electronics

MDPI

Article

Causal Energy Accessibility and Capacity Allocation for Transient Voltage Security in Converter Networks

Jishuo Qin [1,2,*], Fan Li [1], Binqi Ma [1], Rui Shi [3], Dan Wang [1] and Guodong Guo [1]

[1] State Grid Economic and Technology Research Institute Co., Ltd., Beijing 102209, China; lifan658@163.com (F.L.); mabinqi@chinasperi.sgcc.com.cn (B.M.); wangdan@chinasperi.sgcc.com.cn (D.W.); guoguodong@chinasperi.sgcc.com.cn (G.G.)
[2] School of Electrical and Electronic Engineering, North China Electric Power University, Beijing 102209, China
[3] State Grid Corporation of China, Beijing 100031, China; shirui@sgcc.com.cn
* Correspondence: js-qin@ncepu.edu.cn

Abstract

Unused remote headroom may coexist with local overvoltage when transfer power or its timing prevents surplus energy from reaching the remote resource. This study derives a causal energy-transfer envelope retaining initial releasable energy, injection timing, and converter-mode limits. The envelope gives an exact feasibility condition for an ideal two-region system; its constant-capacity specialization separates resource insufficiency from fixed-controller violations, and a bandwidth penalty gives the lag and power-ramp capability a causal controller needs to realize the ideal transfer. A temporal network optimization retains a common storage history, minimizes persistent absorption additions, prices their location through its dual, and translates them into capacitance with coupled energy margins. In 2592 two-region cases, 670 of 2160 aggregate-capacity passes exhibit fixed-droop overvoltage, including 144 with a feasible ideal transfer, of which a causal controller with a 0.5 ms lag realizes 134. In 576 network cases, shared-equipment prefix bounds leave 59 infeasible cases unresolved. In a three-terminal resistance–inductance (R–L) circuit with finite converter–current response, a necessary voltage bound preserving the sign of the impulse response excludes 30 of 60 cases; validated ball arithmetic confirms every exclusion, the extension through the recovery interval changes none, and an available-power ceiling raises the count to 37. Cases that satisfy the bound remain unresolved until a controller is simulated. A switching-function modular multilevel converter model reproduces the averaged peaks within 0.8% at 30 kJ/MVA station energies and within 4.1% at the lower nominal energies, and 1200 randomly perturbed cases contain no exclusion that meets a safe trajectory.

Keywords: converter networks; transient voltage security; energy accessibility; current limiting; mode switching; storage allocation; necessary conditions; modular multilevel converter

Check for updates

Academic Editor: Andrea Ballo

Received: 8 October 2026
Revised: 11 November 2026
Accepted: 13 November 2026
Published: 16 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

A fault can exhaust a local capacitor's voltage headroom while remote devices retain unused absorption capacity. Whether that remote capacity can prevent overvoltage depends on the initial energy states, the timing of source injection, and the transfers permitted by converter limits and operating modes. This paper determines when these restrictions make energy accommodation impossible and when an admissible transfer exists despite a voltage violation under a particular controller.

1.1. Engineering Motivation

When receiving-end power falls below renewable generation, the surplus charges capacitive components until generation decreases, export recovers, or another device absorbs the imbalance. Feltes et al. developed coordinated fault ride-through control for wind farms connected through voltage-source-converter high-voltage direct-current (VSC-HVDC) transmission [1]. Ramtharan et al. compared the fault ride-through of fully rated converter wind turbines with AC and DC transmission, with a DC chopper dissipating the surplus during onshore faults [2]. Tzelepis et al. investigated DC-voltage control for wind farms connected to multi-terminal networks through modular multilevel converters (MMCs) [3]. Cao et al. studied full-bridge dissipation coordinated with DC circuit breakers for wind-power integration [4]. These controls and devices remove surplus power only within their transfer, absorption, and response-time limits.

For a sending region with little local headroom, remote absorption is limited by the power that intervening equipment can transfer before the local voltage reaches its limit. Several transfer paths may share one converter rating, and blocking or restoration changes the available paths. A connection restored after a voltage crossing cannot prevent that crossing, regardless of its rating. Local absorption or earlier generation reduction can instead reduce the energy that must cross the constrained interface.

The initial voltage determines the remaining absorption margin; rated energy alone does not. Li et al. constructed long-term wind and photovoltaic output scenarios that preserve temporal correlations [5]. The present two-region study instead prescribes the initial source power, the transfer availability, the curtailment delay, and the curtailment rate, so its case counts describe parameter combinations rather than probabilities.

1.2. Related Work and Remaining Gap

Ilves et al. related MMC capacitor-storage requirements to power transfer and permissible capacitor-voltage variation [6]. Samimi et al. connected internally stored MMC energy to DC-bus voltage control and demonstrated its effect on transients [7]. Freytes et al. showed that energy-control choices in MMC-based multi-terminal DC grids affect voltage transients and the contribution of converter energy to the DC system [8]. Thus, an absorption resource must be described by both its energy limits and the controlled power exchanges through which it participates in the disturbance.

Bergna-Diaz et al. developed voltage-based MMC state-space equations with constant steady-state equilibria and compared them with detailed converter simulations [9]. Their energy-based formulation provides full and reduced representations of converter dynamics [10]. Freytes et al. identified interactions between DC current and internal stored energy that can produce weakly damped dynamics under circulating-current suppression [11]. These interactions require the physical energy states and their control coupling to be specified before capacitor headroom is used in a network feasibility calculation.

Gnanarathna et al. developed the detailed equivalent MMC model for electromagnetic-transient programs [12], and Saad et al. compared averaged and simplified MMC models with it [13]. Merlin and Green related the cell capacitance to the stored-energy requirement [14], which, with the sizing analysis of Ilves et al. [6], places MMC designs at 30–40 kJ/MVA. The present assessment uses the averaged port for the complete parameter studies and a switching-function MMC model to check its peaks.

Tavakoli et al. examined fault ride-through mechanisms in the DC-voltage and active-power control loops of an MMC-HVDC link, accounting for fault type and pre-fault power direction [15]. Wan et al. reviewed how MMC topology determines DC-fault behavior and fault-blocking capability [16]. A blocking command changes the available conduction

paths; it does not by itself remove capacitor or inductor energy. The corresponding transfer limits must therefore retain the relevant topology, residual currents, and magnetic energy.

Fan et al. reviewed current-limiting control under symmetrical disturbances and its effects on grid-forming operation [17]. Taul et al. developed a fault-mode controller that limits current while preserving voltage-mode characteristics [18]. Qoria et al. derived critical-clearing-time relations for grid-forming converters with virtual-impedance current limitation [19]. Rokrok et al. examined how the saturated current-reference angle affects transient stability and the return to voltage control [20]. These results identify restrictions on the power exchange and recovery behavior that an admissible converter-mode description must preserve.

Fan and Wang developed an equivalent circuit for a grid-forming converter with a circular current limiter and analyzed its synchronization behavior [21]. Kkuni and Yang quantified the effect of current limitation on transient stability and investigated virtual active-power feedback [22]. Baeckeland et al. presented a unified large-signal model for several grid-forming current limiters, evaluated using numerical analysis, electromagnetic-transient models, and hardware [23]. Here, mode descriptions constrain real terminal-power exchange; virtual control quantities do not add to the physical absorption inventory.

Willems defined dissipativity through storage and supply functions, distinguishing available storage from required supply [24]. The present problem constrains the increase in physical storage before an upper limit is reached, under a specified set of admissible transfers. Ford and Fulkerson studied dynamic flows with transfer capacities and a prescribed horizon [25]. Hoppe and Tardos treated multi-source, multi-sink transshipment with capacities and transit times [26]. These established methods provide the temporal transfer framework; the electrical assessment must additionally associate each carried energy state with its voltage-derived absorption and release limits.

Song et al. used cutset properties to explain angle separation and estimate critical energy for power-system transient stability [27]. Dinh et al. combined an improved minimum-cut procedure with storage siting and sizing for expansion planning and congestion management [28]. For fault-energy assessment, regional injections, absorption margins, dissipation, and mode-dependent transfers must refer to the same disturbance and control permissions. Otherwise, a configuration comparison can attribute the benefit of greater source-curtailment authority to added storage or transmission capability.

Liu et al. developed control synthesis with a guaranteed region of attraction for DC microgrids [29]. Chang et al. derived large-signal sufficient conditions for DC grids with controlled converters and constant-power loads [30]. Failure of a sufficient stability condition does not establish instability. The necessary tests considered here use the opposite implication: if a model containing every permitted physical trajectory cannot accommodate the energy, no such trajectory can remain within its limits. A passing test remains unresolved until the omitted dynamics and the selected controller are checked. Mazenc and Bernard bounded the states of linear systems with bounded disturbances by splitting the response into its nonnegative and nonpositive parts [31]; Section 5 applies this construction to the converter network with current envelopes derived from the curtailment authority.

A network assessment must distinguish three causes of unavailable remote headroom: source energy arriving after a useful transfer interval, simultaneous transfers competing for one device, and incompatible storage states assumed by separate regional tests. It must also determine when regional inequalities are exact and when retaining the complete state history changes the minimum resource addition. These distinctions require a feasible construction for the tractable cases and explicit counterexamples to broader sufficiency claims.

1.3. Contributions and Organization

The contributions are as follows:

1. A causal two-region envelope includes source-energy arrival, initially releasable energy, transfer capacity, and finite remote absorption. A constructive transfer gives an exact feasibility condition for the stated ideal model. Its constant-capacity, nonincreasing-injection specialization distinguishes resource infeasibility from voltage violations caused by fixed control.
2. A piecewise network formulation retains energy conservation, persistent absorption capacity, mode-dependent transfers, and shared-device limits. A serial-transfer example passes every fixed-region, all-interval energy inequality but remains infeasible. The allocation problem and its dual quantify the required addition and its location within this feasible set.
3. Matched configurations and structured parameter studies compare storage location, transfer reinforcement, and source reduction. A three-terminal average-value circuit retains finite current response, line magnetic energy, and resistive losses; a necessary voltage bound preserves the sign of its impulse response, is evaluated with a validated numerical enclosure, and covers the recovery interval. A causal finite-bandwidth controller, coupled capacitance sizing, a switching-function MMC model, and randomized parameter studies test the averaged-model conclusions.

The guarantees concern finite-window energy and voltage feasibility under the stated source, transfer, and mode permissions. Table 1 separates the established methods from the derivations of this paper.

Table 1. Origin of each result: established basis and derivation of this paper.

Result	Established Basis	Derived in This Paper
Energy intervals and regional tests, Equations (1)–(9)	Energy conservation; storage and supply functions [24]	Voltage-derived node margins; prefix and interval tests with shared devices
Causal envelope, Equations (11)–(14)	Min-plus service bounds; reflected cumulative process [32]	Releasable energy and remote cap; necessary and sufficient ideal-interface condition
Constant-capacity specialization, Equations (15)–(17)	None	No-discharge construction, exact requirement, realization gap and bandwidth penalty
Temporal optimization and dual, Equations (19)–(21), (23), (27) and (28)	Dynamic flows with storage [25,26]; linear-programming duality	Persistent additions on one storage history; conditions (E1)–(E4) and (N1); coupled sizing, Equation (24)
Moving-region cut, Equation (26)	None	Derivation and reversed-serial counterexample
Circuit voltage bound, Equations (35)–(42)	Comparison theorem; variation of constants; sign-split enclosure [31]; ball arithmetic [33]	Envelopes from the curtailment authority, recovery extension, iterated coupling, validated exclusion

Sections 2–4 define the physical limits, derive the two-region condition, and develop the temporal network formulation. Section 5 describes the studies, the parameter basis, the bound evaluation, and the MMC model; Section 6 reports the results; Sections 7 and 8 state the engineering use, the scope of the models, and the conclusions.

2. Physical Energy Limits and Regional Accessibility

An injected energy surplus raises local stored energy unless an interface transfers it elsewhere or a sink dissipates it. Remote nameplate capacity contributes only through transfers available before the local limit is reached. We assess whether the initial energy, storage limits, and permitted transfers admit a bounded trajectory over the disturbance window.

The accounting uses physical component energies within the storage-and-supply framework [24]. MMC capacitor sizing relates energy requirements to operating limits [6], and internal capacitor energy can support DC-voltage regulation [7]. Energy-based MMC models describe the internal storage dynamics [10]; the control structure determines how much of that energy participates in a multi-terminal transient [8]. The inventory here contains only physical storage, so a virtual-inertia coefficient without an associated energy source contributes no additional capacity.

Measuring storage relative to the initial operating point separates absorption headroom from energy available for early discharge:

$$e_i(t) = E_i(t) - E_{i0}, -R_i \le e_i(t) \le H_i, e_i(0) = 0. \tag{1}$$

Here, E_i and E_{i0} are physical and initial stored energies, H_i is remaining absorption headroom, and R_i is releasable energy above the lower operating limit. Energy is measured in MJ. The initial condition $e_i(0) = 0$ permits nonzero E_{i0}; the additional restriction $R_i = 0$ prohibits discharge below that initial state.

For a capacitor, the upper and lower voltage limits map directly to the two energy margins:

$$\begin{cases} E_{i0} & = \frac{C_i V_{i0}^2}{2}, v_i = \sqrt{1 + \frac{e_i}{E_{i0}}}, \\ H_i & = E_{i0}\left(v_{i,\max}^2 - 1\right), R_i = E_{i0}\left(1 - v_{i,\min}^2\right). \end{cases} \tag{2}$$

Here, C_i is capacitance, V_{i0} is initial voltage, and $v_i = V_i / V_{i0}$ is per-unit voltage. Evaluating the first expression in F and V gives J, which are converted to MJ. At fixed voltage limits and nominal voltage, both H_i and R_i scale with capacitance. Increasing H_i while holding R_i fixed therefore represents added absorption capability with an unchanged discharge allowance, rather than the complete effect of adding capacitance.

For a directed graph of ideal controlled interfaces, equal and opposite terminal powers give the nodal balance

$$\dot{e}(t) = p(t) + Bq(t) - d(t), 0 \le q_e(t) \le \overline{q}_e(t). \tag{3}$$

The incidence matrix B has -1 at an interface tail and $+1$ at its head. The inputs p_i are net injections, $d_i \ge 0$ are dissipation or modeled external withdrawals, and q_e are directed transfers. Powers are in MW and time in s, giving energy integrals in MJ. Two opposite arcs represent a bidirectional device only together with its shared rating; independent arc limits alone allow simultaneous use beyond that rating.

A converter or collection channel shared by several interfaces imposes a common limit on their simultaneous transfers:

$$\sum\nolimits_{e \in \mathcal{G}_g} a_{ge} q_e(t) \le \overline{P}_g(t), a_{ge} \ge 0. \tag{4}$$

The set $\mathcal{G}_g$ contains interfaces using resource g, a_{ge} is their nonnegative usage coefficient, and $\overline{P}_g$ is available power. Normal, limited, blocked, and recovering operation change this power limit. Blocking sets the affected path limit to zero while retaining its stored energy. With uncertain switching times, exclusion of all allowed schedules requires infeasibility over their union; infeasibility for one schedule excludes only that schedule.

For a passive resistance–inductance line, the terminal-power difference accumulates magnetic energy and supplies resistive loss:

$$L_e\dot{i}_e = V_a - V_b - r_e i_e, V_a i_e - V_b i_e = \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{L_e i_e^2}{2}\right) + r_e i_e^2. \tag{5}$$

Here, i_e, L_e, and r_e are line current, inductance, and resistance. In SI units, the terminal powers and losses are in W and magnetic energy in J. Including the line as a storage element preserves this balance. Omitting its permitted energy increase or loss can overstate the surplus assigned to capacitors and produce a false exclusion. In the absence of impulsive currents and voltages, switching preserves capacitor voltages and inductor currents; an imposed state jump requires an accounted energy transfer.

Summing nodal balances over a region δ cancels internal transfers:

$$e_\delta(t) = W_\delta(t) + F_{\mathrm{in},\delta}(t) - F_{\mathrm{out},\delta}(t) - D_\delta(t), W_\delta(t) = \int_0^t p_\delta(\tau)\mathrm{d}\tau. \tag{6}$$

Regional subscripts denote sums. The cumulative incoming, outgoing, and dissipated energies are $F_{\mathrm{in},\delta}$, $F_{\mathrm{out},\delta}$, and D_δ, all in MJ. Since directed incoming transfer is nonnegative, a lower bound $\underline{W}_\delta$ on injection and an upper bound U_δ on outgoing transfer plus dissipation give the necessary storage condition

$$\underline{W}_\delta(t) - U_\delta(t) \le H_\delta, H_\delta = \sum\nolimits_{i\in\delta} H_i. \tag{7}$$

A positive deficit at any time contradicts the regional upper limit, even if the clearing-time deficit is nonpositive. Satisfaction leaves two possible restrictions unresolved: different regional maxima may require incompatible transfer schedules, and energy may arrive after the available transfer interval.

The bounds cover every trajectory in the declared control class: $\underline{W}_\delta \le W_\delta$ and $U_\delta \ge F_{\mathrm{out},\delta} + D_\delta$. For a DC port, the permitted upper voltage times the current rating gives a valid power upper bound. Prescribed reactive-current priority further restricts an AC port's active-power capability. A bound inferred from one controller trajectory does not cover other permitted controls. Under parameter uncertainty, using the largest admissible absorption/export capabilities and smallest unavoidable injection preserves exclusion when the resulting deficit remains positive.

Maximizing outward transfer under shared-equipment limits removes duplicated use of a converter from the regional bound:

$$\overline{F}_{\mathrm{out},\delta,k} = \max_{f_k\in\mathcal{Q}_k} \sum\nolimits_{e\in\delta^+(\delta)} f_{e,k}, U_\delta(t_j) = \sum\nolimits_{k=1}^{j} \overline{F}_{\mathrm{out},\delta,k} + \overline{D}_\delta(t_j). \tag{8}$$

The set $\mathcal{Q}_k$ contains integrated interface and shared-resource limits, $\delta^+(\delta)$ is the outgoing cut, and $\overline{D}_\delta$ bounds dissipation. Maximization omits the exporting terminals' prior energy histories, so it cannot underestimate outward capability when $\mathcal{Q}_k$ contains every admissible interval transfer. Summing independent ratings discards the shared constraint and can only weaken the bound.

Starting the balance at an intermediate time s replaces the known initial state by its lower bound and yields

$$\int_s^t p_\delta(\tau)\mathrm{d}\tau - U_\delta(s,t) \le H_\delta + R_\delta, R_\delta = \sum\nolimits_{i\in\delta} R_i. \tag{9}$$

At $s = 0$, the known zero incremental state tightens the right-hand side to H_δ. At $s > 0$, the condition permits the region to start at $-R_\delta$. It detects injection trapped after a channel closes, but omits any higher starting energy forced by earlier injections. The temporal model retains that accumulated state.

Unlike cutset conditions for synchronization stability [27], these inequalities constrain physical energy accumulation. A bound on total regional energy does not bound each internal capacitor separately. Including individual nodes and their unions adds the corresponding spatial constraints; angle, modulation, and internal converter states remain outside this energy-only representation.

Table 2 distinguishes the MW/MJ energy-network variables from the SI circuit variables. Circuit energies computed in J are converted to MJ for comparison. The node quantity R_i denotes releasable energy, whereas r_e or R_e denotes line resistance; similarly, $C(t)$ is cumulative transfer capability and C_i is capacitance.

Table 2. Principal variables and their physical meanings.

Symbol	Meaning	Unit
E_{i0},e_i	Initial physical energy and its incremental change	MJ
H_i,R_i	Upper absorption headroom and initially releasable energy	MJ
p_i,q_e,d_i	Net injection, directed interface transfer and dissipation	MW
W,C,X	Cumulative source energy, transfer capability and actual transfer	MJ
$\mathcal{S}$, $\mathcal{G}_g$	Selected energy region and interfaces sharing a device	Set
x_k,f_k,u_k,z_k	Interval-end state and integrated energy vectors	MJ
$\overline{F}_{e,k}$, $\overline{G}_{g,k}$	Independent-interface and shared-device interval limits	MJ
h_i,$J^\star$	Persistent added headroom and optimal weighted addition	MJ; weighted MJ for dimensionless costs
π_i,$\mu_{e,k}$	State-headroom and interface-capacity dual multipliers	Objective units per MJ
V_j,C_j,L_e,R_e	Circuit voltage, capacitance, inductance and resistance	V, F, H, ohm
i_s,i_r,i_1,i_2	Source, receiver and two passive-line currents	A
M_V,T	Controller-specific upper-voltage margin and interruption duration	p.u., s
τ_q,r_q,Δ_q	Interface current-loop lag, power-ramp limit and bandwidth penalty	s, MW/s, MJ
γ_i,ΔC_i	Release-to-absorption ratio and added capacitance of node i	Dimensionless, F

Figure 1 passes the component and interface limits to regional exclusion tests and temporal capacity allocation. A feasible ideal schedule then provides a transfer target for validation with the retained converter and line dynamics.

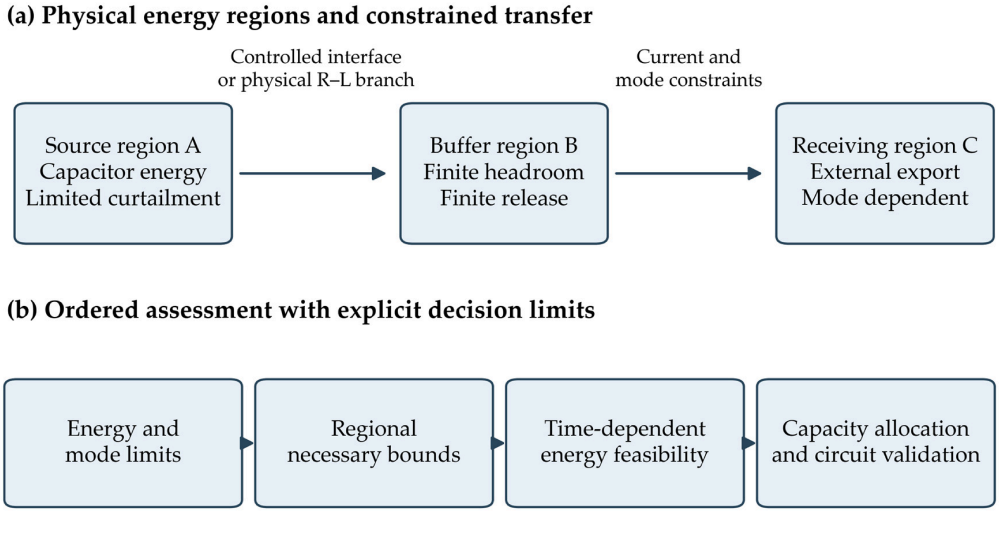


Figure 1. Energy regions and assessment sequence: (**a**) physical storage and restricted transfer; (**b**) regional bounds, temporal feasibility and configuration validation.

3. Causal Accessibility and the Two-Region Energy Boundary

Consider two regions with prescribed nonnegative injection $p(t)$ at A, no independent injection, withdrawal, or dissipation at B, and an ideal one-way interface $0 \le q(t) \le c(t)$. Cumulative injection, available service, and actual transfer determine the two stored-energy increments:

$$W(t) = \int_0^t p(\tau)\mathrm{d}\tau, C(t) = \int_0^t c(\tau)\mathrm{d}\tau, X(t) = \int_0^t q(\tau)\mathrm{d}\tau, e_A = W - X, e_B = X. \tag{10}$$

Cumulative arrival and service bounds follow the established network-calculus framework [32]. The reflected cumulative construction used below is a classical operation, applied here with voltage-derived absorption and release margins and the specified electrical source/interface permissions.

The functions W and C are nondecreasing and absolutely continuous. The separate limits $X \le H_B$ and $X \le C$ do not constrain when transferable energy arrives. For any intermediate time s, remaining interface service limits transfer after s:

$$X(t) \le X(s) + C(t) - C(s) \le W(s) + R_A + C(t) - C(s), 0 \le s \le t. \tag{11}$$

The second inequality substitutes $X(s) \le W(s) + R_A$, obtained from $e_A(s) \ge -R_A$. Thus, early transfer can use initial releasable energy but not future injection. Taking the tightest intermediate-time bound gives the greatest transfer compatible with the interface, the local lower bound, and the remote upper bound:

$$X_{\max}(t) = \min\{H_B,\ C(t),\ \inf_{0\le s\le t}[W(s) + R_A + C(t) - C(s)]\}. \tag{12}$$

The infimum limits export to energy available before each service interval. If a channel closes before a later pulse, its early transfer is bounded by the energy then present, including R_A; increasing an already sufficient early rating does not absorb more of the later pulse.

The following construction attains this upper envelope by subtracting service lost when the local lower energy limit is reached:

$$Y(t) = \max\left\{0, \sup_{0\le s\le t}[C(s) - W(s) - R_A]\right\}, \tilde{X}(t) = C(t) - Y(t), X^\star(t) = \min\left\{H_B, \tilde{X}(t)\right\}. \tag{13}$$

The nondecreasing function Y records unused service. Above the local lower limit, $\dot{Y} = 0$ and $\dot{\tilde{X}} = c$. At that limit, $\dot{Y} = \max(c - p, 0)$, so $\dot{\tilde{X}} = \min(c, p)$ almost everywhere. Therefore, $0 \le \dot{\tilde{X}} \le c$ and $W - \tilde{X} \ge -R_A$. Capping this nondecreasing transfer at H_B preserves the interface and local lower limits and enforces the remote upper limit. Substitution gives $X^\star = X_{\max}$.

This maximal transfer minimizes local energy pointwise. The remaining local upper constraint is therefore

$$W(t) - X_{\max}(t) \le H_A \text{for all } t \in [0, T]. \tag{14}$$

No admissible transfer can exceed $X_{\max}$, which establishes necessity. Conversely, $q = \dot{X}^\star$ satisfies the interface limit, both lower energy limits, and the remote upper limit; the displayed condition supplies the local upper limit. The condition is thus necessary and sufficient for the specified ideal one-way interface.

For constant capacity $c(t) = \bar{q}$ and nonnegative, nonincreasing injection, the simpler cumulative bound also admits a constructive realization without initial discharge:

$$X_0(t) = \min\{W(t), H_B, \bar{q}t\}, e_{A,0}(t) = \max\{0, W(t) - H_B, W(t) - \bar{q}t\}. \tag{15}$$

The function W is concave with $W(0) = 0$. Wherever $X_0 = W \le \bar{q}t$, concavity gives $p(t) \le W(t)/t \le \bar{q}$; the other branches have derivatives $\bar{q}$ or zero. Hence, $0 \le \dot{X}_0 \le \bar{q}$ and

both regional energy increments are nonnegative. When $R_A > 0$, X_0 is a feasible no-discharge construction, although $X_{\max}$ may permit additional early transfer.

Maximizing the constructed local energy over the disturbance window gives the exact headroom requirement for this subclass:

$$H_{A,\mathrm{req}} = \max_{0\le t\le T}[W(t) - \min\{H_B, \bar{q}t\}]_+, H_A \ge H_{A,\mathrm{req}}. \tag{16}$$

Here, $[z]_+ = \max(z,0)$. The regional necessary bound equals the largest local energy in the constructed trajectory, establishing sufficiency. A fixed controller that violates the voltage limit while this condition holds therefore has a realization gap relative to another dispatch with the same source and interface limits.

A causal controller that follows the constructed transfer through a current loop with time constant τ_q and a power-ramp limit r_q keeps the local energy within a computable distance of the ideal trajectory. With reference $q_{\mathrm{ref}} = \dot{X}_0$ and actuator $\tau_q\dot{q} = q_{\mathrm{ref}} - q$, $|\dot{q}| \le r_q$, the lag withholds $X_0(t) - X(t) = \tau_q[q(t) - q(0)]$ and a rate-limited rise to $\bar{q}$ at most $\bar{q}^2/(2r_q)$, so

$$e_A(t) \le e_{A,0}(t) + \Delta_q, \Delta_q = \tau_q\bar{q} + \frac{\bar{q}^2}{2r_q}, \tag{17}$$

where Δ_q is the bandwidth penalty in MJ, with τ_q in s and r_q in MW/s. The controller keeps the local voltage within its limit whenever $H_A \ge H_{A,\mathrm{req}} + \Delta_q$, and a predictive stop that sets the reference to zero once e_B plus the energy still in transit reaches H_B preserves the remote limit. For an 80 MW interface, the penalty is 0.04 MJ at a 0.5 ms lag and 0.36 MJ with a 10 MW/ms ramp limit.

A source with reaction delay d, maximum downward ramp r_p, and minimum output P_{floor} has the nonincreasing lower envelope

$$p_{\min}(t) = \max\{P_{\mathrm{floor}}, P_0 - r_p[t-d]_+\}, W(t) = \int_0^t p_{\min}(\tau)\mathrm{d}\tau. \tag{18}$$

The powers P_0 and P_{floor} are in MW, d is in s, and r_p is in MW/s. A source following this envelope satisfies the preceding specialization. With fixed transfer bounds, a positive regional deficit computed from this lower injection excludes every larger permitted injection. Feasibility at minimum injection has a different implication: in a network with fixed withdrawals and lower-voltage constraints, higher injection can prevent shortage, so fixing all sources at their minima is not a general outer relaxation.

4. Temporal Network Feasibility and Resource Allocation

A feasible sequence of interface modes shares one continuous storage history; each interval starts from the energy left by the preceding interval.

Time-dependent network flow supplies the optimization structure [25,26]. Physical storage bounds define the carried inventory, mode-dependent interfaces define available connections, and shared ratings couple simultaneous transfers. A necessary electrical test follows when every permitted physical trajectory projects into this constrained set.

Place time boundaries at changes in source expression, interface availability, and shared capability. Integrating each transfer and input over an interval defines

$$\begin{cases} x_{i,k} & = e_i(t_k), f_{e,k} = \int_{t_{k-1}}^{t_k} q_e(t)\mathrm{d}t, \\ u_{i,k} & = \int_{t_{k-1}}^{t_k} p_i(t)\mathrm{d}t, z_{i,k} = \int_{t_{k-1}}^{t_k} d_i(t)\mathrm{d}t. \end{cases} \tag{19}$$

Here, $k = 1,\ldots,K$, $\Delta t_k = t_k - t_{k-1}$ is in s, and $x_{i,k}, f_{e,k}, u_{i,k}, z_{i,k}$ are in MJ. Integration of the nodal differential equations gives

$$x_k - x_{k-1} = u_k + Bf_k - z_k, x_0 = 0, -R_i \le x_{i,k} \le H_i. \tag{20}$$

The shared variables x_k retain energy already absorbed when a mode changes. The lower limit $-R_i$ also limits pre-discharge before a later injection.

Integrating the instantaneous independent and shared power limits gives interval transfer bounds:

$$0 \le f_{e,k} \le \overline{F}_{e,k} := \int_{t_{k-1}}^{t_k} \overline{q}_e(t)\mathrm{d}t, \sum\nolimits_{e\in\mathcal{G}_g} a_{ge} f_{e,k} \le \overline{G}_{g,k} := \int_{t_{k-1}}^{t_k} \overline{P}_g(t)\mathrm{d}t. \tag{21}$$

The capacities $\overline{F}_{e,k}$ and $\overline{G}_{g,k}$ are in MJ. Admissible source and dissipation sets constrain u and z. For signed powers, opposite directional arcs with a common absolute-power limit preserve the device rating. Relaxing additional physical couplings gives a necessary test only if the resulting set contains every permitted trajectory.

A continuous interface power satisfying $\left|\dot{q}_e\right| \le r_e^q$ also constrains adjacent interval averages:

$$\left|\frac{f_{e,k+1}}{\Delta t_{k+1}} - \frac{f_{e,k}}{\Delta t_k}\right| \le \frac{r_e^q}{2}(\Delta t_k + \Delta t_{k+1}). \tag{22}$$

Averaging $\left|q_e(t) - q_e(s)\right| \le r_e^q(t-s)$ over the later and earlier intervals gives this bound, because their mean time separation is $(\Delta t_k + \Delta t_{k+1})/2$. It is necessary for a continuous ramp-limited trajectory, but does not reconstruct one. A mode event permitting a power jump lies outside this continuity assumption.

The network cases use the integrated-capacity constraints without this optional ramp inequality.

One addition h_i applied to every time layer represents a persistent absorption resource. Minimizing its weighted size gives

$$\begin{cases} J^\star & = \min_{x,f,u,z,h} \sum_i c_i h_i, u \in \mathcal{U}, z \in \mathcal{D}, \\ h_i & \ge 0, x_{i,k} \le H_i + h_i \text{for every } k. \end{cases} \tag{23}$$

The weight c_i is a unit cost, or is dimensionless when the objective is in weighted MJ. The sets $\mathcal{U}$ and $\mathcal{D}$ contain admissible integrated injections and dissipation; the network experiments prescribe u and use $\mathcal{D} = \{0\}$. The original balances, transfer limits, and lower energy bounds remain imposed. Each h_i is counted once, and unit weights make $J^\star$ the minimum total added absorption headroom. Because R_i is fixed here, this objective differs from capacitor sizing that changes both voltage-derived energy margins.

When the addition is realized by capacitance at fixed voltage limits, both margins grow with it. Equation (2) gives the coupled sizing constraints

$$x_{i,k} \le H_i + h_i, x_{i,k} \ge -(R_i + \gamma_i h_i), \gamma_i = \frac{R_i}{H_i} = \frac{1 - v_{i,\min}^2}{v_{i,\max}^2 - 1}, \Delta C_i = \frac{2h_i}{V_{i0}^2(v_{i,\max}^2 - 1)}, \tag{24}$$

where γ_i is the dimensionless release-to-absorption ratio of node i and ΔC_i is the added capacitance in F for h_i in J; at 320 kV with a 1.15 p.u. upper limit, one MJ of upper headroom corresponds to 60.56 µF. The coupled problem remains linear with a feasible set containing that of Equation (23), so its optimum is at most the uncoupled addition; a unit cost per microfarad replaces the unit cost per MJ, and a dissipating device enters through $\mathcal{D}$ with its power rating.

Two statements fix what the optimization establishes. Exactness holds under four conditions: (E1) injections, interface limits and shared limits are constant on each event-aligned interval; (E2) every interface is a lossless, ideally controllable power port without current-loop state or ramp limit; (E3) the admissible states of each node form the energy interval of Equation (1); and (E4) no further dynamic constraint acts. A safe trajectory of this class supplies a feasible point through its interval-end states and integrated powers, and a feasible point reconstructs a continuous trajectory as

$$q_e(t) = \frac{f_{e,k}}{\Delta t_k}, e_i(t) = (1-\theta)x_{i,k-1} + \theta x_{i,k}, \theta = \frac{t - t_{k-1}}{\Delta t_k}. \tag{25}$$

The interpolation satisfies each energy interval by convexity, and the constant powers satisfy the independent and shared limits, so feasibility is necessary and sufficient under (E1)–(E4). Necessity extends to a wider physical class under one inclusion condition, (N1): the interval-end states and integrated powers of every admissible trajectory of that class satisfy the constraints. (N1) holds for a ramp-limited interface once Equation (22) is added. It holds for passive lines when line magnetic energy and dissipation are carried by variables bounded by $\frac{1}{2}L_e I_{e,\max}^2$ and $r_e I_{e,\max}^2 \Delta t_k$, or by equivalent headroom allowances. It holds for a bidirectional device represented by opposite arcs under a common absolute-power limit, and a one-way exclusion extends to that class because the reverse direction only raises the local energy. For compulsory withdrawals with lower energy limits, the injection enters as a decision variable, because fixing it at its minimum can create a shortage that no admissible trajectory has. Subdividing a constant segment preserves the exact-subclass optimum; in a valid outer relaxation, additional endpoint constraints can shrink the feasible set.

A region $\mathcal{S}_k$ that changes between intervals captures storage obligations missed by fixed-region tests. Summing its interval balances gives

$$\sum_{k=1}^{K} u_{\mathcal{S}_k,k} \le H_{\mathcal{S}_K} + \sum_{k=1}^{K-1} \left(H_{\mathcal{S}_k \setminus \mathcal{S}_{k+1}} + R_{\mathcal{S}_{k+1} \setminus \mathcal{S}_k} \right) + \sum_{k=1}^{K} \overline{F}_{\text{out},\mathcal{S}_k,k}. \tag{26}$$

The outgoing-energy bound includes any permitted dissipation. In the telescoping storage sum, a node leaving the region contributes at most H_i, while a newly included node contributes at most R_i through its negated boundary state. For a chain whose downstream channel opens first, early energy remains upstream, and the constrained region later expands to include the intermediate buffer. Fixed-region interval tests can separately choose incompatible starting states; the shared temporal states prevent that combination. With shared-device constraints, the dual inequalities of the linear program can also contain nonbinary regional weights.

Let $\lambda_{i,k} \ge 0$ multiply $x_{i,k} \le H_i + h_i$. Increasing existing headroom changes the right-hand side at every layer, giving the accumulated dual price

$$\pi_i = \sum\nolimits_k \lambda_{i,k}, J^\star(H + \delta H) \ge J^\star(H) - \pi^{\mathrm{T}} \delta H. \tag{27}$$

An optimal dual solution gives this supporting bound with all other problem data fixed. At a differentiable point, $\partial J^\star / \partial H_i = -\pi_i$; unit addition costs give $0 \le \pi_i \le 1$. At a degenerate optimum, alternative optimal multipliers can differ from a selected one-sided finite difference. Moving upper headroom from j to i has local benefit proportional to $\pi_i - \pi_j$ when release limits and electrical parameters remain fixed.

For rating Q_e with availability $a_{e,k}$, the constraint $f_{e,k} \le Q_e a_{e,k} \Delta t_k$ gives the marginal value

$$\frac{\partial J^\star}{\partial Q_e} = -\sum\nolimits_k \mu_{e,k} a_{e,k} \Delta t_k, \tag{28}$$

Here, $\mu_{e,k} \ge 0$ is the capacity multiplier, and the derivative applies where the optimal value is differentiable. Its unit is MJ/MW for a unit-MJ objective. Intervals with $a_{e,k} = 0$ contribute zero to the rating sensitivity. The rating derivative is zero when no available interval has a positive capacity multiplier.

A positive optimal addition, supported by a dual lower bound exceeding numerical tolerance, excludes feasibility with the original resources. A zero-addition solution supplies the transfer and state schedule for ideal reconstruction or subsequent dynamic evaluation. Solver failure leaves the feasibility status unresolved.

5. Study Design and Physical Validation Model

Three model classes are tested: a two-region controlled interface for the effect of resource location, a four-region energy network for mode timing and capacity allocation,

and a three-terminal resistance–inductance (R–L) circuit for finite current response. Each uses constructed parameters positioned in representative ranges; a switching-function MMC model tests the averaged circuit ports. The section specifies the comparison structure, the network cases, the circuit, the parameter basis, the necessary voltage bounds, the MMC model, and the numerical settings.

5.1. Comparison Structure and Numerical Reproducibility

Table 3 maps each experiment to its cases and supported conclusion. Ideal-interface schedules are tested for energy feasibility, whereas circuit trajectories are tested for node-voltage limits under specified controllers.

Table 3. Experimental structure and scope of the resulting conclusions.

Study	Cases and Independent Checks	Supported Conclusion
Two-region controlled interface	2592 parameter combinations; 77 integration checks	Exact ideal energy feasibility for the specified source class; controller realization gap
Causal two-region identity	300 varying-input and varying-capacity cases	Algebraic and numerical agreement within the ideal model
Four-region transfer network	576 parameter combinations; two optimization variants per case	Increment from shared equipment and chronological energy availability
Reversed serial interfaces	One analytically resolved three-region example	A changing limiting region can invalidate fixed-region completeness
Physical R–L network	60 parameter combinations; 18 representative and boundary cases	Controller-specific dynamic behavior and configuration trends
Circuit integration convergence	24 numerical convergence runs	Numerical implementation consistency for the average-value port model
Causal finite-bandwidth controller	1634 feasible cases at six settings; 144 gap cases at twelve	Realizability of the ideal transfer
Coupled capacitance sizing	576 network cases; circuit mapping	Headroom additions as equipment
Switching-function MMC model	4 principal, 60 grid and 36 MMC-consistent cases; station-energy sweep of 252 cases	Fidelity of the averaged port over station energy
Randomized robustness	1200 circuits, 600 perturbed principal and 3000 two-region samples	Classification under parameter uncertainty
Validated bound evaluation	60 cases in 256-bit ball arithmetic	Rigor of every exclusion

The two-region sweep uses every combination of local headroom fraction 0.1–0.9 in steps of 0.1; interface rating 0, 10, 20, 40, 60, 80, 100 or 150 MW; interruption duration 12, 24, 36 or 48 ms; maximum curtailment rate 1000, 2000 or 4000 MW/s; and delay 0, 4 or 8 ms. Total upper headroom is 3 MJ. All 2592 combinations, including voltage violations, are retained.

The two-region source starts at 100 MW and decreases to a 20 MW floor according to the prescribed delay and ramp; its energy is integrated analytically on each segment. Both initial voltages are 1 p.u., the upper limit is 1.15 p.u., and the voltage base is 320 kV. Moving headroom between the two independent capacitor buses changes their capacitances while preserving total capacitance, initial voltages, and the upper-voltage limit. A lossless controlled conversion interface connects the buses.

The fixed controller applies voltage droop subject to current and rated-power limits:

$$q(t) = sat_{[-q_c(t),q_c(t)]}\{K_v[v_A(t) - v_B(t)]\}, q_c(t) = \bar{q}\min\{1, v_A(t), v_B(t)\}. \quad (29)$$

where $K_v = 1000$ MW/p.u., q_c is the instantaneous interface limit in MW, and $\bar{q}$ is its rating in MW. The ideal transfer construction and the droop controller use the same source trajectory and power ceiling; they differ only in the selected transfer. The ideal interface permits power steps because it has no current-loop state.

5.2. Network Cases and Fair Regional Comparisons

The four-region network has injecting regions A and B and buffering regions C and D. The radial topology contains directed interfaces A–C, B–C and C–D; the meshed topology adds A–D and B–D. Each case spans four 5 ms intervals, with 1.6 MJ total injection and 2 MJ total upper headroom. A and B have equal upper headroom, selected from 0.1, 0.2, 0.3, and 0.4 MJ; C has 0.25 MJ, and D receives the remainder. All four regions use the same initial release allowance, either 0 or 0.1 MJ. Upper headroom and release allowance are specified independently in this controlled-interface model.

Uniform injection supplies 40 MW to each source throughout the window; delayed injection supplies zero during the first two intervals and 80 MW during the last two. Early availability is one during the first two intervals and zero afterward; late availability reverses that sequence; constant availability is 0.5 throughout. These schedules have equal integrated availability. Interface ratings are 40, 80, or 120 MW, with a factor of 0.8 for interfaces terminating at D. The sum of all source-originating transfers is also limited by a shared converter rating of 60 or 120 MW, multiplied by the same availability.

The design contains $2 \times 4 \times 3 \times 2 \times 3 \times 2 \times 2 = 576$ cases. Temporal optimization is compared with whole-network energy, independent-port fixed-region prefix bounds, shared-equipment fixed-region prefix bounds, and fixed-region bounds over every event-aligned interval. For the shared-equipment tests, regional outward transfer is maximized under the same shared limit as the temporal optimization. Any remaining difference therefore cannot be attributed to counting the converter rating separately at each port.

At a nonzero interval start, the fixed-region tests allow the regional state to reach its lower energy limit. With piecewise constant injection and capacity, the deficit is affine between event times, so all continuous-time extrema are covered by the event pairs. Temporal optimization additionally requires the states assumed by different intervals to belong to one trajectory.

The allocation objective assigns unit cost to each added MJ of upper headroom. Node sensitivities are checked by separate 0.00001 MJ perturbations. The finite-release example, shared-converter example, and radial case used for headroom sensitivity are subdivided by factors 1, 2, 4, 8, 16, and 32; their constant segment data should leave the ideal optimum unchanged. Solver timing uses directed ring-and-chord networks with 4, 8, 16, 32 or 64 nodes and 10, 40, or 100 intervals.

5.3. Three-Terminal Circuit with Passive Line Dynamics

Source node S connects to receiver R through one R–L branch and to passive buffer B through another. Each node has an independent shunt capacitor. Source and receiver converter currents have first-order dynamics and bounded references. This topology retains line-current establishment and magnetic-energy exchange during export interruption and restoration.

Kirchhoff's laws give the capacitor and branch equations:

$$\begin{cases} C_S\dot{V}_S = i_s - i_1 - i_2, & C_R\dot{V}_R = i_1 - i_r, & C_B\dot{V}_B = i_2, \\ L_1\dot{i}_1 = V_S - V_R - R_1 i_1, & L_2\dot{i}_2 = V_S - V_B - R_2 i_2. & \end{cases} \quad (30)$$

where voltages are in V, currents in A, capacitances in F, inductances in H, and resistances in ohms. Currents i_1 and i_2 are positive from S toward R and B and follow the branch differential equations. Converter limits act on the reference dynamics of i_s and i_r; they do not clip the line-current states.

The source power command follows the delayed ramp used in the ideal model. Actual injection is $V_S i_s$, with converter currents governed by the following:

$$\tau_s \dot{i}_s = sat_{[0,I_{s,\max}]}\left(\frac{p_{\text{ref}}}{V_S}\right) - i_s, \tau_r \dot{i}_r = i_r^\star - i_r. \tag{31}$$

where τ_s and τ_r are time constants in s, $I_{s,\max}$ is the source current limit in A, and p_{ref} is the power command in W. The implementation floors the voltage denominator at 0.3 p.u.; this floor is inactive on every reported trajectory. Setting $i_r^\star = 0$ during the interruption makes the receiving current decay continuously. This averaged port represents interrupted active-power transfer.

After restoration, the receiver uses a power feedforward term and local voltage droop:

$$i_r^\star = sat_{[0,I_{r,\max}]}\left(\frac{p_{\text{ref}}(t-d_c) - R_1\left[p_{\text{ref}}(t-d_c)/V_b\right]^2 + K_r(V_R - V_{R0})}{V_R}\right). \tag{32}$$

where $I_{r,\max}$ is the receiver current limit in A, V_b is the nominal voltage in V, V_{R0} is the pre-interruption receiver voltage, K_r has units W/V, and d_c is the delay with which the source curtailment command reaches the receiver, zero in the principal cases and varied in the sensitivity study.

The resistive voltage drop and equal source and receiver currents determine the initial operating point:

$$i_0 = P_0/V_b, V_{S0} = V_{B0} = V_b, V_{R0} = V_b - R_1 i_0, i_1(0) = i_s(0) = i_r(0) = i_0, i_2(0) = 0. \tag{33}$$

where P_0 is source power in W. Subtracting initial resistive loss in the receiver feedforward makes all seven physical derivatives zero at this point. A separate small-signal run over these seven states applies a 0.001 p.u. source-voltage perturbation without export interruption and records its decay.

Table 4 specifies the circuit and matched interventions. The 2 H base buffer-line inductance separates early receiver response from remote absorption. An additional short-line sensitivity sets $L_1 = 0.015$ H and $L_2 = 0.020$ H.

Table 4. Physical circuit parameters and principal matched interventions.

Quantity	Base Value	Intervention or Interpretation
Nominal DC voltage and upper limit	320 kV; 1.15 p.u.	Common voltage base at all nodes
Total capacitor headroom	3 MJ	Computed from actual initial voltages
Capacitor fractions $S/R/B$	0.15/0.10/0.75	Relocation: 0.75/0.10/0.15
Base capacitances $S/R/B$	27.2364/18.1576/136.1821 μF	Same total capacitance under S/B exchange
Branch resistance R_1,R_2	1/1 ohm	Lossless sensitivity sets both to zero
Branch inductance L_1,L_2	0.15/2 H	Buffer improvement: 0.15/0.15 H
Source initial/floor power	100/20 MW	Actual power is $V_S i_s$
Curtailment delay and rate	4 ms; 2000 MW/s	Rate sensitivity: 1000 and 4000 MW/s
Current response τ_s,τ_r	2/0.5 ms	Slow-source sensitivity: 8 ms
Converter limits $I_{s,\max}$,$I_{r,\max}$	360/500 A	Limit sensitivities: 315/330 A
Receiver voltage gain K_r	5000 W/V	Fixed for the principal comparisons
Interruption and observation	24 ms; an additional 100 ms	Voltage checked throughout both periods

Capacitances are scaled at their equilibrium voltages to give 3 MJ total upper headroom: 0.449728, 0.301633 and 2.248639 MJ at S, R and B. Exchanging S and B capacitances preserves the total because their initial voltages are equal; the corresponding lower voltage reserves and dynamic time scales change with the capacitances.

The energy ledger includes capacitor storage, line magnetic storage, and resistive dissipation:

$$\mathcal{E}(t) = \frac{1}{2}\sum\nolimits_{j\in\{S,R,B\}} C_j V_j^2(t) + \frac{1}{2}L_1 i_1^2(t) + \frac{1}{2}L_2 i_2^2(t), \dot{\mathcal{E}} = V_S i_s - V_R i_r - R_1 i_1^2 - R_2 i_2^2. \quad (34)$$

where $\mathcal{E}$ is in J. Source and export energies are integrated from actual port powers, while line magnetic energy and resistive loss are accounted for separately. Using the source command instead of $V_S i_s$ would violate this balance during finite current response.

5.4. Parameter Basis and Representative Ranges

The matched comparisons require values that can be varied one at a time across complete grids under one disturbance and one set of control permissions; manufacturer models and field records are tied to specific installations and controls. The grid is therefore built from constructed values, and Table 5 positions each of them against the quantity it represents, a normalized interpretation, and the range of installed equipment.

Table 5. Constructed parameters, their normalized meaning and representative ranges.

Quantity	Constructed Value	Normalized Meaning	Representative Range and Basis
DC voltage and voltage box	320 kV; 0.90–1.15 p.u.	Common base at every node	Symmetric-monopole voltage class; box fixed by the study
Node energies $S/R/B$	1.394/0.930/6.972 MJ	13.9/9.3/69.7 kJ/MVA at 100 MVA	MMC designs 30–40 kJ/MVA [6,14]; buffer above that range
Converter current limits	360/500 A	1.15/1.60 of the rated 312.5 A	Current limiting at 1.1–1.5 p.u. [17]
Current-loop time constants	2/0.5 ms	Closed-loop bandwidths 80/320 Hz	Inner current loops of a few hundred hertz [18]
Curtailment delay and rate	4 ms; 1000–4000 MW/s	10–40 p.u./s; floor reached in 24–84 ms	Communication-based ride-through power reduction [1,2]
Branch inductance L_1	0.15 H	Current build-up time $t_b = L_e i_0/(0.05V_b) = 2.9$ ms	Series reactors of 0.1 H at line ends in DC-grid protection designs [34,35]
Buffer-line inductance L_2	0.1, 0.5, 2, 5 H	$t_b = 2.0, 9.8, 39, 98$ ms against the 24 ms interruption	0.1–0.5 H reactor-terminated lines; 2–5 H stress values
Branch resistances	1 ohm	0.1 MW loss at rated current	Short reactor-terminated connections
Interruption	24 ms (12–48 ms)	One to two cycles	Protection clearing times

The build-up time t_b in Table 5 is the time a 0.05 p.u. voltage difference needs to establish the rated current. The 0.1 and 0.5 H members of the grid respond within the 24 ms interruption and form the representative subset of 30 cases, whose counts are reported with the complete grid; 2 and 5 H are stress values at which the buffer path responds more slowly than the interruption.

5.5. Necessary Voltage Bounds Retaining Passive-Network Dynamics

The voltage requirement is $[V_-, V_+] = [0.90, 1.15]V_b$ at every node, with $V_b = 320$ kV. The following bounds test whether a trajectory can remain in this box under the specified source-current and receiver-blocking restrictions.

Let $p_{\min}(t)$ denote the fastest admissible source-curtailment envelope. Admissible source controls satisfy $p_c(t) \geq p_{\min}(t)$, with current reference $sat_{[0,I_{s,\max}]}[p_c(t)/V_S(t)]$. The scalar comparison theorem yields the following:

$$\begin{cases} \tau_s \dot{\underline{i}}_s & = \min\{I_{s,\max}, p_{\min}/V_+\} - \underline{i}_s, \\ \tau_s \dot{\bar{i}}_s & = I_{s,\max} - \bar{i}_s, \\ \underline{i}_s(0) & = \bar{i}_s(0) = i_0, \\ \underline{i}_s(t) & \leq i_s(t) \leq \bar{i}_s(t). \end{cases} \tag{35}$$

where $\underline{i}_s$ and $\bar{i}_s$ are current envelopes in A. They cover all source controls satisfying the stated curtailment, current, and bandwidth limits. During receiver blocking, $i_r(t) = i_0 \exp(-t/\tau_r)$, giving a source-power lower bound $V_- \underline{i}_s$ and export-power upper bound $V_+ i_r$.

An elementary energy bound applies the voltage box separately to each branch. Define $\Delta V = V_+ - V_-$ and $i_e^{\pm}(t) = i_{e0} e^{-R_e t/L_e} \pm (\Delta V/R_e)\left(1 - e^{-R_e t/L_e}\right)$ for branch e. The corresponding absolute-current bound is $I_e(t) = \max(|i_e^+(t)|, |i_e^-(t)|)$, using the continuous limit if $R_e = 0$. Region SR contains its two capacitors and the complete internal inductor L_1. Conservation gives the lower bound on the energy retained in the region,

$$J_{SR}(t) = \int_0^t [V_- \underline{i}_s(\tau) - V_+ i_r(\tau)] \mathrm{d}\tau - \int_0^t [R_1 I_1^2(\tau) + V_+ \max(i_2^+(\tau), 0)] \mathrm{d}\tau, \tag{36}$$

where J_{SR} is in J. Every trajectory that remains inside the box then satisfies

$$J_{SR}(t) \leq H_S + H_R + \frac{L_1}{2}[I_1^2(t) - i_{10}^2], \tag{37}$$

where H_S and H_R are expressed in J. The right side includes capacitor headroom and the largest line magnetic-energy increase allowed by the current envelope. Its endpoint-voltage extremes need not occur simultaneously, which can make the bound loose.

To retain passive-network coupling, write the circuit in SI states $z = \left[V_S, V_R, V_B, i_1, i_2\right]^{\mathrm{T}}$:

$$\begin{cases} \dot{z} & = Az + b_s i_s - b_r i_r, \\ A & = \begin{bmatrix} 0 & 0 & 0 & -C_S^{-1} & -C_S^{-1} \\ 0 & 0 & 0 & C_R^{-1} & 0 \\ 0 & 0 & 0 & 0 & C_B^{-1} \\ L_1^{-1} & -L_1^{-1} & 0 & -R_1/L_1 & 0 \\ L_2^{-1} & 0 & -L_2^{-1} & 0 & -R_2/L_2 \end{bmatrix}, \\ b_s & = \left[C_S^{-1}, 0, 0, 0, 0\right]^{\mathrm{T}}, \\ b_r & = \left[0, C_R^{-1}, 0, 0, 0\right]^{\mathrm{T}}. \end{cases} \tag{38}$$

where voltage and current derivatives have units V/s and A/s; capacitances and inductances are in F and H. The known response is $z_f(t) = e^{At} z_0 - \int_0^t e^{A(t-\tau)} b_r i_r(\tau) \mathrm{d}\tau$. The impulse response $h_k(t) = e_k^{\mathrm{T}} e^{At} b_s$ selects node k through e_k and has units V/(A · s). Because an oscillatory R–L–C response can change sign, its positive part must multiply the lower

current envelope and its negative part the upper envelope to obtain a voltage lower bound. With $h_k^+ = \max(h_k, 0)$ and $h_k^- = \min(h_k, 0)$, variation of constants gives:

$$\begin{cases} V_k(t) & = e_k^{\mathrm{T}} z_f(t) \\ & \quad + \int_0^t h_k^+(t-\tau) i_s(\tau)\mathrm{d}\tau \\ & \quad + \int_0^t h_k^-(t-\tau) \bar{i}_s(\tau)\mathrm{d}\tau, \\ \overline{V}_k(t) & = e_k^{\mathrm{T}} z_f(t) \\ & \quad + \int_0^t h_k^+(t-\tau) \bar{i}_s(\tau)\mathrm{d}\tau \\ & \quad + \int_0^t h_k^-(t-\tau) i_s(\tau)\mathrm{d}\tau, \\ V_k(t) & \le V_k(t) \le \overline{V}_k(t). \end{cases} \tag{39}$$

Every trajectory remaining inside the box must consequently satisfy:

$$\begin{cases} V_k(t) & \le V_+, \overline{V}_k(t) \ge V_-, \\ k & \in \{S, R, B\}, 0 \le t \le T. \end{cases} \tag{40}$$

Violation excludes every box-safe trajectory in the declared control class. Passing remains unresolved because the interval calculation relaxes temporal and voltage-feedback correlations in the source current.

After restoration, the receiver reference may take any value in $[0, I_{r,\max}]$, so its current lies between the continued decay and the first-order rise toward the limit:

$$i_r(t) = i_r(T)e^{-(t-T)/\tau_r}, \bar{i}_r(t) = I_{r,\max} + [i_r(T) - I_{r,\max}]e^{-(t-T)/\tau_r}, t \ge T, \tag{41}$$

where $i_r(T) = i_0 e^{-T/\tau_r}$ is the blocked current at restoration in A. With $h_k^r(t) = e_k^{\mathrm{T}} e^{At} b_r$, the receiver enters Equation (39) as a second channel with the opposite sign (its positive part multiplies $\bar{i}_r$ in the lower bound), so one envelope covers the interruption and the recovery window; during blocking it coincides with the single-channel form.

The envelopes of Equation (35) use V_+ in the lower command and $I_{s,\max}$ as the upper command. Since every box-safe trajectory satisfies $\max\{V_S, V_-\} \le V_S \le \min\{\overline{V}_S, V_+\}$, the commands can use the voltage envelopes of the previous evaluation:

$$\tau_s \dot{i}_s = \min\left\{ I_{s,\max}, \frac{p_{\min}(t)}{\min\{\overline{V}_S(t), V_+\}} \right\} - i_s, \tau_s \dot{\bar{i}}_s = \min\left\{ I_{s,\max}, \frac{P_0}{\max\{V_S(t), V_-\}} \right\} - \bar{i}_s. \tag{42}$$

The first equation alone defines Class I, the declared class with iterated voltage coupling; the second adds the available-power ceiling $p_c(t) \le P_0$ of a renewable source and defines Class II. Every iterate is a necessary bound for its class, and the lower envelopes are nondecreasing; six iterations change the peaks by less than 10^{-7} p.u.

Equation (39) is evaluated twice. The floating-point survey uses trapezoidal convolution on a 5 µs grid and accepts an exclusion above V_+ plus 1 V. The validated evaluation encloses every quantity in midpoint–radius ball arithmetic [33] at 256-bit precision. In energy-normalized coordinates $\tilde{z} = P^{1/2} z$ with $P = \mathrm{diag}(C_S, C_R, C_B, L_1, L_2)$, the matrix $\tilde{A} = P^{1/2} A P^{-1/2}$ is skew-symmetric plus $-\mathrm{diag}(0,0,0,R_1/L_1, R_2/L_2)$, so $\|e^{\tilde{A}s}\|_2 \le 1$ and $\left|h_k^{(m)}(s)\right| \le \|\tilde{A}^m\|_F / \sqrt{C_k C_S}$. The matrix exponentials, the free response, and the closed-form current envelopes are balls, and the convolution uses the midpoint rule with the remainder $(\Delta^3/24)\sup|f''|$ where the integrand is twice differentiable and interval hulls where the kernel may change sign. An exclusion is validated when the lower end of the enclosure exceeds V_+.

5.6. Switching-Function MMC Model

The averaged source and receiver ports of Equation (31) are replaced by three-phase half-bridge MMCs [12,13] with N submodules per arm and individual capacitor voltages;

each station keeps its nominal stored energy through $C_{\mathrm{sm}} = NC_j/6$. The arms carry L_{arm} and R_{arm}, and the station terminal is an inductive node solved from Kirchhoff's current law over the three legs and the connected lines. Each station connects through L_g and R_g to a stiff three-phase source of peak phase voltage E_g. Its dq current loops are designed for the time constants τ_s and τ_r of the averaged ports, a synchronous-reference-frame phase-locked loop supplies the angle, and a proportional circulating-current suppression with gain K_{cc} acts on each leg. The modulator inserts submodules in sorted order until the accumulated voltage reaches the arm target, with the last submodule in pulse-width modulation; the DC component of the target scales with the measured arm-average voltage, so the terminal voltage follows the stored energy as in the averaged circuit. The receiving-grid fault is a dip of the onshore source to 5% for the interruption, during which the receiver's active-current reference is zero; after clearance, it follows Equation (32). Table 6 lists the parameters; the nominal energy mapping (node capacitances of Table 4) and an MMC-consistent mapping with 30 kJ/MVA at both stations are both evaluated. A sweep of the common station energy from 10 to 40 kJ/MVA in steps of 5 kJ/MVA repeats the 36 MMC-consistent cases at each energy.

Table 6. Parameters of the switching-function MMC model.

Quantity	Value	Basis
Submodules per arm N; nominal submodule voltage	20 (10 and 40 in the sensitivity); 16 kV	Switching-function model with individual capacitors
Submodule capacitance	Nominal mapping 90.79 µF at S, 60.53 µF at R; MMC-consistent mapping 195.3 µF at both	$NC_j/6$; 58.59 µF node equivalent at 30 kJ/MVA
Arm inductance and resistance	50 mH; 0.5 ohm	0.11 p.u. on the station base
Grid-side inductance, resistance and source voltage	80 mH; 0.3 ohm; $E_g = 110$ kV peak phase	Transformer equivalent; modulation index 0.8
Current loops, circulating-current gain, phase-locked loop	$\tau_s = 2$ ms, $\tau_r = 0.5$ ms; $K_{cc} = 10$ ohm; $k_p = 50\ \mathrm{s}^{-1}$, $k_i = 2000\ \mathrm{s}^{-2}$ per unit	Time constants of the averaged ports; lock within 0.1 s
Fault, integration and control	Source amplitude 5% for the interruption; 10 µs fourth-order steps; 100 µs control period	200 ms pre-run; 100 ms recovery window

5.7. Voltage Boundaries and Error Measures

For a fixed controller and interruption duration, the voltage margin covers the interruption and 100 ms after export restoration:

$$M_V(T) = V_{\mathrm{max,pu}} - \max_{j\in\{S,R,B\}} \max_{0\le t\le T+0.1\ \mathrm{s}} \frac{V_j(t;T)}{V_b}. \tag{43}$$

where $V_{\mathrm{max,pu}} = 1.15$ and M_V is in p.u. A negative margin denotes an upper-voltage violation under that controller, not control-independent infeasibility or a synchronization critical clearing time. No overvoltage trip is modeled; states after the first crossing are unprotected mathematical continuations.

For each of four principal configurations, 33 uniformly spaced durations over 1–80 ms locate the first sampled transition to voltage violation. Fifteen bisections refine its bracket to approximately 0.0000754 ms; durations are reported to three decimals. The coarse sequence is checked for monotonicity, but a narrower violating interval between earlier samples may remain undetected, so the result is the first detected and refined boundary rather than a globally established critical duration. Runs at 98% and 102% of that boundary check the adjacent classifications.

The 60-case circuit grid combines source-capacitance fractions of 0.10, 0.25, 0.45, 0.65, and 0.80; buffer-line inductances of 0.1, 0.5, 2, and 5 H; and curtailment rates of 1000, 2000, and 4000 MW/s. The receiver fraction remains 0.10, the buffer receives the remaining capacitance, and the interruption lasts 24 ms; all other settings follow Table 4. Each equilibrium is recomputed. For the four principal configurations, DOP853 and RK45 with maximum steps of 100, 40, and 20 microseconds give 24 convergence runs. The reference uses DOP853, relative tolerance 10^{-11} and a 5 microsecond maximum step. Voltage extrema include derivative roots, event times, and observation endpoints.

Table 7 specifies the integration, optimization, and cross-check settings.

Table 7. Numerical settings and independent verification measures.

Computation	Method and Settings	Verification
Two-region trajectories	Event-segmented RK45; relative tolerance 10^{-9}; maximum step $T/100$	77 adaptive and fixed-step comparisons; exact injection integral
Network feasibility	Sparse linear optimization using HiGHS 1.8.0 through SciPy 1.16.3; primal and dual tolerances 10^{-9}	Equality and bound residuals; independent interval-cut computation
Causal envelope	Piecewise-affine cumulative-energy evaluation	300 comparisons with independently formulated optimization
Capacity derivative	Dual multiplier sum over time at each node	Separate 10^{-5} MJ capacity perturbation
Circuit trajectories	DOP853; relative tolerance 2×10^{-9}; maximum step 20 microseconds	Complete energy balance and derivative-root peaks
Circuit convergence	DOP853 and RK45; 24 runs	Peak-voltage and classification agreement
Duration refinement	33-point initialization and 15 bisections	Safe and violating cases on both sides of each boundary
Validated bound enclosure	Arb ball arithmetic at 256 bits through Python-FLINT 0.9.0; 5 μs midpoint rule with rigorous remainder	Enclosure radius and exclusion margin of every case
Switching-function MMC model	Fixed-step fourth-order integration at 10 μs; 100 μs control period; 200 ms pre-run	Energy ledger residual; pre-fault power and ripple
Randomized robustness	1200 circuits, 600 perturbed principal and 3000 two-region samples; random-number generator initial state 20261105	Necessary bound and tested trajectory for every sample

The software versions are Python 3.11, NumPy 2.3.5, SciPy 1.16.3 with the bundled HiGHS 1.8.0, and python-flint 0.9.0; the processor is an Intel Core Ultra 7 258V (Intel, Santa Clara, CA, USA). Reported optimization times include only the single-process solver call, excluding model assembly, plotting, and file output.

6. Results and Engineering Interpretation

Matched two-region cases separate the effect of headroom location from total capacity. The network comparisons then distinguish shared-device constraints from transfer chronology, and the R–L circuit quantifies the additional restrictions imposed by current dynamics and recovery. The section also quantifies the realizability of the ideal transfer with a finite-bandwidth controller, translates headroom additions into capacitance, reports the effect of

omitted dynamics on the classification, validates and extends the voltage bounds, compares the averaged circuit with the MMC model, and closes with the randomized studies.

6.1. Equal Total Headroom Does Not Give Equal Voltage Performance

Figure 2 compares three cases with the same 3 MJ total headroom, 100 MW initial source power, 4 ms curtailment delay, 2000 MW/s reduction rate, and 24 ms interruption. The voltage trajectories therefore correspond to the same 2 MJ injected energy and 1 MJ aggregate absorption margin.

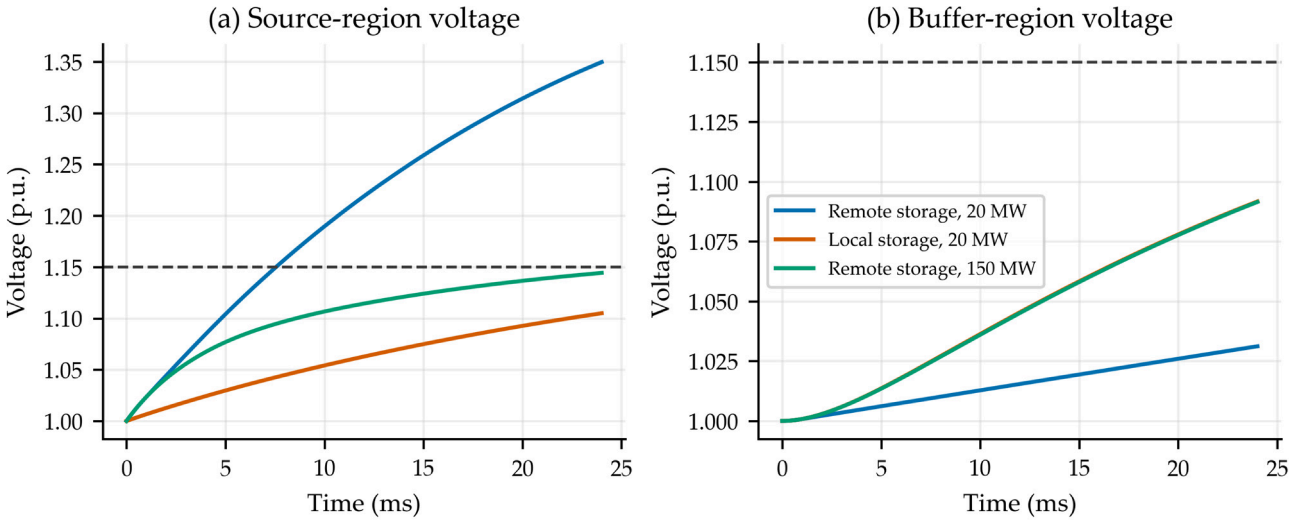


Figure 2. Matched two-region trajectories with identical total headroom and injected energy: (**a**) source-region voltage; (**b**) buffer-region voltage. The grey dashed line marks the 1.15 p.u. upper voltage limit.

Table 8 compares the voltage peaks with the required local headroom. A 20 MW interface transfers at most 0.48 MJ in 24 ms, leaving at least 1.52 MJ at the source. With only 0.6 MJ local headroom, the resulting 0.92 MJ deficit excludes every transfer satisfying that interface ceiling and source trajectory, independently of the droop gain.

Table 8. Matched configurations with 3 MJ total headroom and 2 MJ disturbance injection.

Source/Buffer Headroom (MJ)	Interface Rating (MW)	Required Source Headroom (MJ)	Source Voltage Peak (p.u.)	Interpretation
0.6/2.4	20	1.52	1.349598	Ideal energy infeasibility
2.4/0.6	20	1.52	1.104885	Tested droop remains within limit
0.6/2.4	150	0	1.144254	Tested droop remains within limit

At 150 MW, the ideal port can transfer the full incoming power without accumulating local energy. Fixed droop still requires a voltage difference to initiate transfer, which explains its positive voltage excursion despite zero ideal local-headroom requirement.

In the 2592-case study, 670 of 2160 aggregate-capacity passes violate the voltage limit under fixed droop (Table 9). The regional condition excludes 958 cases, none of which have a safe droop trajectory. For the remaining 1634 cases, Equation (15) supplies an admissible ideal transfer; fixed droop violates the limit in 144 of them. Those 144 cases therefore identify a controller-realization gap within the same source and interface permissions.

Table 9. Complete two-region classification and whole-window checks.

Measure	Number of Cases	Meaning
Complete parameter design	2592	All prescribed combinations retained
Aggregate headroom passes	2160	Total energy alone does not locate the bottleneck
Aggregate passes with droop voltage violation	670	Spatial and control limitations remain
Ideal energy infeasible	958	Exact boundary fails for this ideal source/interface class
Ideal infeasible but droop safe	0	No contradiction of the exclusion condition
Ideal feasible but droop violates	144	Explicit ideal trajectory establishes a realization gap
Ideal feasible and droop safe	1490	Both resource and tested-control checks pass
Clearing-only bound misses an earlier deficit	16	Full-window evaluation is necessary
Earlier voltage violation disappears by clearing	22	A later acceptable state does not erase the earlier crossing

Case E0190 has source/buffer headrooms of 0.3/2.7 MJ, an 80 MW interface, zero curtailment delay, a 1000 MW/s reduction rate, and a 24 ms interruption. The ideal construction requires 0.2 MJ locally and reaches 1.102270 p.u.; fixed droop reaches 1.170297 p.u. Figure 3 compares these trajectories and identifies all 144 such cases in the full scan.

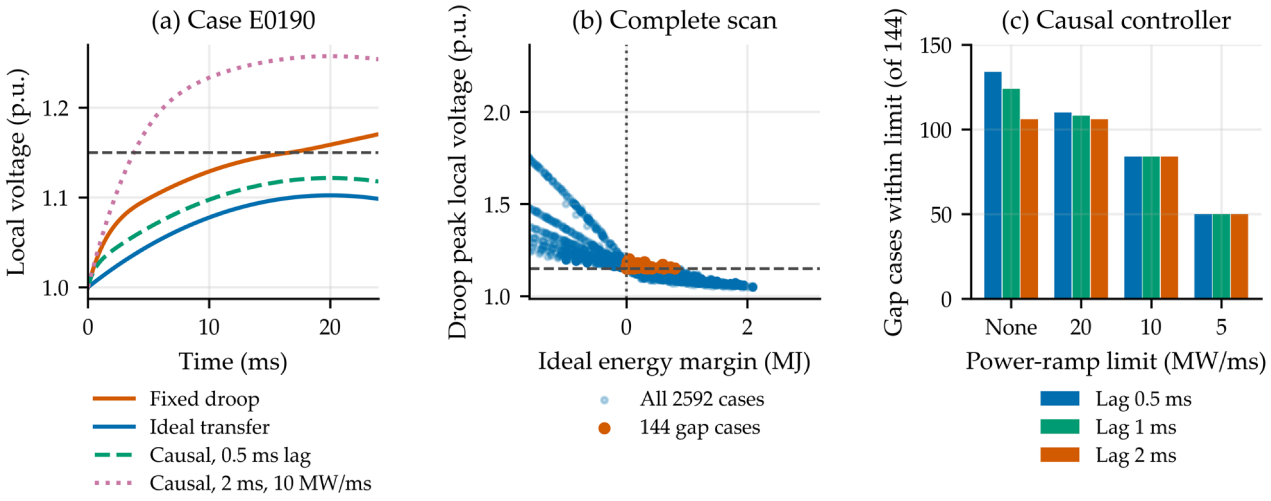


Figure 3. Ideal feasibility and controller realization: (**a**) source voltage in case E0190 under the ideal transfer (blue), fixed droop (orange), the causal controller with a 0.5 ms lag (green dashed) and with a 2 ms lag and 10 MW/ms ramp limit (purple dotted); (**b**) fixed-droop voltage peaks versus ideal energy margin in the complete scan, with the 144 gap cases in orange; (**c**) gap cases kept within the limit by the causal controller for three lags and four power-ramp limits. In (**a**) and (**b**), the grey dashed line marks the 1.15 p.u. upper voltage limit. The grey dotted line in (**b**) marks zero ideal energy margin.

The causal controller of Equation (17) tracks the constructed transfer with the same rating and current derating as the droop law. Table 10 and Figure 3c report its outcome for the 144 gap cases and all 1634 ideal-feasible cases. With a 0.5 ms lag and no ramp restriction, 134 gap cases and 1624 feasible cases remain within the local limit, and the penalty bound guarantees 118 and 1606 of them before any simulation. The local energy never exceeds the ideal trajectory by more than the bound (largest excess 0.050 MJ against 0.075 MJ), and the predictive stop holds the remote energy within 4 J of its limit in every run. Power-ramp

limits govern the remainder: at 20, 10 and 5 MW/ms the gap counts are 106–110, 84 and 50 for the three lags, because the step to the rating withholds $\bar{\eta}^2/(2r_q)$ of transfer. In case E0190, the 0.5 ms controller peaks at 1.121606 p.u. with an excess equal to its 0.040 MJ bound, and the 2 ms, 10 MW/ms controller at 1.257164 p.u. (Figure 3a). The realization gap of the fixed droop law is therefore closed by a causal controller whose power-ramp capability reaches the rating within a few milliseconds; Equation (17) gives the margin a slower converter needs.

Table 10. Causal finite-bandwidth controller against the ideal construction: cases kept within the local voltage limit.

Lag τ_q (ms)	Ramp Limit r_q (MW/ms)	Gap Cases Within Limit (of 144)	Gap Cases Guaranteed by Equation (17)	Feasible Cases Within Limit (of 1634)	Largest Excess Over the Ideal Trajectory (MJ)
0.5	None	134	118	1624	0.050
0.5	10	84	15	1510	0.498
1	None	124	89	1614	0.100
1	10	84	15	1509	0.500
2	None	106	51	1590	0.197
2	10	84	5	1507	0.507

6.2. Chronology Creates Restrictions Beyond Integrated Capability

The source injects zero during the first 10 ms and 100 MW during the next 10 ms, whereas the interface is available at 100 MW only during the first interval. Local and remote upper headrooms are 0.2 and 1 MJ. With 0.3 MJ initially releasable at the source, at most 0.3 MJ can be transferred before the interface closes, leaving at least 0.7 MJ locally. Prefix conditions count the earlier 1 MJ transfer opportunity and pass; the causal condition instead requires 0.5 MJ additional local absorption.

Figure 4 compares the prefix estimate with the attainable minimum local energy. Increasing the early interface rating cannot raise pre-transfer above the 0.3 MJ release allowance because source injection begins only after the interface closes.

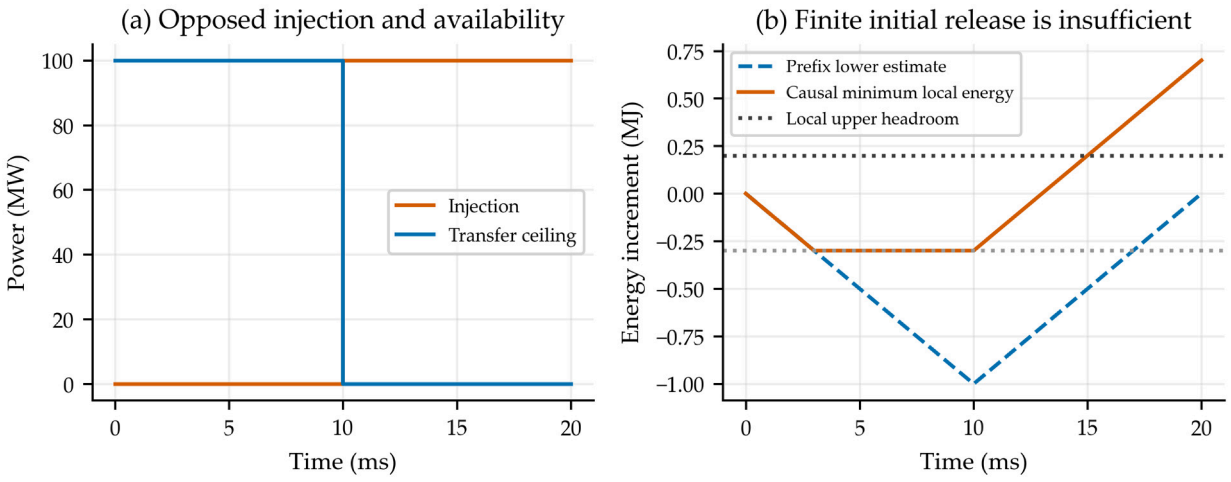


Figure 4. A channel available before source energy arrives: (**a**) injection and transfer ceiling; (**b**) prefix estimate and causal minimum local energy with finite initial release. In (**b**), the dark dotted line marks the 0.2 MJ local upper headroom, and the light grey dotted line marks the lower energy limit of −0.3 MJ set by the initial release.

For a capacitor with 0.2 MJ upper headroom and a 0.9–1.15 p.u. voltage range about a 1 p.u. initial voltage, the initial energy is 0.620155 MJ and the releasable energy is 0.117829 MJ.

The same timing example then requires 0.882171 MJ local headroom, or a 0.682171 MJ addition; independent optimization gives the same value. Across 300 piecewise-varying source and interface cases, the maximum discrepancy between the causal envelope and independent optimization is 6.66×10^{-16} MJ.

The directed chain *A*–*B*–*C* has upper headrooms of 1, 0.1 and 1 MJ and releasable energy of 0.05 MJ at each node. *A* receives 100 MW in both 10 ms intervals. Only *B*–*C* is available at 100 MW in the first interval, and only *A*–*B* is available at 100 MW in the second. Every fixed-region condition over every interval passes, with maximum deficit zero. However, *A* retains its first 1 MJ, while *B* can release only 0.05 MJ before receiving energy. *B* can therefore absorb at most 0.15 MJ in the second interval, forcing *A* to end with at least 1.85 MJ against its 1 MJ limit.

Temporal optimization requires 0.85 MJ additional upper headroom. Selecting *A* in the first interval and *A* with *B* in the second in Equation (26) yields the same deficit: 2 MJ of injection exceeds 1.15 MJ of compatible absorption and release margins. The fixed-region tests miss this restriction because their independently favorable intermediate states cannot belong to one continuous energy trajectory.

The shared-converter example in Table 11 has two sources, each injecting 0.6 MJ over 10 ms with 0.2 MJ local headroom and a 0.5 MJ individual transfer limit. A shared 50 MW converter limits combined transfer to 0.5 MJ, leaving a 0.3 MJ aggregate local deficit that independent-port accounting misses.

Table 11. Mechanism examples with independently checked analytical and optimization results.

Mechanism	Incomplete Condition that Passes	Minimum Added Upper Headroom (MJ)	Missing Restriction
Late injection; zero initial release	Fixed-region prefix conditions	0.800000	Early unused capacity cannot move future energy
Late injection; 0.3 MJ release	Fixed-region prefix conditions	0.500000	Finite initial discharge allowance
Late injection; capacitor-derived release	Fixed-region prefix conditions	0.682171	Lower voltage limit restricts pre-transfer
Two ports sharing one converter	Independent-port conditions	0.300000	Simultaneous converter rating
Reversed serial interface availability	Every fixed-region interval condition	0.850000	A single compatible state history

6.3. Complete Network Comparison and Capacity Allocation

All 576 network cases pass aggregate accounting because 2 MJ headroom exceeds 1.6 MJ injection (Table 12). Independent-port prefix tests exclude 256 cases; imposing the shared converter increases that count to 436. Temporal optimization excludes another 59 of the 140 shared-prefix passes, isolating restrictions due to energy history after shared-device capacity has been accounted for.

Every regional exclusion is also infeasible in temporal optimization. The all-interval fixed-region test and the temporal formulation give identical classifications for these 576 cases. Their agreement on this grid does not extend to the reversed-serial example, whose 0.85 MJ deficit requires a changing region.

Figure 5 shows the required additions for delayed injection, 0.1 MJ initial release per node, and a 120 MW shared rating. Increasing nominal port ratings reduces the deficit only when their availability overlaps the arriving surplus. In the early-availability panels,

additional rating provides no reduction; in the late-availability panels, the direct source–*D* interfaces reduce the meshed-network requirement relative to the radial case.

Table 12. Network comparison using the same injection and equipment permissions.

Assessment	Cases Excluded	Cases Passing	Passing Cases Infeasible Under Temporal Optimization
Aggregate energy	0	576	495
Independent-port regional prefix bounds	256	320	239
Shared-equipment regional prefix bounds	436	140	59
Shared-equipment fixed-region interval bounds	495	81	0
Temporal energy feasibility	495	81	0

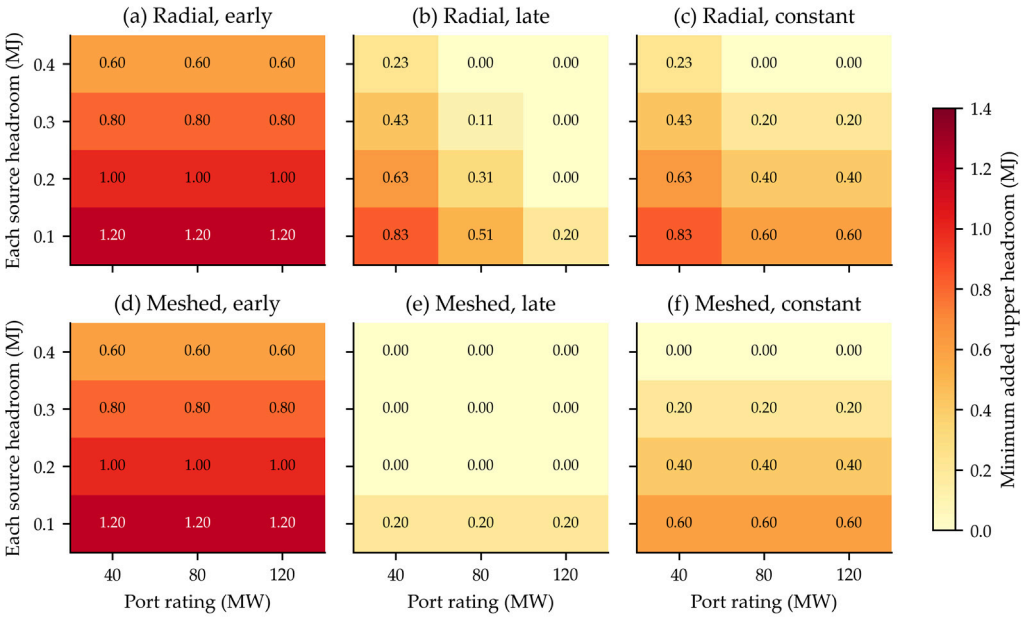


Figure 5. Minimum supplemental upper headroom across topology, resource location and availability timing: (**a**–**c**) radial networks; (**d**–**f**) meshed networks. All panels use delayed injection and identical total installed headroom.

The largest required addition is 1.4 MJ. Removing the shared-device constraint changes the optimum in 255 cases. These additions increase upper absorption headroom at fixed release limits; Equation (24) converts them to capacitance with both energy margins.

For the delayed-injection, early-availability radial case with 0.2 MJ headroom at each source, 80 MW nominal interfaces, a 60 MW shared limit and 0.1 MJ release per node, the headroom prices at *A*, *B*, *C* and *D* are 1, 1, 0 and 0. A 0.00001 MJ addition at either source reduces the optimum by 0.00001 MJ to numerical precision; the same addition at *C* or *D* gives no reduction (Figure 6).

The channel sensitivity case uses uniform injection, constant availability, 0.1 MJ headroom per source, and 40 MW nominal interfaces. Increasing the actual *C*–*D* ceiling by 0.001 MW in every interval reduces the required addition by 0.00002 MJ, giving −0.020000 MJ/MW, equal to the dual derivative. This perturbation has unit availability factor and is not a derivative with respect to the original nominal rating scale. Both source–*C* derivatives are zero because the *C*–*D* constraint remains binding.

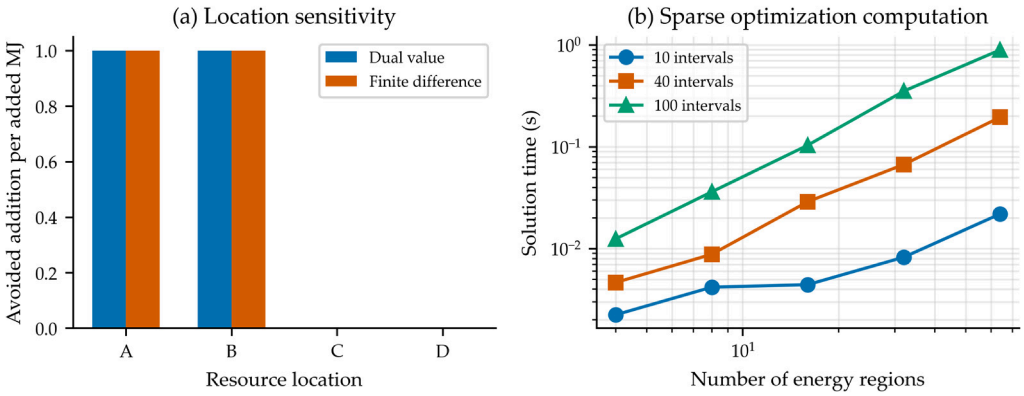


Figure 6. Configuration sensitivity and computation: (**a**) node headroom prices and finite differences; (**b**) solver time versus network size for three temporal resolutions.

Subdivision of constant-power intervals by factors up to 32 changes the tested optima only by numerical roundoff. Across the 576 cases, the largest balance residual is 2.78×10^{-16} MJ and the largest primal–dual objective gap is 3.33×10^{-16} MJ. Median solver time is 1.31 ms. A 64-region, 128-interface, 100-interval case has 19,264 variables and 6400 balance equalities and solves in 0.89 s on the stated processor.

6.4. Coupled Capacitance Sizing and Equipment Mapping

Equation (24) re-solves the 576 cases with the release margin coupled to the addition. In the 288 cases with 0.1 MJ initial release, the coupled addition is smaller than the uncoupled one in 95 cases, by 0.066 MJ on average and at most 0.6 MJ; the cases without release are unchanged because $\gamma_i = 0$, and the classification of all 576 cases is unchanged. For the capacitor-derived example of Table 11, the addition is 0.429268 MJ instead of 0.682171 MJ, in agreement with the closed form $(1 - R_A - H_A)/(1 + \gamma)$, corresponding to 26.00 µF instead of 41.31 µF added to an initial 12.11 µF at 320 kV.

Table 13 translates the remote-buffer circuit case, whose peak of 1.205091 p.u. occurs at node R, into equipment, with the circuit re-simulated for each candidate. Added capacitance of 29.83 µF at R brings the whole-window peak to 1.15 p.u. The same result at S needs 122.49 µF, and no addition at B up to 1500 µF achieves it, because the early peak at R is fed through L_1 before the buffer path responds. A braking resistor conducting above 1.05 p.u. needs 35.8 MW at R or 40.3 MW at S. With unit costs c_C per microfarad and c_P per megawatt, the resistor at R is the cheaper measure when $c_P/c_C < 0.832$ µF/MW; the location result holds for both technologies.

Table 13. Equipment that brings the remote-buffer circuit within the 1.15 p.u. limit over the whole window.

Measure	Node	Size	Added Upper Headroom (MJ)	Added Release (MJ)	Whole-Window Peak (p.u.)
Added capacitance	R	29.83 µF (1.64 times base)	0.495	0.287	1.150000
Added capacitance	S	122.49 µF (4.50 times base)	2.023	1.192	1.150000
Added capacitance	B	1500 µF	24.77	14.59	1.204767
Braking resistor above 1.05 p.u.	R	35.8 MW	0.68 dissipated	—	1.150000
Braking resistor above 1.05 p.u.	S	40.3 MW	0.77 dissipated	—	1.150000

6.5. Omitted Dynamics and the Feasibility Classification

Condition (N1) of Section 4 is tested on the exclusions. With outer allowances for line magnetic energy and dissipation at 500 A, 951 of the 958 two-region exclusions survive for 0.15 H and 1 ohm, 707 for 2 H and 1 ohm, and 694 for 2 H and 5 ohm; the median deficit is 0.654 MJ. A ramp-limited interface class with capability $\int_0^t \min\{\overline{q}, r_q \tau\} \mathrm{d}\tau$ raises the exclusions to 993, 1080, 1181 and 1409 at 20, 10, 5 and 2 MW/ms, and Equation (22) raises the 495 network exclusions to 496, 506 and 540 at 20, 10 and 5 MW/ms. The compulsory-load example withdraws 40 MW at B over two 10 ms intervals, with 0.3 MJ release and a 100 MW interface in the first interval only. Fixing the source at its zero minimum leaves a 0.5 MJ shortage and an infeasible relaxation, whereas every first-interval injection of at least 50 MW is feasible and the optimization with injection as a decision variable finds it.

6.6. Passive-Line Dynamics Change the Configuration Decision

Figure 7 displays the first 80 ms of the three-terminal voltage trajectories; reported maxima are evaluated through 100 ms after restoration. Export resumes at 24 ms, and each capacitor node is checked against the same 1.15 p.u. upper limit.

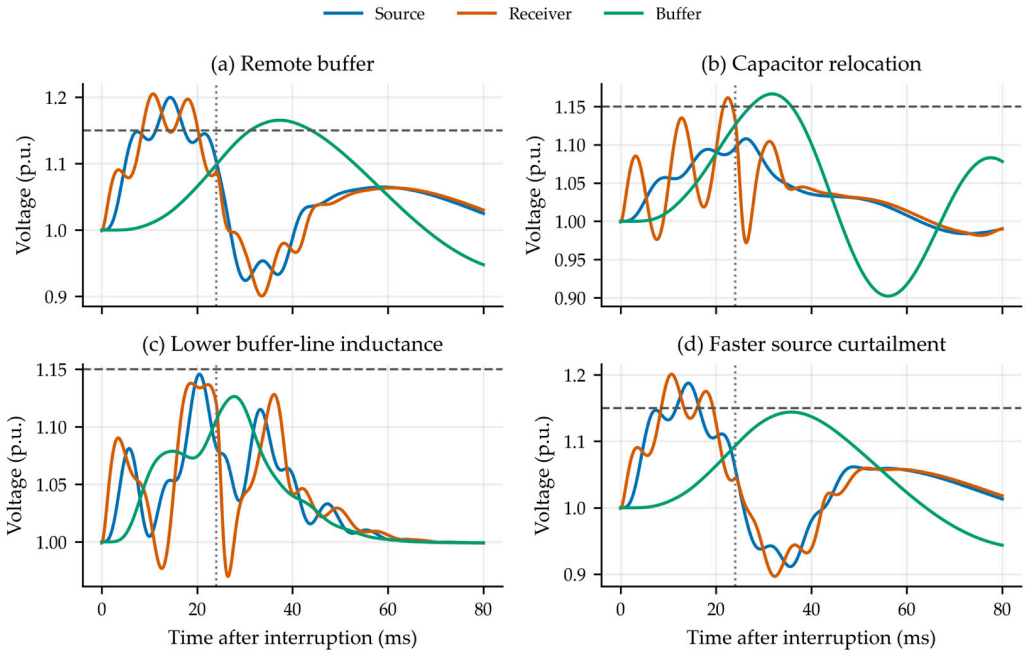


Figure 7. Three-terminal circuit voltages: (**a**) remote buffer; (**b**) relocated capacitance; (**c**) lower buffer-line inductance; (**d**) faster source curtailment. The grey dashed line marks the 1.15 p.u. upper voltage limit. The dotted vertical line marks export restoration at 24 ms.

With 3 MJ total headroom, the remote-buffer circuit reaches 1.205091 p.u. Exchanging source and buffer capacitances reduces the maximum to 1.166424 p.u., still above 1.15 p.u.; the maximum now occurs at the buffer after export restoration. The source improvement therefore leaves a different node constrained. Reducing the buffer-line inductance from 2 to 0.15 H instead gives a maximum of 1.145663 p.u. over the evaluated window.

Table 14 compares these peaks and detected duration boundaries. Doubling the curtailment rate reduces fault-window net injection from 2.048681 to 1.718055 MJ, but reduces the early oscillatory peak only from 1.205091 to 1.201100 p.u. The detected duration boundary changes from 8.327 to 8.354 ms. The early peak occurs before the faster ramp removes much energy, whereas its effect on the integrated surplus accumulates later.

Table 14. Principal physical-circuit results and controller-specific duration limits.

Configuration	Net Injection During 24 ms (MJ)	Full-Window Peak (p.u.)	First Detected Boundary (ms)	24 ms Result
Remote buffer	2.048681	1.205091	8.327	Voltage violation
Source/buffer capacitor exchange	2.037656	1.166424	20.390	Voltage violation
Buffer-line inductance 0.15 H	2.040344	1.145663	31.797	Within upper limit
Curtailment rate 4000 MW/s	1.718055	1.201100	8.354	Voltage violation
Source current time constant 8 ms	2.271525	1.218010	Not evaluated	Voltage violation
Zero line resistance	2.048690	1.206410	Not evaluated	Voltage violation
Line inductances 0.015/0.020 H	2.038628	1.116843	Not evaluated	Within upper limit

Of 60 circuit cases, 20 remain below the voltage ceiling throughout the fault, and 19 also remain below it during recovery; one fault-window pass therefore develops a recovery-period violation. In the representative subset with 0.1 and 0.5 H buffer lines, 18 of 30 cases remain below the ceiling during the fault; in the 2 and 5 H subset, 2 of 30. Cases at 98% and 102% of each detected duration boundary have opposite classifications.

All ten nominal sensitivity configurations have zero operating-point residual at reported precision and physical-state eigenvalues with negative real parts. The remote-buffer circuit has a least-damped pair of $-16.7014 \pm j58.4913$ per s. Its 0.001 p.u. source-voltage perturbation decays to a maximum voltage deviation of 2.19×10^{-8} p.u. in the small-signal check, establishing that the interruption is applied to a locally stable equilibrium.

Figure 8 compares duration-dependent peaks and the energy balance. Total capacitor and magnetic-energy change equals supplied energy minus exported energy and resistive loss. The maximum balance residual across 60 cases is 1.75×10^{-6} J. In 24 convergence runs, the largest peak-voltage change is 1.47×10^{-9} p.u., with unchanged classifications.

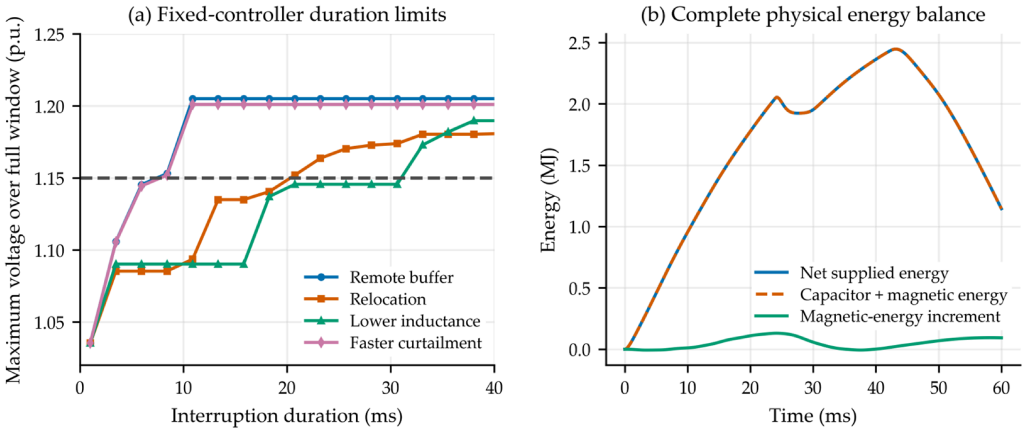


Figure 8. Circuit duration limits and conservation: (**a**) full-window peak voltage versus interruption duration; (**b**) net supplied energy and total physical stored-energy change. The grey dashed line in (**a**) marks the 1.15 p.u. upper voltage limit.

Receiver feedforward delay (0–10 ms), source current-limit logic (reference saturation, voltage-dependent derating, constant-power limit), interruption (20, 24, 28 ms), and a restoration ramp of the receiver reference (0, 2, 5 ms) give 108 combinations per principal configuration. The blocking-window peaks of the remote-buffer, lower-inductance, and faster-curtailment configurations are identical in all 108 runs, because they occur while the

receiver is blocked and the source current stays below every limit variant. The relocated configuration, whose maximum occurs at the buffer after restoration, varies between 1.1474 and 1.1736 p.u. On the 60-case grid, a 5 ms delay and voltage-dependent derating change no classification and contradict no exclusion.

6.7. Necessary Voltage Bounds in the Physical Circuit

Equation (36) leaves all ten representative cases unresolved because the independent endpoint-voltage extremes allow excessive internal current and magnetic headroom. Equation (39) retains the coupled R–L–C response instead of assigning those extremes independently.

Figure 9 plots two voltage-envelope cases over the complete window, the tightened classes of a third case, and the full 60-case comparison. The maximum lower-envelope voltages of the declared class for remote capacitance, relocated capacitance, lower buffer-line inductance, and faster curtailment are 1.197073, 1.153716, 1.108560, and 1.192996 p.u., respectively. The first, second, and fourth exclude box-safe operation throughout the interruption. The lower-inductance case remains unresolved by the bound; its tested controller keeps the trajectory below the upper limit.

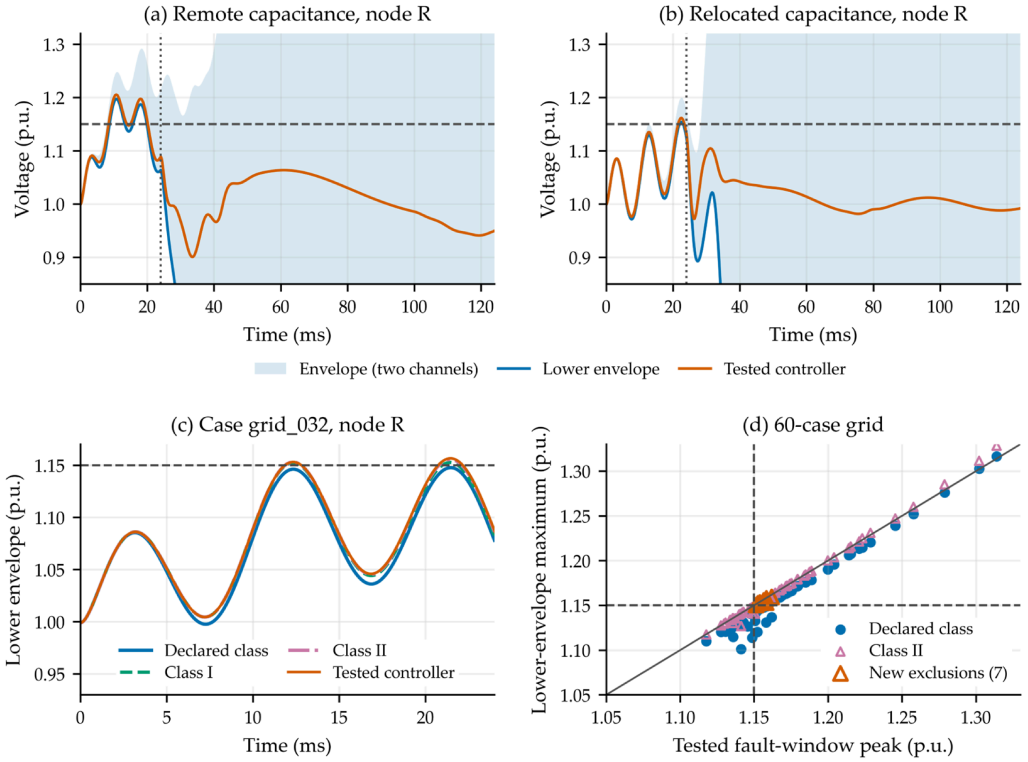


Figure 9. Voltage envelopes of the physical circuit: (**a**) receiving-node envelope of the remote-capacitance case over the interruption and the recovery window (blue band and line) with the tested trajectory (orange); (**b**) the same after capacitance relocation; (**c**) lower envelopes of case grid_032 for the declared class (blue), Class I (green dashed) and Class II (purple dash-dotted) with the tested trajectory (orange); (**d**) lower-envelope maxima of the declared class (blue circles) and Class II (purple triangles) against the tested fault-window peaks of all 60 cases, with the seven Class II exclusions in orange. The dotted vertical line marks restoration. The grey dashed lines mark the 1.15 p.u. upper voltage limit. The solid grey line in (**d**) is the line of equality.

The bound excludes 30 of the 60 predefined cases, all of which exhibit fault-window overvoltage under the tested controller. Of the 30 unresolved cases, 20 remain inside the box, and 10 exceed its upper boundary. No box-safe tested trajectory contradicts an exclusion.

The lowest fault-window voltage is 0.946042 p.u.; thus, the 0.90–1.15 p.u. box and upper-limit-only tests give identical classifications on this grid. In the representative subset with 0.1 and 0.5 H buffer lines, the bound excludes 4 of 30 cases and leaves 8 tested violations unresolved; in the 2 and 5 H subset, it excludes 26 of 30 and leaves 2. These counts describe prescribed parameter combinations, not disturbance probabilities.

The validated evaluation confirms all 30 exclusions: the lower end of the enclosure exceeds 1.15 p.u. in each, with an enclosure radius of at most 5.47 V against a smallest exclusion margin of 1275.7 V (case grid_055); the enclosure midpoints differ from the survey values by at most 2.2×10^{-5} p.u., and no unresolved case becomes excluded.

The two-channel envelope of Equation (41) covers the 124 ms window. After restoration, the admissible receiver current reaches 500 A within a few time constants, and the lower envelope falls below the limit in every case; the latest exclusion time is 41.3 ms, in cases already excluded during blocking. The 30 fault-window exclusions are therefore the full-window exclusions. The one trajectory that violates only after restoration, grid_036 with 1.1541 p.u. at the buffer, is not excluded, because an admissible receiver current removes that surplus (lower envelope 1.1224 p.u. in the recovery window).

Class I raises the lower-envelope maxima by up to 0.0136 p.u. and excludes 34 cases; Class II raises them by up to 0.0281 p.u. and excludes 37. Of the 10 tested violations left unresolved by the declared class, 6 remain after Class I, and 3 after Class II, and no iterate contradicts a box-safe trajectory. The seven new Class II exclusions have tested peaks between 1.1513 and 1.1619 p.u. For the four principal configurations, the Class II maxima are 1.205042, 1.161061, 1.135912 and 1.201016 p.u.; the lower-inductance configuration remains unresolved. In the representative subset, the two classes exclude 6 and 9 of 30 cases and leave 6 and 3 tested violations unresolved; in the 2 and 5 H subset, both exclude 28 of 30 and leave none.

6.8. Switching-Function MMC Model Against the Averaged Circuit

Figure 10 compares the receiving-node voltage of the averaged circuit with the terminal voltage and the station-average submodule voltage of the MMC model. With the nominal energy mapping, the remote-buffer case peaks at 1.2006 p.u. at the terminal and 1.2097 p.u. in the submodule average, against 1.205091 p.u., and the four principal configurations agree within 0.020 p.u. (Table 15). Across the 60-case grid, the submodule-average peaks differ from the averaged peaks by 0.0104 p.u. on average and 0.036 p.u. at most, with 50 of 60 classifications in agreement; the terminal peaks are lower by 0.0135 p.u. on average because the arm inductance lies between the submodule energy and the terminal. Measured at the terminal, the peaks differ by 0.015 p.u. on average and 0.040 p.u. at most, and 41 of 60 classifications agree; in all 19 differences, the averaged circuit reports a violation where the MMC terminal stays within the limit. Of the 30 exclusions, 10 meet MMC terminal peaks of 1.136–1.149 p.u.; for 2 of them, grid_048 and grid_049, the submodule-average peaks of 1.1497 and 1.1478 p.u. also lie below the limit. The receiving station of this mapping holds 9.3 kJ/MVA and its submodule voltages swing between 0.27 and 1.73 p.u., so the comparison tests the averaged port below the operating range of an MMC, where the exclusions remain statements about the averaged port.

With 30 kJ/MVA at both stations, the submodule voltages stay within 0.85–1.16 p.u.; all 36 classifications agree; the peaks differ by 0.0039 p.u. on average and 0.0077 p.u. at most, and the three exclusions of the bound meet no safe MMC trajectory. A sweep of the common station energy over the same 36 cases gives largest terminal differences of 0.025, 0.015, 0.012, 0.011, 0.0077, 0.0077 and 0.013 p.u. at 10, 15, 20, 25, 30, 35 and 40 kJ/MVA; 250 of its 252 classifications agree and all 66 exclusions meet MMC violations. Submodule counts of 10, 20 and 40 change the peaks by less than 0.002 p.u. The averaged port therefore represents

stations with common stored energies of 15–40 kJ/MVA within 0.015 p.u., and a case whose averaged peak lies within this distance of the limit is classified with the MMC model.

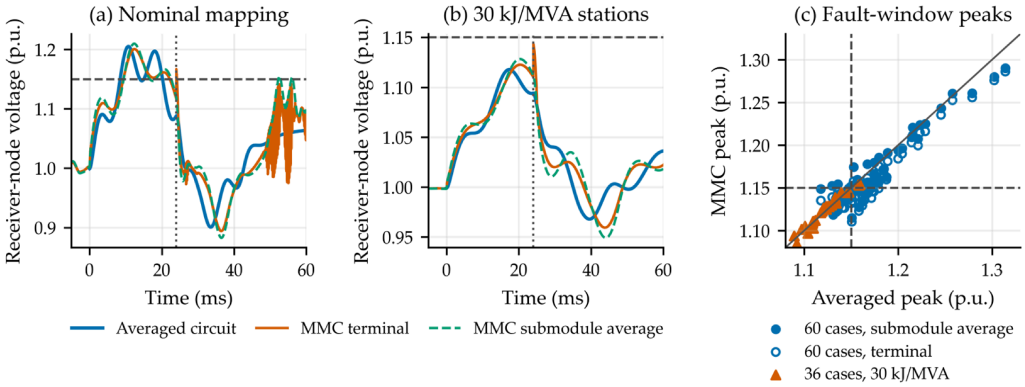


Figure 10. Averaged circuit against the switching-function MMC model: (**a**) receiving-node voltage of the remote-capacitance case with the nominal energy mapping, averaged circuit (blue), MMC terminal (orange) and station-average submodule voltage (green dashed); (**b**) the same configuration with 30 kJ/MVA at both stations; (**c**) MMC fault-window peaks against the averaged peaks for the 60-case grid (submodule average, filled blue; terminal, open blue) and the 36 MMC-consistent cases (orange triangles). The dotted vertical line marks restoration. The grey dashed lines mark the 1.15 p.u. upper voltage limit. The solid grey line in (**c**) is the line of equality.

Table 15. Fault-window peaks of the averaged circuit and the switching-function MMC model.

Configuration	Averaged Circuit (p.u.)	MMC Terminal (p.u.)	MMC Submodule Average (p.u.)	MMC Highest Submodule (p.u.)
Remote buffer, nominal mapping	1.205091	1.2006	1.2097	1.588
Capacitor relocation, nominal mapping	1.161376	1.1419	1.1550	1.538
Buffer-line inductance 0.15 H, nominal mapping	1.145663	1.1437	1.1493	1.502
Curtailment rate 4000 MW/s, nominal mapping	1.201100	1.1958	1.2053	1.584
Remote buffer, 30 kJ/MVA stations	1.118836	1.1227	1.1283	1.206

6.9. Randomized Parameter Studies

Table 16 summarizes the randomized studies. The 1200 circuit samples draw the operating voltage from 0.98 to 1.02 p.u., the interruption from 16 to 32 ms, the source-capacitance fraction from 0.10 to 0.80, the buffer-line inductance from 0.1 to 5 H, L_1 within ±20%, both resistances from 0.5 to 1.5 ohm, the source power from 80 to 120 MW, a signed buffer injection from −20 to 20 MW, a receiver load from 0 to 20 MW, the curtailment delay from 2 to 8 ms, and the source time constant from 1 to 4 ms. In these samples, the bound excludes 299, the tested controller keeps 700 inside the voltage box during the fault, 602 over the whole window, and 201 violations remain unresolved; no exclusion meets a box-safe trajectory, and no sample reaches the lower voltage limit. The Spearman rank correlations of the tested peak are 0.49 for the signed buffer injection, 0.47 for the source power, and 0.35 for the buffer-line inductance, and are negative for the receiver load and the source capacitance fraction; the lower-envelope maximum follows the same ordering at a median

of 0.017 p.u. below the tested peak. Around the principal configurations (±2% operating voltage, ±20% inductances, 0.5–1.5 ohm, 20–28 ms interruptions, ±10 MW buffer injection, up to 10 MW receiver load), the exclusions and tested classifications persist in most of the 150 samples each; the relocated configuration, whose nominal peak lies 0.011 p.u. above the limit, keeps its exclusion in 33. In 3000 random two-region samples with a common operating voltage of 0.97–1.03 p.u., the necessary condition excludes 816 cases; 182 are ideal-feasible realization gaps, and no exclusion meets a safe droop trajectory.

Table 16. Randomized parameter studies: exclusions, tested classifications and agreement.

Study	Samples	Excluded by the Necessary Test	Tested Trajectory Within Limits	Unresolved Violations	Exclusions Meeting a Safe Trajectory
Circuit, wide perturbations	1200	299	700 (fault window); 602 (whole window)	201	0
Circuit, 150 perturbed samples around each principal configuration (remote buffer/relocation/0.15 H/4000 MW/s)	4 × 150	139/33/1/124	2/79/96/8	9/38/53/18	0
Two-region, common operating voltage 0.97–1.03 p.u.	3000	816	2002	182	0

7. Discussion of Scope and Engineering Use

An infeasible energy relaxation excludes every physical trajectory contained in its admissible set. A feasible ideal schedule instead establishes resource sufficiency only for the dynamics retained in that model. Configuration decisions therefore depend on both the active energy constraints and the validity of this trajectory-set inclusion.

7.1. From an Energy Deficit to a Configuration Decision

A positive optimal addition $J^\star$ with a positive dual lower bound excludes zero-addition feasibility. Node prices locate where extra upper headroom reduces that deficit, while interface multipliers identify the useful transfer intervals. In the tested early-availability case, positive prices occur at A and B because surplus reaches those nodes after remote transfer has closed.

Local absorption or earlier source reduction can reduce a deficit accumulated while an interface is closed. Higher independent port ratings cannot increase simultaneous transfer while the shared-device limit remains binding. In the relocation case, the maximum moves to the buffer after restoration, so sizing from the source peak alone would accept a configuration that still exceeds a node-voltage limit.

The computed prices value absorption additions within the specified constraints. Upper headroom at an unchanged lower limit is provided by a higher permissible capacitor voltage within the insulation and submodule ratings, a braking resistor, storage connected through its own converter, or a lower pre-disturbance operating voltage. Added capacitance raises both margins through Equation (24), and the circuit mapping of Table 13 carries the equipment parameters into the sizing comparison.

The assessment enters a design procedure with four steps: temporal screening of a candidate configuration with its mode schedule, returning an exclusion with prices and multipliers or a feasible schedule; the circuit voltage bound with the candidate's line parameters, bandwidths and current limits, whose exclusion ends the evaluation; controller-specific simulation of the passes; and confirmation of the retained candidates with the MMC model. Only an exclusion is final in the first two steps: a feasible ideal schedule or a satisfied voltage bound leaves the candidate unresolved, and it is recorded as safe only after its controller-specific simulation and the MMC confirmation. The prices rank the measures, and Equation (24) or the circuit mapping converts the ranked addition into capacitance or a chopper rating, as in Table 13.

7.2. Validity Under Mode and Parameter Uncertainty

The cases prescribe mode schedules. With uncertain restoration timing, a control-independent exclusion must cover every permitted schedule, either through their union or through an outer capability bound. Failure under one unfavorable sequence excludes only that sequence. Initial-voltage uncertainty also changes both absorption and release margins and therefore belongs in the state constraints. The randomized study carries operating voltage, switching time, line parameters, signed injections, and loads through both the bound and the tested trajectory. No exclusion meets a safe trajectory in 4800 samples; the signed buffer injection and the source power are the strongest drivers of the peak, and the classifications of the principal configurations persist except where the nominal peak lies within 0.011 p.u. of the limit.

Minimum injection is sufficient for the nonnegative-surplus comparison with no compulsory downstream withdrawal. In a network with fixed loads and lower voltage limits, the same substitution creates the 0.5 MJ shortage of the compulsory-load example and excludes the feasible 50 MW injection. Such a network requires an admissible source set with its power, energy and ramp restrictions, which the optimization with injection as a decision variable provides.

7.3. Scope of the Models

Table 17 lists, for each model class, the assumptions under which it is used, the guarantee it provides, and the physical effects outside it. The averaged circuit retains line-current dynamics and first-order converter-port currents, and the switching-function MMC model adds submodule-voltage dispersion, circulating currents, arm inductance, modulation limits, and the phase-locked loop. The circuit uses receiver power-command feedforward and a prescribed restoration event, so its duration boundaries are conditional on those controls; its voltage envelopes cover the interruption and the recovery window, and every exclusion is validated in ball arithmetic.

The distance between the lower envelope and a tested violation has three sources. The lower and upper current envelopes may be attained by different hypothetical trajectories at different times; the receiver's blocked current is known, while any admissible restoration current widens the recovery envelope; and every tested trajectory stays above the box lower limit, so that limit adds no exclusion. Class II removes the part caused by the unbounded upper command and leaves three tested violations (grid_000, grid_014 and grid_024, all with 0.1 H buffer lines) without exclusion; their peaks lie within 0.0023 p.u. of the limit.

The cases use constructed parameters positioned in the ranges of Table 5. Their deterministic counts quantify the selected parameter design and the randomized studies quantify the sampled perturbations; neither estimates disturbance probabilities or annual reliability. Processor timings describe the tested sparse problems; they do not include measurement, control execution, or communication delays.

Table 17. Model classes, their assumptions, guarantees and the effects outside each model.

Model	Assumptions	Guarantee	Effects Outside the Model
Ideal two-region interface (Section 3)	Nonnegative prescribed injection; lossless one-way interface; energy intervals	Necessary and sufficient, Equations (14) and (16); penalty, Equation (17)	Line storage and losses; voltage-dependent current limits
Temporal optimization (Section 4)	(E1)–(E4) for exactness; (N1) for necessity	Exact under (E1)–(E4); necessary for ramp-limited, lossy and bidirectional classes	Voltage-dependent current limits; internal converter states
Averaged R–L circuit (Section 5)	First-order port currents; fixed droop with feedforward; prescribed restoration	Controller-specific trajectories and duration boundaries	Submodule ripple, arm inductance, modulation limits, AC synchronization
Circuit voltage bound (Section 5)	Source class $p_c \geq p_{\min}$ (Class I) or $p_{\min} \leq p_c \leq P_0$ (Class II); receiver blocked, then within $[0, I_{r,\max}]$	Necessary exclusion over the full window, validated numerically; a satisfied bound leaves the case unresolved	Source current–voltage correlation beyond the iterated envelopes
Switching-function MMC model (Section 5)	N submodules per arm; direct DC modulation; stiff AC sources; dq control and phase-locked loop	Averaged terminal peaks within 0.0077 p.u. at 30 kJ/MVA and 0.015 p.u. for 15–40 kJ/MVA; up to 0.040 p.u. in the nominal mapping	Receiver blocking; AC-side phase jumps; protection

8. Conclusions

Three classes of results follow, each with its own guarantee: necessary and sufficient within the ideal model, necessary only for the circuit voltage bound, and specific to the tested controller for the simulations. Exact ideal-model feasibility: the causal transfer envelope combines finite initial release, energy-arrival timing, and remote headroom into a necessary and sufficient condition for the ideal two-region network. In the 2592-case study, it identifies 958 infeasible cases and, among the 670 of 2160 aggregate-capacity passes that violate the limit under fixed droop, 144 controller-realization gaps; a causal controller with a 0.5 ms lag closes 134 of them, and Equation (17) gives the margin for a slower converter. Temporal optimization extends the exact condition to shared converters and changing interface availability under (E1)–(E4) and excludes 495 of 576 network cases, 59 of them beyond shared-equipment prefix bounds. It remains necessary for ramp-limited, lossy, and bidirectional classes, and coupled capacitance sizing reduces the required addition in 95 of the 288 cases with initial release.

Necessary physical-model exclusion: in the three-terminal R–L circuit, the coupled voltage envelope excludes 30 of 60 fault-window cases, each confirmed by validated arithmetic with an enclosure radius below 5.5 V against margins above 1275 V. Extended through the recovery interval, it adds and removes none; iterated voltage coupling and the available-power ceiling raise the count to 34 and 37 and leave three tested violations within 0.0023 p.u. of the limit. Across 1200 perturbed circuit samples and 3000 two-region samples, no exclusion meets a safe trajectory. A satisfied bound leaves a case unresolved: 10 of the 30 grid cases and 201 of the 901 randomized circuit samples that satisfy it violate the limit under the tested controller, and only the controller-specific simulation establishes safe operation.

Controller-specific results: with unchanged total headroom, relocation shifts the detected duration boundary from 8.327 to 20.390 ms and lower buffer-line inductance to 31.797 ms, while doubling the curtailment rate changes it only to 8.354 ms. Added

capacitance of 29.8 µF at the receiving node or a 35.8 MW braking resistor brings the remote-buffer case within the limit, whereas the source needs 122.5 µF and the buffer cannot. The switching-function MMC model reproduces the averaged terminal peaks within 0.0077 p.u. at 30 kJ/MVA station energies and within 0.015 p.u. for common station energies of 15–40 kJ/MVA, where every exclusion is confirmed. With the nominal energies, whose receiving station holds 9.3 kJ/MVA, the peaks differ by up to 0.040 p.u., 41 of 60 classifications agree, and 10 of 30 exclusions meet MMC terminal peaks below the limit. The averaged port therefore carries the configuration conclusions for stations in the MMC energy range, and peaks within 0.015 p.u. of the limit are decided by the MMC model.

Three refinements extend the assessment. Energy controllers that decouple the terminal voltage from the submodule energy enter the envelope through an additional bounded channel for the controlled DC current. Receiver blocking and AC-side phase jumps add discrete modes whose union the exclusion covers. Measured distributions of curtailment delay, current-limit settings, and restoration logic replace the ranges of Table 5 and turn the grid counts into disturbance-weighted counts through the randomized study.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/electronics15225321/s1, Supplementary Material S1: Reproducibility package with the complete case definitions, the computation scripts, the switching-function MMC model, and the numerical results of every table and figure.

Author Contributions: Conceptualization, J.Q. and F.L.; methodology, J.Q.; software, J.Q. and B.M.; validation, J.Q. and B.M.; formal analysis, J.Q. and F.L.; investigation, R.S. and D.W.; resources, R.S.; data curation, B.M.; writing—original draft preparation, J.Q. and F.L.; writing—review and editing, F.L. and R.S.; visualization, B.M. and D.W.; supervision, F.L.; project administration, G.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Technology Service Project of State Grid Corporation of China (Research on Power Grid Development Planning for the 15th Five-Year Plan Period—Research on Planning Schemes), grant number B3440825G027. The funders had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Data Availability Statement: The original contributions presented in this study are included in the article and in Supplementary Material S1. This reproducibility package contains the complete case definitions, the computation scripts, the switching-function MMC model, and the numerical results of every table and figure. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: Authors Jishuo Qin, Fan Li, Binqi Ma, Dan Wang, and Guodong Guo were employed by the company State Grid Economic and Technology Research Institute Co., Ltd. Author Rui Shi was employed by the company State Grid Corporation of China. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest. The authors declare that this study received funding from the Technology Service Project of State Grid Corporation of China (Research on Power Grid Development Planning for the 15th Five-Year Plan Period—Research on Planning Schemes). The funder was not involved in the study design, collection, analysis, interpretation of data, the writing of this article, or the decision to submit it for publication.

References

1. Feltes, C.; Wrede, H.; Koch, F.W.; Erlich, I. Enhanced Fault Ride-Through Method for Wind Farms Connected to the Grid Through VSC-Based HVDC Transmission. *IEEE Trans. Power Syst.* **2009**, *24*, 1537–1546. [CrossRef]
2. Ramtharan, G.; Arulampalam, A.; Ekanayake, J.B.; Hughes, F.M.; Jenkins, N. Fault ride through of fully rated converter wind turbines with AC and DC transmission systems. *IET Renew. Power Gener.* **2009**, *3*, 426–438. [CrossRef]
3. Tzelepis, D.; Rousis, A.O.; Dyśko, A.; Booth, C.; Strbac, G. A new fault-ride-through strategy for MTDC networks incorporating wind farms and modular multi-level converters. *Int. J. Electr. Power Energy Syst.* **2017**, *92*, 104–113. [CrossRef]

4. Cao, S.; Xiang, W.; Lu, X.; Lin, W.; Zhang, K.; Wen, J.; Zhang, X. Energy dissipation of MMC-HVDC based onshore wind power integration system with FB-DBS and DCCB. *IET Renew. Power Gener.* **2020**, *14*, 222–230. [CrossRef]
5. Li, H.; Yu, H.; Liu, Z.; Li, F.; Wu, X.; Cao, B.; Zhang, C.; Liu, D. Long-term scenario generation of renewable energy generation using attention-based conditional generative adversarial networks. *Energy Convers. Econ.* **2024**, *5*, 15–27. [CrossRef]
6. Ilves, K.; Norrga, S.; Harnefors, L.; Nee, H.-P. On Energy Storage Requirements in Modular Multilevel Converters. *IEEE Trans. Power Electron.* **2014**, *29*, 77–88. [CrossRef]
7. Samimi, S.; Gruson, F.; Delarue, P.; Colas, F.; Belhaouane, M.M.; Guillaud, X. MMC Stored Energy Participation to the DC Bus Voltage Control in an HVDC Link. *IEEE Trans. Power Deliv.* **2016**, *31*, 1710–1718. [CrossRef]
8. Freytes, J.; Akkari, S.; Rault, P.; Belhaouane, M.M.; Gruson, F.; Colas, F.; Guillaud, X. Dynamic Analysis of MMC-Based MTDC Grids: Use of MMC Energy to Improve Voltage Behavior. *IEEE Trans. Power Deliv.* **2019**, *34*, 137–148. [CrossRef]
9. Bergna-Diaz, G.; Freytes, J.; Guillaud, X.; D'Arco, S.; Suul, J.A. Generalized Voltage-Based State-Space Modeling of Modular Multilevel Converters With Constant Equilibrium in Steady State. *IEEE J. Emerg. Sel. Top. Power Electron.* **2018**, *6*, 707–725. [CrossRef]
10. Bergna-Diaz, G.; Suul, J.A.; D'Arco, S. Energy-Based State-Space Representation of Modular Multilevel Converters with a Constant Equilibrium Point in Steady-State Operation. *IEEE Trans. Power Electron.* **2018**, *33*, 4832–4851. [CrossRef]
11. Freytes, J.; Bergna, G.; Suul, J.A.; D'Arco, S.; Gruson, F.; Colas, F.; Saad, H.; Guillaud, X. Improving Small-Signal Stability of an MMC With CCSC by Control of the Internally Stored Energy. *IEEE Trans. Power Deliv.* **2018**, *33*, 429–439. [CrossRef]
12. Gnanarathna, U.N.; Gole, A.M.; Jayasinghe, R.P. Efficient Modeling of Modular Multilevel HVDC Converters (MMC) on Electromagnetic Transient Simulation Programs. *IEEE Trans. Power Deliv.* **2011**, *26*, 316–324. [CrossRef]
13. Saad, H.; Peralta, J.; Dennetière, S.; Mahseredjian, J.; Jatskevich, J.; Martinez, J.A.; Davoudi, A.; Saeedifard, M.; Sood, V.; Wang, X.; et al. Dynamic Averaged and Simplified Models for MMC-Based HVDC Transmission Systems. *IEEE Trans. Power Deliv.* **2013**, *28*, 1723–1730. [CrossRef]
14. Merlin, M.M.C.; Green, T.C. Cell capacitor sizing in multilevel converters: Cases of the modular multilevel converter and alternate arm converter. *IET Power Electron.* **2015**, *8*, 350–360. [CrossRef]
15. Tavakoli, S.D.; Prieto-Araujo, E.; Gomis-Bellmunt, O. AC Fault Ride Through in MMC-Based HVDC Systems. *IEEE Trans. Power Deliv.* **2022**, *37*, 2775–2786. [CrossRef]
16. Wan, Y.; Mao, M.; Zhou, L.; Xi, X.; Xie, B.; Zhou, S. Review on topology-based dc short-circuit fault ride-through strategies for MMC-based HVDC system. *IET Power Electron.* **2020**, *13*, 203–220. [CrossRef]
17. Fan, B.; Liu, T.; Zhao, F.; Wu, H.; Wang, X. A Review of Current-Limiting Control of Grid-Forming Inverters Under Symmetrical Disturbances. *IEEE Open J. Power Electron.* **2022**, *3*, 955–969. [CrossRef]
18. Taul, M.G.; Wang, X.; Davari, P.; Blaabjerg, F. Current Limiting Control With Enhanced Dynamics of Grid-Forming Converters During Fault Conditions. *IEEE J. Emerg. Sel. Top. Power Electron.* **2020**, *8*, 1062–1073. [CrossRef]
19. Qoria, T.; Gruson, F.; Colas, F.; Denis, G.; Prevost, T.; Guillaud, X. Critical Clearing Time Determination and Enhancement of Grid-Forming Converters Embedding Virtual Impedance as Current Limitation Algorithm. *IEEE J. Emerg. Sel. Top. Power Electron.* **2020**, *8*, 1050–1061. [CrossRef]
20. Rokrok, E.; Qoria, T.; Bruyere, A.; Francois, B.; Guillaud, X. Transient Stability Assessment and Enhancement of Grid-Forming Converters Embedding Current Reference Saturation as Current Limiting Strategy. *IEEE Trans. Power Syst.* **2022**, *37*, 1519–1531. [CrossRef]
21. Fan, B.; Wang, X. Equivalent Circuit Model of Grid-Forming Converters With Circular Current Limiter for Transient Stability Analysis. *IEEE Trans. Power Syst.* **2022**, *37*, 3141–3144. [CrossRef]
22. Kkuni, K.V.; Yang, G. Effects of current limit for grid forming converters on transient stability: Analysis and solution. *Int. J. Electr. Power Energy Syst.* **2024**, *158*, 109919. [CrossRef]
23. Baeckeland, N.; Yang, B.; Seo, G.-S. Unified Model of Current-Limiting Grid-Forming Inverters for Large-Signal Analysis. *IEEE Trans. Power Syst.* **2026**, *41*, 198–213. [CrossRef]
24. Willems, J.C. Dissipative dynamical systems part I: General theory. *Arch. Ration. Mech. Anal.* **1972**, *45*, 321–351. [CrossRef]
25. Ford, L.R., Jr.; Fulkerson, D.R. Constructing Maximal Dynamic Flows from Static Flows. *Oper. Res.* **1958**, *6*, 419–433. [CrossRef]
26. Hoppe, B.; Tardos, É. The Quickest Transshipment Problem. *Math. Oper. Res.* **2000**, *25*, 36–62. [CrossRef]
27. Song, Y.; Hill, D.J.; Liu, T. Characterization of Cutsets in Networks With Application to Transient Stability Analysis of Power Systems. *IEEE Trans. Control Netw. Syst.* **2018**, *5*, 1261–1274. [CrossRef]
28. Dinh, N.S.; Duong, T.L.; Truong, V.A.; Nguyen, T.T. Determine the location and capacity of energy storage in the power system using the improved Min-Cut algorithm. *Sci. Technol. Dev. J. Eng. Technol.* **2020**, *3*, 339–351. [CrossRef]
29. Liu, J.; Zhang, Y.; Conejo, A.J.; Qiu, F. Ensuring Transient Stability With Guaranteed Region of Attraction in DC Microgrids. *IEEE Trans. Power Syst.* **2023**, *38*, 681–691. [CrossRef]
30. Chang, F.; Cui, X.; Wang, M.; Su, W.; Huang, A.Q. Large-Signal Stability Criteria in DC Power Grids With Distributed-Controlled Converters and Constant Power Loads. *IEEE Trans. Smart Grid* **2020**, *11*, 5273–5287. [CrossRef]

31. Mazenc, F.; Bernard, O. Interval observers for linear time-invariant systems with disturbances. *Automatica* **2011**, *47*, 140–147. [CrossRef]
32. Cruz, R.L. A calculus for network delay. I. Network elements in isolation. *IEEE Trans. Inf. Theory* **1991**, *37*, 114–131. [CrossRef]
33. Johansson, F. Arb: Efficient Arbitrary-Precision Midpoint-Radius Interval Arithmetic. *IEEE Trans. Comput.* **2017**, *66*, 1281–1292. [CrossRef]
34. Li, R.; Xu, L.; Yao, L. DC Fault Detection and Location in Meshed Multiterminal HVDC Systems Based on DC Reactor Voltage Change Rate. *IEEE Trans. Power Deliv.* **2017**, *32*, 1516–1526. [CrossRef]
35. Wang, W.; Barnes, M.; Marjanovic, O.; Cwikowski, O. Impact of DC Breaker Systems on Multiterminal VSC-HVDC Stability. *IEEE Trans. Power Deliv.* **2016**, *31*, 769–779. [CrossRef]

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