

Temporal synchronization amplifies power-system frequency stress from computational-load flexibility

Wei Tai

Southeast University

Qi Wang

wangqi@seu.edu.cn

Southeast University

Yi Tang

Southeast University

Research Article

Keywords: computational loads, data-centre flexibility, temporal synchronization, frequency stability, inverter-based resources, power-system resilience

Posted Date: October 9th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10951121/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Temporal synchronization amplifies power-system frequency stress from computational-load flexibility

Wei Tai^a, Qi Wang^{b,*}, Yi Tang^b

^a School of Cyber Science and Engineering, Southeast University, Nanjing 210000, China

^b School of Electrical Engineering, Southeast University, Nanjing 210000, China

* Corresponding author. E-mail address: wangqi@seu.edu.cn (Qi Wang).

Abstract

Large computational loads can provide rapid grid flexibility, but closely timed constituent actions may convert that flexibility into a short-duration power imbalance. Here we introduce temporal synchronization as a distinct response dimension and isolate its effect using paired equal-power, equal-energy simulations on the IEEE 39-bus, RTS-GMLC and ACTIVSg2000 systems. Increasing synchronization consistently raised the maximum-frequency and short-window rate-of-change-of-frequency (RoCoF) increments, with nonlinear sensitivity conditioned by response participation, system inertia and inverter-based-resource penetration. Under matched response magnitude, action energy, service deadline and support-resource limits, coordinated timing and fast support reduced the maximum-frequency increment by 81.5–83.0% relative to an equal-resource uncoordinated response, while RoCoF reductions exceeded 75% within each system-specific metric. The effect direction and control ordering persisted across network locations and scales. These results show that computational-load flexibility should be characterized by power, energy and synchronization because temporal organization can materially alter its grid impact.

Keywords. computational loads; data-centre flexibility; temporal synchronization; frequency stability; inverter-based resources; power-system resilience

1. Introduction

Large computational facilities, including artificial-intelligence and cloud data centres, are emerging as a distinct class of bulk-power-system load. Their significance extends beyond rapid growth in electricity consumption^{1–3}. Individual campuses can connect at transmission scale, much of their demand is interfaced through power electronics, and software, protection and backup-power controls can change grid exchange on time scales much shorter than those of conventional industrial loads⁴. This dynamic behaviour has become an explicit reliability concern. In July 2026, the US Federal Energy Regulatory Commission directed the North American Electric Reliability Corporation (NERC) to develop reliability standards for computational-load integration⁵. At the same time, a substantial modelling gap remains: among 33,282 MW of operational data-centre load evaluated by NERC, roughly three quarters lacked a dynamic representation considered adequate for facility behaviour⁶. This leaves a substantial representation gap precisely for the large loads whose dynamic behaviour is becoming increasingly relevant to system operation.

Recent operating events show why this gap matters. During a transmission disturbance in July 2024, customer-side controls transferred approximately 1,500 MW of data-centre load away from the grid following several short voltage depressions. The resulting generation–demand imbalance increased system frequency and local voltage even though the affected load had not been disconnected by the utility⁷. NERC has also highlighted the possibility that interruption of a shared computing process at one facility can initiate closely timed demand reductions at other facilities, including locations that did not directly experience the initiating voltage disturbance⁸.

Such responses can develop within seconds. Their grid impact therefore depends not only on the amount of load that responds, but also on how closely the constituent actions align in time.

The same controllability also makes computational loads potentially valuable sources of grid flexibility^{9, 10}. Workload placement can redistribute demand across locations and scheduling intervals, and server power control can follow fast active-power commands^{11, 12}. Cooling systems provide slower thermal flexibility, while uninterruptible power supplies or other local energy resources can provide short-duration power modulation without interrupting computing service^{13–15}. Existing flexibility studies consequently characterize computational-load response through quantities such as available power, delivered energy, response duration, state of charge, operating cost and service-level constraints^{16–18}. These quantities determine how much flexibility is available and whether a requested service can be completed. Recent work has also developed spatial and temporal representations of workload flexibility for electricity-market and power-system operation^{19–21}. These formulations do not, however, specify how constituent actions aggregate within the short electromechanical response window. Two responses with the same power commitment and the same delivered energy can therefore impose different short-window power imbalances if one is dispersed across many constituent actions and the other is nearly simultaneous.

This distinction is especially relevant in power systems with substantial inverter-based generation. Lower synchronous inertia can accelerate frequency dynamics following an active-power imbalance, while converter control introduces additional dynamic interactions that become increasingly important as inverter-based-resource penetration rises^{22–24}. Grid strength and the location of large electronic loads can further condition their dynamic interaction with the surrounding system²⁵. Dynamic data-centre and large-load models have consequently become increasingly detailed^{26–28}. The Power Electronic Reconnecting and Ceasing (PERC1) load model

captures voltage-dependent active and reactive power, current limiting, cessation and reconnection, while recent electromagnetic-transient data-centre models represent UPS converters, batteries, cooling drives and computing equipment as interacting subsystems^{29–31}. Detailed power-electronic and aggregation models further show that internal converter dynamics and aggregation assumptions can materially affect the represented response^{32, 33}. Such models improve the physical representation of computational facilities, but greater model detail does not by itself identify the effect of response timing. Changes in timing are commonly accompanied by changes in response magnitude, energy, location or control composition, making their individual contributions difficult to separate.

Here we treat temporal synchronization as a distinct dimension of computational-load flexibility, alongside response power and energy. Event synchronization measures the concentration of constituent load actions within a short time window: low synchronization represents temporally dispersed responses, whereas high synchronization represents closely clustered responses. Synchronization can change while the total responding power and delivered response energy remain fixed. This creates a controlled counterfactual in which only the temporal organization of an otherwise equivalent response is changed. We hypothesize that greater synchronization concentrates the grid-facing power imbalance, increasing short-window frequency stress, and that the magnitude of this effect depends on response participation, system inertia and inverter-based-resource penetration. The metric therefore describes the aggregate grid-facing response rather than any particular workload, protection scheme or facility technology.

We test this hypothesis using paired equal-power, equal-energy dynamic simulations in which the computational load is represented by 20 independently timed response blocks. The modified IEEE 39-bus system is used for mechanism identification across response participation,

inertia and inverter-based-resource conditions. A continuous response-surface analysis then characterizes the nonlinear dependence of frequency stress on synchronization and operating conditions, while time-step refinement distinguishes persistent electromechanical effects from numerical artefacts. We further compare timing diversity, fast support and their coordination under matched response magnitude, action energy, service deadline and support-resource limits. Finally, the direction of the synchronization effect and the B6–B5 control ordering are tested across three electrically distinct locations in Reliability Test System–Grid Modernization Lab Consortium (RTS-GMLC) and in the unreduced ACTIVSg2000 network.

Increasing synchronization raises both the maximum-frequency and short-window rate of change of frequency (RoCoF) increments even when response magnitude and energy are unchanged, with a nonlinear sensitivity conditioned by the operating state. Timing diversity reduces this stress, and combining dispersed response timing with fast support substantially outperforms an equal-resource uncoordinated response without relaxing the computing-service requirement. The direction of the synchronization effect and the control ordering persist across the three network models, although their absolute sensitivities remain system dependent. Computational-load flexibility should therefore be characterized not only by how much power can respond and how much energy must be delivered, but also by how tightly that response is concentrated in time.

2. Results

The study design links block-level response timing to network-scale frequency stress and subsequently tests whether this relationship can be shaped by coordinated control and transferred across systems (Fig. 1).

2.1. Temporal synchronization amplifies short-window frequency stress at fixed response

energy

We first isolated the effect of response synchronization while holding the amount of computational flexibility fixed within each comparison. The modified IEEE 39-bus system contained an 828-MW computational facility at bus 39 and four structural operating profiles, combining inverter-based resource (IBR) penetrations of 50% and 65% with system inertia scales of 0.25 and 0.35. Within each profile, either 75% or 90% of the facility participated, corresponding to responses of 621.0 and 745.2 MW. Each response was distributed across 20 computational blocks, delivering 124.20 or 149.04 MW s, respectively.

Synchronization was varied from 0.10 to 1.00 at seven prespecified levels while response amount, duration and energy remained fixed (Fig. 1). Eight timing seeds were shared across the complete factorial design, producing 448 action trajectories and four matched zero-action controls. The timing seed served as the independent sampling unit. Realized synchronization and response energy matched their prescribed values to numerical precision, and the maximum network power-balance residual was 9.78×10^{-9} p.u.

Increasing synchronization consistently increased both the maximum-frequency and mechanism-window maximum-RoCoF increments (Fig. 2a,b), despite identical response amount and energy within each paired comparison. Both outcomes were non-decreasing across all 64 combinations of structural profile, response participation and timing seed. The fully synchronized minus low-synchronization frequency effect was positive in all eight structural-participation cells, ranging from 0.159 Hz in the higher-inertia, higher-IBR, 75% participation condition to 0.267 Hz in the lower-inertia, 50% IBR, 90% participation condition. The corresponding 95% paired-bootstrap confidence intervals were 0.1587–0.1592 Hz and 0.2662–0.2670 Hz, respectively (Fig. 2c).

RoCoF showed the same ordering. Fully synchronized actions increased maximum RoCoF by 1.02–2.01 Hz s⁻¹ across the eight cells. The corresponding 95% confidence intervals were 1.020–1.027 Hz s⁻¹ for the smallest effect and 2.002–2.015 Hz s⁻¹ for the largest (Fig. 2d). Because each contrast preserved facility location, response amount, action energy, structural profile, disturbance and timing seed, the observed increase arose from temporal concentration rather than additional response capacity.




Fig. 1 | Temporal synchronization shapes the grid impact of computational-load flexibility. **a**, Temporal synchronization as a dimension of computational-load flexibility. **b**, Equal response capacity and energy, different temporal organization. **c**, Temporal concentration shapes aggregate power and frequency stress. **d**, Study framework from mechanism identification to cross-system validation.




Fig. 2 | Temporal synchronization amplifies short-window frequency stress at fixed response energy. **a**, representative low- and high-synchronization action schedules and frequency trajectories. **b**, corresponding RoCoF trajectories. **c,d**, paired high-minus-low effects across the eight IEEE 39 operating cells, showing timing-seed estimates, means and 95% paired-bootstrap confidence intervals.

2.2. Synchronization produces a nonlinear stress surface conditioned by operating state

The effect of synchronization was monotonic but not linear. Frequency and RoCoF increments increased more rapidly as the computational actions became more concentrated, with the curvature varying across operating conditions (Fig. 3a,b). This increase was strongest under high response participation and low system inertia. At 50% IBR penetration, an inertia scale of 0.25 and 90% participation, increasing synchronization from 0.10 to 0.25 raised the maximum-frequency increment by 0.0382 Hz, whereas the same 0.15 increase from 0.85 to 1.00 raised it by 0.0538 Hz. The corresponding maximum-RoCoF increases were 0.303 and 0.359 Hz s⁻¹.

We quantified this dependence using response surfaces fitted to the 448 baseline-adjusted trajectories. Synchronization, inertia, IBR penetration and response participation were centered and scaled before fitting. The model included first-order interactions between synchronization and each operating variable, together with a quadratic synchronization term. At the center of the operating domain, the estimated quadratic coefficient was 0.0113 Hz for maximum frequency (95% cluster-bootstrap confidence interval, 0.0111–0.0114 Hz). The corresponding coefficient for RoCoF was 0.0443 Hz s⁻¹ (0.0408–0.0481 Hz s⁻¹; Fig. 3c).

The fitted interactions showed how operating conditions changed sensitivity to synchronization. The synchronization-by-participation interaction was positive for both outcomes, whereas the synchronization-by-inertia interaction was negative. Greater response participation therefore increased the additional stress associated with synchronization, while higher synchronous inertia reduced it. The synchronization-by-IBR interaction was also negative under the modeled converter composition. This result indicates a smaller incremental effect of synchronization at higher IBR penetration in the tested conditions; it does not imply lower overall disturbance severity.

179 Leave-one-seed-out validation showed that the nonlinear relationship was retained for
180 held-out timing realizations. The quadratic surface reduced maximum-frequency prediction error
181 from 0.00481 to 0.00186 Hz, a 61.4% improvement over the linear model. For maximum RoCoF,
182 the error decreased from 0.0294 to 0.0235 Hz s⁻¹, a 20.0% improvement. Both reductions exceeded
183 the prespecified 2% model-selection margin. The corresponding out-of-seed coefficients of
184 determination were 0.9993 and 0.9980 (Fig. 3d,e).

185 The response remained continuous over the sampled synchronization range. All 448 action
186 trajectories recovered, and none activated the overfrequency timer or secondary protection. The
187 results therefore support a continuous nonlinear stress surface rather than a critical synchronization
188 threshold within the tested operating domain. Its slope and curvature depended on response
189 participation and grid conditions.




Fig. 3 | Operating conditions shape the nonlinear synchronization–stress relationship.

a,b, maximum-frequency and maximum-RoCoF surfaces across synchronization, response participation, inertia and IBR penetration. **c**, coefficient forest plot for curvature and interaction terms with cluster-bootstrap confidence intervals. **d**, linear-versus-quadratic leave-one-seed-out errors. Panel **e**, observed-versus-predicted outcomes for held-out timing seeds.

2.3. The short-window synchronization effect persists under numerical refinement

Coarse-grid screening suggested a synchronization-dependent recovery boundary, motivating a time-step refinement test. We repeated 126 action trajectories and three matched controls at 10, 5 and 2.5 ms using six paired timing seeds. Each schedule was constructed in continuous time and averaged over the corresponding integration intervals, preserving response energy to numerical precision.

The apparent recovery boundary was not reproduced on the refined grids. At 10 ms, all 42 action trajectories lost network convergence between 10.3 and 10.6 s, and the matched zero-action control also failed at 10.6 s. By contrast, every action and control trajectory at 5 and 2.5 ms reached 15 s, with identical recovery classifications at both 9 and 15 s. The maximum paired difference in the 15-s terminal margin was 0.00988, below the prespecified tolerance of 0.02 (Fig. 4a–c).

The short-window synchronization effect remained stable across the refined grids. The largest paired difference in mechanism-window RoCoF between 5 and 2.5 ms was 8.02×10^{-4} Hz s^{-1} , below the 1.0×10^{-3} Hz s^{-1} convergence criterion. At 30% participation, the high-minus-low contrast was 0.0165 Hz s^{-1} at 5 ms and 0.0166 Hz s^{-1} at 2.5 ms; at 45% participation, the corresponding values were 0.0295 and 0.0296 Hz s^{-1} (Fig. 4d). Numerical refinement therefore preserved the short-window synchronization effect while removing the apparent coarse-grid recovery boundary. Subsequent experiments used the 2.5-ms grid.




Fig. 4 | Short-window synchronization effects persist under numerical refinement. a,

Representative trajectories at 10, 5 and 2.5 ms. **b**, Simulation completion and recovery across step size, response participation and synchronization. **c**, Paired 15-s terminal-margin differences between the 5- and 2.5-ms grids, with the prespecified tolerance. **d**, High-minus-low synchronization effects on mechanism-window RoCoF across time steps, with individual timing-seed estimates and the convergence criterion.

2.4. Joint timing and fast support reduce frequency stress under equal resource budgets

We next tested whether synchronization-aware coordination could reduce frequency stress without relaxing the computing-service requirement. Eight strategies were designed to separate synchronized response, timing diversity, load-side flexibility, fast support, joint coordination and oracle scheduling (Fig. 5a). Across the active strategies, the requested response, net action energy and 6.5-s service deadline were identical within each cell-seed pair. All support-enabled strategies

shared the same support-power and buffer-energy limits. Eight new timing seeds were applied to each strategy under high- and moderate-stress conditions, producing 128 trajectories.


In the high-stress condition, B2, which used timing diversity without fast support, produced maximum-frequency and maximum-RoCoF increments of 0.0624 Hz and 0.396 Hz s⁻¹, respectively. B3 retained a fully synchronized load response but added fast support, resulting in larger increments of 0.238 Hz and 1.75 Hz s⁻¹. The uncoordinated combination B5 produced increments of 0.145 Hz and 1.05 Hz s⁻¹, whereas joint coordination under B6 reduced them to 0.0246 Hz and 0.112 Hz s⁻¹ (Fig. 5b,c). Relative to B5, the paired B6 reductions were 0.120 Hz (95% paired-bootstrap confidence interval, 0.1197–0.1206 Hz) and 0.943 Hz s⁻¹ (0.936–0.950 Hz s⁻¹).

The same strategy ordering held under moderate stress. Maximum-frequency increments were 0.0402 Hz for B2, 0.159 Hz for B3 and 0.0935 Hz for B5, compared with 0.0160 Hz for B6. Maximum RoCoF was 0.0628 Hz s⁻¹ under B6 and 0.617 Hz s⁻¹ under B5. The paired B6-minus-B5 differences were -0.0775 Hz (95% confidence interval, -0.0777 to -0.0772 Hz) and -0.554 Hz s⁻¹ (-0.558 to -0.550 Hz s⁻¹). Across both conditions, B6 reduced the maximum-frequency increment by approximately 60% relative to B2, 90% relative to B3, 81% relative to B4 and 83% relative to B5 (Fig. 5d).

The B5–B6 comparison isolated the benefit of coordination from resource quantity. Both strategies had identical support-power limits of 112 MW under high stress and 93.2 MW under moderate stress, with corresponding buffer-energy limits of 29.8 and 24.8 MW s. Their realized buffer-energy exchanges differed by less than 2.5%, yet B6 reduced the maximum-frequency increment by approximately 83% and the RoCoF increment by approximately 90%. The gain

therefore came from reshaping the net response within the same resource envelope rather than from additional support capacity.

The oracle strategy B7 further reduced the maximum-frequency increment relative to B6 by 0.0158 Hz under high stress and 0.0103 Hz under moderate stress, indicating that B6 did not attain the idealized performance bound. Every active response strategy delivered the requested computational action by the common deadline, and all 128 trajectories recovered. Together, these comparisons show that coordinating response timing with fast support can substantially reduce frequency stress without increasing the available support resources.



c Resource and service-budget parity

	B0	B1	B2	B3	B4	B5	B6	B7
Response amount / reference	0	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Net action energy / reference	0	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Service deadline (s)	—	6.5	6.5	6.5	6.5	6.5	6.5	6.5
Fast-support cap / P _{action}	0	0	0	0.15	0	0.15	0.15	0.15
Buffer cap / P _{action} (s)	0	0	0	0.04	0	0.04	0.04	0.04

Within-cell normalization; outline links B5 and B6 with identical budgets.




Fig. 5 | Joint coordination reduces frequency stress under equal resource budgets. a,

Strategy-resource matrix for B0–B7. **b,c**, Maximum-frequency and maximum-RoCoF increments with individual timing-seed estimates and 95% paired-bootstrap confidence intervals. **d**, Paired B6-minus-comparator effects relative to B2–B5. **e**, B6 reductions at the three RTS-GMLC sites. The resource strip shows parity in response amount, net action energy, service deadline and support limits for the relevant matched comparisons.

2.5. Synchronization effects and control ordering persist across test systems

We next tested whether the synchronization effect and B6–B5 control ordering transferred beyond IEEE 39. A 100-MW computational facility was evaluated at three sites in the 73-bus RTS-GMLC network and at bus 7229 in the unreduced ACTIVSg2000 network (Fig. 6). RTS buses 104, 203 and 319 represented low-, medium- and high-strength locations, respectively. Comparisons used eight paired timing seeds per RTS site and four in ACTIVSg2000.

Representative trajectories illustrate how these effects developed in time (Fig. 7a). In each system, high synchronization produced a sharper initial frequency excursion than low synchronization, while B6 reduced and spread the initial response relative to B5. Each trajectory was referenced to its matched zero-action baseline; statistical comparisons used all paired timing seeds.

High-minus-low synchronization effects remained positive across all 12 tested conditions (Fig. 7b). Mean maximum-frequency effects ranged from 159 to 267 mHz across the eight IEEE conditions. At RTS buses 104, 203 and 319, the corresponding effects were 10.9, 9.69 and 13.1 mHz. Their magnitudes did not follow the ordering of the local-strength proxy. ACTIVSg2000 also retained a positive effect of 1.08 mHz (95% paired-bootstrap confidence interval, 1.06–1.09 mHz).

282 The B6–B5 ordering also persisted across systems (Fig. 7c). Within each comparison, B5
283 and B6 had identical response amount, net action energy, service deadline, support-power limit
284 and buffer-energy limit. In RTS-GMLC, these were 90 MW, 18 MW s and a 6.5-s deadline, with
285 support limits of 13.5 MW and 3.6 MW s. Across the six cross-system control conditions, B6
286 retained only 17.0–18.5% of the B5 maximum-frequency increment, corresponding to reductions
287 of 81.5–83.0%.

288 RoCoF showed the same B6–B5 ordering under the experiment-specific metric definitions.
289 B6 reduced the mechanism-window RoCoF increment by 89.4–89.8% in IEEE 39. For action-
290 attributable RoCoF, reductions were 76.0–80.3% across the RTS sites and 77.8% in
291 ACTIVSg2000. In ACTIVSg2000, the maximum-frequency increment decreased from 0.665 to
292 0.123 mHz, an 81.5% reduction. Because the mechanism and transfer experiments used different
293 RoCoF definitions, these magnitudes are reported separately.

294 Across the three test systems, the direction of the synchronization effect and the B6–B5
295 control ordering remained consistent despite substantial differences in absolute effect magnitude.
296 The transferable result is therefore the synchronization mechanism and control ordering, rather
297 than a common numerical threshold across systems.




Fig. 6 | Cross-system validation design for IEEE 39, RTS-GMLC and ACTIVSg2000. a,b,c, The three systems support mechanism identification, location transfer and network-scale validation, respectively, with their network size, computational-facility rating and validation role shown for each case.




Fig. 7 | Synchronization stress and coordinated attenuation across test systems. a,

Representative frequency trajectories relative to matched zero-action baselines for IEEE 39, RTS-GMLC bus 104 and ACTIVSg2000 bus 7229. Grey dashed and ochre curves denote low and high synchronization; blue and teal denote B5 and B6. Shading marks the 1.00–1.32-s mechanism window. **b,** Paired high-minus-low synchronization effects across all 12 tested conditions. Points show condition means; 95% confidence intervals are provided in Source Data. **c,** Residual B6 stress relative to the equal-budget B5 reference, expressed as a percentage of the B5 mean. Fine ticks show seed-paired ratios and error bars show 95% percentile-bootstrap intervals. The B5

reference is 100%; lower values indicate greater attenuation. Eight paired seeds were used per IEEE and RTS condition and four for ACTIVSg2000. RoCoF follows the experiment-specific definitions in Methods.

3. Discussion

This study identifies temporal synchronization as a distinct dimension of computational-load flexibility that is not captured by response capacity and delivered energy alone. At fixed response amount and energy, increasing synchronization consistently increased the maximum-frequency and short-window RoCoF increments, with a nonlinear sensitivity shaped by response participation, inertia and IBR penetration. Joint timing and fast-support coordination substantially reduced these effects without relaxing the computing-service deadline or increasing the available support resources. The grid value of computational-load flexibility therefore depends not only on how much power and energy can be shifted, but also on how that response is organized in time.

The mechanism arises from the aggregation of many individually admissible actions. A single server block or power-conversion module may contribute little to system imbalance, but aligned protection thresholds, control settings or supervisory commands can concentrate many such actions into a short-duration power impulse at the point of common coupling. The grid initially responds through the same inertial and converter-mediated channels that govern other rapid active-power imbalances. This accounts for the stronger RoCoF response at higher synchronization and for its amplification as response participation increases and inertia decreases. The relevant quantity is therefore not only the flexible power available, but the fraction of that response concentrated within the grid's short dynamic window.

The IBR interaction also shows that synchronization sensitivity cannot be inferred from IBR penetration alone. Under the modeled grid-forming share and converter controls, higher IBR

penetration modestly reduced the incremental high-minus-low synchronization effect, while disturbance-only stress increased in parts of the operating space. Inertia, grid-forming composition, current-limiting behavior and disturbance location therefore remain separate determinants of the dynamic response. Computational-load interconnection studies should report these quantities alongside IBR penetration when assessing sensitivity to synchronized actions.

Large computational facilities occupy two contrasting roles in the power system. Large computational facilities can provide rapid flexibility through IT workload control, uninterruptible power supplies, cooling systems and converter interfaces. The same class of facilities has also exhibited large, near-simultaneous load reductions during transmission disturbances, with measurable frequency and voltage consequences. Current reliability assessments further report active-load reductions, oscillatory responses and deficiencies in large-load dynamic models. The distinction between flexibility service and additional disturbance therefore lies not only in nominal response capacity, but also in the temporal organization of the response.

Purely aggregate load representations miss part of this timing structure. A constant-impedance, constant-current and constant-power (ZIP) load model captures voltage-dependent operating behavior but cannot represent event timing across computational blocks, supervisory commands or heterogeneous ride-through states. PERC1-type models add ceasing, reconnection and power-electronic response logic, providing a more suitable basis for large computational loads. Even so, a single aggregate parameterization can impose artificial simultaneity when internal timing diversity is omitted. The synchronization metric introduced here links device-level timing to system-level stress without requiring every server, converter or cooling device to be represented explicitly in the transmission model.

Synchronization is also a controllable design variable. Timing diversity alone reduced stress relative to a fully synchronized response, whereas fast support was less effective when the underlying load action remained concentrated. B6 combined timing diversity with causal, energy-neutral shaping at the point of common coupling and outperformed the equal-resource B5 strategy across all tested systems. In the IEEE 39 conditions, B6 reduced the maximum-frequency increment by about 83% and the RoCoF increment by about 90% relative to B5; maximum-frequency reductions remained about 82% across the RTS-GMLC sites and the ACTIVSg2000 operating condition.

A practical implementation could treat synchronization as an operational constraint alongside response capacity, energy and completion time. Facility controllers could diversify ride-through delays, recovery ramps and workload transitions across independently managed blocks, while a supervisory layer limits the short-window aggregate power peak using local buffering³⁴.³⁶ The transmission operator would not need visibility of every internal action; instead, the facility could expose an admissible response envelope defined by power, energy, deadline and maximum temporal concentration³⁷. Such an interface would distinguish, for example, a 500-MW response delivered progressively from the same nominal response delivered as a near-simultaneous step.

Cross-system validation preserved the direction of the effect but not its absolute magnitude. The synchronization effect retained the same direction at three electrically distinct RTS-GMLC locations and in the unreduced 2,000-bus ACTIVSg2000 network, while B6 consistently outperformed the equal-resource B5 strategy. Absolute effects varied with facility-to-system size, topology, disturbance structure and dynamic composition, and were not ordered by local strength alone. The transferable result is therefore the synchronization mechanism and control ordering rather than a common numerical threshold.

381 The time-step study also changes how the apparent recovery boundary should be
382 interpreted. The apparent recovery boundary observed at 10 ms disappeared at 5 and 2.5 ms,
383 whereas the short-window synchronization effect remained stable. The concurrent failure of the
384 10-ms zero-action control further indicated that the late boundary was numerical rather than
385 physical. More generally, recovery and protection claims in fast computational-load studies should
386 be supported by time-step refinement, matched controls and simulation horizons appropriate to the
387 dynamics being interpreted.

388 The present evidence is limited to positive-sequence electromechanical simulations with
389 aggregate grid-following, grid-forming and computational-load dynamics. Switching-frequency
390 interactions, phase imbalance and manufacturer-specific converter controls are outside this model
391 class, and fast support is represented as causal, energy-neutral net-power shaping rather than a
392 specific UPS, battery or grid-forming implementation³⁸. A natural next step is EMT validation
393 using more detailed facility models, including the published ERCOT data-centre representation,
394 to assess event ordering, current limiting, state-of-charge feasibility and recovery timing. Field-
395 informed parameters would then be required to translate the synchronization-stress relationship
396 into interconnection-specific operating envelopes.

397 A practical design principle follows: computational-load flexibility should be specified by
398 power, energy and synchronization. Grid studies should specify not only how many megawatts
399 can respond and how much energy must be delivered or recovered, but also how tightly those
400 actions may cluster in time. Facility controls can then coordinate internal timing diversity and
401 short-window buffering to satisfy an explicit synchronization constraint. This view treats large
402 computational loads as controllable dynamic participants whose grid contribution depends on the
403 joint management of response magnitude, energy and temporal concentration.

4. Methods

4.1. Study design and test systems

We used three benchmark transmission systems with complementary validation roles. The modified IEEE 39-bus system supported the complete factorial experiment for mechanism identification, RTS-GMLC tested transfer across three electrically distinct computational-load locations, and the unreduced ACTIVSg2000 system provided a targeted full-network evaluation^{39–41}. Table 1 summarizes the operating conditions, initiating disturbances and validation role of each system.

Each operating condition was initialized independently using an alternating-current optimal power flow followed by an alternating-current power flow. Initial states were retained only when generator limits, branch ratings and voltage bounds were satisfied and the normalized current-balance residual was below 10^{-6} .

The IEEE 39-bus system contained an 828-MW computational load at bus 39, corresponding to 75% of the original demand at that bus. The resistance and reactance of branches incident on bus 39 were multiplied by 1.5, increasing the local network impedance while preserving an admissible operating point. Four structural profiles combined IBR penetration, denoted by η_{IBR} , at values of 0.50 and 0.65 with system inertia scales of 0.25 and 0.35. Grid-forming (GFM) resources accounted for 10% of inverter-based generation in each profile. Computational-load response participation, m_{resp} , defined as the fraction of the pre-event computational load participating in the response, was 0.75 or 0.90, producing eight structural-participation cells.

RTS-GMLC received a 100-MW computational load at bus 104, 203 or 319 in separate experiments. The three sites occupied different areas and represented low, intermediate and high

pre-event short-circuit capacities of approximately 1.24, 2.02 and 4.23 GVA, respectively. They were selected from 27 load-bus candidates after operating-point and branch-contingency screening. Because no single branch outage was admissible at all three sites in the supplied network snapshot, the most heavily loaded admissible inter-area branch was selected separately for each site using the same prespecified rule.

The ACTIVSg2000 study retained all 2,000 buses and 3,206 branches. A 100-MW computational load with a power factor of 0.96 was connected at external bus 7229. Online generators connected to the same bus were aggregated into 392 dynamic equivalents while retaining the full network equations. One prespecified operating condition was evaluated to test whether the synchronization effect and B6–B5 control ordering persisted at network scale, rather than to reproduce the complete IEEE 39 factorial experiment.

A timing seed defines one paired realization of block-action times and serves as the independent statistical replicate. The inertia scale is defined relative to the all-synchronous reference state of the corresponding network. Computational-load capacities are study parameters and do not represent specific operating facilities.

Table 1 | Test systems, disturbances and validation roles

System	Computational-load configuration	Dynamic structure	Initiating disturbance	Sampling and numerical setting	Validation role
IEEE 39-bus, 39 buses and 46 branches	828 MW at bus 39; $m_{\text{resp}} \in \{0.75, 0.90\}$	$\eta_{\text{IBR}} \in \{0.50, 0.65\}$; inertia scale $\in \{0.25, 0.35\}$; GFM share within IBR = 0.10	Fault at bus 15; $Z_f = j0.10$ p.u.; duration 0.18 s; branch 24 opened at clearing	Seven synchronization levels; eight timing seeds per cell; 2.5-ms time step; 8-s horizon	Mechanism identification, conditional response-surface analysis and complete B0–B7 control comparison
IEEE 39-bus refinement set	Same model family; $m_{\text{resp}} \in \{0.30, 0.45\}$	Fixed structural condition	Fault reactance 0.20 p.u.; remaining disturbance settings matched across time steps	Seven participation–synchronization conditions; six timing seeds; 10-, 5- and 2.5-ms time steps; 15-s horizon	Numerical convergence and recovery audit
RTS-GMLC, 73 buses and 120 branches	100 MW at buses 104, 203 and 319 in separate placements; $m_{\text{resp}} = 0.90$	$\eta_{\text{IBR}} = 0.50$; inertia scale = 0.35; GFM share within IBR = 0.10	Site 104: fault bus 325, branch 118; site 203: fault bus 113, branch 24; site 319: fault bus 325, branch 118; $Z_f = j0.15$ p.u.; duration 0.12 s	Eight timing seeds per site; 2.5-ms time step; 8-s horizon	Multi-site transfer of the synchronization effect and B6–B5 control ordering
ACTIVSg 2000, 2,000 buses and 3,206 branches	100 MW at bus 7229; $m_{\text{resp}} = 0.90$	$\eta_{\text{IBR}} = 0.50$; inertia scale = 0.35; GFM share within IBR = 0.10	Fault at bus 7017; $Z_f = j0.15$ p.u.; duration 0.12 s; branch 2228 opened at clearing	Four timing seeds; 2.5-ms time step; 7-s horizon	Unreduced full-network validation of the synchronization effect and B6–B5 ordering

IBR, inverter-based resource; GFM, grid-forming; Z_f , fault impedance; j , imaginary unit.

4.2. Computational-load representation and synchronization metric

The facility model represents three contributions at the point of common coupling (PCC): baseline demand, block-level computational-load action and short-duration fast support. At time t , the grid-import active power $P_{\text{PCC}}(t)$ is

$$P_{\text{PCC}}(t) = P_{\text{CL}}^0 - p_e(t) + u_f(t) \quad (1)$$

where P_{CL}^0 is the pre-event computational load, $p_e(t) \geq 0$ is the gross load reduction associated with computational-load action event e , and $u_f(t)$ is the fast-support action used to offset that reduction at the PCC. Positive $u_f(t)$ therefore reduces the net change in grid import. The modeled facility does not export active power to the external network.

The grid-facing computational load was divided into $N = 20$ equal response blocks, where N denotes the number of blocks. This representation extends an aggregate PERC1-type action by retaining the timing of each constituent response. In an operating facility, these response times may arise from protection logic, workload control or supervisory commands. In the paired experiments, block-action times were prescribed directly so that temporal concentration could be varied without changing response magnitude, duration, energy, location or the initiating disturbance.

For computational-load action event e and response block $i \in \{1, \dots, N\}$, let $a_{i,e} > 0$ and $t_{i,e}$ denote the response power and start time of block i , respectively. The total response magnitude A_e and response-participation fraction $m_{\text{resp},e}$ are

$$A_e = \sum_{i=1}^N a_{i,e}, \quad m_{\text{resp},e} = \frac{A_e}{P_{\text{CL}}^0} \quad (2)$$

Event synchronization, $\rho_{\text{sync},e}(\Delta)$, was defined as the largest power-weighted fraction of the total response whose block actions begin within any window of duration Δ :

$$\rho_{\text{sync},e}(\Delta) = \max_{\tau} \frac{\sum_{i=1}^N a_{i,e} \mathbf{1}(\tau \leq t_{i,e} < \tau + \Delta)}{A_e} \quad (3)$$

where τ denotes the candidate start time of the synchronization window and $\mathbf{1}(\cdot)$ is the indicator function, equal to one when its argument is true and zero otherwise. The synchronization-

window duration was $\Delta = 0.20$ s, equal to the duration of one block response. For equal block powers, ρ_{sync} ranges from $1/N$ to 1. Because a common scaling of all $a_{i,e}$ leaves ρ_{sync} unchanged, the metric captures temporal concentration independently of response magnitude.

Each block followed a rectangular response of duration Δ , giving the aggregate gross load-reduction trajectory

$$p_e(t) = \sum_{i=1}^N a_{i,e} \mathbf{1}(t_{i,e} \leq t < t_{i,e} + \Delta) \quad (4)$$

For equal-duration, same-direction actions, equations (3) and (4) give

$$\max_t p_e(t) = A_e \rho_{\text{sync},e}(\Delta), \quad \int_{-\infty}^{\infty} p_e(t) dt = A_e \Delta \quad (5)$$

Thus, changing block-action times changes the short-window peak power while preserving both the total response magnitude A_e and the delivered response energy $A_e \Delta$.

The primary experiments evaluated seven synchronization levels, $\rho_{\text{sync}} \in \{0.10, 0.25, 0.40, 0.55, 0.70, 0.85, 1.00\}$. For each target level, the required number of blocks was placed within a compact timing cluster, while the remaining blocks were separated by more than one response duration. The same timing seed was reused across synchronization levels within each operating cell, preserving the paired comparison. As a waveform-shape check, the schedule construction was repeated with trapezoidal actions and confirmed the same response-power and energy invariants.

In the IEEE factorial experiment, response-participation fractions of 0.75 and 0.90 corresponded to response magnitudes of 621 and 745.2 MW, respectively, with action energies of 124.2 and 149.04 MW s. Crossing four structural profiles, two participation levels, seven synchronization levels and eight timing seeds produced 448 action trajectories. Four zero-action controls, one for each structural profile, provided the matched disturbance references.

4.3. Positive-sequence dynamic model

The transmission system was represented by a positive-sequence differential–algebraic model:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{V}, p_{\text{net}}, \mathbf{d}), \quad \mathbf{0} = \mathbf{g}(\mathbf{x}, \mathbf{V}, p_{\text{net}}, \mathbf{d}) \quad (6)$$

Here, $\mathbf{x}$ is the vector of dynamic states containing the synchronous-machine, governor, grid-following and grid-forming states; $\mathbf{V}$ is the complex bus-voltage vector; $\mathbf{f}(\cdot)$ represents the differential-state dynamics; $\mathbf{g}(\cdot)$ represents the algebraic network-balance equations; $\mathbf{0}$ is the corresponding zero-residual vector; $p_{\text{net}} = p_e - u_f$ is the net computational-load action at the PCC; and $\mathbf{d}$ represents the initiating fault and post-fault topology. An overdot denotes differentiation with respect to time.

Each synchronous generator was represented by a voltage source behind transient reactance, a swing equation and a first-order governor:

$$\begin{cases} \dot{\delta}_g = \omega_b \omega_g \\ M_g \dot{\omega}_g = P_{m,g} - P_{e,g} - D_g \omega_g \\ T_g \dot{P}_{m,g} = P_g^{\text{ref}} - K_g \omega_g - P_{m,g} \end{cases} \quad (7)$$

Here, g indexes synchronous generators; δ_g and ω_g are the rotor angle and per-unit speed deviation of generator g , respectively; ω_b is the base angular frequency; $P_{m,g}$ and $P_{e,g}$ are the mechanical and electrical powers; M_g and D_g are the inertia and damping coefficients; T_g is the governor time constant; K_g is the governor speed-feedback gain; and P_g^{ref} is the reference mechanical-power input.

The reference IEEE machine inertia constants were 4.5, 4.0, 3.8, 3.6, 3.5, 3.2, 3.0, 3.5, 4.2 and 5.0 s before structural scaling. Transient reactance was 0.25 p.u. on each machine base.

Damping was 1.0 p.u., the nominal governor droop was 0.05 p.u., and the governor time constant was 0.35 s. RTS-GMLC and ACTIVSg2000 used a 4.0-s reference inertia before the prescribed system-level scaling.

System frequency was obtained from the inertia-weighted centre-of-inertia (COI) speed.

The COI speed deviation was defined as $\omega_{\text{COI}}(t) = \sum_g M_g \omega_g(t) / \sum_g M_g$, giving

$$f_{\text{COI}}(t) = f_0 [1 + \omega_{\text{COI}}(t)], \quad f_0 = 60\text{Hz} \quad (8)$$

where $f_{\text{COI}}(t)$ is the COI frequency and f_0 is the nominal system frequency.

Over time scales shorter than the dominant governor and converter power-control responses, the aggregate frequency balance can be approximated by $M_{\text{sys}} d\Delta f / dt + D_{\text{sys}} \Delta f = p_{\text{net}} - p_g$, where Δf is the system-frequency deviation from nominal, M_{sys} and D_{sys} are the aggregate inertia and damping coefficients, respectively, and p_g denotes the generation-side active-power response. Over the initial response window, combining this balance with equation (5) gives the short-window RoCoF approximation

$$\max |\text{RoCoF}| \approx \frac{A_e \rho_{\text{sync}}}{M_{\text{sys}}} \quad (9)$$

This approximation shows why the short-window response depends jointly on synchronization, response magnitude and system inertia. All reported outcomes were obtained from the full nonlinear network simulation rather than from equation (9).

Grid-following (GFL) resources used reduced phase-tracking, frequency–watt and voltage–var dynamics. The phase-tracking and frequency-state time constants were 0.03 and 0.04 s, respectively, and the active- and reactive-power time constants were 0.08 and 0.05 s. The active-power target was $P_0 - 3\omega_{\text{COI}}$, and the reactive-power target was $Q_0 + 1.5(V_0 - |V|)$, where P_0 and

Q_0 are the pre-disturbance active- and reactive-power operating points, respectively, V_0 is the pre-disturbance terminal-voltage magnitude, and $|V|$ is the instantaneous terminal-voltage magnitude. The resulting current command was constrained by a circular limit of 1.15 times the initial current. Active current was prioritized under normal voltage conditions, with reactive-current priority below 0.90 p.u.

Grid-forming (GFM) resources used a swing-like active-power law with a virtual inertia constant of 4.0 s and frequency damping of 30 p.u. Their internal-voltage target followed reactive-power droop with a gain of 0.08 and a time constant of 0.08 s, subject to an internal-voltage range of 0.85–1.15 p.u. Converter current was limited through a dynamic virtual reactance:

$$X_{\text{eff}} = \max \left(X_v, \frac{|E - V|}{I_{\text{max}}} \right) \quad (10)$$

where X_{eff} is the effective virtual reactance, X_v is its nominal value, E and V are the converter internal- and terminal-voltage phasors, respectively, and I_{max} is the converter current-magnitude limit. This formulation retains a continuous voltage-source representation while enforcing the current limit.

The remaining demand followed a ZIP composition of 0.20 constant impedance, 0.30 constant current and 0.50 constant power. At bus k , let $S_{L,k}$ and $S_{L,k}^0$ denote the instantaneous and pre-disturbance complex load demands, respectively, and let V_k and V_k^0 denote the corresponding instantaneous and pre-disturbance bus-voltage phasors. Defining the voltage-magnitude ratio as $v_k = |V_k| / |V_k^0|$, the complex demand was

$$S_{L,k} = S_{L,k}^0 \left(0.20v_k^2 + 0.30v_k + 0.50 \right) \max(0.5, 1 + \omega_{\text{COI}}) \quad (11)$$

Computational-load actions reduced active and reactive demand at a fixed power factor of 0.96. At each time step, the algebraic network equations included the ZIP demand, converter

current limits, synchronous-machine Norton equivalents, active fault admittance and the current network topology.

The IEEE 39-bus simulations retained two secondary protection channels so that additional switching could arise from the simulated dynamics. The monitored-line relay used a 1.0-p.u. pickup with a 1.0-s persistence delay. The overfrequency relay disconnected the grid-following resources when frequency remained above 60.5 Hz for 0.15 s. Secondary protection was disabled in the RTS-GMLC and ACTIVSg2000 transfer experiments, which evaluated electromechanical response under fixed post-fault topologies.

4.4. Equal-budget coordination

Fast support was represented by an energy-neutral facility buffer that reshaped the scheduled computational-load response at the PCC. The net load-reduction trajectory and buffer-energy state were

$$p_{\text{net}}(t) = p_e(t) - u_f(t), \quad \dot{E}_b(t) = u_f(t) \quad (12)$$

where $E_b(t)$ denotes the buffer-energy state. The buffer was subject to

$$|u_f(t)| \leq P_f^{\max} = 0.15A_e, \quad 0 \leq E_b(t) \leq E_b^{\max} = 0.04sA_e, \quad E_b(0) = E_b(T_d) = 0 \quad (13)$$

where $P_f^{\max}$ is the support-power limit, $E_b^{\max}$ is the buffer-energy limit, and $T_d = 6.5$ s is the common service deadline. The buffer-energy state E_b is measured in MW s. Positive u_f reduces the instantaneous net load reduction and charges the buffer, whereas negative u_f returns stored energy during lower-power intervals. Because $E_b(T_d) = E_b(0)$, the buffer is energy neutral over the service interval.

The controller first determined the minimum average response required to deliver the committed action energy by the deadline. The resulting peak cap was

$$P_{\text{cap}} = \max \left[\frac{E_e}{T_d - t_a}, \max_t p_e(t) - P_f^{\max} \right] \quad (14)$$

where P_{cap} is the target cap on the net load-reduction trajectory, $E_e = \int p_e(t)dt$ is the committed action energy, and t_a is the first block-action time. The support-power limit $P_f^{\max}$ is defined in equation (13).

When $p_e(t)$ exceeded P_{cap} , the buffer was dispatched to reduce the net peak subject to equation (13); stored energy was returned during lower-response intervals while enforcing $E_b(T_d) = 0$. The terminal energy constraint gives

$$\int_0^{T_d} p_{\text{net}}(t)dt = \int_0^{T_d} p_e(t)dt = E_e \quad (15)$$

Equation (15) shows that coordination reshapes the response without changing its net action energy.

Eight strategies were used to separate the effects of response timing, fast support and their coordination. Within each paired cell-seed comparison, B1–B7 had the same response magnitude A_e , net action energy and service deadline T_d . The support-enabled strategies B3, B5, B6 and B7 also shared the same support-power and buffer-energy limits. The strategies differed in temporal concentration, support availability and whether response timing and support were coordinated (Table 2).

Table 2 | Response strategies and equal-resource controls

ID	Strategy	Target ρ_{sync}	Load action	Fast support	Experimental role
B0	Static-load reference	—	None	None	Matched zero-action disturbance control
B1	Passive synchronized response	1.00	Fixed	None	Fully synchronized unmitigated response
B2	Heterogeneous ride-through	0.25	Fixed	None	Timing diversity without fast support
B3	Fast support only	1.00	Fixed	Enabled	Fast support without timing diversity
B4	Load flexibility only	0.50	Fixed	None	Intermediate timing-diversity comparator
B5	Uncoordinated local response	0.70	Fixed	Enabled	Equal-resource uncoordinated comparator
B6	Coordinated resilience control	0.25	Fixed	Enabled	Joint timing and fast-support coordination
B7	Oracle upper bound	0.10	Fixed	Enabled	Idealized performance bound

Within each cell–seed pair, B1–B7 had identical A_e , net action energy and T_d . B3, B5, B6 and B7 additionally shared identical $P_f^{\max}$ and $E_b^{\max}$. Their realized support trajectories could differ because the same resource limits were applied to different block-action schedules.

The IEEE control experiment compared B0–B7 under one high-stress and one moderate-stress condition using eight new timing seeds, producing 128 trajectories. RTS-GMLC transferred B2, B5 and B6 to the three selected sites, producing 72 action trajectories. ACTIVSg2000 evaluated low- and high-synchronization responses together with B5 and B6, producing 16 action trajectories and one matched zero-action control.

4.5. Outcome definitions

System-frequency outcomes were calculated from the centre-of-inertia (COI) frequency defined in equation (8). Let $f_a(t)$ denote the COI-frequency trajectory for a computational-load action and $f_b(t)$ its matched zero-action baseline, where the subscripts a and b denote the action

and baseline cases, respectively. The baseline-adjusted maximum-frequency increment, $\Delta f_{\max}$, was defined as

$$\Delta f_{\max} = \max_{t \in T} f_a(t) - \max_{t \in T} f_b(t) \quad (16)$$

where T denotes the full simulated disturbance and recovery interval. This difference-of-maxima definition isolates the additional peak-frequency excursion associated with the computational-load action relative to the matched disturbance baseline.

Two RoCoF measures were used for experiments with different analysis objectives. For the synchronization-surface analysis, the IEEE 39-bus control comparison and the RTS-GMLC synchronization experiments, the baseline-adjusted mechanism-window maximum-RoCoF increment, $\Delta \dot{f}_{\max}^{\text{mech}}$, was defined as

$$\Delta \dot{f}_{\max}^{\text{mech}} = \max_{t \in [1.00, 1.32]s} |\dot{f}_a(t)| - \max_{t \in [1.00, 1.32]s} |\dot{f}_b(t)| \quad (17)$$

where the superscript “mech” denotes the mechanism-window metric. The 320-ms window captures the immediate rate-of-change response following computational-load activation while excluding the subsequent recovery schedule.

For the transferred control experiments in RTS-GMLC and the targeted ACTIVSg2000 validation, the action-attributable maximum RoCoF, $\dot{f}_{\max}^{\text{act}}$, was calculated from the action-minus-baseline differential trajectory:

$$\dot{f}_{\max}^{\text{act}} = \max_{t \in [1.0, 6.5]s} |\dot{f}_a(t) - \dot{f}_b(t)| \quad (18)$$

The superscript “act” denotes the action-attributable metric. This measure captures the largest action-attributable rate excursion over the complete response and recovery interval. Because the two RoCoF measures apply the matched baseline differently and use different time windows, their absolute magnitudes were reported separately. Cross-system comparisons therefore

focused on the direction of the synchronization effect and the B6–B5 control ordering within each metric definition.

4.6. Conditional response-surface analysis

We fitted response surfaces to the 448 baseline-adjusted IEEE 39 trajectories. Synchronization, inertia scale, IBR penetration and response participation were centered and scaled to $[-1,1]$ over their respective experimental ranges:

$$z_\rho = \frac{\rho_{\text{sync}} - 0.55}{0.45}, \quad z_H = \frac{H_{\text{scale}} - 0.30}{0.05}, \quad z_\eta = \frac{\eta_{\text{IBR}} - 0.575}{0.075}, \quad z_m = \frac{m_{\text{resp}} - 0.825}{0.075} \quad (19)$$

Here, z_ρ , z_H , z_η and z_m denote the centered and scaled forms of synchronization, inertia scale, IBR penetration and response participation, respectively, and H_{scale} denotes the system inertia scale introduced in Section 4.1.

For each modeled outcome Y , with $Y \in \{\Delta f_{\text{max}}, \dot{\Delta f}_{\text{max}}^{\text{mech}}\}$, the quadratic response surface was

$$\begin{aligned} Y = & \beta_0 + \beta_\rho z_\rho + \beta_H z_H + \beta_\eta z_\eta + \beta_m z_m \\ & + \beta_{\rho H} z_\rho z_H + \beta_{\rho \eta} z_\rho z_\eta + \beta_{\rho m} z_\rho z_m \\ & + \beta_{\rho^2} z_\rho^2 + \beta_{\rho^2 H} z_\rho^2 z_H + \beta_{\rho^2 \eta} z_\rho^2 z_\eta + \beta_{\rho^2 m} z_\rho^2 z_m + \varepsilon \end{aligned} \quad (20)$$

Here, β_0 is the intercept, the remaining β terms are fitted regression coefficients whose subscripts identify the associated main, interaction or curvature terms, and ε is the residual. The z_ρ^2 term represents synchronization curvature at the center of the operating domain, while $z_\rho^2 z_H$, $z_\rho^2 z_\eta$ and $z_\rho^2 z_m$ allow that curvature to vary with inertia scale, IBR penetration and response participation, respectively. The linear comparison model retained the same main effects and first-order synchronization interactions but omitted z_ρ^2 , $z_\rho^2 z_H$, $z_\rho^2 z_\eta$ and $z_\rho^2 z_m$.

Both models were fitted by ordinary least squares. Predictive performance was evaluated by eight-fold leave-one-seed-out validation: in each fold, one complete timing seed was withheld, the model was fitted to the remaining seven seeds, and predictions were generated for all operating conditions associated with the held-out seed. The quadratic specification was retained only when its leave-one-seed-out root-mean-square error was at least 2% lower than that of the linear comparison model.

4.7. Statistical analysis

The independent sampling unit was the timing seed, or the timing seed within each network location for multisite experiments. Time samples, buses and converter states were treated as within-run observations rather than independent replicates. Paired comparisons retained the same timing seed, operating point, disturbance, computational-load location, response magnitude and response-energy budget; only the synchronization level or control strategy was changed.

Unless stated otherwise, effects were reported as paired mean differences with 95% percentile bootstrap confidence intervals based on 2,000 resamples of complete timing-seed pairs. The IEEE 39-bus synchronization and control experiments and the RTS-GMLC experiments used eight paired timing seeds per condition. Response-surface coefficient intervals were obtained by cluster bootstrap at the timing-seed level, with all observations associated with a selected seed resampled together. Predictive performance was evaluated using the eight-fold leave-one-seed-out procedure described in Section 4.6. The targeted ACTIVSg2000 experiments used four paired timing seeds per prespecified condition.

The numerical-refinement experiment used six timing seeds paired across the 10-, 5- and 2.5-ms time steps. Numerical agreement was assessed using maximum paired discrepancies and prespecified convergence tolerances rather than population-level inference. Primary contrasts were

defined before the corresponding confirmatory runs. Confidence intervals were used for effect estimation; no null-hypothesis significance tests or multiplicity-adjusted p values were applied.

4.8. Numerical implementation and reproducibility

Dynamic states were advanced on a fixed time grid. Continuous computational-load actions were integrated over each numerical interval before being applied to the network equations, preserving action energy independently of alignment with the simulation grid. At each step, the algebraic network equations were solved by relaxed iteration while converter current-limit states were updated within the same solve. When required, a residual-form nonlinear fallback was accepted only if it reduced the current-balance residual. ACTIVSg2000 used sparse admittance assembly and topology-aware factorization reuse without the nonlinear fallback.

Production experiments used a 2.5-ms time step. Numerical refinement repeated 126 action trajectories and three matched zero-action controls using 10-, 5- and 2.5-ms time steps over a 15-s horizon. Agreement between the two refined grids was required for the recovery classification, with a paired terminal-margin difference below 0.02 and a maximum paired mechanism-window RoCoF difference below $1.0 \times 10^{-3} \text{ Hz s}^{-1}$.

Runs that failed the network-convergence criterion were retained for numerical diagnostics but were not interpreted as physical system outcomes.

All simulations were implemented in MATLAB R2024b. MATPOWER 8.1 supplied the network cases, operating-point solvers and admittance-matrix construction. Random timing schedules were generated using the Mersenne Twister, with all seeds retained for reproducibility. Each run recorded the simulation configuration, code entry point, software version, random seed, scenario identifiers, block-action event log, waveform metrics and numerical acceptance checks. A trajectory was used for physical comparison only when its matched zero-action control and the

corresponding synchronization, response-energy and network-residual checks satisfied the prespecified criteria.

Data availability

The benchmark network data used in this study are publicly available through the canonical IEEE 39-bus (<https://doi.org/10.1109/TPAS.1979.319407>), RTS-GMLC (<https://doi.org/10.1109/TPWRS.2019.2925557>), and ACTIVSg2000 (<https://doi.org/10.1109/TPWRS.2016.2616385>) sources. Simulation data generated in this study are available from the corresponding author upon reasonable request.

Code availability

The simulation and analysis code used in this study is part of an ongoing research project and is available from the corresponding author upon reasonable request.

References

- [1] Masanet, E., Shehabi, A., Lei, N., Smith, S. & Koomey, J. Recalibrating global data center energy-use estimates. *Science* 367, 984–986 (2020).
- [2] Shehabi, A. et al. 2024 United States Data Center Energy Usage Report. LBNL-2001637 (Lawrence Berkeley National Laboratory, Berkeley, CA, USA, 2024).
- [3] International Energy Agency. Energy and AI (IEA, Paris, France, 2025).
- [4] North American Electric Reliability Corporation. Characteristics and Risks of Emerging Large Loads (NERC, Atlanta, GA, USA, 2025).
- [5] Federal Energy Regulatory Commission. Reliability Standard(s) Pertaining to Computational Load Integration. Docket No. RD26-7-000, 196 FERC ¶ 61,031 (2026).

- [6] North American Electric Reliability Corporation. Aggregated Report on NERC Level 2 Industry Recommendation: Large Load Interconnection, Study, Commissioning, and Operations (NERC, Atlanta, GA, USA, 2026).
- [7] North American Electric Reliability Corporation. Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions (NERC, Atlanta, GA, USA, 2025).
- [8] North American Electric Reliability Corporation. Reliability Guideline: Risk Mitigation for Emerging Large Loads (NERC, Atlanta, GA, USA, 2026).
- [9] Han, T. et al. Designing and regulating clean energy data centres. *Nat. Rev. Clean Technol.* 1, 373–374 (2025).
- [10] Colangelo, P. et al. AI data centres as grid-interactive assets. *Nat. Energy* 11, 254–261 (2026).
- [11] Chien, A. A., Gupta, U., Ren, S., Sriraman, A. & Tomlinson, B. Strategies and design for increasing AI sustainability. *Nat. Rev. Clean Technol.* 2, 624–638 (2026).
- [12] Caprara, A. et al. Data center workload flexibility for power system demand response: Evidence from Alibaba traces. *Int. J. Electr. Power Energy Syst.* 178, 111940 (2026).
- [13] Alaperä, I., Honkapuro, S. & Paananen, J. Data centers as a source of dynamic flexibility in smart grids. *Appl. Energy* 229, 69–79 (2018).
- [14] Kaur, K. et al. An adaptive grid frequency support mechanism for energy management in cloud data centers. *IEEE Syst. J.* 14, 1195–1205 (2020).
- [15] Azizi, A. et al. Strengthening data center operations using grid-forming battery energy storage as a line-interactive uninterruptible power supply. *Int. J. Electr. Power Energy Syst.* 175, 111638 (2026).
- [16] Li, J., Bao, Z. & Li, Z. Modeling demand response capability by Internet data centers processing batch computing jobs. *IEEE Trans. Smart Grid* 6, 737–747 (2015).

- [17] Wang, H., Huang, J., Lin, X. & Mohsenian-Rad, H. Proactive demand response for data centers: A win-win solution. *IEEE Trans. Smart Grid* 7, 1584–1596 (2016).
- [18] Chen, M. et al. Internet data center load modeling for demand response considering the coupling of multiple regulation methods. *IEEE Trans. Smart Grid* 12, 2060–2076 (2021).
- [19] Zhang, W. & Zavala, V. M. Remunerating space–time, load-shifting flexibility from data centers in electricity markets. *Appl. Energy* 326, 119930 (2022).
- [20] Riepin, I., Brown, T. & Zavala, V. M. Spatio-temporal load shifting for truly clean computing. *Adv. Appl. Energy* 17, 100202 (2025).
- [21] Han, J. et al. Evaluating the dispatchable capacity of cross-regional data center clusters toward power system operation. *IEEE Trans. Smart Grid* 17, 1180–1193 (2026).
- [22] Gu, Y. & Green, T. C. Power system stability with a high penetration of inverter-based resources. *Proc. IEEE* 111, 832–853 (2023).
- [23] Dörfler, F. & Groß, D. Control of low-inertia power systems. *Annu. Rev. Control Robot. Auton. Syst.* 6, 415–445 (2023).
- [24] He, C., Geng, H., Rajashekara, K. & Chandra, A. Analysis and control of frequency stability in low-inertia power systems: A review. *IEEE/CAA J. Autom. Sinica* 11, 2363–2383 (2024).
- [25] Song, S., Kim, J., Kim, S. W., Kim, J. & Yoon, M. Oscillation assessment of the KPG-193 test system with hyperscale data center interconnections. *Int. J. Electr. Power Energy Syst.* 177, 111811 (2026).
- [26] Gyang, P. P., Chakraborty, P., Meegahapola, L. G. & Yu, X. Dynamic modeling of a data center for power system stability studies. *IEEE Trans. Power Syst.* 41, 2317–2332 (2026).

759 [27] Jiménez-Ruiz, A. & Milano, F. Data centre model for transient stability analysis of power
760 systems. *IET Gener. Transm. Distrib.* 20, e70381 (2026).

761 [28] Energy Systems Integration Group. Large Load Modeling for Dynamic Studies: Current
762 Practices and Recommendations (Energy Systems Integration Group, 2026).

763 [29] Wang, X. et al. Dynamic Modeling of AI Data Center Load in PSCAD. Zenodo,
764 10.5281/zenodo.21004140 (2026).

765 [30] Ross, B. A. & Follum, J. D. Electromagnetic Transient Modeling of Large Data Centers for
766 Grid-Level Studies. PNNL-38817 (Pacific Northwest National Laboratory, Richland, WA, USA,
767 2026).

768 [31] Sun, J., Xu, M., Cespedes, M. & Kauffman, M. Data center power system stability—Part I:
769 Power supply impedance modeling. *CSEE J. Power Energy Syst.* 8, 403–419 (2022).

770 [32] Sun, J., Mihret, M., Cespedes, M., Wong, D. & Kauffman, M. Data center power system
771 stability—Part II: System modeling and analysis. *CSEE J. Power Energy Syst.* 8, 420–438 (2022).

772 [33] Zhu, T., Wang, X., Zhao, F. & Torrico-Bascope, G. V. Impedance-based aggregation of
773 paralleled power factor correction converters in data centers. *IEEE Trans. Power Electron.* 38,
774 5254–5265 (2023).

775 [34] Potts, J. & Castellanos, S. Grid planning insights from near-optimal generation portfolios
776 amidst large-load growth. *npj Clean Energy* 1, 14 (2025).

777 [35] Lu, X. et al. Structural alignment for energy–computation co-design. *Nat. Rev. Electr. Eng.*
778 3, 266–267 (2026).

779 [36] Chen, S. et al. Operational flexibility of active distribution networks with the potential from
780 data centers. *Appl. Energy* 293, 116935 (2021).

- [37] Naoi, H., Delage, R., Nakata, T. & Kozai, M. Integrating location planning and spatio-temporal workload shifting for 100% renewable energy data centers: A case study in Japan. *Energy Strategy Rev.* 61, 101823 (2025).
- [38] Mohammadi, S. et al. Grid integration of AI data centers: A critical review of energy storage solutions. *Adv. Appl. Energy* 23, 100291 (2026).
- [39] Athay, T., Podmore, R. & Virmani, S. A practical method for the direct analysis of transient stability. *IEEE Trans. Power Appar. Syst.* PAS-98, 573–584 (1979).
- [40] Barrows, C. et al. The IEEE Reliability Test System: A proposed 2019 update. *IEEE Trans. Power Syst.* 35, 119–127 (2020).
- [41] Birchfield, A. B., Xu, T., Gegner, K. M., Shetye, K. S. & Overbye, T. J. Grid structural characteristics as validation criteria for synthetic networks. *IEEE Trans. Power Syst.* 32, 3258–3265 (2017).

Acknowledgements

This work was supported by the Smart Grid-National Science and Technology Major Project under grant 2024ZD0801200.

Author contributions

Wei Tai: Writing - original draft, Visualization, Software, Methodology, Investigation, Data curation. **Qi Wang:** Writing - review & editing, Validation, Supervision, Funding acquisition, Conceptualization. **Yi Tang:** Writing - review & editing, Validation, Investigation, Formal analysis.

Competing interests

The author declares no competing interests.

