# Power Management Optimization for Data Centers: A Power Supply Perspective

Huikang Huang, Weiwei Lin, *Senior Member, IEEE*, Jianpeng Lin, and Keqin Li, *Fellow, IEEE*

*(Survey-Tutorial Paper)*

***Abstract*—With the escalating demand for cloud computing services, data centers (DCs) encounter formidable challenges extending beyond capital investment needs to accommodate increasing computational demands and routine infrastructure maintenance. These challenges include substantial electricity costs due to high energy consumption and the environmental issues caused by high carbon emissions. To reduce costs and mitigate environmental impacts, modern DCs not only use energy-efficient technologies to improve the efficiency of common IT and cooling systems, but also actively optimize the direct, indirect and environmental costs of the power supply side, posing significant challenges for DCs power management. Therefore, this paper presents a comprehensive survey of cost-aware optimization from the power supply perspective. First, it reviews the structures and key metrics of the power supply system, modeling methods and supporting techniques for main power and IT system components, establishing a foundation for optimization. Second, traditional (brown) and green energy sources are categorized to survey and compare existing critical works, analyzing the application of power management methods to tackle cost-related challenges. Finally, future research trends in the power supply perspective for DCs are discussed. This survey aims to provide recommendations for power supply side cost optimization to further advance the sustainable development of DCs.**

***Index Terms*—Data center, power management, green energy, power system, cost optimization.**

## I. Introduction

With the prevalence of cloud computing services around the world, and driven by the rapid development of computationally intensive technologies such as artificial intelligence, distributed computing, and blockchain, there is an inevitable need to increase the computing power of data centers (DCs). As high energy-consuming structures, DCs typically consume 10 to 100 times more electricity per floor compared to other standard buildings [1]. This necessitates substantial initial capital investment to deploy the necessary infrastructure to support increased computing power [2]. Additionally, the high energy consumption of DCs results in an annual electricity usage exceeding 200 TWh, accounting for more than 1% of global electricity consumption [3]. It is estimated that by 2030, the electricity usage of DCs will rise to 8% of the global total [4]. The high energy consumption of DCs not only places significant pressure on power grids but also has severe negative environmental impacts. The primary reason is that current electricity generation is still predominantly driven by fossil fuels. Carbon emissions generated by DCs alone exceed 0.3% of the total global carbon emissions and have been on the rise in recent years [4]. Therefore, meeting the demands of cloud computing services while addressing the associated high costs, high energy consumption, and high carbon emissions poses a significant challenge for DC operators. They must balance the total cost of ownership (TCO) and environmental costs.

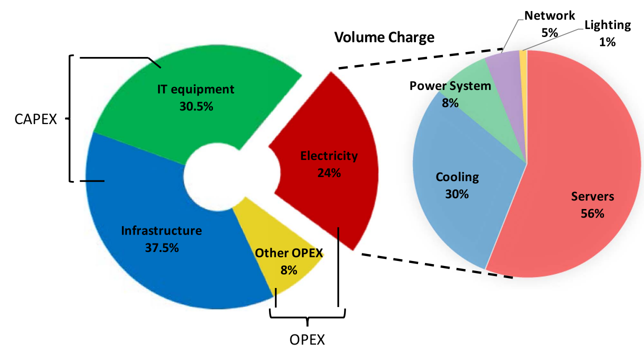


Fig. 1. Amortized monthly cost and energy consumption ratio of the DC.

The TCO of DCs is composed of capital expenditure (CAPEX) and operating expenses (OPEX). The CAPEX for DCs includes the cost of computing equipment and supporting infrastructure, as well as the associated depreciation expenses. The investment in infrastructure accounts for a high percentage of CAPEX, which varies depending on the size and class of the DC [5]. In a typical DC's monthly cost amortization [2], the expenditure on supporting infrastructure can exceed the OPEX of the DC, as illustrated on the left side of Fig. 1. Moreover, DCs frequently experience resource underutilization [6], as substantial capital investments in infrastructure are often made to

Received 29 July 2024; revised 9 December 2024; accepted 9 February 2025. Date of publication 18 February 2025; date of current version 8 August 2025. This work was supported in part by the Guangdong Major Project of Basic and Applied Basic Research under Grant 2019B030302002, in part by Guangzhou Development Zone Science and Technology Project under Grant 2023GH02 and in part by the Major Key Project of PCL, China under Grant PCL2023A09. Recommended for acceptance by Chao Li. *(Corresponding author: Weiwei Lin.)*

Huikang Huang and Jianpeng Lin are with the School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China (e-mail: huikanghuang0321@gmail.com; csjianpenglin@mail.scut.edu.cn).

Weiwei Lin is with the School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China, and also with Pengcheng Laboratory, Shenzhen 518066, China (e-mail: linww@scut.edu.cn).

Keqin Li is with the Department of Computer Science, State University of New York, New Paltz, NY 12561 USA (e-mail: lik@newpaltz.edu).

Digital Object Identifier 10.1109/TSUSC.2025.3542779

accommodate peak demand. In light of the high infrastructure costs and the underutilization of resources, DCs frequently mitigate amortized infrastructure costs by deploying additional servers within existing facilities. Regarding OPEX, one of the primary contributors is high energy consumption, which results in significant electricity bills. In a typical DC, 56% of the energy consumption is generated by the IT system, 30% by the cooling system, and 8% by the power system, and the rest is generated by the network, lighting, and other equipment [7], as shown on the right side of the Fig. 1. The electricity bill consists not only of basic or volume-based charges linked to electricity usage, but also peak demand charges [8], incurred when the DC's peak power consumption exceeds pre-agreed contractual limits. Thus, in order to reduce TCO, it is critical to optimize infrastructure utilization during operations, implement energy efficiency measures, and adopt strategies for managing peak demand. Additionally, DCs can further lower electricity costs by engaging in demand response (DR) or regulation services (RS) within electricity markets [9].

The environmental cost of DC operations is primarily determined by the carbon emissions resulting from energy consumption, as well as the mix of green and brown energy sources used. Achieving true Green IT, as defined by Greenpeace [10] (Green IT = Energy Efficiency + Renewable Energy), requires not only the adoption of advanced energy-efficient technologies but also the integration of renewable energy sources to power DCs. As the world's largest carbon emissions trading market [11], China has implemented a carbon emission quota trading system to incentivize DCs to utilize cleaner energy and develop more advanced energy-saving technologies. For instance, major Chinese companies such as Alibaba [12] and Tencent [13] have committed, in their carbon neutrality reports, to achieving carbon neutrality across their operations and supply chains by 2030. Similarly, in the United States, companies like Google [14] and Facebook [15] have made similar pledges, committing to carbon-free energy operations and net-zero carbon emissions by 2030. Thus, to reduce the environmental cost of sustainable development, the integration of green energy in DCs has become an undeniable matter of fact. However, the inherent instability and intermittency of green energy sources pose significant challenges for DC power management technologies and strategies.

However, current research on optimizing power management in DCs predominantly focuses on IT and cooling systems, due to their high energy consumption. There are relatively few review studies on the power supply side of DCs and all of them have their own focus. For example, Malla et al. [16] focus on secure power over-subscription, discussing how to improve the utilization of power infrastructure from both the operator and tenant perspectives. Kong et al. [17] survey power management in green energy-aware DCs, covering workload scheduling, virtual machine (VM) management, and capacity planning. Cao et al. [18] review optimization research from the perspective of carbon neutrality, combining carbon market regulation and carbon reduction technologies. Compared to these studies, our focus is to provide a more comprehensive and systematic investigation of cost awareness (TCO and environmental costs) from a power supply perspective, emphasizing the application of power management techniques and approaches to DC cost optimization. We review the relevant elements of problem modeling in terms of cost-aware power management optimization problems and discuss how to solve problems with different cost awareness. Fig. 2 shows the overview of this survey, and the main contributions are as follows. (1) We specifically analyze DCs with different power supply structures, profile the power models of different power supply hierarchies and the corresponding power management techniques. Additionally, we summarize key evaluation metrics of the power supply side to provide a foundation for power management optimization. (2) We systematically review the optimization studies of cost-aware DCs, offering a comprehensive analysis and comparison of critical research across three key aspects: direct and indirect costs of traditional energy-based DCs, and the environmental costs associated with green energy DCs. (3) Finally, from the perspective of the power supply side, we analyze and identify the research trends in DC power management optimization.

The paper is organized as follows: Section II presents the power supply structure of DC power systems, key metrics on the power supply side, and the power models and management technologies critical to both IT and power systems. These elements are essential for problem modeling. Section III addresses the optimization of power management in traditional energy systems, while Section IV explores the optimization of power management for green energy. Finally, from the perspective of the power supply side, we highlight the future research trends in power management optimization. It should be emphasized that this paper does not consider HPC and edge scenarios, as these are beyond our research scope.

## II. Power Management Optimization: Problem Modeling Elements

### A. Data Center Power Supply System Structure

The power supply of a traditional DC is brought from the utility power from the power supply company, but it cannot guarantee an uninterrupted power supply. To ensure high availability of the power system, DCs typically source utility power from two different routes. In addition, the DC will be equipped with a Diesel Generator (DG) as backup power in case both utility power roads are unavailable and can be switched to DG via an Automatic Transfer Switching (ATS). However, since it takes about 20-30 seconds to start the DG [19], depends on Energy Storage Devices (ESDs), i.e., a short power supply from the Uninterruptible Power Supply (UPS) in the DC is required to seamlessly connect the power load from utility power to DG. Fig. 3 shows the power supply structure of DCs, which can be roughly classified into centralized power supply structure, distributed power supply structure, and green energy power supply structure.

*1) Centralized Power Supply Structure:* At present, DCs still widely use centralized power supply structure, as shown in Fig. 3(a). The capacity of the UPS is large, and the power transmission process requires double conversion, which generates high power loss. When providing power to IT systems, further power loss is generated by conversion nodes such as

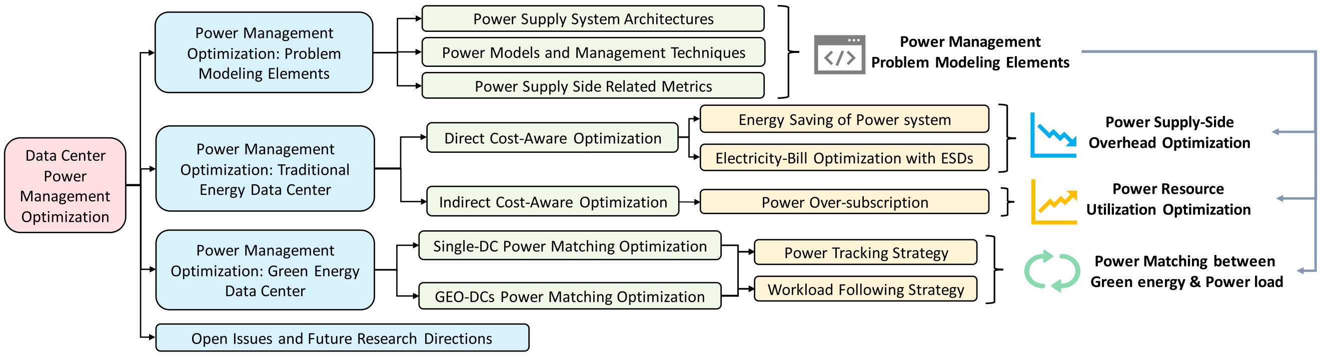


Fig. 2. An overview of power management optimization from the power-supply side.

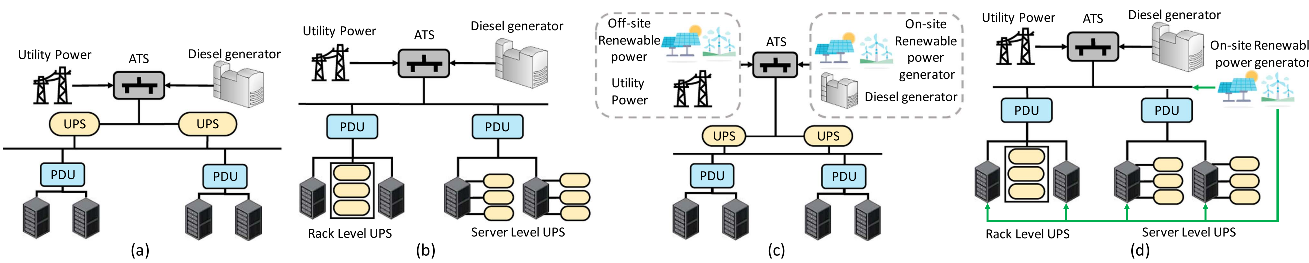


Fig. 3. Architectures of the Power supply system. (a) Centralized power supply structure. (b) Distributed power supply structure. (c) Centralized green energy power supply structure. (d) Distributed green energy power supply structure.

TABLE I
OVERVIEW OF TIER CLASSIFICATION REQUIREMENT [21]

| Items | Tier I | Tier II | Tier III | Tier IV |
|---|---|---|---|---|
| Utility Supply (Connection point) | Single Point | Single Point | Single Point | Dual |
| Backup Generator | Optional | N | N+1 | 2N |
| Backup system (UPS) | N | N+1 | N+1 | 2N |
| Maintenance | outage for maintenance | outage for maintenance | concurrently maintainable | fault tolerant |
| Availability | 0.999947 | 0.9999512 | 0.9999791 | 0.9999976 |

Power Distribution Units (PDU) or Cabinets Power Distribution Units (CDU). In addition, considering the possibility of a single point of failure, according to the DC tier standards classified by UPTIME [20], Tier IV DCs need to be equipped with 2 N UPS to ensure that the two sets of power supply systems are physically isolated, to maximize the availability of the power system in the DC. More specifically, the DC Tier is divided into four levels for evaluating the power availability of DC facilities, as shown in Table I.

It is worth noting that the AC 2 N UPS power modules used in fault-tolerant DCs require a high upfront investment and are less power efficient due to the double conversion. Therefore, new DCs are considering using direct utility power combined with high-voltage direct current (HVDC) for power supply. HVDC eliminates the power loss caused by double conversion [22], which enhances the power supply efficiency. For more information on the power efficiency performance of different power supply structures, refer to [23].

*2) Distributed Power Supply Structure:* The biggest advantage of a distributed power supply structure is its ability to effectively avoid single points of failure in the power transmission process. Earlier, Google proposed an extremely distributed power supply structure [24]. They deployed the UPS on servers, eliminating the double conversion process of centralized UPS and thereby reducing power loss during transmission. However, some of the power distribution modules need to be customized and are not friendly to general DCs. Therefore, Facebook in 2011, launched the Open Compute Project (OCP) [25], to promote collaboration in the design and technical improvement of DC standards. This initiative has further advanced the development of distributed power supply structures. Currently, there are two representative types of distributed power supply structures:

- *Rack-level UPS power supply structure* (the left side of Fig. 3(b)). The rack-level UPS power supply structure of Facebook [25] moves UPS to a special battery cabinet close to the server rack, allowing the stored power to be closer to the IT systems, which helps reduce power loss.
- *Server-level UPS power supply structure* (the right side of Fig. 3(b)). Compared to the rack-level UPS, Microsoft's LES structure [26] moves UPS further backward to the 380 V DC bus of the Power Supply Unit (PSU) inside the server. Google's early proposed extreme UPS [24], which is directly connected to the output of the PSU.

Although the battery capacity and power-sharing domain of distributed UPS are smaller compared to centralized UPS, their flexible deployment characteristics allow for more efficient interaction with IT systems, offering greater optimization potential.

*3) Green Energy Power Supply Structure:* DCs will not only purchase green energy from power supply companies, but some

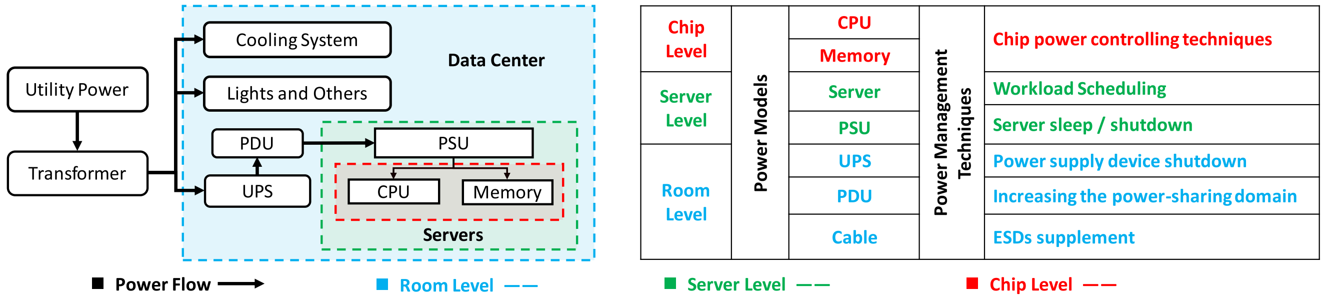


Fig. 4. Data center power consumption components and power management techniques.

DCs will also deploy on-site green energy generation equipment for themselves. We classify the integration of green energy into the power system and its operation modes into the following two types:

- *Centralized green energy power supply structure* [27] (as shown in Fig. 3(c)). Available in both off-site and on-site green energy, the difference is that the off-site green energy is similar to the utility power, which is procured from the power supply companies. While the on-site green energy is similar to the green energy generator located within the DC.
- *Distributed green energy power supply structure* [28] (as shown in Fig. 3(d)). In this structure, DCs integrate on-site green energy, which is stabilized and rectified before being directly connected to the racks or PDUs. However, building such a distributed green energy power supply system is more costly than a centralized green energy power supply system [28].

Currently, constructing a DC solely powered by green energy remains impractical due to the intermittency of green energy sources.

## B. *Power Models and Management Techniques*

To further illustrate the relationship between IT systems and power systems in DCs, we draw a concise architecture as shown in Fig. 4. In this section, we analyze common power models and power management techniques from the chip level (red box), server level (green box), and room level (blue box).

- *Chip Level:* In a typical server, the main components are the CPU, memory, disk, and NIC [29]. The CPU and memory contribute significantly to energy consumption and are equipped with power control technologies to dynamically adjust power usage. Thus, at the chip level, the focus is primarily on the power models of the CPU and memory.
- *Server Level:* Power modeling of servers has been an ongoing exploration in academia and industry. Due to the interaction of internal components, different workloads, and human interference, it is difficult to model accurately in complex states.
- *Room Level:* With a Data Center Infrastructure Management system, such as NetEco [30], operators can monitor critical power equipment, load levels, and more. The white paper [31] identifies UPS, PDU, and cables as the primary sources of power loss at the room level, with losses categorized into three types: (i) load-independent losses, (ii) losses linearly related to load, and (iii) losses proportional to the square of the current.

As our focus is on the power supply side for providing stable power to IT equipment, cooling systems and other modules are not considered. The following subsections introduce the power model and power management technologies.

*1) Power Models:*

- *CPU power model:* In general, CPU is the main energy consuming component of a server, especially dealing with computationally intensive workloads. A general CPU power model was proposed by Shin et al. [32], who modeled the CPU power as:

$$P_{cpu} = P_d + P_s + P_0, \tag{1}$$

where $P_d$, $P_s$ and $P_0$ denote dynamic power, static power and fixed power, respectively. Fixed power $P_0$ related to CPU type. Static power $P_s$ is the amount related to the chip temperature. Dynamic power $P_d$ can be expressed as:

$$P_d = f_{clock} C_{int} V^2, \tag{2}$$

where $f_{clock}$, $C_{int}$, $V$ denote clock frequency, physical capacitance, and the CPU voltage, respectively. It is worth noting that $P_d$ will be limited by dynamic voltage frequency scaling (DVFS) [33].

- *Memory power model:* Memory is the second largest energy-consuming module of servers [34]. An intuitive idea for modeling memory power is using the frequency of memory accesses. Patricia et al. [35] express the power model of memory as a relationship between the memory operating temperature $T_{mem}$ and the memory access frequency $f_{mem}(k)$ when the server uses different DVFS modes $k$:

$$P_{mem}(k) = a_1 T_{mem} + a_2 T_{mem}^2 + a_3 f_{mem}(k), \tag{3}$$

where $a_i$ are obtained from the model fitting.

- *Server power model:* Inspired by the DVFS technique, the server power model is considered to have a cubic function with CPU frequency, Mootaz et al. [36] used a function model based on frequency $f$ to estimate server power

$P_{server}$:

$$P_{server} = c_0 + c_1 f^3, \tag{4}$$

where $c_0$ is a constant that does not vary with the server load level and $c_1$ is a constant related to the CPU capacitance and voltage. The values of $c_0$ and $c_1$ vary depending on the server. Alternatively, in a more common server power model, Fan et al. [37] considered that CPU utilization has a strong correlation with server power, they constructed $P_{server}$ as a linear model with CPU utilization:

$$P_{server} = P_{idle} + (P_{\max} - P_{idle})\, u_{cpu}, \tag{5}$$

where $P_{idle}$ is the idle power of the server, $P_{\max}$ is the power of the server at full load, and $u_{cpu}$ denotes the CPU utilization. This linear model implies that $P_{server}$ varies linearly with $u_{cpu}$. Considering that this server power model is overly simplified, they subsequently propose a nonlinear model that fits $\gamma$ by minimizing the squared error of the model on the training data:

$$P_{server} = P_{idle} + (P_{\max} - P_{idle})\left(2u_{cpu} - u_{cpu}^{\gamma}\right). \tag{6}$$

- *PSU power model:* Although many server power models [38] have been proposed subsequently, they do not mention the PSU. As the power supply module for servers, the loss of the PSU is considered to be related to the size of the load. 80PLUS [39] shows the power supply conversion efficiency of different classes of PSUs at different load levels. Intel models the loss rate $\eta_{psu\_loss}$ and the load level $\eta_{psu\_load}$ of PSUs as a quadratic relationship [40]:

$$\eta_{psu\_loss} = a_0 + a_1 \eta_{psu\_load} + a_2 \eta_{psu\_load}^2, \tag{7}$$

where $a_i$ is selected according to the different types of PSUs. In Tier IV DCs, servers are equipped with dual PSUs to meet fault tolerance requirements. It means that the PSU is often at a lower load level when the wider high-efficiency range of the PSU is more important than the high efficiency at full load.

- *UPS power model:* In traditional DCs, PDUs for voltage conversion and diverting are deployed between the UPS and the rack, facilitating management and maintenance. Pelley et al. [7] modeled the relationship between UPS losses and PDU loads $P_{PDU}$ as:

$$P_{ups\_loss} = P_{ups\_idle} + \pi_{ups} * \left(\sum_{M} P_{pdu}\right), \tag{8}$$

where $P_{ups\_idle}$ denotes the idle power of the UPS, $\pi_{ups}$ is needed to be fitted as a parameter, and $M$ denotes the number of PDUs powered by the UPS. In [41], the authors point out that the idle loss of the UPS accounts for more than 40% of the total loss of the UPS.

- *PDU power model:* The PDU can take the three-phase power from the UPS and output it as multiple sets of single-phase power suitable for the server. The is some energy loss during the conversion process. Pelley et al. [7] suggest that the power model between the PDU and the server can be modeled as:

$$P_{pdu\_loss} = P_{pdu\_idle} + \pi_{pdu} * \left(\sum_{N} P_{srv}\right)^2, \tag{9}$$

where $P_{pdu\_idle}$ denotes the idle power consumption of the PDU, and $\pi_{pdu}$ is a parameter need to be fitted. Intel models the PDU loss as a quadratic function between load level $\eta_{pdu\_load}$ and loss rate $\eta_{pdu\_loss}$ [40]:

$$\eta_{pdu\_loss} = 0.0026 - 0.0054\eta_{pdu\_load} + 0.0343\eta_{pdu\_load}^2. \tag{10}$$

Meanwhile, PDUs are capable of supplying power directly to the rack, or they can be further divided by a remote power panel (RPP). The power enters the rack and is supplied to the server by the CDU.

- *Cable power model:* The loss at the room level cannot be ignored because of cable loss. In a power system, all devices need to be connected through cables. Ahmed et al. [21] pointed out that the cable losses between PSU and server do not need to be considered due to the short distance. The remaining part of the cable loss is calculated from the UPS current and cable resistance as shown in the following equation:

$$P_{Cable}^{Loss} = \left(\frac{P_{rack} + P_{PSU}^{Loss} + P_{PDU}^{Loss}}{V_{norm} * PF}\right)^2 * R_{Cable}, \tag{11}$$

where $P_{Rack}$ is the power consisting of all servers in the rack, $P_{PSU}^{Loss}$ and $P_{PDU}^{Loss}$ is the power loss of PSU and PDU, respectively, $V_{norm}$ is the nominal voltage, PF is the power factor, and $R_{Cable}$ is the resistance of the cable, usually provided by the manufacturer.

- *ESDs lifetime degradation model:* Battery life is indicated by State-of-Health (SoH), and replacement is required when $SoH$ reaches $SoH_{dead}$. Manufacturers typically set $SoH_{dead}$ to 80% of the rated capacity $C_R$. The $SoH$ decreases gradually with increased battery cycles and depth-of-discharge (DoD) [42]:

$$SoH = SoH - \frac{100 - SoH_{dead}}{Cycles_{DoDfinal}} * \frac{C_R}{C_{eff}}, \tag{12}$$

where $Cycles_{DoDfinal}$ is the number of $DoD$, which is determined by the battery type and the $DoDfinal$ of each discharge. Specifically, during the battery discharge, the State-of-charge (SoC) of the battery minus DoD, and after the end of discharge, $DoDfinal$ is expressed as (100 - $SoC$)%. A higher $DoDfinal$ results in fewer cycles, so shallow discharging/charging is generally recommended. Additionally, $C_{eff}$ represents the effective capacity of the battery, calculated as follows:

$$C_{eff} = C_R * \left(\frac{C_R}{I_{discharge} * H}\right)^{k-1} * \frac{SoH}{100}, \tag{13}$$

where $H$ denotes the rated discharge time and $I_{discharge}$ is the discharge current. $k$ is Peukert's exponent, which takes different values for different battery types. In addition, a higher discharge current will lead to a reduction in battery capacity, called Peukert's law [43].

*2) Power Management Techniques:*

- *Chip power controlling techniques:* For CPU power controlling, the common method is to use DVFS [33] to directly limit the CPU clock frequency, or limit the cores used. For memory power, Intel's RAPL [44] integrates a set of memory power management schemes that can enforce memory power limitations. Recent work has analyzed that memory can be a bottleneck for improving resource utilization in DCs [45], implying that performance matching between chips is important.
- *Workload Scheduling:* Workload scheduling techniques include workload migration, placement, delay, and scheduling using virtualization technologies. Workload scheduling is a common computing resource optimization method at the server level. For power systems, workload scheduling enables the power supply devices to reach high energy efficiency ranges [46]. Moreover, it prevents peak stacking when workloads with the same power behavior are placed together [47].
- *Server sleep/shutdown:* Shutting down or sleeping servers to reduce server idle energy consumption. Several works [48], [49] have pointed out that low load levels or idle servers still generate a lot of idle power $P_{idle}$. In addition, the statistics provided by 80PLUS have shown that PSUs have greater power loss at low load levels.
- *Power supply device shutdown:* When power equipment operates at a low load level or is idle for extended periods, shutting down power modules can reduce idle energy consumption, provided the power-sharing domain allows for it. Fawaz et al. [50], [51] showed that the use of workload scheduling to consolidate servers and further turn off UPS can reduce more energy consumption.
- *Increasing the power-sharing domain:* The expansion of the power-sharing domain is beneficial to increase the power over-subscription level. From a security perspective, DC sets the power-sharing domain at the rack level or PDU level, using circuit breakers to isolate different power hierarchies [52]. However, this limits the scope of power sharing and results in power fragments that cannot be fully utilized. Existing works demonstrate the potential of power over-subscription in different power-sharing domains (rack level [53], PDU level [54], room level [55], [56], and medium voltage distribution level [57]).
- *ESDs supplement:* ESDs in DCs, in addition to being used to provide a brief power transition in case of a utility power failure, can also be used to provide short periods of peak shaving [58], or to enhance server performance [59] when peak demand is encountered and the power budget is insufficient.

### C. Relevant Metrics

It is worthwhile to explore the metrics related to the use of power resources in DCs because evaluating whether a DC is green or efficient requires specific metrics for reference. This subsection summarizes the metrics related to the evaluation of DC power resource usage.

1) *Power Usage Effectiveness (PUE):* PUE [60] is the most widely used metric to evaluate the energy efficiency of DCs. A higher PUE means that the energy consumption of IT system is a lower percentage of the overall DC. PUE is defined as:

$$PUE = \frac{E_{total}}{E_{IT}}. \tag{14}$$

However, it should be noted that people currently advocate energy saving and blindly pursue low PUE, which is not desirable. It is only used as an auxiliary evaluation criterion and is not comprehensive. Such as PUE does not consider the energy efficiency of IT system.

2) *Data Center Infrastructure Efficiency (DCiE):* DCiE [60] is defined as the reciprocal of PUE and is used to represent the energy consumption of IT system as a percentage of total DC energy consumption, as shown below:

$$DCiE = \frac{1}{PUE} \times 100\%. \tag{15}$$

3) *Power Overs-ubscription Level (POL):* At different power supply hierarchy in a DC, such as rack, PDU, or UPS, each hierarchy has a corresponding power capacity determined by upstream power supply equipment and associated circuit breakers. Nowadays, DCs often use power over-subscription to deploy more servers to increase the DC throughput and thus increase profit. The metric POL for measuring power over-subscription is shown in the following equation [61]:

$$POL = \frac{P_{\max} - P_{limit}}{P_{limit}}, \tag{16}$$

where $P_{\max}$ represents the theoretical maximum power of the current power supply node, and $P_{limit}$ represents the power limit of the node.

4) *Effective Power Utilization (EPU):* Compared to PUE, which is used to measure the energy consumption percentage of IT systems in DC, EPU can reflect the energy efficiency used to generate throughput. It can be used to evaluate the effectiveness of power control strategies and is defined as [62]:

$$EPU = \frac{\sum P_{throughput}}{\sum P_{supply}}, \tag{17}$$

where $\sum P_{throughput}$ indicates the power directly used to generate workload throughput and $\sum P_{supply}$ indicates the current total power supply. As the EPU approaches 1, it means that more power is directly used to generate workload throughput.

5) *Green Energy Coefficient (GEC):* The GEC used to quantify the percentage of renewable energy consumed by DCs, was introduced by Green Grid in 2012. GEC is defined as [60]:

$$GEC = \frac{E_{green}}{E_{total}}. \tag{18}$$

6) *Carbon Usage Effectiveness (CUE):* The CUE is proposed by Green Grid, and is the first widely adopted metric

related to DC carbon emissions to reflect the total carbon footprint. CUE is shown below [60]:

$$CUE = \beta * \frac{E_{total}}{E_{IT}} = \beta * PUE, \tag{19}$$

where $\beta$ is the carbon emission factor of the grid, representing the emissions required to generate 1 kWh of electricity. When calculating the CUE, the electricity generated by green energy should be subtracted from the total consumption. CUE should be as small as possible.

7) *Carbon Free Energy Score (CFE):* Since green energy acquired by DC operators through power purchase agreements (PPAs) or renewable energy certificates (RECs) is not counted in the calculation of CUE, it leads to an unfair assessment of carbon emissions. Therefore, Google proposed the CFE that integrates off-site green energy to score green DCs [63]:

$$CFEScore = \frac{CFE_{contracted} + CFE_{grid}}{E_{total}}, \tag{20}$$

where $CFE_{contracted}$ is the CFE provided by the green energy provider through a PPA or REC, $CFE_{grid}$ is the CFE from the grid, and $E_{total}$ is the total power consumed by the DC. The CFE combines the CFE from the contract with the CEF from the grid and provides a more comprehensive evaluation criteria.

8) *Grid Usage Effectiveness (GUE):* GUE is the Grid Usage Efficiency [64] determined by the On-site Energy Matching metric (OEM), On-site Energy Fraction (OEF). GUE shows the power grid dependence of the DC relative to the IT load. which is defined as:

$$GUE = \frac{\left(\frac{1}{OEM} - OEF\right) * E_{total}}{E_{IT}}. \tag{21}$$

OEM and OEF represent on-site renewable energy and power demand at specific time steps, with values ranging from 0 to 1. When both are 1, the DC's energy demand is fully met by on-site generation.

## III. Power Management Optimization of Traditional Energy Data Center

This section focuses on the study of the power supply side in terms of both direct and indirect costs in the context of traditional energy sources. For direct costs, we focus on overhead optimization with the involvement of power supply equipment. For indirect costs, we focus on utilization optimization of power capacity.

### A. Direct Cost-Aware Optimization

*Background:* During the entire lifecycle of a DC, from the power supply perspective, the direct cost mainly includes the power transmission loss and the electricity bill of the DC. A white paper [65] identified seven critical factors causing DC outages, as shown in Fig. 5. Excluding unforeseeable factors, the battery life of UPS systems heavily depends on the number of charge-discharge cycles and the DoD [66]. Therefore, in the following sections, in addition to discussing the energy efficiency optimization of the power system, we will focus on the study of electricity bill optimization with power supply devices consolidation (i.e., ESDs).

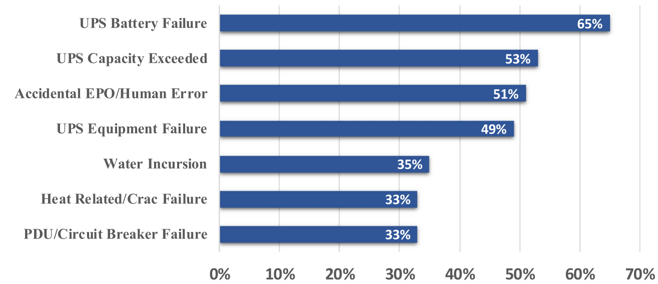


Fig. 5. Key elements lead to DC downtime.

TABLE II
THE POWER EFFICIENCY OF 80PLUS CITIFICATION FOR PSUS

| Loading | 10% | 20% | 50% | 100% |
|---|---|---|---|---|
| Bronze | - | 81% | 85% | 81% |
| Silver | - | 85% | 89% | 85% |
| Gold | - | 88% | 92% | 88% |
| Platinum | - | 90% | 94% | 91% |
| Titanium | 90% | 94% | 96% | 91% |

*1) Energy Saving Optimization in Power System:* To reduce the power losses in power supply systems, Ref. [31] has summarized several energy-saving strategies for power supply devices:

1) Upgrading power supply devices to improve conversion efficiency.
2) Matching power supply devices with workload to ensure high-efficiency operation.
3) Reducing the use of power supply devices to remove idle power.

In fact, energy-saving research often focuses more on strategy i. Examples include flywheel UPS [67] and supercapacitors (SC) [68] with higher conversion efficiency. However, the power management technologies related to strategies ii. and iii. are the primary focus of this subsection.

For the energy consumption of internal power supply devices in DCs, such as UPS, PDU, cables and PSU, most of the losses are due to thermal dissipation caused by the square law of current. Among these, PSU and UPS are the main components that can be adjusted using power control knobs. The PSU directly powers the motherboard of server. Its power conversion efficiency is related to the load level and PSU type [39], as shown in Table II. In fault-tolerant DCs, servers are equipped with dual PSUs that can switch the power mode (Active / Standby Mode and Load Balance Mode) to improve efficiency according to the load level [69]. With regard to UPS energy-saving optimization, the traditional approach is to use the Eco mode [70]. In addition, turning off / sleeping some UPS modules can reduce idle energy consumption, which is common in the early stages of DC operations. The above description is based on the mode switching of the power supply devices themselves for energy saving.

Regarding the improvement of power efficiency of power supply devices through power management techniques, Zhang et al. [46] based on the power efficiency curve of centralized UPS, through workload scheduling to make the UPS in the high energy efficiency area, which can achieve more energy saving under the same throughput. Similarly, Ye et al. [71] studied VM consolidation relying on the centralized UPS power efficiency curve. They formulated a hierarchical bin-packing problem for VM scheduling and used the best-fit decreasing algorithm to group VMs. Then, they designed a genetic algorithm (GA) for optimal VM scheduling decisions based on the high power efficiency area of UPS. Additionally, due to the flexibility of distributed UPS, energy savings can be achieved by consolidating UPS during low-load periods. Al-Hazemi et al. [50] employed a micro-ATS device to connect two adjacent rack-level UPS to the same server, expanding the power-sharing domain. Their approach consolidates UPS in addition to server consolidation, reducing energy consumption and mitigating PUE degradation. We know that PUE is expressed as follows:

$$PUE = \frac{P_{total}}{\sum_{u=1}^{U} P_{IT,u}}. \tag{22}$$

Among the non-IT equipment contains UPS, cooling systems, and so on. When exploring the relationship between UPS energy consumption and PUE, the PUE can be expressed as:

$$PUE = \sum_{u\in U} \frac{P_{IT,u}}{P_{IT,u}} + \frac{\sum_{u\in U} P_{UPS,u}}{\sum_{u\in U} P_{IT,u}} + \frac{P_{other} + P_{Cooling}}{\sum_{u\in U} P_{IT,u}}. \tag{23}$$

When IT devices are integrated for energy saving, $\sum_{u\in U} P_{IT,u}$ becomes smaller. Keeping the other variables constant, and using energy-saving methods for the UPS, then the value of the second term is reduced, and the degradation of PUE is mitigated. Furthermore, they improved the distributed UPS consolidation scheme [51], where the consolidation of three UPSs can better reduce UPS idle power compared to the consolidation of two adjacent UPSs [50].

*2) Electricity Bill Optimization With ESDs:* The electricity bills of a DC are divided into two parts: basic electricity expenses and peak demand charges. Basic electricity expenses refer to the total energy consumption over the operation period, charged by the electricity company based on the electricity price. Peak demand charges are incurred when the DC's maximum power exceeds the contractual budget. The total electricity bills can be expressed as:

$$J = \lambda_{elec}(t) * \sum_{t=1}^{T} s(t) + \lambda_{peak} * \max_{1,\ldots,T} s(t), \tag{24}$$

where $\lambda_{elec}(t)$ denotes the basic electricity price, which is determined by the electricity market. $\lambda_{peak}$ denotes the peak demand price, which is specified in the contract between the DC and the power supply company. $s(t)$ is the power consumption at time $t$. Optimization of the basic electricity expenses can be performed by adjusting the time-varying relationship between $s(t)$ and $\lambda_{elec}(t)$, while the peak demand charge is optimized by reducing the peak behavior of $s(t)$.

Additionally, DCs can fully utilize ESDs to actively participate in RS for generating revenue. In order to obtain RS support from the DC, the grid operator pays a fee $\lambda_c$ at intervals to the resource provided by the DC with a reserve power capacity of $C$. During RS, if there is a discrepancy between the resources procured by the grid operator and the actual response $b(t)$, a penalty is applied based on the penalty coefficient $\lambda_{mis}$. Let $r(t)$ be the normalized frequency regulation signal, i.e., how much frequency regulation the grid wishes to perform within the reserved power capacity $C$. The DC's revenue from participating in RS can be defined as:

$$R = \lambda_c C * T - \lambda_{\min} \sum_{t=1}^{T} |b(t) - Cr(t)|. \tag{25}$$

For the study of optimizing basic electricity expenses, Sun et al. [72], based on the Lyapunov optimization framework, constructed workloads and ESDs as virtual queues under electricity price and workload uncertainty. They proposed a distributed online algorithm for workload execution decisions, aiming to minimize electricity expenses while respecting load delay constraints. Lasemi et al. [73] focused on optimal workload scheduling and energy management for GEO-DCs with ESDs, considering time-varying electricity prices at different locations. They incorporated QoS, dynamic electricity pricing, and battery life management to optimize both battery lifespan and energy costs. The authors modeled the problem as a mixed-integer linear optimization and solved it using GAMS software.

Optimizing peak demand charges using ESDs, Nasiriani et al. [74] modeled the uncertainty in DC power demand as a Markov chain to assess the risks of overcharging or undercharging batteries due to this stochasticity. They balanced cost risks, including peak demand charges and battery degradation, based on power infrastructure and workload characteristics, using a Markov decision process (MDP) for online dynamic peak shaving. Since ESDs alone may struggle to meet the demand for peak shaving to optimize peak demand charges, Dabbagh et al. [75] added additional workload delay control variables, i.e., delaying workloads during peak periods. They simultaneously considered actual energy storage losses and battery constraints, proposing an integrated peak-shaving strategy that combines ESDs and workload management, making optimal decisions with full knowledge of future demands and outperforming existing technologies with limited future information.

With respect to the involvement of RS, Chen et al. [76] introduced EnergyQARE, which includes bidding strategies for RS participation and runtime strategies. For bidding, they constructed a power model for DCs and used numerical methods to optimize the strategy. Runtime strategies involve power management and workload QoS feedback, enabling DCs to accurately track frequency regulation signals and adjust power dynamically to meet QoS constraints. Furthermore, in addition to using the ESD of the DC to participate in RS, Wang et al. [77] combined ESDs with DGs to enhance benefits. They proposed a hierarchical IDC UPS scheduling strategy using model predictive control (MPC), integrating UPS state and workload uncertainties. Their approach ensures safe IDC operation with an MPC-driven energy model, while coordinating DG and ESD outputs via an upper-level power scheduling model and minimizing frequency regulation errors with a lower-level strategy. Moreover, ESDs can simultaneously optimize RS participation and peak

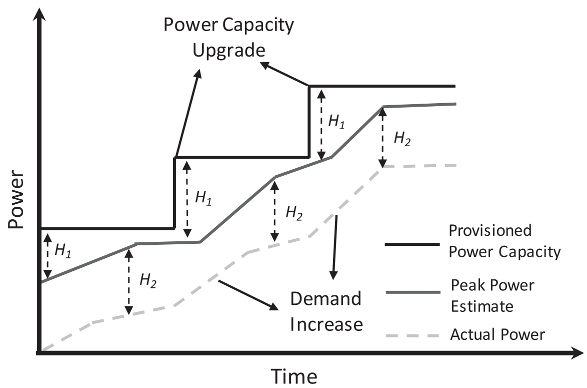


Fig. 6. The evolution of power capacity and demand in a DC.

demand charges. For instance, Shi et al. [78] showed that using ESDs for both RS participation and peak demand reduction yields super-linear gains. They addressed battery degradation, uncertainties in workloads, and frequency regulation signals, formulating convex optimization problems and implementing a threshold-based real-time control method, solved using CVX software.

### B. Indirect Cost-Aware Optimization

*Background:* To account for overloads and surges at each power hierarchy, and considering the need to expand computing capacity, DC power systems are designed with slack space at each power level, and circuit breakers are used for protection. However, conservative server configurations and diurnal patterns lead to low power utilization over extended periods, indirectly increasing the cost per unit of power supplied. Specifically, there are two types of headroom between the power demand and actual power budget in DCs, as illustrated in Fig. 6. The first headroom $H_1$ is the gap between the power budget and the estimated peak power, while the second headroom $H_2$ is the difference between the actual power and the estimated peak power. Fan et al. [37] examined the total power consumption characteristics of DCs and identified a significant gap between the actual and theoretical peak power at the cluster level, indicating potential capacity for deploying additional IT devices.

Therefore, appropriate capacity planning is essential during the pre-design phase to meet the power demand of IT systems. For operating DCs, power over-subscription can be an effective strategy to reduce the unit power supply cost (i.e., indirect cost) by deploying more servers than the power capacity limit. While power over-subscription improves resource utilization, it also creates vulnerabilities that malicious actors can exploit to launch power attacks [79]. This subsection explores research on optimizing power over-subscription and mitigating power attacks through advanced power management technologies.

*1) Power Over-Subscription:* Power over-subscription enhances power capacity utilization, thereby optimizing indirect costs. This subsection discusses the research on power over-subscription from the perspectives of power adjustment, workload scheduling, ESDs supplement, and power attack defense.

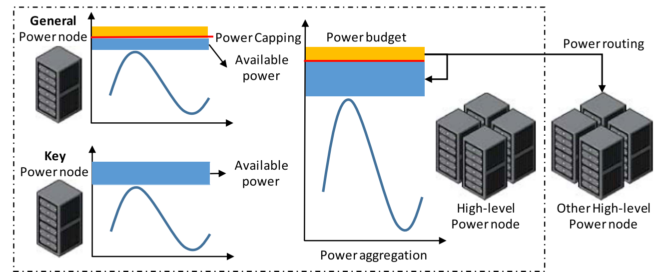


Fig. 7. Power Adjustments.

*Power Adjustments:* At power supply hierarchy with power over-subscription, to handle sudden peak power and prevent power resource competition, power control techniques (RAPL [44], DVFS [33], Thunderbolt [80], or computing resource allocation) are employed to achieve power adjustments. Additionally, power budget reduction can be further realized by enhancing the power-sharing domain. A simple example of power adjustments is shown in Fig. 7. For nodes executing general workloads can be power capped, allowing more power budget to be allocated to critical nodes at the same level or to another cluster node within the power-sharing domain.

In power-constrained scenarios, competition for power resources among servers requires precise power allocation. Wu et al. [53] introduced Precise Power Capping, using a Fine-Grained Differential approach to accurately assess performance degradation due to power capping. Therefore, historical peak power and the degradation model are used to guide power allocation to meet QoS of applications. Similarly, Patel et al. [81] analyzed power patterns in Large Language Model (LLM) training and inference scenarios, finding that LLM inference clusters have low average and peak power utilization. They proposed POLCA, a power over-subscription framework for LLM inference clouds, which uses a priority-based power reclamation strategy. POLCA also analyzes historical power data to set upper threshold values for different priority loads. For heterogeneous workloads, both Pelican [82] and CuttleSys [83] implement power-constrained control by core allocation. Baidu's Pelican power scheduling system is based on a greedy policy that prioritizes power reduction for execution of latency-tolerant workloads and high dynamic power servers. While CuttleSys, gathers profiling samples and power capping settings, using stochastic gradient descent to estimate power consumption and performance across various core configurations and cache allocations. Then it uses dynamic dimension search to quickly find the optimal configuration that can achieve a trade-off between performance and power consumption.

In cross-layer power resource optimization, several systems and strategies have been developed to improve power allocation and efficiency. Facebook's Dynamo [55] manages the power budget of its distributed system through collaboration between the agent, Leaf controller, and Upper-Level controller. The Leaf controllers use a three-band algorithm, a greedy policy based on high-bucket prioritization, to determine the capping power and servers, thus minimizing the impact on critical server performance. Also employing a greedy strategy, Piga et al. [84]

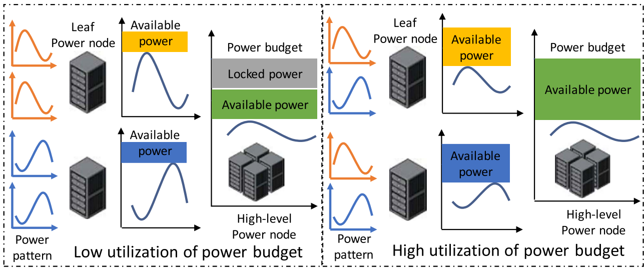


Fig. 8. Workload scheduling.

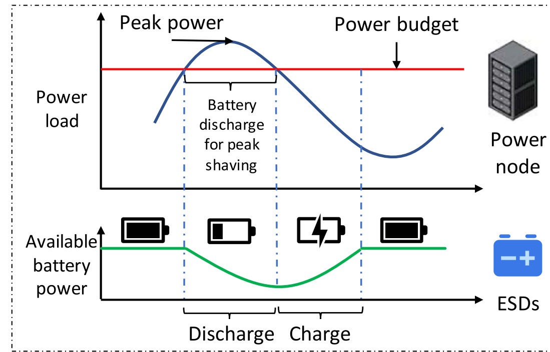


Fig. 9. ESDs supplement.

proposed a power over-allocation technique for heterogeneous servers and large-scale clusters, DVFS boosting, which utilizes performance counters and machine learning to predict which servers can provide higher performance per watt. However, IBM's CapMaestro [56] introduced a global priority-aware power allocation policy, a workload-balancing-based policy that routes power based on server priority and uses a proportional-integral feedback controller to ensure that PSUs do not exceed their allocated power budget. The above studies all consider workload priority to perform power capping. Azimi et al. [85] addressed actuation latency in cross-layer power management by proposing a decentralized power capping scheme that allows servers to quickly respond to workload throughput and priority, making local power capping decisions through cluster-to-cluster information exchange.

By altering the power-sharing domain to optimize power over-subscription, Pelley et al. [54] addressed power waste in 2 N fault-tolerant DCs, where PDU utilization is only 50%, by optimizing power over-subscription through a shuffled PDU topology. This topology alters the connections between PDUs and cabinets, reducing backup power capacity from X/2 to X/N, thereby lowering the need for backup resources. Zhang et al. [86] adopted a 4 N/3 connection method between UPS and PDU, increasing server deployment by 33% compared to the N+1 redundancy design. At a higher system level, Google's medium voltage power plane (MVPP) [57] enhances power over-subscription by enabling power sharing at a medium voltage distribution level, achieving an over-subscription rate of over 25%. Unlike RAPL [44], which is limited to Intel platforms, MVPP's power capping method is not platform-restricted.

*Workload Scheduling:* Workload scheduling involves placing, migrating, or delaying workloads to achieve load balancing, which directly affects power budget usage across different levels of the power supply system. In power over-subscribed DCs, balancing workloads is crucial to minimize risks. As shown in Fig. 8, optimizing workload scheduling with varying power patterns frees locked power budgets, allowing power surpluses to deploy additional servers or enhance server performance.

Considering that different workloads have different priorities, Microsoft proposed a criticality- and utilization-aware VM prediction method to guide over-subscription in public clouds [87], where RAPL reduces power for non-critical VMs when the power budget is exceeded, and lifts power capping when sufficient resources are available. They also introduced Flex [86] for *Zero Reserved Power* DCs, which uses the Flex-Offline workload placement algorithm based on integer linear programming to ensure workloads meet power reduction requirements by shutting down redundant software. To ensure availability and security, Flex's Power telemetry pipeline leverages redundant and diverse monitoring and transmission methods to ensure high availability and low latency, serving as a guarantee for risk response. Meanwhile, Pang et al. [88] studied the scheduling of high-priority latency-sensitive (LS) services and low-priority best-effort (BE) applications. They proposed Sturgeon, which evaluates the expected throughput of BE applications and LS services when co-located, and optimizes performance by colocating LS services with preferred BE applications.

To avoid peak power from workload stacking, Wang et al. [89] designed the Ampere power management system, which efficiently allocates power budgets between row-level cabinets, enabling the deployment of more servers without performance interference. Using statistical analysis, they identified servers unable to accept incoming jobs, scheduling tasks within the same cabinet or placing them in a queue to maintain row-level power consumption within budget. When considering workloads with long-running characteristics, such as service instances or VMs, Hsu et al. [47] proposed SmoothOperator, a peak-aware placement framework that uses the K-means algorithm to classify service instances with similar power characteristics, reducing peak stacking by allocating similar workloads to different nodes. For VM scheduling, Sheng et al. [90] introduced C2MARL, a chance-constrained multi-agent reinforcement learning approach for power over-subscription in VM scheduling. C2MARL uses probabilistic constraints from safe RL to optimize VM placement, mitigating power risks from colocating VMs.

*ESDs supplement:* In production DC environments, ESDs are rarely used to handle emergencies caused by utility power interruptions [91]. This makes ESDs ideal candidates for peak shaving, particularly in DCs with high levels of power over-subscription. However, using ESDs to reduce peak power requires careful control strategies, as frequent deep discharges and recharges can shorten battery lifespan and decrease the availability of the power supply system. As shown in Fig. 9, ESDs can be employed for short-term peak shaving when a node's peak power exceeds its pre-allocated power budget.

Earlier, ESDs were mostly centralized UPS. Govindan et al. [58] regarded centralized UPS as power buffers, marking the first study to use ESDs for peak shaving. Li et al. [92] proposed WattValet, a solution to reduce peak power using centralized heterogeneous UPS. WattValet considers the different power efficiencies of heterogeneous UPS, optimizing for increased use of high-efficiency UPS. For a given power demand, it uses a greedy strategy to search for a battery discharge power sequence that minimizes performance impact. Due to the inflexibility of centralized UPS, subsequent research has mainly focused on distributed UPS.

Google was the first to design a peak shaving solution based on server-level distributed UPS [5], focusing on how different battery characteristics affect peak shaving capabilities. In further explorations of distributed UPS, Alanazi et al. [93] compared server-level and rack-level distributed UPS, designing a management framework that includes VM placement and UPS power distribution strategies. The placement strategy reduces active servers and minimizes unused power for peak shaving, while the power distribution strategy optimizes UPS charging and discharging to minimize wasted power. Simulations showed that rack-level UPS, with larger power-sharing domains, reduces resource fragmentation compared to server-level UPS. Compared to the server or rack-level UPS mentioned above, Ref. [2] showed that placing ESDs at multiple levels of the power supply hierarchy offers better economic benefits. Thus, Wang et al. [59] proposed vPower, a software system that virtualizes power resources across different UPS levels and uses heuristic methods to optimize UPS selection, minimizing server performance degradation. To address power safety concerns in over-subscribed DCs, Malla et al. [94] found that battery charging could trip circuit breakers. They proposed a variable battery charger mechanism that reduces charging power by up to 80%, along with a priority-aware algorithm that allocates power based on server application priority. Additionally, they integrated Dynamo with lightweight agents and distributed controllers for real-time monitoring and adjustment of charging currents, preventing circuit breaker overloads.

*Power Attack Defense:* Power over-subscription can effectively improve the utilization of power resources, but it also provides opportunities for malicious power attacks. Xu et al. [79] systematically studied the three mainstream cloud service business models: PaaS, IaaS, and SaaS. They launched power attacks through carefully designed workloads and demonstrated its feasibility. It can be seen that for the power over-subscription scenario, it is crucial not only to implement peak shaving mechanisms but also to consider effective strategies for preventing malicious attacks.

Li et al. [95] proposed Power Attack Defense (PAD) to mitigate malicious power attacks in distributed energy storage scenarios. PAD addresses visible peaks through software scheduling, while using ESDs at server and rack levels to handle hidden peaks. Rack-level ESDs create virtual ESDs for power sharing, concealing the attacked server's status, and server-level ESDs manage hidden and sudden peaks. To further improve defense capabilities, they proposed Integrated PAD (IPAD) [96] to improve the power attack defense capability. IPAD integrates the Speculative Performance Scaling (SPS) strategy to reduce hidden power peaks in real time, based on monitored server-level ESD conditions. This minimizes the negative impact on server performance. The SPS mechanism triggers power capping, making the attacker believe their attack has been detected, which helps mitigate the risk of power overload. Hou et al. [97] addressed traffic flood attacks in power oversubscription scenarios with Anti-DOPE, a request-aware power management framework. It includes power-driven forwarding (PDF) and request-driven power management (RPM). PDF first identifies suspicious requests and directs them to isolated servers based on offline analysis. RPM then monitors power resource pressure and uses power capping to adjust the execution of these requests, preventing power peaks.

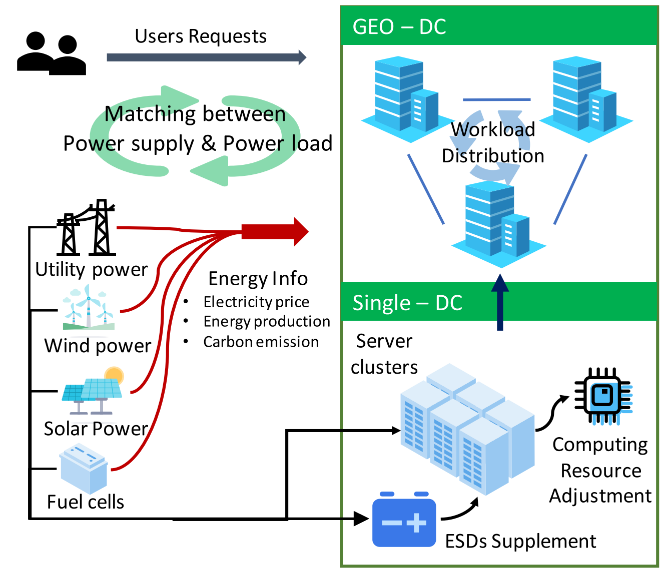


Fig. 10. Power matching optimization for green energy powered data centers.

Actually, participating in the electricity market is an effective approach to deal with power attacks. Specifically, Hou et al. [98] consider that users are cost-conscious and propose a flexible power capacity management framework CFP. CFP uses power bidding to increase the cost for power attackers and adopts an incentive mechanism to compensate users who save power capacity, further raising the cost for attackers.

## IV. Power Management Optimization of Green Energy Data Center

*Background:* Powering data centers (DCs) with green energy is a key strategy in achieving carbon neutrality and has become a promising solution [99], [108]. However, the intermittent nature of green energy can create mismatches between the power demand of DCs and the availability of renewable energy. In contrast to single data centers, geo-distributed data centers (GEO-DCs) offer greater spatial flexibility, enabling workloads to be shifted to DCs with lower electricity costs and abundant renewable energy. This flexibility provides a more adaptable solution to address power supply-demand mismatches, as illustrated in Fig. 10. However, Sukprasert et al. [100] presented an analytical article arguing that there are certain limitations to carbon-aware spatiotemporal workload shifting. Temporally, it is constrained by fluctuations in carbon emission intensity, while spatially, it is hindered by resource limitations. Even if achieving net-zero carbon emissions is the ultimate goal, most DC operators must also consider factors such as electricity prices, service quality, and the maintenance costs of ESDs. These considerations influence both the TCO and the environmental impact of DC operations, requiring a delicate balance between environmental costs and profit optimization.

In the context of green energy supply, the problem can be viewed as curve matching between the *Power Generation-Time* curve $g(t)$ and the *Power Consumption-Time* curve $c(t)$ [101]. Theoretically, power management technologies could achieve arbitrary transformations of the curves to make them perfectly aligned. However, in practice, due to constraints such as efficiency optimization, service quality, electricity costs, and maintenance costs, DC operators prefer the approach of matching the *Balanced Power Generation-Time* $bg(t)$ and *Balanced Power Consumption-Time* curves $bc(t)$. Because the balance between TCO and environmental cost is obviously more important than only considering environmental cost. We define the function $p(t) = \max[g(t) - c(t),\ 0]$ as the renewable energy wasted at time $t$, and the function $q(t) = -\min[0,\ g(t) - c(t)]$ as the brown energy wasted at time $t$, then the objective function of the optimization problem can be expressed as:

$$C = \int_0^T [\omega_\alpha(t)p(t) + \omega_\beta(t)q(t)]\, dt, \tag{26}$$

where $\omega_\alpha(t)$ represents the generation cost of renewable energy, and $\omega_\beta(t)$ represents the generation cost of brown energy, both as functions of time $t$. This objective function equals 0 if $g(t) = c(t)$. However, when applying power management strategies like power tracking for reshaping $c(t)$ and load following for reshaping $g(t)$, relevant constraints must be considered. We define the power generation and consumption curve functions to be able to be transformed to $bg(t)$ and $bc(t)$ based on power management techniques:

$$bg(t) = M_{bg}(t)g(t), \tag{27}$$

$$bc(t) = M_{bc}(t)c(t). \tag{28}$$

The cost transformation operators $M_{bg}(t)$ and $M_{bc}(t)$ that concern the operators can be seen as tools to reshape the power generation and consumption curves to meet various operational and environmental goals. These transformation operators can take into account the following factors: the utilization of green energy, the resource constraints of the ESD, or the electricity price at different times, etc. If the mathematical expressions or problem models for these operators can be found, the constraints of the transformation become clear. The problem then turns into solving the following optimization problem:

$$C_{total} = \int_0^T [\omega_\alpha(t)bg(t) + \omega_\beta(t)bc(t)] + C\left(M_{bg}(t), M_{bc}(t)\right) dt, \tag{29}$$

where $C(M_{bg}(t), M_{bc}(t))$ represents the cost incurred due to the power reshaping transformations and the objective function $C_{total}$ accounts for both environmental costs and TCO.

Power management strategies can be divided based on whether they integrate ESDs, specifically into power tracking design and load following design [102].

- *Power tracking design:* DCs actively control the power demand of loads in order to reshape the power curve, including power allocation techniques (power resource routing, power limitation, computing resource allocation, etc.) or load limitation techniques (workload scheduling, latency, etc.) in order to change the power curve of loads.
- *Load following design:* Different from the power tracking design, beyond requiring power allocation and load limitation techniques, load following design also uses ESDs within the DC to reshape the $g(t)$ to $bg(t)$.

In the following, we will categorize and discuss Single-DC and GEO-DC, and further explore the research of the two power management modes to give a clear and complementary discussion.

### A. Single-DC Power Matching Optimization

This section primarily discusses how the two strategies of power tracking design and load following design achieve the fitting of $bg(t)$ and $bc(t)$ from the perspective of power management technology. Specifically, three types of control knobs (power adjustment, workload scheduling, and ESDs supplement) are used as decision variables and discussed separately according to their combinations.

*1) Power Tracking Design:* Research on the use of power tracking design in a Single-DC has focused on either power adjustment techniques or workload limiting techniques for the purpose of reshaping the power curve. Studies related to both techniques are discussed separately below. A summary of related work about power tracking design in Single-DC can be viewed in the Table III.

When considering research on power adjustment techniques, the most direct application is power routing. DCs with distributed green energy sources, each of which has different power production characteristics. When green energy supply is insufficient, switching between sources can impact the QoS of workloads, so it is crucial to align the supply and demand as closely as possible. For example, Gao et al. [103] used LSTM to predict the probability that renewable energy production would meet a certain threshold, while estimating power demand based on CPU utilization of each PM group. Then DQN solves the power matching problem and outputs the mapping decision for the allocation of renewable energy to PM groups. However, this approach assumes accurate resource availability predictions. To address potential instability in energy forecasts, they later proposed a more robust renewable energy allocation system [104], which introduced ECRA, a heuristic algorithm for adjusting computational resources based on job deadlines. Similarly, Xing et al. [105] introduced Carbon Responder, a performance-aware power allocation framework for DR participants. By defining SLOs for both online and batch workloads, Carbon Responder

TABLE III
GREEN ENERGY OPTIMIZATION IN A SINGLE DATA CENTER (POWER TRACKING DESIGN)

| Ref. | Decision Variables | Objectives | Formulation / Problem | Algorithm / Solution |
|---|---|---|---|---|
| [104] | Power routing | Cost, Carbon emission | Markov decision process | Deep reinforcement learning |
| [105] | Computing resource allocation | Cost, QoS, Carbon emission | Computing resource allocation problem | Heuristic algorithms based on ddl |
| [106] | Computing resource allocation | QoS, Carbon emission | Computing resource allocation problem | Heuristic algorithms based on power-performance modeling |
| [107] | VM scheduling | Brown energy | Bin-packing problem | Heuristic algorithms based on energy supply |
| [108] | Job scheduling | Carbon emission, Energy consumption | Markov decision process | Deep reinforcement learning |
| [100] | Workload scheduling | Carbon emission | Stochastic optimization problem | Heuristic algorithms of model prediction |

TABLE IV
GREEN ENERGY OPTIMIZATION IN A SINGLE DATA CENTER (LOAD FOLLOWING DESIGN)

| Ref. | Decision Variables | Objectives | Formulation / Problem | Algorithm / Solution |
|---|---|---|---|---|
| [109] | ESDs supplement | QoS, EDSs lifetime | Heuristic based control framework | Control algorithm based on heuristic strategy |
| [112] | Workload scheduling ESDs supplement | Renewable energy utilization, EDSs lifetime | Multi-variable nonlinear scheduling problems | Greedy strategy |
| [113] | Workload scheduling ESDs supplement | Renewable energy utilization, Cost | Workload optimization control problem | Control algorithm based on greedy policy |
| [114] | Workload scheduling ESDs supplement | Cost | Mixed integer linear programming | Column-and-constraint generation algorithm |
| [111] | Power routing ESDs supplement | Cost, EDSs lifetime | Large-scale linear programming | Matlab CPLEX Solver |
| [62] | Power limitation Computing resource allocation ESDs supplement | QoS | Quadratic curve of server performance and power allocation ratio | Power allocation based on Quadratic curve Solver |
| [110] | Power limitation ESDs supplement | Cost, QoS | Stochastic optimization problem with Chance- and Risk-Constrained | IBM ILOG CPLEX |
| [99] | Workload scheduling ESDs supplement | Carbon emission | Stochastic optimization problem | Carbon aware greedy algorithm |
| [117] | Workload scheduling Computing resource allocation ESDs supplement | Carbon emission | Stochastic optimization problem | Rule-based algorithm |

uses machine learning to train models that predict performance and power losses, then adjusts power allocation based on the workloads that incur the least marginal performance loss.

In the research on adjusting $c(t)$ using workload limiting techniques, Liu et al. [107] addressed large-scale job scheduling, considering job dependencies, heterogeneity, and QoS. They formulated the scheduling problem as MDP and proposed a DRL-based approach to achieve energy-aware, online scheduling. Google's Carbon-Intelligent Computing Management System [108] employed a data-driven approach, gathering and analyzing data on server power consumption, workload forecasts, power contracts, and carbon intensity. The system constructs a carbon-aware virtual capacity curve (VCC) to optimize resource allocation, deferring execution for flexible workloads. Additionally, for VM scheduling problem, Chakraborty et al. [106] proposed a framework for elastic power utilization. They leveraged energy information to drive VM overbooking, migration, and consolidation, thereby matching the power consumption of workloads with the supply of renewable energy. Specifically, when the supply of green energy is insufficient, the framework uses heuristics based on monitored renewable energy supply information to overbook VMs, reducing the number of active hosts and lowering energy consumption.

*2) Load Following Design:* In load following design of a Single-DC, the introduction of new control variables ESD gives more optimization space for the power matching problem but also introduces more constraints. A summary of related work about load following design in Single-DC can be viewed in the Table IV.

When decision variables are limited to the charging and discharging of ESDs, the power matching problem presents significant challenges. Liu et al. [109] introduced the HHEB energy buffering technology, deploying multiple heterogeneous ESDs (SC and UPS) at different levels (PDU-level, rack-level, server-level) and using a triple exponential forecasting method to predict peak situations at each level. For small peaks, peak shaving is conducted using SC with high efficiency and rapid charge/discharge capability. For large peaks, SC and UPS are dynamically coordinated for power distribution. When single-level ESDs cannot meet power demands, higher-level ESDs are coordinated to assist in managing power imbalances.

Some studies explore joint optimization of power adjustments and ESDs. For instance, Kwon [110] optimized server configurations and electricity procurement in DCs by integrating DR decisions. He proposed a two-stage stochastic program: the first stage sets server on/off schedules and day-ahead electricity commitments based on historical data, while the second stage optimizes real-time DVFS adjustments, solar power use, and ESDs, then chance and risk constraints are introduced for decision-making to ensure QoS. Naturally, due to the inevitable updates in DC equipment, heterogeneous conditions often arise between IT systems and power infrastructure. For heterogeneous ESDs, Gu et al. [111] proposed the GreenFlowing power scheduling scheme to minimize total power costs. They developed performance models for different ESD types and solved charge/discharge and power allocation decisions using large-scale linear programming. For heterogeneous servers, Cai et al. [62] introduced GreenHetero, a dynamic power allocation framework optimized with rack-level UPS. They used a Power-Performance relational database and quadratic functions to quickly determine optimal power consumption ratios. Obviously, modeling is very important for heterogeneous server or heterogeneous ESD scenarios.

Since the ability to use ESDs to smooth green energy is always limited, joint optimization is often combined with load-limiting techniques, primarily for delay-tolerant workloads. Some studies will be based on the greedy strategy. For instance, the Carbon Explorer of Meta [99] focuses on carbon emissions, using a carbon-aware greedy policy that delays workloads to perform at times of low carbon intensity. Similarly, Yang et al. [112] first judge the fluctuation of power request through two-stage low-pass filters. For low-frequency fluctuations, ESDs are used to provide short-time clipping. For high-frequency fluctuations, due to the limited peak shaving capability of ESDs, workload scheduling with maximum user satisfaction is used to deal with them within the acceptable delay latency. Additionally, some studies adopt two-stage frameworks for more comprehensive optimization. For example, the Smoother proposed by Liu et al. [113] will first forecast the green energy and use ESDs to actively smooth the green energy. Then, employing a greedy strategy to select greener and lower-cost electricity time slots for delay-tolerant workload execution. While Zhou et al. [114] introduced a two-stage optimal operation model based on Distribution Robust Optimization to minimize DC costs. In the first stage, they formed fuzzy sets by combining norm-1 and norm-inf to capture uncertain probability distributions of green energy, minimizing operational and carbon emission costs through delayed workload execution. In the second stage, they solved a mixed-integer linear programming problem involving the coordination of ESDs, DGs, and dynamic electricity prices using column-and-constraint generation methods to optimize workload rescheduling costs.

Joint optimization of ESDs, power adjustments and load limiting techniques has also been explored. Recent work [115], [116], [117] highlights that the energy system and its information are often hidden within power management systems, preventing applications from directly interacting with the energy system to optimize carbon efficiency. To tackle this issue, Prashant Shenoy's team proposed the Ecovisor software-defined control system [117] and the Carbon Containers tool [115]. These innovations virtualize energy systems and expose them directly to containerized applications, enabling applications to optimize green energy utilization based on the flexibility and fault tolerance of their own software. By providing carbon emission rate settings to containerized applications, they empower applications to enforce carbon emission rate settings through virtual battery charging/discharging, container scaling, migration, pause, or resume actions, based on renewable energy availability and carbon intensity variations.

### *B. GEO-DCs Power Matching Optimization*

In GEO-DCs, the simultaneous consideration of utilizing green energy across multiple DCs while minimizing total costs increases the complexity of optimization. Compared to Single-DC, where workloads can only be managed within one location, distributing workloads across different DCs introduces a larger optimization space in the spatial dimension. Therefore, spatial optimization is essential, otherwise, the study would regress to that of a Single-DC. While this expanded optimization space offers potential benefits, it also imposes higher demands on algorithm design. From a classification of algorithms, which can be divided into three types:

- *Traditional Heuristic Algorithms:* These rely on intuitive logic to make decisions on constructed optimization objectives, offering fast solutions but often not achieving optimal results.
- *Intelligent Optimization Algorithms:* These simulate the decision-making logic of intelligent agents, encompassing swarm intelligence, evolutionary algorithms, and reinforcement learning.
- *Numerical Optimization Algorithms:* These include dedicated optimization solvers and decomposition algorithms for large-scale optimization problems, capable of achieving precise optimal solutions or approximations.

Therefore, this subsection discusses the problem of matching the $bg(t)$ and $bc(t)$ curves through algorithmic approaches.

*1) Power Tracking Design:* To address the power matching problem, a heuristic approach places workloads in DCs with higher green energy availability. For example, Sharma et al. [118] employed the ARIMA forecasting algorithm to predict solar energy availability. They devised a Renewable Energy-Aware Worst-fit greedy strategy to identify the DC with the most green energy at the current time and allocate workloads to that DC. Furthermore, in addition to considering the availability of green energy, workload priorities can also be considered to make decisions. Kaur et al. [119] used the Boruta random forest algorithm to select job features and classify priorities using a locally sensitive hashing SVM. Then they formulated a multi-objective optimization problem for job scheduling and VM placement, using an augmented greedy heuristic for scheduling decisions.

Some studies use intelligent algorithms for power tracking decisions. For example, Ammari et al. [120] modeling the multi-task scheduling problem among GEO-DCs as a nonlinearly constrained problem. They proposed an improved Firefly Algorithm to maximize renewable energy use while strictly meeting task latency constraints. While the bi-objective optimization problem was considered by Khalid et al. [121]. Their solution optimizes both revenue (by offering more services) and expenses (due to dynamic power consumption and electricity prices), aiming for Pareto-optimal solutions that balance profit maximization and cost minimization. Whereas in research based on RL, CFWS [122] and DeepScale [123] were proposed to solve the VM scheduling and container expansion problems in virtualization technology. Specifically, CFWS incorporates the adaptive threshold adjustment method to assess the probability

TABLE V
GREEN ENERGY OPTIMIZATION IN GEOGRAPHICALLY DISTRIBUTED DATA CENTERS (POWER TRACKING DESIGN)

| Ref. | Decision Variables | Objectives | Formulation / Problem | Algorithm / Solution |
|---|---|---|---|---|
| [118] | Job scheduling | Renewable energy utilization | Job scheduling decision problem | Control algorithm based on greedy policy |
| [119] | Workload scheduling VM scheduling | Cost, QoS, Renewable energy utilization | Multi-objective optimization for job scheduling and VM placement | Enhanced heuristic approach based on greedy strategy |
| [125] | Task migration | Carbon emission | Mixed integer programming | Benders decomposition algorithm |
| [126] | Computing resource allocation Workload scheduling | QoS | Nonlinear programming problem | $fmincon$ |
| [121] | Computing resource allocation | Profit | Constrained multi-objective optimization problem | Evolutionary algorithm-based higher-level heuristic |
| [122] | VM scheduling | Renewable energy utilization, QoS | Bin-packing problem | Deep reinforcement learning |
| [124] | Workflow scheduling | Energy consumption, Brown energy | NP-hard problem | Multi agent reinforcement learning |
| [120] | Task scheduling | Renewable energy utilization | Constrained nonlinear problem | Improved Firefly algorithm |
| [123] | Container scaling | Cost | Location-aware container scaling problem | Deep reinforcement learning |

of host overload, preventing unnecessary VM migrations. It utilizes DQN networks to learn VM consolidation strategies, optimizing the dual objectives of energy cost and carbon footprint in GEO-DCs. DeepScale achieves automatic expansion of containers by predicting workload patterns and dynamically adjusting vCPUs according to demand. However, existing single-agent RL algorithms are ineffective for handling the decentralization and adaptive control challenges of GEO-DCs. Therefore, Jayanetti et al. [124] addressed workflow scheduling problems and proposed an enhanced multi-agent hierarchical scheduling framework. At the top level, the Global Scheduler is responsible for task submission to executing DCs, while at the bottom level, the Local Scheduler manages server selection. This framework improves training efficiency by sharing experiences among local agents, addressing the curse of dimensionality.

The following studies demonstrate the application of numerical optimization methods to address complex decision-making scenarios in GEO-DCs. Yang et al. [125] studied scenarios where green energy sources in GEO-DCs complement each other at specific times. They modeled GEO-DCs interconnected by optical networks and formulated a large-scale mixed-integer programming problem for task migration decisions in both temporal and spatial dimensions. Using the Benders decomposition algorithm, they optimized decisions for carbon reduction. Similarly, for transactional workload scheduling, Cheng et al. [126] used queuing theory to model transactional workload performance and formulated a nonlinear programming problem that integrates weather data, workload arrival rates, and service rates to optimize placement and resource allocation. The *fmincon* function in Matlab was then used to find the optimal solution. A summary of related work about power tracking design in GEO-DCs can be viewed in the Table V.

*2) Load Following Design:* In GEO-DCs with load following design, there are necessarily two or more decision variables. This brings challenges to traditional heuristics to optimize the power matching problem. Pahlevan et al. [127] explored DCs participating in RS, leveraging power adjustments and ESDs for profit. They proposed ECOGreen, a cost-effective online strategy that divides decision-making into two stages: bidding for RS and VM allocation. The strategy uses the forecasts of VM workload to quickly analyze optimal power and reserve bidding values, followed by a VM allocation method that minimizes correlation and peak stacking. Similarly, Nadalizadeh et al. [128] proposed GreenPacker, a renewable and fragmentation-aware VM placement algorithm, considering the resource fragmentation that can result from VM placement decisions. GreenPacker quantifies the suitability of a DC for a given VM request by introducing a cost metric, which is based on the availability of green energy, dynamic electricity price, PUE, and fragmentation caused by the placement. Furthermore, Yang et al. [129] presented a two-stage heuristic power regulation algorithm tailored for scenarios with fluctuating power demands across GEO-DCs. This method enhances the responsiveness of GEO-DCs to power demand. It includes workload scheduling and UPS group control algorithms. For workload scheduling, a low-pass filter is employed to smooth power demand control objectives across GEO-DC interconnections, addressing high-frequency power fluctuations. Simultaneously, a battery control model is established to make charging and discharging decisions while meeting UPS soft constraints, effectively managing low-frequency fluctuations.

For the study using intelligent optimization algorithms, Guo et al. [130] investigated the use of GEO-DCs to participate in energy trading in electricity markets. They proposed a two-stage framework. In the first stage, a preemptive bidding model is formulated as a Mixed Integer Non-Linear Programming problem, where a Meta-heuristic Natural Aggregation Algorithm manages computational requests and ESDs. In the second stage, they developed an energy balance model based on electricity market regulation signals, using a Mixed Integer Linear Programming formulation to optimize the migration of computational requests in response to market conditions. In the research involving RL, a single agent is insufficient to effectively handle cases involving two or more decision variables. To address this, Sarkar et al. [131] explored the application of MARL in optimizing multiple decision variables in GEO-DCs. They introduced the DC-CFR MARL framework, creating three MDPs for cooling, workload migration, and battery storage. The framework facilitates interaction among three agents, each managing an MDP, with collaborative rewards synthesizing the impacts of their decision actions.

TABLE VI
GREEN ENERGY OPTIMIZATION IN GEOGRAPHICALLY DISTRIBUTED DATA CENTERS (LOAD FOLLOWING DESIGN)

| Ref. | Decision Variables | Objectives | Formulation / Problem | Algorithm / Solution |
|---|---|---|---|---|
| [132] | Workload scheduling ESDs supplement | Carbon emission, Cost | Lyapunov stochastic optimization | ADMM |
| [128] | VM scheduling ESDs supplement | Cost, QoS | Multi dimensional online bin packing problem | Renewable and Fragmentation aware greedy Algorithm |
| [133] | Power routing Workload scheduling ESDs supplement | Cost | Mixed integer quadratic programming (MIQP) | Branch-and-cut method |
| [129] | Workload scheduling ESDs supplement | Cost, ESDs lifetime | Constrained nonlinear programming problem | Two-stage heuristic algorithm |
| [127] | VM scheduling ESDs supplement | Cost | Nonlinear integer programming problem | Rule-based algorithm |
| [130] | Workload scheduling ESDs supplement | Cost | MINLP, MILP | Natural aggregation algorithm, MIP Solver |
| [131] | Workload migration ESDs supplement | Carbon emission Energy consumption | Markov decision process | Multi agent reinforcement learning |

Within the research using numerical optimization methods, Zhang et al. [132] advanced the integration of ESDs in GEO-DCs for spatio-temporal workload scheduling. They formulated a stochastic optimization problem to minimize the total weighted cost, including electricity charges, water consumption, and carbon emissions. Using Lyapunov techniques, they developed the LYA-OACM online algorithm to balance cost and QoS for delay-tolerant workloads. To reduce computational complexity and communication overhead between DCs, they proposed the ADMM-DACM algorithm, allowing each DC to make independent control decisions. In contrast to existing works that primarily consider the lifespan of ESDs, Ye et al. [133] leveraged the nonlinear characteristics of UPS power losses to propose a novel approach to optimizing the joint operation of GEO-DCs. They integrated local electricity prices, renewable energy generation, conventional generators, ESDs, and UPS nonlinearities into a Mixed Integer Quadratic Programming model, solving day-ahead and intra-day scheduling tasks with branch-and-cut methods. A summary of related work about load following design in GEO-DCs can be viewed in the Table VI.

## V. OPEN ISSUES AND FUTURE RESEARCH DIRECTIONS

With the prevalence of cloud computing services, the demand for computing power in DCs is rising. As a highly energy-intensive building, it needs to be equipped with a reliable power management system to ensure power supply safety and service quality. Especially when green energy is integrated into the power supply, advanced power management technologies are needed to overcome the challenges brought by green energy. Based on the observation of the existing works, we give the open issues and future trends of DCs when it comes to power management optimization from the perspective of the power supply side.

### A. Safety Evaluation of Power Over-Subscription

Through power over-subscription, data centers (DCs) can deploy more servers to meet computing capacity demands, but measuring the associated power supply risks is challenging. Existing safety evaluation models rely on static overload probabilities [61] and fail to account for the relationship between the magnitude and duration of overloads [52]. In reality, even short periods of over-peak can pose safety risks. Therefore, it is essential to develop a more scientific and robust safety evaluation model.

### B. Availability Assessment of Distributed Power Supply Structures

DCs can categorize their availability levels based on criteria established by UPTIME [20] for centralized UPS architectures. However, assessing the availability of DCs with distributed UPS structures, such as rack-level and server-level UPS setups, poses a significant question for discussion. A critical question is whether it is possible to easily transpose the availability assessment approach of centralized power structures to distributed power structures? Especially with distributed power supply structures becoming a future trend, availability as a label affects customer choice and the DC revenue indirectly.

### C. Green Energy Management for AI Workload Integration

With the development of green energy and the gradual increase of AI-enabled scenarios (especially those of LLM integration), DC operators, as the intermediary between power supply companies and DC customers, will inevitably need to consider the optimization of power resources between green energy and AI workloads. AI training is a batch task that becomes more power-hungry with accelerators, while AI inference is latency-sensitive and its power consumption varies across different inference modes [134]. Thus, effective power management is crucial for optimizing power matching in AI scenarios.

### D. Virtualized Energy System

Similar to Software-Defined Networking (SDN), DC power systems can also be software-defined, i.e., virtualized energy systems [116], [117]. Specifically, the energy accessed by DCs, whether brown or green energy, can obtain relevant energy information through monitoring systems (such as electricity prices, carbon emission factors, power budgets, and capacities). Virtualizing energy systems through a software-defined approach can provide greater power management flexibility. The key challenge is how to leverage virtualized energy information with different characteristics of applications (e.g., scalable containers, lightweight microservices, etc.) to exploit greater optimization potential.

## VI. Conclusion

To satisfy the demand for cloud computing services, DCs have been developed and become representative buildings with high cost, high energy consumption and high carbon emission. Therefore, for sustainable development, modern DCs have to pay attention to cost optimization (both TCO and environmental costs) on the power supply side, except for optimizing IT or cooling systems. To provide a full view of the problem, we systematically investigate the study of cost-aware DC power supply-side resource optimization. We first present basic knowledge about the power supply perspective, including the power supply system structure and related metrics, as well as modeling approaches for key components and power management techniques. Then, we analyze and compare crucial works including direct, indirect, and environmental costs, emphasizing the ways in which power management techniques and methods are involved. Finally, as part of the survey, we point out open challenges and future research directions for power management in power supply-side resource optimization. We hope this survey will provide useful assistance to researchers and engineers, aiming to offer valuable guidelines for building greener DCs.

## References


[1] A. Shehabi, S. J. Smith, E. Masanet, and J. Koomey, "Data center growth in the United States: Decoupling the demand for services from electricity use," *Environ. Res. Lett.*, vol. 13, no. 12, Dec. 2018, Art. no. 124030, doi: 10.1088/1748-9326/aaec9c.

[2] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: What, where, and how much?," *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 1, pp. 187–198, Jun. 2012, doi: 10.1145/2318857.2254780.

[3] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.aba3758

[4] N. Jones, "How to stop data centres from gobbling up the world's electricity," *Nature*, vol. 561, no. 7722, pp. 163–167, 2018.

[5] V. Kontorinis et al., "Managing distributed UPS energy for effective power capping in data centers," in *Proc. 39th Annu. Int. Symp. Comput. Archit.*, 2012, pp. 488–499.

[6] S. Govindan, J. Choi, B. Urgaonkar, A. Sivasubramaniam, and A. Baldini, "Statistical profiling-based techniques for effective power provisioning in data centers," in *Proc. 4th ACM Eur. Conf. Comput. Syst.*, New York, NY, USA, 2009, pp. 317–330, doi: 10.1145/1519065.1519099.

[7] S. Pelley, D. Meisner, T. F. Wenisch, and J. W. VanGilder, "Understanding and abstracting total data center power," in *Proc. Workshop Energy-Efficient Des.*, 2009, pp. 1–6.

[8] L. Alt, *Energy Utility Rate Setting*. Morrisville, NC, USA: Lulu.com, 2006.

[9] M. Chen, C. Gao, M. Song, S. Chen, D. Li, and Q. Liu, "Internet data centers participating in demand response: A comprehensive review," *Renewable Sustain. Energy Rev.*, vol. 117, 2020, Art. no. 109466. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032119306744

[10] How clean is your cloud? Apr. 2012. [Online]. Available: https://www.greenpeace.org/static/planet4-international-stateless/2012/04/e7c8ff21-howcleanisyourcloud.pdf

[11] B. Nogrady, "China launches world's largest carbon market: But is it ambitious enough?," *Nature*, vol. 595, no. 7869, pp. 637–637, Jul. 2021. [Online]. Available: https://ideas.repec.org/a/nat/nature/v595y2021i7869d10.1038_d41586-021-01989-7.html

[12] Alibaba group carbon neutrality action report, Dec. 2021. [Online]. Available: https://sustainability.alibabagroup.com/en

[13] Tencent carbon neutrality target and roadmap report, 2021. [Online]. Available: https://static.www.tencent.com/attachments/TencentCarbonNeutralityTargetandRoadmapReport.pdf

[14] 24/7 carbon-free energy: Methodologies and metrics, Feb. 2021. [Online]. Available: https://www.gstatic.com/gumdrop/sustainability/24x7-carbon-free-energy-methodologies-metrics.pdf

[15] Facebook's net zero commitment, Sep. 2020. [Online]. Available: https://sustainability.fb.com/wp-content/uploads/2020/12/FB_Net-Zero-Commitment.pdf

[16] S. Malla and K. Christensen, "A survey on power management techniques for oversubscription of multi-tenant data centers," *ACM Comput. Surv.*, vol. 52, no. 1, pp. 1–31, Feb. 2019, doi: 10.1145/3291049.

[17] F. Kong and X. Liu, "A survey on green-energy-aware power management for datacenters," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 1–38, Nov. 2014, doi: 10.1145/2642708.

[18] Z. Cao, X. Zhou, H. Hu, Z. Wang, and Y. Wen, "Toward a systematic survey for carbon neutral data centers," *IEEE Commun. Surv. Tut.*, vol. 24, no. 2, pp. 895–936, Second Quarter, 2022.

[19] D. Wang, S. Govindan, A. Sivasubramaniam, A. Kansal, J. Liu, and B. Khessib, "Underprovisioning backup power infrastructure for datacenters," in *Proc. 19th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2014, pp. 177–192, doi: 10.1145/2541940.2541966.

[20] W. P. Turner IV, J. PE, P. Seader, and K. Brill, "Tier classification define site infrastructure performance," *Uptime Inst.*, vol. 17, 2006. Accessed: Jul. 9, 2024. [Online]. Available: http://www.wilsonengineered.com/Tech/TUI705CTierClassification_WP.pdf

[21] K. M. U. Ahmed, M. Alvarez, and M. H. Bollen, "Reliability analysis of internal power supply architecture of data centers in terms of power losses," *Electric Power Syst. Res.*, vol. 193, 2021, Art. no. 107025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378779621000055

[22] Z. Cai and X. Tian, "The research on data center energy saving technology based on DC power supply technology," in *Proc. 4th Int. Conf. Inf. Technol. Manage. Innov.*, Atlantis Press, 2015/10, pp. 76–84, doi: 10.2991/icitmi-15.2015.15.

[23] S. Qi, W. Sun, and Y. Wu, "Comparative analysis on different architectures of power supply system for data center and telecom center," in *Proc. IEEE Int. Telecommun. Energy Conf.*, 2017, pp. 26–29.

[24] W. Whitted et al., "Data center uninterruptible power distribution architecture," Jun. 2006, Patent No. CA2653778A, Filed Jun. 27th, 2006, Issued Dec. 13th, 2007. Accessed: Jul. 9, 2024. [Online]. Available: https://patents.google.com/patent/US7560831B2/en

[25] A. Michael, "Facebook: The open compute project," in *Proc. IEEE Hot Chips Symp.*, Los Alamitos, CA, USA, 2011, pp. 1–41. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/HOTCHIPS.2011.7477483

[26] "OCP summit 2015 - Microsoft," Mar. 2018. [Online]. Available: https://cloud.tencent.com/developer/article/1063539

[27] C. Li, W. Zhang, C.-B. Cho, and T. Li, "Solarcore: Solar energy driven multi-core architecture power management," in *Proc. IEEE 17th Int. Symp. High Perform. Comput. Archit.*, 2011, pp. 205–216.

[28] C. Li et al., "Enabling datacenter servers to scale out economically and sustainably," in *Proc. 46th Annu. IEEE/ACM Int. Symp. Microarchit.*, New York, NY, USA, 2013, pp. 322–333, doi: 10.1145/2540708.2540736.

[29] L. A. Barroso, U. Hölzle, and P. Ranganathan, "The datacenter as a computer: Designing warehouse-scale machines," *Synth. Lectures Comput. Archit.*, vol. 13, no. 3, pp. i–189, 2018.

[30] "Imanager neteco6000 data center infrastructure management system," 2022. [Online]. Available: https://digitalpower.huawei.com/en/data-center-facility/product_solution/dce_intelligent_nms/detail/126.html

[31] N. Rasmussen, Electrical efficiency modeling of data centers, 2006. Accessed: Jul. 9, 2024. http://media.comtec.com/pdf/apc/ElectricalEfficiencyModelling.pdf

[32] D. Shin, J. Kim, N. Chang, J. Choi, S. W. Chung, and E.-Y. Chung, "Energy-optimal dynamic thermal management for green computing," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Des. Dig. Tech. Papers*, 2009, pp. 652–657.

[33] T. Kidd, "Power management states: P-states, C-states, and package C-states," *Intel Develop. Resour.*, 2014. Accessed: Jul. 9, 2024. https://www.ilinuxkernel.com/files/CPU.Power/pwr_mgmt_states_r0.pdf

[34] L. Minas and B. Ellison, *Energy Efficiency for Information Technology: How to Reduce Power Consumption in Servers and Data Centers*. Intel Press, 2009.

[35] P. Arroba, J. L. Risco-Martín, M. Zapater, J. M. Moya, J. L. Ayala, and K. Olcoz, "Server power modeling for run-time energy optimization of cloud computing facilities," *Energy Procedia*, vol. 62, pp. 401–410, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S187661021403433X

[36] E. N. M. Elnozahy, M. Kistler, and R. Rajamony, "Energy-efficient server clusters," in *Power-Aware Computer Systems*, B. Falsafi and T. N. Vijaykumar, Eds, Berlin, Germany: Springer, 2003, pp. 179–197.

[37] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *Proc. 34th Annu. Int. Symp. Comput. Archit.*, New York, NY, USA, 2007, pp. 13–23, doi: 10.1145/1250662.1250665.

[38] W. Lin, F. Shi, W. Wu, K. Li, G. Wu, and A.-A. Mohammed, "A taxonomy and survey of power models and power modeling for cloud servers," *ACM Comput. Surv.*, vol. 53, no. 5, Sep. 2020, doi: 10.1145/3406208.

[39] "80 plus certified power supplies and manufacturers," 2022. [Online]. Available: https://www.clearesult.com/80plus/

[40] A. Pratt, P. Kumar, and T. V. Aldridge, "Evaluation of 400v DC distribution in telco and data centers to improve energy efficiency," in *Proc. 29th Int. Telecommun. Energy Conf.*, 2007, pp. 32–39.

[41] R. L. Sawyer, "Making large ups systems more efficient," *Elektron J.-South Afr. Inst. Elect. Eng.*, vol. 23, no. 6, 2006, Art. no. 65.

[42] B. Aksanli, E. Pettis, and T. Rosing, "Architecting efficient peak power shaving using batteries in data centers," in *Proc. IEEE 21st Int. Symp. Modelling Anal. Simul. Comput. Telecommunication Syst.*, 2013, pp. 242–253.

[43] D. Baert and A. Vervaet, "Lead-acid battery model for the derivation of Peukert's law," *Electrochimica Acta*, vol. 44, no. 20, pp. 3491–3504, 1999. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0013468699001176

[44] H. David, E. Gorbatov, U. R. Hanebutte, R. Khanna, and C. Le, "RAPL: Memory power estimation and capping," in *Proc. ACM/IEEE Int. Symp. Low-Power Electron. Des.*, 2010, pp. 189–194.

[45] J. Guo et al., "Who limits the resource efficiency of my datacenter: An analysis of Alibaba datacenter traces," in *Proc. Int. Symp. Qual. Serv.*, New York, NY, USA, 2019, doi: 10.1145/3326285.3329074.

[46] Q. Zhang and W. Shi, "Energy-efficient workload placement in enterprise datacenters," *Computer*, vol. 49, no. 2, pp. 46–52, 2016.

[47] C.-H. Hsu, Q. Deng, J. Mars, and L. Tang, "Smoothoperator: Reducing power fragmentation and improving power utilization in large-scale datacenters," in *Proc. 23rd Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2018, pp. 535–548, doi: 10.1145/3173162.3173190.

[48] K. Haghshenas, A. Pahlevan, M. Zapater, S. Mohammadi, and D. Atienza, "MAGNETIC: Multi-agent machine learning-based approach for energy efficient dynamic consolidation in data centers," *IEEE Trans. Serv. Comput.*, vol. 15, no. 1, pp. 30–44, Jan./Feb. 2022.

[49] C. Gu, Z. Li, H. Huang, and X. Jia, "Energy efficient scheduling of servers with multi-sleep modes for cloud data center," *IEEE Trans. Cloud Comput.*, vol. 8, no. 3, pp. 833–846, Third Quarter, 2020.

[50] F. AL-Hazemi et al., "Dynamic allocation of power delivery paths in consolidated data centers based on adaptive UPS switching," *Comput. Netw.*, vol. 144, pp. 254–270, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128618307138

[51] F. AL-Hazemi, J. Lorincz, A. F. Y. Mohammed, and F. Salamh, "Reducing data center power losses through UPS serial consolidation," in *Proc. Int. Conf. Softw. Telecommun. Comput. Netw.*, 2019, pp. 1–6.

[52] X. Fu, X. Wang, and C. Lefurgy, "How much power oversubscription is safe and allowed in data centers," in *Proc. 8th ACM Int. Conf. Autonomic Comput.*, New York, NY, USA, 2011, pp. 21–30, doi: 10.1145/1998582.1998589.

[53] S. Wu et al., "Precise power capping for latency-sensitive applications in datacenter," *IEEE Trans. Sustain. Comput.*, vol. 6, no. 3, pp. 469–480, Third Quarter, 2021.

[54] S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch, and J. Underwood, "Power routing: Dynamic power provisioning in the data center," in *Proc. 15th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2010, pp. 231–242, doi: 10.1145/1736020.1736047.

[55] Q. Wu et al., "Dynamo: Facebook's data center-wide power management system," in *Proc. 43rd Int. Symp. Comput. Archit.*, 2016, pp. 469–480, doi: 10.1109/ISCA.2016.48.

[56] Y. Li et al., "A scalable priority-aware approach to managing data center server power," in *Proc. IEEE Int. Symp. High Perform. Comput. Archit.*, 2019, pp. 701–714.

[57] V. Sakalkar et al., "Data center power oversubscription with a medium voltage power plane and priority-aware capping," in *Proc. 25th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2020, pp. 497–511, doi: 10.1145/3373376.3378533.

[58] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and limitations of tapping into stored energy for datacenters," in *Proc. 38th Annu. Int. Symp. Comput. Archit.*, 2011, pp. 341–351.

[59] D. Wang, C. Ren, and A. Sivasubramaniam, "Virtualizing power distribution in datacenters," in *Proc. 40th Annu. Int. Symp. Comput. Archit.*, New York, NY, USA, 2013, pp. 595–606, doi: 10.1145/2485922.2485973.

[60] V. D. Reddy, B. Setz, G. S. V. R. K. Rao, G. R. Gangadharan, and M. Aiello, "Metrics for sustainable data centers," *IEEE Trans. Sustain. Comput.*, vol. 2, no. 3, pp. 290–303, Third Quarter, 2017.

[61] S. Malla and K. Christensen, "The effect of server energy proportionality on data center power oversubscription," *Future Gener. Comput. Syst.*, vol. 104, pp. 119–130, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X19304145

[62] H. Cai, Q. Cao, H. Jiang, and Q. Wang, "GreenHetero: Adaptive power allocation for heterogeneous green datacenters," in *Proc. IEEE 41st Int. Conf. Distrib. Comput. Syst.*, 2021, pp. 160–170.

[63] Moving toward $24 \times 7$ carbon-free energy at google data centers: Progress and insights, 2018. [Online]. Available: https://resource-platform.eu/wp-content/uploads/files/statements/Google%201.pdf

[64] T. Van De Voort, V. Zavrel, I. T. Galdiz, and J. Hensen, "Analysis of performance metrics for data center efficiency," *REHVA J.*, vol. 2017, no. February, pp. 37–43, 2017.

[65] Addressing the leading root causes of downtime, Oct. 2012. [Online]. Available: https://gallery.mailchimp.com/454283c8686ebec7a9c03bec9/files/Addressing_Leading_Root_Causes_of_Downtime.pdf

[66] H. Chen, T. N. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding, "Progress in electrical energy storage system: A critical review," *Prog. Natural Sci.*, vol. 19, no. 3, pp. 291–312, 2009. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S100200710800381X

[67] M. Milad and M. Darwish, "Comparison between double conversion online UPS and flywheel UPS technologies in terms of efficiency and cost in a medium data centre," in *Proc. 50th Int. Universities Power Eng. Conf.*, 2015, pp. 1–5.

[68] W. Zheng, K. Ma, and X. Wang, "Hybrid energy storage with supercapacitor for cost-efficient data center power shaving and capping," *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 4, pp. 1105–1118, Apr. 2017.

[69] M. Muccini and W. Cook, "Power consumption reduction: Hot spare," *Dell Tech. White Paper*, 2012.

[70] X. Liu, D. Teng, D. Wang, Q. Zhu, and Z. Liu, "Application of ECO mode UPS in data center," in *Proc. IEEE Int. Telecommun. Energy Conf.*, 2017, pp. 30–34.

[71] G. Ye, F. Gao, and J. Fang, "A mission-driven two-step virtual machine commitment for energy saving of modern data centers through ups and server coordinated optimizations," *Appl. Energy*, vol. 322, 2022, Art. no. 119467. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261922007942

[72] J. Sun, S. Chen, P. You, Q. Yang, and Z. Yang, "Battery-assisted online operation of distributed data centers with uncertain workload and electricity prices," *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 1303–1314, Second Quarter, 2023.

[73] M. A. Lasemi et al., "Energy cost optimization of globally distributed internet data centers by copula-based multidimensional correlation modeling," *Energy Rep.*, vol. 9, pp. 631–644, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352484722026348

[74] N. Nasiriani and G. Kesidis, "Optimal peak shaving using batteries at datacenters: Charging risk and degradation model," in *Proc. Int. Conf. Comput. Netw. Commun.*, 2018, pp. 58–62.

[75] M. Dabbagh, B. Hamdaoui, A. Rayes, and M. Guizani, "Shaving data center power demand peaks through energy storage and workload shifting control," *IEEE Trans. Cloud Comput.*, vol. 7, no. 4, pp. 1095–1108, Fourth Quarter, 2019.

[76] H. Chen, Y. Zhang, M. C. Caramanis, and A. K. Coskun, "EnergyQARE: QoS-aware data center participation in smart grid regulation service reserve provision," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 4, no. 1, pp. 1–31, Jan. 2019, doi: 10.1145/3243172.

[77] K. Wang et al., “A hierarchical dispatch strategy of hybrid energy storage system in internet data center with model predictive control,” *Appl. Energy*, vol. 331, 2023, Art. no. 120414. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261922016713
[78] Y. Shi, B. Xu, D. Wang, and B. Zhang, “Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains,” *IEEE Trans. Power Syst.*, vol. 33, no. 3, pp. 2882–2894, May 2018.
[79] Z. Xu, H. Wang, Z. Xu, and X. Wang, “Power attack: An increasing threat to data centers,” in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2014.
[80] S. Li et al., “Thunderbolt: Throughput-Optimized, Quality-of-Service-Aware power capping at scale,” in *Proc. 14th USENIX Symp. Operating Syst. Des. Implementation*, 2020, pp. 1241–1255. [Online]. Available: https://www.usenix.org/conference/osdi20/presentation/li-shaohong
[81] P. Patel et al., “Characterizing power management opportunities for LLMs in the cloud,” in *Proc. 29th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2024, pp. 207–222, doi: 10.1145/3620666.3651329.
[82] B. Luo, W. Chen, X. Liu, X. Li, L. Zhang, and W. Shi, “Pelican: Power scheduling for QoS in large-scale data centers with heterogeneous workloads,” in *Proc. 10th Int. Green Sustain. Comput. Conf.*, 2019, pp. 1–8.
[83] N. Kulkarni, G. Gonzalez-Pumariega, A. Khurana, C. A. Shoemaker, C. Delimitrou, and D. H. Albonesi, “CuttleSys: Data-driven resource management for interactive services on reconfigurable multicores,” in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2020, pp. 650–664.
[84] L. Piga et al., “Expanding datacenter capacity with DVFS boosting: A safe and scalable deployment experience,” in *Proc. 29th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2024, pp. 150–165, doi: 10.1145/3617232.3624853.
[85] R. Azimi, M. Badiei, X. Zhan, N. Li, and S. Reda, “Fast decentralized power capping for server clusters,” in *Proc. IEEE Int. Symp. High-Perform. Comput. Archit.*, Los Alamitos, CA, USA, 2017, pp. 181–192. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/HPCA.2017.49
[86] C. Zhang et al., “Flex: High-availability datacenters with zero reserved power,” in *Proc. ACM/IEEE 48th Annu. Int. Symp. Comput. Archit.*, 2021, pp. 319–332.
[87] A. G. Kumbhare et al., “Prediction-Based power oversubscription in cloud platforms,” in *Proc. USENIX Annu. Tech. Conf.*, 2021, pp. 473–487. [Online]. Available: https://www.usenix.org/conference/atc21/presentation/kumbhare
[88] P. Pang, Q. Chen, D. Zeng, and M. Guo, “Adaptive preference-aware co-location for improving resource utilization of power constrained datacenters,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 2, pp. 441–456, Feb. 2021.
[89] G. Wang et al., “Increasing large-scale data center capacity by statistical power control,” in *Proc. 11th Eur. Conf. Comput. Syst.*, New York, NY, USA, 2016, pp. 1–15, doi: 10.1145/2901318.2901338.
[90] J. Sheng et al., “Learning cooperative oversubscription for cloud by chance-constrained multi-agent reinforcement learning,” in *Proc. ACM Web Conf.*, New York, NY, USA, 2023, pp. 2927–2936, doi: 10.1145/3543507.3583298.
[91] B. Aksanli, “Data center peak power management with energy storage devices,” *IEEE Internet Comput.*, vol. 21, no. 4, pp. 26–33, 2017.
[92] S. Li, S. Huneycutt, S. Wang, S. Gu, C. Pan, and T. Abdelzaher, “WattValet: Heterogenous energy storage management in data centers for improved power capping,” in *Proc. 11th Int. Conf. Autonomic Comput.*, 2014, pp. 273–279. [Online]. Available: https://www.usenix.org/conference/icac14/technical-sessions/presentation/li_shen
[93] S. Alanazi, M. Dabbagh, B. Hamdaoui, M. Guizani, and N. Zorba, “Reducing data center energy consumption through peak shaving and locked-in energy avoidance,” *IEEE Trans. Green Commun. Netw.*, vol. 1, no. 4, pp. 551–562, Dec. 2017.
[94] S. Malla et al., “Coordinated priority-aware charging of distributed batteries in oversubscribed data centers,” in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2020, pp. 839–851.
[95] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, “Power attack defense: Securing battery-backed data centers,” in *Proc. 43rd Int. Symp. Comput. Archit.*, 2016, pp. 493–505, doi: 10.1109/ISCA.2016.50.
[96] X. Hou, C. Li, J. Yang, W. Zheng, X. Liang, and M. Guo, “Integrated power anomaly defense: Towards oversubscription-safe data centers,” *IEEE Trans. Cloud Comput.*, vol. 10, no. 3, pp. 1875–1887, Third Quarter, 2022.
[97] X. Hou, M. Liang, C. Li, W. Zheng, Q. Chen, and M. Guo, “When power oversubscription meets traffic flood attack: Re-thinking data center peak load management,” in *Proc. 48th Int. Conf. Parallel Process.*, New York, NY, USA, 2019, pp. 1–10, doi: 10.1145/3337821.3337856.
[98] X. Hou, L. Hao, C. Li, Q. Chen, W. Zheng, and M. Guo, “Power grab in aggressively provisioned data centers: What is the risk and what can be done about it,” in *Proc. IEEE 36th Int. Conf. Comput. Des.*, 2018, pp. 26–34.
[99] B. Acun et al., “Carbon explorer: A holistic framework for designing carbon aware datacenters,” in *Proc. 28th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2023, pp. 118–132, doi: 10.1145/3575693.3575754.
[100] T. Sukprasert, A. Souza, N. Bashir, D. Irwin, and P. Shenoy, “On the limitations of carbon-aware temporal and spatial workload shifting in the cloud,” in *Proc. 19th Eur. Conf. Comput. Syst.*, New York, NY, USA, 2024, pp. 924–941, doi: 10.1145/3627703.3650079.
[101] S. Jie, S. Zong-Zhe, L. Hui, B. Yu-Bin, and Y. Ge, “Research advance on energy consumption optimizatino of hyper-powered data center,” *Chin. J. Comput.*, vol. 41, no. 12, pp. 2670–2688, 2018.
[102] L. X. Ye Ran and L. Chao, “Energy storage system in green data centers: Architecture and management,” *J. Comput. Res. Develop.*, vol. 53, no. 2, 2016, Art. no. 326. [Online]. Available: https://crad.ict.ac.cn/EN/abstract/article_3105.shtml
[103] J. Gao, H. Wang, and H. Shen, “Smartly handling renewable energy instability in supporting a cloud datacenter,” in *Proc. IEEE Int. Parallel Distrib. Process. Symp.*, 2020, pp. 769–778.
[104] H. Shen, H. Wang, J. Gao, and R. Buyya, “An instability-resilient renewable energy allocation system for a cloud datacenter,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 34, no. 3, pp. 1020–1034, Mar. 2023.
[105] J. Xing et al., “Carbon responder: Coordinating demand response for the datacenter fleet,” 2023, *arXiv:2311.08589*.
[106] T. Chakraborty, A. N. Toosi, and C. Kopp, “Elastic power utilization in sustainable micro cloud data centers,” *IEEE Trans. Sustain. Comput.*, vol. 8, no. 3, pp. 465–478, Third Quarter, 2023.
[107] W. Liu et al., “Online job scheduling scheme for low-carbon data center operation: An information and energy nexus perspective,” *Appl. Energy*, vol. 338, 2023, Art. no. 120918. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261923002829
[108] A. Radovanović et al., “Carbon-aware computing for datacenters,” *IEEE Trans. Power Syst.*, vol. 38, no. 2, pp. 1270–1280, Mar. 2023.
[109] L. Liu, H. Sun, C. Li, T. Li, J. Xin, and N. Zheng, “Exploring highly dependable and efficient datacenter power system using hybrid and hierarchical energy buffers,” *IEEE Trans. Sustain. Comput.*, vol. 6, no. 3, pp. 412–426, Third Quarter, 2021.
[110] S. Kwon, “QoS-aware data center operations based on chance- and risk-constrained optimization,” *IEEE Trans. Cloud Comput.*, vol. 10, no. 4, pp. 2887–2896, Fourth Quarter, 2022.
[111] C. Gu, W. Li, C. Shen, and S. Cui, “Greenflowing: A green way of reducing electricity cost for cloud data center using heterogeneous ESDs,” in *Proc. IEEE Symp. Comput. Commun.*, 2019, pp. 1–6.
[112] T. Yang, Y. Hou, Y. C. Lee, H. Ji, and A. Y. Zomaya, “Power control framework for green data centers,” *IEEE Trans. Cloud Comput.*, vol. 10, no. 4, pp. 2876–2886, Fourth Quarter, 2022.
[113] X. Liu, Y. Hua, X. Liu, L. Yang, and Y. Sun, “Design and implementation of smooth renewable power in cloud data centers,” *IEEE Trans. Cloud Comput.*, vol. 11, no. 1, pp. 85–96, First Quarter, 2023.
[114] K. Zhou, Z. Fei, and X. Lu, “Optimal energy management of internet data center with distributed energy resources,” *IEEE Trans. Cloud Comput.*, vol. 11, no. 3, pp. 2285–2295, Third Quarter, 2023.
[115] J. Thiede, N. Bashir, D. Irwin, and P. Shenoy, “Carbon containers: A system-level facility for managing application-level carbon emissions,” in *Proc. ACM Symp. Cloud Comput.*, New York, NY, USA, 2023, pp. 17–31, doi: 10.1145/3620678.3624644.
[116] N. Bashir et al., “Enabling sustainable clouds: The case for virtualizing the energy system,” in *Proc. ACM Symp. Cloud Comput.*, New York, NY, USA, 2021, pp. 350–358, doi: 10.1145/3472883.3487009.
[117] A. Souza et al., “Ecovisor: A virtual energy system for carbon-efficient applications,” in *Proc. 28th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, New York, NY, USA, 2023, pp. 252–265, doi: 10.1145/3575693.3575709.
[118] D. Sharma and S. Rao, “Scheduling computing loads for improved utilization of solar energy,” *Sustain. Comput., Inform. Syst.*, vol. 32, 2021, Art. no. 100592. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210537921000810

[119] K. Kaur, S. Garg, G. S. Aujla, N. Kumar, and A. Y. Zomaya, "A multi-objective optimization scheme for job scheduling in sustainable cloud data centers," *IEEE Trans. Cloud Comput.*, vol. 10, no. 1, pp. 172–186, First Quarter, 2022.
[120] A. C. Ammari, W. Labidi, F. Mnif, H. Yuan, M. Zhou, and M. Sarrab, "Firefly algorithm and learning-based geographical task scheduling for operational cost minimization in distributed green data centers," *Neurocomputing*, vol. 490, pp. 146–162, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925231222000716
[121] S. Khalid and I. Ahmad, "Dual optimization of revenue and expense in geo-distributed data centers using smart grid," *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 1622–1635, Second Quarter, 2023.
[122] D. Zhao, J.-T. Zhou, and K. Li, "CFWS: DRL-based framework for energy cost and carbon footprint optimization in cloud data centers," *IEEE Trans. Sustain. Comput.*, vol. 10, no. 1, pp. 95–107, Jan./Feb. 2025.
[123] T. Shi, H. Ma, G. Chen, and S. Hartmann, "Auto-scaling containerized applications in geo-distributed clouds," *IEEE Trans. Serv. Comput.*, vol. 16, no. 6, pp. 4261–4274, Nov./Dec. 2023.
[124] A. Jayanetti, S. Halgamuge, and R. Buyya, "Multi-agent deep reinforcement learning framework for renewable energy-aware workflow scheduling on distributed cloud data centers," *IEEE Trans. Parallel Distrib. Syst.*, vol. 35, no. 4, pp. 604–615, Apr. 2024.
[125] T. Yang, H. Jiang, Y. Hou, and Y. Geng, "Carbon management of multi-datacenter based on spatio-temporal task migration," *IEEE Trans. Cloud Comput.*, vol. 11, no. 1, pp. 1078–1090, First Quarter, 2023.
[126] D. Cheng, X. Zhou, Z. Ding, Y. Wang, and M. Ji, "Heterogeneity aware workload management in distributed sustainable datacenters," *IEEE Trans. Parallel Distrib. Syst.*, vol. 30, no. 2, pp. 375–387, Feb. 2019.
[127] A. Pahlevan, M. Zapater, A. K. Coskun, and D. Atienza, "ECOGreen: Electricity cost optimization for green datacenters in emerging power markets," *IEEE Trans. Sustain. Comput.*, vol. 6, no. 2, pp. 289–305, Second Quarter, 2021.
[128] Z. Nadalizadeh and M. Momtazpour, "Greenpacker: Renewable-and fragmentation-aware VM placement for geographically distributed green data centers," *J. Supercomput.*, vol. 78, no. 1, pp. 1434–1457, 2022.
[129] T. Yang, Y. Hou, S. Cai, J. Yu, and H. Pen, "Multi-data center tie-line power smoothing method based on demand response," *IEEE Trans. Cloud Comput.*, vol. 12, no. 4, pp. 983–995, Fourth Quarter, 2024.
[130] C. Guo, F. Luo, J. Yang, and Z. Cai, "Transactive operational framework for internet data centers in geo-distributed local energy markets," *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 1133–1143, Second Quarter, 2023.
[131] S. Sarkar et al., "Carbon footprint reduction for sustainable data centers in real-time," in *Proc. AAAI Conf. Artif. Intell.*, 2024, pp. 22 322–22 330. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/30238
[132] G. Zhang, S. Zhang, W. Zhang, Z. Shen, and L. Wang, "Distributed energy management for multiple data centers with renewable resources and energy storages," *IEEE Trans. Cloud Comput.*, vol. 10, no. 4, pp. 2469–2480, Fourth Quarter, 2022.
[133] G. Ye, F. Gao, J. Fang, and Q. Zhang, "Joint workload scheduling in geo-distributed data centers considering UPS power losses," *IEEE Trans. Ind Appl.*, vol. 59, no. 1, pp. 612–626, Jan./Feb. 2023.
[134] V. J. Reddi et al., "MLPerf inference benchmark," in *Proc. ACM/IEEE 47th Annu. Int. Symp. Comput. Archit.*, 2020, pp. 446–459.

**Huikang Huang** received the MS degree in computer science and theory from the South China Agricultural University, Guangzhou, China, in 2021. Now, he is working toward the PhD degree with the School of Computer Science and Engineering, South China University of Technology. His research interests mainly focus on cloud computing and reinforcement learning.

**Weiwei Lin** (Senior Member, IEEE) received the BS and MS degrees from Nanchang University in 2001 and 2004, respectively, and the PhD degree in computer application from the South China University of Technology in 2007. Currently, he is a professor with the School of Computer Science and Engineering, South China University of Technology. His research interests include distributed systems, cloud computing, and AI application technologies. He has published more than 150 papers in refereed journals and conference proceedings. He has been a reviewer for many international journals, including *IEEE Transactions on Parallel and Distributed Systems*, *IEEE Transactions on Services Computing*, *IEEE Transactions on Cloud Computing*, *IEEE Transactions on Computers*, *IEEE Transactions on Cybernetics*, etc. He is a distinguished member of CCF.

**Jianpeng Lin** received the MS degree in computer science from the Guangdong University of Technology, in 2019. He is currently working toward the PhD degree with the School of Computer Science and Engineering, South China University of Technology, Guangzhou, China. His research interests include cloud computing, thermal management, and deep reinforcement learning.

**Keqin Li** (Fellow, IEEE) is a SUNY distinguished professor of computer science with the State University of New York. He is also a National distinguished professor with Hunan University, China. His current research interests include cloud computing, fog computing and mobile edge computing, energy-efficient computing and communication. He has authored or coauthored more than 870 journal articles, book chapters, and refereed conference papers. He is among the world's top 5 most influential scientists in parallel and distributed computing in terms of both single-year impact and career-long impact based on a composite indicator of Scopus citation data-base. He is currently an associate editor of the *ACM Computing Surveys* and the *CCF Transactions on High Performance Computing*. He has served on the editorial boards of *IEEE Transactions on Parallel and Distributed Systems*, *IEEE Transactions on Computers*, *IEEE Transactions on Cloud Computing*, the *IEEE Transactions on Services Computing*, and *IEEE Transactions on Sustainable Computing*. He is an AAIA fellow. He is also a member of Academia Europaea (Academician of the Academy of Europe).