Operational Integration of Data Centers into Wholesale Markets: A Power-Trajectory Framework

Menghan Zhang and Caisheng Wang
Department of Electrical and Computer Engineering
Wayne State University
Detroit, USA
{mhzhang, cwang}@wayne.edu

Wenyuan Tang
Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, USA
wtang8@ncsu.edu

***Abstract*—With the rapid expansion of the AI industry, large-scale electricity consumers such as data centers are emerging as a new class of large loads in modern power systems. Their concentrated capacity and abrupt connection/disconnection events produce high-magnitude power steps, which strain system reliability and complicate coordination with wholesale-market schedules. This paper proposes a power-trajectory framework that represents large loads as predefined power trajectories and coordinates generation trajectories to track them. By treating large loads as time-varying power trajectories rather than reacting to power step changes, the proposed framework enables smooth transitions between dispatch points and alleviates ramping stress on generating units. Case studies demonstrate that the proposed framework mitigates deterministic intra-hourly imbalances and improves operational consistency between existing market schedules and the physical behavior of large loads.**

***Index Terms*--Large Loads, Data Centers, Power Trajectories, Operational Coordination, Wholesale Markets.**

I. Introduction

The rapid proliferation of AI applications has driven the emergence of large-scale electricity consumers—notably data centers—which now represent a fast-growing class of demand in modern power systems. Recent NERC assessments indicate that their increasing scale and the timing of their connection/disconnection pose new challenges for system planning, market participation, and real-time operations [1]. Unlike conventional demand, these large loads exhibit concentrated capacity and abrupt, high-magnitude power step changes during connection/disconnection events, imposing substantial operational stress on the grid and wholesale markets.

Large loads, typically ranging from tens of megawatts to over a gigawatt [2], can induce high-magnitude power step changes that must be accommodated by system operations. While system reserves and other ancillary services are generally adequate for managing such variations [3], they are primarily designed to address uncertainty-driven net-load variations (e.g., renewable forecast errors) and unpredictable events. In contrast, the large-load connection/disconnection is largely deterministic and known ex ante. Treating such predictable actions solely through post-event measures may lead to avoidable ramping stress and deterministic intra-hourly imbalances.

Traditionally, wholesale energy markets have been designed around generating units, with limited mechanisms to regulate the physical behavior of large loads. Existing interconnection standards, such as FERC Order No. 2003 [4], clearly specify procedures for interconnecting large generating units to the grid. However, no equivalent regulatory framework exists for demand-side resources of comparable scale. Although large loads can, in principle, be modeled as price-responsive or dispatchable resources, their physical characteristics and operational constraints differ fundamentally from those of generating units. The rapid emergence of AI-driven facilities further amplifies the difficulty of coordinating their physical behavior with wholesale-market schedules.

While existing demand-response mechanisms provide tools to guide large-load behavior, they are not designed to manage the operational stress associated with connection/disconnection events. From a system-operation perspective, the main challenge lies in coordinating load and generation trajectories. This coordination facilitates smooth connection/disconnection, while preserving energy balance and mitigating ramping stress on generating units.

To address these challenges, this paper proposes a power-trajectory framework that represents large loads as predefined power trajectories and serves as time-varying dispatch targets. This trajectory-based approach is conceptually analogous to the controlled connection/disconnection of large-scale generators, which transition through trajectories [5]. Crucially, the proposed framework is positioned as an operational coordination layer—rather than a redesign of market-clearing or settlement mechanisms—aimed at improving the operational integration of large loads into wholesale markets.

The main contributions are as follows. 1) We introduce a power-trajectory framework for integrating large loads (e.g., data centers) into wholesale markets by representing them as predefined power trajectories. 2) We develop a continuous-time, trajectory-based dispatch model that tracks these predefined power trajectories while ensuring smooth and physically consistent generation trajectories. 3) Case studies demonstrate that the proposed framework mitigates deterministic intra-hourly imbalances and highlights key operational differences between intra-hourly and multi-hourly coordination.

979-8-3315-8138-1

II. A Power-Trajectory Framework

This section develops a power-trajectory framework that integrates large loads by operationally coordinating predefined load trajectories with continuous-time generation trajectories. It also evaluates the consistency between market schedules and physical execution across different coordination timescales, including intra-hourly and multi-hourly perspectives.

A. Problem Description

In U.S. wholesale electricity markets, deterministic intra-hourly imbalances can arise even when demand forecasts are perfect [6]. These imbalances originate from a structural mismatch between market schedules and the physical ramping behavior of generating units. Generally, the day-ahead market (DAM) produces hourly schedules, whereas the real-time market (RTM) operates every five minutes within a rolling 60-minute horizon [7], [8]. When large loads connect or disconnect according to DAM schedules, they are represented as instantaneous step changes in power.

Fig. 1 illustrates this structural mismatch using a specific connection event. The baseline system demand $D(t)$ is constant over the interval of interest (bottom-right panel), and a large load $L(t)$ is scheduled to connect at 17:00 in the DAM schedule (top-right panel, black line). Historically, loads have been assumed to have limited operational impact. They are typically not subject to generator-like ramping/trajectory requirements, and are therefore modeled as step changes in market schedules.

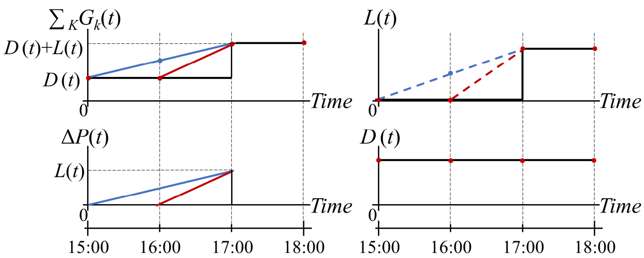

Fig. 1. Deterministic intra-hourly imbalance $\Delta P(t)$ caused by the mismatch between stepwise load schedules and the ramping behavior of generating units.

In contrast, generating units are subject to well-defined interconnection and ramping constraints; therefore, they cannot instantaneously follow the scheduled step change. As shown in the top-left panel of Fig. 1, the aggregate generation trajectory $\Sigma_K G_k(t)$ must ramp gradually (red or blue solid lines) from $D(t)$ toward $D(t)+L(t)$, beginning prior to the scheduled connection time (17:00). This unavoidable mismatch between scheduled load and physical generation produces the deterministic intra-hourly imbalance $\Delta P(t)$, shown in the bottom-left panel. $\Delta P(t)$ at time t is quantified as (1).

$$\Delta P(t) = \sum\nolimits_K G_k(t) - D(t) - L(t) \tag{1}$$

where k indexes generating units and K is the total number of generating units.

Unlike forecast errors, this imbalance is fully predictable and results exclusively from the mismatch between market schedules and the physical ramping behavior. Although such deterministic intra-hourly imbalances have long existed in power system operations, they become more pronounced with emerging large, concentrated loads such as data centers. Traditional DAM formulations enforce power balance only at discrete dispatch points and implicitly assume piecewise-constant loads and linear generation trajectories. In the absence of guidance between two dispatch points, generating units naturally follow linear ramps, thus giving rise to $\Delta P(t)$.

Increasing the temporal resolution mitigates—but cannot eliminate—deterministic intra-hourly imbalances. Moreover, denser dispatch points may increase ramping stress as generators attempt to follow step changes over shorter intervals. To address this structural mismatch, the proposed framework replaces the step-function representation of $L(t)$ with a continuous-time, predefined power trajectory $L(t)$ (top-right panel in Fig. 1, dashed curves). While both representations describe the same total load at the scheduled connection time (17:00), they differ fundamentally in their temporal evolution before 17:00. Under this framework, generation trajectories are dispatched to continuously track $D(t)+L(t)$, ensuring alignment between physical generation behavior and load evolution (top-left panel). By enforcing continuous-time coordination, the proposed trajectory-based approach can eliminate deterministic intra-hourly imbalances arising from mismatches between market schedules and the physical ramping behavior of generating units.

B. Continuous-time Dispatch Formulation

A predefined power trajectory $L(t)$ is specified to represent the connection/disconnection of a large load. This trajectory is not restricted to being linear. Additionally, the baseline system demand $D(t)$ represents the aggregate demand excluding the large load and is modeled either as a constant value or as a time-varying power profile. The analysis considers a system-wide model without nodal differentiation, as transmission constraints are beyond the scope of this study. Both $L(t)$ and $D(t)$ are treated as system-level load trajectories.

The large-load trajectory is defined as a smooth function that increases from zero at a starting time S to a target load level at time T. For notational convenience, S is assumed to coincide with a DAM dispatch point, although the trajectory may begin at any instant. The terminal time T is chosen to align with a DAM dispatch point. The objective is to determine continuous-time generation trajectories $\mathbf{G}(t)$ throughout the period $\mathcal{T}=[S, T]$ that maintain power balance while respecting the operational limitations of generating units. Formally, the continuous-time dispatch problem is formulated as follows.

$$\min_{\mathbf{G}(t)} \int_{\mathcal{T}} \left(C\left(\mathbf{G}(t)\right)\right) dt \tag{2}$$

$$\text{s.t.} \sum\nolimits_K G_k(t) = D(t) + L(t), \forall t \tag{3}$$

$$G_k(t) \in \mathcal{G}_k, \forall k, t \tag{4}$$

$$\mathcal{G}_k = \left\{ G_k(t) \,|\, G_k^{\min} \le G_k(t) \le G_k^{\max}; \; R_k^{\text{down}} \le G_k^{'}(t) \le R_k^{\text{up}} \right\} \tag{5}$$

In this model, the total production cost is given by $C(\mathbf{G}(t))=\Sigma_K C_k(G_k(t))$, where $C_k(\,)$ denotes the cost function of generating unit k. For each generating unit k, $G_k^{\min}/G_k^{\max}$ are the minimum/maximum generation outputs. $R_k^{\text{down}}/R_k^{\text{up}}$ denote the downward/

979-8-3315-8138-1

upward ramping capabilities (in MW per time unit). $G_k'(t)$ is the time derivative of $G_k(t)$, representing the ramping rate.

The objective function (2) minimizes the total production cost over $\mathcal{T}$ by determining the continuous-time generation trajectories $\mathbf{G}(t)=(G_1(t),\ldots,G_K(t))^\top$. Constraint (3) enforces the continuous-time power balance. Constraints (4)-(5) impose operational limitations $\mathcal{G}_k$ on each generating unit k, including generation capacity constraints and ramping constraints.

The DAM schedule is used to set the boundary conditions $\mathbf{G}(S)$ and $\mathbf{G}(T)$, ensuring consistency between generation trajectories and DAM schedules at the endpoints of $\mathcal{T}$. To solve (2)-(5), we adopt the iterative dispatch methodology developed in [9], [10], which enforces continuous-time power balance without requiring a priori spline-based approximations of trajectories during the solution process. As summarized in Fig. 2, this methodology constructs generation trajectories, verifies feasibility under continuous-time constraints, and iteratively refines them until feasible solutions $\mathbf{G}(t)$ are obtained.

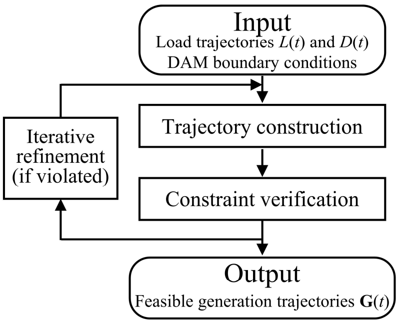

Fig. 2. Overview of the iterative dispatch methodology [9], [10].

Here, feasibility refers to the existence of continuous-time generation trajectories $\mathbf{G}(t)$ that satisfy all unit-specific operational limitations while maintaining continuous-time power balance with the total load trajectory $L(t)+D(t)$. Feasible generation trajectories $\mathbf{G}(t)$ do not always exist if $L(t)$ and $D(t)$ exceed available generation capacity or ramping capability. In such cases, feasibility can be restored by specifying a smoother trajectory for $L(t)$, for example by extending the ramp duration (i.e., enlarging $\mathcal{T}$), thereby alleviating generation capacity and binding ramping limitations.

C. Consistency with Market Schedules

The previous section focused on the physical coordination between a predefined power trajectory and generation trajectories. However, different predefined power trajectories do not guarantee consistency with existing DAM schedules. In particular, predefined power trajectories under different coordination timescales can affect consistency with DAM schedules in fundamentally different ways. This section examines the conditions under which the proposed framework preserves consistency with DAM schedules.

In existing U.S. wholesale electricity markets, operational processes such as reliability unit commitment (RUC) and look-ahead dispatch (LAD) already act as intermediate operational coordination layers between the DAM and RTM [7], [8]. Although these processes do not constitute market clearing, they adjust operating instructions and may introduce out-of-market corrections [11]. The proposed framework, referred to as additional dispatch in Fig. 3, plays a similar operational role but is specifically structured to integrate large loads.

As illustrated in Section II.A, the DAM represents large loads as step changes at hourly dispatch points, whereas generating units must physically ramp continuously. This mismatch creates deterministic intra-hourly imbalances that propagate into the RTM. The proposed framework addresses this issue by assigning physically feasible trajectories to both large loads and generators between DAM dispatch points. Fig. 3 summarizes the operation of the proposed framework. After the DAM clears, a predefined power trajectory $L(t)$ —either intra-hourly or multi-hourly—is established. It serves as a dispatch target for the additional dispatch, which computes the corresponding generation trajectories. This framework specifies the physical behavior between DAM hourly dispatch points, which is not defined in existing market procedures.

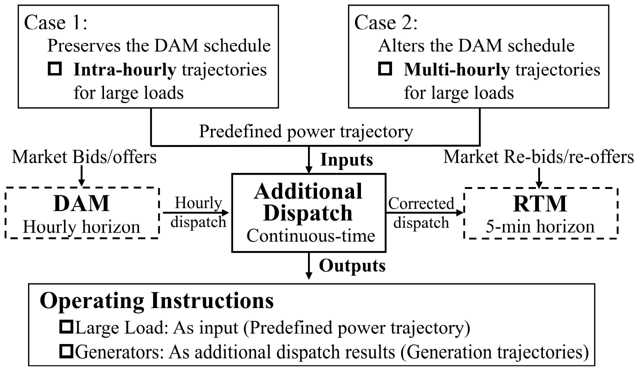

Fig. 3. Operational placement of the proposed framework between the DAM and RTM under different coordination timescales.

The consistency with market schedules depends on whether the predefined power trajectory alters DAM boundary conditions. Two cases can be distinguished:

1) Intra-hourly trajectories (e.g., red dashed line in Fig. 1): These trajectories remain within a single DAM hour and preserve DAM boundary conditions. Since they do not modify hourly DAM schedules, they maintain consistency with DAM schedules and act purely as an operational refinement.

2) Multi-hourly trajectories (e.g., blue dashed line in Fig. 1): These trajectories extend across multiple DAM hours to produce smoother ramps. Although they alleviate ramping stress, they modify DAM boundary conditions and require adjustments to DAM schedules that must be propagated to the RTM as out-of-market corrections.

In summary, intra-hourly trajectories maintain consistency with DAM schedules and function as operational refinements, whereas multi-hourly extensions adjust DAM schedules and introduce out-of-market corrections. Because such corrections affect settlements and reduce transparency [12], intra-hourly trajectories are preferred, and multi-hourly extensions are used only when required to restore feasibility.

D. Trajectory Declaration and Operational Coordination

Predefined power trajectories $L(t)$ are envisioned to be declared ex ante by large-load participants based on their internal operational constraints. These trajectories are subject to validation and approval by the system operator to ensure consistency with grid reliability requirements. Alternatively, predefined power trajectories may be specified directly by the system operator through operator-led coordination mechanisms when reliability issues require intervention. A key distinction is that trajectories are defined over continuous time, requiring explicit consideration of intra-hourly operating behavior.

At present, wholesale markets do not explicitly specify the intra-hourly evolution of load following the DAM, leaving flexibility in how physical behavior is executed between dispatch points. The proposed framework leverages this flexibility by introducing explicit large-load trajectories that serve as operational dispatch targets rather than new market-clearing variables. Consistent with current U.S. market practice, there is no market clearing process between the DAM and RTM; deviations from scheduled quantities are settled in the RTM. Here, deviations from the predefined power trajectory are envisioned to be managed through existing real-time imbalance settlement mechanisms, without introducing new penalty mechanisms. In this paper, we therefore focus exclusively on the operational coordination layer between the DAM and RTM.

III. Case Studies

The proposed framework is demonstrated using the iterative dispatch methodology and the illustrative 5-bus system described in [9]. Two types of large-load trajectories—intra-hourly and multi-hourly—are examined to illustrate the operational behaviors analyzed in Section II.C. For illustration, the predefined power trajectory $L(t)$ is modeled as a linear ramp over the period $\mathcal{T}=[S, T]$, while the baseline system demand $D(t)$ is modeled as a nonlinear trajectory. Hourly DAM dispatch points are indicated by red points in the figures. The disconnection process follows the same principles and is therefore omitted.

A. Results for Intra-Hourly Predefined Power Trajectories

Under the conventional DAM representation, the large load is scheduled as 0 MW at $S=-60$ min (16:00) and 50 MW at $T=0$ min (17:00), without any specification of its evolution between these two dispatch points. Consequently, the DAM implicitly represents the large load $L(t)$ as remaining at 0 MW over $[S, T)$ and jumping to 50 MW only at $t=T$, as shown by the black trajectory in Fig. 4(c). Accordingly, generating units are dispatched to supply only the baseline system demand $D(t)$ over $[S, T)$, and are scheduled to supply $D(t) + L(t)$ at $t=T$, as shown by the black trajectory in Fig. 4(a).

From a physical perspective, generating units are limited by ramping constraints and cannot instantaneously follow this DAM-scheduled step change. To reach the required output at t=0, feasible generation trajectories must involve ramping prior to the scheduled connection time. Because the DAM specifies only the boundary conditions at $t=-60$ and $t=0$ and provides no intra-hourly operating instructions, generating units typically follow a feasible physical ramping path, shown by the black trajectory in Fig. 5. This reveals a fundamental temporal mismatch between the DAM's step-change representation and the continuous-time physical capabilities of generating units. As analyzed in Section II.A, this mismatch produces a deterministic intra-hourly imbalance.

Under the proposed framework, the step change is replaced by a predefined power trajectory $L(t)$, shown by the red dashed curve in Fig. 4(c), which increases smoothly from 0 MW to 50 MW. Consequently, the total load $D(t) + L(t)$, shown by the red dashed curve in Fig. 4(a), differs significantly from the step change in the DAM schedule. Using the continuous-time dispatch model (2)-(5), generating units are dispatched to track $D(t) + L(t)$. The resulting generation trajectories, red curves in Fig. 5, maintain continuous-time power balance throughout $[S, T]$. Because the predefined power trajectory is intra-hourly here, the DAM boundary conditions at $t=-60$ and $t=0$ remain unchanged, and no adjustment to the DAM schedule is required.

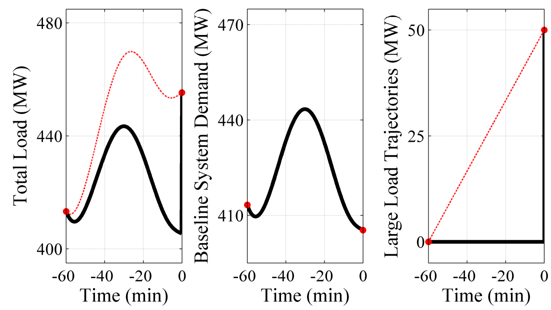

Fig. 4. Continuous-time load profiles under the intra-hourly trajectory.

When the predefined power trajectory $L(t)$ and generation trajectories $\mathbf{G}(t)$, shown as red curves in Fig. 5, are issued as operating instructions, the deterministic intra-hourly imbalance is avoided, provided participants follow their assigned trajectories. Any remaining mismatch arises from stochastic deviations of $D(t)$ and $L(t)$, which are naturally handled by the RTM. From the RTM perspective, the proposed framework removes the deterministic component of intra-hourly imbalance caused by the DAM's coarse temporal resolution. This framework effectively decouples systematic schedule-induced mismatches from stochastic forecast errors.

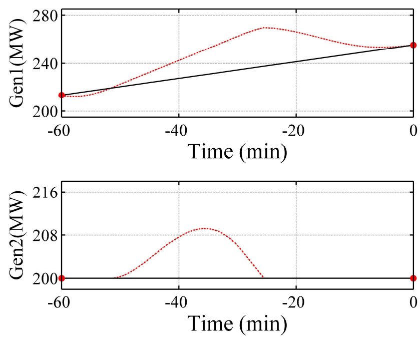

Fig. 5. Continuous-time generation trajectories corresponding to the intra-hourly trajectory.

B. Results for Multi-Hourly Predefined Power Trajectories

The previous case shows that an intra-hourly predefined power trajectory is feasible when sufficient ramping resources exist. We now examine a multi-hourly predefined power trajectory to illustrate the effects of extending the ramp duration and its implications for consistency with market schedules.

979-8-3315-8138-1

In this case, the predefined power trajectory $L(t)$ spans 120 minutes ($\mathcal{T}=[S, T] = [-120, 0]$). The DAM specifies the large load as 0 MW at 15:00 ($S=-120$) and 50 MW at 17:00 ($T=0$). The intermediate DAM dispatch at 16:00 (t=−60) is not preserved, allowing the multi-hourly trajectory to adjust the DAM schedule.

Fig. 6 and Fig. 7 compare intra-hourly and multi-hourly trajectories over the final hour for load and generation, respectively. $D(t)$ in Fig. 4(b) and Fig. 6(b) remains unchanged across the comparison. Under the multi-hourly trajectory (blue curve in Fig. 6), $L(t)$ reaches 25 MW at $t = -60$, whereas the intra-hourly trajectory (red curve in Fig. 6) remains at 0 MW. This smoother $L(t)$ reduces the ramping stress over [−60, 0] but requires an additional 25 MW of generation at $t = -60$, thereby adjusting the DAM schedule and introducing an out-of-market correction for generating units (blue points in Fig. 7).

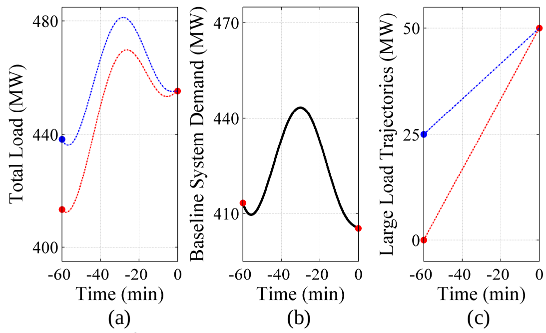

Fig. 6. Comparison of intra-hourly and multi-hourly large-load trajectories and resulting total load profiles over the final hour.

Fig. 7 shows the corresponding generation trajectories. Under the multi-hourly trajectory, Gen 1 must begin at a higher output (25 MW) at $t = -60$, resulting in smoother subsequent ramping trajectories (blue curve in Fig. 7). By distributing the total ramp over a longer duration, the multi-hourly coordination alleviates ramping stress and allows lower-cost generation (Gen1) to contribute more energy in this illustrative case.

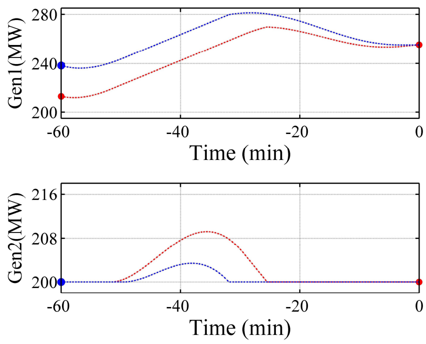

Fig. 7. Comparison of continuous-time generation trajectories under intra-hourly and multi-hourly coordination over the final hour.

In summary, multi-hourly coordination yields physically smoother transitions and enhances operational flexibility by leveraging a longer ramp duration. However, these benefits arise at the cost of adjusting DAM schedules, thereby introducing out-of-market corrections. By contrast, intra-hourly coordination preserves DAM schedules and functions as an operational refinement. Consistent with market practice, intra-hourly trajectories should therefore be prioritized, while multi-hourly extensions serve as a feasibility-restoring mechanism when ramping resources are insufficient.

IV. Conclusion and Future Work

This paper proposes a power-trajectory framework to integrate large loads, such as data centers, into wholesale electricity markets. In the proposed framework, large loads are represented as predefined power trajectories, and a continuous-time dispatch model provides generation trajectories that track them. Case studies demonstrate that the proposed framework mitigates deterministic intra-hourly imbalances, reduces ramping stress during large-load connection/disconnection, and improves coordination between market schedules and the physical behavior of large loads.

By specifying the trajectory-based behavior between discrete dispatch points—currently undefined in market procedures—the framework acts as an operational coordination layer that bridges the gap between market schedules and physical execution. This framework is complementary to existing day-ahead market scheduling processes, focusing on operational feasibility without necessitating a redesign of market-clearing formulations. Future work will examine regulatory and economic implications, including settlement and incentive design, as well as extensions to network-constrained systems and stochastic load uncertainty.

References

[1] NERC, "2024 Long-Term Reliability Assessment," Atlanta, GA, USA, 2024. [Online]. Available: https://www.nerc.com/globalassets/programs/rapa/ra/nerc_long-term-reliability-assessment_2024.pdf

[2] Grid Strategies LLC, "Strategic Industries Surging: Driving U.S. Power Demand," Washington, DC, USA, 2024. [Online]. Available: https://gridstrategiesllc.com/wp-content/uploads/National-Load-Growth-Report-2024.pdf

[3] Y. V. Makarov, C. Loutan, J. Ma, and P. de Mello, "Operational Impacts of Wind Generation on California Power Systems," *IEEE Trans. Power Syst.*, vol. 24, no. 2, pp. 1039–1050, 2009.

[4] FERC, "Standardization of Generator Interconnection Agreements and Procedures," Order No. 2003, Washington, DC, USA, 2003. [Online]. Available: https://www.ferc.gov/sites/default/files/2020-04/E-1_71.pdf

[5] G. Morales-España, J. M. Latorre, and A. Ramos, "Tight and Compact MILP Formulation of Start-Up and Shut-Down Ramping in Unit Commitment," *IEEE Trans. Power Syst.*, vol. 28, no. 2, pp. 1288–1296, 2013.

[6] M. Parvania and A. Scaglione, "Unit commitment with continuous-time generation and ramping trajectory models," *IEEE Trans. Power Syst.*, vol. 31, no. 4, pp. 3169–3178, 2016.

[7] MISO, "BPM 002: Energy and Operating Reserve Markets," Carmel, IN, USA, 2023. [Online]. Available: https://www.misoenergy.org/legal/business-practice-manuals/

[8] E. Litvinov, F. Zhao, and T. Zheng, “Electricity markets in the United States: Power industry restructuring processes for the present and future,” *IEEE Power Energy Mag.*, vol. 17, no. 1, pp. 32–42, Jan.–Feb. 2019.

[9] M. Zhang, Y. Zhou, M. Xiang, J. Yu, W. Li, and Z. Yang, “A methodology for building generation trajectories to balance continuous-time load profiles,” *IEEE Trans. Power Syst.*, vol. 40, no. 2, pp. 1609–1622, 2025.

[10] M. Zhang and C. Wang, “Dual analysis of continuous-time economic dispatch and its price implications,” arXiv:2601.19095, 2026. [Online]. Available: https://doi.org/10.48550/arXiv.2601.19095

[11] Y. M. Al-Abdullah, M. Abdi-Khorsand, and K. W. Hedman, “The role of out-of-market corrections in day-ahead scheduling,” *IEEE Trans. Power Syst.*, vol. 30, no. 4, pp. 1937–1946, 2015.

[12] M. Zhang, Z. Yang, W. Lin, J. Yu, and W. Li, “Internalization of reliability unit commitment in day-ahead market: Analysis and interpretation,” *Appl. Energy*, vol. 326, Art. no. 119926, 2022.

979-8-3315-8138-1