mathematics

Article

Energy-Auditable Distributed Virtual Asynchronous Machine Control for Thermal-Energy-Storage-Based Virtual Energy Storage Systems

Wentao Yang, Yibo Wang *, Yuhan Guo and Runze Zhang

Department of Electrical Engineering, School of Information Science and Engineering, Northeastern University, Shenyang 110819, China
* Correspondence: 2271028@stu.neu.edu.cn

Abstract

Converter-dominated power systems increasingly require flexible resources that can support frequency while preserving a physically interpretable energy trajectory. Thermal-energy-storage (TES)-based virtual energy storage systems (VESSs) can shift electrical demand within thermal-energy and comfort constraints, and have therefore attracted extensive interest. However, existing studies commonly coordinate requested or normalized power without fully connecting it to actuator execution and the electrical-to-thermal energy path. Command-level sharing cannot be directly equated with executed physical power, and controller storage cannot be combined with joule-valued hardware energy without dimensional separation. Therefore, this paper proposes a physically coupled and energy-traceable virtual asynchronous machine (VAM) control method for TES-based VESSs. First, a loss-resolved averaged model establishes the point-of-common-coupling (PCC)–converter–DC-link–actuator–TES physical chain and separates the hardware Hamiltonian from the dimensionless control Lyapunov function. Second, a neighbor-coupled marginal controller is embedded in a command–projection–execution chain so that frequency regulation and weighted sharing are evaluated using the executed service. Third, a constraint-handling mechanism combines directional headroom gating, actuator saturation and ramp limits, thermal comfort bounds, and request-inactive state reset to maintain executable trajectories under the declared constraints. Simulations under a sustained 120 kW disturbance show that primary-only control retains a -0.08883 Hz steady-state offset, whereas the proposed nominal case restores frequency. In the constrained case, the 30 s terminal trend remains above the prescribed limit, while both terminal windows of the 60 s run satisfy the restoration criterion; the final-window mean error and dimensionless eligible-unit sharing spread are -6.317×10^{-5} Hz and 6.564×10^{-5}, respectively. The model-internal electrical–thermal balance achieves a dimensionless relative RMS residual of 6.2361×10^{-8}. Because the PCC voltage/current pair is reconstructed from the same power source, this residual quantifies model-internal consistency rather than independent measured closure. These results demonstrate constrained frequency restoration, executed-power coordination, and energy traceability within the averaged-model scope.

Keywords: virtual asynchronous machine; thermal energy storage; neighbor-coupled secondary control; physical energy balance; energy traceability; DC-link energy; constrained power sharing

MSC: 93C10; 93D05; 93D15; 93A30; 34H05

Academic Editor: Aydin Azizi

Received: 27 June 2026
Revised: 24 July 2026
Accepted: 4 August 2026
Published: 6 August 2026

1. Introduction

Converter-dominated power systems must regulate frequency with fewer directly coupled rotating masses, making grid-forming and virtual-machine control an important interface technology. Grid-forming converters can provide inertia-like and damping response in converter-rich networks [1,2]. TES-based VESSs provide complementary flexibility by shifting electrical demand within stored-energy and comfort envelopes [3,4]. Their combined value depends on more than forming a power reference: electrical boundary power, DC-link storage, actuator execution, heat delivery, losses, and TES energy must remain connected by an interpretable physical path.

Research relevant to this problem can be organized into four connected streams: TES flexibility and controllable-load aggregation, converter-interface and virtual-machine control, distributed secondary coordination, and energy-based physical modeling. In the first stream, Arteconi et al. analyzed thermal storage as a demand-side-management resource [5]. Kim and Norford coordinated commercial-building thermal storage with price response and distribution-network operation [6]. Hurtado et al. quantified how structural thermal storage and comfort management determine available demand flexibility across climate zones [7]. Finck et al. evaluated power-to-heat and TES flexibility in terms of energy, power, and cost [8]. Tang and Wang incorporated TES operation and thermal comfort into a model-predictive demand-response controller [9], while Chen et al. experimentally compared passive thermal mass with active storage in buildings [10]. These studies show that thermal flexibility is shaped by storage capacity, service duration, ambient conditions, and comfort limits rather than by a power command alone.

Aggregate load models translate resource flexibility into grid services. Callaway and Hiskens established a control-oriented basis for coordinating flexible electrical loads [11]. Generalized battery models capture the power and energy limits of aggregated thermostatically controlled loads (TCLs) [4]. Geometric formulations describe their feasible aggregate power trajectories [12]. Compressor lock time changes the available regulation capacity [13], while network and coupling constraints further tighten the feasible battery approximation [14]. These models establish aggregate feasibility, but they do not identify how converter power becomes executed heat and stored thermal energy in a VAM-controlled unit.

At the converter interface, the synchronverter concept established a direct way to reproduce synchronous-machine behavior with power electronics [15]. Schiffer et al. derived reduced microgrid models from component-level physics and made the assumptions behind voltage-source representations explicit [16]. Arghir et al. matched averaged converter dynamics to a synchronous machine by using the DC-link voltage as a power-imbalance signal [17]. For integrated electrical–thermal systems, Wang et al. coupled VAM dynamics with real-time power balancing and adaptive inertia [18]. A subsequent damping–droop study identified a practical small-signal stability region for VAM-controlled interlinking converters [19]. Bidirectional source/load operation was then incorporated into a VAM/VSG power-balancing strategy [20]. These developments establish interface dynamics, model-reduction principles, and operating conditions, while the loss-resolved path from the electrical boundary to TES execution remains outside their principal scope.

Distributed secondary control supplies the next layer of coordination. Simpson-Porco et al. used distributed averaging to restore microgrid frequency and voltage with neighboring information [21]. The stability properties of DAPI frequency control were subsequently characterized in a power-system setting [22]. Dörfler et al. connected distributed secondary regulation with economic optimality in microgrids [23], and Trip et al. used a distributed internal-model design for optimal frequency regulation with voltage dynamics [24]. Zhang et al. extended distributed secondary restoration and weighted sharing to heterogeneous VAM-based VESSs with thermal and communication constraints [25]. These results demonstrate the effectiveness of neighbor-based coordination, but convergence of an abstract allocation variable does not by itself show that constrained converter or TES powers execute the same allocation.

Energy-based modeling provides a complementary route to physical consistency. Interconnection and damping assignment formalized energy shaping for port-controlled Hamiltonian systems [26]. Control by interconnection further clarified the relation between power ports, storage, and closed-loop dynamics [27]. The port-Hamiltonian overview of van der Schaft and Jeltsema consolidated the storage, interconnection, and dissipation structure used across physical domains [28]. Stegink et al. coupled physical power-network and market dynamics through a unified energy-based formulation [29]. At the broader complex-systems level, Bucolo et al. showed numerically and experimentally that chaos-modulated coupling can direct energy transport in networks of capacitively coupled RLC oscillators [30]. This result reinforces the need to identify the physical channels and coupling mechanisms through which energy is redistributed, while the present TES–VAM framework addresses a different plant and control objective. Monshizadeh et al. further developed a power-controlled Hamiltonian representation for electrical systems whose inputs act directly on the power balance [31]. Applied to a VAM–TES system, this structure requires a dimensional distinction between joule-valued hardware energy and normalized controller storage, together with an explicit link between the control output and the physical actuator port.

VAM–TES research has progressed from real-time power balancing and damping–droop parameter screening [18,19] to bidirectional source/load operation and four-unit distributed secondary restoration with weighted sharing, thermal comfort, and communication effects [20,25]. These advances establish that VAM units can restore frequency and share power. They leave unresolved whether the commanded service is carried through one loss-resolved, dimensionally consistent electrical–thermal implementation.

Taken together, the literature establishes flexible-resource aggregation, converter-based frequency support, distributed coordination, and energy-structured analysis. However, aggregate or commanded power is not equivalent to executed physical power; a controller Lyapunov coordinate is not a physical energy state; and separate control and energy-accounting models do not establish one traceable implementation. TES-based VESSs therefore require a framework that constructs a unique electrothermal path, connects secondary control to constrained actuator execution, and separates physical energy balance from control-stability analysis.

Therefore, this paper proposes a physically coupled and energy-traceable VAM control method for TES-based VESSs. The main contributions are as follows:

- A loss-resolved averaged model connects PCC power, converter loss, DC-link energy, dynamic actuator power, actuator efficiency, separated thermal withdrawals, and TES energy. Raw command, projected command, executed power, delivered heat, and stored energy are assigned distinct physical roles, and the joule-valued hardware Hamiltonian is separated from the dimensionless control Lyapunov function.
- A neighbor-coupled marginal secondary controller is embedded in the command–projection–execution–PCC–frequency chain. Frequency restoration and weighted sharing are formulated using the executed service, and conditional equilibrium results are established for a connected graph and a feasible fixed active set.
- A constraint-handling and traceability mechanism combines directional headroom gating, actuator saturation and ramp limits, thermal comfort bounds, request-inactive state reset, and a source–unit–sign–equation map. This mechanism converts controller

outputs into physically executable trajectories while preserving the origin and role of every energy term.

Validation uses one four-unit averaged Simulink plant for sustained-disturbance controller comparison, constrained execution, native-grid energy-balance checks, and direct correspondence with an independently coded same-equation MATLAB ODE. Separate sensitivity studies examine thermal time scales, communication conditions, parameter uncertainty, and numerical post-processing.

2. System Architecture

Figure 1 summarizes the causal organization of the four-unit averaged VAM–TES system.

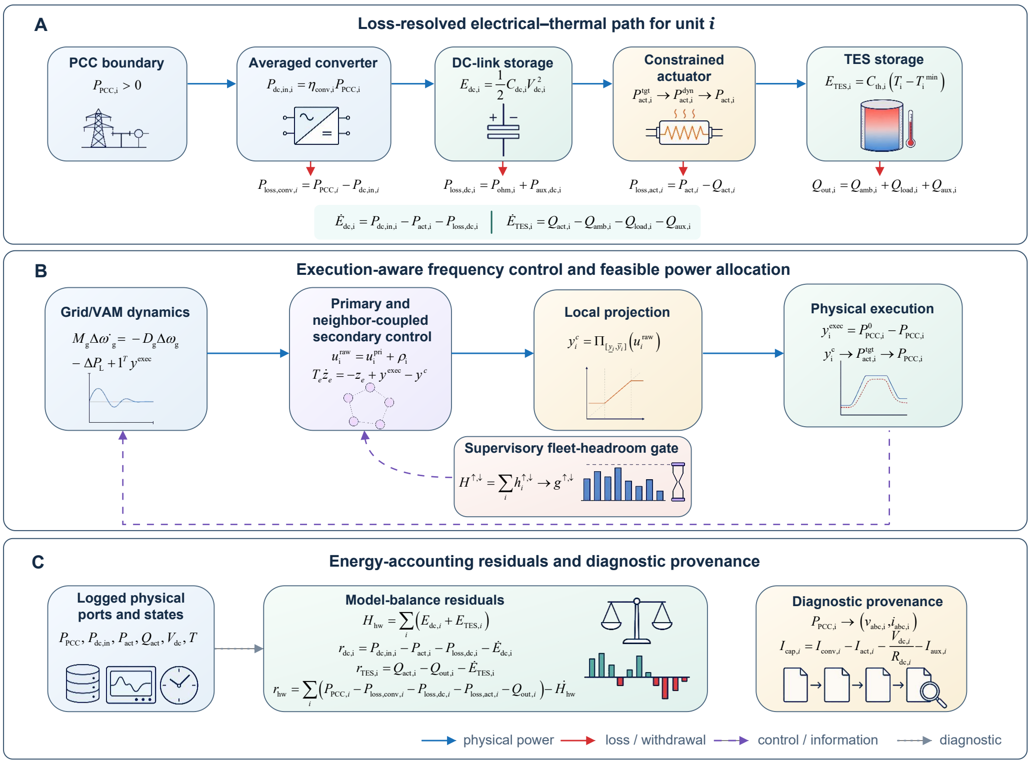

Figure 1. Physical, control, and diagnostic architecture of the four-unit averaged virtual asynchronous machine (VAM)–thermal energy storage (TES) system. PCC denotes the point of common coupling, and DC denotes direct current.

Panel A follows physical power from the PCC boundary through conversion, DC storage, constrained actuation, and TES storage, with each loss or withdrawal shown on the corresponding branch. Panel B connects grid/VAM dynamics, neighbor-coupled secondary control, local projection, physical execution, and the supervisory fleet-headroom gate. Panel C maps the logged ports and states to the model-balance residuals and identifies the provenance of the diagnostic channels. Component symbols provide visual orientation; the equations, arrows, and line styles define the model semantics.

The three-panel separation prevents a diagnostic reconstruction from being interpreted as an additional physical port. The averaged PCC phase variables are reconstructed from the modeled boundary power, and the capacitor-current channel follows from the declared DC current balance. These channels therefore evaluate model-internal consistency rather than independent sensor closure. Panel B also makes the coordination architecture explicit: the secondary law is neighbor-coupled, whereas the directional fleet-headroom sum is supplied by a supervisory aggregator. Sections 3–5 develop the physical chain, control loop, and diagnostic scope in the same order.

3. Unified Electrical–Thermal Model and Signal Taxonomy

The model follows the three roles separated in Figure 1. Panel A fixes the power directions, unique stored energies, and loss-resolved electrical–thermal chain. Panel B introduces the command, projection, and execution coordinates developed in Section 4. Panel C reserves reconstructed quantities for diagnostics so that they cannot be counted as additional physical inputs. Within this structure, the secondary state ρ_i is expressed in watts, physical sharing is evaluated from y_i^{exec}/w_i, and directional fleet headroom is supplied by the supervisory aggregation path.

3.1. Power Directions and Disjoint Signal Roles

For each of the $N = 4$ units, electrical power is positive from the grid into the unit. Thermal actuator power is positive into the TES, and ambient, served-load, and auxiliary heat flows are positive out of the TES. Grid support is represented by the service coordinate

$$y_i = P^0_{\text{PCC},i} - P_{\text{PCC},i}, \tag{1}$$

so that $y_i > 0$ means a reduction in grid consumption relative to the declared equilibrium baseline $P^0_{\text{PCC},i}$. This service coordinate is not an additional energy input; it is calculated from the physical PCC power and classified in Table 1.

Table 1. Disjoint signal classes in the averaged model.

Class	Examples	Model Role	Traceability Role
Modeled boundary variable	$P_{\text{PCC},i}$	Averaged AC boundary	Supply term in the model identity
Same-source reconstruction	$v_{\text{abc},i}$, $i_{\text{abc},i}$	Synthetic pair derived from $P_{\text{PCC},i}$	Trace check only; not independent metering
Physical states	$V_{\text{dc},i}$, T_i	DC and TES storage	Enter H_{hw}
Executed powers	$P_{\text{act},i}$, $Q_{\text{act},i}$	Modeled physical transfer	Internal port, counted once
Loss/withdrawal ports	$P_{\text{loss,conv}}$, $P_{\text{loss,dc}}$, $P_{\text{loss,act}}$, Q_{amb}, Q_{load}, Q_{aux}	Dissipation or declared withdrawal	Subtracted once
Commands	u_i^{raw}, y_i^c	Controller output and projection	Never an independent supply
Controller states	ρ_i, $z_{e,i}$	Marginal and execution-filter states	Stability coordinates only
Derived diagnostics	$I_{\text{cap},i}$, y_i, y_i^{exec}/w_i, residuals	Algebraic channel or postprocessing	Not independent measurements

3.2. Loss-Resolved Electrical-to-Thermal Chain

The physical path for unit i is

$$P_{\text{PCC},i} \to P_{\text{dc,in},i} \to E_{\text{dc},i} \to P_{\text{act},i} \to Q_{\text{act},i} \to E_{\text{TES},i}. \tag{2}$$

The averaged converter has no unreported energy state. Its input and static loss are

$$P_{\text{dc,in},i} = \eta_{\text{conv},i} P_{\text{PCC},i}, \tag{3}$$

$$P_{\text{loss,conv},i} = P_{\text{PCC},i} - P_{\text{dc,in},i}. \tag{4}$$

The unique DC-link storage state and its separated losses satisfy

$$E_{\mathrm{dc},i} = \frac{1}{2} C_{\mathrm{dc},i} V_{\mathrm{dc},i}^2, \tag{5}$$

$$\dot{E}_{\mathrm{dc},i} = P_{\mathrm{dc,in},i} - P_{\mathrm{act},i} - P_{\mathrm{ohm},i} - P_{\mathrm{aux,dc},i}, \tag{6}$$

$$P_{\mathrm{ohm},i} = \frac{V_{\mathrm{dc},i}^2}{R_{\mathrm{dc},i}}, \qquad P_{\mathrm{aux,dc},i} = V_{\mathrm{dc},i} I_{\mathrm{aux},i}. \tag{7}$$

Equivalently, the capacitor-current channel gives

$$\dot{E}_{\mathrm{dc},i} = C_{\mathrm{dc},i} V_{\mathrm{dc},i} \dot{V}_{\mathrm{dc},i} = V_{\mathrm{dc},i} I_{\mathrm{cap},i}. \tag{8}$$

The DC branch currents used in the implementation are explicitly

$$I_{\mathrm{act},i} = \frac{P_{\mathrm{act},i}^{\mathrm{dyn}}}{\max(V_{\mathrm{dc},i}, V_{\mathrm{div}})}, \qquad P_{\mathrm{act},i} = V_{\mathrm{dc},i} I_{\mathrm{act},i}, \tag{9}$$

$$I_{\mathrm{cap},i} = I_{\mathrm{conv},i} - I_{\mathrm{act},i} - \frac{V_{\mathrm{dc},i}}{R_{\mathrm{dc},i}} - I_{\mathrm{aux},i}, \qquad V_{\mathrm{div}} = 100 \text{ V}. \tag{10}$$

The fixed voltage floor is used only in division; it does not clip or reset the physical DC-voltage state.

The raw service command is projected onto instantaneous physical bounds,

$$y_i^c = \Pi_{[\underline{y}_i, \overline{y}_i]}(u_i^{\mathrm{raw}}), \qquad u_i^{\mathrm{raw}} = u_i^{\mathrm{pri}} + \rho_i. \tag{11}$$

The bounds combine actuator power, converter current, voltage validity, and a continuous comfort margin. The projected command maps to the actuator target through the physical PCC and DC losses,

$$P_{\mathrm{act},i}^{\mathrm{tgt}} = \mathrm{clip}\Big[\eta_{\mathrm{conv},i}\left(P_{\mathrm{PCC},i}^0 - y_i^c\right) - P_{\mathrm{loss,dc},i}, P_i^{\min}, P_i^{\max}\Big]. \tag{12}$$

Command execution is dynamic and rate-limited:

$$\dot{P}_{\mathrm{act},i}^{\mathrm{dyn}} = \mathrm{clip}\left(\frac{P_{\mathrm{act},i}^{\mathrm{tgt}} - P_{\mathrm{act},i}^{\mathrm{dyn}}}{\tau_{a,i}}, -r_i^{\downarrow}, r_i^{\uparrow}\right). \tag{13}$$

The converter current is also a continuous limited state. In normal operation its target is

$$I_{\mathrm{conv},i}^{\mathrm{tgt}} = \mathrm{clip}\left[\frac{P_{\mathrm{act},i}^{\mathrm{dyn}} + P_{\mathrm{loss,dc},i}}{\max(V_{\mathrm{dc},i}, V_{\mathrm{div}})} + k_{v,i}(V_{\mathrm{dc},i}^{\star} - V_{\mathrm{dc},i}), I_i^{\min}, I_i^{\max}\right], \tag{14}$$

followed by a first-order current response and a declared current-ramp limit. The executed electrical actuator power is the modeled, logged port product in Equation (9), rather than the command or the unqualified actuator state alone.

The thermal interface and TES satisfy

$$Q_{\mathrm{act},i} = \eta_{\mathrm{act},i} P_{\mathrm{act},i}, \qquad P_{\mathrm{loss,act},i} = P_{\mathrm{act},i} - Q_{\mathrm{act},i}, \tag{15}$$

$$E_{\mathrm{TES},i} = C_{\mathrm{th},i}(T_i - T_i^{\min}), \tag{16}$$

$$\dot{E}_{\mathrm{TES},i} = Q_{\mathrm{act},i} - Q_{\mathrm{amb},i} - Q_{\mathrm{load},i} - Q_{\mathrm{aux},i}, \tag{17}$$

$$Q_{\mathrm{amb},i} = U_i A_i (T_i - T_{\mathrm{amb}}). \tag{18}$$

Here $T_i^{\min}$ is the declared lower comfort bound and the TES-energy reference, and $\eta_{\mathrm{act},i} \leq 1$ is a conversion efficiency, not a heat-pump coefficient of performance. The signed ambient term is a loss only when $T_i > T_{\mathrm{amb}}$. Equations (3)–(18) contain one copy of every physical power and do not add a command or sharing coordinate as a second input.

3.3. Explicit Initialization and Precharge Modes

The model uses four contiguous modes rather than a voltage-threshold mask: initialization on $[0, 0.05)$ s, controlled precharge on $[0.05, 0.75)$ s, continuous handover on $[0.75, 2.0)$ s, and normal operation from 2.0 s. During precharge, the actuator is disabled and the converter follows a limited precharge-current target. During handover, converter and actuator targets are continuously blended; no storage state is reset at a transition. Controller state derivatives are held at zero until normal operation. Voltage thresholds are retained only as diagnostics and never define or delete scored samples.

3.4. Variables, Units, and Balance Membership

The supplementary signal dictionary maps every exported variable to one source block, unit, sign, sample policy, role, and manuscript equation. The averaged plant exports raw and projected commands, executed electrical and thermal powers, converter and actuator efficiencies, saturation and ramp flags, capacitor current, separated losses, TES boundary flows, controller states, sharing diagnostics, and communication-status fields. The base plant's communication-status fields describe its nominal configuration; dynamic delay and packet loss are evaluated with the receiver-state extension defined in Section 5. The balanced three-phase PCC v–i reconstruction is generated from the same averaged power source and therefore checks sign, scaling, and logging consistency rather than independent hardware metering. Section 7 interprets the numerical results within these modeling limits. The principal variables and their dimensional roles are summarized in Table 2.

Table 2. Principal variables and their dimensional roles.

Symbol	Unit	Meaning	In Hardware Balance?
$V_{\mathrm{dc}}, I_{\mathrm{cap}}$	V, A	DC-link voltage and capacitor current	Via E_{dc}
$I_{\mathrm{conv}}, I_{\mathrm{act}}$	A	Converter and executed-actuator branch currents	Via physical DC ports
C_{dc}	F	DC-link capacitance	Yes
T, C_{th}	K, J/K	TES temperature and capacitance	Via E_{TES}
$\Delta\omega_g, s_i$	rad/s	Grid deviation and VAM slip	No; dynamic coordinates
δ_i	rad	VAM virtual angle	No
M_i, D_i, K_i	W s^2/rad, W s/rad, W/rad	Power-domain VAM coefficients	No physical joule sum
$u_i^{\mathrm{raw}}, y_i^c, y_i^{\mathrm{exec}}$	W	Raw, projected, and executed service	No direct entry; y_i^{exec} is derived from $P_{\mathrm{PCC},i}$ and only the underlying absolute physical ports are counted
$\rho_i, z_{e,i}$	W	Marginal state and execution-mismatch filter	No; controller coordinates
w_i	1	Allocation weight	No
k_c, k_e	s^{-1}	Consensus rates	No
K_I^F	W/rad	Fleet integral coefficient	No
S_b, τ_b	state-dependent; s	Diagonal physical state bases and time base	No
$\bar{x}, \bar{A}_{\mathrm{cl}}, Q_c, P_c, V_{\mathrm{ctrl}}$	1	Normalized stability quantities	No

4. Hardware Energy Balance and Physically Coupled Secondary Control

The analysis follows the causal separation in Figure 1. Panel A supplies a joule-valued hardware Hamiltonian and its physical power ports. Panel B supplies the power-domain VAM and secondary-control dynamics, together with a dimensionless control Lyapunov function. Panel C supplies the residual used to evaluate the implemented balance. Only the hardware Hamiltonian enters the physical energy identity. The VAM coefficient M_i has unit W s^2/rad in the implemented power-domain equation; consequently, $M_i s_i^2/2$ is not, without an additional physical scaling, DC-link or TES energy in joules.

4.1. Assumptions and Claim Domain

Assumption 1 (Physical ports and storage)**.** *For every unit,* $C_{\mathrm{dc},i} > 0$, $C_{\mathrm{th},i} > 0$, *the physical states are finite on the simulated interval, and every power in Equations* (6) *and* (17) *has the sign and source declared in Table 1. A command affects the hardware balance only through its executed physical power.*

Assumption 2 (Communication and controller coordinates)**.** *The nominal communication graph is undirected and connected, with Laplacian* $L = L^\top \succeq 0$, $L\mathbf{1} = 0$, *positive allocation weights* w_i, *and positive controller gains. The nominal analysis assumes zero communication delay and no packet dropout. The corresponding logged status fields do not, by themselves, establish robustness to either effect.*

Assumption 3 (Fixed feasible equilibrium)**.** *For the equilibrium statements below, the disturbance is constant, the service request lies in the aggregate box constraints, the active set and comfort-dependent bounds are fixed locally, and the electrical, converter, actuator, and thermal states admit a steady solution. At that solution,* $y^{\mathrm{exec}} = y^c$ *and the execution-filter state is zero.*

These assumptions deliberately distinguish an equilibrium characterization from a global convergence proof. Switching active sets, infeasible requests, time-varying thermal bounds, delays, and packet loss are outside the theorem claims unless they are separately analysed.

4.2. Joule-Valued Power-Port Hamiltonian

For unit i, collect the physical energy states as

$$x_{\mathrm{hw},i} = \mathrm{col}(E_{\mathrm{dc},i}, E_{\mathrm{TES},i}), \qquad H_{\mathrm{hw},i} = \mathbf{1}^\top x_{\mathrm{hw},i}. \tag{19}$$

The associated power-port vector is

$$p_{\mathrm{hw},i} = \mathrm{col}(P_{\mathrm{PCC},i}, P_{\mathrm{loss,conv},i}, P_{\mathrm{ohm},i}, P_{\mathrm{aux,dc},i}, P_{\mathrm{act},i}, Q_{\mathrm{act},i}, Q_{\mathrm{amb},i}, Q_{\mathrm{load},i}, Q_{\mathrm{aux},i}). \tag{20}$$

Using energy rather than voltage or temperature as the state gives the block-level power-port representation

$$\dot{x}_{\mathrm{hw},i} = B_{\mathrm{hw}} p_{\mathrm{hw},i}, \qquad B_{\mathrm{hw}} = \begin{bmatrix} 1 & -1 & -1 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & -1 & -1 \end{bmatrix}. \tag{21}$$

The conjugate port selector is $B_{\mathrm{hw}}^\top \nabla H_{\mathrm{hw},i}$, so its product with $p_{\mathrm{hw},i}$ has unit watts. This formulation does not insert a fictitious electrical resistance into an energy coordinate. Converter, DC-link, actuator, ambient, service-load, and auxiliary losses remain explicit power ports. The internal actuator transfer cancels when $P_{\mathrm{loss,act},i} = P_{\mathrm{act},i} - Q_{\mathrm{act},i}$ is used.

Define the fleet hardware Hamiltonian by

$$H_{\mathrm{hw}} = \sum_{i=1}^{N}(E_{\mathrm{dc},i} + E_{\mathrm{TES},i}), \qquad [H_{\mathrm{hw}}] = \mathrm{J}. \tag{22}$$

Theorem 1 (Loss-resolved hardware energy identity)**.** *Under Assumption 1, the model in Section 3 satisfies*

$$\dot{H}_{\mathrm{hw}} = \sum_{i=1}^{N} \big(P_{\mathrm{PCC},i} - P_{\mathrm{loss,conv},i} - P_{\mathrm{ohm},i} - P_{\mathrm{aux,dc},i} - P_{\mathrm{loss,act},i} - Q_{\mathrm{amb},i} - Q_{\mathrm{load},i} - Q_{\mathrm{aux},i}\big). \tag{23}$$

Neither a raw command nor a controller state is a term in this identity.

Proof. Add Equations (6) and (17), substitute $P_{\mathrm{dc,in},i} = P_{\mathrm{PCC},i} - P_{\mathrm{loss,conv},i}$ and $Q_{\mathrm{act},i} = P_{\mathrm{act},i} - P_{\mathrm{loss,act},i}$, and sum over the units. The internal $P_{\mathrm{act},i}$ terms cancel once, leaving Equation (23). □

The numerical hardware residual used later is therefore

$$r_{\mathrm{hw}} = \sum_{i=1}^{N} \big(P_{\mathrm{PCC},i} - P_{\mathrm{loss,conv},i} - P_{\mathrm{ohm},i} - P_{\mathrm{aux,dc},i} - P_{\mathrm{loss,act},i} - Q_{\mathrm{amb},i} - Q_{\mathrm{load},i} - Q_{\mathrm{aux},i}\big) - \dot{H}_{\mathrm{hw}}. \tag{24}$$

Equation (23) is an exact model identity. A small residual computed from signals generated by those same equations tests implementation, signs, scaling, sampling, and numerical differentiation; it does not become an independent hardware experiment. In particular, the present balanced PCC voltage–current trace is synthesized from the same averaged power channel. This evidence boundary is stated explicitly in the results.

4.3. Power-Domain VAM and Executed-Service Coupling

Let $\Delta\omega_g$ be grid angular-frequency deviation, δ_i the virtual angle, and s_i the VAM slip coordinate. The common averaged plant uses

$$M_g \Delta\dot{\omega}_g = -D_g \Delta\omega_g - \Delta P_L + \mathbf{1}^{\top} y^{\mathrm{exec}}, \tag{25}$$

$$\dot{\delta} = s, \tag{26}$$

$$M\dot{s} = -Ds - K\delta - k_f \Delta\omega_g, \tag{27}$$

$$u^{\mathrm{pri}} = K\delta + Ds. \tag{28}$$

Thus, the controller does not inject an abstract sharing coordinate into the grid equation. It produces the command in Equation (11); the electrical and actuator dynamics then generate y^{exec}, which is the only service term fed back to Equation (25).

Let $W = \mathrm{diag}(w_i)$ and define the marginal coordinate $q = W^{-1}\rho$. The proposed candidate controller is

$$y^c = \Pi_{[\underline{y},\overline{y}]}\Big(u^{\mathrm{pri}} + \rho\Big), \tag{29}$$

$$T_e \dot{z}_e = -z_e + y^{\mathrm{exec}} - y^c, \tag{30}$$

$$\dot{\rho} = -K_I^F \alpha\, \phi(\Delta\omega_g) - k_c L W^{-1}\rho - k_e L W^{-1} z_e, \tag{31}$$

where $T_e = \mathrm{diag}(\tau_{e,i})$, $\alpha_i > 0$, $\mathbf{1}^\top \alpha = 1$, and K_I^F has unit W/rad. In every reported Proposed-KKT run, $\alpha = W\mathbf{1}/(\mathbf{1}^\top W\mathbf{1})$. The directional gate is defined by

$$H^\uparrow = \sum_i \max(\overline{y}_i - y_i^c, 0), \qquad g^\uparrow = \frac{H^\uparrow}{H^\uparrow + H_0}, \tag{32}$$

$$H^\downarrow = \sum_i \max(y_i^c - \underline{y}_i, 0), \qquad g^\downarrow = \frac{H^\downarrow}{H^\downarrow + H_0}, \tag{33}$$

$$\phi(\Delta\omega_g) = g^\downarrow \max(\Delta\omega_g, 0) - g^\uparrow \max(-\Delta\omega_g, 0). \tag{34}$$

The z_e term communicates transient command–execution mismatch; it is not labeled as a proven anti-windup law. Steady allocation follows from marginal consensus and projection. Moreover, $H^\uparrow$ and $H^\downarrow$ are exact fleet sums in the present model. Neighbor disagreement is distributed, but the directional gate is supervisory; hence, the implementation is not claimed to be fully distributed until a dynamic aggregate estimator or an explicit aggregator channel is modelled and tested.

Theorem 2 (Conditional frequency restoration)**.** *Suppose Assumptions 2 and 3 hold and the applicable directional headroom gate is strictly positive. Every equilibrium of Equations (25)–(31) satisfies*

$$\Delta\omega_g^\star = 0, \qquad \mathbf{1}^\top y^{\mathrm{exec},\star} = \Delta P_L. \tag{35}$$

Proof. Left-multiply the equilibrium form of Equation (31) by $\mathbf{1}^\top$. Since $\mathbf{1}^\top L = 0$ and $\mathbf{1}^\top \alpha = 1$, one obtains $K_I^F \phi(\Delta\omega_g^\star) = 0$. Positive directional headroom makes ϕ zero only at $\Delta\omega_g^\star = 0$. Substitution into the equilibrium grid balance (25) gives the second equality. □

The theorem is conditional on existence of a feasible equilibrium. It does not assert restoration for the infeasible-request experiment, nor does it establish convergence through all projection and comfort-boundary switches.

4.4. Constrained Weighted Allocation at Equilibrium

Theorem 3 (Conditional KKT allocation)**.** *Under Assumptions 2 and 3, and at a restored equilibrium with $u^{\mathrm{pri},\star} = 0$, the controller equilibrium satisfies*

$$\rho_i^\star = w_i\lambda, \qquad y_i^{\mathrm{exec},\star} = \mathrm{clip}(w_i\lambda, \underline{y}_i, \overline{y}_i), \tag{36}$$

for a multiplier λ that enforces $\mathbf{1}^\top y^{\mathrm{exec},\star} = \Delta P_L$. The executed service is the unique solution of

$$\begin{aligned} \underset{y}{\text{minimize}} \quad & \frac{1}{2} y^\top W^{-1} y, \\ \textit{subject to} \quad & \mathbf{1}^\top y = \Delta P_L, \qquad \underline{y} \le y \le \overline{y}. \end{aligned} \tag{37}$$

Proof. At the stated equilibrium, $z_e^\star = 0$ and $\Delta\omega_g^\star = 0$, so Equation (31) gives $LW^{-1}\rho^\star = 0$. Connectivity implies $W^{-1}\rho^\star = \lambda\mathbf{1}$. The projection relation then gives Equation (36). Together with fleet balance, these are the KKT conditions of Equation (37). Since $W \succ 0$, the objective is strictly convex and the feasible executed-service vector is unique. □

The experimentally reported sharing variable is therefore y_i^{exec}/w_i, not ρ_i/w_i alone. When some units are on a bound, equality is assessed only among the remaining free units, while fleet balance and feasibility are reported separately.

4.5. Control Lyapunov Function and Local Stability Scope

Let x_{cl} contain every differentiable state of the selected closed-loop mode, including grid, VAM, controller, converter, actuator, DC-link, and TES states. Around an equilibrium with a fixed active set, choose a diagonal matrix of positive physical base values S_b, select the declared time base $\tau_b = 1$ s, and define the dimensionless state, time, and Jacobian

$$\bar{x} = S_b^{-1}(x_{\mathrm{cl}} - x_{\mathrm{cl}}^{\star}), \quad \bar{t} = \frac{t}{\tau_b}, \quad \frac{d\bar{x}}{d\bar{t}} = \bar{A}_{\mathrm{cl}}\bar{x} + o(\|\bar{x}\|), \quad \bar{A}_{\mathrm{cl}} = \tau_b A_{\mathrm{cl}}. \tag{38}$$

Theorem 4 (Local Lyapunov certificate)**.** *If the complete fixed-mode Jacobian A_{cl} (equivalently, $\bar{A}_{\mathrm{cl}}$) in Equation (38) is Hurwitz, then for every $Q_c = Q_c^{\top} \succ 0$ there is a unique $P_c = P_c^{\top} \succ 0$ satisfying*

$$\bar{A}_{\mathrm{cl}}^{\top} P_c + P_c \bar{A}_{\mathrm{cl}} = -Q_c. \tag{39}$$

The function

$$V_{\mathrm{ctrl}} = \frac{1}{2}\bar{x}^{\top} P_c \bar{x}, \tag{40}$$

is dimensionless and proves local exponential stability of that smooth closed-loop mode.

Proof. The matrices $\bar{A}_{\mathrm{cl}}$, Q_c, and P_c are dimensionless, so the Lyapunov equation and V_{ctrl} are dimensionally closed. The equation has a unique positive-definite solution for a Hurwitz matrix. Along the normalized linearized dynamics, $dV_{\mathrm{ctrl}}/d\bar{t} = -\bar{x}^{\top} Q_c \bar{x}/2 < 0$ for $\bar{x} \neq 0$; the standard indirect method extends the conclusion locally to the smooth nonlinear mode. □

The mathematical function V_{ctrl} is never added to H_{hw}. The theorem also does not convert a mode-specific Jacobian into a global hybrid, delay-robust, or dropout-robust guarantee. The numerical stability evidence is reported with exactly this scope in Section 6.

5. Materials and Methods: Unified Model and Validation Setup

The validation setup instantiates the three paths in Figure 1 within a common numerical framework. The four-unit averaged Simulink plant resolves the physical chain in Panel A; each controller acts through the projection and execution loop in Panel B; and the logged ports feed the residual and provenance checks in Panel C. An independently coded MATLAB implementation of the same state equations provides the correspondence test and complete-state linearizations. Targeted long-duration and communication studies examine the slower thermal dynamics and network assumptions. The models, scenarios, parameters, metrics, and claim limits are defined below.

5.1. Four-Unit Averaged Plant

The primary plant is implemented in MATLAB/Simulink R2024b Update 9. Its 33 continuous states comprise the common grid-frequency state $\Delta\omega_g$ and, for each unit, the VAM angle and slip (δ_i, s_i), DC voltage and converter current $(V_{\mathrm{dc},i}, I_{\mathrm{conv},i})$, dynamic actuator state $P_{\mathrm{act},i}^{\mathrm{dyn}}$, TES temperature T_i, secondary state ρ_i, and execution-mismatch filter z_i. Thus,

$$\begin{aligned} x = \mathrm{col}\big(&\Delta\omega_g, \delta_{1:4}, s_{1:4}, V_{\mathrm{dc},1:4}, I_{\mathrm{conv},1:4}, \\ &P_{\mathrm{act},1:4}^{\mathrm{dyn}}, T_{1:4}, \rho_{1:4}, z_{1:4}\big) \in \mathbb{R}^{33}. \end{aligned} \tag{41}$$

Here, $P_{\mathrm{act},i}^{\mathrm{dyn}}$ is the rate-limited actuator state. The executed electrical power used in the service calculation and physical balance is the logged port product $P_{\mathrm{act},i} = V_{\mathrm{dc},i} I_{\mathrm{act},i}$, where $I_{\mathrm{act},i}$ and the capacitor-current split follow Equations (9) and (10). The passive physical-port convention takes $P_{\mathrm{PCC},i} > 0$ from the grid into unit i, whereas the service coordinate $y_i = P^0_{\mathrm{PCC},i} - P_{\mathrm{PCC},i}$ is positive when electrical consumption is reduced. The grid-frequency equation therefore uses the executed service y_i^{exec}, while the energy balance uses the absolute PCC, converter, DC-link, actuator, and thermal powers.

Four controller configurations act on the same plant, initial condition, disturbance, and constraint map: primary-only droop, distributed averaging proportional–integral (DAPI) control, the closest previous VAM consensus method, and the proposed execution-aware KKT controller. The proposed controller exchanges neighboring information on the connected line graph. Its directional feasibility calculation uses the exact fleet-headroom sum provided by a supervisory aggregator; accordingly, the realization combines neighbor-coupled secondary control with supervisory feasibility coordination.

5.2. Parameters, Initial Conditions, and Numerical Configuration

Table 3 gives the common scalar, graph, controller, and solver parameters. The VAM coefficients are generated from the ratings S_i, inertia constants H_i, natural frequencies $\omega_{n,i}$, and damping ratios ζ_i as

$$M_i = \frac{2H_i S_i}{\omega_0}, \qquad K_i = M_i \omega_{n,i}^2, \qquad D_i = 2\zeta_i M_i \omega_{n,i}. \tag{42}$$

The baseline actuator powers are initialized at the electrical–thermal equilibrium,

$$P^0_{\mathrm{act},i} = \frac{U_i A_i (T_i^0 - T_{\mathrm{amb}}) + Q_{\mathrm{load},i} + Q_{\mathrm{aux},i}}{\eta_{\mathrm{act},i}}. \tag{43}$$

The corresponding converter currents include the ohmic and auxiliary DC losses.

Table 3. Common grid, communication, control, and solver parameters.

Quantity	Value	Meaning or Unit
f_0, S_{base}	$50, 2.0 \times 10^6$	Hz; W
M_g, D_g	$2(4.0)S_{\mathrm{base}}/\omega_0, 1.0 \times 10^5$	W s^2/rad; W s/rad
Communication graph	1–2–3–4	Undirected line graph
Pinning vector	$[1, 0, 0, 1]^\top$	End-node frequency pinning
Allocation weights w_i	$[1.00, 0.90, 1.20, 1.10]^\top$	Positive relative weights
Frequency distribution α	$W\mathbf{1}/(\mathbf{1}^\top W\mathbf{1})$	Weight-normalized fleet integral allocation
DAPI local integral gains	$[20, 20, 20, 20] \times 10^3$	W/rad; effective fleet gain 80 kW/rad
Previous-VAM main gains	$[40, 40, 40, 40] \times 10^3$	W/rad; pins 1 and 4 give an 80 kW/rad effective fleet gain
Previous-VAM sensitivity gains	$[20, 20, 20, 20] \times 10^3$	W/rad; pins 1 and 4 give a 40 kW/rad effective fleet gain
Proposed aggregate integral gain	80×10^3	W/rad
k_c, k_e	0.60, 0.80	s^{-1}
$\tau_{e,i}$	0.40 for every unit	s; execution-mismatch filter
k_{reset}	0.80	s^{-1}; request-inactive ρ leakage
Headroom transition	1.0×10^3	W
DAPI droop gains	$[28, 30, 27, 31] \times 10^3$	W s/rad
Random seed	4,433,632	Fixed seed for stochastic extensions
Reference solver	`ode15s`	Variable-step stiff solver
Relative/absolute tolerance	10^{-7} / 10^{-8}	Solver tolerances
Maximum step	0.01	s

DAPI, gain-matched Previous-VAM, and Proposed-KKT use the same effective fleet aggregate integral gain of 80 kW/rad. The 40 kW/rad Previous-VAM setting is included as a literature-based sensitivity rather than in the equal-gain comparison. The coordination gain $k_c = 0.60\ \mathrm{s}^{-1}$, execution-mismatch gain $k_e = 0.80\ \mathrm{s}^{-1}$, and filter time constant $\tau_{e,i} = 0.40$ s are held constant across the declared scenarios. This gain-matching rule isolates controller structure but does not imply globally optimal tuning.

The heterogeneous unit parameters are reported in Table 4. Temperatures are absolute, and the ambient temperature is applied as $T_{\mathrm{amb}} = 278.15$ K. The actuator model represents resistance heating, so $\eta_{\mathrm{act},i}$ is an electrical-to-thermal efficiency rather than a heat-pump coefficient of performance.

Table 4. Per-unit parameters of the unified averaged model.

Parameter (Unit)	**Unit 1**	**Unit 2**	**Unit 3**	**Unit 4**
VAM rating S_i (kW)	150	160	140	155
Inertia constant H_i (s)	2.0	2.4	1.8	2.2
$\omega_{n,i}$ (rad/s)	2.0	1.8	2.2	1.9
ζ_i (–), $k_{f,i}$ (kW s/rad)	0.75; 25	0.80; 30	0.70; 28	0.85; 32
T_i^0 (K)	294.15	293.95	294.35	294.05
$T_i^{\min}/T_i^{\max}$ (K)	291.15/297.15	291.15/297.15	291.15/297.15	291.15/297.15
$C_{\mathrm{th},i}$ (10^8 J/K)	1.30	1.55	1.20	1.45
U_iA_i (kW/K)	2.6	3.0	2.4	2.8
$Q_{\mathrm{load},i}/Q_{\mathrm{aux},i}$ (kW)	45/0.8	50/0.9	55/0.7	48/0.8
$\eta_{\mathrm{act},i}$ (–)	0.980	0.975	0.985	0.978
$P_i^{\min}/P_i^{\max}$ (kW)	20/160	25/180	15/150	30/170
$\tau_{a,i}$ (s)	0.30	0.50	0.40	0.60
Up/down ramp limit (kW/s)	120	100	90	110
$V_{\mathrm{dc},i}^{\star}$ (V)	800	820	780	810
$C_{\mathrm{dc},i}$ (F)	0.080	0.090	0.075	0.085
$R_{\mathrm{dc},i}$ (kΩ)	200	220	190	210
$I_{\mathrm{aux},i}$ (A)	0.50	0.55	0.45	0.50
$\eta_{\mathrm{conv},i}$ (–)	0.980	0.975	0.985	0.978
$\tau_{\mathrm{conv},i}$ (s)	0.018	0.022	0.020	0.025
$k_{v,i}$ (A/V)	4.0	3.8	4.2	3.9
$I_{\mathrm{conv},i}^{\max}$ (A)	300	320	285	305
$I_{\mathrm{conv},i}^{\min}$ (A)	0	0	0	0
Division floor V_{div} (V)	100	100	100	100
Current ramp limit (A/s)	6000	6000	5500	5800
Precharge current limit (A)	180	205	165	190

5.3. Fast-Control and Constraint Scenarios

The principal frequency test runs for 60 s with a 1 ms output grid. A 120 kW load increase is applied at $t = 3$ s and remains present to the end of the simulation. Primary-only, DAPI, gain-matched Previous-VAM, and Proposed-KKT are evaluated under nominal parameters and under a common constraint scenario. In the latter, Unit 2 starts 0.15 K above its lower comfort boundary. Unit 1 retains its announced command bounds, while its downstream actuator minimum is set to $P_1^0 - 2.000$ kW and both physical ramp limits are reduced to 0.500 kW/s. The electrical and thermal equilibrium is recomputed after these changes so that every controller starts from the same physical state.

A separate event-removal test evaluates recovery after downstream actuator derating. The 120 kW request is active over $[3, 25)$ s, and the simulation continues to 45 s. Unit 1 retains an announced minimum of 20 kW but has a downstream actuator minimum of 87.18 kW; Unit 2 again starts 0.15 K above its lower comfort boundary. The execution-redistribution gain is $0.80\ \mathrm{s}^{-1}$. After request removal, the recovery implementation adds $-k_{\mathrm{reset}}\rho$ to the secondary-state derivative when $|\Delta P_L| \leq 1$ W; the comparison case sets

$k_{\text{reset}} = 0$. This test examines state recovery under a specific derating and does not constitute a global anti-windup result.

5.4. Signal Provenance and Electrical–Thermal Balance

Each logged signal is assigned a unique source block, engineering unit, sign convention, signal role, sample-time behavior, and associated balance equation. Commands, projected commands, executed powers, states, losses, and offline derived quantities are kept as distinct roles. The PCC phase variables used in the trace test are synthesized from the modeled averaged PCC power, and I_{cap} is calculated from the DC current balance. Consequently, these two comparisons assess model-internal consistency rather than independent sensor closure.

The normal-equilibrium energy-balance test uses a 20 s horizon and a 1 ms output grid. Storage derivatives are calculated from the exported energy states; state-equation derivatives are used only as same-equation diagnostics. With $P_{\text{loss,dc},i} = P_{\text{ohm},i} + P_{\text{aux,dc},i}$, the DC and TES residuals are

$$r_{\text{dc},i} = P_{\text{dc,in},i} - P_{\text{act},i} - P_{\text{loss,dc},i} - \frac{\mathrm{d}E_{\text{dc},i}}{\mathrm{d}t}, \tag{44}$$

$$r_{\text{TES},i} = Q_{\text{act},i} - Q_{\text{amb},i} - Q_{\text{load},i} - Q_{\text{aux},i} - \frac{\mathrm{d}E_{\text{TES},i}}{\mathrm{d}t}. \tag{45}$$

Converter and actuator bookkeeping is evaluated through

$$r_{\text{conv},i} = P_{\text{PCC},i} - P_{\text{dc,in},i} - P_{\text{loss,conv},i}, \tag{46}$$

$$r_{\text{act},i} = P_{\text{act},i} - Q_{\text{act},i} - P_{\text{loss,act},i}. \tag{47}$$

The fleet hardware balance is

$$r_{\text{hw}} = \sum_{i=1}^{4} \left(P_{\text{PCC},i} - P_{\text{loss,conv},i} - P_{\text{loss,dc},i} - P_{\text{loss,act},i} - Q_{\text{amb},i} - Q_{\text{load},i} - Q_{\text{aux},i} \right) - \frac{\mathrm{d}H_{\text{hw}}}{\mathrm{d}t}, \tag{48}$$

$$H_{\text{hw}} = \sum_{i=1}^{4} (E_{\text{dc},i} + E_{\text{TES},i}). \tag{49}$$

The VAM and controller quadratic storage used in the stability analysis is not included in the joule-valued H_{hw}. PCC trace consistency is reported separately as

$$r_{\text{PCC}} = P_{\text{PCC,log}} - v_{abc}^{\top} i_{abc}. \tag{50}$$

Peak and RMS residuals, relative peak and RMS values, and integrated mismatch are calculated over the full record. The smooth-window peak excludes only the first and last derivative samples and the interval within 0.02 s of the ideal load step. The declared limits are 10^{-4} for relative RMS, 10^{-3} for smooth-window relative peak, and 10^{-3} for integrated mismatch normalized by the associated boundary energy.

5.5. DC-Link Initialization and Precharge

The precharge simulation runs for 2.5 s on a 1 ms grid and starts with unenergized DC links. Four contiguous operating modes replace a voltage-threshold mask:

$$m(t) = \begin{cases} 0, & 0 \le t < 0.05 \text{ s} \quad \text{(initialization)}, \\ 1, & 0.05 \le t < 0.75 \text{ s} \quad \text{(precharge)}, \\ 2, & 0.75 \le t < 2.00 \text{ s} \quad \text{(handover)}, \\ 3, & t \ge 2.00 \text{ s} \quad \text{(normal operation)}. \end{cases} \tag{51}$$

All samples are retained, including those immediately before, at, and after each mode transition. The unfiltered voltage- and current-based capacitor energy derivatives are compared as

$$\left.\frac{\mathrm{d}E_{\mathrm{dc},i}}{\mathrm{d}t}\right|_V = C_{\mathrm{dc},i} V_{\mathrm{dc},i} \frac{\mathrm{d}V_{\mathrm{dc},i}}{\mathrm{d}t}, \qquad \left.\frac{\mathrm{d}E_{\mathrm{dc},i}}{\mathrm{d}t}\right|_I = V_{\mathrm{dc},i} I_{\mathrm{cap},i}. \tag{52}$$

The integrated route mismatch is normalized by the final precharge energy and limited to 10^{-3}; the normal-operation derivative NRMSE is limited to 1%.

5.6. Frequency Restoration and Executed-Power Sharing

Because the load increase persists from 3 to 60 s, the final frequency cannot return to zero merely through disturbance removal. Frequency performance is reported using steady-state error, nadir, 100 and 500 ms windowed RoCoF, integral absolute error, overshoot, settling time, final-window trend, controller effort, actuator electrical energy, executed-support energy, and constraint duration. Secondary restoration is evaluated over two adjacent terminal windows:

$$\left|\overline{\Delta f}_{50:55}\right|, \left|\overline{\Delta f}_{55:60}\right| \le 0.002 \text{ Hz}, \qquad \left|\frac{\mathrm{d}\widehat{\Delta f}_{50:55}}{\mathrm{d}t}\right|, \left|\frac{\mathrm{d}\widehat{\Delta f}_{55:60}}{\mathrm{d}t}\right| \le 10^{-4} \text{ Hz/s}. \tag{53}$$

The primary-only case is expected to retain its droop offset. A common controller-effort measure is defined from the raw unit commands:

$$\mathcal{E}_u = \sqrt{\frac{1}{4(60-3)} \int_3^{60} \sum_{i=1}^{4} \left(u_i^{\mathrm{raw}}(t)\right)^2 \mathrm{d}t}. \tag{54}$$

Power sharing is evaluated from the physical executed service rather than an independent controller coordinate. A unit belongs to the final-window eligible set $\mathcal{F}$ when the union of command projection, actuator saturation and ramp limiting, converter-current saturation and ramp limiting, and comfort-ramp activity occupies no more than 5% of the final 5 s. The relative executed-power spread is

$$\epsilon_{\mathrm{sh}} = \frac{\max_{i\in\mathcal{F}} \overline{(y_i^{\mathrm{exec}}/w_i)} - \min_{i\in\mathcal{F}} \overline{(y_i^{\mathrm{exec}}/w_i)}}{\max\left(\left|\mathrm{mean}_{i\in\mathcal{F}} \overline{(y_i^{\mathrm{exec}}/w_i)}\right|, 1 \text{ W}\right)}. \tag{55}$$

For a feasible request, the declared sharing limit is $\epsilon_{\mathrm{sh}} \le 0.02$, and the final fleet execution error is limited by

$$\epsilon_P^{\max} = \max(0.01\Delta P_L, 0.001 S_{\mathrm{base}}). \tag{56}$$

The online announced-bound feasibility flag is compared with the offline allocation condition

$$\sum_i \underline{y}_i - \varepsilon_{\mathrm{sat}} \le \Delta P_L \le \sum_i \overline{y}_i + \varepsilon_{\mathrm{sat}}, \qquad \varepsilon_{\mathrm{sat}} = 1 \text{ W}. \tag{57}$$

A second offline calculation incorporates the actual actuator and current limits, voltage-dependent DC losses, and comfort factors. Announced feasibility and downstream physical feasibility are therefore reported separately.

5.7. Simulink–ODE Correspondence

The correspondence test compares the saved Simulink plant with a separately coded 33-state MATLAB right-hand side. Both implementations use the Proposed-KKT controller, the same normal-voltage equilibrium, parameters, initial state, graph, 3 s sustained disturbance, solver configuration, and 2 ms output grid over 0–20 s. Compared outputs include grid frequency, executed actuator power, DC and TES energies, DC voltage, TES temperature, secondary states, raw and projected commands, executed service, and the frequency/headroom signal. Level and initial-excursion NRMSE, maximum normalized and absolute error, time lag, and initial error are reported. Control-output NRMSE is limited to 5%, energy NRMSE to 2%, maximum normalized error to 10%, and lag to 50 ms. This experiment tests two implementations of the same averaged equations; it does not compare the averaged plant with a switching converter model.

5.8. Long-Duration TES and Communication Studies

The long-duration TES study runs paired service and no-service cases over 0–14,400 s. Service is active over $[900, 10{,}800)$ s, followed by a one-hour recovery interval. The scenario includes time-varying ambient temperature and served heat load, heterogeneous thermal parameters, actuator efficiencies, capacities and ramps, and fixed-seed regulation requests. The low-bandwidth request consists of 300 s setpoint intervals and 30 s causal ramps arranged in 11 balanced 900 s blocks. Paired 60% and 95% headroom cases use integration steps of 0.05 and 0.025 s, respectively. With $[t_a, t_b) = [900, 10{,}800)$ s, the continuous-time tracking score is

$$\mathrm{NRMSE}_{\mathrm{CT}} = \sqrt{\frac{\int_{t_a}^{t_b} [y^{\mathrm{exec}}(t) - r(t)]^2 \, dt}{\max\left\{ \int_{t_a}^{t_b} r^2(t) \, dt, \, [0.01 P_{\mathrm{base}}]^2 T_{\mathrm{svc}} \right\}}}. \tag{58}$$

Here, $\mathrm{NRMSE}_{\mathrm{CT}}$ is the continuous-time normalized root-mean-square tracking error; $y^{\mathrm{exec}}(t)$ and $r(t)$ are the total executed fleet service and the requested service, respectively; $P_{\mathrm{base}} = 2.0$ MW is the system base power; and $T_{\mathrm{svc}} = t_b - t_a = 9900$ s is the active-service duration, with $t_a = 900$ s and $t_b = 10{,}800$ s. Tracking NRMSE and energy bias are compared with limits of 5% and 2%. Temperature extrema, comfort violations, throughput, duty cycle, saturation duration, cumulative heat loss, and post-service recovery are also reported. The thermal inputs and parameters are Class-C assumptions, and the four-hour profile has not been run through the full Simulink plant. The long-duration results therefore characterize the declared low-bandwidth averaged thermal model rather than a calibrated building or hardware installation.

Communication robustness is studied using the complete averaged electrical, thermal, and control dynamics. The 33-state plant is linearized over an 8×8 damping–droop grid at the all-free and Unit-1-bound equilibria. A two-stage Erlang receiver is added to each directed link for the marginal and execution coordinates, increasing the connected-path model to 57 states. The receiver holds its state during packet loss, and permanent edge loss removes both directed channels. Six mean delays from 0.01 to 0.25 s, four dropout levels from 10% to 50%, three physical edge removals, and combined stress are evaluated in both operating modes. A fixed-seed uncertainty study uses 300 joint draws with $\pm 20\%$ variation in M, D, K, τ_a, and C_{th}; ± 2 percentage points in converter and actuator efficiencies; $\pm 15\%$ in load; 0.01–0.15 s mean delay; and 0–30% packet loss.

Stability metrics include spectral abscissa, minimum oscillatory damping ratio, modal participation, and nonlinear time response. The nonthermal spectral-abscissa limit is $-0.02\ \mathrm{s}^{-1}$, the minimum oscillatory damping ratio is 0.05, and the four thermal eigenvalues are compared with $-U_i A_i / C_{\mathrm{th},i}$ using a 1% relative-error limit. The line graph is a tree, so removal of any physical edge disconnects the network; edge-loss cases are interpreted with this connectivity loss rather than as N–1 redundancy.

The composite Q2 gate requires an exact equilibrium, strictly stable poles, a nonthermal spectral abscissa below $-0.02\ \mathrm{s}^{-1}$, a minimum oscillatory damping ratio of at least 0.05, a maximum relative thermal-pole error of 0.01, and passage of the finite-difference (FD) Jacobian-convergence check. The FD check compares relative perturbation steps of 10^{-6} and 3×10^{-7}; its relative Frobenius-difference and nearest-eigenvalue-difference limits are 10^{-6} and 10^{-5}, respectively.

A corresponding deterministic Simulink receiver study uses the reference plant, static communication bypass, heterogeneous and uniform receiver delays, temporary loss of the middle edge over $[4, 8)$ s, permanent middle-edge removal, and a maximum-step-halved repeat. During temporary loss, the affected receiver states are held and then resume causally when communication returns. The finite-state Erlang receiver represents a mean-delay channel; it is not an exact transport-delay or switching-network model.

5.9. Numerical Sensitivity and Post-Processing

Numerical sensitivity is evaluated over 30 s with Proposed-KKT and a common 2 ms requested output grid. Nominal and constrained cases are repeated with maximum-step halving, an alternate solver, a different zero-crossing setting, and two deterministic initial-condition perturbations. An infeasible case is run with the reference and halved maximum steps. The initial perturbations satisfy $|\Delta V_{\mathrm{dc}}| \le 0.5\%$, $|\Delta I_{\mathrm{conv}}| \le 1\%$, $|\Delta P_{\mathrm{act}}| \le 1\%$, and $|\Delta T| \le 0.05$ K.

Key-control changes are limited to 1% under step halving and 2% under solver or zero-crossing changes. The corresponding limits are 0.2% for integrated energy and 10% for residual-RMS changes when both residuals remain below 10^{-4}. Frequency restoration, executed sharing, support delivery, and energy-balance status are also checked under the two initial perturbations. The complete scenario matrix is provided in the Supplementary Materials.

Post-processing sensitivity is evaluated by decimating the native 2 ms outputs by five and reconstructing them with linear and shape-preserving cubic interpolation. A five-point central derivative is compared with the native differentiation route; zero-order hold and first-order backward difference are included as stress controls. The common evaluation mask excludes only declared discontinuities and two endpoint samples. No smoothing, clipping, or result-dependent sample deletion is applied. Residual checks use the 10^{-4} absolute criterion and a 0.1 limit on the floor-normalized change. The independently logged $C_{\mathrm{dc}} V_{\mathrm{dc}} \dot{V}_{\mathrm{dc}}$ and $V_{\mathrm{dc}} I_{\mathrm{cap}}$ routes are compared using a 1% worst-unit NRMSE limit; the algebraic origin of I_{cap} remains part of the interpretation.

5.10. Reproducibility and Scope

Every simulation stores the model identity, parameter set, initial condition, disturbance definition, solver and output grid, raw MAT data, processed tables, and processing-source identity. The complete signal dictionary, source–consumer mapping, run-level acceptance tables, artifact hashes, version chronology, and unsuccessful numerical variants are provided in the Supplementary Materials. The main text reports the physical model, scenarios, metrics, and scientific results needed to interpret the figures and tables. The resulting claims apply to the declared averaged models and internal energy traceability; independent

PCC metering, switching-converter dynamics, and calibrated hardware validation remain outside this numerical scope.

6. Validation Results

The results follow the causal order of Figure 1. The first study evaluates the Panel B closed loop under a sustained disturbance and physical constraints. The second combines the Panel A storage balances with the Panel C residuals and diagnostic provenance. The remaining studies test model correspondence and sensitivity to thermal, numerical, and communication conditions. Each result is interpreted within its measurement and modeling scope: internal balance tests assess consistency of the declared equations, while communication robustness requires the graph to remain connected.

6.1. Sustained Frequency Restoration and Executed Power Sharing

At $t = 3$ s, a 120 kW load increment is applied and remains present until the end of the 60 s simulation. The final frequency error therefore measures restoration under a sustained imbalance rather than recovery after disturbance removal. Figure 2 compares Primary control, DAPI, gain-matched Previous-VAM, and Proposed-KKT using the same averaged plant, initial condition, solver, and disturbance. The three secondary controllers have the same effective aggregate integral gain of 80 kW/rad; the original 40 kW/rad Previous-VAM setting is examined separately as a sensitivity case. The corresponding metrics are reported in Table 5.

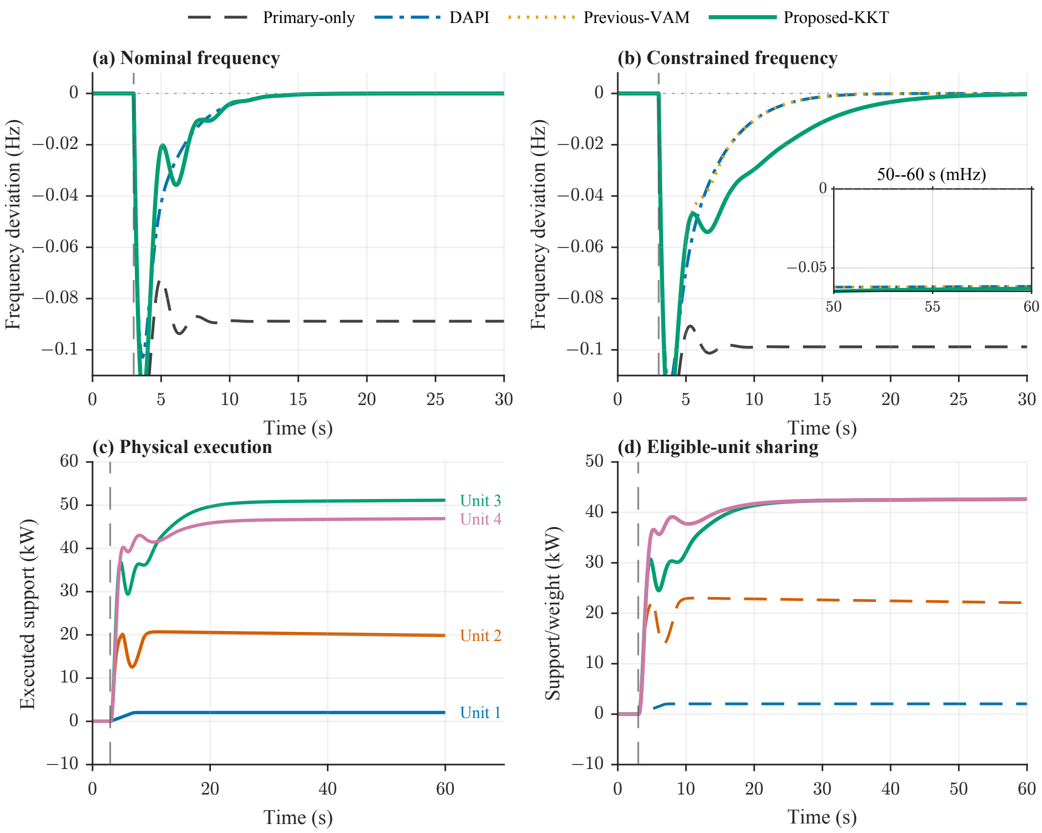

Figure 2. Closed-loop response to a sustained 120 kW load increase under nominal and constrained conditions. Panels (**a**,**b**) compare the first 30 s of the frequency responses; the inset extends the three secondary-control responses to 50–60 s under constraints. Panels (**c**,**d**) show the full 60 s executed support and weight-normalized support of Proposed-KKT. Colors and line styles identify the controllers and units. Panel (**d**) uses blue dashed, orange dashed, green solid, and magenta solid traces for Units 1–4, respectively.

Table 5. Controller performance under the sustained load disturbance. The steady-state error and trend are evaluated over 55–60 s, and the adjacent 50–55 s window is also checked for continued convergence.

(a) Frequency Restoration and Control Effort

Scenario	Controller	SS Error (Hz)	Trend (Hz/s)	IAE (Hz s)	Raw RMS (W)
Nominal	Primary	-8.883×10^{-2}	-3.606×10^{-17}	5.064	16,181
Nominal	DAPI	3.568×10^{-11}	-1.121×10^{-11}	0.2387	29,794
Nominal	Previous (matched)	3.351×10^{-11}	-1.077×10^{-11}	0.2387	29,876
Nominal	Proposed-KKT	-1.207×10^{-12}	5.113×10^{-13}	0.2400	29,855
Stress	Primary	-9.881×10^{-2}	-5.239×10^{-17}	5.638	18012
Stress	DAPI	-6.165×10^{-5}	5.366×10^{-8}	0.3562	43,951
Stress	Previous (matched)	-6.166×10^{-5}	5.380×10^{-8}	0.3562	44,044
Stress	Proposed-KKT	-6.317×10^{-5}	1.053×10^{-7}	0.5446	73,779

(b) Executed Sharing and Physical Constraint Activity

Scenario	Controller	Sharing Spread	Execution Error (W)	U1 Sat. (s)	U1 Ramp (s)	Eligible
Nominal	Primary	0.3611	−55814	0	0	4
Nominal	DAPI	9.259×10^{-9}	1.868×10^{-5}	0	0	4
Nominal	Previous (matched)	8.880×10^{-9}	1.745×10^{-5}	0	0	4
Nominal	Proposed-KKT	2.723×10^{-11}	-5.809×10^{-7}	0	0	4
Stress	Primary	0.2197	−62084	56.83	3.685	2
Stress	DAPI	2.552×10^{-6}	−38.72	56.97	3.699	2
Stress	Previous (matched)	3.067×10^{-4}	−38.73	56.87	3.689	2
Stress	Proposed-KKT	6.564×10^{-5}	−39.66	56.85	3.687	2

Primary control retains steady-state offsets of −0.08883 Hz in the nominal case and −0.09881 Hz under physical constraints. This behavior follows directly from droop control: it balances the persistent power deficit at a displaced frequency and does not coordinate normalized unit powers. The three secondary controllers reduce the final error to the 10^{-5}–10^{-4} Hz range or below and coordinate the eligible executed powers. Thus, zero error after a cleared event and coordinated sharing are recognized as distinct properties; the sustained test evaluates both without relying on disturbance removal.

The constrained Proposed-KKT trajectory also reveals a slow terminal mode. At 30 s, its five-second mean error is -5.767×10^{-4} Hz, whereas its positive trend of 1.298×10^{-4} Hz/s exceeds the prescribed 1.0×10^{-4} Hz/s settling criterion. With the same model, gains, estimator, and limits, the mean errors over 50–55 s and 55–60 s are -6.398×10^{-5} and -6.317×10^{-5} Hz, and the corresponding trends decrease to 2.429×10^{-7} and 1.053×10^{-7} Hz/s. The result therefore demonstrates eventual convergence over 60 s but does not support a 30 s settling claim under this constrained condition.

The comparison does not indicate uniform dynamic superiority. Proposed-KKT gives the smallest nominal sharing spread. Under stress, DAPI produces lower IAE, raw-command RMS, sharing spread, and execution-error magnitude than Proposed-KKT. Gain-matched Previous-VAM also has lower IAE, raw-command RMS, and execution-error magnitude, but its sharing spread is larger. The raw control effort is evaluated as

$$\left[\frac{1}{4(60-3)}\int_{3}^{60}\sum_{i=1}^{4}(u_i^{\mathrm{raw}})^2\,\mathrm{d}t\right]^{1/2}.$$

These results position the proposed controller as a physically coupled and traceable allocation method rather than as the optimum for every transient performance index.

With the original 40 kW/rad Previous-VAM integral setting, the nominal and constrained final errors are -1.433×10^{-6} and -1.422×10^{-4} Hz, respectively, and the corresponding IAEs are 0.4775 and 0.7119 Hz s. Both cases continue to converge by 60 s, but the lower gain produces a slower response and is therefore retained as a sensitivity case rather than included in the equal-gain comparison. The complete lower-gain metrics are provided in the Supplementary Materials.

The stress condition sets Unit 1's actual minimum to $P_1^0 - 2.0$ kW and its upward and downward ramp limits to 0.5 kW/s; Unit 2 starts at $T_{\min} + 0.15$ K after the equilibrium state is refreshed. For Proposed-KKT, Units 3 and 4 form the final eligible sharing set, with a normalized spread of 6.564×10^{-5} and a fleet execution error of -39.66 W. Unit 1 remains saturated for 56.85 s and ramp-limited for 3.687 s. The minimum comfort margin is 0.1431 K and the comfort-violation duration is zero. Figure 2c,d shows how mismatch feedback transfers the remaining service to the eligible units. The nonzero executed-to-projected mismatch means that this result supports constraint-aware compensation and eligible-unit sharing, but not global KKT optimality for arbitrary active-set switching.

6.2. Recovery After Request Removal

A separate experiment evaluates post-event recovery when a constraint is located downstream of command projection. Unit 1 declares a minimum of 20 kW to the controller, whereas the modeled actuator minimum is 87.18 kW. The service request is removed at 25 s and the response is observed until 45 s. Table 6 compares conditional leakage of the inactive-request controller states with an otherwise identical no-leakage ablation.

Table 6. Post-removal recovery with and without inactive-request state leakage.

(a) Event Response and Physical Constraints

Case	Sat. (s)	Ramp (s)	Event e_f (Hz)	Support Error (W)	Sharing Spread	Feasible
Leakage on	27.89	8.213	-1.448×10^{-3}	-795.4	0.002991	1
Leakage off	41.85	3.687	-1.448×10^{-3}	-795.4	0.002991	1

(b) Post-Removal State Recovery

Case	Final e_f (Hz)	Final Support (W)	Final z RMS (W)	Final ρ RMS (W)
Leakage on	-1.777×10^{-7}	-0.731	1.703	1.469
Leakage off	3.077×10^{-4}	160.7	34,277	45,471

With leakage enabled, the final frequency and support magnitudes fall to 1.777×10^{-7} Hz and 0.731 W, respectively. The minimum comfort margin is 0.1474 K, at least two units remain eligible, and no comfort or command/execution sign violation occurs. Without leakage, the final controller-state RMS values increase to 3.428×10^4 and 4.547×10^4 W, leaving 160.7 W of support after the request is removed. The contrast explains the benefit of resetting inactive-request controller memory. Because this is a specific downstream derating and event-removal test, it does not establish a universal anti-windup result or replace the sustained-disturbance analysis.

6.3. Electrical–Thermal Energy Balance and Precharge Transition

The energy-balance experiment evaluates the saved averaged Simulink model using separately exported DC-link and TES energy states. The companion precharge experiment retains all samples through four contiguous intervals: initialization (0–0.049 s), precharge (0.050–0.749 s), handover (0.750–1.999 s), and normal operation from 2.000 s. Figure 3 links the system residual, the voltage trajectory, and the two independent numerical routes used to evaluate the DC-energy derivative.

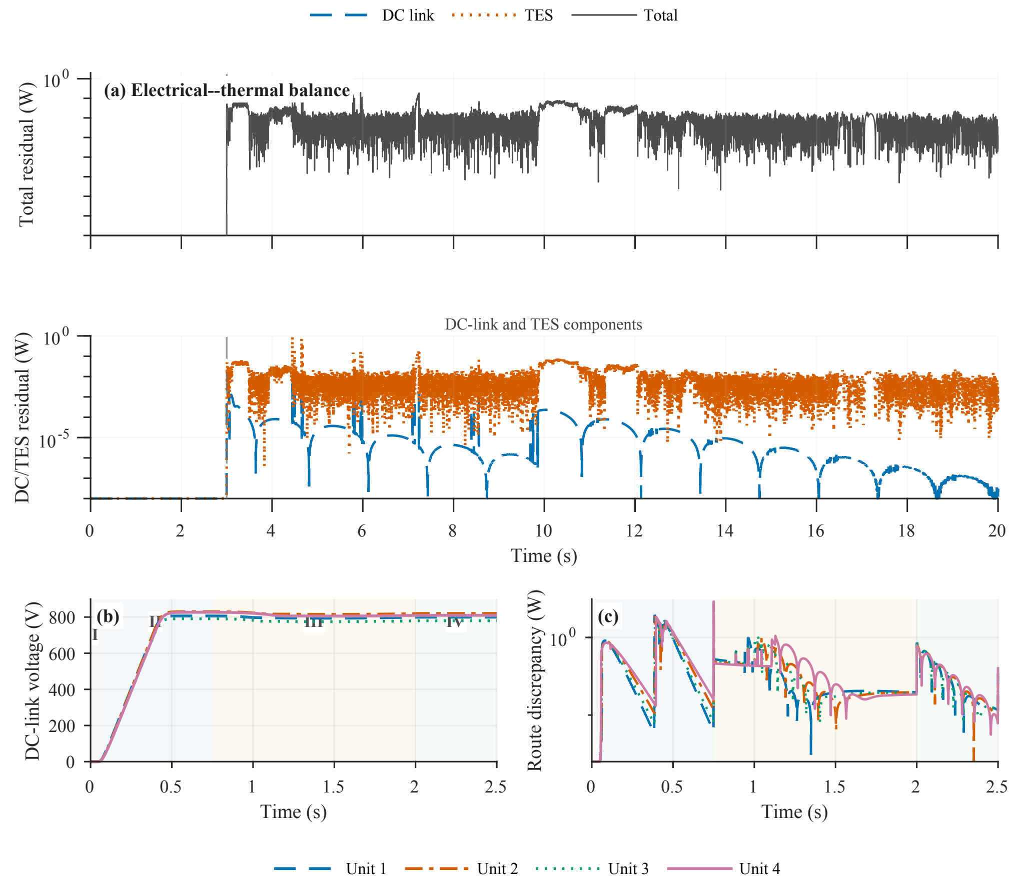

Figure 3. Internal electrical–thermal balance and DC-link precharge transition. Panel (**a**) uses a shared time axis with the total residual shown in the upper logarithmic plot and the DC-link and TES residuals resolved in the lower plot; values below 10^{-8} W are displayed at the plotting floor. Panels (**b**,**c**) show the DC-link voltage and the difference between the $C_{dc}V_{dc}\dot{V}_{dc}$ and $V_{dc}I_{cap}$ routes. Shaded intervals I–IV denote initialization, precharge, handover, and normal operation, respectively.

The fleet electrical–thermal residual has a peak of 0.9375 W, an RMS value of 0.02112 W, a dimensionless relative RMS value of 6.236×10^{-8}, and a dimensionless relative integrated mismatch of 4.823×10^{-9}. The residual peaks coincide with fast state changes, whereas the DC-link residual decreases toward the plotting floor during smooth operation. The same-equation residuals remain diagnostic channels rather than independent measurements.

All 36 samples within 1 ms of the three precharge boundaries are retained without smoothing. Across the four units, the maximum voltage is 829.8 V, the maximum current-route storage-power magnitude is 147.8 kW, and the maximum raw route difference is 228.2 W. In the normal interval, the largest NRMSE between

$$r_V = C_{dc}V_{dc}\dot{V}_{dc}, \qquad r_I = V_{dc}I_{cap}$$

is 3.724×10^{-4}. Gain, tap, current-bias, and omitted-loss perturbations change their corresponding diagnostics, demonstrating equation-level fault sensitivity without deleting the boundary samples.

The balanced three-phase PCC current is synthesized from aggregate PCC power, and I_{cap} is computed from the declared current balance. The results therefore establish internal consistency of the modeled energy path but do not constitute closure between independent electrical sensors.

6.4. Correspondence Between Simulink and the MATLAB ODE Model

The averaged Simulink model and an independently coded, same-equation 33-state MATLAB ODE implementation are compared over a common 0–20 s nominal-voltage scenario. Both implementations use the same initial state, parameters, 120 kW step at 3 s, `ode15s` settings, and 2 ms output grid. All decisive traces show nonzero event excursions and zero estimated lag. Figure 4 compares a representative frequency trajectory and the worst excursion NRMSE in each signal family after normalization by its declared criterion.

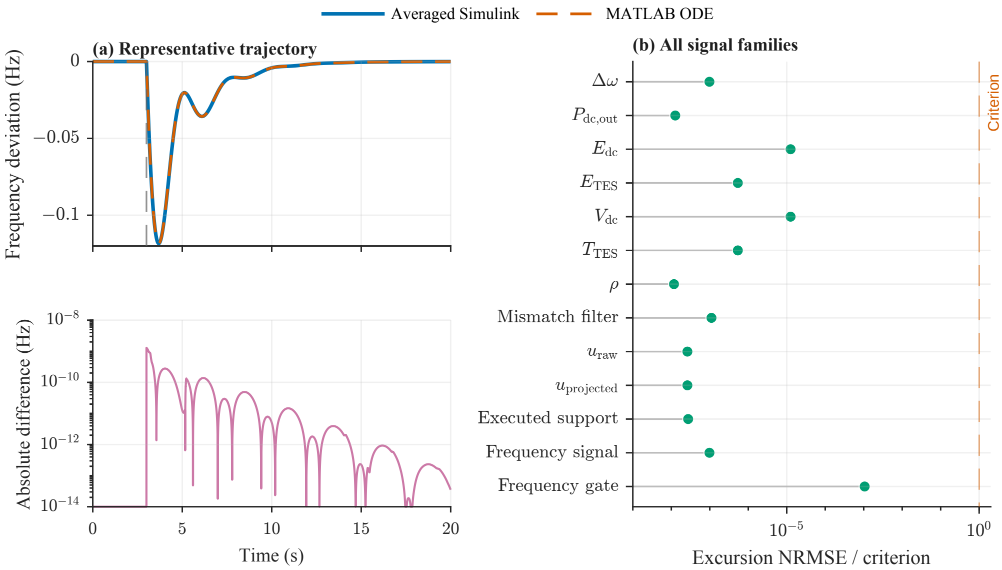

Figure 4. Correspondence between the averaged Simulink model and the independently coded 33-state MATLAB ODE implementation. Panel (**a**) compares representative frequency trajectories and their absolute difference. Panel (**b**) reports the worst excursion NRMSE in each signal family, normalized by the corresponding criterion; values to the left of the vertical line satisfy that criterion.

The close agreement confirms that the two averaged implementations realize the same state equations for this normal-voltage operating condition. The largest excursion NRMSE is 5.290×10^{-5}, the estimated lag is zero for every signal family, and all normalized values remain below their criteria. The complete family-wise error table is provided in the Supplementary Materials. Direct V_{dc}/T_{tes} and E_{dc}/E_{tes} comparisons are same-equation state representations rather than independent measurements. Accordingly, this experiment supports implementation correspondence for the declared averaged scenario; it does not extend the evidence to a switching converter model, precharge dynamics, independent PCC sensing, or arbitrary operating points.

6.5. Numerical Sensitivity and Post-Processing Dependence

The nominal, constrained, and infeasible cases are repeated with a halved maximum step, an `ode15s`-to-`ode23t` solver change, alternative zero-crossing settings, and two deterministic electrical/thermal initial-state perturbations. The frequency, sharing, support-delivery, and internal-energy conclusions remain unchanged under these variations. The initial-state tests also reveal 3.340–3.450 aggregate unit-seconds of transient command–execution sign opposition. This behavior is caused by relaxation of the non-equilibrium actuator states and shows that full-window sign invariance cannot be assumed for perturbed initial conditions. The complete normalized sensitivity matrix and sign-opposition comparison are reported in the Supplementary Materials.

Post-processing is more influential than the solver variations. Linear and shape-preserving cubic reconstruction are applied after decimating the native 2 ms output to 10 ms, while alternative differentiation is evaluated directly on the native grid. Table 7 shows that linear interpolation alters TES and total residuals beyond the declared tolerances. Shape-preserving cubic interpolation removes most of this error, but its infeasible-TES floor-normalized change is 0.1041, just above the 0.1 limit. In contrast, the five-point derivative and the two DC-energy routes satisfy their respective criteria. Hence seven predefined post-processing comparisons fail, and the native-grid energy residuals should not be interpreted as invariant to resampling.

Table 7. Sensitivity of energy residuals to reconstruction and derivative methods.

Method	Worst Metric	Outcome
Linear interpolation	7.245×10^{-4}	TES/total limits exceeded
`pchip` interpolation	1.110×10^{-5}	One normalized TES limit exceeded
Five-point derivative	5.587×10^{-8}	Within criterion
DC-energy derivative routes	5.094×10^{-4} NRMSE	Within 1% criterion

6.6. Four-Hour Thermal Response and Communication Robustness

The long-horizon study uses 300 s setpoints, 30 s causal ramps, and paired no-service, 60%-headroom, and 95%-headroom simulations over 14,400 s. Both service cases keep the continuous-time tracking NRMSE below 5% and the energy bias below 2%. The conclusions are unchanged after halving the integration step and repeating the fixed-seed input. Table 8 summarizes the thermal response.

Table 8. Four-hour low-bandwidth thermal-service results.

Case	CT NRMSE	Energy Bias	Temperature (°C)	Bound/Ramp (Unit-s)	Violation/Recovery (s)
60% headroom	0.04690	−0.006327	20.742–21.297	0/0	0/0
95% headroom	0.04682	−0.006886	20.714–21.322	837/0	0/0

The simulated temperature remains between 20.714 and 21.322 °C, with no comfort violation; the 95%-headroom case accumulates 837 unit-seconds at a power bound. These results demonstrate numerical feasibility for the specified low-bandwidth inputs. The thermal parameters and inputs are Class-C assumptions, the paired supervisor has not been independently validated, and direct full/slow-model event-window comparisons were not performed. The temperature result therefore supports a four-hour numerical case study, not a calibrated real-building comfort claim.

Communication robustness is evaluated around the actual 33-state averaged right-hand side. Two receiver stages for each of two communicated coordinates increase the connected model to 57 states. The study includes deterministic mean delay and dropout, fixed-seed binary packet trajectories, every physical edge outage, 300 joint-uncertainty draws evaluated at two operating modes, and nonlinear, modal, solver, and operating-point checks. Table 9 separates the connected and disconnected outcomes.

Table 9. Closed-loop communication results for connected and disconnected graphs.

Evaluation Set	Q2 Pass	Strictly Unstable	Marginal	Connected	FD Pass
Deterministic, all	22/30	0/30	8/30	22/30	26/30
Deterministic, connected	22/22	0/22	0/22	22/22	22/22
Monte Carlo, all mode evaluations	406/600	0/600	194/600	406/600	512/600
Monte Carlo, connected	406/406	0/406	0/406	406/406	406/406
Single-edge-loss modes	0/6	0/6	6/6	0/6	2/6

In Table 9, Q2 denotes the composite mode-separated stability-and-physics gate, and FD denotes the finite-difference Jacobian-convergence check defined in Section 5.8.

Every connected deterministic case and all 406 connected Monte Carlo mode evaluations satisfy the stability criterion. The current four-node path is not edge-redundant: removal of any physical edge disconnects the graph, so all six single-edge-loss modes retain marginal consensus poles. Binary dropout rates of 10%, 30%, and 50% preserve the declared nonlinear classification and performance for the connected cases. However, a whole-state speed criterion is not met in any of 128 evaluations, and the combined nonlinear branch/classification criteria are met in 9 of 12 cases. Thus, the results support stability under the tested connected delays, dropouts, and uncertainties, but not connected N–1 resilience or global hybrid convergence.

Linearization of the saved averaged state equations agrees closely with the MATLAB closed-loop matrix at the all-free and Unit-1-bound equilibria. For the 33-state plant and the 57-state receiver model, relative Frobenius errors range from 2.321×10^{-10} to 3.862×10^{-10}. The maximum nearest- eigenvalue discrepancy is 1.891×10^{-9} s^{-1} for the plant and 2.501×10^{-6} s^{-1} for the receiver model, while the spectral-abscissa discrepancies are approximately 1.17×10^{-15} s^{-1}. This correspondence links the stability matrix to the saved averaged Simulink equations without changing the connectivity limitations above.

The deterministic receiver is also evaluated directly in the averaged Simulink model. Under connected communication, the maximum mean frequency error and trend are 2.070×10^{-11} Hz and 8.388×10^{-12} Hz/s, the support error is 1.012×10^{-5} W, and the normalized executed-power spread is 3.416×10^{-10}. The worst smooth relative RMS and peak energy residuals are 6.000×10^{-8} and 3.051×10^{-7}. Permanent removal of the middle edge gives Laplacian rank 2, two components, and $\lambda_2 = 0$, so the connected-case restoration result does not extend to global consensus after that outage. The complete receiver table is provided in the Supplementary Materials.

A temporary loss of the middle edge over $[4, 8)$ s produces a post-reconnection receiver-derivative peak of 1.166×10^5 W/s and a state change of 1.347×10^4 W, confirming that the receiver continues to evolve after communication is restored. This deterministic averaged-model result does not represent stochastic packet transport or restore N–1 connectivity; it characterizes the response of the implemented two-stage receiver within the tested communication conditions.

7. Discussion

7.1. Control Performance Under Sustained Disturbances

The unified plant connects the secondary-control output to converter, DC-link, actuator, and thermal dynamics before the delivered service enters the grid frequency equation. This connection changes the interpretation of both restoration and sharing: the relevant quantity is the executed service rather than an independent allocation coordinate. Under the sustained 120 kW disturbance, primary-only control retains a steady-state offset of -8.883×10^{-2} Hz. The three secondary controllers restore the nominal frequency, and the proposed controller gives the smallest nominal executed-power sharing spread, 2.723×10^{-11}.

Physical limits slow the response without changing this basic mechanism. When Unit 1 is derated and ramp-limited and Unit 2 approaches its lower comfort boundary, the proposed controller reaches a final-window mean error of -6.317×10^{-5} Hz and a trend of 1.053×10^{-7} Hz/s at 60 s. The corresponding sharing spread among the eligible units is 6.564×10^{-5}, the fleet execution error is -39.66 W, and the minimum comfort margin remains 0.1431 K. The same trajectory has a 30 s terminal trend of 1.298×10^{-4} Hz/s, slightly above the prescribed 10^{-4} Hz/s limit. The two consecutive terminal windows of the 60 s run satisfy the restoration criterion, showing that constrained actuator dynamics require a longer convergence interval.

The controller comparison reveals a performance tradeoff rather than uniform dominance. DAPI produces lower integrated frequency error, command RMS, sharing spread, and execution error in the constrained case. The gain-matched previous-VAM controller also gives lower integrated error, command RMS, and execution error, although its sharing spread is larger than that of the proposed controller. The principal benefit of the proposed structure is therefore the explicit coupling of marginal coordination, physical execution, and energy tracing, while transient speed and control effort remain dependent on gain selection and active constraints.

7.2. Constraint Handling and Executed-Power Sharing

The constrained experiments separate supervisory projection from downstream physical execution. Unit 1 encounters an actuator limit that is not available to the command projection, so the controller must compensate for persistent command–execution mismatch through the remaining eligible units. Unit 2 is excluded from final sharing when its comfort-related constraint is active, and Units 3–4 converge in executed service per unit weight. This behavior supports feasible redistribution for the tested active set; arbitrary constraint switching and global optimality over all saturation combinations remain outside the equilibrium analysis.

After the regulation request is removed, conditional reset of the marginal and mismatch states prevents the residual controller energy from being observed without reset. The reset case returns the frequency and support close to zero while preserving the comfort boundary. This result concerns recovery after request removal and complements, rather than replaces, the sustained-disturbance experiment. Together, the two cases show how physical headroom, ramp limits, thermal eligibility, and controller-state recovery affect different stages of the closed-loop response.

7.3. Energy Traceability and Numerical Correspondence

The loss-resolved model assigns converter loss, DC-link storage, actuator conversion, ambient heat exchange, served thermal load, auxiliary loss, and TES energy to distinct channels. The fleet electrical–thermal balance has a peak residual of 0.9375 W, an RMS residual of 0.02112 W, and a dimensionless relative RMS residual of 6.236×10^{-8}. The precharge analysis uses explicit initialization, precharge, handover, and normal-operation intervals. All 36 samples adjacent to the three boundaries are retained without smoothing, and the two DC-energy derivative routes agree in normal operation with a worst NRMSE of 3.724×10^{-4}.

These residuals establish internal consistency of the averaged equations and logged states. The three-phase PCC voltage/current pair is reconstructed from the same averaged power channel, and the capacitor current follows from the declared current balance. The corresponding tests can reveal equation, sign, unit, and numerical-processing errors within the model, but they do not provide independent measured closure or independent sensor-fault detection.

The independently coded 33-state MATLAB ODE reproduces the averaged Simulink implementation over the common 0–20 s scenario. The worst reported excursion NRMSE is 5.290×10^{-5} and the estimated lag is zero. This agreement supports equation-to-implementation correspondence at the declared operating point. It does not extend to switching dynamics, independently sensed boundaries, or arbitrary operating points.

Numerical sensitivity is likewise method dependent. Step-size reduction, solver changes, zero-crossing settings, and the tested initial-condition perturbations preserve the principal control and energy conclusions. In contrast, decimation followed by linear reconstruction changes several residual metrics, and shape-preserving cubic reconstruction still exceeds the tolerance in one comparison. The five-point central derivative and the two algebraic DC-energy routes remain consistent. Energy-residual accuracy is therefore tied to native-grid state exports and the declared differentiation procedure rather than to arbitrary resampling.

7.4. Thermal and Communication Scope

The four-hour thermal study uses 300 s setpoints and 30 s causal ramps. Across the two service levels, continuous-time tracking NRMSE is approximately 4.68×10^{-2}, energy bias remains between -0.00689 and -0.00633, and temperature stays within 20.714–21.322 °C with zero numerical comfort violation. These results demonstrate bounded behavior for the assumed low-bandwidth, Class-C parameter set. The thermal capacitance, envelope conductance, ambient profile, served load, and conversion efficiency have not been calibrated against a specific installation, and direct correspondence between the fast and slow models has not been established. The multihour result is consequently a numerical assessment of the assumed model rather than a building-scale performance prediction.

The communication study confirms stable behavior for the tested connected delay and dropout conditions. During a temporary link interruption, the affected receiver states are held and resume coordination after reconnection. Permanent removal of the middle edge splits the four-node path into two components, with Laplacian rank 2 and $\lambda_2 = 0$. Because every single-edge outage disconnects this graph, the present topology has no N–1 communication redundancy. The directional headroom calculation also uses an exact fleet sum, so the implementation combines neighbor-based coordination with supervisory feasibility aggregation.

7.5. Applicability and Further Development

The present framework is applicable to averaged, aggregated VAM–TES plants in which converter and actuator dynamics are represented by continuous states and the communication graph remains connected. Its main contribution is a common physical and numerical chain for frequency control, executed-power sharing, and energy tracing. Practical deployment requires independently measured electrical and thermal boundary channels, identified thermal parameters, explicit sensor uncertainty, and a communication topology with redundant connectivity. Further development will replace the exact fleet-headroom sum with a distributed estimator and examine switching-level correspondence after the averaged-model parameters and boundary measurements have been calibrated.

8. Conclusions

This paper developed a physically coupled VAM control framework for a four-unit TES-based VESS. The averaged model links the PCC–converter–DC-link–actuator–TES power transfers, assigns distinct roles to commanded, projected, executed, and stored quantities, and separates the joule-valued hardware Hamiltonian from the dimensionless control Lyapunov function. The neighbor-coupled secondary controller is embedded in the command–projection–execution chain, so frequency restoration and weighted sharing are determined by executed service under connected-graph and feasible fixed-active-set conditions.

Under the sustained 120 kW disturbance, primary-only control retains a steady-state frequency offset of -8.8831×10^{-2} Hz, whereas the proposed nominal case reaches -1.2069×10^{-12} Hz and a dimensionless sharing spread of 2.7227×10^{-11}. In the common constrained case, Unit 1 is subject to physical derating and a 500 W/s ramp limit, while Unit 2 starts 0.15 K above the lower comfort boundary. At 60 s, the final-window mean frequency error, trend, and eligible-unit executed-power sharing spread are -6.3169×10^{-5} Hz, 1.0527×10^{-7} Hz/s, and 6.5644×10^{-5}, respectively. The minimum comfort margin remains 0.1431 K and the comfort-violation duration is zero. The same constrained trajectory has a 30 s terminal trend of 1.2978×10^{-4} Hz/s, which exceeds the prescribed 10^{-4} Hz/s limit and indicates a longer convergence time. DAPI and the gain-matched previous-VAM controller also satisfy the 60 s restoration criterion. In the constrained comparison, DAPI gives lower IAE, raw-command RMS, sharing spread, and execution-error magnitude; the gain-matched previous-VAM controller gives lower IAE, raw-command RMS, and execution-error magnitude but a larger sharing spread. The proposed controller gives the smallest nominal sharing spread. The comparison therefore characterizes a tradeoff between transient performance and a physically traceable constrained-execution structure.

The native-grid electrical–thermal balance yields a dimensionless relative RMS residual of 6.2361×10^{-8}. The precharge analysis retains all 36 samples around the initialization, precharge, handover, and normal-operation boundaries without smoothing. Direct comparison of the averaged Simulink model with an independently coded same-equation 33-state MATLAB ODE reproduces all compared signal families with a worst excursion NRMSE of 5.290×10^{-5} and zero estimated lag. These results establish internal consistency between the physical equations, logged states, and the two numerical implementations. The PCC voltage/current pair is reconstructed from the same power source and the capacitor current is algebraic; accordingly, the energy result represents model-internal traceability rather than independent measured closure.

The applicability of the framework is bounded by the averaged and aggregated plant representation, the supervisory fleet-headroom sum, the connected fixed-active-set analysis, and uncalibrated Class-C thermal parameters. The four-node line graph has no single-edge redundancy, and the residuals remain sensitive to interpolation after decimation, so the reported accuracy is tied to the native grid and declared differentiation route. Future work will focus on independent electrical and thermal boundary measurements, calibrated thermal-parameter identification, a distributed headroom estimator on a redundantly connected graph, and switching-level model correspondence.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/math14152859/s1. The supplementary materials contain the complete numerical protocol, acceptance thresholds, internal experiment identifiers, version provenance, full-precision results, retained unsuccessful cases, batch-check summaries, model and script hashes, signal dictionaries, raw and processed data, and reproduction manifests.

Author Contributions: Conceptualization, W.Y. and Y.W.; methodology, W.Y. and Y.W.; software, W.Y. and Y.G.; validation, W.Y., Y.W., Y.G. and R.Z.; formal analysis, W.Y.; investigation, W.Y., Y.G. and R.Z.; data curation, W.Y., Y.G. and R.Z.; writing–original draft, W.Y.; writing–review and editing, Y.W., Y.G. and R.Z.; visualization, W.Y., Y.G. and R.Z.; supervision, Y.W.; project administration, Y.W.; funding acquisition, Y.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article and Supplementary Materials. Further inquiries can be directed to the corresponding author.

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

Abbreviations

The following abbreviations are used in this manuscript:

VAM	Virtual asynchronous machine
VSM	Virtual synchronous machine
VESS	Virtual energy storage system
TES	Thermal energy storage
PCC	Point of common coupling
PH	Port-Hamiltonian
DAPI	Distributed averaging proportional-integral control

References

1. Lasseter, R.H.; Chen, Z.; Pattabiraman, D. Grid-Forming Inverters: A Critical Asset for the Power Grid. *IEEE J. Emerg. Sel. Top. Power Electron.* **2020**, *8*, 925–935. [CrossRef]
2. Rosso, R.; Wang, X.; Liserre, M.; Lu, X.; Engelken, S. Grid-Forming Converters: Control Approaches, Grid-Synchronization, and Future Trends–A Review. *IEEE Open J. Ind. Appl.* **2021**, *2*, 93–109. [CrossRef]
3. Alva, G.; Lin, Y.; Fang, G. An overview of thermal energy storage systems. *Energy* **2018**, *144*, 341–378. [CrossRef]
4. Hao, H.; Sanandaji, B.M.; Poolla, K.; Vincent, T.L. Aggregate Flexibility of Thermostatically Controlled Loads. *IEEE Trans. Power Syst.* **2015**, *30*, 189–198. [CrossRef]
5. Arteconi, A.; Hewitt, N.J.; Polonara, F. State of the art of thermal storage for demand-side management. *Appl. Energy* **2012**, *93*, 371–389. [CrossRef]
6. Kim, Y.; Norford, L.K. Optimal Use of Thermal Energy Storage Resources in Commercial Buildings through Price-Based Demand Response Considering Distribution Network Operation. *Appl. Energy* **2017**, *193*, 308–324. [CrossRef]
7. Hurtado, L.A.; Rhodes, J.D.; Nguyen, P.H.; Kamphuis, I.G.; Webber, M.E. Quantifying Demand Flexibility Based on Structural Thermal Storage and Comfort Management of Non-Residential Buildings: A Comparison between Hot and Cold Climate Zones. *Appl. Energy* **2017**, *195*, 1047–1054. [CrossRef]
8. Finck, C.; Li, R.; Kramer, R.; Zeiler, W. Quantifying Demand Flexibility of Power-to-Heat and Thermal Energy Storage in the Control of Building Heating Systems. *Appl. Energy* **2018**, *209*, 409–425. [CrossRef]
9. Tang, R.; Wang, S. Model Predictive Control for Thermal Energy Storage and Thermal Comfort Optimization of Building Demand Response in Smart Grids. *Appl. Energy* **2019**, *242*, 873–882. [CrossRef]
10. Chen, Y.; Xu, P.; Chen, Z.; Wang, H.; Sha, H.; Ji, Y.; Zhang, Y.; Dou, Q.; Wang, S. Experimental Investigation of Demand Response Potential of Buildings: Combined Passive Thermal Mass and Active Storage. *Appl. Energy* **2020**, *280*, 115956. [CrossRef]
11. Callaway, D.S.; Hiskens, I.A. Achieving Controllability of Electric Loads. *Proc. IEEE* **2011**, *99*, 184–199. [CrossRef]
12. Zhao, L.; Zhang, W.; Hao, H.; Kalsi, K. A Geometric Approach to Aggregate Flexibility Modeling of Thermostatically Controlled Loads. *IEEE Trans. Power Syst.* **2017**, *32*, 4721–4731. [CrossRef]
13. Wang, P.; Wu, D.; Kalsi, K. Flexibility Estimation and Control of Thermostatically Controlled Loads with Lock Time for Regulation Service. *IEEE Trans. Smart Grid* **2020**, *11*, 3221–3230. [CrossRef]
14. Wang, G.; Li, Z.; Wang, F. Enhanced Sufficient Battery Model for Aggregate Flexibility of Thermostatically Controlled Loads Considering Coupling Constraints. *IEEE Trans. Sustain. Energy* **2021**, *12*, 2493–2496. [CrossRef]
15. Zhong, Q.C.; Weiss, G. Synchronverters: Inverters That Mimic Synchronous Generators. *IEEE Trans. Ind. Electron.* **2011**, *58*, 1259–1267. [CrossRef]
16. Schiffer, J.; Zonetti, D.; Ortega, R.; Stanković, A.M.; Sezi, T.; Raisch, J. A Survey on Modeling of Microgrids—From Fundamental Physics to Phasors and Voltage Sources. *Automatica* **2016**, *74*, 135–150. [CrossRef]
17. Arghir, C.; Jouini, T.; Dörfler, F. Grid-Forming Control for Power Converters Based on Matching of Synchronous Machines. *Automatica* **2018**, *95*, 273–282. [CrossRef]
18. Wang, R.; Li, M.J.; Wang, Y.; Sun, Q.; Zhang, P. A study of novel real-time power balance strategy with virtual asynchronous machine control for regional integrated electric-thermal energy systems. *Sci. China Technol. Sci.* **2024**, *67*, 2074–2086. [CrossRef]
19. Wang, Y.; Wang, R.; Li, M.J.; Zhang, P. Damp and droop coefficient stability region analysis for interlinking converters with virtual asynchronous machine controllers in virtual energy storage systems. *J. Energy Storage* **2024**, *99*, 113306. [CrossRef]
20. Wang, Y.; Wang, R.; Sun, Q.; Li, M.J.; Zhang, P.; Wang, P. A Power Balancing Strategy of Virtual Asynchronous Machine/Synchronous Generator for STEM During Full Operating Conditions. *IEEE Trans. Smart Grid* **2025**, *16*, 4395–4407. [CrossRef]
21. Simpson-Porco, J.W.; Shafiee, Q.; Dörfler, F.; Vasquez, J.C.; Guerrero, J.M.; Bullo, F. Secondary Frequency and Voltage Control of Islanded Microgrids via Distributed Averaging. *IEEE Trans. Ind. Electron.* **2015**, *62*, 7025–7038. [CrossRef]

22. Simpson-Porco, J.W. On Stability of Distributed-Averaging Proportional-Integral Frequency Control in Power Systems. *IEEE Control. Syst. Lett.* **2021**, *5*, 677–682. [CrossRef]
23. Dörfler, F.; Simpson-Porco, J.W.; Bullo, F. Breaking the Hierarchy: Distributed Control and Economic Optimality in Microgrids. *IEEE Trans. Control. Netw. Syst.* **2016**, *3*, 241–253. [CrossRef]
24. Trip, S.; Bürger, M.; De Persis, C. An Internal Model Approach to (Optimal) Frequency Regulation in Power Grids with Time-Varying Voltages. *Automatica* **2016**, *64*, 240–253. [CrossRef]
25. Zhang, R.; Wang, R.; Li, M.J.; Sun, Q.; Zhang, P.; Wang, Y.; Wang, P. Distributed Secondary Frequency Restoration Control Based on Virtual Asynchronous Machines for Virtual Energy Storage Systems. *J. Mod. Power Syst. Clean Energy* **2026**, *14*, 945–955. [CrossRef]
26. Ortega, R.; van der Schaft, A.; Maschke, B.; Escobar, G. Interconnection and damping assignment passivity-based control of port-controlled Hamiltonian systems. *Automatica* **2002**, *38*, 585–596. [CrossRef]
27. Ortega, R.; van der Schaft, A.; Castanos, F.; Astolfi, A. Control by Interconnection and Standard Passivity-Based Control of Port-Hamiltonian Systems. *IEEE Trans. Autom. Control.* **2008**, *53*, 2527–2542. [CrossRef]
28. van der Schaft, A.; Jeltsema, D. Port-Hamiltonian Systems Theory: An Introductory Overview. *Found. Trends Syst. Control.* **2014**, *1*, 173–378. [CrossRef]
29. Stegink, T.; De Persis, C.; van der Schaft, A. A Unifying Energy-Based Approach to Stability of Power Grids with Market Dynamics. *IEEE Trans. Autom. Control.* **2017**, *62*, 2612–2622. [CrossRef]
30. Bucolo, M.; Buscarino, A.; Famoso, C.; Fortuna, L. Chaos Addresses Energy in Networks of Electrical Oscillators. *IEEE Access* **2021**, *9*, 153258–153265. [CrossRef]
31. Monshizadeh, P.; Machado, J.E.; Ortega, R.; van der Schaft, A. Power-Controlled Hamiltonian Systems: Application to Electrical Systems with Constant Power Loads. *Automatica* **2019**, *109*, 108527. [CrossRef]