

Spatial Load Correlation in AI Data-Center-Dominated Power Systems

Chandan Chaudhary, *Student Member, IEEE*, Alaaeldin Abdelkader, *Student Member, IEEE*, Yansong Pei, *Member, IEEE*, Mohammed Benidris, *Senior Member, IEEE*, and Joydeep Mitra, *Fellow, IEEE*
Electrical and Computer Engineering, Michigan State University, East Lansing, MI 48824, USA
Emails: chaud152@msu.edu, abdelk15@msu.edu, peiyanso@msu.edu, benidris@msu.edu, and mitraj@msu.edu

Abstract—The proliferation of large-scale data centers introduces spatially correlated demand profiles that challenge the long-standing assumption of statistical independence of loads in power system analysis. This paper examines the emergence of such load correlations and evaluates their impact on data-center-dominated grids. Analytical derivations reveal that correlated load fluctuations amplify aggregate stochastic disturbances, reduce voltage stability margins through weakened reactive power stiffness, and degrade frequency stability margin by erosion of natural load diversity effects. Real-time digital simulation studies confirm that moderate spatial correlation in distributed data centers produces simultaneous frequency deviations and voltage fluctuations across multiple buses. The findings offer transmission system operators a physics-based perspective to interpret emerging oscillatory phenomena and establish stability planning criteria grounded in measurable load-correlation structures rather than traditional diversity assumptions.

Index Terms—AI data center, correlated load, inter-area oscillations, load correlation, power system dynamics, RTDS

I. INTRODUCTION

The rapid expansion of AI and high-performance computing (HPC) has fundamentally reshaped electricity demand across major interconnections. Data centers have emerged as one of the most consequential load classes directly interfacing with bulk power system. Recent projections reveal that global data centers could demand up to 945 TWh annually by 2030, with AI-driven workloads accounting for up to 12% of U.S. electricity consumption by 2028 [1], [2]. Modern hyperscale facilities operate at extreme power densities, often exceeding 100 MW per site [2], [3], and some campuses are now planned at the gigawatt scale. Such facilities exhibit sharp power fluctuations arising from synchronized compute and communication phases in large-scale training jobs [3], [4].

A defining characteristic of this evolution is geographic concentration. Northern Virginia’s “Data Center Alley” now exceeds 2.5 GW of active demand, and the addition of new AI-data-center campuses produces highly localized, synchronized load patterns that intensify transmission stress [2], [3]. This clustering forms demand hubs that impose intense stress on regional transmission infrastructure and invalidate the load diversity assumptions of power system planning and operation.

Data centers interface with the grid primarily through voltage-source converters, which include double-conversion UPS systems, rack-level rectifiers, and high-density server power supplies [5]–[7]. This converter-dominated architecture creates loads with ultra-low inertia, weak damping, and negative impedance across multiple frequency bands, which decouple demand from inertial buffering and differ from elec-

tronic loads. Field measurements reveal low-frequency resonances near 11 Hz and high-frequency oscillations in the 5–10 kHz range at rack, triggered by control delays and DC bus dynamics [8], [9]. Real-world events, such as 14.7–14.8 Hz oscillations in Dominion Energy, show that unplanned excitations from data center operations destabilize hidden dynamic modes in weak grids [10].

AI and HPC load profiles demonstrate that these loads generate megawatt-scale ramps within seconds during synchronized compute-communication cycles, with fluctuation spectra capable of exciting inter-area electromechanical modes when aligned with weakly damped grid resonances [3], [6], [11]. Multiple data centers operating under shared orchestration platforms, thermal constraints, or workload scheduling exhibit correlated power variations across distinct network buses. These loads induce voltage-sensitive reductions, trigger frequency excursions, and exacerbate stability risks in converter-rich grids [2], [6].

Although existing studies provide insight into single-site data-center dynamics, analytical tools that capture correlation across geographically dispersed facilities and translate those effects to system-level behavior remain scarce. Planning criteria assume load independence and overlook impacts of correlated load variations in stability thresholds or reserve requirements. This paper addresses these gaps in three ways: (1) defining spatial load correlation mathematically and identifying the physical mechanisms that create it; (2) deriving analytical relationships that link load correlation to voltage stability, frequency response, and oscillation amplification; and (3) validating the theoretical framework through real-time digital simulations that show even moderate correlation can produce simultaneous disturbances across multiple buses.

The remainder of this paper is organized as follows. Section II introduces the concept of correlated load. The physical mechanisms giving rise to load correlation are detailed in Section III. Section IV quantifies the impacts of correlated load on voltage stability, frequency dynamics, inter-area oscillations, and reserve planning requirements. Section V presents real-time digital simulation studies on the IEEE 39-bus system. Finally, Section VI concludes the paper and provides direction for future research in this sector.

II. CORRELATED LOAD

Power system analysis has relied on the premise that electrical loads at different network locations fluctuate independently. This consideration, valid for dispersed and uncoordinated residential or industrial demand, offers analytical simplicity

and justifies the use of diversity factors in planning and stability studies. Each bus-level demand is modeled as an uncorrelated stochastic process $\{P_i(t)\}_{i=1}^N$, which ensures aggregate variance grows linearly with the number of loads. This naturally mitigates system-wide fluctuations.

This independence simplifies forecasting, modal analysis, and infrastructure planning as each load acts as an isolated stochastic source. This allows linearized small-signal models and probabilistic methods to treat buses separately. Mathematically, independence requires that for $i \neq j$,

$$\mathbb{E}[P_i(t)P_j(t')] = \mathbb{E}[P_i(t)]\mathbb{E}[P_j(t')], \text{Corr}(P_i(t), P_j(t')) = 0 \quad (1)$$

a condition increasingly violated by the rise of synchronized and digitally managed facilities.

A. Spatially Correlated Load Behavior

Modern grids are comprised of converter-dominated entities such as data centers, EV fleets, and smart campuses, that operate under centralized control or shared environmental stimuli. Their coordinated behavior induces temporal correlation among power injections at geographically distinct buses, which gives rise to *spatially correlated load*. From a system perspective, such correlation emerges only when fluctuations at multiple nodes evolve coherently in time.

A single, time-varying load at one bus, regardless of the internal synchronization among its subsystems, does not constitute spatial correlation because the grid perceives only one aggregated injection. A facility may be composed of multiple subsystems that operate in synchrony. If the facility connects to the transmission network through a single node i , the grid observes only the total power injection as a single stochastic process. Internal synchronization among subsystems, such as GPU racks or HVAC controls, remains electrically invisible, thus yields $\text{Corr}(P_i(t), P_j(t')) = 0$ for all $j \neq i$.

Spatial correlation instead occurs when loads at distinct buses i and j (with $\mathbf{Z}_{ij} \neq 0$, $\mathbf{Z}$ is the network impedance matrix) fluctuate coherently. The correlation strength is captured by the cross-correlation function,

$$\rho_{ij}(\tau) = \frac{\mathbb{E}[(P_i(t) - \bar{P}_i)(P_j(t + \tau) - \bar{P}_j)]}{\sigma_i \sigma_j}, \quad (2)$$

and a significant correlation exists when $|\rho_{ij}(\tau)| > \epsilon$ for some $\tau \in [-\tau_{\max}, \tau_{\max}]$. The resulting correlation matrix

$$C_{ij} = \max_{\tau \in [-\tau_{\max}, \tau_{\max}]} |\rho_{ij}(\tau)| \quad (3)$$

is symmetric with $C_{ii} = 1$ and $C_{ij} \in [0, 1]$. Strong off-diagonal terms indicate synchronized fluctuations among distinct buses, a phenomenon increasingly observed in data-center-dominated grids. As illustrated in Fig. 1, synchronized facilities connected to distinct buses can generate strong off-diagonal correlation terms C_{ij} , even when neighboring loads remain independent. This spatial structure, characterized by selective non-zero off-diagonal elements in $\mathbf{C}$, fundamentally differentiates correlated load phenomena from single-site temporal variability and necessitates analytical frameworks that extend beyond conventional power flow formulations.




Fig. 1. Representation of multi-site correlated data-center loads with rack-level and HVAC synchronization.

III. PHYSICAL MECHANISMS OF LOAD CORRELATION

Correlated load in modern power systems stems from physical, operational, and environmental couplings among large digitally managed facilities. This section examines the physical mechanisms that induce such correlations across grid buses. In data-center-dominated networks, these mechanisms fall into three categories: operational synchronization, converter-mediated electrical coupling, and environmental forcing.

A. Cross-Facility Workload Orchestration

Hyperscale cloud and AI operators routinely distribute computational workloads across multiple sites for redundancy, throughput, and latency optimization. Centralized orchestration frameworks (e.g., DeepSpeed, Kubernetes, or internal schedulers) synchronize training phases, checkpoint updates, and communication cycles across data centers connected to different transmission nodes [12]–[14].

When compute clusters at buses i and j initiate or pause tasks nearly simultaneously, their active power demand follows a phase-aligned temporal pattern, that can be approximated by:

$$P_i(t) = P_{i,0} + A_i \cos(\omega_c t + \phi_i), \quad (4)$$

$$P_j(t) = P_{j,0} + A_j \cos(\omega_c t + \phi_j), \quad (5)$$

where ω_c is the characteristic frequency of the workload cycle. When $|\phi_i - \phi_j| \ll 1$, the cross correlation

$$\mathbb{E}[\Delta P_i(t) \Delta P_j(t)] = \frac{A_i A_j}{2} \cos(\phi_i - \phi_j) \approx \frac{A_i A_j}{2} \quad (6)$$

is positive, which implies spatial correlation. This effect is particularly plausible in regions such as Northern Virginia, where a dense concentration of hyperscale data centers shares common transmission corridors with large commercial and residential loads. In such clusters, synchronized AI training schedules can create temporally correlated power fluctuations observable across multiple substations.

B. Grid-Mediated Electromagnetic Coupling

Modern AI data centers are *converter-dominated loads* that interface with the grid through high-capacity AC/DC rectifiers, UPS systems, and DC–DC converters [7]. These converters exhibit fast voltage–current dynamics, low effective inertia, and nonlinear impedance characteristics [5]. As a result, local power disturbances can travel through the network as electromagnetic effects that couple nearby substations [15].

Under small-signal conditions, the relation between nodal voltage and power perturbations can be approximated by

$$\Delta \mathbf{V} = \mathbf{Z} \Delta \mathbf{I}, \quad \Delta I_k \approx \frac{\Delta P_k}{V_k^*}. \quad (7)$$

The voltage deviation at bus j due to active-power variations across all buses is

$$\Delta V_j = \sum_{k=1}^N Z_{jk} \frac{\Delta P_k}{V_k^*}. \quad (8)$$

When a local power perturbation ΔP_i occurs at bus i , converters at adjacent buses respond through grid-mediated coupling. The induced active-power response at bus j is

$$\Delta P_j^{\text{resp}} = -\frac{E'V \cos \theta}{X} \cdot \frac{\Delta P_i}{S_b}, \quad (9)$$

where E' and V are the internal and terminal voltages, θ is the power angle, X is the line reactance, and S_b is the system base power [15]. The disturbance travels through the network as an *electromechanical wave* with approximate speed

$$c \approx \sqrt{\frac{E'V \cos \theta \omega_s}{2HS_bX}}, \quad (10)$$

where H is the system inertia. This links the fast voltage and current dynamics of spatially separated converter stations [15].

In regions with high data-center density, such as AI clusters sharing 230–500 kV corridors, this propagation mechanism synchronizes converter responses across substations. The resultant voltage–current alignment reinforces correlated load oscillations generated by synchronized computational workloads and reduces voltage stability margins.

C. Thermal and Environmental Coupling

Beyond fast electrical interactions, slower environmental processes also generate correlated load patterns. Data centers located in the same geographical area experience similar weather conditions. Thus, their HVAC systems follow comparable heating and cooling patterns. The thermal dynamics of each facility is defined by a lumped-capacitance model [16].

$$C_{\text{th}} \frac{dT_i}{dt} = P_{\text{IT}}^{(i)}(t) - \frac{T_i - T_{\text{amb}}}{R_{\text{th}}} - Q_{\text{cool}}^{(i)}(t), \quad (11)$$

where $Q_{\text{cool}}^{(i)}$ is the HVAC cooling rate. When data centers experience similar ambient temperatures T_{amb} , their cooling controllers respond in a comparable manner. This leads to correlated cooling demand and electrical load patterns. Such effect is most pronounced in geographically clustered facilities but can also arise in large commercial or institutional buildings with similar HVAC control strategies.

The mechanisms above act on distinct timescales. GPU training induces high-frequency power fluctuations [4], converter interactions cause mid-frequency dynamics [5], and HVAC systems add low-frequency modulation [4]. When data centers are electrically and geographically proximate, these mechanisms may combine to form correlated load patterns.

IV. IMPACTS OF CORRELATED LOAD ON POWER SYSTEMS

Spatial and temporal correlations among large electrical loads reshape system dynamics and reliability margins. When load variations align across locations, aggregate disturbances rise, and system modes show stronger oscillations that reduce voltage stability, frequency balance, and available reserves.

A. Voltage Stability Margin Degradation

Voltage variability near an operating point is governed by the linearized power–flow sensitivities:

$$\begin{bmatrix} \Delta P \\ \Delta Q \end{bmatrix} = \begin{bmatrix} J_{PP} & J_{PV} \\ J_{QP} & J_{QV} \end{bmatrix} \begin{bmatrix} \Delta \theta \\ \Delta V \end{bmatrix}, \quad \Delta V \approx -J_{QV}^{-1} \Delta Q, \quad (12)$$

so the smallest singular value $\sigma_{\min}(J_{QV})$ quantifies the weakest voltage–restoring stiffness. The coupling between active- and reactive-power variations can be expressed through a local small-signal transfer function $H_{QP,i}(j\omega)$ that captures converter admittance and network voltage sensitivity after the elimination of voltage perturbations [8], [9]. At each bus i ,

$$\Delta Q_i(j\omega) = H_{QP,i}(j\omega) \Delta P_i(j\omega). \quad (13)$$

Expanding $H_{QP,i}$ around the steady-state set point at $(\omega=0)$:

$$H_{QP,i}(j\omega) = H_{QP,i}(0) + j\omega H'_{QP,i}(0) + \mathcal{O}(\omega^2). \quad (14)$$

Neglecting higher-order terms and defining

$$\alpha_i = H_{QP,i}(0), \quad \beta_i = H'_{QP,i}(0), \quad (15)$$

the first-order approximation from (13), (14) and (15) becomes

$$\Delta Q_i(j\omega) \approx (\alpha_i + j\omega\beta_i) \Delta P_i(j\omega). \quad (16)$$

Transforming to the time domain yields

$$\Delta Q_i(t) \approx \alpha_i \Delta P_i(t) + \beta_i \frac{d\Delta P_i(t)}{dt}. \quad (17)$$

The coefficient α_i represents the quasi-static dependence of reactive power on active power at the converter node, determined primarily by the DC-link and PLL control gains, while β_i quantifies the dynamic phase shift introduced by control and network time constants [8], [9], [15].

Let the active-power fluctuations be spectrally correlated over a dominant frequency band Ω_c , with correlation strength λ_1 and total spectral power S_c at frequency ω_c . The corresponding voltage spectrum (for the most sensitive mode) can be approximated as

$$|\Delta V(j\omega)| \approx \frac{|\Delta Q(j\omega)|}{\sigma_{\min}(J_{QV})} \approx \frac{|\alpha + j\omega_c\beta| S_c \lambda_1}{\sigma_{\min}(J_{QV})}. \quad (18)$$

Voltage instability occurs when the correlated disturbance energy in the most sensitive mode equals the voltage–restoring capability, i.e., $|\Delta V(j\omega_c)| \gtrsim 1$. Rearranging (18) yields the critical correlation threshold,

$$\lambda_1^{\text{crit}} = \frac{\sigma_{\min}(J_{QV})}{\sqrt{\alpha^2 + \omega_c^2 \beta^2} S_c}. \quad (19)$$

Voltage instability is expected when $\lambda_1 > \lambda_1^{\text{crit}}$. Equation (19) expresses an energy balance between correlated stochastic excitation and voltage–restoring stiffness. Higher load correlation (λ_1), dominant low-frequency content (ω_c), strong active–reactive coupling (α, β), or greater spectral power (S_c) all lower λ_1^{crit} , thereby reducing the voltage stability margin.

B. Frequency Deviation Amplification

Correlated load ramps reduce the effectiveness of inertial buffering [17]. When data centers impose synchronized active-power changes, the aggregate disturbance amplifies the initial frequency deviation.

Let L_i denote zero-mean load fluctuations at bus i with variance σ^2 . For analytical tractability, the N load fluctuations L_i are assumed to have identical variances $\text{Var}(L_i) = \sigma^2$. The uniform pairwise correlation is

$$\rho_0 = \frac{\text{Cov}(L_i, L_j)}{\sigma^2} \quad (20)$$

so $\text{Cov}(L_i, L_j) = \sigma^2$ if $i = j$, and $\rho_0 \sigma^2$ otherwise. The aggregate load $L_{\text{agg}} = \sum L_i$ then has variance

$$\begin{aligned} \sigma_{\text{agg}}^2 &= \text{Var}\left(\sum_{i=1}^N L_i\right) = \sum_{i=1}^N \text{Var}(L_i) + 2 \sum_{i < j} \text{Cov}(L_i, L_j) \\ &= N\sigma^2 + N(N-1)\rho_0\sigma^2 = N\sigma^2[1 + (N-1)\rho_0] \end{aligned} \quad (21)$$

For a step $\Delta P_{\text{agg}} = \sqrt{1 + (N-1)\rho_0} \Delta P_0$, the swing equation gives

$$\Delta f_0 \approx -\frac{\Delta P_0}{2H_{\text{sys}}f_0} \sqrt{1 + (N-1)\rho_0}. \quad (22)$$

where f_0 is the nominal system frequency and H_{sys} is the system inertia constant. For $N = 5$, $\rho_0 = 0.8$, the factor $\Delta f_0 \approx 2.05$, doubling the expected dip. This reduces frequency margins under low-inertia conditions, where fast converter control contains electromechanical waves [15] but synchronized action worsens RoCoF and nadir.

C. Inter-Area Oscillation Amplification

Spatially correlated load variations can couple into electromechanical oscillations of the interconnected system. When the spatial correlation mode of load fluctuations, say $\mathbf{u}_1$, aligns with a poorly damped inter-area mode shape ϕ_k , the modal excitation strength is proportional to the participation factor $|\mathbf{u}_1^\top \phi_k|^2$. High values indicate constructive alignment between the correlated disturbance pattern and the oscillatory mode. Correlated activity from large-scale data centers or synchronized computing clusters, often concentrated around ~ 1 Hz, may interact with inter-area modes typically occurring in the 0.2–0.8 Hz range, potentially amplifying tie-line oscillations and degrading damping [11], [17]. Spectral overlap can induce beat phenomena linking data center power electronics and AI workloads to wide-area oscillatory behavior [10], [11]. These effects are pronounced when the spectral content of the correlated load fluctuations overlaps the dominant electromechanical mode frequencies.

D. Planning and Operating Reserve Implications

Traditional diversity factors assume that individual load fluctuations are statistically independent, which yields an aggregate variance that scales linearly with the number of buses. Spatially distributed data center loads, on the other hand exhibit correlated behavior. Using the covariance structure derived in (21), positive correlation ($\rho_0 > 0$) increases aggregate variability relative to the independent-load aggregate.

For illustration, with $N=10$ and a moderate correlation level $\rho_0=0.3$ (30% correlation), the aggregate variance is $37\times$ larger than that of a single bus and $3.7\times$ larger relative to the independent-load baseline $N\sigma^2$, corresponding to a $\sqrt{3.7} \approx 1.9\times$ increase in aggregate standard deviation. This enlarged variability necessitates proportionally higher operating and contingency reserves to maintain equivalent reliability margins under correlated demand fluctuations [17].

V. REAL-TIME DIGITAL SIMULATION STUDIES

Real-time simulations were conducted on the IEEE 39-bus test system using the RTDS platform to assess the dynamic impact of spatially correlated data center loads. RTDS was selected because it maintains real-time execution and synchronous phase relationships across loads (which offline tools cannot guarantee) and provides sub-microsecond timestep resolution to capture converter dynamics. Three representative AI data centers were placed at buses 4, 12, and 15, each supported by dedicated battery energy storage systems (BESS), with ratings summarized in Table I. Dynamic load models with average-value converter representations were employed for the data centers following [6], rather than detailed switching models, as switching frequencies (5–20 kHz) are far above the correlation time scales of interest. The active and reactive power demands evolved according to multi-cycle AI workload patterns, with a correlation coefficient of $\rho \approx 0.4$ imposed across the three sites to model partial synchronization driven by shared computational scheduling. This moderate correlation represents a conservative middle ground between statistical independence ($\rho = 0$) and perfect synchronization ($\rho = 1$).

TABLE I
AI DATA CENTER LOADS AND BESS RATING

Bus	P [MW]	Q [MVar]	P_{BESS} [MW]	E_{BESS} [MWh]
4	200	40.6	200	50.0
12	225	45.7	225	56.3
15	250	50.8	250	62.5

Figure 2a shows the active and reactive power trajectories at the data center buses. The traces reveal synchronized ramp-up and ramp-down sequences, with active power peaks occurring nearly concurrently across all sites. Reactive power exhibits analogously but with minor phase shifts from independent converter voltage regulation. This temporal coherence in load evolution generates amplified aggregate power swings relative to uncorrelated scenarios, consistent with theoretical expectations. The system frequency response, shown in Fig. 2b, displays clear sensitivity to the correlated load ramps. During coincident demand surges, frequency deviates downward by up to 0.15 Hz from the 60 Hz nominal. Recovery occurs promptly as load decreases, supported by generator droop action and BESS discharge. The distributed BESS units limit the total frequency excursion to within 0.3 Hz peak-to-peak. This further signifies the need for dedicated ESS for effective suppression of frequency deviations induced by correlated load fluctuations. Figure 2c illustrates the RMS voltage response at buses 4, 12, and 15. Voltage magnitudes remain within




Fig. 2. Dynamic responses of the IEEE 39-bus system under correlated data-center load profiles

$\pm 2\%$ of the nominal value but show distinct fluctuations during the correlated load ramps. The nearly simultaneous voltage dips observed across all three buses reveal a shared electromechanical mode driven by spatially correlated load variations. This behavior demonstrates inter-bus coupling in the voltage domain and underscores the influence of coherent load fluctuations on network-wide voltage stability.

Real-time simulations confirm that moderate load correlation ($\rho \approx 0.4$) markedly diminishes the natural diversity effect that underpins classical load independence assumptions. Incorporating spatial load correlation into transient analysis becomes essential to capture realistic grid behavior. Coordinated BESS operation proves vital to maintain frequency and voltage stability in data-center-dominated networks.

VI. CONCLUSION

This paper provides a scientific rationale for characterizing and analyzing spatially correlated loads in data-center-dominated power systems. Analytical derivations quantify how load correlation amplifies aggregate disturbances, weakens voltage stability margins, degrades frequency response, and excites inter-area oscillation modes. Real-time digital simulations on the IEEE 39-bus system validate theoretical predictions and demonstrate that moderate spatial correlation produces simultaneous frequency and voltage deviations across multiple buses. The results confirm that traditional diversity assumptions no longer hold in converter-dominated grids with synchronized computational workloads.

Future research will focus on empirical validation with data from data-center-intensive regions to calibrate correlation models and refine stability thresholds. The development of real-time load correlation estimation algorithms for system operators represents an operational need. Investigation of coordinated control strategies for distributed energy storage systems to mitigate correlation-induced oscillations offers an important direction for grid management. Finally, integration of load correlation metrics into long-term transmission planning tools will support infrastructure investment decisions in the evolution of AI-driven power systems.

ACKNOWLEDGMENT

The research is supported in part by Sandia National Lab, in part by MSU Research Foundation and in part by the U.S. National Science Foundation under grant 2408615.

REFERENCES

- [1] A. Shehabi, A. Newkirk, S. J. Smith, A. Hubbard, N. Lei, M. A. B. Siddik, B. Holecek, J. Koomey, E. Masanet, and D. Sartor, "2024 United States Data Center Energy Usage Report," Tech. Rep. LBNL-2001637, Lawrence Berkeley National Laboratory, Berkeley, CA, Dec. 2024. Energy Analysis and Environmental Impacts Division.
- [2] NERC, "Characteristics and risks of emerging large loads," task force white paper, North American Electric Reliability Corporation, 2025.
- [3] R. Quint, K. Thomas, J. Zhao, A. Isaacs, and C. Baker, "Practical Guidance and Considerations for Large Load Interconnections," tech. rep., Elevate Energy Consulting and GridLab, May 2025.
- [4] S. Go, J. Park, S. More, H. Wu, I. Wang, A. Jezghani, T. Krishna, and D. Mahajan, "Characterizing the Efficiency of Distributed Training: A Power, Performance, and Thermal Perspective," in *Proceedings of the 58th IEEE/ACM International Symposium on Microarchitecture*, pp. 626–642, 2025.
- [5] L. Xiong, X. Liu, Y. Liu, and F. Zhuo, "Modeling and stability issues of voltage-source converter-dominated power systems: A review," *CSEE Journal of Power and Energy Systems*, vol. 8, no. 6, pp. 1530–1549, 2020.
- [6] C. Chaudhary, A. Abdelkader, M. Egan, E. Udren, M. Benidris, and J. Mitra, "Impact of data center load modeling on power system stability," in *Grid of the Future Symposium*, CIGRE US, (Denver, Colorado, USA), Nov. 2025.
- [7] J. Sun, S. Wang, J. Wang, and L. M. Tolbert, "Development of a converter-based data center power emulator," in *2021 IEEE Applied Power Electronics Conference and Exposition (APEC)*, pp. 126–133, IEEE, 2021.
- [8] J. Sun, M. Xu, M. Cespedes, and M. Kauffman, "Data center power system stability—Part I: Power supply impedance modeling," *CSEE Journal of Power and Energy Systems*, vol. 8, no. 2, pp. 403–419, 2022.
- [9] J. Sun, M. Mihret, M. Cespedes, D. Wong, and M. Kauffman, "Data center power system stability—Part II: System modeling and analysis," *CSEE Journal of Power and Energy Systems*, vol. 8, no. 2, pp. 420–438, 2022.
- [10] C. Mishra, L. Vanfretti, J. Delaree Jr, T. Purcell, and K. D. Jones, "Understanding the inception of 14.7 Hz oscillations emerging from a data center," *Sustainable Energy, Grids and Networks*, p. 101735, 2025.
- [11] M.-S. Ko and H. Zhu, "Wide-Area Power System Oscillations from Large-Scale AI Workloads," *ArXiv*, vol. abs/2508.16457, 2025.
- [12] T. V. Kumar, "Scalable Kubernetes Workload Orchestration for Multi-Cloud Environments," *TRJ*, vol. 11, Jan-March 2025.
- [13] H. B. Patel and N. Kansara, "Dynamic Orchestration of Multi-Cloud Resources for Scalable and Resilient AI/ML Workloads: Strategies and Frameworks," *Journal*, 2024.
- [14] S. Sharma, N. Kumar, Y. Dash, A. Dubey, and K. Devi, "Intelligent multi-cloud orchestration for AI workloads: enhancing performance and reliability," in *2024 7th International Conference on Contemporary Computing and Informatics (IC3I)*, vol. 7, pp. 1421–1426, IEEE, 2024.
- [15] H. Cui, S. Konstantinopoulos, D. Osipov, J. Wang, F. Li, K. L. Tomsovic, and J. H. Chow, "Disturbance propagation in power grids with high converter penetration," *Proceedings of the IEEE*, vol. 111, no. 7, pp. 873–890, 2022.
- [16] H. S. Erden, *Experimental and analytical investigation of the transient thermal response of air cooled data centers*. PhD thesis, Syracuse University, 2013.
- [17] P. Kundur et al., "Power system stability," *Power system stability and control*, vol. 10, no. 1, pp. 7–1, 2007.