A Unified Dynamic Modeling and Data-Driven Workload Framework for Large Electric Loads

Siyu Lu, **Chenhan Xiao**, **Yang Weng**
Department of Electrical, Computer and Energy Engineering, Arizona State University, AZ, USA
Email: {siyulu12, cxiao20, yang.weng}@asu.edu

***Abstract*—Large Electric Loads (LELs) such as data centers, cryptocurrency mining facilities, and hydrogen electrolyzers introduce fast, workload-driven, and protection-dependent dynamics that challenge traditional load modeling in modern power systems. To enable consistent analysis across different LEL types and provide a scalable foundation for incorporating emerging large-load technologies into transient stability studies, we develop a unified dynamic modeling framework that captures their shared characteristics. These include duty-idle workload surges, cooling and auxiliary demand, and UPS or dc-link based ride-through behavior. To accurately represent the server component, which dominates total LEL demand, we introduce a data-driven workload model obtained by calibrating an Ornstein-Uhlenbeck (OU) stochastic process to real CPU and GPU utilization traces from the MIT Supercloud dataset. The OU-based model preserves burstiness, autocorrelation, and mean-reverting behavior, offering a substantial improvement over traditional Poisson-based workload representations, and can be re-parameterized by system operators for site-specific tuning. The unified LEL model is integrated into the ANDES simulation platform and evaluated through extensive studies. The simulations reveal system-level phenomena, including persistent ride-through disturbances, prolonged recovery in larger grids, and characteristic voltage-frequency interactions. These results provide a more realistic basis for assessing the stability implications of emerging large-load deployments.**

***Index Terms*—Large electric loads, data centers, dynamic load modeling, transient stability**

I. Introduction

The rapid growth of *Large Electric Loads (LELs)*, including data centers, cryptocurrency mining facilities, and hydrogen electrolyzers [1], [2], has introduced new modeling and operational challenges in modern power systems [1]–[4]. Unlike traditional large loads (e.g., industrial induction-motor loads) that vary smoothly, LELs can change their power demand by tens of MW within milliseconds and often operate at capacities exceeding 75 MW [1], [2]. Their power consumption is strongly workload-driven, as CPU/GPU clusters respond almost instantaneously to computation job arrivals and task scheduling events. LELs also employ proactive protection mechanisms, such as uninterrupted power supply (UPS) systems, dc-link energy storage, and fast ride-through controls, which significantly shape their interactions with the grid during voltage sags, frequency dips, and other disturbances [5]–[7]. Recent analyses from ERCOT and other utilities show that these protection mechanisms can lead to large, abrupt load losses during grid disturbances, and can further induce reactive power swings, delayed recovery, and even oscillatory behavior [8]–[11]. As LELs continue to expand and cluster geographically, they pose stability concerns that cannot be adequately addressed using traditional load models [7], [12].

Traditional load models, including ZIP, induction-motor, and static exponential formulations, assume slowly varying aggregate behavior and do not capture the fast, workload-driven, and protection-dependent dynamics exhibited by LELs. Recent advances have improved data center models by integrating CPU/GPU load surges with cooling demand, auxiliary consumption, and UPS-based ride-through mechanisms [5], [7], and complementary studies have examined their system-level impacts, including forced oscillations, reactive-power sensitivity, and degraded low-frequency stability margins [8]–[11]. However, most existing work focuses solely on data centers, even though NERC and ERCOT report that cryptocurrency mining facilities and hydrogen electrolyzers exhibit similar duty-idle workload patterns and employ comparable protection mechanisms, including UPS or dc-link buffering [1], [2]. These reports also highlight that cryptocurrency mining currently accounts for the largest share of LEL demand (see Fig. 1). Moreover, existing modeling studies typically rely on manually constructed workload patterns rather than aligning them with real CPU/GPU utilization data.

To address these gaps, this paper first develops a unified dynamic load modeling that captures three major categories of LELs (data centers, cryptocurrency mining facilities, and hydrogen electrolyzers [1], [2]) using a common set of electrodynamic components and protection behaviors. The framework captures their shared structural elements, including duty-idle workload surges, cooling and auxiliary demand, and UPS or dc-link based ride-through mechanisms, while allowing type-specific customization through parameter differentiation. This unified framework enables consistent analysis across LEL types and provides a scalable foundation for incorporating emerging LEL technologies into transient stability studies.

A key contribution of this work is the introduction of a data-driven workload model for the server component, which is the dominant portion of LEL demand and typically accounts for 60-80% percent of total power consumption. To accurately represent the fast and bursty CPU/GPU utilization surges that drive server power draw, we calibrate an Ornstein-Uhlenbeck (OU) stochastic process [13], [14] to real utilization traces obtained from the MIT Supercloud dataset [15]. Compared with traditional Poisson-based workload models [5] that assume independent increments and limited temporal correlation, the OU process captures mean-reversion, autocorrelation, and continuous-valued fluctuations that closely match the temporal structure of practical computational workloads. The calibration is achieved by mapping the empirical mean, variance, and autocorrelation statistics of the Supercloud traces directly

979-8-3315-8138-1

to the OU drift, diffusion, and mean reversion parameters. Importantly, the same procedure allows system operators to re-estimate these parameters using their own server data, including data subject to privacy constraints, enabling site-specific customization of the workload model.

To evaluate the proposed LEL modeling framework, we integrate all components into the ANDES dynamic simulation platform [16] and conduct extensive time-domain studies on representative transmission networks. The simulations reveal several system-level phenomena that are not captured by traditional load models, including persistent ride-through disturbance effects, prolonged recovery times in larger grids, and characteristic voltage-frequency interactions driven by workload variability and protection behavior. These observations demonstrate that the proposed modeling approach provides a more realistic foundation for assessing the stability implications of emerging LEL deployments. The remainder of the paper is organized as follows. Section II presents the unified dynamic modeling framework for the three LEL types. Section III evaluates the proposed models through time-domain simulations in the ANDES platform. Section IV concludes the paper.

II. Unified Dynamic Modeling of LELs

This section develops a unified, data-guided dynamic representation of Large Electric Loads (LELs), including data centers (DC), cryptocurrency mining (CM) facilities, and hydrogen electrolyzers (HE). The goal is to construct load models that accurately capture both their steady-state characteristics and their dynamic responses to voltage and frequency disturbances. Although these LEL types differ in purpose and internal processes, they exhibit three common behavioral patterns that motivate a unified modeling approach (see Fig. 1). First, all three categories contain a core load (mining rigs for CM, IT servers for DC, and the electrolyzer load for HE) that follows a duty-idle or process-driven workload structure in which power demand can change rapidly due to computation jobs, scheduling events, or process ramps (Section II-A). Second, each includes supporting subsystems such as cooling, pumps, or auxiliary electronics that monitor, stabilize, and supply the core load (Section II-B). Third, all LELs interface with the grid through local energy storage or fast protection mechanisms that shield the core load from voltage and frequency disturbances (Section II-C). These mechanisms include power supply unit (PSU)-based electronic protection in CM, UPS-based ride-through in DC, and converter and dc-link control in HE. These shared characteristics provide the foundation for a unified dynamic modeling framework.

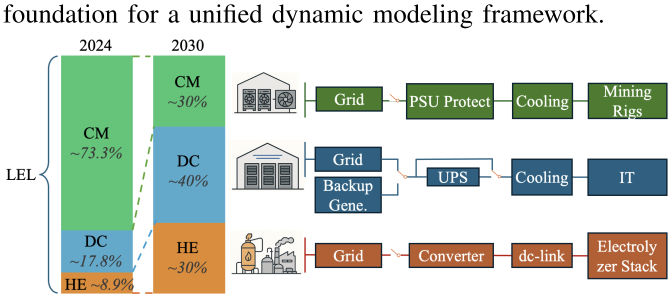

Fig. 1. 2024 and 2030 LEL load ratios reported from ERCOT [1], [2].

A. Duty-Idle Workload Modeling

The workload dynamics of LELs exhibit alternating periods of high activity and partial idling, driven by computation job arrivals in DC and CM facilities or process ramps in HE systems. These fluctuations occur on fast timescales and account for the majority of variability in instantaneous power draw, making workload-driven behavior a dominant contributor to LEL electrodynamics. To explicitly represent this structure, we model the core load using a duty-idle formulation that maps workload intensity to active power consumption using

$$p_{\text{work}}(t) = p_{\text{base}} + \eta(t)\left(p_{\text{full}} - p_{\text{base}}\right), \quad \eta(t) \in [0,1], \tag{1}$$

where p_{base} is the idle/standby draw, p_{full} the rated draw at full duty, and $\eta(t)$ represents the instantaneous utilization level. In DC and CM facilities, $\eta(t)$ reflects CPU/GPU or ASIC duty cycles, while in HE systems it captures process-driven ramping of the electrolyzer load. Analysis of the MIT Supercloud dataset [15] further shows that server power consumption is strongly correlated with CPU and GPU utilization, providing empirical support for the duty-idle formulation in Eq. (1).

To represent the temporally correlated and statistically bounded nature of workload fluctuations, we model $\eta(t)$ as an Ornstein-Uhlenbeck (OU) process, a formulation widely used for load variability in power-system dynamics [13], [14]. Prior workload models often rely on purely deterministic or Poisson-driven duty-cycle patterns [5] , which generate abrupt and memoryless fluctuations that can be overly stochastic and difficult to tune for realistic variability. Moreover, Poisson-based formulations provide no direct mechanism for aligning the workload trajectory with empirical server logs. In contrast, the OU process incorporates mean reversion, temporal autocorrelation, and continuous-valued perturbations [13], [14], enabling both smoother trajectories for differential-algebraic simulation and a principled approach for parameter calibration using real workload traces. The OU process is

$$\tau_\eta \dot{\eta} = -(\eta - \mu_\eta) + \xi(t) + \sum_k A_k\, \delta(t - t_k), \tag{2}$$

where $\mu_\eta \in [0,1]$ is the nominal utilization level, $\tau_\eta > 0$ determines the rate at which τ_η reverts toward μ_η. $\xi(t)$ is a zero-mean stochastic perturbation capturing small continuous fluctuations in workload intensity. The event times $\{t_k\}$ follow a Poisson process with rate λ, and each Dirac impulse $\delta(t - t_k)$ introduces an instantaneous workload burst of magnitude A_k/τ_η, after which $\eta(t)$ is clipped to $[0,1]$.

When real workload utilization $\eta(t)$ is available—such as the CPU/GPU traces provided by the MIT Supercloud dataset [15]—the OU model enables calibration of its parameters for site-specific customization. To implement this calibration, we work with a discrete-time approximation of the OU process. Let Δt denote the sampling interval and define $\theta = 1/\tau_\eta$. The continuous-time dynamics in Eq. (2) are then equivalent to the first-order autoregressive model $\eta_{k+1} = a\,\eta_k + b + \varepsilon_k$, where $a = e^{-\theta \Delta t}, b = \mu_\eta(1-a), \varepsilon_k \sim \mathcal{N}(0, \sigma_d^2)$. Given a utilization time series $\{\eta_k\}$ extracted from server logs (for example, CPU/GPU utilizations from the MIT Supercloud dataset [15]), we fit this AR(1) model by linear

979-8-3315-8138-1

regression of η_{k+1} on η_k and obtain estimates $\hat{a}$, $\hat{b}$, and residual variance $\hat{\sigma}_d^2$. The OU parameters are then recovered via $\hat{\theta} = -\frac{1}{\Delta t}\ln\hat{a}, \hat{\mu}_\eta = \frac{\hat{b}}{1-\hat{a}}, \hat{\sigma}_\xi = \sqrt{\frac{2\hat{\theta}\,\hat{\sigma}_d^2}{1-\hat{a}^2}}$, which determine the mean-reversion rate, nominal utilization, and noise intensity in Eq. (2). This AR(1)-based calibration procedure provides a simple data-driven adapter that maps empirical utilization logs to OU parameters, enabling the workload model to be tuned to specific LEL facilities. When no empirical workload trace $\eta(t)$ is available, we provide representative OU parameters for the three LEL categories, as summarized in Table I.

TABLE I
OU PARAMETERS ADAPTED FOR DIFFERENT LEL TYPES.

Para.	DC	CM	HE
$\mu_\eta(t)$	Daily variation	Near const.	Renewable/market signal
τ_η	10–60 s	200–600 s	60–300 s
$\xi(t)$	Moderate noise	Very low noise	Moderate noise
A_k	$\sim$ LogN	Small $A_k < 0$	$\sim$ Unif
λ	0.5–3/min	0.2–2/hr	1–6/hr

B. Cooling and auxiliary systems

To support safe operating temperatures for core load components, cooling units are essential across all LEL types. These systems are predominantly motor-driven (e.g., chillers, compressors, pumps) [5] and are modeled as aggregated induction motors following the WECC CMPLDW templates [17]. Motor variables are expressed in the standard synchronous dq-frame, where v_{ds}, v_{qs} and i_{ds}, i_{qs} denote the stator voltages and currents. The cooling active power is

$$p_{\text{cool}}(t) = v_{ds}(t)\, i_{ds}(t) + v_{qs}(t)\, i_{qs}(t), \tag{3}$$

Under voltage sags, the motors exhibit increased reactive power demand and may stall when $|V| < V_{\text{stall}}$ for τ_{stall} or when slip exceeds its protection threshold; stalled units are then disconnected or replaced by a high-impedance block until voltage recovers. All parameters and protection settings follow the WECC CMPLDW documentation [17], with initialization at steady-state slip and bus voltage.

In addition to cooling systems, auxiliary loads such as lighting, control units, security systems, and small electronic devices are modeled as static ZIP [18] components:

$$p_{\text{aux}}(V) = p_{\text{aux},0}(\alpha_Z(\tfrac{V}{V_0})^2 + \alpha_I(\tfrac{V}{V_0}) + \alpha_P), \tag{4}$$

where $p_{\text{aux},0}$ denotes the nominal auxiliary active power at the reference voltage V_0. The coefficients α_Z, α_I, and α_P represent the fractional contributions of the *constant-impedance*, *constant-current*, and *constant-power* components, respectively, satisfying $\alpha_Z + \alpha_I + \alpha_P = 1$. The total LEL demand combines workload, cooling, and auxiliary components:

$$p_{\text{load}} = \beta_{\text{work}}\, p_{\text{work}} + \beta_{\text{cool}}\, p_{\text{cool}} + \beta_{\text{aux}}\, p_{\text{aux}}, \tag{5}$$

where fractions β_{work}, β_{cool}, and β_{aux} differ across DC, CM, and HE facilities.

C. Local Protection Mechanisms

All LELs interface with the grid through fast local protection or embedded energy-storage mechanisms designed to shield sensitive electronics from voltage and frequency disturbances and to manage controlled reconnection. Despite technology-specific implementations, these mechanisms share the goals of reducing electrical stress, preventing uncontrolled shutdowns, and coordinating post-disturbance recovery.

Data centers employ UPS systems [1], [4] that instantaneously transfer IT loads to battery or flywheel reserves when voltage or frequency exceed tolerance bands, effectively dropping grid-side demand to near zero. Cryptocurrency mining facilities—which typically operate without dedicated energy storage—use PSU-level electronic protection and feeder relays; voltage sags trigger rapid throttling or sectional tripping of mining rigs, producing abrupt real-power reductions. Hydrogen electrolyzers rely on converter-integrated and PLC-supervised protections: disturbances activate current limiting and dc-link buffering, and if deviations persist, a contactor disconnects the stack, followed by a controlled ramp-up once nominal conditions are restored.

Despite these technological differences, the protective response follows a consistent hierarchical pattern that can be represented by a unified hybrid load model. When voltage V or frequency ω violates the ride-through envelope, the load reduces or suspends its grid-side consumption according to

$$\tilde{p}_{\text{load}} = \kappa\, p_{\text{load}}, \text{ if } |V - V_{\text{ref}}| > \Delta V \text{ or } |\omega - \omega_{\text{ref}}| > \Delta\omega, \tag{6}$$

where $\kappa \in [0,1]$ is the retained load fraction, and $(V_{\text{ref}}, \omega_{\text{ref}}, \Delta V, \Delta\omega)$ define the protection thresholds. The specific parameters differ by technology: data centers exhibit $\kappa \approx 0$ due to complete UPS isolation, mining farms typically retain $\kappa \in [0.3, 0.7]$ under throttling or partial feeder trips, and electrolyzers adjust κ continuously through converter current limiting before full disconnection.

III. NUMERICAL RESULTS

This section describes the dynamic simulation setup used to assess the proposed LEL models and their impacts on grid stability. All simulations are performed in the open-source `ANDES` platform [16], which supports user-defined differential-algebraic load models and time-domain transient stability analysis with configurable disturbance events. The unified LEL representations from Section II are embedded into standard benchmark systems, including the IEEE 14-bus, 39-bus, NPCC 140-bus, and WECC 179-bus grids, and subjected to controlled voltage and frequency perturbations to examine their dynamic responses.

Implementation details. Simulations are conducted in `Python 3.12` using `ANDES` version 1.9.3. The time-step is fixed at $\Delta t = 2$ ms and total simulation horizon is $T_{\text{sim}} = 40$ s. For each test system, a subset of buses in each system are randomly assigned as LEL sites. The fraction of LEL buses is denoted as ξ_{LEL}. Each selected bus is randomly assigned one of the three LEL types (DC, CM, or HE). Unless otherwise specified, we set $\xi_{\text{LEL}} = 20\%$; Section III-C examines sensitivity to this parameter. To evaluate system resilience, a three-phase short-circuit fault is applied to a selected transmission bus at time $t = 5.0$ s and cleared after 100 ms. This disturbance produces voltage depressions and transient oscillations across the network, enabling assessment of LEL dynamic behavior and ride-through performance.

979-8-3315-8138-1

A. Data-Driven OU Workload Modeling vs. Poisson Baseline

We first validate the proposed OU-based duty-idle workload model against real CPU/GPU utilization traces from the MIT Supercloud dataset [15]. Fig. 2 compares four workload trajectories: (i) the real MIT Supercloud power draw, (ii) a Poisson-based arrival model following [5], (iii) an OU process using the default (uncalibrated) parameters in Table I, and (iv) an OU process with parameters calibrated from the MIT traces. The MIT data exhibit pronounced duty-idle behavior in which the power draw remains near a baseline for extended periods and undergoes short, workload-driven surges. The Poisson-based model captures the occurrence of surges but produces overly stochastic and memoryless fluctuations. In contrast, the calibrated OU model reproduces the empirical workload amplitude, the duration of duty-idle intervals, and the observed temporal smoothness, while the uncalibrated OU model follows the same qualitative structure but deviates in variance and correlation time. These results confirm that the OU formulation can be effectively calibrated to real workload traces, enabling data-driven customization and improved accuracy in LEL dynamic simulations.

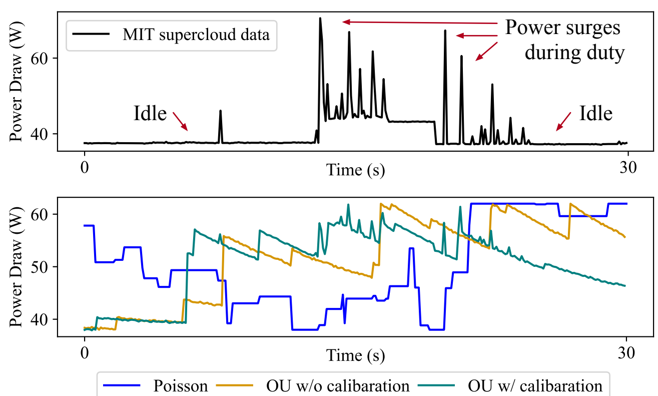

Fig. 2. MIT Supercloud workload traces versus simulated workload models.

B. System-Level Impact of LEL Ride-Through

This subsection investigates how LEL protection actions (e.g., UPS-driven trip) and reconnection behavior shape post-fault recovery under different operating conditions, LEL sizes, and system scales. We consider a single LEL site in the network and fix its reconnection attempt at $t = 20$s to isolate the impact of its protection logic on system dynamics.

We first compare the IEEE 39-bus system with and without an LEL at Bus 4. In the baseline case (no LEL), voltages and frequency recover smoothly after fault. With a 75 MW LEL, however, the post-fault voltage at Bus 4 drops below the 0.7 p.u. ride-through threshold, triggering a UPS-driven disconnection. The sudden loss of active and reactive demand causes frequency overshoot and local voltage swell, producing oscillations that exceed operational limits for several seconds. This illustrates how LEL protective actions, while locally beneficial, can induce secondary systemwide disturbances.

To study how LEL size influences these effects, we increase the Bus-4 LEL from 75 MW to 225 MW and adopt a stricter voltage ride-through threshold (0.9 p.u.). As shown in Fig. 4, at 75 MW the reconnection produces only a modest voltage dip, allowing successful resynchronization. At 225 MW, however, the reconnection introduces a much larger step change in power draw, producing an immediate voltage depression that again violates the protection threshold and triggers an instantaneous retrip. These results show that larger LELs, especially when paired with sensitive ride-through settings, can turn an benign reconnection into a destabilizing event.

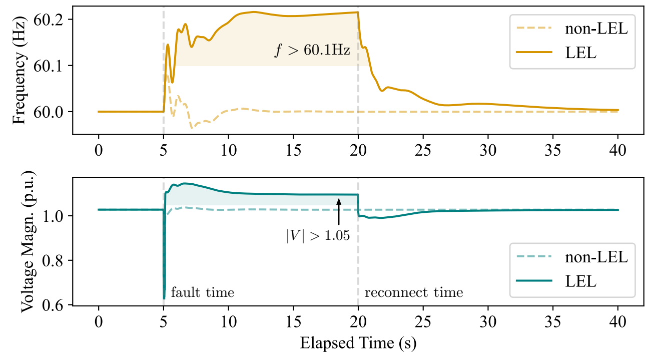

Fig. 3. Post-fault frequency and voltage trajectories in IEEE 39-bus system.

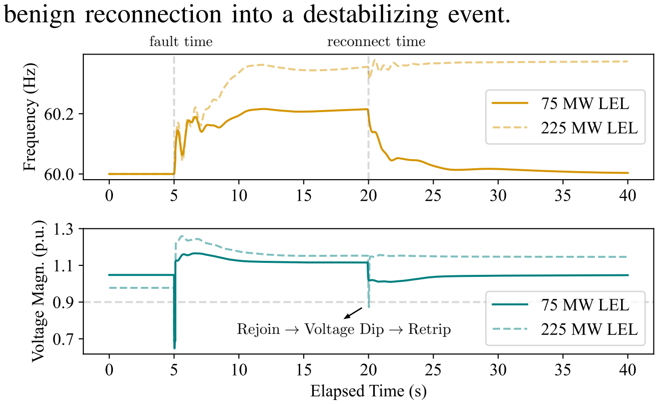

Fig. 4. Comparison of a 225 MW LEL to a 75 MW LEL.

To understand how these behaviors scale with system size, we simulate identical LEL trip-reconnect events across four networks. Results in Fig. 5 show that recovery times grow noticeably with system scale. In the NPCC 140-bus and WECC 179-bus systems, the LEL-induced trip and reconnection excite multiple local and inter-area modes in the 0.1-1 Hz range, many of which exhibit slow damping. As a result, the system experiences prolonged voltage and frequency oscillations long after the fault has cleared. This behavior highlights that LEL protective actions have a more persistent and diffuse footprint in large, interconnected grids.

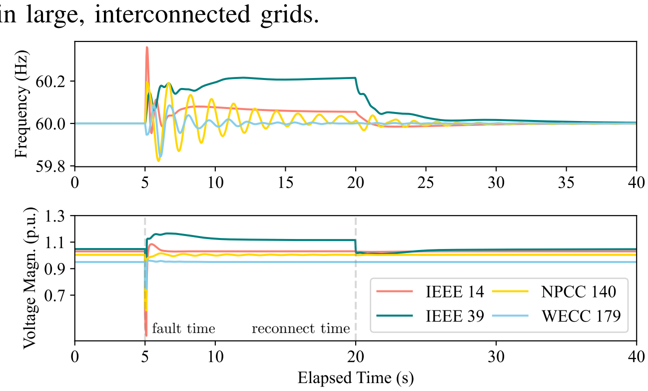

Fig. 5. Comparison of post-fault recovery across systems.

979-8-3315-8138-1

C. Effect of LEL Ratios on Grid Stability

As LELs grow in size and geographic concentration, their penetration level can significantly influence system stability. This subsection examines how increasing LEL ratios affect the severity and propagation of disturbance-driven events.

To quantify the extent of LEL-driven propagation, we employ two complementary indices. The first is the *disconnection propagation index* (DPI), which measures how much additional LEL load is lost after an initial failure. Let i be the initial LEL disconnected at time $t_0 \approx 5$ (due to injected fault), and let the reconnection delay be $t_r = 5s$. Let $\mathcal{T}$ denote the set of all LELs that subsequently disconnect within the time window $[t_0, t_0 + 120\text{s}]$. The DPI is defined as $\text{DPI} = \frac{\sum_{k \in \mathcal{T} \setminus \{i\}} P_k}{P_i}$, where P_k is the pre-trip active power of LEL k. Larger DPI values indicate stronger cascading behavior and higher susceptibility of the system to LEL-to-LEL propagation. To assess the grid-level severity of the disturbance, we use the *frequency violation index* (FVI), which measures the magnitude and duration of frequency deviations beyond the nominal value. The FVI is defined as $\text{FVI} = \sum_{t=t_0}^{t_0+120} |\omega(t) - \omega_{\text{norm}}|$ where $\omega(t)$ is the system frequency and $\omega_{\text{norm}} = 60\text{Hz}$ is its nominal value. A larger FVI corresponds to longer and deeper frequency violations, indicating greater stress imposed on the grid during the disturbance.

Fig. 6 compares the proposed LEL model with a baseline version in which the OU-based duty-idle dynamics are replaced by a Poisson-driven utilization pattern. Across a range of LEL penetration levels, the OU-based model produces a clearer and more monotonic trend: as LEL share increases, both DPI and FVI consistently rise, indicating greater vulnerability to secondary LEL trips and longer post-disturbance frequency violations. In contrast, the Poisson-based model often obscures this relationship because its memoryless and overly abrupt fluctuations mask the underlying coupling between workload surges and grid dynamics. This analysis highlights that higher LEL penetration can significantly amplify disturbance propagation and recovery burdens, and that accurate workload modeling is essential for anticipating these risks and for planning mitigation or protection strategies.

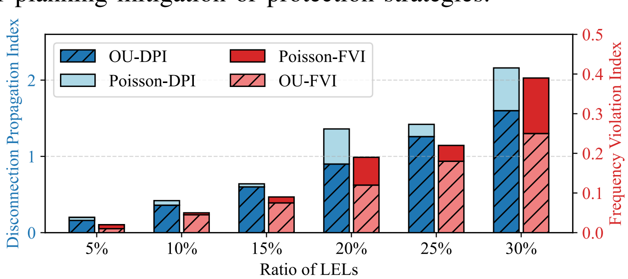

Fig. 6. Comparison of DPI and FVI under increasing LEL penetration for the proposed OU-based workload model and a Poisson-based baseline.

IV. Conclusion

This paper presented a unified dynamic modeling framework for Large Electric Loads (LELs), including data centers, cryptocurrency mining facilities, and hydrogen electrolyzers. A data-driven workload model was developed using an Ornstein-Uhlenbeck (OU) process calibrated to real CPU/GPU utilization traces, capturing the burstiness and autocorrelation of practical computational demand. Integrated into the ANDES simulation platform, the unified model reproduced key system-level behaviors such as disturbance-induced LEL trips, reconnect-retrip cycles, and prolonged recovery in larger networks. The results demonstrate that realistic workload dynamics and protection behavior are essential for accurately assessing the stability implications of growing LEL penetration and for guiding future protection and reconnection strategies.

References

[1] North American Electric Reliability Corporation (NERC). (2024) LLTF Kickoff Presentations. Accessed: Nov. 7, 2025. [Online]. Available: https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/lltf/lltf_kickoff_presentations.pdf

[2] Electric Reliability Council of Texas (ERCOT). (2024) ERCOT Large Load Loss/Reduction Events 2020-2024. Accessed: Nov, 2025. [Online]. Available: https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/lltf/lltf_presentations_december_12_2024.pdf

[3] E. Meier and S. Basumallik, "Reliability challenges and solutions for large load integration in bulk power systems," *Authorea Preprints*, 2025.

[4] X. Chen, X. Wang, A. Colacelli, M. Lee, and L. Xie, "Electricity demand and grid impacts of ai data centers: Challenges and prospects," *arXiv preprint arXiv:2509.07218*, 2025.

[5] A. Jimenez-Ruiz and F. Milano, "Data center model for transient stability analysis of power systems," *arXiv preprint arXiv:2505.16575*, 2025.

[6] J. Conto, Y. Cheng, J. Rose, and J. Schmall, "Texas loads ride toward grid stability: Voltage ride through of large power electronic loads," *IEEE Power and Energy Magazine*, vol. 23, no. 5, pp. 56–67, 2025.

[7] Y. Xie, W. Cui, and A. Wierman, "Enhancing data center low-voltage ride-through," *arXiv preprint arXiv:2510.03867*, 2025.

[8] M.-S. Ko and H. Zhu, "Wide-area power system oscillations from large-scale ai workloads," *arXiv preprint arXiv:2508.16457*, 2025.

[9] S. Biswas, A. C. Varghese, K. Chatterjee, S. Nekkalapu, B. Ross, and J. Follum, "Evaluating the risk to bulk power system reliability from large load induced oscillations," *Authorea Preprints*, 2025.

[10] G. Valverde, G. Pierrou, J. Krajacic, and G. Hug, "Forced oscillations in power systems induced by ai workloads in data centers," *Authorea Preprints*, 2025.

[11] K.-B. Kwon, S. Mukherjee, and V. Adetola, "Operational risks in grid integration of large data center loads: Characteristics, stability assessments, and sensitivity studies," *arXiv preprint arXiv:2510.05437*, 2025.

[12] P. Parikh, M. Ali, N. Virani, and S. C. Evans, "Strain to synergy: Powering ai data centers with sustainable and smart power grid infrastructures," in *2025 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 2025, pp. 1–6.

[13] C. Roberts, E. M. Stewart, and F. Milano, "Validation of the ornstein-uhlenbeck process for load modeling based on μpmu measurements," in *2016 Power Systems Computation Conference (PSCC)*. IEEE, 2016, pp. 1–7.

[14] J. L. Castillo-Arellano, M. Badaoui, and D. Ruiz-Vega, "Angular and voltage power system stability analysis considering the uncertainty of load magnitude," in *2024 IEEE PES Generation, Transmission and Distribution Latin America Conference and Industrial Exposition (GTDLA)*. IEEE, 2024, pp. 1–6.

[15] S. Samsi, M. L. Weiss, D. Bestor, B. Li, M. Jones, A. Reuther, D. Edelman, W. Arcand, C. Byun, J. Holodnack *et al.*, "The mit supercloud dataset," in *2021 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 2021, pp. 1–8.

[16] H. Cui, F. Li, and K. Tomsovic, "Hybrid symbolic-numeric framework for power system modeling and analysis," *IEEE Transactions on Power Systems*, vol. 36, no. 2, pp. 1373–1384, 2020.

[17] D. Kosterev, A. Meklin, J. Undrill, B. Lesieutre, W. Price, D. Chassin, R. Bravo, and S. Yang, "Load modeling in power system studies: Wecc progress update," in *2008 IEEE Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century*. IEEE, 2008, pp. 1–8.

[18] F. Milano, *Power system modelling and scripting*. Springer Science & Business Media, 2010.

979-8-3315-8138-1