and IEEE-69. This audit only converts the determined technical configurations into economic values and does not participate in re-optimizing the BESS location or capacity.

Let the unit BESS power cost and unit energy cost be c_p^{inv} and c_E^{inv}, respectively. When P^{cap} is expressed in MW and E^{cap} in MWh, while the cost parameters are expressed in USD/kW and USD/kWh, the initial investment cost is:

$$C_{CAPEX} = 10^3(c_p^{inv} P^{cap} + c_E^{inv} E^{cap}) \tag{123}$$

This study uses $c_p^{inv} = 597$ USD/kW, $c_E^{inv} = 192$ USD/kWh, a discount rate of $r_{disc} = 7\%$, and an economic lifetime of $N_{life} = 15$ years. The capital-recovery factor is:

$$CRF = \frac{r_{disc}(1 + r_{disc})^{N_{life}}}{(1 + r_{disc})^{N_{life}} - 1} \tag{124}$$

This corresponds to $CRF \approx 0.10979$. Fixed operation and maintenance expenses are calculated as 2.5% of the initial CAPEX per year and are set to C_{VOM} in the present audit. The annualized ownership cost is therefore:

$$C_{own}^{ann} = CRF \cdot C_{CAPEX} + 0.025 C_{CAPEX} + C_{VOM} \tag{125}$$

Table 27 shows that a stricter minimum-duration requirement materially increases annualized ownership cost. For IEEE-33, the STRICT scheme relative to MAIN increases this cost by 23.365%; for IEEE-69, the corresponding increase is 28.391%. In both systems, this increase results mainly from the larger energy capacity, whereas the change in power capacity is relatively small, indicating that the duration requirement changes not only the technical configuration but also the corresponding scale of capital investment. This result explains the economic implications of the technical selection and does not reinterpret the current SG-ENAR-OPF as a lifecycle economic-optimization model.

Table 27. Post hoc annualized ownership-cost audit of the final BESS configurations.

System	Case	Bus	P/MW	E/MWh	Initial CAPEX/Million 2019 USD	Annual Ownership/Million 2019 USD·y^{-1}
IEEE-33	MAIN	30	10.6613	10.6613	8.4118	1.1339
IEEE-33	STRICT	30	10.5782	21.1565	10.377	1.3988
IEEE-69	MAIN	19	7.4552	7.4552	5.8821	0.79288
IEEE-69	STRICT	18	7.6984	15.3968	7.5521	1.0180

To further characterize the cycling exposure of the final BESS during actual operation, the charged and discharged energy of the IEEE-33 MAIN scheme is quantified over 46 d of 1104 h fixed-policy execution. The cumulative gross throughput is defined as:

$$E_{gross} = \sum_t (P_{ch,t}^{act} + P_{dis,t}^{act})\Delta t \tag{126}$$

The number of equivalent full cycles is defined as:

$$N_{EFC} = \frac{E_{gross}}{2E^{cap}} \tag{127}$$

Over the 46 d validation period, cumulative charging energy is 392.216 MWh, cumulative discharging energy is 353.975 MWh, and total throughput is 746.192 MWh, corresponding to 34.995 EFCs and a mean daily EFC of 0.761. Relative to rated energy capacity, the actual maximum charging and discharging C-rates are, respectively:

$$C_{ch}^{peak} = P_{ch}^{peak} / E^{cap} = 3.8903/10.6613 \approx 0.3649C \tag{128}$$

$$C_{dis}^{peak} = P_{dis}^{peak} / E^{cap} = 4.7820/10.6613 \approx 0.4485C \tag{129}$$

Thus, for MAIN, although the 1–4 h technical band permits a power-to-energy ratio of up to approximately 1C in rated terms, the peak charging and discharging rates during the tested actual fixed-policy operation are substantially below this upper limit. The preceding throughput and EFC values describe only operating cycling exposure and are not used to infer cell-level degradation, capacity fade, battery life, or replacement cost; these factors are not explicitly included in the current planning model.

4. Discussion

4.1. Functional Distinction Between Sensitivity Screening and BESS Siting

Bus sensitivity reflects the marginal network response to a unit active- or reactive-power net injection near the base operating point. It is suitable for identifying buses with potential value for voltage support, loss reduction, and power-flow relief, but cannot be equated directly with the composite configuration value of a BESS under multi-period, multiscenario conditions. Final siting is also affected jointly by technical duration, power/energy-capacity requirements, responses to uncertain scenarios, and network constraints. Differences between the sensitivity ranking and the final composite OPF ranking are therefore reasonable.

The revised full-node audit uses common normalization bounds over all buses to evaluate the main candidate-set and full-node results on a consistent basis. Under the eight preset evaluation weights used in this study, for IEEE-33 MAIN and STRICT, the full-node Top5 buses are Bus 30, Bus 28, Bus 29, Bus 31, and Bus 32, all of which are retained in the main candidate set. This indicates that sensitivity-guided screening compresses the subsequent search space under the current evaluation settings while retaining high-ranking planning buses.

However, RobustScore does not have a unique ranking independent of the evaluation weights. Analyses using 5000 random weight sets, equal weights, drop-one, and one-at-a-time all show that the final ranking changes markedly with multi-criteria preferences. Therefore, Bus 30 as the "best audited bus" means only that it ranks first under the current model, common full-node normalization bounds, and explicit weight settings, not that it is a universally globally optimal location under all evaluation preferences. The primary role of sensitivity analysis is thus to reduce the scope of rigorous optimization at low cost, whereas the full-node audit serves as a one-time benchmark for verifying whether candidate compression omits high-value buses.

4.2. Coupling Among BESS Location, Capacity, and Network Performance

The BESS connection location, power capacity, energy capacity, and technical duration are strongly coupled; sizing therefore cannot be discussed independently of the network location and technical feasible region. The original independent P/E formulation, in the IEEE-69 system, produced an extremely short-duration configuration of 10 MW/0.4496 MWh, approximately 0.045 h. After adding explicit duration constraints of 1–4 h and 2–4 h and repeating the full-node planning, both the preferred locations and capacity compositions change in IEEE-33 and IEEE-69. This shows that rated duration is not a post hoc interpretation parameter, but a technical assumption that directly changes the planning feasible region and capacity structure.

IEEE-33 shows that both MAIN and STRICT select Bus 30, but when the minimum duration increases from 1 h to 2 h, the principal change is an increase in energy capacity. In IEEE-69, under MAIN the preferred location is Bus 19, whereas under STRICT it is Bus 18. This further indicates that location and capacity are not separable but adjust jointly with the duration requirement.

It should be distinguished that the planning objective used in this study is an engineering scalarization of network safety, operating performance, response intensity, and capacity scale, rather than lifecycle economic optimization in a common monetary unit. Therefore, the capacity-size regularization coefficients in Equation (64) cannot be interpreted directly as capital cost. The newly added annualized ownership-cost audit only quantifies the investment scale of the final technical configurations. Under STRICT relative to MAIN, the results for IEEE-33 and IEEE-69 show that annualized ownership cost increases by approximately 23.37% and 28.39%, respectively. If different capital-cost, degradation, and replacement assumptions are to be considered, the lifecycle economic objective must be reconstructed and the problem re-optimized.

4.3. Symmetry and Asymmetry in Disturbance-Space and Temporal-Path Construction

Static scenarios on the two-dimensional load–renewable disturbance plane use paired opposite policy coordinates to prevent the day-ahead policy from responding only to a single direction, such as high net load or high renewable output. Here, central symmetry is a geometric property of the standardized-disturbance coordinates and does not imply spatial symmetry in the physical locations of the distribution network or in multisite forecast errors. The temporal stress paths further introduce SOC-path differences caused by disturbance duration, occurrence order, and switching pattern.

The ablation experiment shows that, within the tested range, multidirectional disturbance-space scenarios reduce instantaneous charging/discharging-power and operating-mode violations, while adding temporal paths and a terminal-energy constraint further reduces intertemporal energy violations and terminal deviations. The two types of scenarios therefore play complementary roles in covering disturbance magnitude and temporal evolution. This result does not imply that a nodal or multisite spatial uncertainty model has been established, nor can the 25 static points and seven paths be interpreted as complete coverage of the continuous disturbance domain and all time series.

The further path-design sensitivity audit examines changes in peak duration, segment length, alternating frequency, ramp form, and number of paths. All 1242 cross-window replay runs remain engineering-feasible, with a maximum change in objective of 0.411482%. Thus, the current seven-path design exhibits low engineering sensitivity among the tested finite variants, but its specific durations and shapes are not theoretically unique. In particular, when the number of paths is reduced to three, the objective changes most, indicating that the coverage of the temporal scenario set still affects the day-ahead policy. The current results cannot be interpreted as a complete reconstruction of actual forecast-error trajectories.

4.4. Fixed-Policy Execution Evidence and Validation Scope

Actual validation keeps the day-ahead BESS policy parameters fixed, and the active-power dispatch is not re-optimized. For an actual physical trajectory outside the prediction interval, the standardized policy input is mapped according to the bounded day-ahead disturbance domain, whereas the BESS charging/discharging commands produced by the frozen policy are not subjected to clipping or projection. In addition to daily validation, IEEE-33 is subjected to a 46 d interday SOC continuity test to examine interday storage-state continuity of the final technical configuration over the tested continuous operating sequence.

The current 46-window actual validation is entirely within the chronological forecast test segment and, relative to forecast-model training and validation, constitutes out-of-sample evaluation. However, the available DK1 2019 dataset does not provide an independent subsequent period after 12/31, so this sequence is not interpreted as a genuine fresh post-period holdout. More rigorous subsequent validation should use genuinely future-period data after the model, interval-calibration rules, and scenario-construction parameters have all been frozen.

IEEE-69 provides supplementary validation as a second radial test feeder. After duration-aware replanning and a capacity-boundary audit, its final configurations are no longer directly restricted by the original 20 MW/80 MWh capacity upper bounds.

BESS engineering feasibility is 46/46 and network engineering feasibility is 45/46 for both technical-duration schemes. Thus, this result provides supplementary evidence of method execution on a second test feeder under the common uncertainty-input protocol, but is insufficient to demonstrate generalized cross-system applicability across different load levels, renewable penetration levels, and uncertainty distributions.

Regarding the prediction intervals, without post hoc coverage recalibration, this study uses q^{05}–q^{95} nominal 90% conditional quantile intervals. In the final test samples, the empirical coverage rates for Load, PV, and WT are 84.06%, 82.34%, and 76.63%, respectively. For Load/PV, the tail-specific results show predominantly upper-tail undercoverage, whereas WT exhibits predominantly lower-tail undercoverage. Therefore, the prediction interval cannot be interpreted as an uncertainty set with a 90% finite-sample coverage guarantee. Independent actual fixed-policy validation is used to examine operating performance in the presence of this forecast error, but operating validation cannot substitute for calibration of the probabilistic forecasts themselves. On an independent calibration set, future work could use conformal calibration or quantile recalibration to improve interval reliability, followed by reevaluation on truly future data not used in calibration.

4.5. Modeling Limitations and Future Extensions

From the perspective of uncertainty modeling, this study uses 25 static scenarios and seven temporal stress paths to form a finite scenario set that covers multidirectional disturbances and selected typical temporal-accumulation patterns. However, neither a max–min robust counterpart nor a separation procedure over a continuous uncertainty set is formulated. Accordingly, finite-scenario feasibility cannot be interpreted as a formal robustness guarantee for all disturbance trajectories. The actual rolling validation, temporal-path sensitivity, and other stress tests provide only empirical support under the current data and finite disturbance set.

From the perspective of BESS physical modeling, the revised model describes engineering constraints at the system/battery-pack level through rated duration/C-rate, SOC, charging/discharging power bounds, efficiency, and inverter reactive-power capability, and quantifies cycling exposure through throughput and EFC. It does not, however, explicitly consider cell current, voltage, temperature, electrochemical degradation, capacity fade, or replacement cost. Accordingly, the composite objective at the planning stage remains an engineering scalarization, while annualized ownership cost is only a post hoc economic audit. More complete lifecycle economic optimization would require explicit degradation and replacement mechanisms.

Another limitation arises from aggregation of the renewable policy dimensions. PV and WT outputs are retained separately in the physical network model and actual validation, whereas the BESS affine policy uses a capacity-weighted aggregate renewable coordinate to control the policy dimension and finite-scenario MILP size. This representation captures the net effect of PV/WT forecast errors on BESS control requirements but cannot uniquely retain all independent PV–WT combinations, especially errors in opposite directions. In the actual data, approximately 51.91% of valid PV/WT standardized errors have opposite signs, further demonstrating that a single renewable policy coordinate cannot be interpreted as a complete multisource joint uncertainty model. Independent PV/WT targeted stress tests provide only limited supplementary evidence and cannot replace explicit multidimensional dependence modeling.

In addition, the current uncertainty representation remains low-dimensional and aggregated: load is mapped to all buses using a common system-level multiplier, and PV and WT correspond to fixed connection points. It therefore does not characterize spatially heterogeneous forecast errors and their dependence structure across multiple buses and renewable sites. The numerical experiments are based on the standard radial IEEE-33 and IEEE-69 test feeders, with test inputs constructed from public time-series data and preset renewable connection locations, rated capacities, and scaling rules; some operating states, such as synthetic reverse-flow, are used only for the robustness audit. Accordingly, the conclusions are limited to the tested feeders, input construction, and parameter settings and are not extrapolated as proof of general applicability to actual distribution networks or other network types. Future work may introduce independent PV/WT policy coordinates, data-driven joint scenario generation, or correlated/distributionally robust uncertainty models, together with actual multisite feeder data and longer-term validation, to investigate more complete multisource joint errors and cross-system applicability.

5. Conclusions

For the problems of identifying bus regulation value, BESS siting and sizing, and day-ahead operation under forecast uncertainty in distribution networks with high wind and photovoltaic penetration, this study develops a sequential planning-to-operation workflow comprising "sensitivity screening—SG-ENAR-OPF siting and sizing—interval-forecast-driven finite-scenario scheduling—actual fixed-policy validation." The day-ahead stage uses 25 static disturbance combinations and seven temporal stress paths, with a shared affine policy to respond to the modeled disturbances nonanticipatively. The robustness claim is limited to joint constraint satisfaction over the constructed finite scenario set. The rolling-window results provide separate fixed-policy validation evidence on the tested actual trajectories under the stated engineering tolerances; neither result constitutes a formal guarantee for arbitrary disturbance trajectories over a continuous uncertainty domain. The principal conclusions are as follows:

1. Sensitivity screening substantially compresses the subsequent planning search space while retaining key high-value buses.

Bus net-injection sensitivity distinguishes system-level regulation buses from support buses in target weak-voltage areas. When the system-level truncation depth is $k = 8$, Jacobian-assisted AC finite-difference refinement requires eight buses to undergo nodal P/Q perturbation, whereas a complete 32-bus analysis requires 65 AC power-flow calls in this substage; the refinement requires 17, a reduction of 73.85%, while fully retaining the system's composite P/Q Top5 buses. On this basis, predefined topology/stress safeguards add seven buses to form a 15-bus main candidate set, without increasing nodal AC finite-difference runs; all 15 candidate buses are then evaluated using the same SG-ENAR-OPF formulation.

2. Technical-duration constraints directly affect the BESS capacity structure and may change the final siting result.

In the IEEE-33 system, both the 1–4 h MAIN and 2–4 h STRICT schemes select Bus 30, with capacities of 10.6613 MW/10.6613 MWh and 10.5782 MW/21.1565 MWh, respectively. Under the eight evaluation weights and common full-node normalization bounds adopted in this study, the full-node Top 5 buses for both schemes are Bus 30, Bus 28, Bus 29, Bus 31, and Bus 32, all of which are retained in the main candidate set. Further analyses using random weights, equal weights, and individual-metric perturbations show that the RobustScore ranking clearly depends on the evaluation preference. Bus 30 should therefore be understood as the best audited bus under the current model, technical assumptions, and explicit weight settings, not as a universally optimal location independent of evaluation preferences.

3. Disturbance-space scenarios and temporal stress paths supplement information on disturbance magnitude and temporal evolution, respectively.

Paired static scenarios in the two-dimensional load–renewable disturbance plane cover opposite disturbance directions, while temporal stress paths and the terminal-energy safeguard further characterize the cumulative SOC effects of duration, occurrence order, and switching pattern. The ablation results show that increasing static disturbance-space coverage materially improves instantaneous power response, but the principal energy and terminal-state violations are removed only after temporal paths are added. In the sensitivity audit of path count, peak duration, switching time, alternating frequency, and transition shape, 27 frozen policies over 46 actual windows remain within the engineering-feasibility tolerances in all 1242 cross-window replay runs, and the maximum absolute change in objective is only 0.411482%. Thus, the current seven-path design exhibits low engineering sensitivity over the tested variants, but its role should be understood as finite temporal stress testing rather than complete reconstruction of actual forecast-error trajectories. The two-dimensional disturbance space in this study is likewise a BESS policy-coordinate space, not a nodal or multisite physical-spatial uncertainty model.

4. Multiday actual validation and validation on a second test feeder provide fixed-policy execution evidence for the final configurations under the tested settings while identifying the applicability boundaries of the model.

The IEEE-33 MAIN configuration achieves 46/46 BESS and LinDistFlow network engineering feasibility over the 46 d, 1104 h carryover-fixed rolling validation, with continuity across all 45 interday state boundaries and no additional BESS capacity. Independent nonlinear AC post-validation further shows that the fixed uncertainty-aware policy has lower network losses, fewer constraint violations, and a better composite operating metric than the q^{50} baseline, but its AC voltage-feasibility rate is 95.833%, demonstrating that LinDistFlow feasibility cannot be equated directly with complete nonlinear AC feasibility. Under the 1–4 h and 2–4 h schemes, the IEEE-69 system selects Bus 19 and Bus 18, respectively; both configurations achieve 46/46 BESS engineering feasibility and 45/46 network engineering feasibility. These results provide supplementary evidence of method execution under different network scales, radial topologies, and renewable connection locations, but are insufficient to demonstrate general cross-system applicability across different load levels, renewable penetration levels, and uncertainty distributions.

Author Contributions: Conceptualization, J.M., J.P., H.H. and L.S.; Methodology, J.M., J.P., H.H. and L.S.; Software, J.M.; Validation, H.H.; Formal analysis, J.M.; Data curation, H.H. and H.L.; Writing—original draft, J.M.; Writing—review & editing, J.M. and H.H.; Visualization, J.M.; Supervision, J.P. and L.S.; Project administration, J.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the 2026 Basic Scientific Research Project of the Department of Education of Liaoning Province, grant number LJ212610147015.

Data Availability Statement: The dataset used in this study is publicly available on Kaggle. The hourly load, wind power, solar power, and electricity price dataset is available at https://www.kaggle.com/datasets/eugeniyosetrov/load-wind-and-solar-prices-in-hourly-resolution (accessed on 18 June 2026).

Acknowledgments: We confirm that a generative artificial intelligence tool was used during the preparation of the manuscript. Tool used: OpenAI ChatGPT. Model/version: GPT-5.6 Sol. The tool was used exclusively for language-related assistance, including English translation, language polishing, grammar checking, and improvement of clarity and academic expression. We would like to emphasize that the AI tool was used only to improve the linguistic presentation of the manuscript. The research ideas, methodological framework, mathematical models, software implementation, data analysis, interpretation of the results, and conclusions were independently developed and completed by the authors. All AI-assisted text was carefully reviewed, revised, and verified by the authors.

The authors take full responsibility for the accuracy, originality, integrity, and scientific content of the manuscript.

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

References

1. Karve, G.M.; Thakare, M.S.; Vaidya, G.A. Optimal Siting and Sizing of Battery Energy Storage System in Distribution System in View of Resource Uncertainty. *Energies* **2025**, *18*, 2340. [CrossRef]
2. Eftekharnejad, S.; Vittal, V.; Heydt, G.T.; Keel, B.; Loehr, J. Impact of increased penetration of photovoltaic generation on power systems. *IEEE Trans. Power Syst.* **2013**, *28*, 893–901. [CrossRef]
3. Gao, X.; Zhang, J.; Sun, H.; Liang, Y.; Wei, L.; Yan, C.; Xie, Y. A review of voltage control studies on low voltage distribution networks containing high penetration distributed photovoltaics. *Energies* **2024**, *17*, 3058. [CrossRef]
4. Yi, J.H.; Cherkaoui, R.; Paolone, M. Dispatch-aware planning of energy storage systems in active distribution network. *Electr. Power Syst. Res.* **2021**, *190*, 106644. [CrossRef]
5. Das, C.K.; Bass, O.; Kothapalli, G.; Mahmoud, T.S.; Habibi, D. Overview of energy storage systems in distribution networks: Placement, sizing, operation, and power quality. *Renew. Sustain. Energy Rev.* **2018**, *91*, 1205–1230. [CrossRef]
6. da Silva, D.J.; Belati, E.A.; López-Lezama, J.M. Enhancing Distribution Networks with Optimal BESS Sitting and Operation: A Weekly Horizon Optimization Approach. *Sustainability* **2024**, *16*, 7248. [CrossRef]
7. Zicmane, I.; Kovalenko, S.; Sahnovskis, A.; Petrichenko, R.; Junghans, G. Node Symmetry Analysis as an Early Indicator of Locational Marginal Price Growth in Network-Constrained Power Systems with High Renewable Penetration. *Symmetry* **2026**, *18*, 547. [CrossRef]
8. Prabpal, P.; Kongjeen, Y.; Bhumkittipich, K. Optimal battery energy storage system based on var control strategies using particle swarm optimization for power distribution system. *Symmetry* **2021**, *13*, 1692. [CrossRef]
9. Saboori, H.; Hemmati, R.; Ghiasi, S.M.S.; Dehghan, S. Energy storage planning in electric power distribution networks–A state-of-the-art review. *Renew. Sustain. Energy Rev.* **2017**, *79*, 1108–1121. [CrossRef]
10. Das, C.K.; Bass, O.; Kothapalli, G.; Mahmoud, T.S.; Habibi, D. Optimal placement of distributed energy storage systems in distribution networks using artificial bee colony algorithm. *Appl. Energy* **2018**, *232*, 212–228. [CrossRef]
11. Wichitkrailat, K.; Premrudeepreechacharn, S.; Siritaratiwat, A.; Khunkitti, S. Optimal sizing and locations of multiple BESSs in distribution systems using crayfish optimization algorithm. *IEEE Access* **2024**, *12*, 94733–94752. [CrossRef]
12. Li, Y.; Feng, B.; Wang, B.; Sun, S. Joint planning of distributed generations and energy storage in active distribution networks: A Bi-Level programming approach. *Energy* **2022**, *245*, 123226. [CrossRef]
13. Mumtahina, U.; Alahakoon, S.; Wolfs, P. Optimal allocation and sizing of battery energy storage system in distribution network using mountain gazelle optimization algorithm. *Energies* **2025**, *18*, 379. [CrossRef]
14. Liu, X.; Zhao, P.; Qu, H.; Liu, N.; Zhao, K.; Xiao, C. Optimal Placement and Sizing of Distributed PV-Storage in Distribution Networks Using Cluster-Based Partitioning. *Processes* **2025**, *13*, 1765. [CrossRef]
15. Al-Mahammedi, M.T.F.; Onat, M. Optimal siting, sizing, and energy management of distributed renewable generation and storage under atmospheric conditions. *Sustainability* **2025**, *17*, 300. [CrossRef]
16. Bertsimas, D.; Litvinov, E.; Sun, X.A.; Zhao, J.; Zheng, T. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Trans. Power Syst.* **2013**, *28*, 52–63. [CrossRef]
17. Jing, Z.; Gao, L.; Mu, Y.; Liang, D. Flexibility-Constrained Energy Storage System Placement for Flexible Interconnected Distribution Networks. *Sustainability* **2024**, *16*, 9129. [CrossRef]
18. Zhang, H.; Zandehshahvar, R.; Tanneau, M.; Van Hentenryck, P. Weather-informed probabilistic forecasting and scenario generation in power systems. *Appl. Energy* **2025**, *384*, 125369. [CrossRef]
19. Hong, T.; Fan, S. Probabilistic electric load forecasting: A tutorial review. *Int. J. Forecast.* **2016**, *32*, 914–938. [CrossRef]
20. Telle, J.-S.; Upadhaya, A.; Schönfeldt, P.; Steens, T.; Hanke, B.; von Maydell, K. Probabilistic net load forecasting framework for application in distributed integrated renewable energy systems. *Energy Rep.* **2024**, *11*, 2535–2553. [CrossRef]
21. Dumas, J.; Wehenkel, A.; Lanaspeze, D.; Cornélusse, B.; Sutera, A. A deep generative model for probabilistic energy forecasting in power systems: Normalizing flows. *Appl. Energy* **2022**, *305*, 117871. [CrossRef]
22. Ma, C.; Xiong, W.; Tang, Z.; Li, Z.; Xiong, Y.; Wang, Q. Distributed mpc-based voltage control for active distribution networks considering uncertainty of distributed energy resources. *Electronics* **2024**, *13*, 2748. [CrossRef]
23. Van der Meer, D.W.; Widén, J.; Munkhammar, J. Review on probabilistic forecasting of photovoltaic power production and electricity consumption. *Renew. Sustain. Energy Rev.* **2018**, *81*, 1484–1512. [CrossRef]
24. Li, J.; Zhao, T.; Sun, D.; Ma, J.; Yu, H.; Yan, G.; Zhu, X.; Li, C. Multi-layer optimization method for siting and sizing of distributed energy storage in distribution networks based on cluster partition. *J. Clean. Prod.* **2025**, *501*, 145260. [CrossRef]

25. Xu, Z.; Tang, Z.; Chen, Y.; Liu, Y.; Gao, H.; Xu, X. Optimal robust allocation of distributed modular energy storage system in distribution networks for voltage regulation. *Appl. Energy* **2025**, *388*, 125625. [CrossRef]
26. Zhao, J.; Zhai, Q.; Zhou, Y.; Cao, X. Multi-stage robust scheduling of battery energy storage for distribution systems based on uncertainty set decomposition. *J. Energy Storage* **2024**, *92*, 112026. [CrossRef]
27. Yu, S.; Dong, X.; Wang, C.; Lu, T. Distributionally robust optimal scheduling of flexible distribution networks considering dynamic spatio-temporal correlation of renewable energy. *Int. J. Electr. Power Energy Syst.* **2025**, *169*, 110770. [CrossRef]
28. Hosseini, H.; Nezhad, A.E.; Estahbanati, M.S.J. Optimal siting and sizing of distributed generation units and energy storage systems for reconfigurable distribution networks. *J. Energy Storage* **2026**, *156*, 121555. [CrossRef]
29. Xu, R.; Jiang, Q.; Li, B.; Liu, Y.; Zhu, Y.; Gu, Y.; Liu, T. Frequency-Spectra Impedance Margin Ratio for Stability Analysis of IBR-Penetrated Systems. *IEEE Trans. Power Syst.* **2026**, 1–16. [CrossRef]