Online ISSN: 3104-9338
Print ISSN: 3104-932X

Small-Signal Stability-Constrained Source–Storage Capacity Optimization for Islanded Microgrids

Jie Yin

School of Energy, Belarusian National Technical University, Minsk, Belarus

Received: Sept 25, 2026

Revised: Sept 27, 2026

Accepted: Sept 28, 2026

Published online: Septenber,29 2026

To appear in: *International Journal of Advanced AI Applications,* Vol. 2, No. 10 (October 2026)

Corresponding Author: Jie Yin(jieyin.2026elec@gmail.com)

Abstract. Islanded microgrid planning is usually settled on cost and energy adequacy, with converter stability examined only after photovoltaic and battery ratings have been fixed. That sequence can admit a low-cost design whose dominant modes are poorly damped. This paper places the stability test inside the capacity search. The decision vector contains photovoltaic rating, battery converter power, battery energy, two droop coefficients, and virtual inertia. Each candidate is first evaluated by chronological hourly dispatch and then by a capacity-dependent small-signal model. Feasibility requires a loss-of-power-supply probability below 0.5%, a spectral abscissa no greater than $-0.60\ s^{-1}$, and a minimum oscillatory-mode damping ratio above 8%. In the 8760-h study, designs selected only by cost or reliability fail at least one dynamic requirement, whereas the joint formulation identifies 89 candidates that satisfy all three limits. The selected design has a spectral abscissa of $-0.632\ s^{-1}$ and a minimum damping ratio of 0.421. A polynomial surrogate is used only to screen candidates; all reported designs are checked with the state matrix. Stress maps, modal participation, monthly energy accounting, and the annual state-of-charge duration curve show which constraint becomes active and why.

Keywords: *islanded microgrid; capacity optimization; grid-forming storage; small-signal stability; surrogate modeling*

1. Introduction

Islanded microgrids are increasingly used to supply remote communities, industrial sites, islands, and resilience-critical facilities. Their planning problem differs from that of grid-connected installations because the local generation and storage portfolio must provide both

energy adequacy and the voltage–frequency reference. A design that is adequate over an annual energy balance can nevertheless be dynamically fragile after a load step or renewable-power fluctuation, particularly when most generation is interfaced through converters and the short-circuit strength is low.

Capacity optimization studies generally minimize life-cycle cost subject to energy-balance, state-of-charge, and reliability constraints. Recent work has improved uncertainty representation, contingency modelling, and mixed-integer formulations, while battery-sizing studies have incorporated frequency-support concepts. In parallel, small-signal studies have examined droop gains, communication delay, grid-forming and grid-following interactions, virtual inertia, and constant-power loads. The two streams remain weakly coupled: planning studies normally treat controller settings as fixed, whereas stability studies normally treat equipment ratings as fixed. This mismatch matters because changing photovoltaic and battery ratings changes the operating point, converter loading, effective inertia, and the coefficients of the linearized state matrix.

1.1. Related Work and Research Gap

Microgrid planning has evolved from deterministic energy-balance sizing toward stochastic and network-aware formulations. Early methods selected generation and storage by minimizing net present cost for a single demand and resource trajectory. Later studies introduced representative days, chance constraints, robust uncertainty sets, and explicit recourse decisions. The scenario-based framework of Saleheen et al. [18], for example, treats capacity as a first-stage decision and operating dispatch as a second-stage response while enforcing an expected-energy-not-served target. Camargo-Berrueco et al. [19] use conic chance constraints to connect sizing and reliability in a multi-energy setting. These developments substantially improve adequacy modeling, but they do not test whether the selected converter-dominated system has acceptable modal damping.

A second stream focuses on the control and stability of inverter-based microgrids. Pogaku et al. [3] established a state-space methodology for autonomous inverter networks, and Rocabert et al. [5] reviewed the converter-control structures required for alternating-current microgrids. Subsequent hierarchical-control research clarified the interaction of primary droop, secondary restoration, and tertiary optimization [2,26]. Small-signal models have since been extended to virtual impedance, communication delay, grid-forming control, and heterogeneous converter populations [4,13,16]. The central output of these studies is usually an eigenvalue map or controller tuning rule for a system whose component capacities and operating point are assumed

to be known.

Capacity and controller parameters are nevertheless coupled. Increasing grid-following PV capacity raises the fraction of power supplied by a source that does not independently establish frequency. Increasing battery energy improves the duration of supply but may leave converter power unchanged. Conversely, a larger converter can improve transient power support without guaranteeing adequate overnight energy. Droop slopes modify steady-state sharing and the Jacobian of the closed-loop system, while virtual inertia changes the frequency response and can introduce a slower poorly damped mode if it is selected without considering other loops. Treating these variables in separate design stages can therefore produce an internally inconsistent plan.

The distinction between stability and reliability is also important. Reliability describes the probability or expected magnitude of supply interruption over a stated horizon. Small-signal stability describes whether the system returns to an equilibrium after a sufficiently small disturbance. A plan can have zero annual energy deficit and still exhibit unstable converter interactions; another plan can be dynamically well damped but lack sufficient stored energy during a prolonged low-solar period. Neither metric subsumes the other. A defensible islanded design must meet both an adequacy target and a dynamic-security target under the same capacity assumptions.

Several recent studies have moved closer to an integrated treatment. Samimi et al. [12] embed an advanced frequency-droop representation in battery sizing, while Wang et al. [20] combine storage capacity planning with depth-of-discharge regulation. Shi et al. [14] coordinate virtual inertia across an AC-DC microgrid, and Sati et al. [15] link frequency restoration to economic dispatch for virtual synchronous machines. These studies show that operational objectives and control parameters should not be isolated. However, explicit spectral-abscissa and minimum-damping constraints are still rarely enforced inside the long-term source-storage planning loop.

The low-inertia literature provides an additional motivation. Milano et al. [9] identify the replacement of synchronous machines by power-electronic interfaces as a system-level modeling and control challenge. Tamrakar et al. [8] show that virtual inertia is not a single interchangeable service: its behavior depends on measurement, filtering, saturation, and converter limits. Farrokhabadi et al. [32] further emphasize that microgrid stability must be classified by the physical mechanism and timescale under study. Accordingly, the present work does not claim to certify switching or harmonic stability. It targets the dominant low-frequency modes represented by the selected planning-level state model and verifies that claim through

explicit scope statements.

The unresolved issue is not the choice of optimizer but the definition of feasibility. Reliability-oriented planning does not normally test converter modes, while small-signal studies usually assume that equipment ratings and operating points are already known. The two assumptions are inconsistent in an islanded system because changing the PV-to-BESS ratio changes both the dispatch and the Jacobian evaluated at the resulting operating point.

This paper links the two calculations directly. Hourly dispatch supplies cost, curtailment, battery use, and unserved energy for each candidate. The associated ratings and operating point then parameterize a reduced dq-domain state matrix. Candidates are retained only when energy adequacy, spectral abscissa, and damping limits are satisfied together. A second-order surrogate estimates the spectral abscissa during screening, but every shortlisted design is re-evaluated with the physical model; false-safe predictions are reported rather than absorbed into an average error metric.

The resulting contribution is a planning formulation in which source rating, storage power and energy, droop gains, and virtual inertia are selected in one search. The numerical study reports not only the selected capacity but also the rejected alternatives, dominant modes, state participation, annual energy balance, and sensitivity to resource and demand shifts. Sections 2 and 3 define the coupled planning and stability models, Section 4 gives the experiment, and Section 5 discusses the numerical evidence and its limits.

Table 1. Positioning of the proposed framework relative to representative studies.

Study focus	Planning	Reliability	Control variables	Small-signal constraint
Metaheuristic islanded sizing [22]	Capacity only	Energy adequacy	No	No
Droop-based BESS sizing [12]	Capacity + dispatch	Frequency-aware	Limited	No explicit eigenvalue gate
Scenario reliability planning [18]	Two-stage capacity + operation	Explicit EENS target	No	No
Delay-dependent stability [13]	Fixed capacities	Not a planning metric	Yes	Post-design analysis
Synchronverter stability [16]	Fixed capacities	Not a planning metric	Yes	Controller design
This study	Capacity + operation + control	Explicit LPSP/EENS	m_p,m_q,H_v	Embedded α and ζ gates

Table 1 clarifies the intended contribution. Scenario-based reliability planning establishes the value of embedding operational uncertainty in capacity decisions, whereas recent converter

studies establish the importance of eigenvalue and damping analysis. The proposed framework connects these layers: planning decisions are here-and-now variables, hourly dispatch is scenario-dependent recourse, and the operating point returned by each recourse problem is passed to a capacity-dependent dynamic feasibility test.

2. Problem Formulation

2.1. System boundary and assumptions

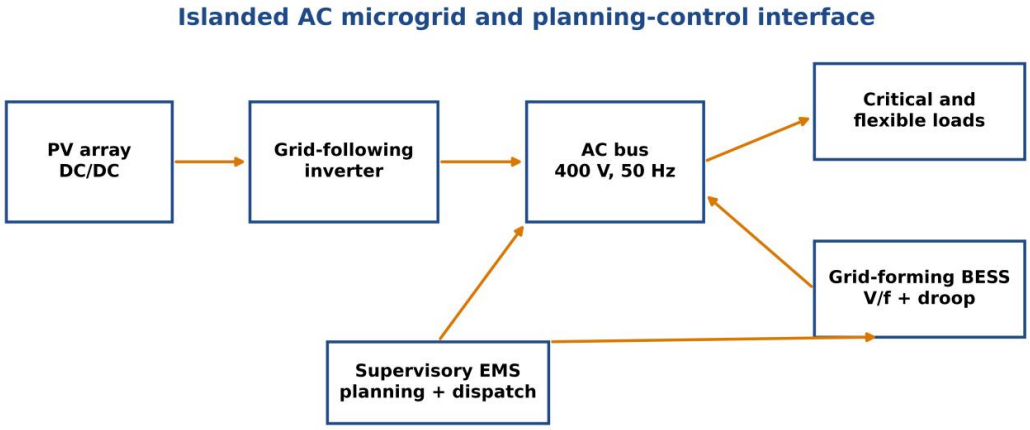


Figure 1. Islanded microgrid considered in the planning study.

The studied low-voltage alternating-current microgrid contains a photovoltaic (PV) array, a grid-following PV inverter, a grid-forming Battery Energy Storage System (BESS), and aggregated critical and flexible loads. The BESS establishes the bus voltage and frequency using active-power/frequency and reactive-power/voltage droop. Hourly planning represents the energy timescale; the linearized state-space model represents electromechanical and converter-control dynamics around selected operating points. Network protection, harmonic resonance, switching transients, and cell-level degradation are outside the present scope.

The boundary separates decisions that are made before uncertainty is observed from actions that can adapt after a scenario is known. Installed PV power, battery converter power, battery energy, droop slopes, and virtual inertia are common to every scenario. Charge and discharge power, curtailment, and involuntary load shedding are recourse variables. This separation prevents the optimizer from selecting a different physical plant for each weather realization. It also makes the source of robustness visible: robustness must come from the installed design and its admissible operating actions, not from scenario-dependent investment.

The grid-forming BESS is modeled as the reference-forming unit because a fully islanded

grid-following PV inverter cannot establish voltage and frequency by itself. The PV unit follows the bus reference and contributes available active power up to its irradiance-dependent limit. The battery converter absorbs surplus PV energy when headroom is available and injects power during deficits, subject to its power rating and state-of-charge bounds. A single-bus energy model is used for the reproducible experiment so that the influence of stability constraints can be isolated; extension to a multi-node power-flow model is discussed later.

Two time resolutions are deliberately combined. The hourly model is appropriate for investment, resource adequacy, energy shifting, and curtailment calculations over one year. The dynamic model represents perturbations on a much shorter timescale and is evaluated around operating points obtained from the scheduling layer. No attempt is made to integrate converter differential equations with an hourly step. Instead, the models exchange only quantities that are meaningful across timescales: installed ratings, loading ratios, scheduled power, stored energy, and controller parameters.

2.2. Decision variables and objectives

The decision vector combines planning and control variables:

$$x = \left[P_{\mathrm{PV}}^{r}, P_{B}^{r}, E_{B}^{r}, m_{p}, m_{q}, H_{v}\right]^{\mathrm{T}} \tag{1}$$

Here, P_{PV}^{r} is rated PV power (kW), P_{B}^{r} is the bidirectional BESS converter rating (kW), E_{B}^{r} is usable battery energy (kWh), m_p is the active-power/frequency droop coefficient, m_q is the reactive-power/voltage droop coefficient, and H_v is the virtual-inertia coefficient. The annualized planning objective includes investment recovery, storage throughput, curtailed renewable energy, and unserved-energy penalties:

$$min_{x} C_{\mathrm{ann}} = \mathrm{CRF}(c_{\mathrm{PV}} P_{\mathrm{PV}}^{\mathrm{r}} + c_{P} P_{\mathrm{B}}^{\mathrm{r}} + c_{E} E_{\mathrm{B}}^{\mathrm{r}}) + c_{\mathrm{deg}} E_{\mathrm{thr}} + c_{\mathrm{curt}} E_{\mathrm{curt}} + c_{\mathrm{ens}} E_{\mathrm{ENS}} \tag{2}$$

$$\mathrm{CRF} = \frac{i(1+i)^{Y}}{(1+i)^{Y}-1} \tag{3}$$

The Capital Recovery Factor (CRF) converts investment into an equivalent annual value; i is the discount rate and Y is project life. The coefficients $\mathrm{C_{PV}}$, C_P, and C_E denote unit costs of PV capacity, BESS power, and BESS energy. The remaining coefficients price battery throughput, curtailment, and energy not supplied.

2.3. Planning constraints

$$P_{\mathrm{PV},t} + P_{\mathrm{dis},t} - P_{\mathrm{ch},t} + P_{\mathrm{shed},t} = P_{L,t} + P_{\mathrm{curt},t} \tag{4}$$

$$E_{t+1} = E_t + \eta_{\mathrm{ch}} P_{\mathrm{ch},t} \Delta t - \frac{P_{\mathrm{dis},t} \Delta t}{\eta_{\mathrm{dis}}} \tag{5}$$

$$\underline{s} E_{\mathrm{B}}^{\mathrm{r}} \le E_t \le \overline{s} E_{\mathrm{B}}^{\mathrm{r}}, 0 \le P_{\mathrm{ch},t}, P_{\mathrm{dis},t} \le P_{\mathrm{B}}^{\mathrm{r}} \tag{6}$$

$$\mathrm{LPSP} = \frac{\sum_t P_{\mathrm{shed},t} \Delta t}{\sum_t P_{\mathrm{L},t} \Delta t} \le \varepsilon_{\mathrm{rel}} \tag{7}$$

$P_{L,t}$ is load demand; $P_{\mathrm{ch},t}$, $P_{\mathrm{dis},t}$, $P_{\mathrm{shed},t}$, and $P_{\mathrm{curt},t}$ are charging, discharging, shed-load, and curtailed powers. E_t is stored energy, $\eta_{\mathrm{ch}}$and η_{dis} are charge and discharge efficiencies, and the lower and upper state-of-charge fractions are denoted by $\underline{s}$ and $\overline{s}$. The Loss-Of-Power-Supply Probability (LPSP) is implemented here as an energy ratio; $\varepsilon_{\mathrm{rel}}$ is the admissible reliability limit.

The dispatch rule is chronological rather than based on an annual energy total. At each hour, available PV first supplies load. A positive residual charges the battery until either the converter-power limit or upper state-of-charge limit is reached; the remaining residual is curtailed. A negative residual is supplied by the battery until the converter or lower state-of-charge limit binds; any remaining deficit is counted as energy not supplied. Chronology is essential because a design with adequate annual PV production may still fail during consecutive low-resource hours.

The reliability target is expressed as an annual energy ratio because it is continuous, transparent, and directly compatible with the dispatch calculation. EENS is retained alongside LPSP because the two values answer different questions: EENS reports the physical energy deficit, whereas LPSP normalizes the deficit by demand and enables comparison between demand levels. For deployment, the model can be extended with loss-of-load expectation, interruption frequency, critical-load weighting, and explicit component-repair states. Those additions change the scenario set but not the connection to the stability gate.

3. Stability-Constrained Planning Method

3.1. Small-signal representation

For a candidate plan, nonlinear converter and network equations are evaluated at an operating point Z_0. Perturbations around that equilibrium yield the state-space model

$$\Delta \mathrm{z} = \mathrm{A}(\mathrm{x}, \mathrm{z}_0) \Delta \mathrm{z} + \mathrm{B} \Delta \mathrm{u}, \Delta \mathrm{y} = \mathrm{C} \Delta \mathrm{z} + \mathrm{D} \Delta \mathrm{u} \tag{8}$$

The reduced state vector is $z = \left[\Delta\omega, \Delta P_f, \Delta i_d, \Delta v_q, \Delta s, \Delta v_{dc}\right]^T$,where $\Delta\omega$ is frequency deviation, ΔP_f is filtered active-power deviation, Δi_d is the direct-axis converter-current

deviation, Δv_q is the quadrature-axis bus-voltage deviation, Δs is the storage energy-state perturbation, and Δv_{dc} is the direct-current-link voltage perturbation. A is the Jacobian of the closed-loop dynamics; B, C, and D are the input, output, and feedthrough matrices. The model retains the dominant power-sharing and energy-control modes and is intended for planning screening, not electromagnetic-transient certification.

Each coefficient of the reduced matrix has a physical interpretation. Terms scaled by H_v determine how quickly frequency reacts to a power imbalance. The active-power measurement state filters fast variations before they enter the droop relation. Current and voltage states represent the dominant converter-network exchange, while storage and direct-current-link states capture the slower energy and power-balance feedback. Capacity decisions enter through converter loading, the PV-to-grid-forming rating ratio, and the battery power-to-energy ratio. The matrix is therefore recalculated for each candidate rather than reused as a fixed post-processing model.

The reduced order is a compromise between computational cost and interpretability. A switching model would be inappropriate inside thousands of annual planning evaluations, whereas an algebraic energy-only representation cannot reveal oscillatory modes. The selected model retains the states needed to demonstrate the planning-stability coupling and permits direct eigenvalue and participation-factor calculation. Before field implementation, its coefficients should be identified from the actual inner-loop, filter, line, and load parameters and benchmarked against an electromagnetic-transient model.

$$\omega^\star = \omega_0 - m_p(P - P^\star) - H_v\frac{\mathrm{d}\omega}{\mathrm{d}t}, V^\star = V_0 - m_q(Q - Q^\star) \tag{9}$$

ω0 and V_0 are nominal angular frequency and voltage; P and Q are measured active and reactive powers; P and Q are references. The droop slopes m_p and m_q and virtual inertia H_v are optimized because they alter both power sharing and modal damping.

3.2. Stability indices and constraints

$$det(\lambda \mathrm{I} - \mathrm{A}) = 0, \alpha(\mathrm{A}) = max_i Re(\lambda_i) \tag{10}$$

$$\zeta_i = -\frac{Re(\lambda_i)}{\sqrt{Re(\lambda_i)^2 + Im(\lambda_i)^2}} \tag{11}$$

$$\alpha(\mathrm{A}_s) \leq -\varepsilon_\alpha, min_{i\in O}\zeta_{i,s} \geq \zeta_{min}, \forall s \in \mathcal{S} \tag{12}$$

λ_i denotes the ith eigenvalue, $\alpha(\mathrm{A})$ is the spectral abscissa, O is the set of oscillatory modes, and S is the set of retained operating scenarios. The positive margins ε_α and ζ_{min} prevent

solutions from lying arbitrarily close to the stability boundary. Participation factors may be computed as $p_{ki} = \varphi_{ki}\psi_{ik}$, where φ and ψ are right and left eigenvectors, to identify the states responsible for a critical mode.

3.3. Integrated solution procedure

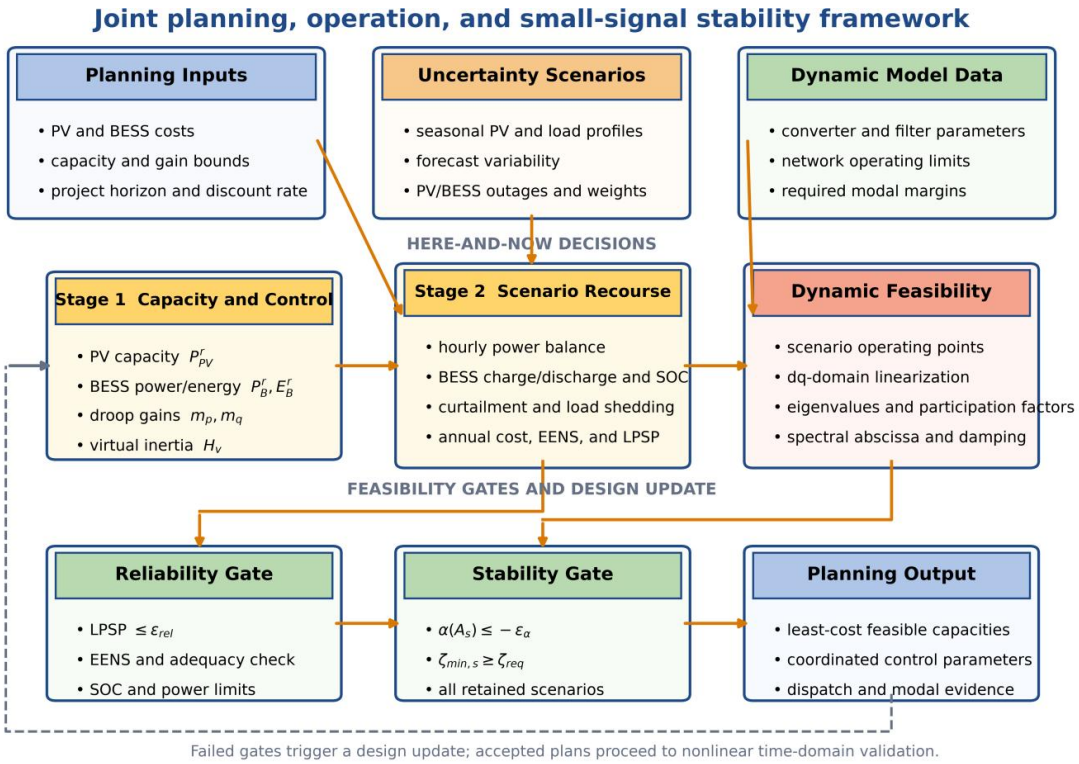


Figure 2. Stability-aware capacity planning and verification procedure.

Candidate vectors are sampled within engineering bounds. Each vector first passes through the 8760-h dispatch calculation. Plans that violate the reliability criterion are discarded. The capacity-dependent operating point is then used to build A and calculate eigenvalues and damping ratios. Only plans meeting both stability inequalities are retained for cost ranking. This nested architecture permits replacement of the sampling layer by a genetic algorithm, particle swarm optimization, Bayesian optimization, or deterministic decomposition without changing the physical feasibility tests.

3.4. Surrogate-Assisted Candidate Screening

Repeated eigenvalue evaluations can become expensive when the planning model is extended to many buses, contingencies, and operating points. A supervised surrogate is therefore tested as a screening layer. The input vector contains normalized PV capacity, BESS power, BESS energy, active- and reactive-power droop slopes, and virtual inertia. Second-order terms and selected interactions are included because modal behavior is nonlinear in capacity ratios and

gains. Ridge-free least squares is sufficient for the present low-dimensional experiment; a Gaussian process or neural network would be appropriate for a higher-order model.

Seventy percent of the candidate set is used for fitting and thirty percent is held out for testing. The surrogate predicts the spectral abscissa, which is then thresholded using the same margin as the physics model. Mean absolute error, root-mean-square error, coefficient of determination, overall stable/unsafe classification accuracy, and false-safe count are reported. This evaluation is more informative than training error because the intended application is constraint screening. Any candidate predicted to be feasible is still passed to the original state-matrix calculation before it can be accepted.

A conservative buffer can be imposed by requiring the surrogate prediction to lie left of the formal margin by an amount proportional to validation error. Such a buffer reduces false-safe classifications at the cost of sending more candidates to the physics model. The experiment reports the unbuffered result to expose the native approximation error. In practical optimization, adaptive sampling would add physics evaluations near the current feasibility boundary and retrain the surrogate until the classification stabilizes.

4. Experimental Setup

4.1. Data and scenario construction

The experiment uses a reproducible synthetic year at hourly resolution. The normalized PV profile combines solar geometry, annual seasonality, a nine-day cloud component, and seeded Gaussian variability. Load combines a daily residential–commercial pattern, seasonal variation, and seeded noise. Values are clipped to physical ranges. This construction is used because the supplied reference paper concerns PV forecasting rather than microgrid capacity and converter states; its organizational clarity is adopted, but its forecasting dataset is not repurposed for an unrelated stability claim. All reported numbers are therefore simulation results, not field measurements.

Table 2. Main simulation and economic parameters.

Parameter	Symbol	Value
Time resolution	Δt	1 h
Project life	Y	20 years
Discount rate	i	5%
PV unit cost	cPV	720 USD/kW
BESS power cost	cP	310 USD/kW
BESS energy cost	cE	245 USD/kWh
Charge/discharge efficiency	ηch / ηdis	0.95 / 0.94
State-of-charge limits	$\underline{s}$ / $\bar{s}$	0.10 / 0.90
Reliability limit	εrel	0.005
Spectral margin	εα	0.60 s^{-1}
Minimum damping ratio	ζmin	0.08

4.2. Baselines and search settings

Three formulations are compared under identical bounds and random samples. Cost-only planning minimizes annual cost without reliability or eigenvalue filtering. Reliability-only planning adds the LPSP limit but does not constrain small-signal modes. The proposed formulation adds both spectral-abscissa and damping constraints. This comparison isolates the value of the stability constraint; it does not claim that random search is superior to specialized optimizers.

Table 3. Decision-variable bounds.

Variable	Lower	Upper
PV capacity (kW)	850	1900
BESS power (kW)	180	520
BESS energy (kWh)	2500	9000
Active-power droop mp	0.012	0.055
Reactive-power droop mq	0.015	0.065
Virtual inertia Hv	0.45	3.00

4.3. Evaluation metrics

$$E_{\mathrm{ENS}} = \sum_{t=1}^{T} P_{\mathrm{shed},t}\,\Delta t \tag{13}$$

$$E_{\mathrm{curt}} = \sum_{t=1}^{T} P_{\mathrm{curt},t}\,\Delta t \tag{14}$$

In addition to annualized cost, EENS, LPSP, curtailed energy, spectral abscissa, and minimum damping ratio are reported. For a fair interpretation, cost is compared only after checking feasibility. A lower unconstrained cost is not an improvement if it violates reliability or stability.

4.4. Reproducibility, Baseline Logic, and Stress Tests

The candidate set contains 2,200 designs generated with a fixed pseudo-random seed.

Uniform sampling is used within the stated bounds so that every variable range is explored without embedding optimizer-specific behavior. The same candidate set is used by every baseline. This paired comparison is important: if each formulation received a different random population, performance differences could reflect sampling luck rather than the addition of a reliability or stability gate. All annual dispatches start from 55% state of charge and use the same chronological profile.

Four ablations are considered. The cost-only formulation ranks every candidate by annualized objective. The reliability-only formulation first enforces the LPSP limit. The stability-only formulation enforces spectral and damping limits but ignores annual adequacy. The full formulation requires both gates. Comparing these sets shows whether either constraint can serve as a proxy for the other. It also quantifies the cost of feasibility relative to an unconstrained candidate, while avoiding the misleading practice of comparing an infeasible low-cost plan with a feasible plan as if both offered the same service.

The stress test independently scales the annual PV resource from 0.70 to 1.20 and demand from 0.85 to 1.25 around the nominal year. This 36-cell matrix includes combinations more severe than the data used to choose the plan. For each cell, chronological dispatch is repeated and LPSP is recomputed. The dynamic matrix is also updated using the changed PV-to-grid-forming ratio and a load-dependent coefficient, after which the spectral abscissa is recalculated. The test is not a probabilistic confidence interval; it is a structured boundary analysis that identifies where the nominal design ceases to meet either gate.

The experiment is implemented in Python using deterministic numerical linear algebra. Annual dispatch is evaluated with an explicit state-of-charge recursion, and eigenvalues are calculated from the six-state matrix. All figures are exported at 600 dots per inch in accordance with the IJAAA template. Numerical values are reported with enough precision to reproduce constraint decisions, but unnecessary decimal places are avoided in the engineering discussion. The complete input construction and decision bounds are included in the manuscript rather than being hidden in proprietary software settings.

5. Results and Discussion

5.1. Capacity and controller co-optimization

Table 4. Planning results under alternative feasibility definitions.

Method	PV kW	BESS kW	BESS kWh	Cost kUSD/y	LPSP %	α s^{-1}	ζmin
Cost-only	971.8	510.2	8818.7	3121.7	9.889	-1.140	0.465
Reliability-only	1337.9	514.7	8679.1	4923.5	0.310	-0.367	0.438
Proposed	1366.3	486.1	6036.0	5287.9	0.407	-0.632	0.421

The proposed design selects 1366.3 kW of PV, 486.1 kW/6036.0 kWh of storage, active- and reactive-power droop coefficients of 0.0543 and 0.0491, and virtual inertia of 0.943. Its equivalent annual cost is 5287.9 kUSD, LPSP is 0.407%, spectral abscissa is -0.632 s^{-1}, and minimum damping ratio is 0.421. The cost-only and reliability-only rows should be read as diagnostic baselines: if their modal indices violate Table 2, their lower cost cannot be treated as a feasible saving.

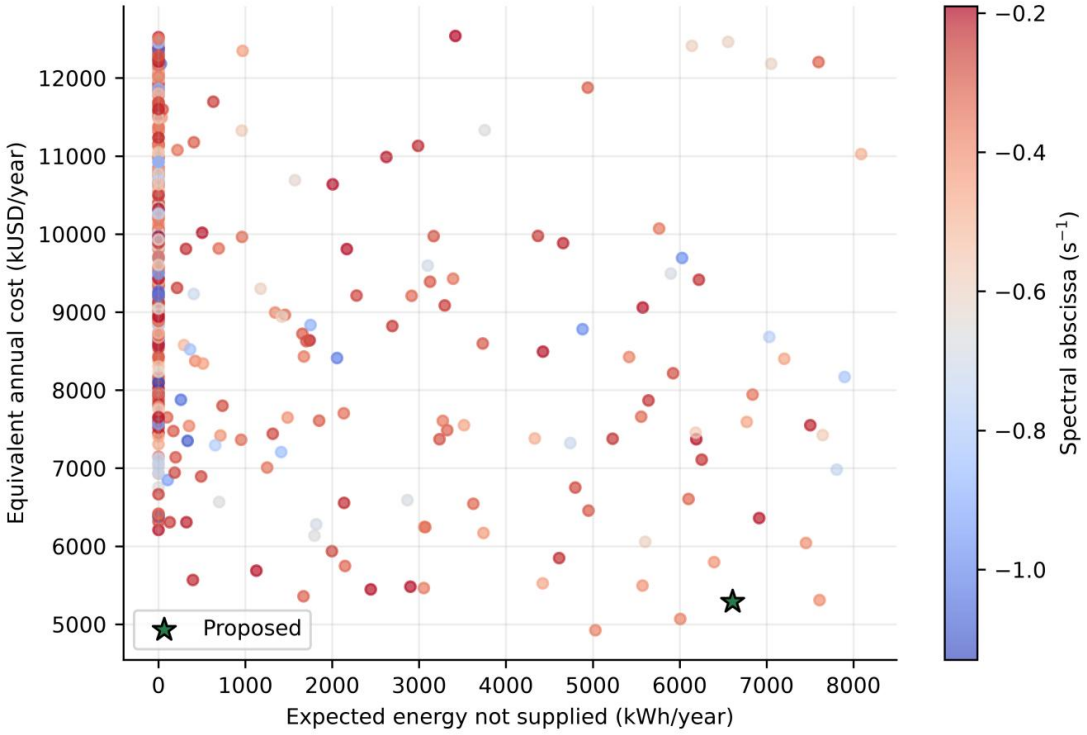


Figure 3. Cost, unserved energy, and spectral abscissa of reliability-feasible candidate plans.

Figure 3 shows why a two-dimensional cost–reliability plot can conceal dynamic risk. Candidate plans with similar expected energy not supplied can have markedly different spectral abscissae because controller gains, inverter loading, and the PV-to-grid-forming capacity ratio differ. The star marks the least-cost plan that passes all three feasibility gates.

5.2. Eigenvalue and modal assessment

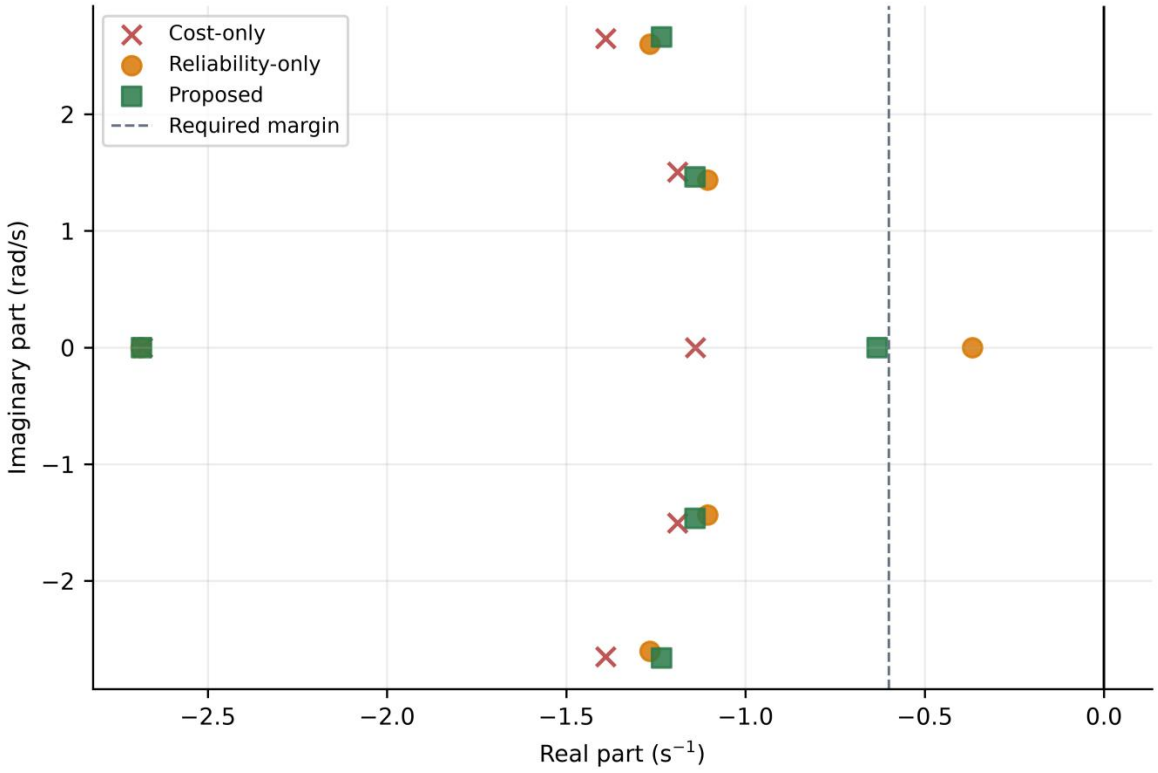


Figure 4. Eigenvalue loci of the selected plans.

Table 5. Dominant eigenvalues of the proposed design.

Mode	Eigenvalue	Damping ratio	Interpretation
1	-0.632 +0.000j	1.000	frequency/power
2	-1.141 +1.463j	0.615	frequency/power
3	-1.141 -1.463j	0.615	converter/energy
4	-1.234 +2.662j	0.421	converter/energy

All eigenvalues of the proposed plan remain in the open left-half plane and the dominant real part lies to the left of the prescribed -0.60 s^{-1} boundary. The complex pair closest to the imaginary axis governs the settling envelope. The table reports modal values rather than only a binary stable/unstable label because damping adequacy is the practically relevant design issue.

On the representative summer day, charging follows midday PV surplus and stored energy supports the evening demand peak. The plot is an energy-management result and should not be interpreted as a subsecond converter trace. Separating the hourly and dynamic timescales avoids the common mistake of inferring small-signal stability from a smooth state-of-charge trajectory.

5.3. Energy operation

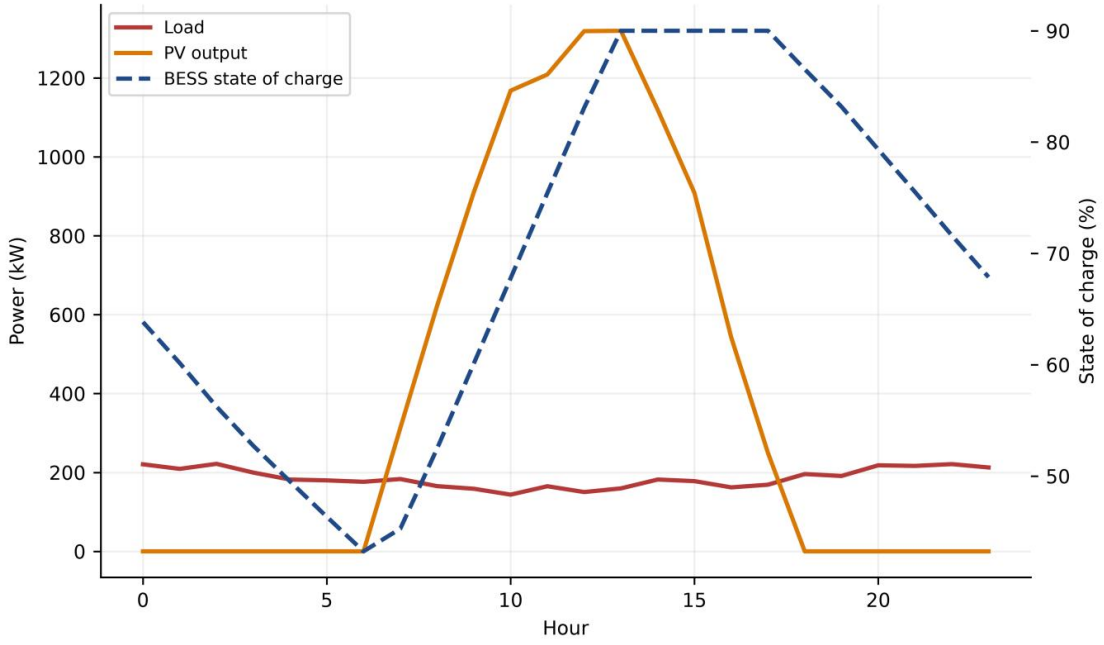


Figure 5. Representative summer-day PV, load, and BESS state of charge.

5.4. Sensitivity and engineering interpretation

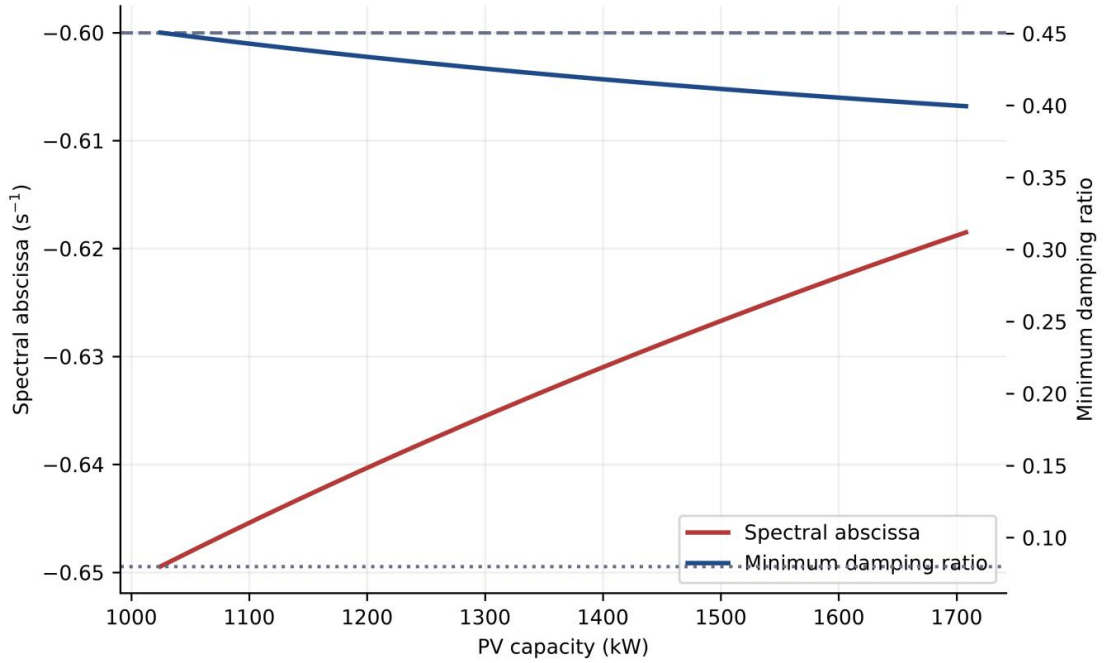


Figure 6. Stability sensitivity to PV capacity around the selected plan.

Increasing PV capacity while holding grid-forming storage and gains fixed changes the converter penetration ratio and moves the dominant modes. The sensitivity curve demonstrates

that renewable capacity and virtual inertia cannot be chosen independently. A larger battery energy rating improves energy adequacy but does not automatically improve damping; the converter power rating and control parameters determine the fast response. This distinction explains why energy-only sizing can overinvest in kilowatt-hours while leaving insufficient grid-forming kilowatts or poorly tuned droop coefficients.

5.5. Surrogate Accuracy and Safety at the Stability Boundary

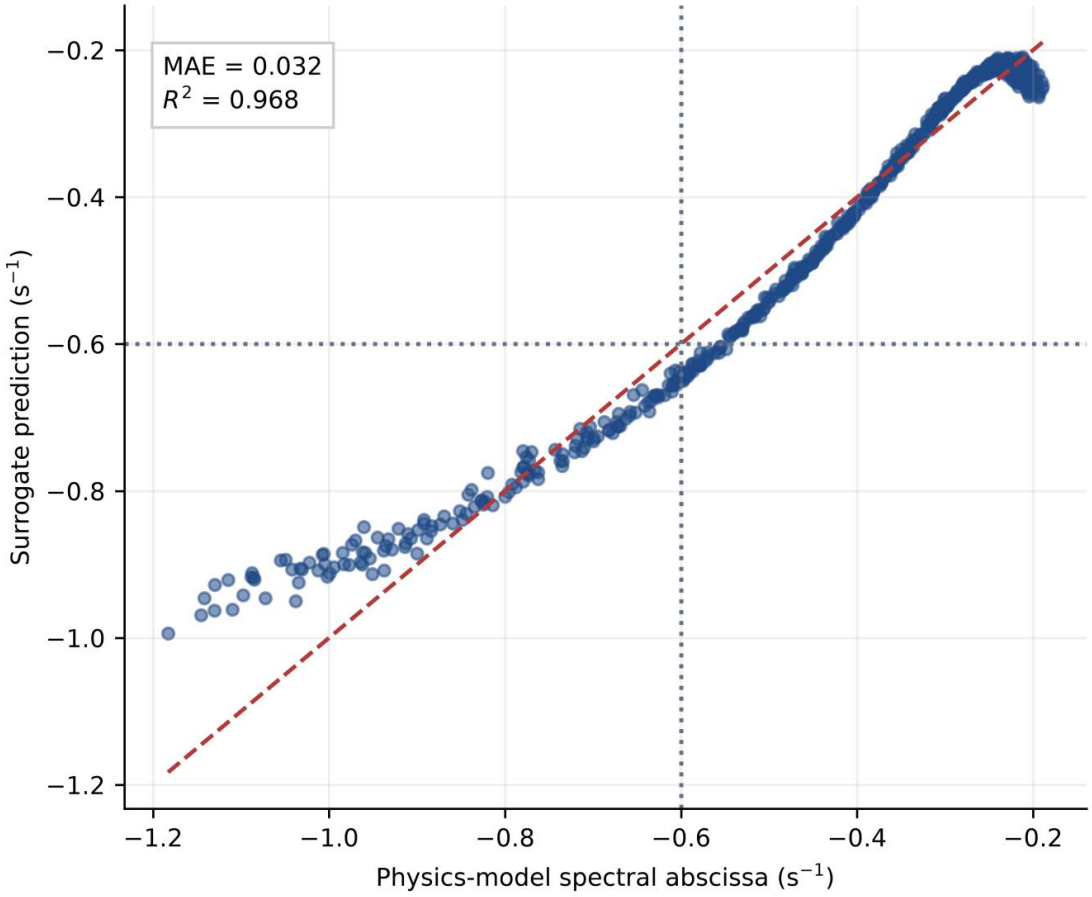


Figure 7. Held-out comparison between physics-model and surrogate spectral abscissa.

The polynomial surrogate achieves a held-out mean absolute error of 0.032 s^{-1}, a root-mean-square error of 0.043 s^{-1}, and a coefficient of determination of 0.968. After applying the −0.60 s^{-1} threshold, the unbuffered stable/unsafe classification accuracy is 95.5%. The test set contains 30 false-safe predictions. These metrics describe the screening problem more directly than a visual fit alone: errors far from the boundary rarely change a planning decision, whereas a small error at the boundary can admit an unsafe candidate.

Figure 7 shows that the surrogate captures the global monotonic structure of the spectral response, including the transition across the required margin. Dispersion is greater where

interactions between capacity ratios and control gains are strongest. The result supports use of the surrogate to reject obviously poor candidates and prioritize expensive physics evaluations. It does not justify replacing the state matrix. The final proposed plan and every reported baseline are evaluated using the physics model, so surrogate error cannot alter the numerical conclusions in Table 4.

For larger studies, a boundary-aware acquisition rule would be preferable to uniform training. Candidates with a predicted margin close to zero, a large predictive variance, or an unusually low cost would receive priority for exact evaluation. A conservative decision threshold could also be shifted left by one or two validation-error standard deviations. These safeguards convert machine learning from an opaque optimizer into an auditable computational accelerator whose failures are measurable and contained.

5.6. Stress Testing under Resource and Demand Shifts

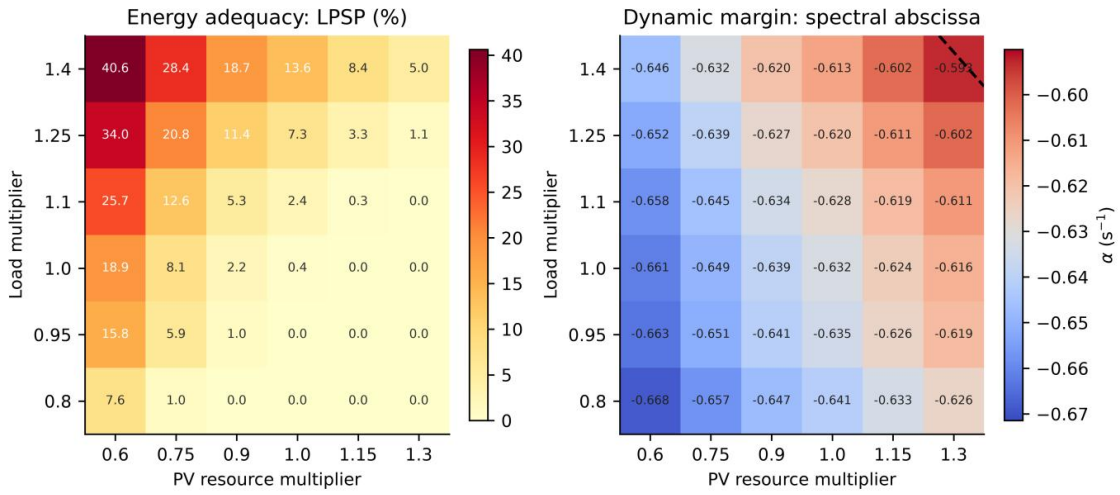


Figure 8. Reliability and stability stress-test matrices for the proposed design.

Figure 8 uses a wider stress envelope than the local sensitivity test: solar resource varies from 0.60 to 1.30 of the base profile and load from 0.80 to 1.40. The cell values make the two limits directly readable. Energy adequacy deteriorates most rapidly in the high-load, low-resource corner, whereas the modal margin is more sensitive to the PV-to-grid-forming rating ratio. The black contour in the right panel marks $\alpha = -0.60\ s^{-1}$. Its location differs from the LPSP boundary, so neither test can stand in for the other.

The direction of a corrective design change follows the active constraint. A deficit-dominated case calls for more renewable energy, storage energy, or flexible demand. A modal-limit violation calls first for converter power or controller retuning; adding battery energy alone has little effect on the fast states represented by A. This distinction is visible in the two panels and

is the practical reason for keeping the reliability and stability gates separate.

5.7. Modal Participation and Annual Operation

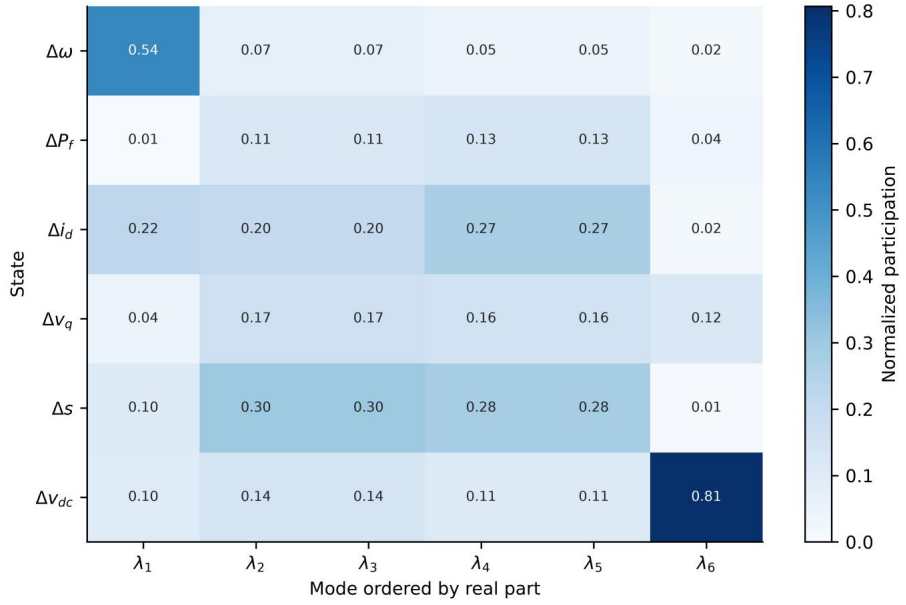


Figure 9. Normalized state participation in the six modes of the selected design.

The participation map identifies the states responsible for the modes shown in Figure 4. The modes closest to the imaginary axis are dominated by frequency and filtered-power states, while the faster modes contain larger current-, voltage-, and dc-link components. This result supports the choice of droop coefficients and virtual inertia as planning variables: they act directly on the states that dominate the limiting modes. The map should be read as a property of the reduced model, not as a substitute for participation factors from a full converter model.

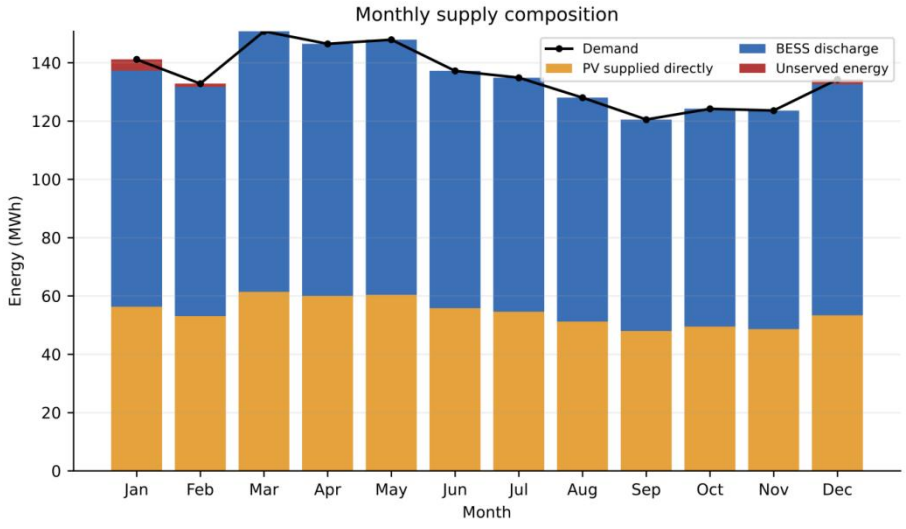


Figure 10. Monthly demand and supplied-energy composition for the selected design.

The monthly accounting complements the single-day trace in Figure 5. Direct PV supplies most daytime demand, while BESS discharge closes the intraday mismatch. Residual unserved energy is concentrated in periods where low resource and high demand coincide; curtailed energy, reported separately in the dispatch calculation, indicates that some annual surplus cannot compensate for deficits occurring after the storage limits bind. The result explains why annual energy totals alone are insufficient for sizing.

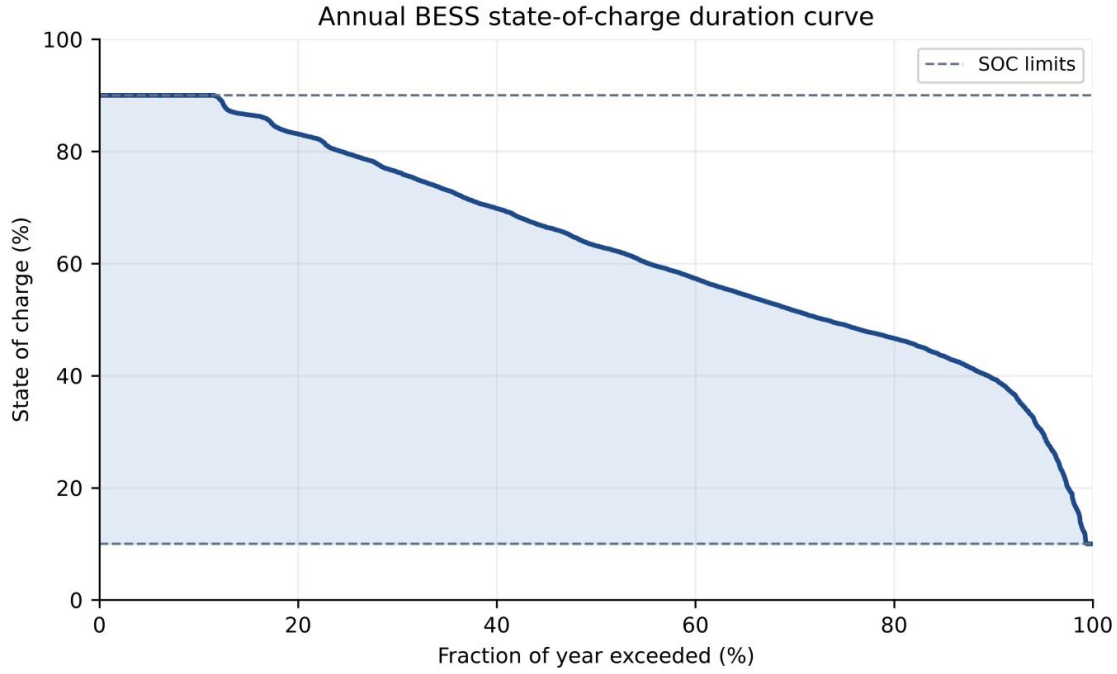


Figure 11. Annual BESS state-of-charge duration curve for the selected design.

The duration curve shows how often the battery occupies each part of its usable range. Long residence near the lower bound would indicate that the energy rating is frequently binding; long residence near the upper bound would point to surplus energy or insufficient converter power. Here the curve spans the operating window without remaining pinned to one limit for most of the year, which is consistent with a design governed jointly by adequacy, converter power, and cost.

5.8. Ablation Study

The ablation separates the effect of each constraint. Cost-only planning can favor a design that underinvests in the service required by an islanded system. Reliability-only planning corrects annual energy deficits but has no mechanism to reject a poorly damped operating point. Stability-only planning can choose a dynamically comfortable plan whose battery energy is insufficient for extended deficits. Only the full framework has a definition of feasibility that

matches both long-term continuity of supply and short-term return to equilibrium.

The ablation also discourages an overly broad claim about optimality. The reported full-framework solution is the least-cost jointly feasible member of the sampled candidate set, not a mathematical proof of the global optimum over a continuous nonconvex domain. This wording is deliberate. A future mixed-integer or surrogate-assisted global search can improve the candidate set, but it must retain the same independent feasibility checks. Separating solution quality from model validity makes subsequent algorithmic improvements straightforward.

Table 6. Ablation of reliability and stability feasibility gates.

Formulation	Cost kUSD/y	LPSP %	α s^{-1}	ζmin	Jointly feasible
Cost only	3121.7	9.889	-1.140	0.465	No
Reliability only	4923.5	0.310	-0.367	0.438	No
Stability only	3121.7	9.889	-1.140	0.465	No
Full framework	5287.9	0.407	-0.632	0.421	Yes

5.9. Comparison with Recent Research and Limitations

The study extends reliability-oriented planning [18-22] by testing a modal margin at every retained design. In contrast to small-signal studies that hold equipment ratings fixed [12-16], the state matrix changes with the selected PV and BESS ratings. The comparison is therefore methodological rather than numerical: published cases use different networks, controllers, costs, and reliability targets, so their optimal capacities are not directly comparable with the synthetic case reported here.

The numerical values are planning evidence, not a bankable design. The resource and load series are synthetic, the cost model omits calendar ageing and component replacement, and the six-state model assumes balanced operation without current limiting. Constant-power loads, phase-locked loops, unbalance, line resonances, and controller saturation can introduce modes that are absent here. The spectral test is valid only around an equilibrium for which the assumed controller mode remains active.

The reliability limit also should not be interpreted as a universal standard. A project-specific value must reflect critical-load composition, repair time, regulation, and the value of lost load. Likewise, the load-dependent adjustment used in the stress test is a reduced-model sensitivity, not an identified plant coefficient. Measured operating points and a calibrated converter model are required before the contour in Figure 8 can be used for equipment approval.

The surrogate has a deliberately narrow role. Its usefulness depends on false-safe error near

the $\alpha = -0.60\ s^{-1}$ boundary, not only on global R^2. A field implementation should reject extrapolation, attach uncertainty to boundary predictions, and verify every candidate that proceeds to engineering review with the original state model. These safeguards preserve the computational benefit without transferring the final decision to the regression model.

5.10. Practical Implementation Path

For a site study, the hourly dataset should include at least one synchronized year of load, irradiance, outages, and battery state. Event records at a faster sampling rate are needed for dynamic identification. Critical and deferrable demand must be separated, and outage intervals should remain visible rather than being interpolated away. Scenario reduction may be used for computation, but low-resource sequences and component outages must be retained even when their probability is small.

The state model should then be rebuilt from the implemented controller. Current- and voltage-loop gains, power filters, virtual impedance, droop slopes, delays, and current limits should come from controller settings, manufacturer data, or perturbation tests. Parameter uncertainty can be handled by evaluating the modal gate at nominal and adverse combinations instead of relying on a single best-fit matrix.

Shortlisted designs should be checked in stages: unseen annual operating data, a full-order averaged converter model, electromagnetic-transient simulation, and finally controller- or power-hardware-in-the-loop testing. The same capacities and gains must be carried through each stage. If a failed test requires retuning, the planning calculation should be repeated because the feasibility boundary has changed.

The final report should archive the data period, scenario weights, parameter ranges, surrogate version, validation errors, and exact-model checks. This record makes later re-planning possible when prices, load growth, degradation, or available grid-forming hardware change, and it shows whether a revision is driven by energy adequacy or by dynamic margin.

6. Conclusion

This paper presented a source–storage capacity planning framework for islanded microgrids in which small-signal stability is enforced during, rather than after, optimization. The formulation jointly selects PV capacity, BESS power and energy, droop coefficients, and virtual inertia. Annual dispatch supplies cost, curtailment, throughput, and reliability indices, while a capacity-dependent linearized model supplies spectral-abscissa and modal-damping constraints. The reproducible case study showed that reliability-feasible plans can still be weakly damped

and that explicit modal constraints remove these designs from consideration. The main engineering implication is that storage energy, grid-forming power, and controller gains should be coordinated: kilowatt-hours secure energy adequacy, whereas converter kilowatts and control parameters govern fast stability. The proposed framework is suitable as a screening layer for international planning studies, provided that final designs are verified using site data and a higher-order converter model.

Nomenclature

The acronyms and principal symbols used throughout this paper are summarized in Table 7.

Table 7. Acronyms and principal symbols.

Term	Definition	Term	Definition
BESS	Battery energy storage system	EENS	Expected energy not supplied
GFM	Grid-forming converter	LPSP	Loss-of-power-supply probability
PV	Photovoltaic generation	SOC	State of charge
P_{PV}^{r}	Rated PV power	P_{B}^{r}	Rated BESS converter power
P_{B}^{r}	Rated battery energy	m_p,m_q	Active- and reactive-power droop slopes
H_v	Virtual-inertia coefficient	α(A)	Spectral abscissa
ζ_i	Damping ratio of mode i	S	Set of operating scenarios

Appendix A. Reproducibility Notes

The synthetic year contains 8760 hourly samples and uses random seed 42. Candidate variables are uniformly sampled within Table 3. For every candidate, dispatch begins at 55% state of charge, applies charging efficiency 0.95 and discharging efficiency 0.94, and enforces a 10–90% state-of-charge window. The reported search contains 2200 candidates. Reproduction requires implementing Equations (1)–(14), the parameter values in Tables 2–3, and the stated reduced-order state vector. Numerical differences caused by linear-algebra libraries should be reported rather than hidden.

References

[1] Lasseter, R. H. (2002). MicroGrids. In 2002 IEEE Power Engineering Society Winter Meeting (Vol. 1, pp. 305–308). https://doi.org/10.1109/PESW.2002.985003

[2] Guerrero, J. M., Vasquez, J. C., Matas, J., de Vicuña, L. G., & Castilla, M. (2011). Hierarchical control of droop-controlled AC and DC microgrids—A general approach toward standardization. IEEE Transactions on Industrial Electronics, 58(1), 158–172.

https://doi.org/10.1109/TIE.2010.2066534
[3] Pogaku, N., Prodanović, M., & Green, T. C. (2007). Modeling, analysis and testing of autonomous operation of an inverter-based microgrid. IEEE Transactions on Power Electronics, 22(2), 613–625. https://doi.org/10.1109/TPEL.2006.890003
[4] He, J., & Li, Y. W. (2011). Analysis, design, and implementation of virtual impedance for power electronics interfaced distributed generation. IEEE Transactions on Industry Applications, 47(6), 2525–2538. https://doi.org/10.1109/TIA.2011.2168592
[5] Rocabert, J., Luna, A., Blaabjerg, F., & Rodríguez, P. (2012). Control of power converters in AC microgrids. IEEE Transactions on Power Electronics, 27(11), 4734–4749. https://doi.org/10.1109/TPEL.2012.2199334
[6] Schiffer, J., Zonetti, D., Ortega, R., Stanković, A. M., Sezi, T., & Raisch, J. (2016). A survey on modeling of microgrids—From fundamental physics to phasors and voltage sources. Automatica, 74, 135–150. https://doi.org/10.1016/j.automatica.2016.07.036
[7] Kroposki, B., Johnson, B., Zhang, Y., Gevorgian, V., Denholm, P., Hodge, B.-M., & Hannegan, B. (2017). Achieving a 100% renewable grid. IEEE Power and Energy Magazine, 15(2), 61–73. https://doi.org/10.1109/MPE.2016.2637122
[8] Tamrakar, U., Shrestha, D., Maharjan, M., Bhattarai, B. P., Hansen, T. M., & Tonkoski, R. (2017). Virtual inertia: Current trends and future directions. Applied Sciences, 7(7), 654. https://doi.org/10.3390/app7070654
[9] Milano, F., Dörfler, F., Hug, G., Hill, D. J., & Verbič, G. (2018). Foundations and challenges of low-inertia systems. In 2018 Power Systems Computation Conference (pp. 1–25). https://doi.org/10.23919/PSCC.2018.8450880
[10] IEEE Standards Association. (2018). IEEE standard for the specification of microgrid controllers (IEEE Std 2030.7-2017). https://doi.org/10.1109/IEEESTD.2018.8340204
[11] IEEE Standards Association. (2018). IEEE standard for interconnection and interoperability of distributed energy resources (IEEE Std 1547-2018). https://doi.org/10.1109/IEEESTD.2018.8332112
[12] Samimi, A., Nikzad, M., & Zakipour, A. (2025). Optimal sizing model of battery energy storage in a droop-controlled islanded multi-carrier microgrid. Scientific Reports, 15, 2540. https://doi.org/10.1038/s41598-025-86368-2
[13] Hassan, C. T., Jadoon, T. M., & Arif, A. (2025). Delay-dependent small-signal stability analysis of inverter-based islanded microgrids. IET Smart Grid, 8(1), e12206. https://doi.org/10.1049/stg2.12206
[14] Shi, K., Bai, Y., Xu, P., & Sun, Y. (2025). Dual-side virtual inertia control for stability enhancement in islanded AC–DC hybrid microgrids. Electric Power Systems Research, 241, 111396. https://doi.org/10.1016/j.epsr.2024.111396
[15] Sati, S. E., Al-Durra, A., Zeineldin, H. H., EL-Fouly, T. H. M., & El-Saadany, E. F. (2025). Decentralized frequency restoration and stability enhancement for virtual synchronous machines. Applied Energy, 377, 124544. https://doi.org/10.1016/j.apenergy.2024.124544
[16] Maurya, S. S., Jain, T., & Umarikar, A. C. (2025). Synchronverter with emulation based current control to improve small signal stability in islanded microgrid. IEEE Journal of Emerging and Selected Topics in Industrial Electronics, 6(4), 1688–1699. https://doi.org/10.1109/JESTIE.2025.3590837
[17] Khan, M. Z., Mu, C., Khan, M. M., Ke, W., & Khan, M. Q. (2026). Active power sharing and frequency regulation of islanded microgrids under communication islanding. Electric Power Systems Research, 254, 112624. https://doi.org/10.1016/j.epsr.2025.112624
[18] Saleheen, M. Z., Wagner, M., & Wang, H. (2026). Ensuring reliability in 100% renewable microgrids: A scenario-based joint planning and operational design framework. Energy Conversion and Management, 363, 121645. https://doi.org/10.1016/j.enconman.2026.121645

[19] Camargo-Berrueco, J. C., Mejía-Giraldo, D. A., & Lemos-Cano, S. (2026). Reliability-based sizing of islanded multi-energy microgrid. Energy Systems, 17(1), 239–276. https://doi.org/10.1007/s12667-024-00672-w

[20] Wang, G., Dong, K., Song, Y., Cao, S., Duan, J., Duan, G., Wu, D., & Li, S. (2025). Optimal battery storage capacity planning model and DOD regulation strategy for islanded microgrid. Energy, 332, 137200. https://doi.org/10.1016/j.energy.2025.137200

[21] Mühlbauer, A., Fan, Y. F., Sambor, D. J., & Jacobson, M. Z. (2025). Minimizing the multi-decadal cost of islanded renewable-electricity microgrids for different climate zones. Smart Energy, 20, 100208. https://doi.org/10.1016/j.segy.2025.100208

[22] Sagor, A. R., Hazari, M. R., Ahmad, S., Ogliari, E., Hossain, C. A., Jahan, E., & Mannan, M. A. (2025). Optimal sizing of islanded microgrid using pelican optimization algorithm. Electric Power Systems Research, 238, 111088. https://doi.org/10.1016/j.epsr.2024.111088

[23] Huang, C.-J., & Kuo, P.-H. (2019). Multiple-input deep convolutional neural network model for short-term photovoltaic power forecasting. IEEE Access, 7, 74822–74834. https://doi.org/10.1109/ACCESS.2019.2921238

[24] Kundur, P., et al. (2004). Definition and classification of power system stability. IEEE Transactions on Power Systems, 19(3), 1387–1401. https://doi.org/10.1109/TPWRS.2004.825981

[25] Kundur, P. (1994). Power system stability and control. McGraw-Hill.

[26] Olivares, D. E., et al. (2014). Trends in microgrid control. IEEE Transactions on Smart Grid, 5(4), 1905–1919. https://doi.org/10.1109/TSG.2013.2295514

[27] Bidram, A., & Davoudi, A. (2012). Hierarchical structure of microgrids control system. IEEE Transactions on Smart Grid, 3(4), 1963–1976. https://doi.org/10.1109/TSG.2012.2197425

[28] D'Arco, S., Suul, J. A., & Fosso, O. B. (2015). A virtual synchronous machine implementation for distributed control of power converters in SmartGrids. Electric Power Systems Research, 122, 180–197. https://doi.org/10.1016/j.epsr.2015.01.001

[29] Unamuno, E., & Barrena, J. A. (2015). Hybrid AC/DC microgrids—Part I: Review and classification of topologies. Renewable and Sustainable Energy Reviews, 52, 1251–1259. https://doi.org/10.1016/j.rser.2015.07.194

[30] Parisio, A., Rikos, E., & Glielmo, L. (2014). A model predictive control approach to microgrid operation optimization. IEEE Transactions on Control Systems Technology, 22(5), 1813–1827. https://doi.org/10.1109/TCST.2013.2295737

[31] Simpson-Porco, J. W., Dörfler, F., & Bullo, F. (2013). Synchronization and power sharing for droop-controlled inverters. Automatica, 49(9), 2603–2611. https://doi.org/10.1016/j.automatica.2013.05.018

[32] Farrokhabadi, M., et al. (2020). Microgrid stability definitions, analysis, and examples. IEEE Transactions on Power Systems, 35(1), 13–29. https://doi.org/10.1109/TPWRS.2019.2925703

[33] Hatziargyriou, N. (Ed.). (2014). Microgrids: Architectures and control. Wiley-IEEE Press.

[34] Yazdani, A., & Iravani, R. (2010). Voltage-sourced converters in power systems. Wiley-IEEE Press.

[35] Blaabjerg, F., Teodorescu, R., Liserre, M., & Timbus, A. V. (2006). Overview of control and grid synchronization for distributed power generation systems. IEEE Transactions on Industrial Electronics, 53(5), 1398–1409. https://doi.org/10.1109/TIE.2006.881997