Middlebox: Unlocking Datacenter Growth and Grid Decarbonization

Liuzixuan Lin
University of Chicago
Chicago, IL, USA
lzixuan@uchicago.edu

Andrew A. Chien
University of Chicago & Argonne National Lab
Chicago, IL, USA
aachien@uchicago.edu

Abstract

Datacenter growth is constrained by power grids because the constant power required by datacenters (DCs) is difficult to balance with variable solar and wind generation. DC load flexibility is the key to solving these grid problems, but it conflicts with the stable capacity needed for compute efficiency.

We propose "power Middlebox", a new system architecture, to bridge the gap. Middlebox decouples datacenter capacity and grid load, reconciling the conflict between DC capacity needs and grid load flexibility requirements. It provides each with the freedom to meet their divergent objectives. We define the Middlebox system architecture, frame its objectives, and explore designs (type and quantity of energy resources, extent of decoupling, management) in varied power grid settings.

Evaluation shows that Middlebox unlocks 460% or 170% datacenter growth with grid reliability or decarbonization constraints in a wind-dominated grid. We study decoupling in growing DC scenarios. Decoupling reconciles the conflict between grid and datacenter needs, enabling constant DC power capacity on 99.9% of days, for a cost equal to 37% of a DC's annual power bill (or 3–5% of DC TCO). Future technologies could reduce Middlebox cost by up to 70%. Further, workload flexibility can be exploited to reduce cost. These results are robust across growth scenarios and grid types. Overall, the results show that Middlebox can be deployed in small to large datacenters economically with today's technology.

CCS Concepts

• Applied computing → Data centers; **• Hardware → Power and energy**; **• Social and professional topics → Sustainability**.

Keywords

Datacenter energy management, Grid load adaptation, Sustainable computing

ACM Reference Format:
Liuzixuan Lin and Andrew A. Chien. 2025. Middlebox: Unlocking Datacenter Growth and Grid Decarbonization. In *ACM Symposium on Cloud Computing (SoCC '25), November 19–21, 2025, Online, USA.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3772052.3772220

SoCC '25, Online, USA

ACM ISBN 979-8-4007-2276-9/25/11
https://doi.org/10.1145/3772052.3772220

1 Introduction

The growth of datacenter energy consumption and carbon emissions is an ongoing and growing concern [13, 66]. The U.S. datacenter energy consumption grew by 13% annually in 2018-2022, which was already fast [83]. Since late 2022, the excitement created by chatbots such as ChatGPT, and more broadly a boom of generative AI applications such as Dall-E have significantly accelerated the growth of datacenter power demand. The explosive growth is driven by both application popularity (users and uses/day) as well as their compute-intensity per use [97].

In 2023–2028, the annual growth rate of DC power is projected at 27% [83]. A major part of this is the huge investments of the hyperscalers that run large datacenter networks. The total capital expenditures of the 11 largest hyperscalers led by Amazon, Google, Meta, and Microsoft are forecasted to be $392B in 2025, growing 38% vs. 2024 and exceeding 2022 and 2023 combined [26]. In response, grid operators are raising load forecasts: Dominion (Northern Virginia), "datacenter alley", projects datacenter load will grow from 25% of load to 49% by 2034 [24]. Similar growth is projected across the U.S., increasing DC power consumption to perhaps 12% of total U.S power consumption in 2028 [15, 18, 43, 83]. Globally, the International Energy Agency (IEA) projects datacenter consumption may reach 1,250 TWh in 2030—3 times vs. 2024 [42]!

The growth of datacenter load has overwhelmed power grids, leading to refusal (or delay of many years) for grid connection in order to maintain grid reliability [48–50]. The growth also retards grid decarbonization, intensifying the public concern over computing's negative environmental impact [65, 86].

Why is datacenter growth so hard for the grid? DCs are constant loads. However, as grids are decarbonized, they are adding large amounts of fluctuating renewable generation that depends on natural dynamics (mainly solar and wind). Reliable power requires these to be matched—all of the time. The two main approaches to overcoming the long-tail statistics for high reliability are costly—energy storage or generation over-provisioning. Growing datacenter load exacerbates the pressure on this difficult "balancing" problem. In contrast, flexible loads that align their demand with renewable generation are much easier and cheaper to support [28, 71].

One way to support rapid datacenter growth is to delay the planned decommissioning of fossil-fuel generation or build new natural gas generators, which is being adopted [24, 65, 80, 86]. Such generation is dispatchable, and can be used to balance constant DC loads. However, this is undesirable due to its carbon emissions, slowing grid decarbonization and making compute "brown".

Therefore, if DC growth is to be rapid to meet the demand of cloud and artificial intelligence, datacenters must be flexible loads, allowing them to match renewable generation. Otherwise, either the DC growth or the grid's ability to decarbonize (and computing along with it) is at risk.

Why is commercial cloud load not flexible? Much academic research has explored temporal and even spatial shifting to align datacenter power use with renewable generation [17, 21, 29, 36, 60, 62, 64, 85, 87, 90, 96, 111] or similarly to coordinate power use to help the power grid [57, 59, 61, 68, 104, 110]. These ideas have seen little adoption by commercial clouds for two reasons. The first is workload inflexibility. In clouds, workloads consist of virtual machines (VMs) whose performance abstraction is generally continuous-rate service. This QoS (quality of service) promise is difficult to flex [105, 107]. The second is that datacenters are large capital assets with extreme business pressure for maximum utilization. These limit DC flexibility. For example, Google's Carbon-aware Computing only affects datacenter power by less than 2%, far from what grid needs to decarbonize [57, 78].

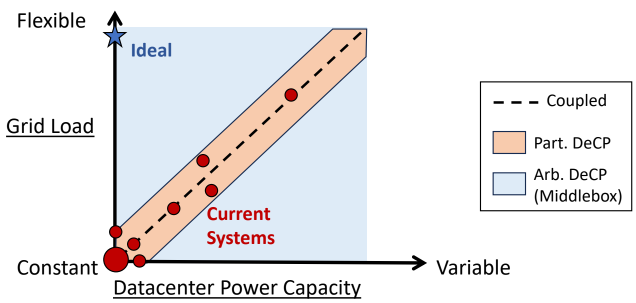

Figure 1: Load flexibility required by the grid conflicts with datacenter's need of stable power capacity when they are coupled or decoupled partially (Part. DeCP). Middlebox enables arbitrary decoupling (Arb. DeCP) that resolves the conflict.

The conflict between datacenter capacity needs and grid load flexibility requirements is difficult to solve because they are coupled with a direct grid connection (diagonal line in Figure 1), which makes datacenter power capacity equal to grid load. Commercial clouds are located at the left end of this line (the biggest red dot), with some recent efforts that explore varying capacity [78, 105]. Some systems such as [1, 29, 87] partially decouple capacity and grid load with co-located renewable generators or energy storage for one-sided benefits (e.g. carbon emissions, power cost), extending the reach (peach colored space). In this paper, we propose "power Middlebox", a key system architecture innovation that *decouples* datacenter power from grid load flexibly. This allows the grid and datacenter to operate anywhere in the blue space, much broader reach than other systems.

Our power Middlebox idea is analogous to the middlebox in Computer Networking that transforms network traffic [6]. We define the concept precisely, and explore its use for decoupling, and its effectiveness for enabling datacenter productivity and high grid load flexibility. We consider several approaches to realizing Middlebox, including generation and storage. We explore the benefits that Middlebox confers on datacenters and grid by reconciling their conflict. Specific contributions include:

- The Middlebox architecture, a system component that decouples datacenter power capacity and grid load, and its implementation design space, including a composition of energy resources and their management.
- Evaluation of Middlebox that explores decoupling, benefits for DCs and grid, and implementation costs:
  - Middlebox enables 460% more stable DC capacity growth (same grid reliability) or 170% more growth (same grid decarbonization) in a wind-dominated grid. This increase is robust across varied grid conditions.
  - Cost-performance studies that show Middlebox can deliver these benefits: with medium growth, Middlebox enables stable DC capacity on 99.9% of days at a cost equal to 37% of DC power cost or 3–5% of DC total cost of ownership (TCO). Future technologies could reduce the cost by 70%.
  - Exploiting DC workload flexibility can reduce Middlebox cost by 40%, to 2–3% increase in DC TCO.

Overall, Middlebox allows datacenters to concurrently achieve their service goals and grow sustainably. Best of all, Middlebox ideas do not require major policy or regulatory reform, and can be deployed immediately.

In Section 2, we describe the background. In Section 3, we show the tension between grid needs for load flexibility and datacenter's need for stable capacity, and introduce the Middlebox approach. We explore Middlebox design in Section 4, following with experimental methodology in Section 5. The evaluation of Middlebox benefits and costs is covered in Section 6. Finally, we discuss related work and summarize.

2 Background

2.1 Grid Decarbonization and Need for Grid Load Flexibility

Power grid decarbonization is progressing all over the world in response to the climate change. Many power grids include 10–40% renewables as of 2024 and expect 40–90% by 2030 [5, 20, 22, 25, 41]. Globally, IEA projects that renewable capacity is expected to increase over 5,520 GW during 2024–2030, 2.6 times more than the deployment of 2017–2023 [41]. The majority of increasing renewable capacity is wind and solar. Their generation costs are already lower than both fossil-fuel and non-fossil alternatives in most countries [41], but they produce time-varying generation adequacy and carbon intensity in the grid, as the timing and amount of generation depend on natural dynamics of wind and sun.

To reliably balance power demand and the increasingly variable power supply, grid operators have been calling for load flexibility, summarized as "Shape, Shift, Shed, and Shimmy" [28]. Recent studies show that even modest and occasional load reduction (shed) can enable more datacenter loads to be integrated with existing and planned grid capacity [58, 71]. As an example, Ireland, which expects DCs to account for 30% of power use by 2032, requires new DC projects be equipped with equivalent generation capacity [34]. Microsoft needed to build a 170 MW natural gas power plant to support its DCs there [51]. As grid carbon budget begins to be enforced [11, 34], "shift" that matches load and renewable generation will become more important. A range of novel programs involving datacenters, electric vehicles, intelligent appliances, and even virtual plants have been proposed [27, 37, 45, 73].

2.2 Realizing Datacenter Grid Load Flexibility

Datacenter grid load can be flexed in coupled or decoupled ways. When power capacity and grid load are coupled, changes in DC operation are needed. Typical examples include dynamic resource scaling [36] and batch workload scheduling [78, 101]. However, these approaches are rarely if ever used in commercial cloud datacenters as they directly affect workload QoS and thus can result in revenue loss. Another approach is to flex DC cooling capacity [4, 9]. The decoupled way of using generators or energy storage to flex datacenter grid load avoids impacts on DC operation. Because of the limitations of diesel backup generators (pollutant and carbon-heavy generation) and UPS batteries (limited capacity [99]) currently deployed in datacenters, this way will introduce deployment of additional energy resources that incur cost to assess. The Middlebox is a new approach in this direction.

3 Problem and Approach

Problem. The challenges caused by datacenter growth could be overcome if grids controlled datacenter load. Such flexible DC load would allow much greater DC load and improve grid reliability [52, 58, 59]. But that approach would be catastrophic for datacenters, creating high capacity variation and harming DC productivity as in Figure 2 [105, 107]. DCs under grid control would experience many capacity shortfalls that disrupt compute services. To get a sense of the business damage, a single 150 MW, 1 hour outage costs \$11.9 million [46], so the shortfalls in Figure 2 correspond to a 1,500 MWh, \$119 million outage (and that's just one day!).

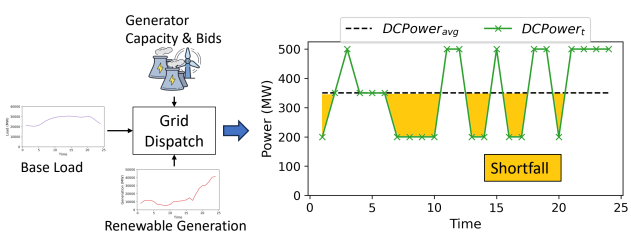

Figure 2: Datacenter capacity under grid control sees frequent capacity shortfalls that harm compute efficiency and incur huge business costs.

Clearly such frequent outages make grid control alone not feasible. The difficulty is the conflict between datacenter and power grid needs. Directly, the problem is that datacenter power is coupled to grid load. Succinctly:

$$DCPower_t = gridLoad_t, \forall t$$

To escape the conflict, the key research question is: **how can we decouple datacenter power capacity and grid load efficiently?** That is, how to support both stable computing capacity and grid load flexibility to enable datacenter growth? Underlying research questions include:

- How much decoupling is required (e.g. in power and energy)?
- How can decoupling be realized inexpensively (e.g. generators, storage, and management algorithms)?
- Can datacenter workload flexibility reduce decoupling cost? How much?
- How do these answers change with the type of power grid?

Approach. Our approach introduced a power "Middlebox" in the system architecture between datacenter and grid. The Middlebox's role is to decouple datacenter power and grid load. This decoupling can help the DC to meet capacity stability objectives AND, enable the grid reliability and renewable absorption. Middlebox redefines the coupling constraint as:

$$DCPower_t + x_t = gridLoad_t, \ \forall t \tag{1}$$

where x_t is Middlebox's power input/output that decouples datacenter power and grid load.

Middlebox decouples the two objectives, giving both sides additional freedom to pursue and satisfy independent goals. While many different uses of this freedom are possible, we focus on a single case, simultaneously achieving high datacenter productivity (meet capacity needs) and enabling grid dispatch for low cost/high renewables (flexible grid load).[1]

We consider a variety of Middlebox implementations. For simplicity, we assume Middlebox is on-premise, and "behind-the-meter" as in Figure 3, but different arrangements are possible [12, 37].

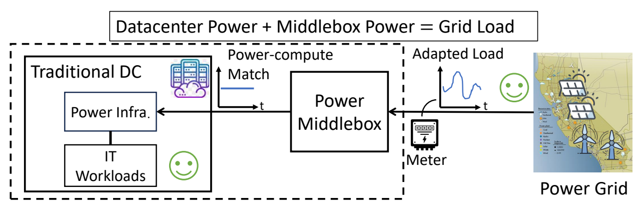

Figure 3: Middlebox decouples datacenter power capacity and grid load, resolving the conflict.

In following sections, we proceed as follows:

(1) Define the Middlebox, its functional role and requirements. This includes the temporal requirements for datacenter capacity, grid load, and the induced decoupling requirements.
(2) Define the implementation design space for Middleboxes with types of energy resources (e.g. storage, generators), composition, and management.
(3) Evaluate Middlebox performance and cost with various datacenter capacity requirements and grid settings.

Next we start from the abstract design of Middlebox and possible implementations.

4 Middlebox Design and Implementations

4.1 Abstract Design

Figure 4 shows the internal design of Middlebox consisting of control and power planes. The control plane takes information from the grid and datacenter to determine power x_t flowing into (+, from the grid) or out of (-, to the datacenter) Middlebox. This realizes decoupling of datacenter power and grid load, i.e. $DCPower_t + x_t = gridLoad_t$.

[1] Some other interesting use cases include power load smoothing for AI training [56], shaving the peak grid load, and more.

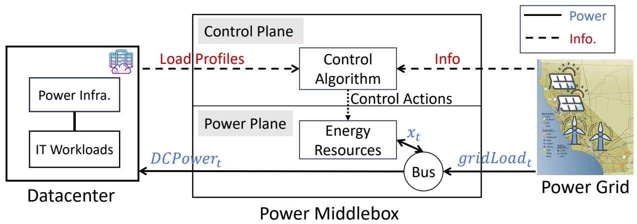

Figure 4: Middlebox Design.

4.1.1 Goals and Metrics. Decoupling enables the goals of datacenter power capacity (*DCPower*) and grid load (*gridLoad*) to be achieved concurrently.

DCPower Goal. The primary goal is **constant capacity**, which is consistent with current practice in cloud datacenters and thus require no change in other layers (e.g. scheduler). Formally:

$$DCPower_t = DCPower_{avg},\ \forall t = 1, ..., 24$$

where $DCPower_{avg} = DCPower_{max} * util_{avg}$ (average utilization). This handles inflexible capacity, $inflexPower_t$.

A variant of constant power is when there is some temporal workload flexibility (e.g. time-shiftable workloads with weak QoS needs) such as that in cloud workloads [2, 16, 78, 94]. We assume such flexible workloads account for a fraction (*tsFrac*) of daily load. The corresponding flexible capacity, $flexPower_t$, can be scheduled within 24 hours at no cost. Overall this produces:

$$DCPower_t = inflexPower_t + flexPower_t$$

The respective goals are:

$$\sum_{t=1}^{24} flexPower_t = tsFrac * DCPower_{max} * 24$$

$$inflexPower_t = (util_{avg} - tsFrac) * DCPower_{max},\ \forall t$$

In both cases, we can assess how well the *DCPower* goal is met with the percentage of time when these equations hold (power QoS), similar to the commonly used availability metric.

gridLoad Goal. For *gridLoad*, the goal is to enable **flexible datacenter grid load adapted to grid dynamics**. This meets needs of renewable-based grids to improve reliability and decarbonization [28, 59]. Formally,

$$util_{min} \leq gridLoad_t / DCPower_{max} \leq util_{max} \tag{2a}$$

$$backlog_t = backlog_{t-1} + (DCPower_{avg} - gridLoad_t) \tag{2b}$$

$$backlog_t \geq 0,\ backlog_0 = backlog_{24} = 0 \tag{2c}$$

where *gridLoad* can be set within a dynamic range $[util_{min}, util_{max}]$ with $DCPower_{avg}$ as average load throughout a day. Non-negative backlog increases with $gridLoad_t < DCPower_{avg}$ and decreases with $gridLoad_t > DCPower_{avg}$.

The benefits of flexible *gridLoad* can be measured with increase in grid capacity for datacenters under specific grid requirements.

Without Middlebox, *DCPower* would need to be highly flexible (contradicting cloud QoS needs), or *gridLoad* would need to be constant (contradicting grid reliability and decarbonization needs). Middlebox reconciles this conflict.

4.1.2 Architecture.

Control Plane. The control algorithm takes the power plane state, datacenter load profiles, and grid information as input, and then it determines x_t, the power input/output of Middlebox power plane. We use the following formulation to achieve the goals:

$$\textbf{min} \quad \max\{|DCPower_t - DCPower_{t-1}| \mid t = 2, ..., 24\} \tag{3a}$$

$$\textbf{s.t.} \quad \sum_{t=1}^{24} DCPower_t = DCPower_{avg} * 24$$

$$DCPower_t + x_t = gridLoad_t,\ \forall t$$

$$\{gridLoad_t\} \text{ given by load adaptation approach,}$$

$$\text{power plane constraints}$$

If the power plane capability allows, the control algorithm produces x_t that makes $DCPower_t = DCPower_{avg}$, $\forall t$, achieving the constant capacity goal; otherwise it will smooth the capacity curve around the average to reduce shortfalls and capacity changes destructive to compute efficiency [107]. *gridLoad* is fixed as a constraint to guarantee meeting grid needs[2]. Power plane constraints depend on Middlebox implementation.

When the temporal flexibility is considered, the control objective changes to:

$$\textbf{min} \quad \max\{|inflexPower_t - inflexPower_{t-1}| \mid t = 2, ..., 24\} \tag{4}$$

with the total power requirements $\sum_{t=1}^{24} flexPower_t = tsFrac * DCPower_{max} * 24$ (for flexible workloads) and $\sum_{t=1}^{24} inflexPower_t = (util_{avg} - tsFrac) * DCPower_{max} * 24$ (for inflexible workloads) as additional constraints.

The algorithm can be online (e.g. repeatedly make decisions every hour with techniques like Lyapunov optimization [84]) or plan-based (e.g. make a day-ahead plan based on predictions) depending on information available.

Power Plane. The power plane consists of energy resources that realize decoupling by providing power x_t. The selection of energy resource results in various implementations of Middlebox varying in capability, cost, physical footprint, etc.

4.2 Implementations with Different Types of Energy Resources

There are a variety of energy resources to build Middlebox. The implementations can be classified as generator-based, energy storage-based, and hybrid ones combining both, whose capability of handling potential *DCPower* shortfalls or surpluses varies (Figure 5). We explore various implementations to find those that can achieve *DCPower* and *gridLoad* goals at low costs (financial, carbon).

4.2.1 Generator-based Implementation. The generator is a simple solution to the datacenter capacity shortfalls by generating $x_t = DCPower_t - gridLoad_t$. With enough power capacity and continuous fuel supply, the shortfalls can be eliminated—100% coverage of constant *DCPower*. However, the generator cannot utilize surplus grid power when $gridLoad_t$ is required to be high. For Middlebox, the generator should have controllable power output

[2]The complementary strategy would be constant *DCPower*, best-effort flex *gridLoad*.

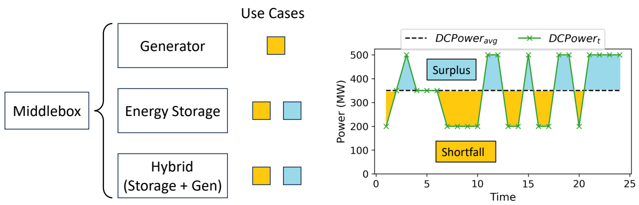

Figure 5: Middlebox can be implemented with various types of energy resources. The resources affect Middlebox's capability to handle *DCPower* shortfalls or surpluses.

and fast-ramping capability, which grid "peaker plants" satisfy. In many datacenters, backup generators are capable, and can be dual-purposed for Middlebox, saving capital expense. We mainly consider the gas turbine for its high ramp rates and lower carbon emission rate vs. other fossil fuels.

Sizing. The power capacity (maximum power output) of generator determines the deepest capacity shortfall that can be filled. When the power capacity $\geq DCPower_{avg} - DCPower_{max} * util_{min}$, all the capacity shortfalls can be eliminated.

4.2.2 Energy Storage-based Implementation. A natural way to shift datacenter's consumption of grid power is adding energy storage to each datacenter. The storage is discharged during the shortfalls and charged with power surpluses to realize decoupling, subject to:

$$s_0 = s_{24} = SoC_{max} * sEnergyCap \tag{5a}$$

$$SoC_{min} * sEnergyCap \leq s_t \leq SoC_{max} * sEnergyCap \tag{5b}$$

$$-sPower_{max} \leq x_t \leq sPower_{max} \tag{5c}$$

$$s_t = s_{t-1} + x_t \tag{5d}$$

where $sPower_{max}$ and $sEnergyCap$ denote power and energy capacity respectively. s_t denotes the stored energy at the end of t and x_t denotes charging rate (<0 means discharging). Eq. (5a) makes state of charge (SoC) circular every day for continuous operation, and we set it to the maximum (SoC_{max}). Constraints (5b)–(5d) are typical for energy storage, including depth of discharge (DoD), power rate limit, and stored energy transition.

We omit charging and discharging energy losses here for simplicity, which makes $\sum_t x_t = 0$ (required by constraints (1), (2a)–(2c)) and Eq. (5a) hold at the same time. In reality the losses need to be compensated by charging with additional grid power consumption. Therefore, this approximation requires high round-trip efficiency (e.g. 85–90%) that produces small energy losses. We primarily consider the Li-ion battery, which is mature and used in many utility-scale battery energy storage (BES) systems. There could be cheaper alternatives in the future. For example, gravity energy storage that stores energy in a heavy lifted mass is thought to be cheaper long-duration energy storage with similar round-trip efficiency [98].

Sizing. The capacity of energy storage systems is broken down into power capacity and energy capacity—maximum charge or discharge rate (MW) and amount of stored energy (MWh). The major task of sizing for energy storage is to find the minimum energy capacity (for reducing costs) that satisfies datacenter capacity goals in a set of scenarios, i.e. robust decoupling. Similar to other resource sizing processes (e.g. server provisioning), we formulate an offline optimization problem over multiple historical scenarios with known adapted $\{gridLoad_t\}$. The optimization for one scenario is:

$$\begin{aligned} \textbf{min} \quad & sEnergyCap \\ \textbf{s.t.} \quad & constraint\ (1), (5a)\text{–}(5d), \\ & DCPower\ constaints \end{aligned} \tag{6a}$$

With $energyCap_i^*$ denoting the solution of scenario i, a scenario set (n days) produces a solution set $\{energyCap_i^* \mid i = 1, ..., n\}$. The p-th percentile corresponds to $p\%$ power QoS, i.e.

$$Prob(DCPower \text{ goals met}) \geq p\%$$

Relaxing this statistical guarantee (coverage) reduces needed energy capacity (and cost) with more power QoS violations.

4.2.3 Hybrid Implementation. The energy storage and generator can also be combined to form a hybrid implementation, which opens up more possibilities. For example, cheaper storage technologies with lower round-trip efficiency (vs. Li-ion battery), such as thermal energy storage and hydrogen energy storage, become feasible options with a complementary generator to fill the energy losses. One control strategy is to charge the storage with all excess power when $DCPower_t < gridLoad_t$, and provide $x_t = DCPower_t - gridLoad_t$ together with the generator when $DCPower_t > gridLoad_t$. Specifically, we consider hydrogen storage with fuel cells, which is a clean power solution but cannot serve as an energy storage-based Middlebox because of its low round-trip efficiency (typically 31% [98]). With hydrogen transported from external sources, it becomes a feasible hybrid implementation.

Sizing. The required power and energy capacity vary with the combination. For the hydrogen system, the fuel cell can be sized as a generator, and the electrolyzer + hydrogen storage (e.g. salt cavern) can be sized as energy storage.

5 Experiment Methodology

5.1 Power Grid Simulation

We incorporate a power grid model in experiments to 1) capture the grid's dynamic load flexibility requirements; 2) accurately assess the grid impacts of datacenter growth and load adaptation. The power grid model is based on a reduced California power system (CAISO) [75]. We add solar generators based on real sites and vary the wind-solar ratio to study a variety of power grids.

The 4 seasons and weekday/weekend form 8 day types, each with a set of profiles including base load, import, and renewable generation (wind and solar excluded). We use 40 wind and solar generation scenarios for each season to model their day-to-day variation. We model three types of power grids: (1) 60%[3] wind, 0% solar (**wind-dominated**); (2) 30% wind, 30% solar (**balanced**); (3) 0% wind, 60% solar (**solar-dominated**). Each covers some grids in reality. For example, SPP (Southwest U.S.) is wind-dominated, ERCOT is balanced, and CAISO is solar-dominated. Renewable penetration exceeding 60% is expected in many grids in 2030–2040 [22, 34, 72]. Figure 6 shows the grid's base load and average renewable generation on spring days. All three types of grids exhibit

[3]Penetration: ratio of generation to grid demand.

mismatch between load and renewable generation, especially with solar generation concentrated in daytime.

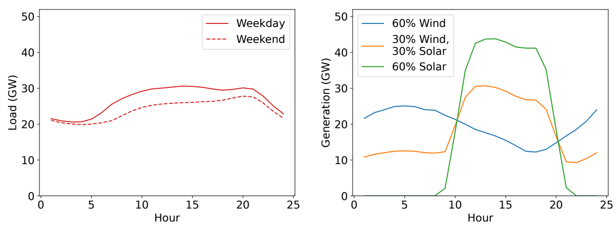

(a) Base Load. **(b) Wind + Solar Generation.**

Figure 6: Grid Base Load and Wind/Solar Generation (Spring Example). Flexible load is easier to add given the load-generation mismatch.

30 datacenters, each with same capacity $DCPower_{max}$ and 70% average utilization, are added at random buses in the grid[4]. We study two growth scenarios for datacenters: medium (500 MW DCs) and maximum (800 MW DCs). They correspond to the midpoint and maximum of datacenter growth that respects grid reliability and decarbonization requirements (explained in Section 6.1.1). These scenarios are unlocked by load adaptation with Middlebox. Fixed-load datacenters violate grid requirements with the same growth.

Table 1: Summary of Power Grid and Datacenter Settings

Power Grid	
Base Load	27,289 MW on average
(Wind%, Solar%)	(60, 0), (30, 30), (0, 60)
Datacenter Growth Scenarios ($util_{avg}$ = 70%)	
Medium	30 $*$ 500 MW
Maximum	30 $*$ 800 MW

With the grid topology and load/generation profiles, we simulate grid operation by solving the direct-current optimal power flow (DC-OPF) problem [39], the predominant method for dispatching generators and transmission. The reliability metric—amount of load shed—is part of the solutions, and the sustainability metric—carbon emissions—is calculated from dispatched generation and generation emission rates. For reviewer convenience, we provide more details about DC-OPF in Appendix A.

5.2 Decoupling Model

5.2.1 Grid Load Adaptation Approaches. DC load levels are decided by the grid dispatch (DC-OPF) and known by the DCs at the start of the day. Load can vary over a **large dynamic range of [40%, 100%]**. This is similar to the "controllable load resources" program in ERCOT that enables faster grid-connection for large loads.

5.2.2 Management of Middlebox. We discuss control without and with workload flexibility separately as the decision space changes.

[4]One datacenter corresponds to multiple buildings on a campus.

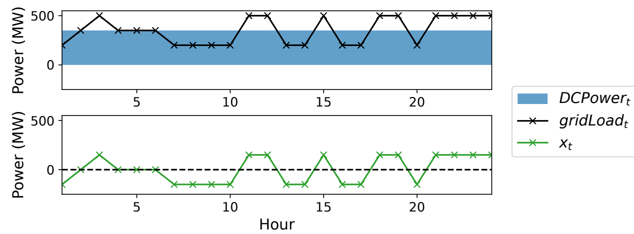

Figure 7: Energy Storage-based Middlebox Daily Operation with Constant $DCPower$ **goal.**

Constant Capacity. The control algorithm aims to keep $DCPower$ constant. With $\{gridLoad_t \mid t = 1, ..., 24\}$ known in advance from the grid, for generator-based and hybrid Middlebox implementations, the algorithm schedules generation when there is a shortfall ($gridLoad_t < DCPower_t$). For energy storage-based Middlebox, it's more complicated due to finite capacity. This requires a linear program (Optimization (3)) to decide charge schedule $\{x_t \mid t = 1, ..., 24\}$, where $gridLoad_t - x_t$ meets the constant $DCPower_t$ target when storage capacity allows (see Figure 7 example).

Flexible Capacity. We assume that the DC scheduler (e.g. Borg [94]) will schedule flexible load in response to available power (as in [78]). Thus, Middlebox management can exploit power scheduling freedom expressed in $flexPower_t$.

Our algorithm consider several goals, ordered to maximize $DCPower$ QoS, using the corresponding objective functions:

- (O1) Capacity requirements of inflexible and flexible workloads: **min** $\max\{|inflexPower_t - inflexPower_{t-1}| \mid t = 2, ..., 24\}$ with the total power constraints.
- (O2) Reduce discharged energy of Middlebox to reduce fuel cost or improve storage lifetime: **min** $\sum_{t=1}^{24} |x_t|$ for $x_t < 0$.
- (O3) Minimize changes in power level of Middlebox to improve lifetime: **min** $\sum_{t=1}^{24} x_t^2$.

The algorithm solves O1–O3 in sequence using linear (O1, O2) and non-linear (O3) solvers:

(1) Optimize O1 with the total power and power plane constraints. This determines $\{inflexPower_t\}$.
(2) With $\{inflexPower_t\}$ fixed from step #1, optimize O2. This determines total discharged energy.
(3) With $\{inflexPower_t\}$ and total discharged energy fixed from step #2 , optimize O3. This determines the final schedule of charge $\{x_t\}$ and flexible capacity $\{flexPower_t\}$.

An example of the output of the algorithm is depicted in Figure 8. This produces the best possible $DCPower$ QoS (blue), given the Middlebox energy resources. And subject to that, tries to minimize operational costs, such as fuel/discharged energy (green). Finally, charging schedule (if needed) is optimized for storage lifetime (green) and flexible compute schedule (orange) is determined.

5.3 Evaluation Methodology

Evaluation covers both design (sizing) and effectiveness of operation (Figure 9) for various implementations (Table 2).

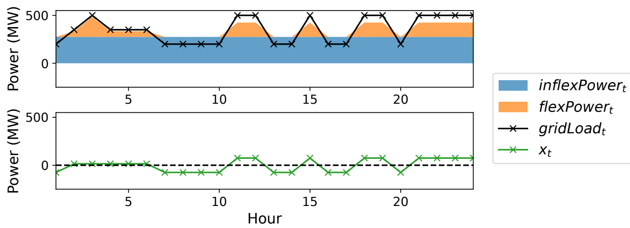

Figure 8: Energy Storage-based Middlebox Daily Operation with Flexible Capacity ($DCPower = inflexPower + flexPower$). Middlebox use is reduced.

Table 2: Middlebox Implementations to Evaluate

Energy Resource	Category
Natural Gas Turbine	Generator
Li-ion Battery	Energy Storage
Gravity Energy Storage	Energy Storage
Hydrogen Storage w/ Fuel Cell	Hybrid

5.3.1 *Design Process.* For each grid type, we create a 320-day datacenter grid load adaptation trace (40 wind and solar generation scenarios $*$ 8 day types) simulating grid total load, renewable generation, and dispatch. This trace indicates the desired adapted $gridLoad$ for the DC that achieves maximum grid decarbonization. The first 240 days (75%) of trace is used to size Middlebox energy resources for each power QoS goal.

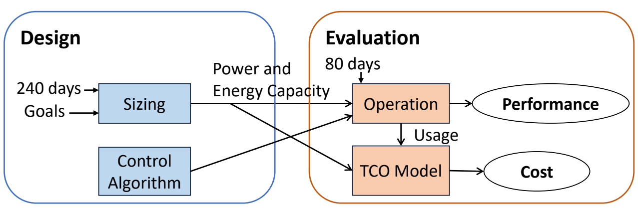

Figure 9: Middlebox Evaluation.

5.3.2 *Operation Evaluation.* To evaluate Middlebox operation, we use the remaining 80-day trace (25%) and capture the performance of designed Middlebox (energy resources and control algorithm). We also assess Middlebox cost, considering the designed capacity and fuel/energy usage.

5.3.3 *Middlebox TCO Models.* Total cost of ownership (TCO) is a widely used cost metric consisting of capital expenses (Capex) depreciation and operational expenses (Opex) incurred periodically. We calculate the TCO of various Middlebox implementations for comparison. The values used are listed in Appendix C.

Generator. The levelized cost of energy (LCOE) measures the average cost of electricity generation for a generator in its lifetime [54]. All the costs, including generator facility Capex, fuel, maintenance, etc. are included in LCOE. Therefore, for generator-based Middlebox implementations, the TCO can be calculated as annualized generation $*$ LCOE. For the natural gas turbine, we refer to the LCOE for "gas peaking" in [54], which is \$200/MWh (2025 midpoint). In the case of using an existing backup generator, only the fuel cost is counted.

Energy Storage. Despite the existence of levelized cost of storage (LCOS), a metric similar to LCOE, we do bottom-up cost assessment for a more accurate estimation because the usage as Middlebox differs from classical assumptions of LCOS. The Capex of energy storage can be broken down into energy component (EC) and power component (PC) costs. EC mainly includes energy blocks (e.g. batteries, cavern), while PC mainly includes electrical infrastructure and interconnection. Their costs are proportional to the energy and power capacity respectively.

The composition of Opex varies with storage technology. For Li-ion batteries, the Opex covers operation, maintenance, and capacity augmentation to mitigate battery degradation. We assume linear battery degradation by cycle[5]. Overall, BES-based Middlebox TCO (\$/year) can be calculated by:

$$\frac{power_{max} * unitCapex_{PC} + energyCap * unitCapex_{EC}}{DeprPeriod_{MB}}$$
$$+ power_{max} * unitCapex_{sys}(0.01 * cycle_{avg} + 0.015) + lossCost$$

where $unitCapex_{sys} = unitCapex_{PC} + unitCapex_{EC} \cdot duration$ is a system cost metric. $lossCost$ is calculated based on typical round-trip efficiency, daily discharge, and power price. We fix $sPower_{max}$ at $DCPower_{max}$ (more than sufficient for power requirement). And then the BES TCO can be regarded as a function of designed energy capacity and discharge in operation.

For the gravity storage, the TCO calculation is the same except that Opex is just proportional to the power capacity.

Hybrid. For the hydrogen system, the TCO calculation is based on the energy storage TCO model above. The Opex is proportional to the fuel cell capacity. The fuel cell and electrolyzer—two power components—are separately sized to $DCPower_{max}$ and $DCPower_{max} - DCPower_{avg}$ to match the maximum power output and input. Energy capacity is based on the maximum single-day hydrogen use. $lossCost$ is the cost of the transported hydrogen for compensating round-trip energy losses.

5.3.4 *Datacenter TCO Model.* We calculate datacenter TCO for context of Middlebox cost:

$$\frac{Capex_{IT}}{DeprPeriod_{IT}} + Opex_{IT} + \frac{Capex_{DC}}{DeprPeriod_{DC}} + Opex_{DC} + Power\ Cost$$

where $Capex_{IT}$ and $Opex_{IT}$ are IT equipment's (mainly servers) capital and operational expenses calculated based on $ITPower = DCPower_{max}/PUE$ (power usage effectiveness). $Capex_{DC}$ denotes DC construction and non-IT equipment costs. $Opex_{DC}$ denotes DC operational costs such as maintenance and security[6]. The value settings are based on industry reports (listed in Appendix C).

5.3.5 *Metrics.* Performance metrics include potential DC growth and power QoS:

- Grid Capacity for Datacenters (MW): maximum datacenter capacity that can be added to the grid at the same reliability or decarbonization level as the base grid.

[5]A cycle is counted when accumulated discharged energy reaches $energyCap * DoD$.
[6]These costs are also expressed as dollar per MW IT power.

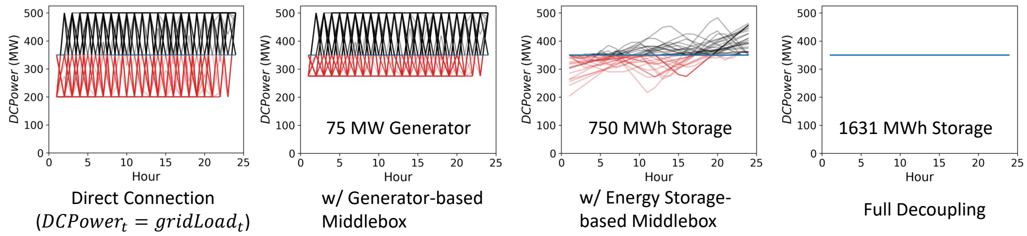

Figure 10: Middleboxes provide different quantities of decoupling, depending on type and quantity of energy resources. Capacity shortfalls and surpluses are highlighted with red and black respectively. (Wind-dominated Grid Example)

- Satisfying the power QoS: fraction of scenarios in which the overall capacity needs are satisfied.

In all simulations, there are multiple DCs at different sites in the power grid, so we aggregate the statistics across them. For each metric, we report the average across the day scenarios, weighted for weekday/weekend and season of the year.

Cost Metrics include:

- Middlebox TCO (total cost of ownership)
- Middlebox TCO as percentage of Annual Power Cost
- Middelbox TCO as percentage of Annual Datacenter TCO

5.4 Simulation Implementation

Simulation of the entire system is implemented in Julia with JuMP package [23]. The Gurobi optimizer is used for solving the optimization problems [35].

6 Middlebox Evaluation

We first explore how the grid-controlled load adaptation unlocked by Middlebox enables DC growth. Second, we study the cost of Middlebox to decouple DC power and grid load with different implementations. Finally, we explore how DC workload flexibility can reduce decoupling cost.

6.1 Middlebox Benefits

6.1.1 More Grid Capacity for Datacenters. Middlebox can enable larger feasible grid DC load, enabling more growth! First, grid load growth is typically constrained by the commitment to reliable power, so we hold reliability constant. Second, an important goal is to at least avoid retarding power decarbonization, so we hold average grid carbon intensity constant. In both cases, we explore the DC growth that Middlebox enables.

The results for these two scenarios are presented in Figure 11. For constant reliability, Middlebox decoupling enables much greater DC grid capacity, 5.6x and 2.7x in wind-dominated and balanced grids respectively. The *gridLoad* flexibility is sufficient to match generation fluctuations, but ultimately grid datacenter capacity is limited by transmission (and ultimately generation). Middlebox decoupling is greater in the wind-dominated grid. In such grids, load shedding first occurs at grid load peaks (e.g. summer daytime), so the ability to reduce DC *gridLoad* at these times ensures reliability. In the solar-dominated grid, without Middlebox, solar generation matches grid load peaks better, so DC capacity is greater. Adding Middlebox provides little benefit.

For constant decarbonization, Middlebox decoupling also enables much greater DC grid capacity, 2.7–5 fold by reducing the mismatch between added load and renewable generation. Because it can cover all kinds of mismatches—both in peak and other times—Middlebox produces large benefits in all of the grid scenarios.

There is widespread reporting that the availability of power is limiting DC growth in North America and Europe [34, 44, 92]. Because their customers are plentiful, each DC increases their profits, making the appeal of unlocking DC grid capacity clear. For example, 10,000 MW capacity of CPU servers ($200B capex) with 20–30% return on invested capital (ROIC) [88], could mean an additional $40–60B annual profit!

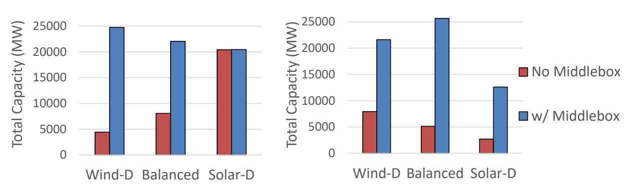

(a) Constant Reliability. **(b) Constant Carbon Intensity.**

Figure 11: Datacenter Grid Capacity w/o Middlebox Decoupling (fixed reliability or decarbonization).

To cover the wide range of potential DC growth just discussed, we use two scenarios: 1) medium growth (30 500-MW datacenters) and 2) maximum growth (30 800-MW datacenters). We study the benefits of Middlebox decoupling in both.

6.1.2 Robust Decoupling for Stable Datacenter Capacity in Decarbonized Grid. Not all Middleboxes can achieve constant DC capacity in the face of changing *gridLoad*, optimized for grid efficiency. Consider Figure 10 that illustrates such limits. The first graph (left), shows 80 *gridLoad* profiles produced for a datacenter by a wind-dominated grid. The objective for stable DC capacity would be a flat line. The second graph shows a generator-based Middlebox can fill the shortfalls but cannot utilize the power surpluses. The third and fourth graphs (right) show the decoupling possible with two storage-based Middlebox designs, with different quantities of storage. Clearly, achieving stable DC power requires both the right type and quantity of energy resources.

Because capacity for energy storage-based Middleboxes is designed statistically, it can fall short of the decoupling target. We validate our sizing by comparing the achieved coverage with the target

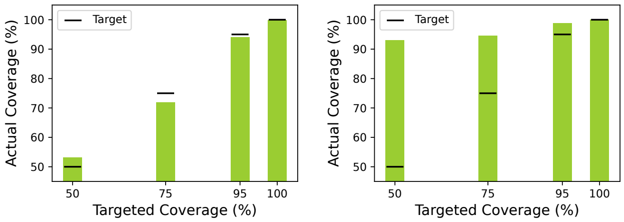

(a) **Wind-dominated Grid.** (b) **Solar-dominated Grid.**

Figure 12: Actual Coverage vs. Middlebox Designed Coverage.

coverage in Figure 12 (medium-growth). Achieved coverage generally reaches the target with slight deviation in the wind-dominated grid. The deviation is due to day-to-day wind variation. For 100% coverage target, achieved coverage is 99.9%, a level that exceeds the reliability of many power grids. In the solar-dominated grid, coverage of 100% can be achieved. Because the scenarios are highly correlated, the transition to nearly 100% coverage is an abrupt at a particular storage capacity threshold.

The solar-dominated grid's results can represent the balanced grid, as the adapted *gridLoad* profiles are similar. In the maximum-growth scenario, our sizing approach achieves similarly high coverage of constant power capacity (99.3% and 99.9% in wind-dominated and solar-dominated grids respectively), showing the robustness against changing grid conditions.

6.2 Middlebox Cost and Limitations

6.2.1 Cost of Different Extents of Decoupling. For the gas generator, full decoupling requires power capacity ≥ the maximum capacity shortfall due to grid control ($DCPower_{avg} - DCPower_{max} * util_{min}$). Smaller power capacity partially fills the shortfalls as Figure 10 shows. We vary the power capacity from 50% to 100% (75–150 MW) of the full decoupling requirement, showing the corresponding average daily generation in Figure 13 (medium growth). As expected, a greater extent of decoupling does require larger generators and more generation. In the meantime, the generator's inability to utilize surplus grid power leads to wasted energy, the amount of which is independent of generator capacity.

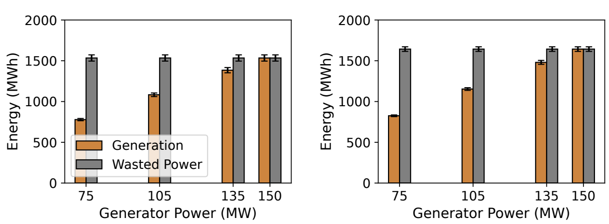

(a) **Wind-dominated Grid.** (b) **Solar-dominated Grid.**

Figure 13: Average Daily Generation and Wasted Power for Gas Generation-based Middlebox. Error bars reflect variation (stdev) across DC sites (same use in following figures).

For battery-energy storage-based (BES) Middlebox, we calculate its TCO with the designed energy capacity and average daily discharge cycles in operation (Figure 14) for various coverages of constant power capacity. The two bars of each group correspond to the medium growth and maximum growth. In the wind-dominated grid (Figure 14a), as the targeted coverage increases from 50% to 100%, the needed energy storage capacity increases from 831 MWh to 1,768 MWh in the medium-growth scenario. The average daily discharge cycles decrease at the same time, slowing the TCO growth. At 100% coverage (full decoupling), the BES experiences 1.1 cycles/day—around typical BES usage of 1 cycle/day [10, 70]. The average Middlebox TCO is $88 million/year for a 500 MW datacenter (Figure 15a). Slightly relaxing the targeted coverage to 95% can reduce the cost by 17%. This results in rare capacity variation, which may be handled by DC management (e.g. power capping [102, 103, 108]) with little workload impact. For larger datacenters (maximum growth), the energy capacity needs to be scaled up by a ratio slightly lower than the datacenter capacity growth.

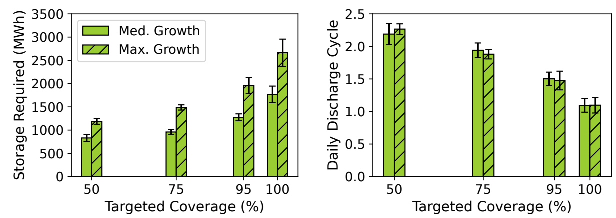

(a) **Wind-dominated Grid.**

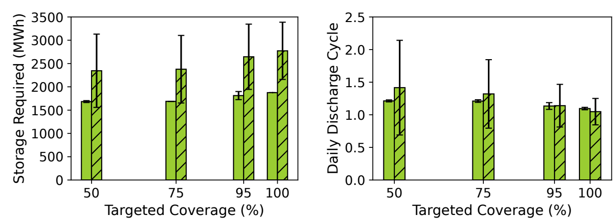

(b) **Solar-dominated Grid.**

Figure 14: Required Energy Storage Capacity (left) and Average Daily Discharge Cycles (right) for Targeted Coverages of Constant Power.

For solar-dominated grid, solar generation concentrated in daytime produces more similar *gridLoad* profiles from day to day, making the storage sizing results (Figure 14b) less sensitive to the targeted coverage change, as well as the associated cost (Figure 15b). Similar results are observed in the balanced grid because of similar adapted *gridLoad* profiles. Storage capacity required for full decoupling (1,875 MWh) is slightly higher than the wind-dominated grid. Comparing the growth scenarios, transmission constraints with higher datacenter load (max. growth) produce larger site variation.

We then compare the cost of full decoupling with the battery storage and natural gas generator. The generator TCO is calculated as LCOE ∗ annualized generation. The facility and fuel costs (based on the U.S. average price of $37/MWh) are separated as the facility cost might be excluded when an existing backup generator is used

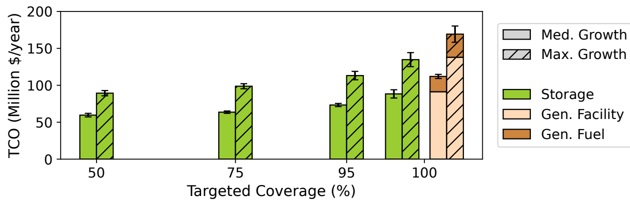

(a) Wind-dominated Grid.

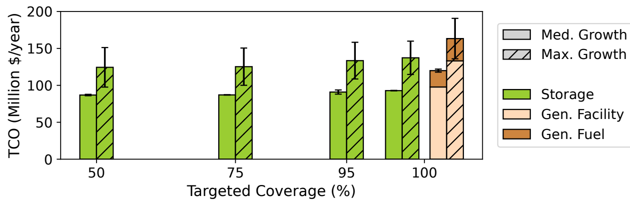

(b) Solar-dominated Grid.

Figure 15: Energy Storage TCO vs. Gas Generator TCO.

as Middlebox. In the medium-growth scenario, the total cost of gas generator is about 130% or 25% of the storage cost depending whether the whole system or only fuel cost is counted.

6.2.2 Middlebox Cost in Context. To put Middlebox cost in context, we compare it to DC TCO and outage costs.

In Figure 16, we compare BES-based Middlebox and DC (500 MW) TCO for two configurations: all CPU servers and all AI/GPU servers. To be conservative, for Middlebox, we use full decoupling (100%) in the most expensive setting (the solar-dominated grid). Compared to DC TCO, Middlebox TCO is small, only 2.7–5.6%. Compared to power cost, Middlebox TCO would be a 39% increase.

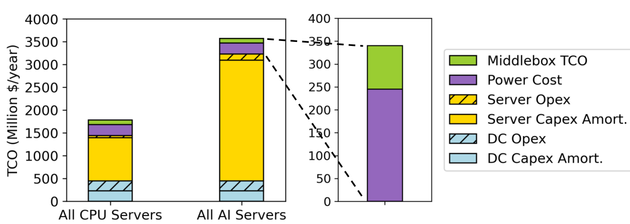

Figure 16: Middlebox TCO vs. Datacenter TCO.

Comparing Middlebox TCO to outage cost, the picture is more extreme. Outage cost is much greater, so we show the ratio of avoided outage cost (relative to direct connection) to Middlebox TCO in Figure 17. At all targets, the avoided outage cost is more than 500x Middlebox cost. This shows the cost effectiveness of Middlebox in reducing outage cost.

6.2.3 A Broader Comparison of Middlebox Implementations. Cost and carbon footprint of Middlebox implementations based on various storage types (gravity, battery), generation (natural gas), and hybrid hydrogen storage-fuel cell are plotted in Figure 18. We show the 100%-coverage implementations with medium DC growth. Annualized cost is from our TCO models or LCOE. The annualized carbon footprint is based on NREL's life cycle assessment [53].

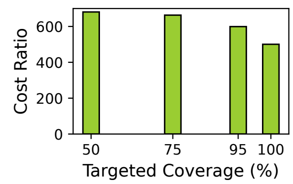

(a) Wind-dominated Grid.

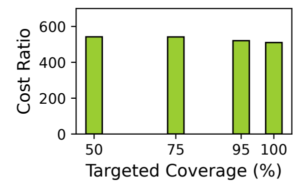

(b) Solar-dominated Grid.

Figure 17: Avoided Outage Cost-Middlebox TCO Ratio for Different Designed Coverages of Constant Power.

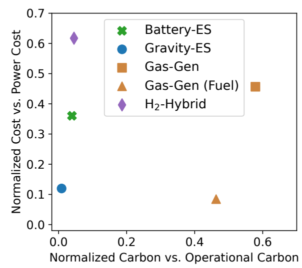

(a) Wind-dominated Grid.

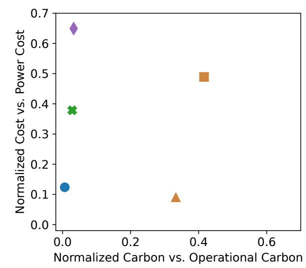

(b) Solar-dominated Grid.

Figure 18: Cost and Carbon Footprint of Various Middlebox Implementations.

Natural gas generators (fuel only) are the cheapest solution with gravity energy storage (future technology) close to it. Battery storage (Li-ion) is 4 times more expensive. H_2 hybrid storage systems are nearly 6 times more expensive. Of course, the gas generators produce significant carbon emissions, while the other approaches are as much as 10-52x lower. Several Middlebox implementations will continue to improve. For example, battery costs are expected to be 30% lower by 2030, and some project H_2 hybrid system TCO to drop 50% in the same period. Fueling the gas generator with H_2 or adding carbon capture system can make it greener at higher LCOE.

6.3 Exploiting Workload Flexibility to Reduce Decoupling Cost

Many academic researchers and recently commercial cloud researchers have noted that datacenter workloads can have significant flexibility [76, 78]. Here we consider how DC workload flexibility can be exploited to reduce Middlebox decoupling cost.

We assume a fraction of workloads (5–15%) and corresponding datacenter energy consumption is flexible, which can be executed any time within a 24-hour period, matching estimates from industry publications and optimistic future projections (e.g. growing ML batch jobs or changing service-level agreements) [2, 76, 94].

Using the modified Middlebox control algorithms (Section 5.2.2) produces the energy storage, operation, and TCO results (medium growth, 100% targeted coverage) in Figure 19. Comparing to 100% inflexible compute workload, Middlebox storage requirement and thereby TCO can be reduced linearly (fraction of flexible load). This benefit applies to all datacenter sites and grid types uniformly. With 15% time-shiftable load, the TCO is reduced by about 40%.

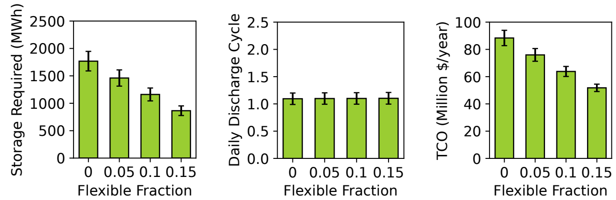

Figure 19: Left to Right: Required Energy Storage Capacity, Average Daily Discharge Cycles, and Energy Storage TCO for Various Flexible Workload fractions. (Wind-dominated Grid)

7 Discussion and Related Work

7.1 Enabling Datacenter Load Growth

Hyperscalers have signed power purchase agreements (PPA) for renewable generation to "power" their datacenters [14, 31, 33, 69]. Renewable PPAs' contribution do not solve the conflict between variable renewables and constant DC load that is the focus of this paper, and limits DC grid capacity. Further, hyperscalers' 15–20% annual growth of location-based carbon emissions shows DC load is heavily dependent on fossil fuel-based generation [14, 31, 69]. Better hourly matching requires both increased DC *gridLoad* flexibility and more renewable generation [2].

Recent press about hyperscalers signing contracts to restart nuclear reactors [81] or build new ones [93] reflects the premium they will pay for grid power capacity and decarbonization [54, 81]. Microsoft's contract for restarting a plant reveals a $30/MWh premium over U.S. average retail price of $80/MWh. The premium for new plants is the difference between LCOE and retail.

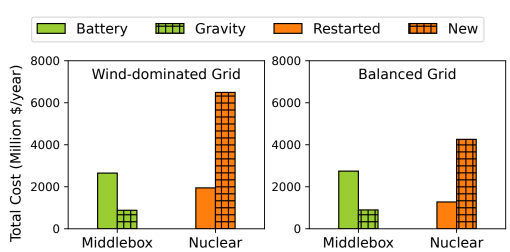

Figure 20: Middlebox Cost vs. Nuclear Power for Same Datacenter Growth.

We compare Middlebox TCO with such a scenario (nuclear PPAs) in Figure 20. The result shows that for 15,000 MW datacenter growth. The storage-based Middlebox has cost comparable to a restarted nuclear plant of which there is limited supply. In the future, gravity storage may be much cheaper. In all cases, the Middlebox approach is much cheaper than new nuclear and has obvious advantages such as low-risk & faster deployment, easy scalability, and avoids nuclear's safety and environmental concerns [38, 47].

DC operators may prefer solutions they can deploy incrementally, and it appears power grids prefer load flexibility solutions [37]. Grid capacity studies show DC load growth can be accelerated if it is flexible during peak periods [58, 71], but provide no solutions for how to make it flexible. Middlebox addresses these needs.

7.2 Datacenter Grid Load Adaptation

Temporal load shaping or shifting—scheduling datacenter workload to exploit variation of power grids' carbon and price dynamics is well-studied [2, 21, 29, 36, 60, 62, 64, 90, 91]. Recently, commercial operations exploit workload flexibility to increase hourly matching of datacenter energy consumption and renewable generation ("24×7" [31] or "100/100/0" [14]). For example, Google's "carbon-aware computing" [78] creates day-ahead, carbon-optimized capacity plan, which shapes real-time power consumption. Realistic workload flexibility and poor server energy proportionality limits load shifting to a few percent [82, 105, 107].

Most DC load adaptation proposals assume DC load changes do not affect grid dispatch, failing to capture grid dynamics. This is increasingly unrealistic as DC load exceeds 15% and even 25% in a growing number of grids [24, 34, 44]. Recent studies combining grid modeling with DC load adaptation, propose cooperative approaches, sharing information in two directions [32, 57, 59, 61, 68, 104, 110]. For example, one shares DC load plans in-advance, to allow grid optimization (e.g. unit commitment) based on more accurate load profiles [57]. This produces good results for the grid, but large penalties for DCs (13% capacity variation).

Demand-response efforts typically focus on peak reduction to reduce grid capital infrastructure costs, studying how to best respond to grid signals [55, 63, 67, 106] and incentives [8, 109, 112]. These peaks only affect a few hours on several days each year.

In contrast to all of above, Middlebox proposes a robust approach to decoupling DC power capacity and grid load. Further, Middlebox enables DC builders to take the initiative, without major regulatory reform, to accelerate DC load growth, while maintaining high compute resource efficiency.

7.3 Datacenter Power Management

The traditional objectives of datacenter power management are improving infrastructure utilization or power efficiency subject to workload QoS. For example, power capping [102, 103, 108] can handle shortfalls in power capacity while overclocking [89] can utilize the headroom. Intelligent cooling management can also flex the power capacity needed [4, 100]. These methods are complementary to our decoupling framework for handling occasional insufficient decoupling or reducing decoupling cost like workload flexibility.

DCs include uninterrupted power supply (UPS) with ≈10 minutes storage capacity [99, 108]. UPS energy storage can be used for ancillary services and peak reduction, reducing power cost and aiding the grid [74, 85, 96, 106, 111]. However their capacity is far too low to meet the 60% capacity changes and hours-long power draws required by Middlebox decoupling.

Some prior work assessed energy storage needs to improve renewable matching (ultimately 24×7) for a given DC [2, 79]. In contrast, Middlebox targets benefits of both DCs and the grid.

8 Summary and Future Work

DC growth is constrained by power grids. We propose Middlebox, a system architecture innovation that decouples DC power capacity and grid load, unlocking DC capacity growth across varied grid conditions. As an example, in a wind-dominated grid, Middlebox unlocks 460% (same grid reliability) or 170% (same grid decarbonization) DC capacity growth. Exploration of Middlebox implementations shows that decoupling achieves compute efficiency under flexible grid load cost-effectively. With medium growth, Middlebox provides 99.9% guarantee of constant power capacity at a cost equal to 37% of DC annual power bill. Future Middlebox technologies would reduce this cost by more than two-thirds.

The framework of Middlebox opens up many research opportunities. There are certainly many Middlebox implementations to study. For cloud researchers, Middlebox may increase the value of workload flexibility and energy proportionality. What about different types of flexibility? From the DC-grid coordination perspective, does Middlebox enable new approaches to arrangement of large loads and cooperative control?

Acknowledgments

We thank the reviewers for their insightful reviews. This work is supported in part by NSF Grants OAC-2442555, CNS-2325956, and the VMware University Research Fund. We also thank the Large-scale Sustainable Systems Group members for their support!

A Direct-current Optimal Power Flow (DC-OPF)

We borrow a typical DC-OPF formulation from [52]. The objective of DC-OPF is to minimize the daily dispatch cost:

$$\min \sum_{t\in\mathcal{T}} \Bigg(\sum_{i\in\mathcal{G}} C_i p_{i,t} + \sum_{i\in\mathcal{ND}} C_i^{nd} d_{i,t}^{nd} + \sum_{i\in\mathcal{DC}} C_i^{dc} d_{i,t}^{dc} + \sum_{i\in\mathcal{I}} C_i^{m} m_{i,t} + \sum_{i\in\mathcal{W}} C_i^{w} w_{i,t} + \sum_{i\in\mathcal{S}} C_i^{s} s_{i,t} + \sum_{i\in\mathcal{R}} C_i^{r} r_{i,t} \Bigg)$$

which consists of generation cost ($C_i p_{i,t}$), load shedding penalties ($C_i^{nd} d_{i,t}^{nd},\ C_i^{dc} d_{i,t}^{dc}$), and import/renewable generation curtailment penalties ($C_i^{m} m_{i,t},\ C_i^{w} w_{i,t},\ C_i^{s} s_{i,t},\ C_i^{r} r_{i,t}$) aggregated across locations and time. Because of the zero generation cost and curtailment penalty of renewables, the grid carbon emissions are also optimized.

The constraints include balancing at each bus, transmission and generator capacity, etc. Datacenter $gridLoad_{i,t}$ is part of the load in the balancing constraint:

$$power_{in} + power_{out} + generation = load - loadShed,\ \forall bus$$

In the grid-controlled load adaptation, DC-OPF decides all datacenters' 24-hour load levels $\{gridLoad_{i,t}\}$ respecting the flexibility constraints. This enables better grid reliability (less load shed) and decarbonization (lower carbon emissions) than fixed-load DCs.

B Datacenter Outage Cost Estimation

Compared with a constant-capacity datacenter, the shortfalls (i.e. $DCPower_t < DCPower_{avg}$) are regarded as outages. The cost estimation is based on an industry survey [46] that considers direct (e.g. equipment cost), indirect (e.g. investigation), and opportunity costs (e.g. lost revenue). With the reported \$50/sqft cost for 95-minute outages in large-scale datacenters [46] and power density of 400 W/sqft [7], the outage cost *unitOCost* is about \$0.079/(W*hr). With that, the outage cost for a datacenter in a day can be calculated as $\sum_t (DCPower_{avg} - DCPower_t)^+ \cdot unitOCost$.

C TCO Analysis Settings

For energy resources, the cost of power components is in million \$/MW, and the cost of energy components is in million \$/MWh.

Table 3: Li-ion BES Attributes (Sources: [10, 70])

Capex (80% DoD, 85% RTE, $DeprPeriod_{MB}$=15 years)	
Power Components	0.36 (2023), 0.32 (2030)
Energy Components	0.39 (2023), 0.25 (2030)
Opex (million \$/(MW·year))	
Capacity Augmentation	1%·$unitCapex_{sys}$·#cycles/day
Others	1.5%·$unitCapex_{sys}$

Table 4: Gravity Energy Storage Attributes (Source: [98])

Capex (80% DoD, 85% RTE, $DeprPeriod_{MB}$=49 years)	
Power Components	0.701 (2021), 0.695 (2030)
Energy Components	0.198 (2021), 0.168 (2030)
Opex (million \$/(MW·year))	
Fixed O&M	0.014 (2021), 0.013 (2030)

Table 5: Hybrid Hydrogen System Attributes (Source: [40, 98])

Capex (80% DoD, 31% RTE, $DeprPeriod_{MB}$=30 years)	
Fuel Cell (Power)	1.518 (2021), 0.549 (2030)
Electrolyzer (Power)	1.355 (2021), 0.389 (2030)
Energy Components	0.006 (2021, 2030)
Opex (million \$/(MW·year)	
Fixed O&M	0.017 (2021), 0.008 (2030)
Others	
Power Generation Efficiency	52 kg H_2/MWh
Green H_2 Price	\$5.5/kg (2023), \$3.3/kg (2030)

Table 6: Datacenter TCO Components

Datacenter Facility ($DeprPeriod_{DC}$=20 years [3], PUE=1.1 [30])	
Capex	\$10 million/MW [95]
Opex	\$0.48 million/(MW·year) [3]
IT Equipment ($DeprPeriod_{IT}$=6 years [77])	
Capex (CPU servers)	\$12.5 million/MW [77]
Capex (AI servers)	\$35 million/MW [77]
Opex	5% of Capex (\$/(MW·year)) [3]
Power Cost	
Electricity Price	\$80/MWh [19]

References

[1] 2018. Lancium. https://www.lancium.com. A startup company, building zero-carbon cloud computing resources..

[2] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. 2023. Carbon explorer: A holistic framework for designing carbon aware datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*. 118–132.

[3] Ranganathan Barroso, Holzle. 2018. *The Datacenter As A Computer: Designing Warehouse-Scale Machines, Third Edition* (third ed.). Morgan-Claypool.

[4] Alyssa Bersine. 2025. Reducing Data Center Peak Cooling Demand and Energy Costs With Underground Thermal Energy Storage. https://www.nrel.gov/news/program/2025/reducing-data-center-peak-cooling-demand-and-energy-costs-with-underground-thermal-energy-storage.html

[5] New York State Energy Planning Board. 2015. The Energy to Lead: 2015 New York State Energy Plan. http://energyplan.ny.gov/Plans/2015.aspx.

[6] Brian Carpenter and Scott Brim. 2002. *Middleboxes: Taxonomy and issues*. Technical Report.

[7] Vantage Data Centers. 2020. What a High-Density Data Center Means Today. https://vantage-dc.com/wp-content/uploads/2020/04/Vantage-what-a-high-density-data-center-means-today.pdf

[8] Niangjun Chen, Xiaoqi Ren, Shaolei Ren, and Adam Wierman. 2015. Greening multi-tenant data center demand response. *Performance Evaluation* 91 (2015), 229–254.

[9] Andrew A Chien, Chaojie Zhang, Liuzixuan Lin, and Varsha Rao. 2022. Beyond PUE: Flexible datacenters empowering the cloud to decarbonize. *USENIX Hot Carbon* (2022).

[10] Wesley Cole and Akash Karmakar. 2023. *Cost projections for utility-scale battery storage: 2023 update*. Technical Report. National Renewable Energy Lab.(NREL), Golden, CO (United States).

[11] California Public Utilities Commission. 2021. CPUC Drives California Toward a Clean Energy Future. https://www.cpuc.ca.gov/news-and-updates/all-news/cpuc-drives-california-toward-a-clean-energy-future-2024

[12] Ireland Commission for Regulation of Utilities. 2021. CRU Direction to the System Operators related to Data Centre grid connection processing. https://cruie-live-96ca64acab2247eca8a850a7e54b-5b34f62.divio-media.com/documents/CRU21124-CRU-Direction-to-the-System-Operators-related-to-Data-Centre-grid-connection-.pdf

[13] Gary Cook and Elizabeth Jardim. 2019. *Clicking Clean Virginia: The Dirty Energy Powering Data Center Alley*. Technical Report. Greenpeace. https://www.greenpeace.org/usa/reports/click-clean-virginia/.

[14] Microsoft Corporation. 2025. Microsft Environmental Sustainability Report 2025. https://www.microsoft.com/en-us/corporate-responsibility/sustainability/report/

[15] PG&E Corporation. 2024. 2024 Investor Update. https://s1.q4cdn.com/880135780/files/doc_presentations/2024/June/2024-Investor-Update-Presentation_Final.pdf

[16] Eli Cortez, Anand Bonde, Alexandre Muzio, Mark Russinovich, Marcus Fontoura, and Ricardo Bianchini. 2017. Resource Central: Understanding and Predicting Workloads for Improved Resource Management in Large Cloud Platforms. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17)*. 153–167. https://doi.org/10.1145/3132747.3132772

[17] Wei Deng, Fangming Liu, Hai Jin, Bo Li, and Dan Li. 2014. Harnessing renewable energy in cloud datacenters: opportunities and challenges. *iEEE Network* 28, 1 (2014), 48–55.

[18] PJM Resource Adequacy Planning Department. 2024. PJM Load Forecast Report. https://www.pjm.com/planning/resource-adequacy-planning/load-forecast-dev-process.

[19] Statista Research Department. 2024. Average retail electricity price for industrial consumers in the United States from 1970 to 2023. https://www.statista.com/statistics/190680/us-industrial-consumer-price-estimates-for-retail-electricity-since-1970/

[20] European Commission. Climate Action DG. 2019. *Going Climate-neutral by 2050: A Strategic Long-term Vision for a Prosperous, Modern, Competitive and Climate-neutral EU Economy*. Publications Office of the European Union.

[21] Hui Dou, Yong Qi, Wei Wei, and Houbing Song. 2017. Carbon-aware electricity cost minimization for sustainable data centers. *IEEE Transactions on Sustainable Computing* 2, 2 (2017), 211–223.

[22] William Driscoll. 2022. California law would target 90% renewable and zero-carbon electricity by 2035. https://pv-magazine-usa.com/2022/09/06/california-law-would-target-90-renewable-and-zero-carbon-electricity-by-2035/

[23] Iain Dunning, Joey Huchette, and Miles Lubin. 2017. JuMP: A modeling language for mathematical optimization. *SIAM review* 59, 2 (2017), 295–320.

[24] Virginia Electric and Power Company. 2023. 2023 Integrated Resource Plan. https://www.dominionenergy.com/about/our-company/irp#.

[25] Department for Energy Security, Net Zero, and Energy & Industrial Strategy Department for Business. 2021. Net Zero Strategy: Build Back Greener.

[26] Christopher Gannatti. 2024. This Is What AI Commitment Looks Like: $392 Billion and Rising. https://www.wisdomtree.com/investments/blog/2025/05/21/this-is-what-ai-commitment-looks-like-392-billion-and-rising

[27] Pacific Gas and Electric Company. 2024. PG&E Flex Connect Pilot. https://www.pge.com/assets/pge/docs/clean-energy/electric-vehicles/flexible-service-connection-pilot-overview.pdf

[28] Brian F Gerke, Sarah Josephine Smith, Samanvitha Murthy, Sunhee Baik, Shreya Agarwal, Peter Alstone, Aditya Khandekar, Cong Zhang, Richard E Brown, Jingjing Liu, and Mary Ann Piette. 2024. The California Demand Response Potential Study, Phase 4: Report on Shed and Shift Resources Through 2050.

[29] Íñigo Goiri, William Katsak, Kien Le, Thu D Nguyen, and Ricardo Bianchini. 2013. Parasol and greenswitch: Managing datacenters powered by renewable energy. In *ACM SIGARCH Computer Architecture News*. ACM, 51–64.

[30] Google. 2024. Google Data Centers Efficiency. https://www.google.com/intl/us-EN/about/datacenters/efficiency/

[31] Google. 2025. Google 2025 Environmental Report. https://sustainability.google/google-2025-environmental-report/

[32] Joseph Gorka, Noah Rhodes, and Line Roald. 2025. Electricityemissions. jl: A framework for the comparison of carbon intensity signals. In *Proceedings of the 16th ACM International Conference on Future and Sustainable Energy Systems*. 19–30.

[33] Alibaba Group. 2025. 2025 Alibaba Group Environmental, Social and Governance (ESG) Report. https://www.alibabagroup.com/en-US/esg

[34] EirGrid Group. 2024. Generation Capacity Statement 2023–2032. https://cms.eirgrid.ie/sites/default/files/publications/19035-EirGrid-Generation-Capacity-Statement-Combined-2023-V5-Jan-2024.pdf.

[35] Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual. https://www.gurobi.com

[36] Walid A Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. Carbonscaler: Leveraging cloud workload elasticity for optimizing carbon-efficiency. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 3 (2023), 1–28.

[37] Tim Horger. 2025. PJM Large Load Additions Workshop Presentation. https://www.pjm.com/-/media/DotCom/committees-groups/workshops/llaw/2025/20250509/20250509-item-02---large-load-additions-workshop---presentation.pdf

[38] Ethan Howland. 2024. FERC rejects interconnection pact for Talen-Amazon data center deal at nuclear plant. https://www.utilitydive.com/news/ferc-interconnection-isa-talen-amazon-data-center-susquehanna-exelon/731841/

[39] M Huneault and FD Galiana. 1991. A survey of the optimal power flow literature. *IEEE transactions on Power Systems* 6, 2 (1991), 762–770.

[40] McKinsey & Company Hydrogen Council. 2023. Hydrogen Insights December 2023. https://hydrogencouncil.com/wp-content/uploads/2023/12/Hydrogen-Insights-Dec-2023-Update.pdf

[41] International Energy Agency (IEA). 2024. Renewables 2024. https://www.iea.org/reports/renewables-2024/electricity

[42] International Energy Agency (IEA). 2025. Energy and AI. https://www.iea.org/reports/energy-and-ai

[43] Electric Power Research Institute. 2024. Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption. https://www.epri.com/research/products/3002028905

[44] Electric Power Research Institute. 2024. Utility Experiences and Trends Regarding Data Centers: 2024 Survey. https://www.epri.com/research/products/000000003002030643

[45] Electric Power Research Institute. 2025. DCFlex Initiative. https://dcflex.sf.epri.com/

[46] Ponemon Institute. 2016. Cost of Data Center Outages. https://www.vertiv.com/globalassets/documents/reports/2016-cost-of-data-center-outages-11-11_51190_1.pdf

[47] Sean James. 2024. The Good, Bad, and Ugly of Directly Powering Datacenters with Nuclear Power. https://www.linkedin.com/pulse/good-bad-ugly-powering-datacenters-nuclear-power-sean-james-ogs2c/?trackingId=zZsZm1K2FmLPsk8DPr6fWA%3D%3D

[48] Peter Judge. 2022. Dominion Energy admits it can't meet data center power demands in Virginia. https://www.datacenterdynamics.com/en/news/dominion-energy-admits-it-cant-meet-data-center-power-demands-in-virginia/

[49] Peter Judge. 2022. EirGrid pulls plug on 30 Irish data center projects. https://www.datacenterdynamics.com/en/news/eirgrid-pulls-plug-on-30-irish-data-center-projects/

[50] Peter Judge. 2022. Singapore lifts data center moratorium - but sets conditions. https://www.datacenterdynamics.com/en/news/singapore-lifts-data-center-moratorium-but-sets-conditions/

[51] Peter Judge. 2023. Microsoft granted permission to run its Dublin data center on gas. https://www.datacenterdynamics.com/en/news/microsoft-granted-permission-to-run-its-dublin-data-center-on-gas/

[52] Kibaek Kim, Fan Yang, Victor Zavala, and Andrew A. Chien. 2016. Data Centers as Dispatchable Loads to Harness Stranded Power. *IEEE Transactions on*

Sustainable Energy (2016). DOI 10.1109/TSTE.2016.2593607.

[53] National Renewable Energy Laboratory. 2021. Life Cycle Greenhouse Gas Emissions from Electricity Generation: Update. https://docs.nrel.gov/docs/fy21osti/80580.pdf

[54] Lazard. 2025. Lazard's 2025 Levelized Cost Of Energy+. https://www.lazard.com/media/uounhon4/lazards-lcoeplus-june-2025.pdf

[55] Tan N Le, Zhenhua Liu, Yuan Chen, and Cullen Bash. 2016. Joint capacity planning and operational management for sustainable data centers and demand response. In *Proceedings of the Seventh International Conference on Future Energy Systems*. 1–12.

[56] Yuzhuo Li, Mariam Mughees, Yize Chen, and Yunwei Ryan Li. 2024. The unseen AI disruptions for power grids: LLM-induced transients. *arXiv preprint arXiv:2409.11416* (2024).

[57] Liuzixuan Lin and Andrew A Chien. 2023. Adapting Datacenter Capacity for Greener Datacenters and Grid. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 200–213.

[58] Liuzixuan Lin, Rajini Wijayawardana, Varsha Rao, Hai Nguyen, Emmanuel Wedan GNIBGA, and Andrew A Chien. 2024. Exploding AI Power Use: an Opportunity to Rethink Grid Planning and Management. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 434–441.

[59] Liuzixuan Lin, Victor M Zavala, and Andrew A Chien. 2021. Evaluating Coupling Models for Cloud Datacenters and Power Grids. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*. 171–184.

[60] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. 2012. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking* 21, 5 (2012), 1378–1391.

[61] Julia Lindberg, Yasmine Abdennadher, Jiaqi Chen, Bernard C Lesieutre, and Line Roald. 2021. A Guide to Reducing Carbon Emissions through Data Center Geographical Load Shifting. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*. 430–436.

[62] Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. 2012. Renewable and cooling aware workload management for sustainable data centers. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*. 175–186.

[63] Zhenhua Liu, Adam Wierman, Yuan Chen, Benjamin Razon, and Niangjun Chen. 2013. Data center demand response: Avoiding the coincident peak via workload shifting and local generation. *Performance Evaluation* 70, 10 (2013), 770–791.

[64] Jianying Luo, Lei Rao, and Xue Liu. 2013. Temporal load balancing with service delay guarantees for data center energy cost optimization. *IEEE Transactions on Parallel and Distributed Systems* 25, 3 (2013), 775–784.

[65] Kari Lydersen. 2024. Data centers offer energy peril and promise, with the Midwest increasingly in the crosshairs. https://www.minnpost.com/other-nonprofit-media/2024/06/microsoft-data-centers-energy-peril-promise-midwest-energy-water-concerns/

[66] Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. Recalibrating global data center energy-use estimates. *Science* 367, 6481 (2020), 984–986. https://doi.org/10.1126/science.aba3758 arXiv:https://science.sciencemag.org/content/367/6481/984.full.pdf

[67] Varun Mehra and Raiden Hasegawa. 2023. Supporting power grids with demand response at Google data centers. https://cloud.google.com/blog/products/infrastructure/using-demand-response-to-reduce-data-center-power-consumption

[68] Ali Menati, Kiyeob Lee, and Le Xie. 2023. Modeling and analysis of utilizing cryptocurrency mining for demand flexibility in electric energy systems: A synthetic texas grid case study. *IEEE Transactions on Energy Markets, Policy and Regulation* (2023).

[69] Meta. 2025. 2024 Sustainability Report. https://sustainability.atmeta.com/2024-sustainability-report/

[70] Samuel Newell, Michael Hagerty, Johannes Pfeifenberger, Bin Zhou, Travis Carless, Rohan Janakiraman, Sang Gang, Patrick Daou, and Joshua Junge. 2022. *PJM CONE 2026/2027 Report*. Technical Report. PJM Interconnection.

[71] Tyler Norris, Timothy Profeta, Dalia Patino-Echeverri, and Adam Cowie-Haskell. 2025. *Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems*. Technical Report. Nicholas Institute for Energy, Environment & Sustainability, Duke University.

[72] Virginia Department of Environmental Quality. 2024. Renewable Energy. https://www.deq.virginia.gov/our-programs/renewable-energy.

[73] Electric Reliability Council of Texas. [n. d.]. Large Flexible Load Task Force. https://www.ercot.com/committees/tac/lfltf.

[74] Janne Paananen and Ehsan Nasr. 2021. Grid-interactive data centers: enabling decarbonization and system stability. https://www.eaton.com/content/dam/eaton/markets/data-center/eaton-microsoft-grid-interactive-whitepaper-wp153031en.pdf

[75] Anthony Papavasiliou and Shmuel S Oren. 2013. Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. *Operations Research* 61, 3 (2013), 578–592.

[76] Anjaly Parayil, Jue Zhang, Xiaoting Qin, Íñigo Goiri, Lexiang Huang, Timothy Zhu, and Chetan Bansal. 2024. Towards Cloud Efficiency with Large-scale Workload Characterization. *arXiv preprint arXiv:2405.07250* (2024).

[77] Dylan Patel and Daniel Nishball. 2023. GPU Cloud Economics Explained – The Hidden Truth. https://www.semianalysis.com/p/gpu-cloud-economics-explained-the

[78] Ana Radovanović, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyue Xiao, Maya Haridasan, Patrick Hung, Nick Care, et al. 2022. Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems* 38, 2 (2022), 1270–1280.

[79] Chuangang Ren, Di Wang, Bhuvan Urgaonkar, and Anand Sivasubramaniam. 2012. Carbon-aware energy capacity planning for datacenters. In *2012 IEEE 20th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*. IEEE, 391–400.

[80] Goldman Sachs Research. 2024. AI/data centers' global power surge and the Sustainability impact. https://www.goldmansachs.com/images/migrated/insights/pages/gs-research/gs-sustain-generational-growth-ai-data-centers-global-power-surge-and-the-sustainability-impact/sustain-data-center-redaction.pdf

[81] Reuters. 2024. Microsoft may pay Constellation premium in Three Mile Island power agreement, Jefferies says (Sept 23). https://www.reuters.com/markets/deals/microsoft-may-pay-constellation-premium-three-mile-island-power-agreement-2024-09-23/

[82] Ian Schneider and Taylor Mattia. 2024. Carbon accounting in the Cloud: a methodology for allocating emissions across data center users. *arXiv preprint arXiv:2406.09645* (2024).

[83] Arman Shehabi, Sarah Josephine Smith, Alex Hubbard, Alexander Newkirk, Nuoa Lei, Md AbuBakar Siddik, Billie Holecek, Jonathan G Koomey, Eric R Masanet, and Dale A Sartor. 2024. *2024 United States Data Center Energy Usage Report*. Technical Report. Lawrence Berkeley National Laboratory.

[84] Jie Shi, Zuzhao Ye, H Oliver Gao, and Nanpeng Yu. 2022. Lyapunov optimization in online battery energy storage system control for commercial buildings. *IEEE Transactions on Smart Grid* 14, 1 (2022), 328–340.

[85] Yuanyuan Shi, Bolun Xu, Baosen Zhang, and Di Wang. 2016. Leveraging energy storage to optimize data center electricity cost in emerging power markets. In *Proceedings of the Seventh International Conference on Future Energy Systems*. 1–13.

[86] Patrick Sisson. 2024. A.I. Frenzy Complicates Efforts to Keep Power-Hungry Data Sites Green. https://www.nytimes.com/2024/02/29/business/artificial-intelligence-data-centers-green-power.html

[87] Abel Souza, Noman Bashir, Jorge Murillo, Walid Hanafy, Qianlin Liang, David Irwin, and Prashant Shenoy. 2023. Ecovisor: A virtual energy system for carbon-efficient applications. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*. 252–265.

[88] Alpha Spread. 2025. ROIC Across Alphabet Inc and Competitors. https://www.alphaspread.com/security/nasdaq/googl/profitability/ratio/return-on-invested-capital

[89] Jovan Stojkovic, Pulkit A Misra, Íñigo Goiri, Sam Whitlock, Esha Choukse, Mayukh Das, Chetan Bansal, Jason Lee, Zoey Sun, Haoran Qiu, et al. 2024. SmartOClock: Workload-and risk-aware overclocking in the cloud. In *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 437–451.

[90] Thanathorn Sukprasert, Abel Souza, Noman Bashir, David Irwin, and Prashant Shenoy. 2024. On the Limitations of Carbon-Aware Temporal and Spatial Workload Shifting in the Cloud. In *Proceedings of the Nineteenth European Conference on Computer Systems*. 924–941.

[91] Jinghan Sun, Zibo Gong, Anup Agarwal, Shadi Noghabi, Ranveer Chandra, Marc Snir, and Jian Huang. 2024. Exploring the Efficiency of Renewable Energy-based Modular Data Centers at Scale. In *Proceedings of the 2024 ACM Symposium on Cloud Computing*. 552–569.

[92] Darren Sweeney and Jared Anderson. 2023. Power of AI: Surging datacenter load has Dominion bracing for AI's added demand. https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/power-of-ai-surging-datacenter-load-has-dominion-bracing-for-ai-s-added-demand-77792886

[93] Dan Swinhoe. 2024. Amazon invests in nuclear SMR company, signs several SMR deals to power data centers. https://www.datacenterdynamics.com/en/news/amazon-signs-deals-to-invest-in-nuclear-smrs-to-power-data-centers/

[94] Muhammad Tirmazi, Adam Barker, Nan Deng, Md E. Haque, Zhijing Gene Qin, Steven Hand, Mor Harchol-Balter, and John Wilkes. 2020. Borg: The next Generation. In *Proceedings of the Fifteenth European Conference on Computer Systems* (Heraklion, Greece) *(EuroSys '20)*. Association for Computing Machinery, New York, NY, USA, Article 30, 14 pages. https://doi.org/10.1145/3342195.3387517

[95] Turner & Townsend. 2023. Data Centre Cost Index 2023. https://reports.turnerandtownsend.com/dcci-2023/

[96] Rahul Urgaonkar, Bhuvan Urgaonkar, Michael J Neely, and Anand Sivasubramaniam. 2011. Optimal power cost management using stored energy in data centers. In *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*. 221–232.

[97] Jonathan Vanian and Kif Leswing. 2023. ChatGPT and generative AI are booming, but the costs can be extraordinary. https://www.cnbc.com/2023/03/13/chatgpt-and-generative-ai-are-booming-but-at-a-very-expensive-price.html

[98] Vilayanur Viswanathan, Kendall Mongird, Ryan Franks, Xiaolin Li, Vincent Sprenkle, and R Baxter. 2022. 2022 grid energy storage technology cost and performance assessment. *Energy* 2022 (2022).

[99] Di Wang, Chuangang Ren, Anand Sivasubramaniam, Bhuvan Urgaonkar, and Hosam Fathy. 2012. Energy storage in datacenters: what, where, and how much?. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*. 187–198.

[100] Ruihang Wang, Zhiwei Cao, Xin Zhou, Yonggang Wen, and Rui Tan. 2024. Green data center cooling control via physics-guided safe reinforcement learning. *ACM Transactions on Cyber-Physical Systems* 8, 2 (2024), 1–26.

[101] Philipp Wiesner, Ilja Behnke, Dominik Scheinert, Kordian Gontarska, and Lauritz Thamsen. 2021. Let's wait awhile: how temporal workload shifting can reduce carbon emissions in the cloud. In *Proceedings of the 22nd International Middleware Conference*. 260–272.

[102] Qiang Wu, Qingyuan Deng, Lakshmi Ganesh, Chang-Hong Hsu, Yun Jin, Sanjeev Kumar, Bin Li, Justin Meza, and Yee Jiun Song. 2016. Dynamo: Facebook's data center-wide power management system. *ACM SIGARCH Computer Architecture News* 44, 3 (2016), 469–480.

[103] Song Wu, Yang Chen, Xinhou Wang, Hai Jin, Fangming Liu, Haibao Chen, and Chuxiong Yan. 2018. Precise power capping for latency-sensitive applications in datacenter. *IEEE Transactions on Sustainable Computing* 6, 3 (2018), 469–480.

[104] Zhaoyuan Wu, Lin Chen, Jianxiao Wang, Ming Zhou, Gengyin Li, and Qing Xia. 2023. Incentivizing the spatiotemporal flexibility of data centers toward power system coordination. *IEEE Transactions on Network Science and Engineering* (2023).

[105] Jiali Xing, Bilge Acun, Aditya Sundarrajan, David Brooks, Manoj Chakkaravarthy, Nikky Avila, Carole-Jean Wu, and Benjamin C Lee. 2023. Carbon Responder: Coordinating Demand Response for the Datacenter Fleet. *arXiv preprint arXiv:2311.08589* (2023).

[106] Baosen Zhang and Ehsan Nasr Senior. 2021. Distributed redundant integration of data center battery storage with the grid for regulation services. In *2021 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 1–5.

[107] Chaojie Zhang. 2023. *Eliminating the Capacity Variation Penalty for Cloud Resource Management*. Ph. D. Dissertation. The University of Chicago.

[108] Chaojie Zhang, Alok Gautam Kumbhare, Ioannis Manousakis, Deli Zhang, Pulkit A Misra, Rod Assis, Kyle Woolcock, Nithish Mahalingam, Brijesh Warrier, David Gauthier, et al. 2021. Flex: High-availability datacenters with zero reserved power. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 319–332.

[109] Linquan Zhang, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2015. A truthful incentive mechanism for emergency demand response in colocation data centers. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2632–2640.

[110] Weiqi Zhang, Line A Roald, Andrew A Chien, John R Birge, and Victor M Zavala. 2020. Flexibility from networks of data centers: A market clearing formulation with virtual links. *Electric Power Systems Research* 189 (2020), 106723.

[111] Dongyu Zheng, Lei Liu, Guoming Tang, Yi Wang, and Weichao Li. 2023. Power Demand Reshaping Using Energy Storage for Distributed Edge Clouds. *IEEE Transactions on Parallel and Distributed Systems* (2023).

[112] Zhi Zhou, Fangming Liu, Shutong Chen, and Zongpeng Li. 2018. A truthful and efficient incentive mechanism for demand response in green datacenters. *IEEE Transactions on Parallel and Distributed Systems* 31, 1 (2018), 1–15.