Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Contingency limited modular nuclear deployment in island power systems with an illustrative Dominican assessment

Baldo Alberto Luigi Dalporto

baldo.dalporto@intec.edu.do

Instituto Tecnológico de Santo Domingo (INTEC)

Sabine Mary

Instituto Tecnológico de Santo Domingo (INTEC)

Santiago Gallur

Instituto Tecnológico de Santo Domingo (INTEC)

Article

Keywords: Small modular reactors, Island power systems, Common-cause contingencies, Frequency response, Capital-cost frontier, Dominican Republic

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11244884/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Contingency limited modular nuclear deployment in island power systems with an illustrative Dominican assessment

Baldo Alberto Luigi Dalporto[1]; Sabine Mary[1]; Santiago Gallur[1]

[1] Area of Engineering, Instituto Tecnológico de Santo Domingo, Santo Domingo, 10602, Dominican Republic

Corresponding author: Baldo Alberto Luigi Dalporto; baldo.dalporto@intec.edu.do

Coauthor emails: sabine.mary@intec.edu.do; santiago.gallur@intec.edu.do

ORCID: Dalporto 0009-0008-8719-1562; Mary 0009-0004-3488-9511; Gallur 0000-0001-6287-7340

Article type: Analytical research article | Running title: Nuclear Modularity and Grid Security

Abstract

Small reactor modules can reduce the largest credible electrical loss, but their reserve and economic benefits depend on actual electrical independence. This study derives a conditional screening framework linking simultaneous-trip groups, surviving inertia, delayed fast response and an affordable capital premium for separation. A baseline contingency prevents attribution of all system reserves to the new plant. Five hypothetical 400 MW configurations are evaluated over 81 operating states. A separate four-state dynamic reference, using a TGOV1-inspired governor and reheat structure with explicit nonlinear limits, then tests the screen over nine damping and governor-lag settings. With 120 MW of available total fast response, the 3,645 architecture-state-reference comparisons yield 85 analytical admissions rejected by the reference and 42 analytical rejections admitted by it. These are deterministic comparison counts, not reliability estimates. In the central analytical case, independent 200 MW and 100 MW modules differ by 54.28 MW of incremental reserve and support a separation premium of USD 61.08/kW. Dynamic reserve sizing on a 0.25 MW grid changes those values to 52.75 MW and USD 59.36/kW. Close central agreement conceals compensating effects of damping and delayed mechanical response. Reserve savings saturate below the existing contingency, and delayed response cannot repair an initial frequency-rate violation. Independent integration, step refinement and an exact linear-state solution check numerical implementation. Dominican observations provide scale context only. The contribution is a reproducible architecture-to-cost boundary together with a quantitative diagnosis of where its simplified dynamic assumptions change screening decisions; neither model establishes real-grid validity.

Keywords: Small modular reactors; Island power systems; Common-cause contingencies; Frequency response; Capital-cost frontier; Dominican Republic

1 Introduction

For an island electricity system, the relevant size of a nuclear installation is not necessarily the rating of one reactor module. Several modules can share a turbine, switchyard, export circuit, protection scheme or supporting infrastructure. An initiating event can therefore remove more electrical power than the module label suggests. Conversely, a nuclear block smaller than an existing system contingency need not increase the reserve requirement when operating conditions are unchanged. These two observations motivate an economic question: how much additional capital expenditure can independently trippable electrical blocks justify through frequency-reserve savings?

Grid reliability and the nuclear plant interface are established subjects in IAEA guidance [1]. Recent research goes substantially further than comparing reactor capacity with peak demand. Poudel, Upadhyaya Subedi and Gokaraju assess SMR placement and sizing through steady-state, dynamic and station-blackout considerations in Saskatchewan [2]. Liu et al. integrate frequency-nadir constraints into island-system unit commitment using operational and dynamic-simulation information [3]. Cruz-De-Jesús et al. assess battery provision of frequency-regulation reserves specifically for the Dominican Republic [4]. Thus, neither considering frequency security nor applying storage economics to a Dominican system is itself a new contribution.

Nuclear economics also extends beyond component efficiency or nominal plant size. Mignacca and Locatelli synthesize the roles of modularity, construction and financing [5]; Eash-Gates et al. identify mechanisms behind historical nuclear construction-cost escalation [6]. Badrianto et al. analyze nuclear–renewable configurations for Bangka Belitung using system optimization and probabilistic economic assessment [7]. Together, these studies establish the need to identify a specific mechanism linking the electrical architecture to its economic benefit.

This paper addresses a narrower relationship. It derives how an explicitly defined simultaneous-trip group changes the minimum delayed fast response, how much of that response is incremental to the existing system, and how the resulting annual expense limits the affordable capital premium for additional separation. The frequency and financial equations are established building blocks. The proposed contribution is their transparent composition into a conditional architecture-comparison rule, including a saturation boundary and an inadmissibility condition that economic scoring cannot override. The formulation is intended as a transparent preliminary screen that complements those more detailed methods. A separate nonlinear time-domain reference tests whether that screen changes admission decisions under finite reserve availability; it also exposes compensation between damping and delayed mechanical response.

The Dominican Republic provides a meaningful scale context. The Ministry of Energy and Mines reports 23,868.5 GWh injected into SENI in 2025 [8] and a demand observation of 4,077.97 MW on 17 June 2026 [9]. These values do not reveal post-trip inertia, reserve response, minimum-load dispatch or common-cause station contingencies. Accordingly, the numerical experiment is deliberately a parameterized island system, not a calibrated representation of SENI. Keeping this distinction explicit permits reproducible theoretical work without presenting a national procurement conclusion unsupported by the available data.

2 Research scope and relation to prior work

2.1 Evidence and study design

The study combines targeted literature positioning, analytical derivation, deterministic parameter exploration and a separately implemented dynamic reference comparison. Public institutional documents, primary research records and accessible abstracts were checked on 28 September 2026. Searches focused on SMR grid placement, frequency-constrained island operation, Dominican reserve provision, modular-reactor economics and nuclear infrastructure. This was not a systematic database review; no exhaustive search, PRISMA screening, bibliometric completeness or duplicate independent assessment is claimed. The comparisons in Table 1 concern documented objectives and contributions, not assertions about every equation in inaccessible full texts.

Table 1 Position of the present contribution relative to selected literature

Source	Established contribution	Focus added in this study
IAEA [1]	Nuclear plant and grid reliability interface	Explicit architecture and reserve-cost boundary
Poudel et al. [2]	SMR siting and sizing with grid and safety assessments	Closed-form electrical-grouping comparison
Liu et al. [3]	Frequency-constrained island unit commitment	Transparent reduced model before calibrated scheduling
Cruz-De-Jesús et al. [4]	Dominican battery reserve economics	Increment attributable to a specified nuclear contingency
Mignacca and Locatelli [5]	SMR economics and finance synthesis	Capital premium bounded by one quantified service
Badrianto et al. [7]	Island nuclear–renewable optimization and probabilistic costs	Grouping, baseline saturation and a separation premium

Only reactor-design interfaces relevant to the research question are considered. Water-cooled, gas-cooled, liquid-metal and salt-based concepts can differ greatly in reactor physics while having similar electrical contingency sizes. Conversely, two stations using the same reactor technology can differ in electrical independence. Reactor-family descriptions therefore cannot determine the grouping parameter used below. A turbine and switchyard diagram, protection study and credible-event assessment are required.

Lead-cooled fast reactors, including programmes such as newcleo, remain relevant examples because thermal design, fuel qualification and plant integration must be evaluated together. The Generation IV International Forum identifies materials, coolant chemistry, inspection and fuel-cycle development as important LFR issues [10]. The present 200 MW example is an abstract electrical block; its assumed cost, availability and separation are not predictions for newcleo or any other supplier. The analysis does not establish the licensability or operating performance of a particular advanced reactor.

2.2 Three distinct kinds of statement

The manuscript distinguishes published observations, analytical consequences and selected scenarios. The Dominican energy and demand figures are reported observations. The derived inequalities follow from the declared reduced model. Inertia, ramp rate, response delay, security limits, capital costs and capacity factors in the numerical experiment are author-selected scenario inputs. They have no claimed statistical distribution and are not calibrated to Dominican measurements.

The principal research question is whether the value of reducing an independently trippable nuclear block can be bounded before a detailed dispatch model is available. Two secondary questions follow: when does further subdivision cease to reduce the reserve burden, and which frequency violations cannot be repaired by adding a delayed battery? The answers are conditional mathematical results. Testing their quantitative usefulness for a real power system remains a separate empirical task.

3 Analytical framework

3.1 Station architecture and the comparison baseline

Let P be total net station output in MW, N the number of equal operating modules and g the largest number of these modules removed by an event within the specified credible-contingency set. Module output is p = P/N.

The nuclear simultaneous loss is gP/N, while L_0 is the largest pre-existing loss in the comparison system. Under otherwise identical response conditions, the system loss to be screened is

$$L = max(L_0, \frac{gP}{N}), \quad p = \frac{P}{N} \quad (1)$$

The parameter g is an event-based grouping, not the probability of independent failures. For heterogeneous blocks, gP/N is replaced by the sum of the outputs removed by each event, and the maximum is taken over the event set. A whole-site event gives g = N. Assuming g = 1 requires evidence of the relevant independence; simply labelling a station modular is insufficient. Electrical independence also does not establish independence of fuel supply, licensing delays or other causes of prolonged unavailability.

Equation 1 is a controlled comparison. It holds external dispatch and post-event response conditions fixed to isolate the effect of grouping. If nuclear entry displaces synchronous machines, changes the largest existing event or alters reserve headroom, L_0 and the response parameters must be recalculated. Under such conditions, a small nuclear block can still change reserve requirements indirectly. The zero-increment result derived later must not be applied outside the fixed-baseline comparison.

3.2 Post-contingency inertia and the immediate frequency-rate constraint

Let E be the kinetic energy remaining electrically connected after the trip, in MW s. With nominal frequency f_0 and a coherent single-area approximation, define M = 2E/f_0, in MW s/Hz. The magnitude of the initial rate of change of frequency is L/M. A selected ceiling ρ therefore requires

$$M = \frac{2E}{f_0}, \quad \frac{L}{M} \leq \rho, \quad L \leq \frac{2E\rho}{f_0} \quad (2)$$

Only surviving rotating energy belongs in E. Energy associated with a disconnected generator cannot support the remaining system. For a detailed comparison, both E and the lost output must be calculated for each contingency and operating state. The central experiment assigns the same external surviving energy to all configurations and credits no nuclear-site rotating energy. This isolates topology; it is not a guarantee of conservatism if future dispatch reduces external inertia.

This is the instantaneous limit of the reduced swing equation. Actual relays can use filtering and finite measurement windows, and spatially separated frequencies can differ. A site assessment must reproduce those implementations. A fast-response resource that first injects at a positive delay cannot alter the mathematical slope at time zero. Grid-forming controls or additional synchronously connected inertia require their own dynamic representation and are not represented by an instantaneous subtraction of battery MW from L.

3.3 Delayed response and the maximum frequency depression

Assume an aggregate surviving primary response that increases at rate a > 0 MW/s from the event, together with a fast response B MW activated at delay τ and then sustained. Define x as the positive frequency depression below f_0. Until the power deficit is eliminated, the model is

$$M \frac{dx}{dt} = L - at - B\, I_{t \geq \tau}, \quad x(0) = 0 \quad (3)$$

Here the indicator is zero before τ and one thereafter. Load damping, governor deadbands, network modes and converter dynamics are omitted. The ramp a is an assumed response envelope, not the rated ramp capability of one generator. A real application must demonstrate that sufficient surviving primary-response headroom can deliver this trajectory. Ignoring positive load damping makes the chosen response comparison more demanding only when that assumed response envelope is actually maintained.

Let t_d = min(τ, L/a) and [z]+ = max(z, 0). Integrating the positive imbalance before and after fast-response activation gives the maximum deficit area A, in MW s, and maximum depression x_{max}:

$$A(L,B) = Lt_d - \frac{at_d^2}{2} + \frac{[L - B - a\tau]_+^2}{2a}$$

$$x_{max} = \frac{A(L,B)}{M}, \quad t_d = min(\tau, L/a) \quad (4)$$

The construction clips the remaining imbalance at zero after balance is reached. It represents arrest at the frequency minimum, not restoration to nominal frequency. Secondary action and replacement energy must subsequently restore frequency and release fast reserve. Without sufficient sustained response and a controlled withdrawal of B, a second frequency decline can occur and the present screen is incomplete.

For an allowed depression F > 0, B_{min} is obtained in three branches. If $L^2/(2a) \leq MF$, primary response alone meets the depression limit and $B_{min} = 0$. If the deficit area accumulated before activation (the first two terms of Eq. 4) exceeds MF, no nonnegative delayed step can prevent the limit from already being crossed. Otherwise the minimum step is

$$B_{min} = L - a\tau - \sqrt{(a\tau)^2 - 2a\tau L + 2aMF} \quad (5)$$

In the positive-response branch, $\tau < L/a$ and $0 < B_{min} \leq L - a\tau$. The square root is real because the pre-activation deficit does not exceed MF. This branch is obtained by expanding $A = (L - B)^2/(2a) + B\tau$ and selecting the smaller quadratic root. The larger root is outside the relevant minimum-response interval. Appendix A gives the derivation and computational handling of boundary cases.

3.4 Incremental reserve and the saturation boundary

Apply the same response model to L_0. The baseline obligation is $B_0 = B_{min}(L_0)$, while the post-entry obligation is $B_{min}(L)$. If both baseline and candidate meet the initial-rate and delay conditions, the incremental power requirement is

$$\Delta B = max(0, B_{min}(L) - B_{min}(L_0)) \quad (6)$$

The result is a reserve-capacity difference, not an incremental dispatch-cost optimum. It assumes the baseline reserve resource remains available and that its response quality is equivalent to the added resource. Where reserve can be shared across non-simultaneous contingencies, requirements follow the maximum eligible loss; adding their full individual reserve requirements would double count capacity. Where events can occur jointly, they belong in the joint contingency set instead.

Proposition 1 follows immediately. For fixed P, g, E, a, τ, F and L_0, the incremental reserve requirement is non-increasing with N. Once $gP/N \leq L_0$, it is zero and further subdivision provides no additional benefit through this particular reserve-capacity mechanism. For fixed g, saturation begins at

$$N \geq \left\lceil \frac{gP}{L_0} \right\rceil \quad (7)$$

This integer condition is subject to $g \leq N$ and the assumed grouping remaining physically meaningful as the architecture changes. For a whole-site event, g = N, so gP/N = P independently of N: subdivision alone provides no contingency reduction. Similarly, if common equipment causes two 100 MW modules to trip together, their disturbance is the same as one 200 MW block within this model. These results concern the screened frequency service; maintenance scheduling, construction flexibility and other modularity benefits can remain.

3.5 From response capability to an annual reserve expense

For a conservative constant-output bridge lasting T hours, the incremental battery needs delivered power Δ B and nameplate energy Δ W. With discharge efficiency η and guaranteed usable fraction u of installed energy at the event, its energy and capital cost are

$$\Delta W = \frac{\Delta B\, T}{\eta u}, \quad K_b = 1000(k_P \Delta B + k_W \Delta W) \quad (8)$$

The coefficients k_P and k_W are separated power-related USD/kW and energy-related USD/kWh costs. They are illustrative additive components; applying this expression to an all-inclusive USD/kWh quotation would double count equipment. The usable fraction must include the operational state-of-charge requirement and end-of-life available capacity, or the capacity must be augmented separately. Reserves cannot be simultaneously promised to another service without proving co-delivery.

Annual battery expense A_b is capital multiplied by CRF(r, n_b) + β, where n_b is battery life and β is an annual fixed-cost fraction. CRF is the conventional capital-recovery factor:

$$CRF(r,n) = \frac{r}{1 - (1+r)^{-n}}, \quad CRF(0,n) = 1/n$$

$$A_b = [CRF(r,n_b) + \beta]K_b \quad (9)$$

The central comparison assumes replacement every 15 years over a 60-year nuclear life at unchanged real cost, giving four battery cycles and no terminal salvage. Equal annualization at the same real rate is consistent with that repeated replacement assumption. Cycling wear, charging energy, changes in replacement prices, common-mode unavailability and network reinforcement are excluded. Their omission restricts the reserve-capital comparison; it does not establish an all-inclusive reserve service cost.

3.6 A conditional capital ceiling and the value of separation

Let K be the nuclear capital basis at commercial operation in USD/kW, n the operating life, f the annual fixed operating cost in USD/kW-year, v the variable allowance in USD/MWh and CF the annual capacity factor. Annual energy is Q = 8760 P CF MWh. A restricted plant-plus-incremental-reserve cost is

$$C = \frac{K\,CRF(r,n) + f}{8.76\,CF} + v + \frac{A_b}{Q} \quad (10)$$

This metric is not total system cost. It excludes the baseline reserve expense common to the controlled comparison, transmission, public programme costs, long-outage replacement and any liabilities absent from f and v. For a scenario ceiling C^* on this same cost boundary, rearrangement yields

$$K_{max} = \frac{8.76\,CF(C^* - v - A_b/Q) - f}{CRF(r,n)} \quad (11)$$

A negative K_{max} means no positive capital basis can satisfy that scenario ceiling. The ceiling must be supplied by the decision problem; it cannot be inferred from a retail tariff or a variable renewable generator's standalone LCOE. Financial feasibility is evaluated only after the dynamic screen passes. A configuration rejected by the initial-rate condition does not acquire an acceptable ranking because its energy cost is low.

Proposition 2 links separation to its affordable premium. Compare a less separated configuration b with a more separated configuration c, with identical P, CF, n, r, f and v and both dynamically admissible. The maximum increase in nuclear capital basis compensated by the reserve savings is

$$\Delta K \leq \Delta K_{sep,max} = \frac{A_{b,b} - A_{b,c}}{1000\,P\,CRF(r,n)} \quad (12)$$

The numerator must be nonnegative for a positive premium. This is a breakeven condition for one mechanism, not an estimate of the total value of modular construction. The derivation subtracts Eq. 10 between the two cases; the equal energy denominator cancels. If operating costs or output differ, they must be carried through that subtraction rather than hidden in the premium. For uncertain operating states, contingency-specific requirements must be evaluated over a declared set before applying the cost comparison.

4 Numerical experiment and verification

4.1 Inputs and architecture comparisons

Table 2 specifies the central assumptions. All dynamic and financial parameters are hypothetical, including the 150 MW pre-existing contingency. That value is not reported as the current largest Dominican unit, an observed SENI trip or a regulatory reserve requirement. Likewise, 0.5 Hz/s and 0.8 Hz are selected study limits, not asserted Dominican grid-code values. Constant dollars are used consistently; the cost inputs are not supplier offers or a market forecast.

Table 2 Central scenario inputs and units

Parameter	Value	Interpretation
P; L0	400 MW; 150 MW	Total station output; hypothetical baseline loss
f0; E	60 Hz; 16,000 MW s	Nominal frequency; external surviving kinetic energy
a; tau	20 MW/s; 0.5 s	Primary-response ramp; fast-response delay
rho; F	0.5 Hz/s; 0.8 Hz	Assumed initial-rate and depression limits
CF; n	0.90; 60 years	Capacity factor; nuclear economic life
K; r	8,000 USD/kW; 8%	Capital basis at operation; real annual discount rate
f; v	120 USD/kW-year; 15 USD/MWh	Separate fixed and variable operating allowances
T; eta; u	0.25 h; 0.95; 0.80	Reserve duration; discharge efficiency; usable fraction
kP; kW	200 USD/kW; 200 USD/kWh	Additive power and energy cost components
nb; beta	15 years; 2%/year	Battery replacement life; annual fixed-cost fraction
Cstar	125 USD/MWh	Hypothetical ceiling for the restricted cost boundary

Five station arrangements maintain P = 400 MW to isolate the consequences of grouping. A has one 400 MW block; B has two independent 200 MW blocks; C has four independent 100 MW modules; D has four 100 MW modules with a credible pair loss; E has four 100 MW modules with a credible whole-site loss. Independence here refers to the stated event set. No configuration is attributed to a commercial design. Equal CF and capital basis initially hold all non-reserve attributes constant, after which Eq. 12 tests a possible capital premium.

The stress set crosses E = 10,000, 16,000 and 24,000 MW s; a = 10, 20 and 40 MW/s; τ = 0.1, 0.5 and 1.0 s; and F = 0.5, 0.8 and 1.0 Hz. It contains 81 parameter states and 405 architecture-state evaluations. This is a deterministic Cartesian sensitivity exercise. It assigns neither probabilities nor frequencies of occurrence to those states, and it does not model the actual correlation between inertia, dispatch and reserve headroom.

Finance is separately tested at real rates of 4%, 8% and 12% and CF values of 0.75, 0.90 and 0.95. The reserve duration is swept from 0.05 to 4 hours to expose its effect on the affordable separation premium. Extending T changes the conservative reserve-energy budget; it does not establish that the assumed primary and secondary trajectories would remain deliverable in a real prolonged outage.

4.2 Computational verification

Supplementary file S1 contains the standard-library Python implementation, complete inputs, scenario outputs and figure-generation script. The finite B_{min} solutions are checked through an independent midpoint integration of the time-domain positive power deficit, with a maximum 1 ms step and an explicit split at the fast-response

activation. The integration does not substitute Eq. 4 into the numerical calculation. Additional checks cover a delayed-response impossibility case, a zero discount rate, the primary-only branch and equal results for configurations with the same simultaneous loss.

Of the 405 evaluations, 369 admit a finite nadir-response solution. Across those cases, the maximum absolute discrepancy between the analytical and numerically integrated frequency depressions is 1.41×10^{-8} Hz. A small reduction of each positive minimum response crosses the selected depression bound. These checks verify arithmetic within the reduced model. The separate state-space comparison below examines additional dynamics; neither exercise validates measured events, hardware or operating reactors. The 36 other cases fail the pre-activation depression condition and do not receive an artificial finite response value.

4.3 Dynamic reference model and comparison protocol

The reference uses a single-area frequency state coupled to a governor lag, reheat lag and finite battery activation. Its linear governor and turbine structure is grounded in the documented TGOV1 model [11]. Deadband, command clipping and an additional valve slew limiter are explicitly defined extensions; the implementation is a TGOV1-inspired aggregate proxy, not a certified reproduction of a commercial model. NERC guidance motivates explicit treatment of deadbands and available headroom and emphasizes the need to represent actual plant response [12]. These sources support the model structure and modelling concerns; they do not supply the numerical scenario parameters in Table 3.

Let w and z be incremental governor and reheat states in MW, b delivered battery power in MW, D_l load damping in MW/Hz and P_m the surviving fleet's mechanical-response increment. The reference frequency equation is

$$M\,\frac{dx}{dt} = L - D_l x - P_m - b \quad (13)$$

With aggregate gain K_g, deadband d_b, headroom H and governor lag T_g, the incremental governor command c and state obey

$$c = clip(K_g[x - d_b]_+, 0, H)$$
$$\frac{dw}{dt} = clip(\frac{c - w}{T_g}, -a, a) \quad (14)$$

Here clip(y,l,u) = min(u,max(l,y)). The command is clipped before the lag; this chosen limiter arrangement is stated explicitly and is not claimed identical to every TGOV1 implementation. Reheat and battery dynamics are

$$\frac{dz}{dt} = \frac{w - z}{T_r}, \quad P_m = F_H w + (1 - F_H) z$$
$$\frac{db}{dt} = \frac{B_{cmd} I_{t \geq \tau} - b}{T_b} \quad (15)$$

All four incremental states initially equal zero. In the unsaturated, zero-deadband limit without the slew restriction, the mechanical transfer function is $K_g(1 + F_H\ T_r\ \text{s})/[(1 + T_g\ \text{s})(1 + T_r\ \text{s})]$, corresponding to TGOV1 time constants T1 = T_g, T2 = $F_H\ T_r$ and T3 = T_r, with its separate turbine damping set to zero. D_l represents load relief only. The primary response belongs to a hypothetical surviving external fleet, not to a disconnected nuclear unit or an assumed Dominican technology mix. No commercial simulation package or measured grid event was used.

Table 3 Hypothetical dynamic reference parameters

Parameter	Value	Role
Kg; equivalent droop	400 MW/Hz; 5% on 1,200 MW	Aggregate surviving response gain

Parameter	Value	Role
H; db	300 MW; 0.036 Hz	Incremental primary headroom and deadband
Tg	0.5 s central; 0.2, 0.5, 1.0 s grid	Governor lag
Tr; FH	5 s; 0.30	Reheat lag and direct high-pressure fraction
Dload	60 MW/Hz central; 0, 60, 120 grid	Linear frequency-dependent load relief
Tb; tau	0.1 s; same tau as Table 2 or stress state	Battery rise time after commanded delay
a	Same 10, 20 or 40 MW/s as stress state	Governor-state slew ceiling
Bcmd	120 MW classification; 0–400 MW sizing	Total fast response including baseline
Horizon; step	60 s; 5 ms	Checked at 120 s and 2.5 ms

All values are selected for this experiment. The 1,200 MW base only explains the chosen gain: 1,200/(0.05 × 60) = 400 MW/Hz. It is not an observed Dominican responsive capacity.

The 81 states in Section 4.1 are crossed with D_l = 0, 60 and 120 MW/Hz and T_g = 0.2, 0.5 and 1.0 s, producing nine reference settings and 3,645 architecture-state-reference comparisons. Configurations A and E, and B and D, intentionally repeat identical loss conditions; the counts are not independent samples. The same a that describes an ideal mechanical ramp in the analytical model limits the governor-state slew in the reference. Reheat then delays part of that response, while damping reduces the imbalance. This is an explicit structural sensitivity test, not a fit of a to the reference or a claim that these fleets are dynamically identical.

For the classification experiment, B_{cmd} is fixed at a hypothetical 120 MW total available fast response, including baseline provision. Both models must satisfy the same initial-rate ceiling and maximum frequency-depression limit. An analytical admission with a reference rejection is a false admission relative to this reference; the reverse is a false rejection relative to it. The reference is a comparison standard for the experiment, not ground truth. A 0.0001 Hz numerical comparison tolerance is used, with no implied operational safety margin. Scenarios far outside the allowed frequency band are diagnostic failures; their continued trajectories are not predictions of system survival without protection.

The four-state equations are integrated with classical fourth-order Runge–Kutta at 5 ms, splitting exactly at battery activation. Repeating the full grid at 2.5 ms checks nadir convergence and classification stability; extending the horizon from 60 to 120 s checks for later minima. An independent adaptive DOP853 implementation evaluates selected near-boundary and corner cases. A nonsaturating, zero-deadband case is also checked against a matrix-exponential solution of the linear state equations. The supplied code records solver tolerances, software versions and complete comparison outputs.

The central sizing experiment scans B_{cmd} from 0 to 400 MW in 0.25 MW increments for losses of 150, 200 and 400 MW. It reports the smallest passing tested value, separately from the initial-rate requirement, and checks the adjacent passing and failing values with the adaptive solver. It does not assume a closed-form minimum or count the 400 MW case as admissible after a rate failure. Battery power remains sustained throughout the experiment; secondary control, withdrawal, energy exhaustion and network dynamics are outside this comparison. The 60 s horizon is shorter than the assumed 900 s reserve-energy budget.

5 Results

5.1 Frequency constraints distinguish module rating from station loss

For E = 16,000 MW s and f_0 = 60 Hz, M = 533.33 MW s/Hz. The selected initial-rate ceiling admits no more than 266.67 MW of instantaneous loss. A 400 MW event produces 0.750 Hz/s and fails this screen. A 200 MW event produces 0.375 Hz/s, while the 150 MW baseline produces 0.281 Hz/s. Figure 1 separates this initial-rate boundary from the delayed response needed to satisfy the depression limit.

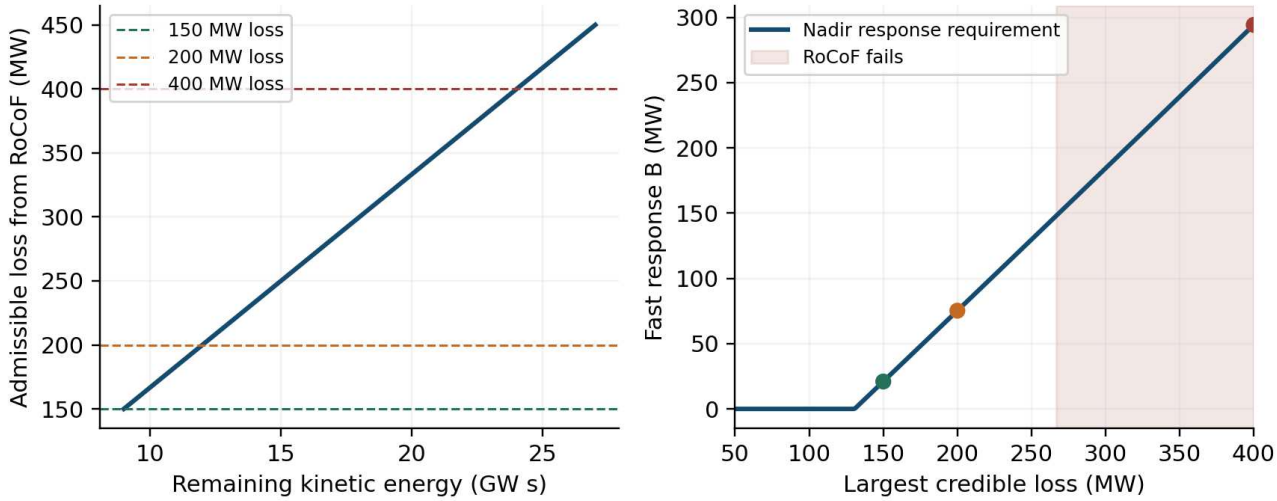


Figure 1 Initial-rate and nadir-response boundaries. Left, admissible loss increases with surviving kinetic energy at the assumed 0.5 Hz/s ceiling. Right, the nadir-only fast-response requirement at E = 16,000 MW s, a = 20 MW/s, τ = 0.5 s and F = 0.8 Hz. The shaded range fails the initial-rate screen despite finite response requirements. Original calculations.

The central baseline requires 20.98 MW of fast response. Configuration B requires 75.25 MW in total and therefore adds 54.28 MW. Configuration C leaves the largest system loss at L_0 = 150 MW, so it requires the baseline provision but adds none. Configuration D reproduces B because two simultaneously lost 100 MW modules remove the same 200 MW. Configurations A and E both fail the initial-rate screen; their finite nadir-only requirement of 294.26 MW must not be interpreted as sufficient for admissibility. Table 4 keeps that distinction visible.

Table 4 Central results for five equal-output station configurations

Case	Modules and group	L MW	Rate Hz/s	B MW	Added B MW	Screen
A	1 × 400; g = 1	400	0.750	294.26	—	Fail rate
B	2 × 200; g = 1	200	0.375	75.25	54.28	Pass
C	4 × 100; g = 1	150	0.281	20.98	0.00	Pass
D	4 × 100; g = 2	200	0.375	75.25	54.28	Pass
E	4 × 100; g = 4	400	0.750	294.26	—	Fail rate

B is the total nadir-only requirement. A finite value does not repair an initial-rate failure. Added capacity and cost are not assigned to rejected configurations. No installed-response power ceiling is imposed.

For fixed g = 1 and P = 400 MW, the reserve-saturation condition is N ≥ ceil(400/150) = 3. Thus, the analytical breakpoint is already reached by three hypothetical equal 133.33 MW blocks. Four blocks are used in the main comparison to provide an exact pair-loss contrast, not because four is an optimized module count. For g = 2 held

fixed, saturation would require at least six equal blocks. A whole-site event retains a 400 MW loss for any subdivision. Figure 2 illustrates these boundaries and the reduced frequency trajectories.

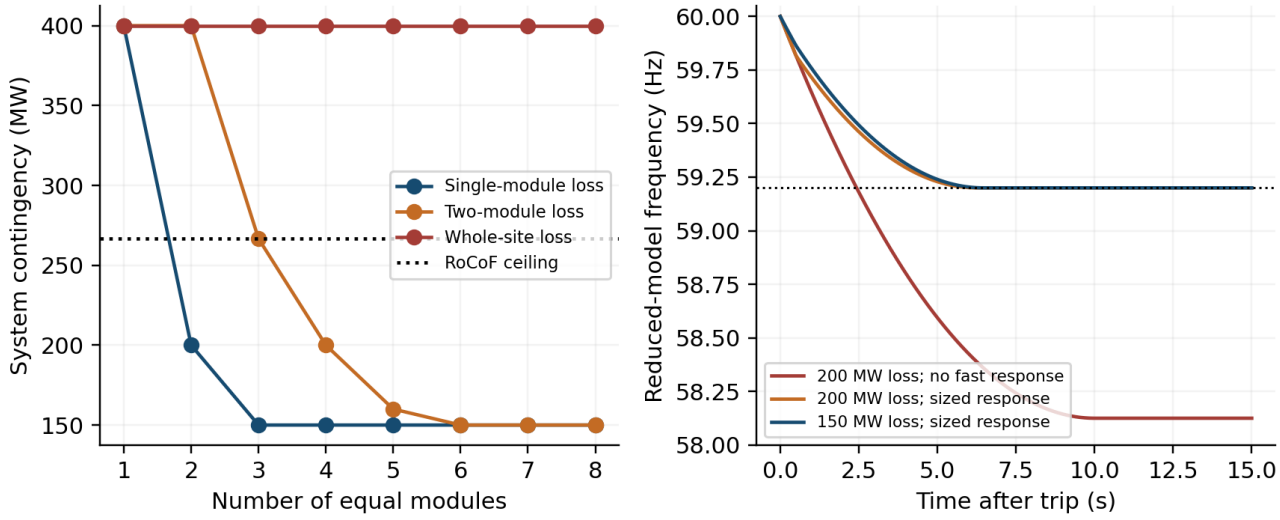


Figure 2 Electrical grouping and reduced frequency response. Left, increasing module count lowers the contingency only under the specified independent or pair-loss grouping; the pre-existing 150 MW loss creates a floor. Right, response-sized trajectories arrest at the selected 59.2 Hz minimum. The plots do not include secondary frequency restoration. Original calculations.

The frequency traces plateau at their calculated minima because the model ends the deficit when response balances the loss. They do not demonstrate recovery to 60 Hz. For the 200 MW case, the minimum occurs approximately 6.24 s after the event with the sized step, compared with 10 s under the unconstrained primary-only ramp. These times follow from the assumed response trajectory and are not measured governor performance.

5.2 The reserve benefit supports a bounded separation premium

At the assumed quarter-hour bridge, configuration B's incremental 54.28 MW requires 17.85 MWh of additional nameplate battery energy, accounting for $\eta = 0.95$ and $u = 0.80$. Its illustrative incremental battery capital is USD 14.43 million and annualized expense USD 1.97 million. These values concern the added reserve resource only; they do not include the baseline reserve plant or claim that a quarter-hour battery can replace a prolonged nuclear outage.

At CF = 0.90, each 400 MW configuration generates 3.1536 TWh annually. Dividing by the published 2025 SENI injection gives 13.21%, while total output is 9.81% of the dated June 2026 demand observation. These ratios normalize scale using two different reference periods. They have no role in the frequency calculation and do not forecast commissioning-year demand.

With K = USD 8,000/kW and r = 8%, the restricted cost is USD 112.83/MWh for B and D and USD 112.21/MWh for C. Reserve savings therefore support a capital premium of only USD 61.08/kW for C relative to B, about 0.76% of the assumed nuclear capital basis. An extra USD 800/kW, corresponding to 10% of that basis, would add about USD 8.20/MWh in annualized nuclear capital cost; the calculated reserve saving is only USD 0.63/MWh. Under these inputs, this single frequency-reserve benefit cannot justify that larger premium.

This comparison is meaningful only for architectures admitted under the specified response model. B can be compared with C in the central reduced-model state; Section 5.4 re-sizes their reserves under the separate reference. A cannot be declared cheaper or more expensive on an admissible-service basis until a specified intervention corrects its initial-rate failure and the intervention cost is included. Omitting this step would compare services with different security outcomes.

5.3 Financial and operating-state frontiers

For the illustrative ceiling of USD 125/MWh, configuration B admits K up to USD 9,187/kW at 8% and CF = 0.90. The same bound is USD 16,815/kW at 4% and USD 6,170/kW at 12%. At 8%, reducing CF to 0.75 lowers the bound to USD 7,398/kW. Table 5 and Figure 3 show that, in these selected scenarios, financing and sustained output move the capital ceiling much more than the central reserve saving.

Table 5 Maximum capital basis for configuration B in USD per kW

Real discount rate	CF 0.75	CF 0.90	CF 0.95
4%	13,545	16,815	17,905
8%	7,398	9,187	9,783
12%	4,967	6,170	6,571

C^* = 125 USD/MWh. All other assumptions follow Table 2, with battery annualization updated at each rate. These are scenario thresholds, not estimated construction costs.

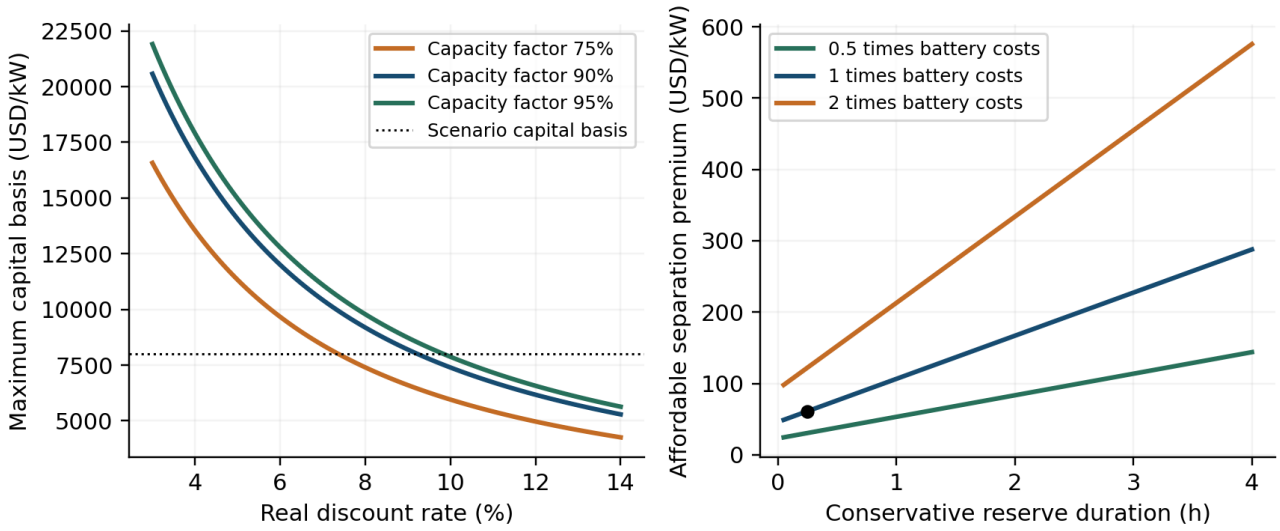


Figure 3 Economic boundaries under hypothetical costs. Left, the maximum nuclear capital basis for configuration B at a 125 USD/MWh restricted-cost ceiling. Right, the separation premium supported by reserve savings between B and C as reserve duration and both battery-cost coefficients vary. The dot marks the central 61.08 USD/kW value. Original calculations.

Configuration C passes the reduced dynamic screen in all 81 states; B and D pass in 54, and A and E in 27. These counts are bookkeeping for the selected grid of assumptions, not estimated reliability or probabilities of success. No finite ceiling on available fast-response power was imposed; a state passes only in the sense that the required response is finite and the initial-rate bound holds. Adding a power limit, network restriction or reserve-availability requirement could remove further states.

If every state in the declared set is mandatory, only C among the five examples remains eligible for subsequent screening. That statement is restricted to this set and this reduced model without a reserve-power ceiling; it is not a conclusion from the finite-power reference experiment below. An expanded set with less external inertia, slower control, lower battery availability or larger common-cause events can change the result. In particular, introducing a whole-site event converts C's contingency into E's and removes its apparent robustness. Event-set justification is therefore more consequential than attaching a generic uncertainty percentage to a nominal module rating.

5.4 Dynamic comparisons reveal the limits of the screening rule

At the central damping and governor settings, the reference gives maximum depressions of 0.8096 Hz for the 150 MW event and 0.7972 Hz for the 200 MW event when supplied with their respective analytical reserve sizes. Thus, the baseline requirement is slightly underestimated even though the larger-event requirement is slightly conservative. Figure 4 shows the time trajectories and response components. The 400 MW event reaches a 0.7675 Hz depression at its analytical reserve size but still fails the initial-rate constraint; nadir agreement cannot remove that failure.

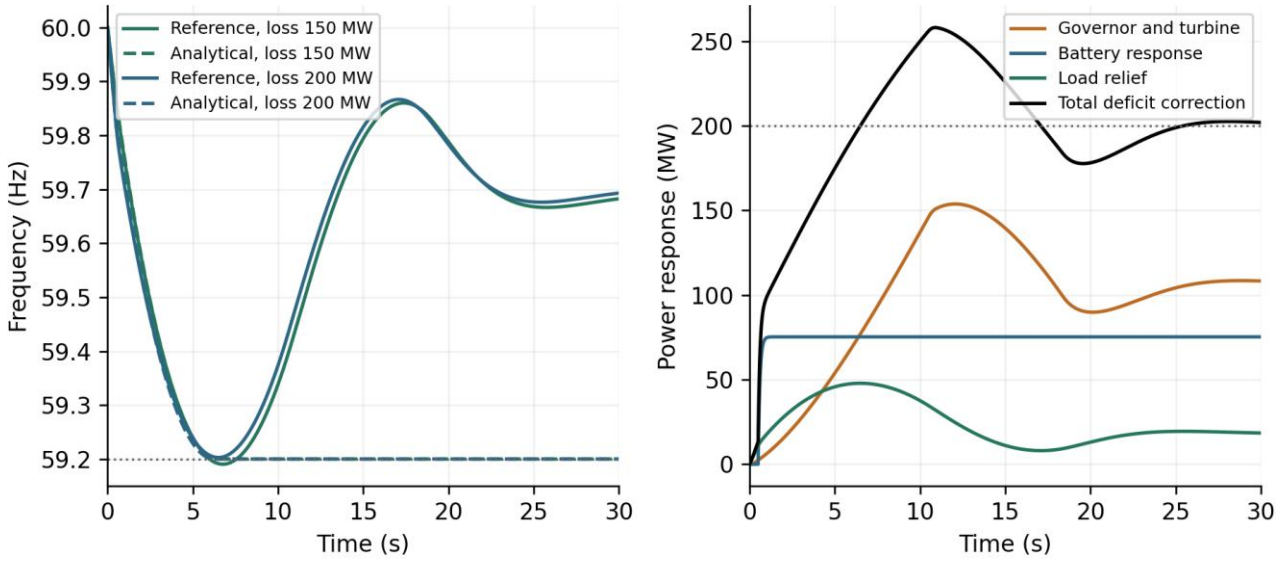


Figure 4 Central time-domain comparison with each event provisioned at its analytical reserve size. Left, the reduced ramp and dynamic reference give similar first minima, but the reference also shows recovery and oscillation. Right, the 200 MW event separates mechanical response, battery injection and frequency-dependent load relief. The dotted lines mark 59.2 Hz and 200 MW. Damping is 60 MW/Hz; T_g is 0.5 s. Full 60 s traces are supplied in S1.

The finite-power classification experiment produces 1,508 joint admissions, 2,010 joint rejections, 85 false admissions and 42 false rejections relative to the reference (Table 6). The reduced model admits 177 of 405 architecture-state combinations at each reference setting, while reference admissions range from 148 to 189. All 85 false admissions occur with zero load damping; positive damping creates the 42 false rejections in the selected set. Figure 5 displays this dependence. These counts describe the chosen grid and its repeated architectural equivalents; they do not estimate actual security, failure likelihood or out-of-sample classification accuracy.

Table 6 Classification comparisons at 120 MW of total fast response

D MW/Hz	Tg s	Both pass	Analytical pass reference fail	Analytical fail reference pass	Both fail
0	0.2	149	28	0	228
0	0.5	149	28	0	228
0	1.0	148	29	0	228
60	0.2	177	0	2	226
60	0.5	177	0	2	226
60	1.0	177	0	2	226

D MW/Hz	Tg s	Both pass	Analytical pass reference fail	Analytical fail reference pass	Both fail
120	0.2	177	0	12	216
120	0.5	177	0	12	216
120	1.0	177	0	12	216
All	—	1508	85	42	2010

Each reference setting contains 405 architecture-state comparisons. Passing requires both the initial-rate condition and the nadir condition. Counts include equivalent architectures and are not probabilities.

Among the 1,593 analytical admissions, the mean absolute difference in maximum depression is 0.106 Hz and the largest absolute difference is 0.389 Hz. Signed errors range from −0.362 to +0.389 Hz, taking reference minus analytical depression. These conditional error metrics describe the screening region, rather than averaging it with trajectories far outside the selected operating limits.

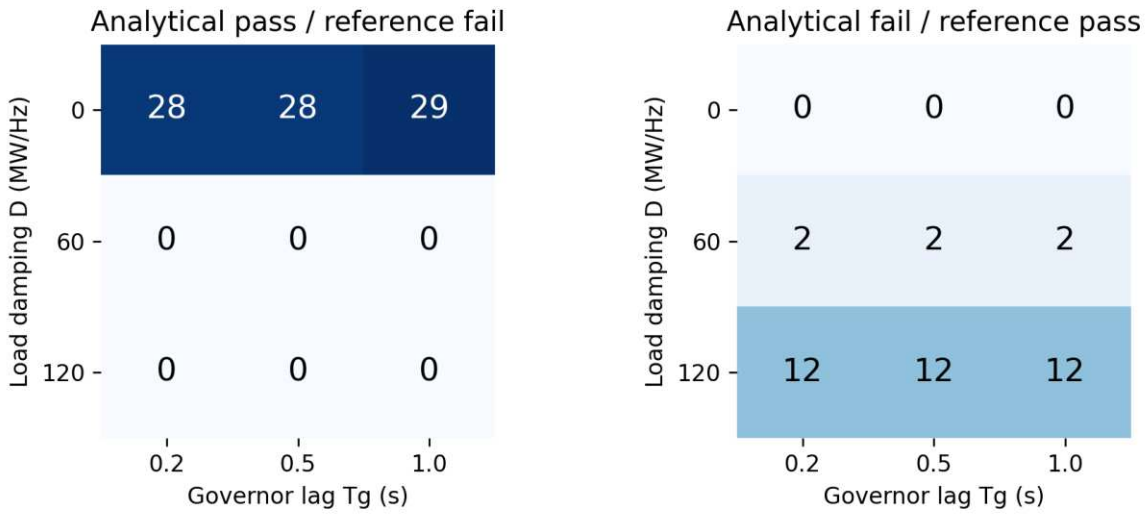


Figure 5 Disagreements between the analytical screen and the dynamic reference at 120 MW total fast response. Each cell contains 405 architecture-state comparisons. Damping can offset slow mechanical response, changing which model is more demanding. Counts describe the selected deterministic grid, including repeated loss-equivalent architectures.

Close agreement at one setting is therefore insufficient evidence of general accuracy. For the central 200 MW event and unchanged 75.2539 MW analytical reserve, removing load damping increases the reference depression from 0.7972 to 1.3356 Hz. Removing the valve slew limit instead lowers it to 0.4831 Hz; bypassing reheat lowers it to 0.5735 Hz. Removing the deadband gives 0.7867 Hz and shortening the battery time constant from 0.1 to 0.01 s gives 0.7906 Hz. Each is a one-factor ablation of the same central proxy, not a claim about a particular machine. Damping offsets delayed mechanical response in the central comparison; that compensation is not preserved across the stress set.

Under the reference, the smallest passing tested total responses are 22.25 MW for the baseline and 75.00 MW for the 200 MW event (Table 7). Their difference is 52.75 MW, and substituting it into the unchanged reserve-cost model yields a separation premium of USD 59.36/kW, compared with USD 61.08/kW from the analytical ramp. The approximately 2.8% difference concerns these hypothetical central inputs only. The two reported adjacent threshold brackets translate into a premium interval of approximately USD 59.08–59.64/kW; parameter and model uncertainty are additional. No assertion about an untested global continuous optimum is made. The

289.75 MW nadir-only value for the 400 MW event exceeds the classification experiment's 120 MW reserve ceiling and cannot cure its rate violation.

Table 7 Central reserve sizing in the two response models

Loss MW	Analytical B MW	Reference B MW	Reference depression at analytical B Hz	Initial rate
150	20.98	22.25	0.8096	Pass
200	75.25	75.00	0.7972	Pass
400	294.26	289.75	0.7675	Fail

Reference B is the smallest passing tested value on a 0.25 MW grid for a 0.8 Hz depression limit, with D = 60 MW/Hz and T_g = 0.5 s. The 400 MW row is a nadir-only diagnostic and fails the rate condition.

The qualitative distinction between module rating and simultaneous loss persists because both models receive the same contingency definition: B and D remain identical, as do A and E, and C shares the baseline 150 MW requirement. That invariance is imposed by the controlled architecture comparison and is not independent evidence that a real plant has electrically independent modules. The economic boundary in Eq. 12 can accept reserve costs from the more detailed model, but its monetary value must be recomputed whenever response assumptions change.

Across the full comparison grid, halving the integration step changes the maximum depression by at most 6.48e-05 Hz and changes no classifications. Extending the horizon adds no deeper sampled minimum. Independent adaptive integration checks 18 selected grid cases, with maximum nadir discrepancy 6.82e-05 Hz; it also confirms the six central threshold-bracket endpoints. The exact linear-state check has maximum componentwise discrepancy 1.08e-10 in the corresponding native state units. These are numerical checks of the disclosed proxy and must not be interpreted as validation of a national grid model.

6 Discussion and application to a Dominican study

6.1 What the combined formulation contributes

The principal result is a chain of auditable decisions: identify simultaneous electrical loss, calculate surviving inertia and response requirements, subtract an equivalent baseline obligation, and translate the remaining expense into an admissible capital premium. Existing frequency-constrained scheduling and SMR economic studies establish the importance of the constituent issues [2–5,7]. This formulation contributes a deliberately reduced, closed-form bridge between them that makes the dependence on electrical grouping and baseline saturation explicit.

The contribution should be judged by that narrow claim. It is not a new swing equation, a replacement for unit commitment, a reactor invention or evidence that nuclear power is the optimal Dominican expansion option. Its value is that a purchaser or system planner can reject internally inconsistent architecture claims and identify which data would change the screening result. A more detailed model may find benefits from sequential commissioning, maintenance diversity, flexibility, avoided fuel exposure or improved adequacy that the present premium omits.

The numerical example also tests a tempting assumption: that a smaller disturbance must produce a large economic advantage. Here it improves the dynamic screen but supports only a modest capital premium. The result can reverse in magnitude with more expensive or longer-duration reserves, different baseline contingencies, or a larger separation-induced reduction in loss. Reporting the parameter dependence is more informative than assigning a universal economic advantage to all SMRs.

6.2 Evidence required before applying the boundary to SENI

The next Dominican study should begin with contingency and operating records from the relevant system institutions. Required inputs include generator-level dispatch and kinetic-energy constants, credible simultaneous-trip groups, recorded frequency events, protection measurement windows and settings, primary-response headroom and trajectories, reserve availability, and battery state-of-charge policies. Low-load high-renewable states deserve explicit coverage. The response set must include failures of the resources being relied upon, not simply their normal operation.

With those inputs, the event-based form replaces the single L and E by contingency-specific quantities. For each time or operating state, calculate the loss, surviving inertia and response after that event, then retain the most demanding feasible requirement. If the baseline and nuclear portfolios have different commitment schedules, both schedules must be solved consistently. Baseline reserve expense is not automatically sunk, and the correct economic comparison is the difference between complete feasible portfolio costs.

Dynamic simulation should then test spatial frequency, governor deadbands, excitation and voltage behavior, inverter control, protection cascades, restoration and reserve withdrawal. The reduced closed-form boundary can be evaluated as a screening classifier against those simulations: report false admissions, unnecessary rejections and error in required response, rather than only correlation. The present reference comparison implements this diagnostic approach for a synthetic single-area proxy. Historical events are still needed for parameter identification, and withheld events or operating states are needed to assess predictive performance in a real network.

6.3 Reactor readiness and national responsibilities remain separate

Passing a frequency and restricted-cost screen cannot compensate for unresolved reactor safety, fuel qualification, waste arrangements or national infrastructure. The IAEA Milestones framework organizes a newcomer programme across phased preparation and infrastructure responsibilities [13]. The IEA also emphasizes the role of financing, delivery capability and supply chains [14]. Those sources support keeping evidence requirements separate from the numerical ranking; they do not establish that the Dominican Republic has met any particular milestone.

For a proposed LFR programme, favorable coolant properties do not establish an operating reference, qualified fuel supply, independent electrical blocks or an enforceable waste pathway [10]. Likewise, conclusions about waste from particular SMR designs cannot be generalized to all concepts [15]. A procurement evidence register should record each obligation as documented, unresolved or failed, with a responsible source and date. Unknown information must remain unknown rather than receiving an assumed passing score.

The corresponding decision is conditional. A plant that fails the numerical screen requires an explicit redesign or system intervention. A plant that passes but lacks essential project evidence remains an option for study. Only candidates with both adequate project evidence and a feasible system representation can enter a full portfolio comparison. No vendor is selected and no licensing conclusion is drawn by this work.

6.4 Limitations and falsifiable implications

The closed-form frequency model neglects location, damping, nonlinear controls and changing topology. The separate reference includes damping, limited governor action and reheat, but remains a single-area aggregate with hypothetical parameters and sustained battery injection. It omits voltage and reactive-power dynamics, protection, spatial modes, inverter controls and secondary restoration. It is a stronger structural comparison than re-integrating the ramp equation, yet is not an independently calibrated benchmark grid. In the closed-form model, the ramp begins without a primary-control dead time, while fast response is an ideal step after a fixed delay. Real response can be slower, limited or unavailable. The framework must therefore be recalculated or replaced when measured delivery falls outside these assumptions. Removing reactor-site inertia from the central comparison does not protect against unmodelled changes in external dispatch.

The economic model omits taxes, escalation, detailed construction cash flows, public infrastructure, reserve opportunity costs, multi-day replacement energy and lifecycle liabilities not included in the hypothetical operating allowances. K is defined at commercial operation, so applying the bound to an overnight quotation would require a construction financing model. The ceiling C^* is an input to a consistent restricted comparison and is not a demonstrated value of electricity. No market-price, sovereign-finance or tariff recommendation follows.

The scenario grid is chosen for transparency and mathematical stress testing, with no empirical probability calibration. Equal total capacity and CF prevent annual output from confounding the topology comparison but suppress possible differences in outages and maintenance. Reserve-energy provision uses a conservative rectangle over T; a validated replacement trajectory could require less battery energy, while delayed restoration could require more. Batteries and nuclear generation may have correlated external-hazard exposure unless their siting and connections are explicitly assessed.

The model nevertheless produces testable implications. Under a fixed operating state, crossing gP/N below L_0 should stop reducing this incremental reserve requirement. Enlarging the credible trip group should remove some apparent benefit of smaller modules. A delayed battery step should not change the instantaneous pre-response slope. Finally, any measured or optimized reserve saving substituted into Eq. 12 should bound the capital premium justified by that saving under equal-output financial assumptions. Failure of a real system to follow these implications identifies which simplifying assumption must be revised.

7 Conclusions

Electrical modularity is valuable for the screened frequency service only to the extent that it reduces the largest credible simultaneous loss. The derived framework makes that dependence explicit and connects it to incremental reserves and an affordable nuclear-capital premium. Two boundaries are decisive: subdivision ceases to reduce this reserve burden once the nuclear contingency is below the existing one, and delayed fast response cannot repair an initial frequency-rate violation.

In the central hypothetical 400 MW comparison, reducing the independent loss from 200 MW to 100 MW eliminates 54.28 MW of incremental response but supports only USD 61.08/kW of additional capital basis under the declared costs. Pair losses erase that advantage, while whole-site losses fail the selected initial-rate condition. The separate dynamic reference gives a central premium of USD 59.36/kW after reserve re-sizing. Nevertheless, the 3,645 finite-power comparisons include 85 analytical false admissions and 42 false rejections relative to that proxy, showing that close central agreement does not establish a uniformly conservative rule. These are reproducible model consequences, not demonstrated performance of SENI or a reactor supplier.

For the Dominican Republic, the framework provides a focused research task: establish credible electrical groupings and measured post-event response, then compare complete feasible portfolios with consistent cost and reliability boundaries. This analytical contribution can be assessed without a reactor prototype, but a national investment recommendation requires calibrated system modelling and independently demonstrated project readiness.

Declarations

Human and animal research

The work reports mathematical analysis, literature interpretation and synthetic computational scenarios. It includes no human-participant or animal interventions.

Author contributions

Baldo Alberto Luigi Dalporto: Conceptualization; Methodology; Project administration; Visualization; Writing—original draft; Writing—review and editing. Sabine Mary: Methodology; Validation; Supervision;

Writing—review and editing. Santiago Gallur: Conceptualization; Formal analysis; Methodology; Validation; Writing—review and editing.

Data availability

Yes. All scenario inputs, analytical formulas and principal results are reported in the manuscript. Supplementary file S1 provides the computational code and numerical outputs described in Appendix A, including analytical and dynamic results, time trajectories, classification records and figure-generation scripts. Published Dominican reference values are identified in references [8] and [9]. The reported calculations generate synthetic scenario outputs, not empirical reactor or grid observations. No proprietary grid records, reactor measurements or personal data were used.

Ethics statement

This analytical and computational study involved no human participants, animals, patient data, or biological specimens and required no participant consent or experimental ethics approval.

Funding

The authors declare that no financial support was received for the research, authorship, and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Declaration of generative AI assistance

OpenAI ChatGPT assisted with literature retrieval, formulation exploration, drafting, translation, code generation, computational checking and document preparation. The reported scenarios are computational constructs and were not represented as empirical reactor or grid data. Responsibility for evaluating the equations, checking the evidence, approving author contributions and accepting the submitted manuscript rests with the human authors.

Appendix A Derivation and reproduction details

Before fast response activates, the accumulated deficit is the integral of L − at up to t_d = min(τ, L/a), giving the first two terms of Eq. 4. If the remaining deficit after activation is positive, it forms a triangle of height L − B − a τ and duration (L − B − a τ)/a. The sum gives Eq. 4. This construction also covers B = 0 and activation after primary balance; in the latter case the post-activation term is zero.

For the positive-response branch, expanding the sum yields A = (L − B)²/(2a) + B τ. Setting A = MF gives B² + (2a τ − 2L)B + L² − 2aMF = 0. The lower root is Eq. 5. Its domain requires the pre-activation area to be no larger than MF. Otherwise the permitted depression has already been exceeded before any B can be delivered. The initial-rate condition must be checked independently even when Eq. 5 has a finite value.

Monotonicity follows directly from the positive-deficit representation. Increasing L cannot decrease the required B when the other parameters are fixed. Since max(L_0, gP/N) is non-increasing in N for fixed g, Eq. 6 is non-increasing until the loss reaches L_0. This proves Proposition 1. Subtracting annualized plant and reserve costs for two equal-output cases gives 1000 P CRF(r,n) Δ K ≤ A_b,b − A_b,c, which proves Proposition 2.

The central calculation uses B_0 = 20.97619286 MW and B_{200} = 75.25390348 MW. Their difference is 54.27771062 MW. The nameplate energy difference is 54.27771062 × 0.25/(0.95 × 0.80) = 17.85451007 MWh. The battery capital is 1000 × (200 × 54.27771062 + 200 × 17.85451007) = USD 14,426,444.14. With CRF(0.08,15) and a 0.02 annual cost fraction, annual expense is USD 1,973,963.79. Dividing by 1000 × 400 × CRF(0.08,60) gives the USD 61.07716275/kW separation premium.

Run reproduce.py with Python 3.10 or later to regenerate results.json using only the standard library. The file contains the five central configurations, all 405 architecture-state evaluations, nine finance cases and verification outputs. Non-finite nadir solutions are stored as null and are flagged as failures; they are not averaged into a result. Run figures.py with NumPy and Matplotlib installed to regenerate Figures 1–3. Run benchmark_dynamics.py with NumPy and SciPy to regenerate dynamic_results.json, dynamic_cases.csv and the three central trajectory archives; dynamic_figures.py generates Figures 4–5. The README specifies all commands and dependencies. Calculations retain full precision until display rounding.

References

[1] International Atomic Energy Agency. Electric Grid Reliability and Interface with Nuclear Power Plants. IAEA Nuclear Energy Series NG-T-3.8. Vienna: IAEA; 2012. https://www-pub.iaea.org/MTCD/Publications/PDF/Pub1542_web.pdf

[2] Poudel B, Upadhyaya Subedi D, Gokaraju R. Generation planning of small modular reactor based nuclear power plants in electrical grids. Annals of Nuclear Energy. 240:112671. Published online 1 August 2026; assigned to the January 2027 issue. https://doi.org/10.1016/j.anucene.2026.112671

[3] Liu X, Fang X, Gao N, Yuan H, Hoke A, Wu H, Tan J. Frequency nadir constrained unit commitment for high renewable penetration island power systems. IEEE Open Access Journal of Power and Energy. 2024;11:141–153. https://doi.org/10.1109/OAJPE.2024.3370504

[4] Cruz-De-Jesús E, Martínez-Ramos JL, Marano-Marcolini A, Gómez-Expósito A. Economic assessment of battery energy storage systems for frequency regulation reserve provision: A case study of the Dominican Republic. Energy for Sustainable Development. 2025;88:101749. https://doi.org/10.1016/j.esd.2025.101749

[5] Mignacca B, Locatelli G. Economics and finance of Small Modular Reactors: A systematic review and research agenda. Renewable and Sustainable Energy Reviews. 2020;118:109519. https://doi.org/10.1016/j.rser.2019.109519

[6] Eash-Gates P, Klemun MM, Kavlak G, McNerney J, Buongiorno J, Trancik JE. Sources of cost overrun in nuclear power plant construction call for a new approach to engineering design. Joule. 2020;4:2348–2373. https://doi.org/10.1016/j.joule.2020.10.001

[7] Badrianto IA, Irnawan R, Gabbar HA, Sarjiya. Techno-economic and probabilistic assessment of grid-connected nuclear–renewable hybrid energy systems for island power systems: A case study of Bangka Belitung. Nuclear Engineering and Technology. 2026;58(11):104482. https://doi.org/10.1016/j.net.2026.104482

[8] Ministerio de Energía y Minas de la República Dominicana. Informe de Desempeño de las Empresas Eléctricas Estatales diciembre 2025. Issued 13 February 2026, pp. 9 and 19. https://mem.gob.do/wp-content/uploads/2026/02/Informe_Desempeno_EEE_diciembre_2025-.pdf

[9] Ministerio de Energía y Minas de la República Dominicana. La demanda eléctrica máxima alcanza niveles históricos de 4,077.97 MW. 18 June 2026; observation dated 17 June 2026. https://mem.gob.do/la-demanda-electrica-maxima-alcanza-niveles-historicos-de-4077-97-mw/

[10] Generation IV International Forum. Lead Fast Reactors LFR. https://www.gen-4.org/generation-iv-criteria-and-technologies/lead-fast-reactors-lfr

[11] PowerWorld Corporation. Governor TGOV1 and TGOV1D. Transient stability model documentation and block diagram. https://www.powerworld.com/WebHelp/Content/TransientModels_PDF/Generator/Governor/Governor%20TGOV1.pdf

[12] North American Electric Reliability Corporation. Reliability Guideline Application Guide for Modeling Turbine-Governor and Active Power-Frequency Controls in Stability Studies. June 2019.

https://www.nerc.com/comm/RSTC_Reliability_Guidelines/Reliability_Guideline-Application_Guide_for_Turbine-Governor_Modeling.pdf

[13] International Atomic Energy Agency. Milestones in the Development of a National Infrastructure for Nuclear Power. IAEA Nuclear Energy Series NG-G-3.1 Rev. 2. Vienna: IAEA; 2024. https://doi.org/10.61092/iaea.zjau-e8cs

[14] International Energy Agency. The Path to a New Era for Nuclear Energy. Paris: IEA; 2025. https://www.iea.org/reports/the-path-to-a-new-era-for-nuclear-energy

[15] Krall LM, Macfarlane AM, Ewing RC. Nuclear waste from small modular reactors. Proceedings of the National Academy of Sciences. 2022;119(23):e2111833119. https://doi.org/10.1073/pnas.2111833119

Public web records were checked on 28 September 2026. Bibliographic issue dates can differ from advance-online dates. Literature-positioning statements based on abstracts or institutional records are limited to their documented scope.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryMaterialEPJQT3009.docx