Grid Aware Cloud
Optimal Sizing and Operation of Battery Energy Storage Systems

GEOFFREY BIAN, The University of British Columbia, Canada

1 Introduction

As of 2025, the rapid expansion of artificial intelligence and cloud computing has led to a significant increase in data center electricity demand. Moreover, grid constraints and renewable energy integration has led to wholesale electricity prices becoming increasingly volatile. These trends motivate new strategies for reducing energy costs and improving grid interaction for these large, energy-intensive data center facilities.

Our project, Grid Aware Cloud, explores two complementary approaches for managing data center energy demand:

(1) A software based workload shifting across time and geography
(2) A physical energy buffer using battery energy storage systems (BESS)

While workload migration can be interpreted as a form of virtual energy storage, physical storage of energy remains essential for short-term power balancing, recovery, and price arbitrage. We explored several storage technologies including mechanical storage (flywheels), thermal storage uninterruptible power supplies (UPS), and the lithium-ion BESS. BESS offers the most flexibility and scalability for grid-operation, and is widely being standardized in data center developments.

This work focuses on explicitly modeling sizing and operation of a co-located BESS serving a data center load. We leverage a linear programming formulation to optimize BESS energy capacity and daily charge-to-discharge operation for a minimal total cost of usage.

All simulation code, data processing scripts, and optimization models used in this study are publicly available at https://github.com/ubc-cirrus-lab/grid-cloud-migration-estimates. This work was conducted as part of the UBC CIRRUS Lab (Cloud Infrastructure Research for Reliability, Usability, and Sustainability) at The University of British Columbia.

2 Problem Definition

In order to set up our optimization problem, we consider a single facility with a fixed electrical load over a 24-hour period. The facility draws power directly from the grid or from a co-located BESS. The BESS can and theoretically should charge from the grid during period of lower electricity cost and discharge during periods of high prices. The BESS charge and discharge is subject to power limits, energy capacity constraints, and efficiency limitations.

Our objective is to minimize the total daily operating cost in which we define as the sum of the electricity procurement costs from the grid and the battery degradation costs associated with the BESS's throughput. We also consider capital expenditure costs indirectly through an amortized degradation model.

A number of data samples were used to test our model. Namely, load data is derived from publicly

Author's Contact Information: Geoffrey Bian, geoffreybian100@gmail.com, The University of British Columbia, Vancouver, BC, Canada.

available Google data center power traces, while electricity price signals are sourced from regional market operators, ComEd and Dominion. The model is evaluated across multiple days to examine how varying demand and price conditions influence the optimal BESS size and dispatch behavior.

3 BESS Cost Modeling

3.1 CAPEX Decomposition

The total installed BESS capital cost is decomposed into power-related and energy-related components:

$$C_{\text{CAPEX}} = C_P P_{\text{BESS}} + C_E E_{\text{BESS}} \tag{1}$$

3.2 Throughput-Based Degradation Cost

Rather than modeling CAPEX as a fixed upfront expense, battery degradation is modeled as a cost proportional to energy throughput:

$$C_{\text{deg,daily}} = c_{\text{cycle}} \sum_{t=1}^{T} P_{\text{ch}}(t)\Delta t \tag{2}$$

where c_{cycle} is derived from the battery's installed cost and expected cycle life.

4 Optimization Model

The Linear Program determines the optimal Battery Energy Storage System (BESS) capacity ($E_c ap$) that minimizes the total daily cost of charging and discharging the BESS. We outline the decision variables, parameters, objective function, and constraints considered in order to achieve our simulation.

4.1 Decision Variables

- $E_{\text{cap}} \geq 0$: Optimal storage capacity (MWh).
- $E_{\text{soc}}(t) \geq 0$: State of Charge at time t (MWh).
- $P_{\text{ch}}(t), P_{\text{dis}}(t) \geq 0$: Charging and discharging power at time t (MW).
- $P_{\text{grid}}(t) \geq 0$: Net power purchased from the grid at time t (MW).

4.2 Parameters

- C_{daily}: Daily capacity cost (\$/MWh/day).
- $\lambda_E(t)$: Electricity price at time t (\$/MWh).
- $P_{\text{load}}(t)$: Original electrical load demand at time t (MW).
- $\eta_{\text{ch}}, \eta_{\text{dis}}$: Charging and discharging efficiencies.
- α: C-Rate factor (Power limit relative to capacity).
- PUE ≥ 1: Power Usage Effectiveness (accounts for auxiliary load for energy drawn from grid).
- Δt: Time step duration (1.0 hour).

4.3 Objective Function

$$\min \sum_{t=1}^{T} \lambda_E(t) P_{\text{grid}}(t)\Delta t + c_{\text{cycle}} \sum_{t=1}^{T} P_{\text{ch}}(t)\Delta t + c_{\text{cap}} E_{\text{cap}} \tag{3}$$

4.4 Constraints

Energy Balance

$$E_{\text{soc}}(t+1) = E_{\text{soc}}(t) + \eta_{\text{ch}} P_{\text{ch}}(t)\Delta t - \eta_{\text{dis}} P_{\text{dis}}(t)\Delta t \tag{4}$$

Power Balance

$$P_{\text{grid}}(t) = P_{\text{load}}(t) + P_{\text{ch}}(t) - P_{\text{dis}}(t) \tag{5}$$

Capacity and Power Limits

$$0 \leq E_{\text{soc}}(t) \leq E_{\text{cap}} \tag{6}$$

$$P_{\text{ch}}(t), P_{\text{dis}}(t) \leq \alpha E_{\text{cap}} \tag{7}$$

Initial and Boundary Conditions

$$E_{\text{soc}}(0) = 0 \tag{8}$$

$$E_{\text{soc}}(T) = E_{\text{soc}}(0) \tag{9}$$

$$E_{\text{shift}}^{\text{out,total}}(0) = 0 \tag{10}$$

$$E_{\text{shift}}^{\text{in,total}}(0) = 0 \tag{11}$$

$$P_{\text{grid}}(t) \geq 0 \tag{12}$$

$$E_{\text{cap}},\ E_{\text{soc}}(t),\ P_{\text{ch}}(t),\ P_{\text{dis}}(t),\ P_{\text{load}}^{\text{eff}}(t),\ P_{\text{shift}}^{\text{out}}(t),\ P_{\text{shift}}^{\text{in}}(t) \geq 0 \tag{13}$$

5 Implementation

The optimization problem is implemented in Python using the PuLP linear programming library with the CBC solver. The model operates on hourly time steps, using real electricity price signals and a representative data center load profile.

6 Results

6.1 Optimization Outcome

The throughput-based battery energy storage system (BESS) optimization was evaluated over a 24-hour horizon. The objective minimized total daily cost, consisting of energy procurement cost and a degradation-based cycling cost proportional to energy throughput.

The solver converged to an optimal solution with the following results:

- **Optimal storage capacity**: $E_{\text{cap}} = 5000$ MWh
- **Total daily cost (TDC)**: \$751.35
- **Daily cycle (degradation) cost**: \$515.50
- **Daily operational energy cost (OPEX)**: \$235.85
- **Baseline grid-only daily cost**: \$1,102.99
- **Net daily savings**: \$351.64 (31.88%)

These results demonstrate that under the given price signal, the optimizer selects the maximum allowable storage capacity to exploit price arbitrage opportunities and minimize grid energy procurement during high-price intervals.

6.2 Optimized Dispatch Behavior

Table 1 presents a snapshot of the optimized hourly operation. The battery charges during low-price periods and discharges to offset grid consumption during peak price intervals. Despite the large installed capacity, the state of charge (SoC) remains relatively low throughout the day, indicating that the solution is driven primarily by power constraints and price arbitrage rather than sustained energy storage.

Table 1. Optimized Hourly Operation Snapshot

Time (h)	Load (MW)	P_{grid} (MW)	P_{dis} (MW)	P_{ch} (MW)	SoC (%)
0	12.92	10.77	0.00	0.00	0
11	15.00	98.42	0.00	103.10	2
12	15.14	0.00	15.14	0.00	2
17	15.10	0.00	15.10	0.00	1
23	14.43	0.00	14.43	0.00	0

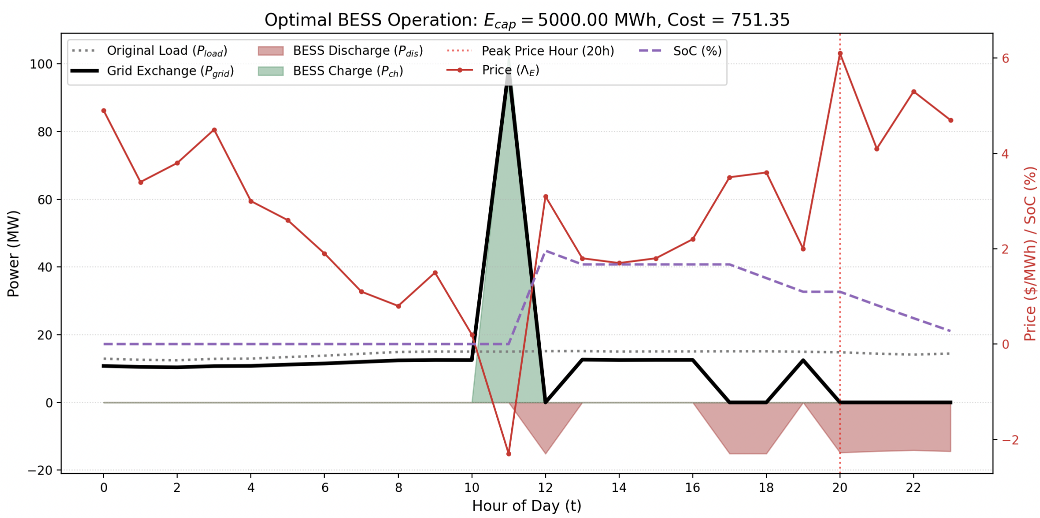

Fig. 1. Optimized daily BESS operation showing grid power exchange and electricity price signal.

6.3 Input Data and Units

6.3.1 Load Profile. The load profile was derived from Google Cluster workload traces and represents aggregate datacenter demand. Load values are expressed in megawatts (MW):

$$P_{\text{load}} = \begin{bmatrix} 12.92 & 12.59 & 12.45 & 12.87 & 12.94 & 13.40 \\ 13.80 & 14.37 & 14.90 & 15.03 & 15.01 & 15.00 \\ 15.14 & 15.16 & 15.02 & 15.06 & 15.06 & 15.10 \\ 15.10 & 14.96 & 14.81 & 14.40 & 14.12 & 14.43 \end{bmatrix}$$

6.3.2 Electricity Price Profile. The electricity price signal was obtained from the ComEd real-time tariff for October 4, 2025:

$$\lambda_E = \begin{bmatrix} 4.9 & 3.4 & 3.8 & 4.5 & 3.0 & 2.6 \\ 1.9 & 1.1 & 0.8 & 1.5 & 0.2 & -2.3 \\ 3.1 & 1.8 & 1.7 & 1.8 & 2.2 & 3.5 \\ 3.6 & 2.0 & 6.1 & 4.1 & 5.3 & 4.7 \end{bmatrix}$$

Prices are reported in cents per kilowatt-hour (¢/kWh) of , consistent with ComEd real-time pricing data. Internally, prices are converted to dollars per megawatt-hour ($/MWh).

6.4 Visualization of Optimized Operation

Figure 1 illustrates the optimized daily operation of the BESS. Charging occurs during low or negative price periods, while discharging offsets grid imports during high-price intervals, consistent with cost-minimizing arbitrage behavior.

6.5 Some Notes

It is important to understand that these highlighted results are only representative of one set of data. Different dates carry different optimization capacities. For instance, we observe many days yielding zero BESS storage capacity and are fully reliant on grid energy. More exploration should be done in order to better understand this problem. The purpose of this section is to highlight that our methodology converges to some comparable result.

7 Discussion

As highlighted, results carry across different dates and price/load profiles. This leads to different optimal BESS sizes. Across all scenarios, the optimization exhibits clear purchasing of grid power to charge the BESS when the cost is lower and discharging the BESS during peak electricity hours. The optimization finds a clear global minimum in the total cost as a funciton of storage capacity across different dates. Small batteries provide insufficient arbitrage benefit, while oversize systems incur excessive CAPEX (degradation, and capacity) costs.

Several simplifying assumptions were made, which may influence these results. Battery degradation is modeled as a linear function of energy throughput over an average BESS lifespan. We ignore factors such as temperature effects, depth of discharge, and environmental factors. Moreover, our model assumes that energy that is shifted does not need to be returned within a fixed time window besides the 24 hour cyclic energy constraint. We also only optimize our BESS capcity of data from a 24 hour period. Future extensions could extrapolate the optimization problem over a greater period, impose temporal shifting constraints, and explicitly couple storage operation with workload migration decisions.

Despite these limitations, our model effectively visualizes a balance between economic benefit and battery energy storage system integration.

8 Conclusion and Future Work

This report presents a linear optimization framework for jointly sizing and operating a battery energy storage system for a data-center-like load. By modeling degradation as a throughput-based cost, the formulation produces economically meaningful storage capacities and operational strategies.

Future work will extend the model to longer time horizons, incorporate annualized cost metrics, and compare physical energy storage against software-based workload shifting. Additional refinements include improved degradation modeling, demand charge representation, stochastic pricing, and tighter integration between energy and compute scheduling decisions.

9 References

References

[1] J. Wilkes, "Google cluster-usage traces v3," tech. rep., Google Inc., Mountain View, CA, USA, Apr. 2020. Available at https://github.com/google/cluster-data/blob/master/ClusterData2019.md.

[2] Commonwealth Edison Company, *ComEd Hourly Pricing Policy and Tariff Schedule*, 2024. Illinois Commerce Commission Filing.
[3] Dominion Energy Virginia, *2024 Integrated Resource Plan*, 2024. Case No. PUR-2024-00184.
[4] I. B. Sperstad, M. Istad, H. Sæle, and M. Korpås, "Cost-benefit analysis of battery energy storage in electric power grids: Research and practices," *Proceedings of 2020 IEEE PES ISGT Europe*, 2020.
[5] U.S. Department of Energy, "Energy storage valuation: A review of use cases and modeling tools," Tech. Rep. PNNL-28565, Pacific Northwest National Laboratory, 2022.

Acknowledgments

We would like to thank Professor Mohammad Shahrad and PhD Student Zihan Pan for their valuable contributions to this project.