A Data-Driven Semi-Analytic Approach for Optimal Control of Data Centers as Grid-Interactive Loads

1st Elinor Ginzburg-Ganz
Electrical & Computer Engineering
Technion
Haifa, Israel
elinor.g@campus.technion.ac.il

2nd Yoash Levron
Electrical & Computer Engineering
Technion
Haifa, Israel
yoashl@ee.technion.ac.il

***Abstract*—We propose an online optimal control method for data centers, treating them as flexible loads within a power grid's demand response mechanism. Our goal is to balance two contradictory objectives: the participation of this controllable load in grid frequency regulation, and the operation of the data center according to its planned workload. Within this framework, we address the computational complexity stemming from inherent uncertainties in both power system loads and data center scheduled workloads. These uncertainties lead to a complex stochastic optimization problem, which poses significant computational challenges. To address these challenges, we develop a semi-analytic solution that can be efficiently integrated with a machine learning estimator to optimally schedule the data center's power demand. This approach significantly reduces the control complexity of the scheduling algorithm, thereby facilitating real-time, online implementation and seamless integration of the proposed policy. The method's efficiency is demonstrated through a real-world case study.**

***Index Terms*—Data centers, Demand response, Machine learning, Optimal control**

I. Introduction

THE power demand from both training and inference of AI workloads introduces a significant challenge due to their immense collective energy footprint. This rapid increase in demand creates unprecedented challenges for the design and operation of power systems, as documented in several recent reports [1]–[3].

Nevertheless, AI workloads exhibit operational flexibility beyond traditional data center tasks. Both training processes, which can be paused or redirected using checkpoints, and inference stages, which tolerate latency and allow for geographic relocation, contribute to this adaptability. This inherent flexibility allows AI data centers to function as controllable loads, capable of adjusting power consumption in response to grid conditions or economic incentives.

Numerous research works explore methodologies to exploit this flexibility through demand response mechanisms. One approach utilizes a fixed optimization horizon, as demonstrated in Work [4], Work [5], Work [6], and Work [7]. Alternatively, Work [8] focuses on instantaneous optimization assuming a known workload. The prior, within the class of Model Predictive Control, requires complete knowledge of future power trajectories and often relies on arbitrary horizons. In contrast, the latter forces a complex hierarchical controller design that cannot account for workload uncertainties.

Considering these challenges, we propose an optimal control approach for operating data centers as flexible loads within a demand response mechanism to regulate grid frequency, as illustrated in Fig. 1. The main result provides a semi-analytic solution to this problem, which a machine learning estimator utilizes for optimal power scheduling. In contrast to Model Predictive Control, this method relies on a single estimated value at each time point instead of the full trajectory, which lowers computational demand. Moreover, the solution structure reveals the effective time horizon. Compared to other approaches, our method enables a simpler, one-level controller that facilitates real-time integration. The proposed method is evaluated using real-world data.

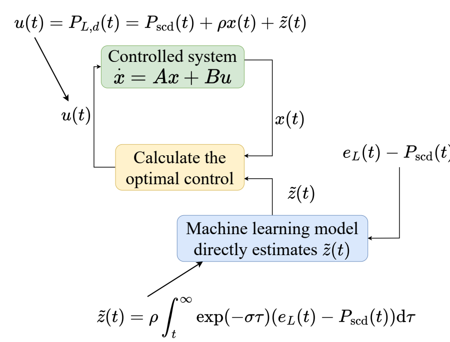

Fig. 1: Illustration of the proposed method. Instead of predicting the future power consumption trajectory, the machine learning estimator predicts $\tilde{z}(t)$ directly, which is then used to calculate the optimal control law, using the proposed closed-form solution.

II. Main Result

We analyze a system featuring a data center with controllable power consumption $P_{L,d}(t) \in \mathbb{R}_{\geq 0}$. This data center is connected to a local distribution network that also serves K other loads with consumption $P_{L,i}(t) \in \mathbb{R}_{\geq 0}$. The total demand is met by G synchronous generators, indexed $j = 1, \dots, G$, providing power $P_i(t) \in \mathbb{R}_{\geq 0}$. Generator buses are labeled $1, \dots, G$, and load buses $(G+1), \dots, (G+1)+(K+1)$.

979-8-3315-8138-1

Our objective is to ensure that for every given time instant $t \in T$, where T is the time horizon, the electrical frequency of each generator $\omega_j(t)$ is close to the nominal synchronous frequency ω_s for all $j = 1, \ldots, G$, or equivalently, the signals $\omega_j(t)$ should track ω_s.

We use the convention that all the generators' phasor angles $\delta_2(t), \delta_3(t), \ldots, \delta_G(t)$ are measured with respect to the first generator, so that $\delta_1(t) = 0$, and define the following vectors:

- $\Delta\omega(t) = [\omega_1(t) - \omega_s, \ldots, \omega_G(t) - \omega_s]^{\mathsf{T}}$,
- $\delta(t) = [\delta_2(t), \ldots, \delta_G(t)]^{\mathsf{T}}$,
- $P(t) = [P_1(t), \ldots, P_G(t)]^{\mathsf{T}}$,
- $P_L(t) = [P_{L,1}(t), \ldots, P_{L,d}, \ldots, P_{L,K+1}(t)]^{\mathsf{T}}$,

Here, $\omega_i(t)$ is the electric frequency of generator i in rad/s, and $\delta_i(t)$ is the relative rotor angle, as defined above. The system is described by the following differential and algebraic equations:

$$\begin{aligned} M\Delta\dot{\omega}(t) + D\omega(t) &= P_{\text{ref}}(t) - P(t), \\ \dot{\delta}(t) &= U\Delta\omega(t), \\ P(t) &= f(\delta(t), P_L(t)). \end{aligned} \tag{1}$$

The diagonal and positive definite matrix M represents the generators' inertia constants, and the diagonal and positive definite matrix D represents the generators' damping coefficients. The vector $P_{\text{ref}}(t)$ contains the reference powers for the generators. The matrix U describes the relations $\dot{\delta}_i(t) = \omega_i(t) - \omega_1(t)$ for $i = 2, \ldots, G$, and is given by $U = \begin{bmatrix} -1_{(G-1)\times 1} & I_{(G-1)\times(G-1)} \end{bmatrix}$, where $1_{(G-1)\times 1}$ is a vector of ones, and $I_{(G-1)\times(G-1)}$ is the identity matrix. The function $f(\cdot)$ depends on the distribution network and the loads. This function is usually not available in closed form, and has to be evaluated numerically, by solving the network's AC power-flow equations. In this work we assume nothing about the distribution network topology, and so $f(\cdot)$ is not given or known.

To continue, we now employ the "central angle" technique, to simplify the power system dynamics[1]. To this end, multiply the first equation in (1) by $\mathbb{1}_{1\times G} = \begin{pmatrix} 1, & \ldots & ,1 \end{pmatrix}^{\mathsf{T}}$ from the left, to obtain

$$\bar{M}\Delta\dot{\omega}(t) + \bar{D}\omega(t) = \sum_{i=1}^{G} P_{\text{ref,i}}(t) - \sum_{i=1}^{G} P_i(t), \tag{2}$$

where $\bar{M} = \begin{pmatrix} M_{11}, M_{22}, \ldots M_{GG} \end{pmatrix}$, and $\bar{D} = \begin{pmatrix} D_{11}, D_{22}, \ldots, D_{GG} \end{pmatrix}$.

Assuming that the distribution network is lossless, and that $\bar{D} = \alpha\bar{M}$, where $\alpha > 0$ is a scalar, this last equation may be written as

$$\dot{x}(t) = -\alpha x(t) + e_L(t) - P_{L,d}(t), \tag{3}$$

[1]The derivation assumes a lossless power distribution network and a uniform ratio between the damping and inertia coefficients for all generators in the system.

where the state variable $x(t) = \bar{M}\Delta\omega$ and the variable $e_L(t)$ are explicitly given by

$$\begin{aligned} x(t) &= M_{11} \cdot (\omega_1(t) - \omega_s) + \ldots + M_{GG} \cdot (\omega_G(t) - \omega_s), \\ e_L(t) &= \sum_{i=1}^{G} P_{\text{ref,i}}(t) - \sum_{i=1, i\neq d}^{K+1} P_{L,i}(t). \end{aligned} \tag{4}$$

Based on these definitions, our objective is to optimally control $P_{L,d}(t)$, the power consumption of the data center, in order to minimize the cost

$$J = \frac{1}{2T}\int_0^T \beta x^2(t) + (P_{L,d}(t) - P_{\text{scd}}(t))^2 \mathrm{d}t, \tag{5}$$

considering the dynamic system (3), where $\beta > 0$ is a constant scalar, and $P_{\text{scd}}(t)$ is a stochastic signal representing the scheduled power consumption of the data center. The cost J encapsulates a trade-off between two critical objectives. One goal is to operate with frequencies $\omega_i(t)$ that remain as close as possible to the nominal grid frequency ω_s, which is essential for maintaining grid stability. Concurrently, the data center's real-time power consumption, $P_{L,d}(t)$, must closely track its scheduled power demand $P_{\text{scd}}(t)$. We treat this consumption as a dynamic variable to exploit its inherent flexibility, without assuming any prior knowledge about it, and deriving it in purely data-driven approach. By intentionally modulating $P_{L,d}(t)$ around its planned schedule, the facility functions as a grid-interactive load that provides frequency regulation services.

This problem, despite featuring a quadratic cost and a linear dynamic system, presents significant challenges, due to the stochastic nature of the signals $e_L(t)$ and $P_{\text{scd}}(t)$ within the linear dynamic system (3), placing it among tracking problems under uncertainty that typically lack closed-form solutions [9]. However, for the specific structure defined by (3) and (5), we derive an analytic solution expressed as a weighted integral of these unknown signals, a solution that significantly reduces the computational complexity compared to general approaches, enabling the development of a low-overhead, real-time scheduling mechanism. Theorem 1 formally presents this solution.

Theorem 1. *Assume that $\alpha > 0$, $\beta > 0$, and the signals $e_L(t), P_{\text{scd}}(t)$ are bounded and piecewise continuous. Then the optimal control law for (3) and (5), at the limit $T \to \infty$, is given explicitly by*

$$\begin{aligned} \lim_{T\to\infty} P^*_{L,d}(t) &= P_{\text{scd}}(t) + \rho x^*(t) + \tilde{z}(t), \\ \tilde{z}(t) &= \rho\int_0^\infty \exp(-\sigma\tau)(e_L(t+\tau) - P_{\text{scd}}(t+\tau))\mathrm{d}\tau, \end{aligned} \tag{6}$$

where $\sigma = \sqrt{\alpha^2 + \beta}$ and $\rho = -\alpha + \sqrt{\alpha^2 + \beta}$.

979-8-3315-8138-1

Proof. Applying Pontryagin's Minimum Principle to (3) and (5) yields necessary conditions for optimality:

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}x^* &= -\alpha x^*(t) + u^*(t), \\ \frac{\mathrm{d}}{\mathrm{d}t}\lambda^* &= \alpha\lambda^*(t) - \beta x^*(t), \\ u^*(t) &= -\lambda^*(t) + e_L(t) - P_{\mathrm{scd}}(t). \end{aligned} \tag{7}$$

with $x^*(0) = x_0$ and $\lambda^*(T) = 0$. This set of equations can be expressed in matrix form,

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} x^* \\ \lambda^* \end{pmatrix} = A_H \begin{pmatrix} x^* \\ \lambda^* \end{pmatrix} + w(t), \tag{8}$$

with

$$A_H = \begin{pmatrix} -\alpha & -1 \\ -\beta & \alpha \end{pmatrix}, \quad w(t) = \begin{pmatrix} e_L(t) - P_{\mathrm{scd}}(t) \\ 0 \end{pmatrix}. \tag{9}$$

The Hamiltonian matrix A_H has two distinct eigenvalues σ_1, σ_2, where $\sigma_1 = \sigma$ and $\sigma_2 = -\sigma$. Thus, A_H is diagonalizable, and may be written as

$$A_H = V\begin{pmatrix} -\sigma & 0 \\ 0 & \sigma \end{pmatrix} V^{-1}, \tag{10}$$

where

$$V = \begin{pmatrix} \frac{\alpha+\sqrt{\alpha^2+\beta}}{\beta} & \frac{\alpha-\sqrt{\alpha^2+\beta}}{\beta} \\ 1 & 1 \end{pmatrix}. \tag{11}$$

The set of equations in (8) can be solved explicitly to obtain

$$\begin{pmatrix} x^*(t) \\ \lambda^*(t) \end{pmatrix} = \exp(A_H t)\left[\begin{pmatrix} x_0 \\ \lambda_0^* \end{pmatrix} + \int_0^t \exp(-A_H\tau)\, w(\tau)\mathrm{d}\tau\right], \tag{12}$$

where $\lambda_0^* = \lambda^*(0)$ is unknown, and x_0 is given.

To find λ_0^* explicitly, we rely on the final condition stated by the PMP formulation, $\lambda^*(T) = 0$, which leads to

$$\begin{pmatrix} 0 & 1 \end{pmatrix} \exp(A_H T)\begin{pmatrix} x_0 \\ \lambda_0^* \end{pmatrix} + r(T) = 0, \tag{13}$$

where

$$r(T) = \int_0^T \begin{pmatrix} 0 & 1 \end{pmatrix} \exp(A_H(T-\tau))\, w(\tau)\mathrm{d}\tau. \tag{14}$$

This may be reformulated and further extended into the form of

$$q(T) + \begin{pmatrix} 0 & 1 \end{pmatrix}\exp(A_H T)\begin{pmatrix} 0 \\ 1 \end{pmatrix}\lambda_0^* + r(T) = 0, \tag{15}$$

where

$$q(T) = \begin{pmatrix} 0 & 1 \end{pmatrix}\exp(A_H T)\begin{pmatrix} 1 \\ 0 \end{pmatrix} x_0. \tag{16}$$

Reformulating the expression further result in

$$\lambda_0^* = \frac{-q(T) - r(T)}{\begin{pmatrix} 0 & 1 \end{pmatrix}\exp(A_H T)\begin{pmatrix} 0 \\ 1 \end{pmatrix}}. \tag{17}$$

Using the diagonal form of A_H, and the definition of exponential matrices, Equation (17) may be expressed as

$$\begin{aligned} \lambda_0^* &= \frac{\beta\left[x_0\sinh(\sigma T) + \mathcal{I}\right]}{\sigma\cosh(\sigma T) + \alpha\sinh(\sigma T)}, \\ \mathcal{I} &= \int_0^T \sinh(\sigma(T-\tau))(e_L(\tau) - P_{\mathrm{scd}}(\tau))\mathrm{d}\tau. \end{aligned} \tag{18}$$

Let us examing the behavior of λ_0^* when T tends to infinity. Using the definitions of $\sinh(\cdot)$ and $\cosh(\cdot)$, and performing several mathematical manipulations, taking T to infinity yields

$$\lambda_0^* = \frac{\beta\left[x_0(1-e^{-2\sigma T}) + \int_0^T(e^{-\sigma\tau} - e^{-\sigma(2T-\tau)})u(\tau)\mathrm{d}\tau\right]}{\sigma(1+e^{-2\sigma T}) + \alpha(1-e^{-2\sigma T})}. \tag{19}$$

Now we take the limit as $T \to \infty$. Since $\sigma > 0$, all the exponential terms with T are zeroed, resulting in

$$\lim_{T\to\infty}\lambda_0^* = \frac{\beta}{\sigma+\alpha}\left[x_0 + \int_0^\infty e^{-\sigma\tau}(e_L(\tau) - P_{\mathrm{scd}}(\tau))\mathrm{d}\tau\right], \tag{20}$$

where the constant $\beta/(\sigma+\alpha)$ is simply ρ. Now, consider the optimal control law at the initial time $t = 0$. Using (7), we have

$$u^*(0) = -\lambda_0^* + e_L(0) - P_{\mathrm{scd}}(0). \tag{21}$$

Substitution λ_0^* from (19) and taking T to infinity yields

$$\lim_{T\to\infty} u^*(0) = -\rho x_0 + \tilde{z}(0) + e_L(0) - P_{\mathrm{scd}}(0), \tag{22}$$

Since the system is time-invariant, the optimal control law for a general time t has the same form, and is given by

$$\begin{aligned} \lim_{T\to\infty} u^*(t) &= P_{\mathrm{scd}}(t) + \rho x^*(t) + \tilde{z}(t), \\ \tilde{z}(t) &= \rho\int_t^\infty \exp(-\sigma\tau)(e_L(\tau) - P_{\mathrm{scd}}(\tau))\mathrm{d}\tau, \end{aligned} \tag{23}$$

Switching the integration variable in (23) yields the optimal solution (6). □

The exponential term in (23) allows to clearly define the length of the integration window, which is proportional to $1/\sigma$, **revealing the effective time-horizon at a point t.**

III. Case Study

We validate the proposed framework through a case study using real-world operational data from a sampling on the rack level, of a 1 GW data center. This data, captured during a large AI model training, reflects the characteristic volatility of such workloads. The simulation study is conducted using MATLAB R2024a and Simulink.

The raw power data, which is sampled at an irregular interval, is first pre-processed. We establish a uniform time grid with a fixed time step of $dt = 2.5$ seconds over a total simulation duration of 1177.5 seconds (472 time points). The original, irregularly spaced power data is mapped onto this uniform grid using linear interpolation to ensure robust data representation without overshoot.

The simulation implements the semi-analytic optimal control solution derived in Theorem 1. We employ a data-driven

979-8-3315-8138-1

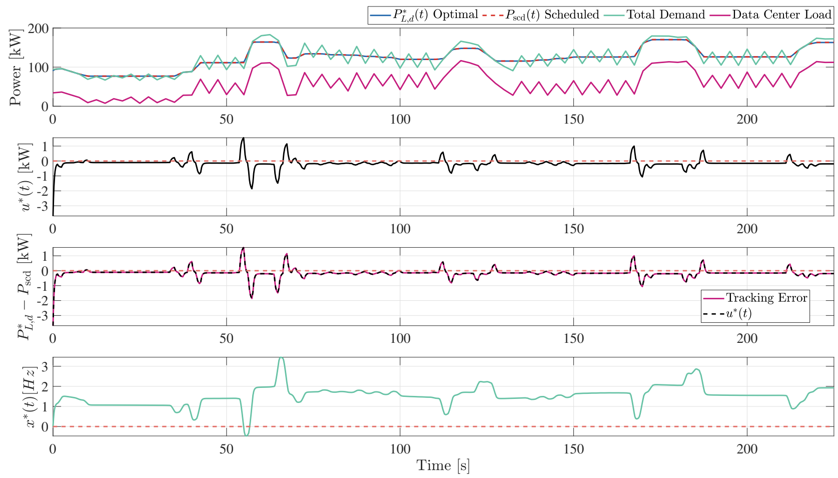

Fig. 2: Time-series performance of the optimal controller with $\beta = 10.0$. From top to bottom: optimal vs. scheduled power, control deviation $u^*(t)$, tracking error, and grid frequency state $x^*(t)$.

approach to implement the control law from Theorem 1, when the signal $e_L(t)$ is *unknown*. Rather than predicting the whole future trajectory of this signal, we predict its weighted integral, $\tilde{z}(t)$.

Examining the optimal control law from Theorem 1, it is clear that direct application of this law is infeasible when $e_L(t)$ is unknown, as $\tilde{z}(t)$ depends on all *future* signal values. While conventional methods like Model Predictive Control are computationally intensive due to the need for future trajectory predictions, our approach circumvents this challenge.

We utilize standard linear regression estimator to directly predict the integral term, $\tilde{z}(t)$, bypassing the high computational overhead. The training data, past and present information of $e_L(t)$ (by which $\tilde{z}(t)$ is calculated), is augmented based on typical consumption patterns. The system damping coefficient α is set to 0.1 s^{-1}, and the trade-off parameter β is set to 10.0. The simulation begins from a nominal state, with the initial frequency deviation $x(0)$ set to zero.

A. Numerical Results

The simulation is executed with the optimal control solution computed in 0.72 seconds. The results, summarized in Table I, demonstrate the controller's effectiveness in balancing the two conflicting objectives.

Workload Tracking Performance: The primary objective for the data center operator is to maintain its scheduled operations. The controller achieves this with high fidelity. The mean of the optimal power $P^*_{L,d}(t)$ is 124.05 kW, which closely tracks the scheduled mean of 124.27 kW. The Root Mean Square (RMS) tracking error, which represents the control deviation $u^*(t)$, is only 0.60 kW. This corresponds to a relative RMS error of just 0.48% of the mean scheduled power. The maximum deviation from the schedule at any point is 3.93 kW, indicating that the controller avoids aggressive, large-scale deviations from the planned workload.

Grid Frequency Regulation: The objective for the grid operator is to minimize the frequency state $x(t)$. The simulation results in an RMS frequency deviation of 1.5796, with a maximum observed deviation of 3.3669.

Cost Analysis: The choice of $\beta = 10.0$ is intended to prioritize grid stability. The cost function analysis confirms this prioritization. The total cost J is 25.31. This is dominated by the state cost (grid stability) component ($\beta\langle x^2\rangle$), which accounts for 24.95, or 98.6% of the total cost. In contrast, the input cost (workload deviation) component ($\langle u^2\rangle$) is only 0.36, or 1.4% of the total. This breakdown confirms that the controller successfully enforces the high priority on frequency regulation, while the operational cost of deviating from the schedule is minimal.

TABLE I: Key Performance Indicators from Simulation

Metric	Description	Value
Workload Tracking		
u_{rms}	RMS Tracking Error	0.60 kW
u_{rel}	Relative RMS Error	0.48%
u_{max}	Max Tracking Error	3.93 kW
Grid Stability		
x_{rms}	RMS Frequency Deviation	1.5796
x_{max}	Max Frequency Deviation	3.3669
Cost Function (with $\beta = 10.0$)		
J_{state}	State Cost ($\beta\langle x^2\rangle$)	24.95 (98.6%)
J_{input}	Input Cost ($\langle u^2\rangle$)	0.36 (1.4%)
J_{total}	Total Cost	25.31
Computation		
T_{comp}	Computation Time	0.72 s

B. Analysis of Results

The results are visualized in Figures 2 – 5. **Figure 2** displays the time-series performance of the controller. The top panel compares the optimal power $P^*_{L,d}(t)$ with the scheduled power $P_{\rm scd}(t)$. The two trajectories are visually almost indistinguishable, reinforcing the low 0.48% RMS tracking error. The second panel shows the control deviation $u^*(t)$ (the tracking error). This signal is centered at zero and exhibits small, brief spikes that are precisely timed to counteract the sharp changes in the total demand, thereby dampening their effect on the grid.

979-8-3315-8138-1

The fourth panel shows the resulting frequency state $x^*(t)$, which deviates when load changes occur but is actively driven back toward zero by the controller's actions. **Figures 3 and 4** provide an analysis of the cost and power distributions. The left panel plots the instantaneous cost components over time. It visually confirms the numerical results from Table I; the total cost (black-dashed line) is composed almost entirely of the state cost (green line), while the input cost (red line) remains near zero. The right panel compares the probability density functions of the optimal and scheduled power. The nearly identical overlap of the histograms and kernel density estimates provides further statistical validation of the controller's exceptional tracking fidelity. **Figure 5** presents a sensitivity analysis of the trade-off parameter β. This plot illustrates the fundamental trade-off between the two primary objectives. As β is increased (x-axis), the priority on grid stability becomes higher. Consequently, the RMS of the frequency deviation presented in the right y-axis decreases, showing improved grid support. This improvement comes at the direct expense of workload fidelity. As β increases, both the RMS of the tracking error, and the maximal power deviation, presented in the left y-axis, increase, as the controller must deviate more significantly from the schedule to stabilize the grid. This plot effectively quantifies the trade-off, allowing an operator to select a β value that aligns with specific operational priorities for grid support and internal performance.

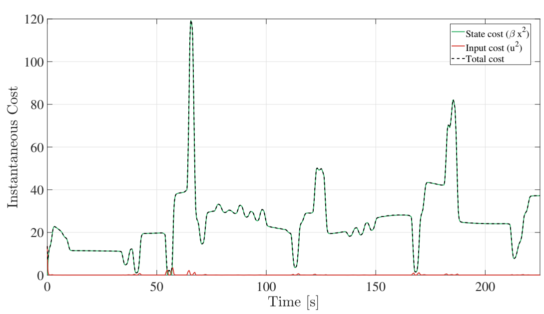

Fig. 3: Analysis of the Instantaneous state, input, and total costs over time.

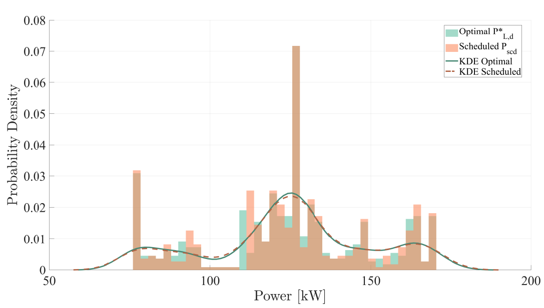

Fig. 4: Analysis of the probability density of optimal power $P^*_{L,d}(t)$ versus scheduled power $P_{\mathrm{scd}}(t)$.

IV. Discussion & Conclusions

Increasing power demand from AI workloads presents new challenges to grid stability. This paper proposes an online

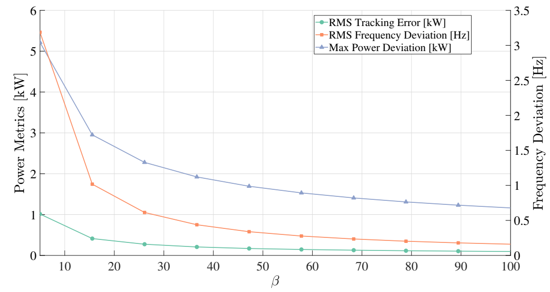

Fig. 5: Sensitivity analysis showing the trade-off between workload tracking (left axis) and grid frequency regulation (right axis) as a function of the weighting parameter β.

optimal control approach that operates data centers as flexible loads within a demand response framework, to aid regulating the grid frequency, despite operational uncertainties. The numerical results validate the proposed method's performance. With the trade-off parameter β set to 10.0, the controller maintains a low relative RMS tracking error of only 0.48%, corresponding to an absolute RMS deviation of 0.60 kW and a maximum deviation of just 3.93 kW from the schedule. Simultaneously, the controller achieves an RMS frequency deviation of 1.5796, with a maximum observed deviation of 3.3669. The cost function analysis confirms the prioritization of grid stability, as the state cost accounted for 24.95 (98.6% of the total cost), while the input cost related to workload deviation was 0.36 (1.4%). The entire optimal solution is computed in 720 ms, which confirms the low computational overhead of the proposed control approach.

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] P. Patel, E. Choukse, C. Zhang, I. n. Goiri, B. Warrier, N. Mahalingam, and R. Bianchini, "Characterizing power management opportunities for LLMs in the cloud," in *ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, (New York, NY, USA), pp. 207–222, 2024.

[3] NERC, "Characteristics and risks of emerging large loads," 2025. [Online] Available https://tinyurl.com/3jw5xyyh, Accessed September 12, 2025.

[4] X. Long, Y. Li, Y. Li, Y. Liu, P. Yang, and Z. Zeng, "Flexible load scheduling of hydrogen energy data centers: An efficiency-preferred non-convex uncertain optimization approach," *IEEE Transactions on Power Systems*, pp. 1–13, 2025.

[5] B. Zou, G. Chen, H. Zhang, and Y. Song, "Coordinating multiple geo-distributed data centers for enhanced participation in frequency regulation services under uncertainty," *Journal of Modern Power Systems and Clean Energy*, pp. 1–12, 2025.

[6] Y. Zhang, H. Tang, H. Li, and S. Wang, "Unlocking the flexibilities of data centers for smart grid services: Optimal dispatch and design of energy storage systems under progressive loading," *Energy*, vol. 316, p. 134511, 2025.

[7] S. Chen, P. Li, H. Ji, H. Yu, J. Yan, J. Wu, and C. Wang, "Operational flexibility of active distribution networks with the potential from data centers," *Applied Energy*, vol. 293, p. 116935, 2021.

[8] K. Kaur, S. Garg, N. Kumar, G. S. Aujla, K.-K. R. Choo, and M. S. Obaidat, "An adaptive grid frequency support mechanism for energy management in cloud data centers," *IEEE Systems Journal*, vol. 14, no. 1, pp. 1195–1205, 2020.

[9] F. Dul, P. Lichota, and A. Rusowicz, "Generalized linear quadratic control for a full tracking problem in aviation," *Sensors*, vol. 20, no. 10, 2020.

979-8-3315-8138-1