Mitigating Rapid AI Data Center Power Fluctuations with Grid-Forming Battery Storage: Accounting for Communication Delays

Xue Lyu, Wei Du, Sheik Mohammad Mohiuddin, Brett A. Ross
Pacific Northwest National Laboratory
Richland, USA

***Abstract*—Artificial intelligence (AI)–centric data centers are growing rapidly, and their training workloads create fast, highly variable power swings that challenge grid stability. This paper designs and compares two grid-forming (GFM) battery energy storage system (BESS) control architectures for smoothing these fluctuations: (1) a point of common coupling (PCC)-feedback broadcast scheme, and (2) a coordinated plant dispatch scheme. The control architecture incorporates plant to device communication delays, and their impact on power smoothing performance is assessed. Electromagnetic transient simulations validate the proposed controls, showing that the GFM BESS effectively damps high frequency power oscillations at the data center PCC and improves frequency/voltage regulation across different grid strengths. When the communication delay is moderate, the PCC-feedback strategy outperforms the plant dispatch strategy, but it may become unstable under large delays. The results demonstrate that GFM BESS is a promising tool for integrating large AI training loads while maintaining grid stability and underline the need to minimize communication latency.**

***Index Terms*—Data center, grid-forming, power fluctuation, energy storage, communication delay.**

I. Introduction

Artificial Intelligence (AI) centric data centers have grown rapidly in scale and deployment rate in recent years. During training, power demand is highly variable: compute-intensive phases draw much more power than communication heavy phases. Because training jobs are large and highly synchronized, they produce high frequency power swings—tens to hundreds of megawatts—that propagate onto the grid and are visible in system measurements [1]. These fast, large magnitude variations can threaten bulk power system reliability by exciting frequency and voltage dynamics, straining stability and control, and degrading power quality [2]. Deploying battery energy storage systems (BESS) at data centers can mitigate these effects by smoothing power demand and thereby enhancing grid stability.

Accurate data center models are essential for characterizing large AI load behavior and assessing potential risks to the grid. Based on impedance models, [3] studies how data center power system stability relates to the design of the building distribution network. Multi level data center architectures—utility, building, backup, and rack—are discussed in [4]. An all-in-one load power emulator on a reconfigurable power-electronics hardware testbed is presented in [5] to emulate data center behavior. An isolated parallel ring bus (IPRB) datacenter microgrid with multiple uninterruptible power supplies (UPSs) is proposed in [6], where each UPS is modeled as an ideal voltage source with grid-forming (GFM) control. Different data center topologies are introduced in [7], and an MVDC data center is shown to achieve better overall efficiency. However, high-fidelity dynamic models that capture the coupled interactions among information technology equipment (ITE), cooling, UPS, BESS, and their controls and protections remain limited and need further study.

GFM control, which enables inverters to establish voltage and frequency like a voltage source, is widely regarded as a key enabler for high penetrations of power electronics interfaced resources. It is particularly well suited to BESS, which already provides a DC-side energy buffer, allowing GFM capability to be implemented largely through control/software updates rather than hardware changes [8]. For mitigating fast load swings from AI training data centers, GFM controlled BESS can provide rapid support; however, plant to device communications latency must be accounted for in controller design and performance analysis. End to end delay depends on telemetry bandwidth, network topology, routing, and the efficiency of the communications technologies and protocols used [9]. In practice, accurately measuring the delay between command issuance and delivery to individual inverter controls is challenging and site specific. It has been reported in [10] that the typical total delays of approximately 450 ms for plants commissioned before 2015, about 200 ms for installations around 2020, and roughly 100 ms for recent plants.

This work develops a data center model that explicitly represents the cooling system, UPS, and BESS, including their power converter and control dynamics. Using this model, we design two control strategies for grid-forming BESS to mitigate fast load fluctuations during AI training. In one approach, the measured point of common coupling (PCC) active power is broadcast to all devices; in the other, a plant level controller computes a coordinated active power command and dispatches it to all devices. Plant to device communication delays are modeled and incorporated into the performance and

The Pacific Northwest National Laboratory (PNNL) is operated by Battelle for the U.S. Department of Energy under Contract DE-AC05-76RL01830. The research described in this paper is based upon work supported by Energy System Co-Design with Multiple Objectives and Power Electronics (E-COMP) Initiative at PNNL. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

979-8-3315-8138-1

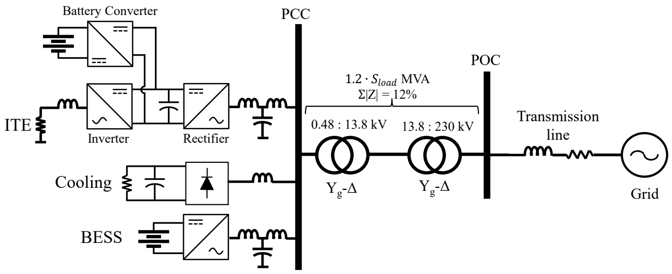

Fig. 1: Configuration of the test system.

stability assessment. Case studies using electromagnetic transient (EMT) simulations evaluate the performance of proposed strategies. Contributions of this work are: (1) a high fidelity dynamic data center model integrating cooling systems, UPS, and BESS with explicit converter and control dynamics; (2) two GFM BESS control strategies that smooth power fluctuations from high frequency variations in AI training loads; and (3) explicit modeling of plant to device communication delays and comparative evaluation of their impacts via EMT simulations.

II. System Description

Figure 1 depicts the test system configuration: the grid is represented by an infinite bus, and the data center is modeled as an aggregate of a cluster of large data centers designed for concurrent maintainability. The data center model includes cooling loads and ITE. Cooling loads are represented by diode rectifier loads supplied via variable frequency drives (VFDs), and the ITE is interfaced to the grid through an UPS. A BESS is also installed at the data center to provide fast acting support. These components connect to the PCC at 0.48 kV. Two step-up transformers interface the PCC with a 230 kV high voltage bus at the point of connection (POC). The transmission line impedance R_g and X_g can be adjusted to represent a strong or weak grid.

The UPS rectifier and inverter are modeled using an average value converter. The rectifier employs grid-following (GFL) control with current vector regulation: i_{dref} regulates the DC-link voltage, and i_{qref} is set by a reactive power controller. The reactive power reference is 0 Mvar, yielding unity power factor operation. Current limiting with i_d priority is applied to enforce the rectifier current magnitude limit. The inverter employs droop based GFM control, with the active power reference set to 0 and the reactive power reference supplied by a proportional integral (PI) regulator that controls the AC voltage. The P–f and Q–V droop coefficients are both 0.01. A hysteresis based current clipping method is employed to limit the inverter current magnitude during faults. The UPS battery is modeled as an ideal voltage source with a 5 $m\Omega$ internal resistance and is interfaced via a bidirectional buck–boost converter. The converter's modulation index is created by PI based active power control combined with a PI DC-link voltage controller. DC-link voltage regulation is achieved by setting the converter's DC reference slightly below nominal (0.99 pu) and activating the DC-voltage loop only in boost mode. With this arrangement, the battery converter contributes to DC-link support whenever the DC voltage falls below 0.99 pu, maintaining tight DC regulation even when the rectifier reaches its current limit and thereby mitigating the impact of deep AC side voltage sags at the rectifier on the inverter AC voltage.

The control strategy for the BESS inverter is described in the next section.

III. Grid-Forming BESS Control Design with Communication Delay Modeling

Each battery module includes a grid tied inverter, and the BESS plant is supervised by a plant level controller that coordinates active power dispatch across all device level units. The device inverters operate in GFM mode, establishing voltage and frequency at the data center PCC via droop control. We develop two control strategies spanning the plant and device level controllers and explicitly model communication delays between these layers.

A. Strategy A

The control architecture for Strategy A—combining plant level dispatch with device level GFM—is shown in Fig. 2. The deviation between the plant level reference P_{ref} and the measured data center PCC power is fed to an integral controller to compensate steady state power losses within the BESS plant; the integral time constant need not be set very small. The corrected reference and the PCC power measurement are then broadcast to all device controllers. The plant to device signal is sampled at rate f_s and transmitted over the communication network; the resulting delay is modeled in the Laplace domain by e^{-sT_p}, where T_p represents the delay. At the device level, controllers receive the plant command and measure the data center PCC active power. Using these signals, the i-th inverter sets its phase angle $\delta_{\text{drp},i}$ via a active power–frequency droop law:

$$\dot{\delta}_{\text{drp},i}(t) = \omega_0 \left(1 + \frac{1}{N} \Big(m_p\, \xi_i(t - T_p) - P_{\text{pcc},f}(t - T_p) \Big) + \Delta\omega_P(t) \right), \tag{1}$$

$$\dot{\xi}_i(t) = k_i \Big(P_{\text{ref}}(t) - P_{\text{pcc},f}(t) \Big). \tag{2}$$

where P_{ref} is the plant level active power reference, P_{pcc} is the measured active power from the data center PCC (after a low pass filter), N is the number of inverters, m_p is the droop coefficient, and ω_0 is the nominal frequency in rad/s. $\Delta\omega_p$ is generated by an anti-windup PI based active power limiting controller (Fig. 3) that ensures the device output remains within its operating limits, $P_{\text{min}} < P_i < P_{\text{max}}$.

A Q–V droop control is used to establish the inverter internal voltage $E_{\text{drp},i}$. The reactive power reference for each device is assumed to be zero in this work. In addition to the droop control voltage reference, a PI controller is implemented to regulate the data-center PCC voltage. To ensure that each inverter's reactive power remains within its capacity limits, an overload mitigation controller enforces the constraintQ_{max} <

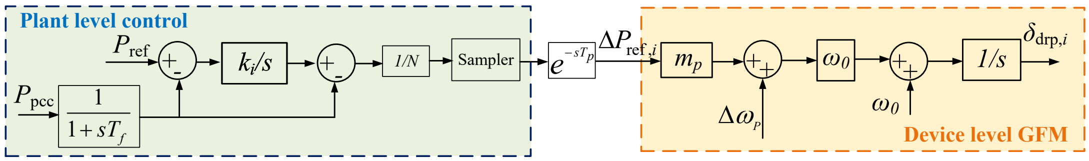

Fig. 2: Control architecture of Strategy A.

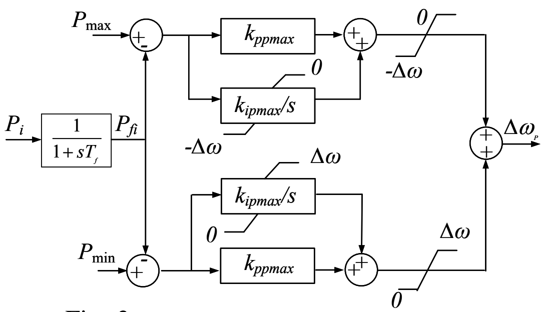

Fig. 3: Active power overload mitigation [11].

$Q_i < Q_{\max}$. In this context, the i-th inverter internal voltage inverter internal voltage is expressed as:

$$E_{\text{drp},i} = (k_{pv} + \frac{k_{iv}}{s})(V_{\text{ref}} - m_q Q_f + \Delta V_Q - V_f), \quad (3)$$

where k_{pv} and k_{iv} are PI control gains, $V_{\text{ref}} = 1\ p.u.$ is the voltage reference, m_q is the droop coefficient, Q_f and V_f are the filtered measurements of reactive power and PCC voltage, respectively, and ΔV_Q is the anti-windup signal from the reactive power limiter.

Direct droop control without an inner current loop is employed for the inverters, where the modulation index is directly generated from $\delta_{\text{drp},i}$ and $E_{\text{drp},i}$. A hysteresis-based current clamping method is used to limit fault currents by blocking inverter switching whenever the instantaneous current magnitude exceeds the limit $I_{\max}$.

B. Strategy B

The control architecture for Strategy B is shown in Fig. 4. At the plant level, the error between the active-power reference and the measured active power at the data center PCC is processed by a PI controller. The resulting power command is then sampled and broadcast to each device, subject to communication delay. At the device level, controllers receive the plant command and measure the active power at the inverter terminal. Using these signals, the $i-$th inverter phase angle is computed from the following equations:

$$\dot{\delta}_{\text{drp},i}(t) = \omega_0\Big(1 + m_p\big(P_{\text{ref},i}(t) - P_i(t)\big) + \Delta\omega_P(t)\Big) \quad (4)$$

$$\dot{x}_i(t) = P_{\text{ref}}(t) - P_{\text{pcc},f}(t), \quad (5)$$

$$P_{\text{ref},i}(t) = \frac{1}{N}\Bigg(k_p\big(P_{\text{ref}}(t-T_p) - P_{\text{pcc},f}(t-T_p)\big) + k_i x_i(t-T_p)\Bigg) \quad (6)$$

where k_p and k_i are the plant-level PI controller gains.

Strategy B adopts the same PCC voltage-magnitude control as Strategy A to maintain the PCC voltage stability.

IV. Case Studies

The test system shown in Fig. 1 is implemented in PSCAD to evaluate the performance of Strategies A and B. The data center model includes 40 MW of cooling load and 400 MW of ITE load. The ITE load consists of a 200 MW resting (baseline) component with superimposed fast (10 Hz) and slow (0.5 Hz) power variation components, reflecting differing demand during compute-intensive and communication heavy phases. The BESS plant is rated at 300 MW and is composed of 2 MW modules. The $P-f$ and $Q-V$ droop coefficients, m_p and m_q, are set to 0.01. We evaluate control performance under varying grid strengths and quantify the impact of communication delays between the BESS plant and device level controllers.

A. Impact of Grid Strength

The transmission grid line impedance X_g and R_g are tuned to achieve a target short circuit ratio (SCR). We first consider a weak grid with SCR = 2. Figure 5 shows the PCC active power demand, PCC voltage magnitude, and PCC frequency under three scenarios: no BESS, BESS with Strategy A, and BESS with Strategy B. The communication delays are neglected for both strategies in these cases. Without BESS, a 95.5 MVar shunt capacitor is installed at the PCC to help maintain voltage stability. Even so, PCC voltage and frequency exhibit large fluctuations driven by the load swings: voltage dips to 0.8 pu, the data center's peak demand is not fully supplied due to voltage stability limits on active power transfer, and frequency excursions reach 62 Hz (peak) and 58 Hz (nadir). A peak near 62 Hz approaches the NERC PRC-024-4 frequency ride-through boundary and may trigger generator protective tripping [12]. With BESS, the shunt capacitor is unnecessary and both PCC voltage and frequency excursions are substantially reduced for Strategies A and B: the voltage minimum improves to 0.95 pu, the frequency peak drops to 60.2 Hz, and the nadir rises to 59.8 Hz. Compared with Strategy A, Strategy B exhibits larger active power transients during the initial cycles when the AI training fluctuation begins and shows slower settling when the fluctuation ends. Under Strategy A, the PCC active power swing (maximum to minimum) is reduced to 25 MW.

The SCR = 3 case is shown in Fig. 6. Without BESS, a 69.5 MVar shunt capacitor is installed at the PCC for voltage support. AI training load swings still cause significant PCC excursions: the voltage nadir drops to 0.86 pu, the frequency nadir to 58.6 Hz, and the frequency peak to 62 Hz. With BESS, both voltage and frequency excursions are substantially reduced under Strategies A and B. Strategy A better suppresses PCC active power variations during the initial cycles of the AI

979-8-3315-8138-1

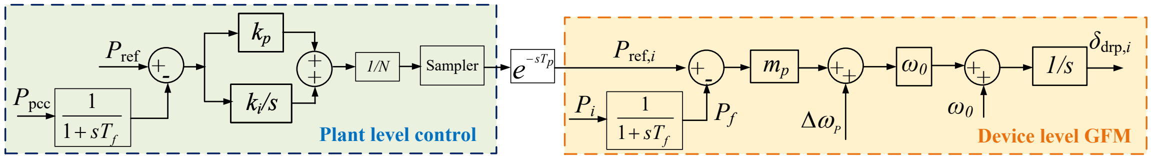

Fig. 4: Control architecture of Strategy B.

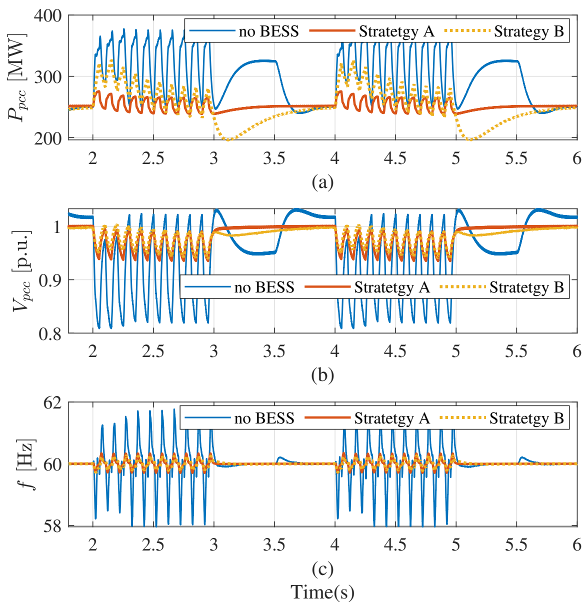

Fig. 5: Simulation results for SCR = 2, (a) active power demand at the data center PCC, (b) PCC voltage magnitude, and (c) PCC frequency.

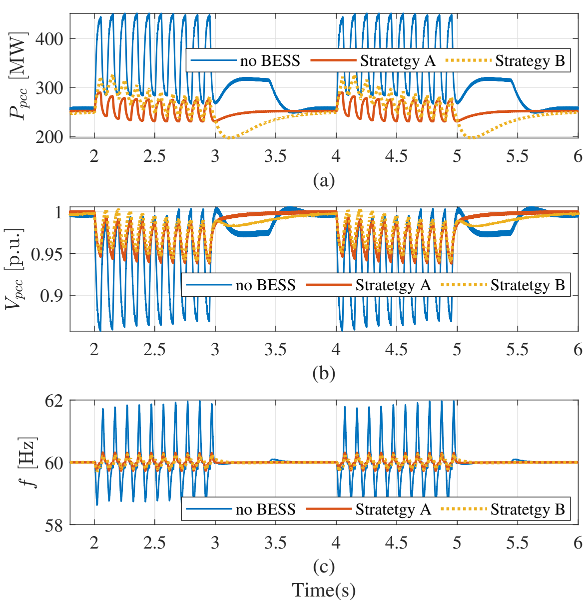

Fig. 6: Simulation results for SCR = 3, (a) active power demand at the data center PCC, (b) PCC voltage magnitude, and (c) PCC frequency.

training transient than Strategy B. Under Strategy A, the PCC active power swing (maximum to minimum) is reduced by approximately 77% relative to the no-BESS case.

The SCR = 10 case is shown in Fig. 7. With a BESS, a 69.5 MVar shunt capacitor is installed at the PCC for voltage support. In this strong grid condition, PCC voltage and

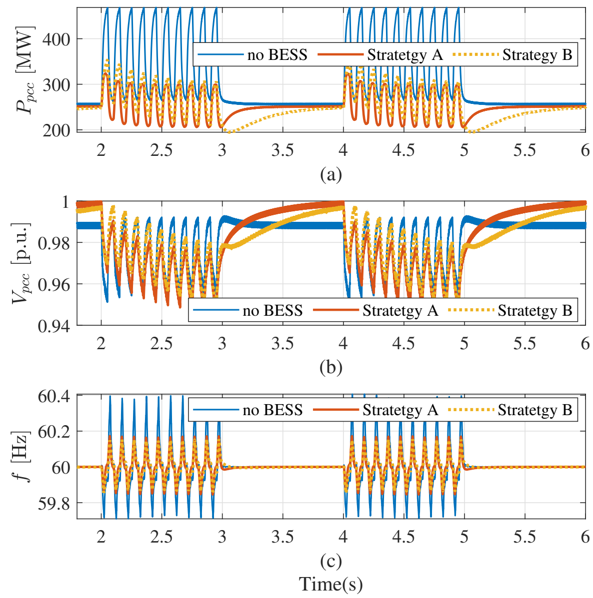

Fig. 7: Simulation results for SCR = 10, (a) active power demand at the data center PCC, (b) PCC voltage magnitude, and (c) PCC frequency.

frequency fluctuations are much smaller than in the SCR = 2 and SCR = 3 cases, as the system can absorb fast load swings more effectively. Installing the BESS further mitigates PCC active-power variations. Strategy A continues to outperform Strategy B during the initial cycles of the AI training transient. Under Strategy A, the PCC active power swing (maximum to minimum) is reduced by approximately 54% relative to the no-BESS case.

Overall, strategy A significantly damp AI training-induced load oscillations, reducing PCC voltage/frequency excursions and smoothing active power injections across grid strengths. Benefits are greatest in weak grids, where BESS control significantly reduces peak-to-peak power swings. In a strong grid, the performance gap between two strategies narrows since the grid's stiffness and damping already suppress oscillations, making both strategies perform similarly.

B. Impact of Communication Delay

We assess communications delay impacts via simulation with SCR = 3. For Strategy A, Fig. 8 compares performance across delay settings. With delays of 0, 50 ms, and 100 ms, the PCC active power demand remains essentially unchanged, indicating that Strategy A maintains effective smoothing under these conditions. At 200 ms, however, the PCC active power becomes oscillatory, indicating loss of delay robustness for large measurement/communication delays. To further interpret this behaviour, we compute the eigenvalues of a Padé(1)-based small-signal model (finite-dimensional approximation of the

979-8-3315-8138-1

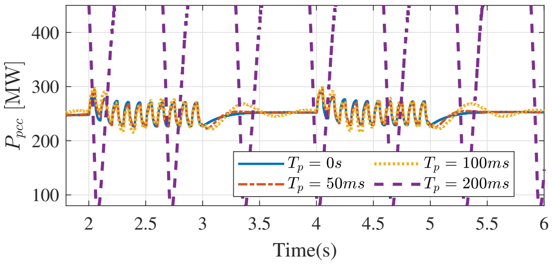

Fig. 8: Active power demand at the data center PCC for Strategy A under varying communication delays.

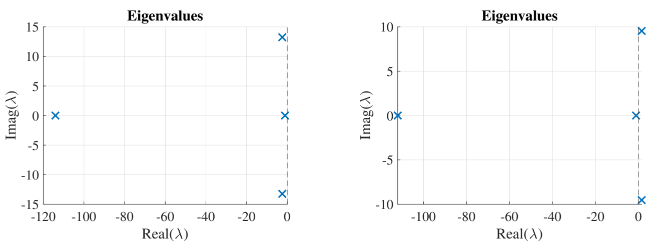

(a) $T_p = 0.1$ s.

(b) $T_p = 0.2$ s.

Fig. 9: Eigenvalues of the Padé(1) small-signal model for Strategy A.

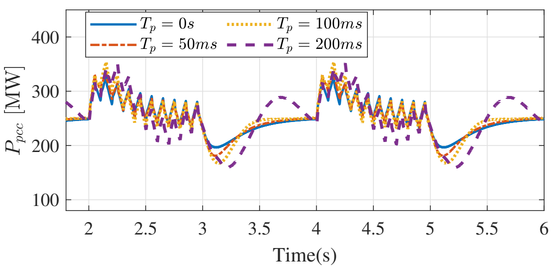

Fig. 10: Active power demand at the data center PCC for Strategy B under varying communication delays.

delay). The model predicts two eigenvalues with positive real part for $T_p = 0.2$ s (unstable) as shown in Fig. 9b. When $T_p = 0.1$ s, all eigenvalues lie in the left-half plane (stable), as shown in Fig. 9a.

For Strategy B, Fig. 10 compares PCC active power demand across delay settings. Increasing plant to device communication delay amplifies initial transients at the onset of AI training load fluctuations and lengthens settling time. Compared with Strategy A, Strategy B is more robust to communication delay. Small-signal eigenvalue analysis (omitted due to space) further confirms stability when T_p=0.2 s.

Overall, both Strategy A and Strategy B deliver robust smoothing at the data center PCC for delays of 50–100 ms, with Strategy A showing minimal change in performance within this range. At 200 ms, the smoothing effect degrades: Strategy A becomes oscillatory, while Strategy B exhibits larger transients and extended settling. These findings highlight the need to minimize BESS plant to device communication delays in data center implementations.

V. Conclusion

This work develops a high-fidelity data center model that explicitly represents the UPS, cooling system, and BESS, including power converter and control dynamics. Using this model, we propose and evaluate two GFM based BESS control strategies to mitigate fast load fluctuations from AI training: Strategy A, where the plant controller broadcasts the PCC active power measurement to all devices; and Strategy B, a plant level dispatch that computes and sends coordinated setpoints to devices. EMT simulations across grid strengths show that BESS control substantially reduces PCC voltage/frequency excursions and smooths active power demands at data center PCC, with the largest benefits in weak grids. Strategy A outperforms Strategy B when the communication delay is small, yielding smaller initial transients and faster settling, with smoothing benefits more significant in weak grids. Delay sensitivity analysis indicates robust smoothing for both strategies with plant to device delays up to 100 ms; at 200 ms, Strategy A becomes oscillatory/unstable, while Strategy B exhibits larger transients and longer settling. These results highlight the importance of minimizing communication latency and support adopting Strategy A where delays can be tightly bounded.

References

[1] E. Choukse, B. Warrier, S. Heath, L. Belmont, A. Zhao, H. A. Khan, B. Harry, M. Kappel, R. J. Hewett, K. Datta *et al.*, "Power stabilization for ai training datacenters," *arXiv preprint arXiv:2508.14318*, 2025.

[2] N. L. L. T. Force, "Characteristics and risks of emerging large loads," *North American Electric Reliability Corporation (NERC), White Paper, July*, 2025.

[3] 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.

[4] P. T. Krein, "Data center challenges and their power electronics," *CPSS Transactions on Power Electronics and Applications*, vol. 2, no. 1, pp. 39–46, 2017.

[5] J. Sun, S. Wang, J. Wang, and L. M. Tolbert, "Dynamic model and converter-based emulator of a data center power distribution system," *IEEE Transactions on Power Electronics*, vol. 37, no. 7, pp. 8420–8432, 2022.

[6] L. L. Qi, H. Suryanarayana, Y. Zhang, T. Jiang, and S. Colombi, "Modeling, simulation, and protection of grid forming inverter-based ring configuration datacenter," in *2022 IEEE Industry Applications Society Annual Meeting (IAS)*. IEEE, 2022, pp. 1–7.

[7] D. Siemaszko, S. Heinig, C. Yuan, J. Iglesias, and M. Mogorovic, "MVDC distribution concept for green data centers: Achieving the sustainability roadmap with highest efficiency," in *PCIM Europe 2023; International Exhibition and Conference for Power Electronics, Intelligent Motion, Renewable Energy and Energy Management*. VDE, 2023, pp. 1–9.

[8] NERC, "White paper: Grid forming functional specifications for BPS-connected battery energy storage systems," https://www.nerc.com/comm/RSTC_Reliability_Guidelines/White_Paper_GFM_Functional_Specification.pdf, 2023.

[9] B. Bhattarai, L. Marinovici, M. Touhiduzzaman, F. K. Tuffner, K. P. Schneider, J. Xie, P. Thekkumparambath Mana, W. Du, and A. Fisher, "Studying impacts of communication system performance on dynamic stability of networked microgrid," *IET Smart Grid*, vol. 3, no. 5, pp. 667–676, 2020.

[10] N. Miller *et al.*, "Diagnosis and mitigation of observed oscillations in ibr-dominant power systems: a practical guide," Technical report. Energy Systems Integration Group, Tech. Rep., 2024.

[11] W. Du, "Model specification of droop-controlled, grid-forming inverters (REGFM_A1)," Pacific Northwest National Laboratory (PNNL), Richland, WA (United States), Tech. Rep., 09 2023. [Online]. Available: https://www.osti.gov/biblio/2229442

[12] NREC, "PRC-024-4: Transmission-connected resources — frequency and voltage protective relay settings," https://www.nerc.com/pa/Stand/202002_Transmissionconnected_Resources_DL/2020-02_PRC-024-4_Clean_03272024.pdf, 2024.

979-8-3315-8138-1