COPA: Highly Cost-Effective Power Back-Up for Green Datacenters

Yan Yin, Junmin Wu, Xu Zhou, Lieven Eeckhout, *Fellow, IEEE*, Amer Qouneh, Tao Li, *Fellow, IEEE*, and Zhibin Yu, *Member, IEEE*

Abstract—Traditional datacenters employ costly diesel generators (DG) and uninterrupted power supplies (UPS) to back up power. However, some or even all racks of a *green* datacenter can still be powered by renewable energy during grid power outages. This makes the utilization of the DGs and UPSs in green datacenters significantly lower than in traditional datacenters. In this paper, we propose a highly cost-effective power back-up (COPA) approach for green datacenters by leveraging the availability characteristics of renewable energy as well as grid power outages. COPA contributes three new techniques. The first technique, called least UPS capacity planning, determines the least rated power capability and runtime of the UPSs to guarantee the normal operations of a green datacenter during grid power outages. The second technique, named cooperative UPS/renewable power supply, employs UPS and renewable energy at the same time to supply power to *each* rack when grid power fails. The last one, dubbed renewable-energy-aware dynamic power management, controls the power consumption dynamically based on the available capacity of renewable energy and UPS. We build an experimental cluster consisting of 10 servers, and use four representative benchmarks as well as verified data about the availability characteristics of solar and wind energy to evaluate COPA. The results show that COPA reduces 47 percent and 70 percent of the power back-up cost for a solar energy powered datacenter and a wind energy powered datacenter, respectively. Moreover, COPA guarantees the application's Service Level Agreement (SLA) for at least 20 minutes (over 79 percent outages) and 56 minutes on average while enabling the back-up power to last for at least 2 hours and for 3 hours on average, which cannot be achieved by other under-provisioning power back-up approaches.

Index Terms—Datacenter power backup, renewable energy, cooperative power distribution, dynamic power management

1 Introduction

The huge amount of grid energy consumed by modern datacenters is exacerbating both energy crisis and global warming because fossil fuel is currently the main source of grid power. On the one hand, a recent report shows that the energy consumption of U.S. datacenters is expected to reach approximately 73 billion kilowatt-hour (KWh) in 2020 [1]. On the other hand, as datacenter footprints continue to expand, greenhouse gas emissions would exceed 100 million tons per year if fossil fuel continues to be the main source for generating energy [2], [3].

Faced by such a burdensome case, Internet giants including Apple [4], Google [5] and Facebook [6] started to use renewable energy powered datacenters (called green datacenters) [7]. Moreover, these giants are still rapidly increasing renewable energy investments in their datacenters [8]. However, due to its instability and intermittence, renewable energy is typically used as a supplement of grid power in a green datacenter.

Power back-up is indispensable for datacenters because grid power outages are unavoidable. Most traditional datacenters employ diesel generators (DG) and uninterrupted power supplies (UPS) to sustain operations when grid power fails [10]. A charged UPS can supply power immediately but it needs to recharge after a limited time interval. A DG, on the contrary, can last for a long time but cannot supply power immediately because it needs time (e.g., 2 minutes) to boot up. Therefore, UPSs typically supply power for a datacenter before DGs boot up. Such power back-up expenditure is over 20 percent of the power infrastructure cost of a datacenter [11], [12]. Consequently, the direct cost of UPSs and DGs amounts to 2.68 millions of dollars per year for a 20 megawatt (MW) datacenter [13].

However, the utilization of such costly UPSs and DGs is extremely low because they are power back-up facilities which are naturally idle most of time. On the one hand, U.S. datacenters experience an average of only 2.66 grid power failures per year according to a recent survey [9]. On the other hand, Fig. 1 shows that 79 and 90 percent of datacenter power outages are shorter than 20 minutes and 2 hours,

- Y. Yin and J. Wu are with the Department of Computer Science, University of Science and Technology of China, Hefei, Anhui 230026, China. E-mail: yy1001@mail.ustc.edu.cn, jmwu@ustc.edu.cn.
- X. Zhou is with the Sangfor Technologies Inc., Shenzhen Institutes of Advanced Technology, Chinese Academy of Science, Shenzhen, Guangdong 518055, China. E-mail: sunupzhou@gmail.com.
- L. Eeckhout is with the Ghent University 9000 Gent, Belgium. E-mail: lieven.eeckhout@ugent.be.
- A. Qouneh is with New Western England University, Springfield, MA 01119. E-mail: aqouneh@hotmail.com.
- T. Li is with the Shenzhen Institutes of Advanced Technology, Chinese Academy of Science, Shenzhen, Guangdong 518055, China. E-mail: tao.li@siat.ac.cn.
- Z. Yu is with the Shenzhen Institutes of Advanced Technology, Chinese Academy of Science, Shenzhen, Guangdong 518055, China. E-mail: zb.yu@siat.ac.cn.

Manuscript received 12 May 2019; revised 28 Aug. 2019; accepted 3 Oct. 2019. Date of publication 21 Oct. 2019; date of current version 10 Jan. 2020.
(Corresponding author: Zhibin Yu.)
Recommended for acceptance by J. Zhai.
Digital Object Identifier no. 10.1109/TPDS.2019.2948336

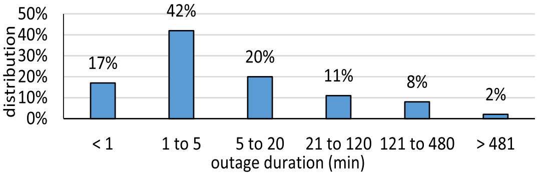

Fig. 1. The frequency and average duration of grid power outages of U.S. datacenters [9].

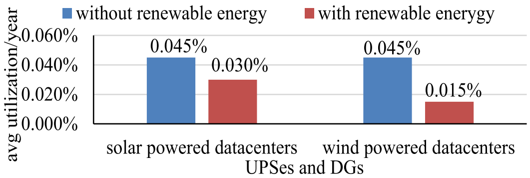

Fig. 2. The comparison of averaged utilizations of UPSs and DGs in datacenters with and without renewable energy (normalized to utilization of the grid power) [9], [16], [17]. Since the renewable energy can power some, or all servers when unplanned grid power fails, the UPSs and DGs utilization of green datacenters is lower.

respectively. This indicates that 79 and 90 percent of datacenters only use their UPSs and DGs for at most 52 minutes and 5.2 hours on average per year, respectively.

When renewable energy is employed in a datacenter, the utilization of the UPSs and DGs becomes even significantly lower, as shown in Fig. 2. The reason is that renewable energy may supply power for some racks when the grid power fails. As such, the UPSs and DGs are not fully utilized even during a grid power outage. Moreover, both UPSs and DGs suffer from aging, which requires them to be replaced after a fixed time even though they are not used (e.g., usually 4 years for acid-lead UPSs and 12 years for DGs [14]). This results in extremely low cost-effectiveness of the power back-up facilities of a green datacenter. To improve the cost-effectiveness, there are three questions to be answered. (1) Are DGs still needed? (2) Given the existing renewable energy, what are the suitable power capacity and runtime of UPSs? (3) What is the power supply strategy (e.g., what is the power back-up share between renewable energy and UPSs)? Although prior work [13] proposed a cost-performance-availability strategy for different levels of power back-up under-provisioning, it does not consider renewable energy in green datacenters. Zhou et al. [15] propose that renewable energy can be used to under-provision the grid power infrastructure. However, they do not consider power back-up cost with renewable energy.

In this paper, we answer these three questions by proposing a highly COst-effective Power bAck-up (COPA) approach that leverages the availability characteristics of renewable energy as well as grid power outages of a green datacenter. COPA removes DGs and contributes three novel techniques. The first technique, called *least UPS capacity planning*, determines the least rated power capacity and runtime, which are the two key parameters, of a UPS that can guarantee the normal operations of a datacenter during grid power outages. The second technique, called *cooperative UPS/renewable power supply*, employs UPS and renewable energy at the same time to power *each* rack during grid power outages. The third technique, called *renewable-energy-aware dynamic power management*, dynamically controls the power consumption based on the available capacity of renewable energy and UPSs. It combines different technologies including power throttling [18], [19], [20] and workload migration [21], [22] for different levels of renewable energy and UPS capacity. The first technique reduces the power back-up cost directly while the second and third techniques decrease the cost indirectly by sustaining the power as long as possible during grid power outages.

In particular, we make the following contributions:

(1) We propose a technique called *least UPS capacity planning* to leverage the availability characteristics of renewable energy and grid power outage characteristics of datacenters to significantly improve the cost-effectiveness of power back-up facilities.

(2) Under the constraint of high cost-effectiveness, we develop two techniques: *cooperative UPS/renewable power supply* and *renewable-energy-aware dynamic power management* to supply power as long as possible during grid power outages, further improving the cost-effectiveness.

(3) We called the proposed techniques COPA, which can reduce 47 and 70 percent of power back-up costs for a solar and a wind energy powered datacenters, respectively. Furthermore, COPA lasts for at least 2 hours and on average 3 hours, which cannot be achieved by other power back-up techniques.

The rest of this paper is organized as follows: Section 2 provides background and motivation. Section 3 depicts the design of COPA. Section 4 presents the experimental setup. Section 5 provides the results and analysis. Section 6 describes related work and Section 7 concludes the paper.

2 Background and Motivation

2.1 Power Back-Up Facilities of a Datacenter

The primary power supply of a traditional datacenter is the grid power. As shown in Fig. 3, line Ⓐ from the grid connects to a number of Power Distribution Units (PDU) providing power to the racks. Datacenters are also equipped with UPSs and DGs which act as a secondary power supply during grid power outages. The distributed UPS units track the supply voltage and frequency variations of the corresponding racks. When a grid power failure occurs, UPSs can quickly turn on their internal DC-AC inverter and supply power to the servers. Meanwhile, an Automatic Transfer Switch (ATS) detects grid failures and

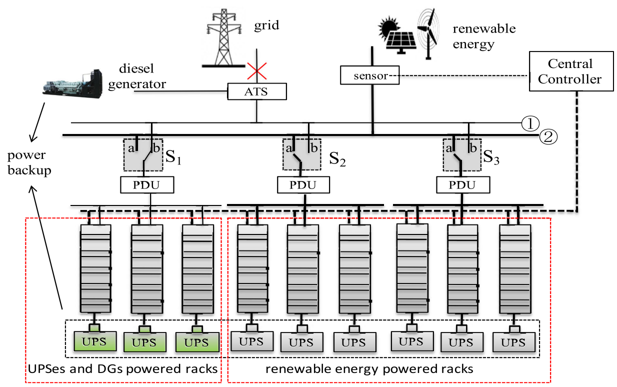

Fig. 3. Power backup facilities of a datacenter.

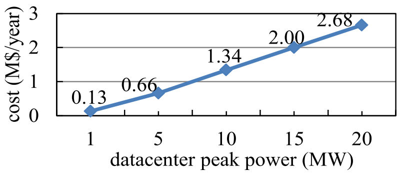

Fig. 4. Annual cost of primary grid power outage.

subsequently switches the power load from the grid to the DGs. It takes about 2 minutes to boot up the DGs and switch the power supply to servers from UPSs to DGs. Subsequently, the DGs power the whole datacenter independently and the UPSs are switched to standby mode by turning off the inverter because power supply voltage level has returned to normal.

When renewable energy is employed in a datacenter, its energy assumes a higher priority than grid power to be fully utilized by the racks, as shown in Fig. 3. The Central Controller (CC) can get the status of the renewable energy periodically through a sensor and determines the number of renewable-energy-powered racks by controlling the relays (S_1, S_2, and S_3) over the Ethernet. Renewable energy supply is not influenced by grid failures because renewable power employs a different transmission line (line ②) from that (line Ⓐ) for the grid power. Therefore, during grid outages, the CC obtains the available renewable power and switches some relays such as S_2 and S_3 to terminal a to let renewable energy power some racks. The UPSs of these racks are still in standby mode. However, UPSs and DGs now power the other racks, previously powered by the grid, by switching S_1 to terminal b. It is obvious that the utilization of DGs and UPSs becomes lower when using renewable energy.

2.2 High Cost of Power Back-Up Solutions

Fig. 4 shows the annual amortized cost for power back-up infrastructure for different datacenter capacities [13]. As can be seen, power back-up cost increases proportionally with the power capacity of datacenters. For instance, for a 20 MW (megawatt) datacenter, the power back-up cost amounts to 2.68 million dollars per year. Moreover, it is predicted that the power of most datacenters can reach up to 100 MW [23], [24] in the near future. This indicates that the power back-up cost of future datacenters would be extremely high. When renewable energy is employed, the back-up cost is further increased. Assuming that the power back-up cost (C_{pb}) can previously handle n grid outages per year. With renewable energy, it only needs to handle m ($m \leq n$) grid power outages per year. The cost therefore relatively increases by $C_{pb}/m - C_{pb}/n$. Therefore, reducing the power back-up cost of datacenters is urgently needed.

2.3 Datacenter Availability Level

Four datacenter availability levels (Tier 1 to 4) have currently been established [11]. Tier 1 is the lowest level while Tier 4 is the highest, with complete DG and UPS systems. Tier 1 specifies the annual grid outage up to 28.8 hours; Tier 2 specifies 22 hours; Tier 3 specifies 1.6 hours; and Tier 4 specifies only 0.4-hour of annual grid outage, or 99.995 percent availability. The tier level drives the design specifications for new datacenters. In the financial industry, high availability and fault

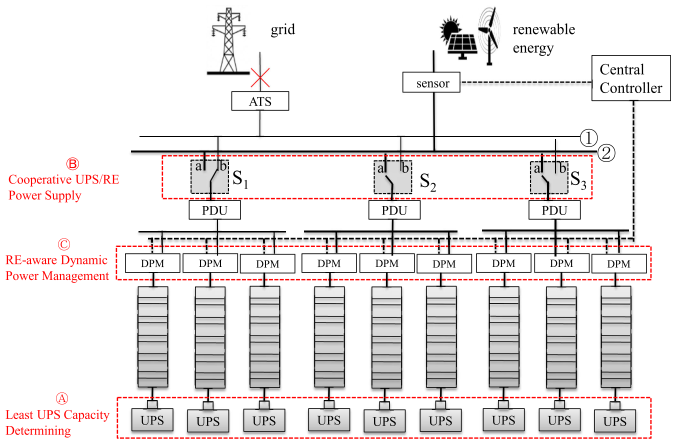

Fig. 5. The COPA architecture and its three key techniques: Ⓐ Least UPS Capacity Planning, Ⓑ Cooperative UPS/RE (renewable energy) Power Supply, and Ⓒ RE-aware Dynamic Power Management (be short for DPM in the figure).

tolerance (that is, typically Tier 3 and Tier 4) are required to support 24×7 financial transaction activities and funds exchanges. For other organizations (such as universities), a lower fault tolerance is acceptable.

3 Design Principles

3.1 COPA Architecture

COPA is a highly cost-effective power back-up approach for green datacenters. COPA aims at a better trade-off between power back-up cost, performance, and availability of datacenters compared to previous solutions. The COPA architecture is shown in Fig. 5. To decrease the impact of voltage transients, frequency distortions, and harmonics on the grid, we directly connect renewable power supplies to PDUs. In other words, we supply power for the racks by the two separated lines (one for grid power and one for renewable power). We adopt distributed UPS design (one UPS for each rack) which is popular in today's datacenters [25], [26] due to its reliability and cost advantage over conventional centralized UPS placement [27].

COPA removes DGs and only employs UPSs for two reasons: (1) The cost of a DG is high and increases almost linearly with its peak power [13]. In contrast, the cost of a UPS rises very slowly with provisioned runtime. For instance, a $20\times$ increase in UPS runtime (from 2 mins to 42 mins) leads to an overall cost increment of only 24 percent (from 1.34 to 1.66 million dollars) for a 10 MW datacenter [13]. (2) Unlike UPS that can supply power immediately, a DG needs tens of seconds to boot up. That is, removing UPSs would make datacenters unavailable during a grid power outage, but removing DGs would not. Therefore, it is feasible as well as cost-effective to remove DGs and only use larger UPSs.

COPA consists of three key techniques as shown in Fig. 5: *least UPS capacity planning*, *cooperative UPS/renewable power supply*, and *renewable-energy-aware dynamic power management*. Least UPS capacity planning determines the least rated power capacity and runtime, which are the two key parameters, of a UPS that can guarantee the normal operations of a datacenter during grid power outages. COPA adaptively switches the power supply between renewable energy and UPSs during a grid power outage. As such, the

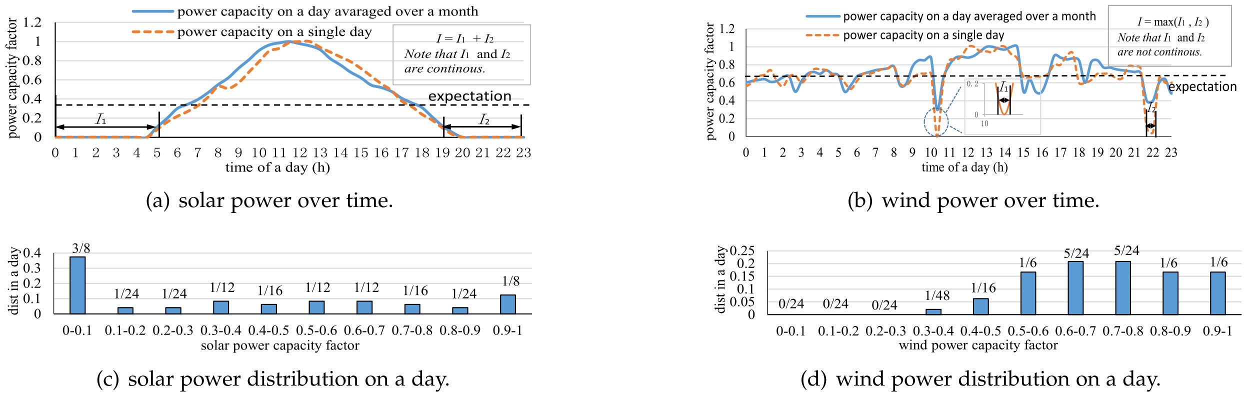

Fig. 6. Solar and wind power on a single day (orange line) and on a day averaged over a month (blue line) for the site #726,798 (United States Air Force number) [28]. The power capacity factor is defined as the real power output divided by the rated power.

cases of power supply can be: (1) renewable energy is abundant and can power the whole datacenter independently. The excess energy can charge the UPSs. (2) The renewable power is insufficient and UPSs need to supplement power. (3) UPSs power the servers independently when renewable power is unavailable. Least UPS capacity planning is designed to configure the exact UPS capacity to effectively cooperate with renewable energy and to reduce the cost of power back-up facilities.

Note that datacenters without DGs may go down due to insufficient power (e.g., the (2) and (3) cases mentioned above) during long-time grid power outages. However, it is acceptable for some datacenters with low fault tolerance (e.g., datacenters for science computation) [11]. For high-availability datacenters, we can configure large UPS capacity according to the actual situation. In this paper, we consider highly cost-effective power back-up configuration in green datacenters and we also propose two techniques to improve the availability and performance of green datacenters as much as possible.

Cooperative UPS/renewable power supply and renewable energy aware dynamic power management address the challenges caused by the time-varying and intermittent characteristics of renewable energy, as well as the limited UPS capacities. Cooperative UPS/renewable power supply employs UPS and renewable energy at the same time to power *each* rack when grid power fails, which leverages the battery discharge characteristics to prolong the runtime of UPSs. Renewable energy aware dynamic power management dynamically controls the power consumption based on the available capacity of renewable energy and UPSs. It combines several technologies such as changing the CPU power state and triggering workload migration by considering the applications' characteristics to keep servers running. In summary, least UPS capacity planning reduces the back-up cost explicitly while cooperative UPS/renewable power supply and renewable energy aware dynamic power management decrease the cost implicitly.

3.2 Least UPS Capacity Planning

3.2.1 Overview

Planning the UPS capacity aiming for high cost-effectiveness is very challenging. First, too large UPS capacity results in low utilization and in turn low cost-effectiveness whereas too small UPS capacity reduces the performance and availability of a datacenter. Second, the grid power outage is inherently unpredictable and renewable energy is generally unstable. Least UPS capacity planning addresses these challenges by leveraging the availability features of renewable energy and grid power outage characteristics.

3.2.2 Pre-Analysis for Planning UPS Capacity

Two factors generally determine the rated capacity of a UPS: power capacity (P_{ups}) and runtime (RL_{ups}). The rated P_{ups} needs to satisfy the power demand for running the datacenter workloads normally. The RL_{ups} refers to the longest duration when a UPS operates normally. Further, the less workload (that is, less than rated P_{ups}) powered by a UPS, the longer the UPS can last (details in Section 3.3). Green datacenters typically prefer to use renewable energy to power their workloads. But if it is insufficient, the UPS would supplement the difference when a grid power outage occurs. Therefore, the availability feature of renewable energy and grid power outage characteristics have a vital impact on UPS capacity.

For the RL_{ups}, although the actual runtime of a UPS varies with the time-varying renewable energy, we only consider the rated RL_{ups}. This indicates the UPS can supply power RL_{ups} hours for the datacenter with normal operations, which is influenced by the outage characteristics of grid power. Grid power outages can be characterized by two factors: outage duration (T_o) and outage duration distribution (D_o). Fig. 1 shows these two factors. For example, the 1-5 min and 5-20 min outages account for 42 and 20 percent of all grid power outages, respectively.

As for the P_{ups}, it is mainly influenced by renewable energy because a UPS is activated to supplement *insufficient* renewable power. The renewable power can also be characterized by two factors: available renewable power (P_{re}) and power availability distribution (D_{re}). We leverage the verified data of the National Solar Radiation Database (NSRDB) provided by National Renewable Energy Laboratory (NREL [28]) to analyze the solar and wind power availability. Figs. 6a and 6b show the solar and wind power of a day averaged over a month for the site #726,798 (United States Air Force number), respectively. We employ *power capacity factor* (*real-power/rated-power*) [17] to represent the available power of a renewable energy equipment for each hour during a day. From these figures, we observe the intermittency and variation of solar and wind power, which makes the determination of P_{re} and D_{re} challenging. Moreover, Figs. 6a and 6b show that the power capacity factor of solar energy is significantly different from that of the wind energy, indicating that the UPS configurations for solar and wind energy should be different.

Fig. 6c shows the accumulated time for each power capacity factor range for solar energy in a day, which intuitively reflects P_{re} and D_{re}. As can be seen, in 9 hours ($24 \times (3/8)$) of a day, the solar power capacity factor is less than 0.1. This indicates the solar energy is almost unavailable for 37.5 percent of a day. We call this duration the *intermittence duration* ($I = I_1 + I_2$, as shown in Fig. 6a) of solar energy. On the contrary, there are only 3 hours when power capacity factor is larger than 0.9, which can support the datacenter to operate normally in practice. Fig. 6d shows that wind power is relatively stable and its power capacity factor is larger than 0.5 for most of the day, which can provide more power than the solar energy. Note that the intermittence duration of wind energy (e.g., 20 minutes) is much shorter than that of solar energy, as shown in Figs. 6a and 6b. This indicates that UPS power capacity should be smaller in case of wind energy compared to solar energy for highly cost-effective power backup facilities.

3.2.3 Planning UPS Capacity

We now determine the two factors of UPS capacity for a datacenter equipped with renewable energy: rated power (P_{ups}) and runtime (RL_{ups}). We use the *mathematical expectation* (ME) of renewable power and grid power outage durations to determine UPS capacity. In statistics, ME is theoretically approaching to the mean value of a random variable if the random experiment generating the variable's values is repeated for unlimited times. In practice, ME represents one's expectation for the value of a random variable to occur most frequently.

We assume that renewable power and grid power outage duration are two random variables. As a datacenter typically operates for a long time (e.g., 10—20 years) after it is constructed, the two variables will repeatedly generate a large number of values. Although this number is not unlimited large in theory, it is large enough in practice to use ME to approximately represent the values that occur the most frequently. That is, we use the ME of the renewable power and that of the grid power outage durations to represent the most frequent power that the renewable energy can provide and the most frequent grid power failure duration that may occur, respectively. Subsequently, we leverage these two MEs to determine the rated power capacity and runtime of the UPS for a datacenter.

Planning the Rated Runtime of a UPS. We calculate RL_{ups} based on the ME of grid power outage durations ($E(T_o)$), as shown in Formula (1):

$$\begin{cases} \boldsymbol{RL_{ups}} = min(E(T_o), I_{re}) \\ E(T_o) = \sum_{T_o} T_o \times D_o \end{cases}, \tag{1}$$

where T_o is the grid power outage duration; D_o is its distribution (in percent); and I_{re} is the intermittency duration of renewable energy that can be observed from statistic power data, as shown in Figs. 6a and 6b.

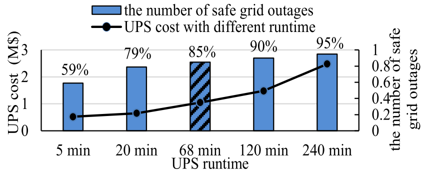

Fig. 7. The cost variation with different rated UPS runtime lengths. Safe grid outages refers to the ones that can be safely handled.

For high cost-effectiveness, we set RL_{ups} to the minimum value of $E(T_o)$ and I_{re}. $E(T_o)$ reflects the characteristics of grid power outages whereas I_{re} reflects the intermittency duration (or availability) feature of renewable energy. When $E(T_o)$ is smaller than I_{re}, it seems that taking $E(T_o)$ as the RL_{ups} would make the datacenter unavailable because the RL_{ups} of the UPS is shorter than the intermittency duration of the renewable energy. This however would not occur because most outage durations are much shorter than $E(T_o)$. Moreover, some power capping techniques can prolong the UPS power supply time, guaranteeing that the datacenter does not shut down while keeping the UPS configuration highly cost-effective. This is the case when using solar energy with long intermittency intervals as the only renewable energy of a datacenter. On the other hand, when $E(T_o)$ is longer than I_{re}, setting RL_{ups} to I_{re} also guarantees a highly cost-effective UPS configuration while keeping the possibility for shutting down a datacenter as low as possible. This is the case when using wind energy as the only renewable energy of a datacenter.

Further, we analyze how UPS cost changes with different rated runtime durations. Fig. 7 shows five cases with different rated UPS runtime durations (all cases with the same 10 MW power capacity) considering the cost and the least number of *safe grid power outage*s that can be successfully handled. Based on Formula (1) and Fig. 1 in a solar energy powered datacenter, $E(T_o)$ is calculated as 68 minutes. With this runtime, the UPS can safely handle at least 85 percent of the grid power outages, which costs 1.05 million dollars.

When the UPS runtime is configured to 20 minutes, it only costs 0.65 million dollars but the UPS can only handle 59 percent of the grid power outages. Compared to the UPS configured with 68 ($E(T_o)$) minutes, its cost is reduced by a factor of $0.65\times$ while the runtime is decreased by a factor of $2.4\times$. On the other hand, compared with the 120-minute or longer time grid power outages, a UPS configured with 68 minutes leads to a cost reduction of at least $2\times$. This indicates that a $2\times$ higher cost only brings a $0.7\times$ runtime improvement. Moreover, this implies that the cost of a UPS increases more sharply when its runtime is configured longer than 68 minutes, and vice versa. In addition, a UPS configured with 68 ($E(T_o)$) minutes can supply power for at least 120 minutes (over 90 percent safe grid power outages) if we leverage dynamical power management techniques (details in Section 3.4) to tune the time and power the UPS can provide. Therefore, $E(T_o)$ is a good trade-off for configuring UPS runtime and cost.

Planning the Rated Power of a UPS. We compute the rated power of a UPS P_{ups} based on the difference between a datacenter's peak power and the ME of renewable power

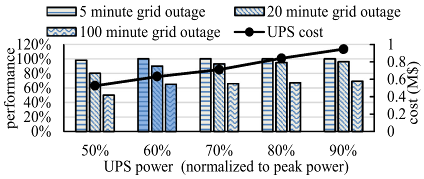

Fig. 8. The performance variation of SEPCjbb (an interactive application) along with the power backup cost for different UPS power capacities. The least UPS power capacity guarantee is mainly for the case of unavailable renewable energy. We therefore evaluate the performance by employing *UPS&Mig&ThrotG* (detailed in Section 3.4).

($E(P_{re})$), as shown in:

$$\begin{cases} \boldsymbol{P_{ups}} = max(P_p - E(P_{re}), 60\% \times P_p) \\ E(P_{re}) = \sum_{P_{re}} P_{re} \times D_{re} \end{cases}, \tag{2}$$

with P_p the peak power needed by a datacenter, P_{re} the available power of renewable energy, $E(P_{re})$ the ME of the available power of renewable energy, and D_{re} the distribution of renewable power. P_{ups} is set to the maximum value among the 60 percent of the peak power (P_p) of a datacenter and the difference between P_p and $E(P_{re})$. Given that a datacenter typically consumes about 60 percent of its peak power [29], the P_{ups} configured with at least 60 percent of the peak power is able to sustain normal operations in a datacenter even if renewable power is unavailable during grid power outages. We therefore select the UPS configured with 60 percent of the peak power as a candidate rated power for the UPS, as Formula (2) shows.

Further, Fig. 8 shows the performance variation of SPECjbb (an interactive application [30]) along with the UPS cost for five different UPS power provisions (all cases with the same $E(T_o)$ runtime). Since the least UPS power capacity guarantee is mainly for the case of unavailable or minimal renewable energy, we adopt the *UPS&Mig&ThrotG* power management technique to evaluate the performance (details in Section 3.4). It guarantees high performance for short outages and high availability for long outages. As can be seen, the performance of the datacenter equipped with a UPS with 50 percent of the datacenter's peak power is obviously lower than the UPS with 60, 70, 80, and 90 percent of the peak power. We therefore should select larger power capacity. On the other hand, the performance of the datacenter with the UPS at 60, 70, 80 and 90 percent of the peak power is very close but their costs increase linearly. Therefore, the UPS with 60 percent of the peak power is a preferred performance-cost trade-off.

We present two examples here for UPS configurations in solar and wind energy powered datacenters. We compute the ME and the intermittency durations of solar energy from Fig. 6c: $E(P_{re}) = 0.34$, I_{re} = 9 hours, and those of the wind energy from Fig. 6d: $E(P_{re}) = 0.66$, I_{re} = 20 minutes. We see that the ME for solar energy is low because of the long-time zero power supply during the night and the ME of wind energy is relatively high because of its continuous power supply. Assuming that the rated renewable power is equal to P_p, we can configure the UPS capacity with *66 percent of the peak power and 68 minutes of runtime* for solar energy powered datacenters, and configure the UPS capacity with *60 percent of the peak power and 20 minutes of runtime* for wind energy powered datacenters. We only determine the least UPS capacity to get a better cost-performance-availability trade-off for green datacenters. For a high-availability requirement in a specific situation, we can properly configure a larger UPS capacity.

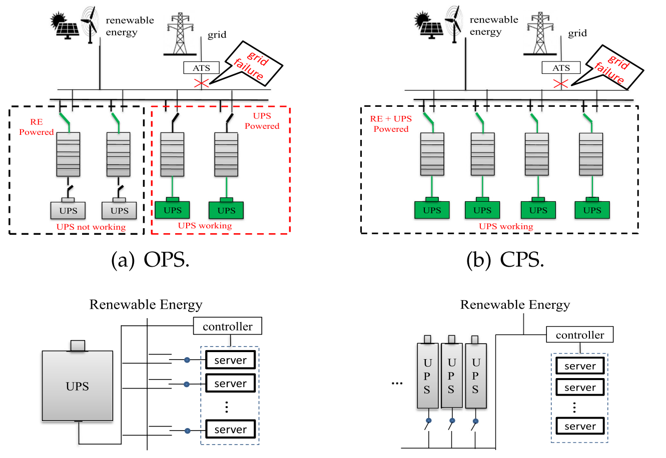

(a) OPS. (b) CPS.

(c) server-end control. (d) power-end control.

Fig. 9. Original Power Supply (be short for OPS) and Cooperative UPS/renewable Power Supply (be short for CPS) is shown in (a) and (b). The two implementations for Cooperative UPS/renewable Power Supply are shown in (c) and (d).

3.3 Cooperative UPS/Renewable Power Supply

When a grid power failure occurs, there are three cases to supply power under the original power supply mode of a green datacenter: (1) renewable energy is sufficient; (2) renewable energy is insufficient but can power some racks and the other racks are powered by UPSs; (3) renewable energy is unavailable and all racks are powered by UPSs. We mainly discuss case (2) which features two simultaneous power sources. In case (2), some racks that are powered by the grid are switched to UPSs to sustain power when grid power fails. The remaining racks are still powered by renewable energy and their corresponding UPSs are in standby mode, as shown in Fig. 9a. This may cause two problems. First, the UPSs are wasted for the racks powered by renewable energy. Second, when the rack-level UPSs are depleted, the corresponding racks would shut down.

We propose cooperative UPS/renewable power supply to utilize all UPSs in a balanced manner and prolong the supply time of the UPSs. During a grid power outage, we configure that *each* rack is simultaneously powered by both renewable energy and UPS, as shown in Fig. 9b. As such, a UPS only needs to power the remaining workload in addition to the workload powered by renewable energy in a rack, which reduces the actual power load on the UPS. According to Peukert's law [31], the actual discharge time (t) of a UPS is much longer for low power draws compared to high power draws, as shown in Formula (3):

$$t = H \times (I_r/I_a)^k = H \times (P_r/P_a)^k, \tag{3}$$

where H is the rated UPS discharge time. I_r and P_r are the rated discharge current and power, respectively. I_a and P_a are the actual discharge current and power, respectively. k is the Peukert constant and depends on the UPS type (e.g., k is between 1.1 and 1.3 for lead-acid battery). Table 1 shows

TABLE 1
Variable Value Description

Variant	OPS	CPS
All racks in datacenter	N	N
RE powered racks	n	N
UPS powered racks	m (m=N-n)	N
RE power	P_{re}	P_{re}
UPS power of the datacenter	P_{ups}	P_{ups}
UPS rated runtime	RL_{ups}	RL_{ups}
each rack demand power	P_d	P_d

N is the Number of Racks of a Datacenter. n is the number of racks powered by renewable energy and m is the number of racks powered by UPS in cooperative UPS/renewable power supply mode. OPS — original Power Supply. CPS — cooperative UPS/renewable power supply.

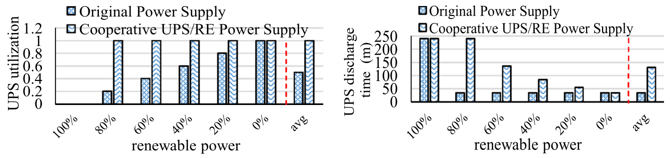

(a) UPS utilization. (b) UPS discharge time.

Fig. 10. UPS utilization and discharge time comparison for Original Power Supply and Cooperative UPS/renewable Power Supply for a 240-minute grid outage.

several variables and their values for original power supply and cooperative UPS/renewable power supply. N is the number of racks of a datacenter. n is the number of racks powered by renewable energy and m is the number of racks powered by UPS in the original power supply mode. In the cooperative UPS/renewable power supply mode, since each rack is powered by renewable energy and UPS simultaneously, the number of renewable energy powered racks is N and that of UPS powered racks is also N. We compare the supply time of UPS for cooperative UPS/renewable power supply (T_{cps}) against that for the original power supply (T_{ops}), as shown in Formulas (4) and (5). As can be seen, since actual power demand for UPS has reduced (from P_d to $P_d - P_{re}/N$), the supply time of UPS for cooperative UPS/renewable power supply is obviously longer than that for the original power supply.

$$T_{ops} = RL_{ups} \times ((P_{ups}/N)/P_d)^k \tag{4}$$

$$\Downarrow T_{cps} = RL_{ups} \times ((P_{ups}/N)/(P_d - P_{re}/N))^k. \tag{5}$$

We propose two possible implementations for cooperative UPS/renewable power supply: (1) in Fig. 9c, a server-end controller (SEC) is designed to separate the workload of a rack to different power sources. Some servers in the rack are powered by renewable energy while others are powered by UPS. For example, there are 10 servers with 2,240 W peak power in a rack; the available renewable energy is 1,344 W, which can power only 6 servers; the other 4 servers need 896 W which are powered by UPSs with rated 1,478 W (66 percent peak power for solar-powered datacenters). (2) In Fig. 9d, a power-end controller (PEC) is designed to configure a number of small-capacity UPSs. The total capacities of the small UPSs are approximately equal to the power capacity in SEC. During a grid power outage, parts of the UPSs are activated to provide the required power. For example, 5 UPSs rated 300 W power can satisfy the power demand of 1,478 W (total 1,500 $\approx$ 1,478 W). If there is a need for an additional 896 W of power, three more UPSs are activated. We only consider SEC in our experimental platform since it is more flexible and easier to implement compared to PEC.

Moreover, we evaluate the *original power supply* mode and *cooperative UPS/renewable power supply* mode (SEC solution) from two aspects for a 4-hour grid power outage: UPS utilization and UPS discharge time in our experiment platform. There are 10 racks, each with 10 servers. The measured peak power of a server is 224 W. We provide 22,400 W of rated power by solar energy and a rack-level UPS with 14,780 W (66 percent of demand power) and 68-minute runtime. Fig. 10a shows that UPS utilization changes with different renewable power for the original power supply mode and only reaches 50 percent on average. The less renewable power (powering fewer racks) results in higher UPS utilization for the original power supply (powering more racks). In contrast, the UPS utilization is 100 percent all the time for cooperative UPS/renewable power supply mode because renewable energy and UPS power each rack together. Note that when renewable power is sufficient to sustain the whole datacenter, the UPS is not used for both power supply modes. Fig. 10b shows the UPS discharge time for both modes. We observe that the cooperative UPS/renewable power supply mode can sustain power for 131 minutes (35 minutes in worst case and 240 minutes in the best case) on average because of the high UPS utilization and UPS discharge characteristics (Formula (3)). For the original power supply mode, the racks that were powered on by rack-level UPSs shut down after only 35 minutes.

3.4 Renewable-Energy-Aware Power Management

Given unpredictable grid power outages, time-varied renewable power, and limited UPS capacity, our *renewable-energy-aware dynamic power management* dynamically employs different power capping techniques based on the discrepancy between power consumption and current power source status. It guarantees the performance and availability of services as much as possible during different durations of grid power outages. We now describe a number of power-capping techniques.

Throttling. Throttling is very effective for exploring the trade-off between performance and power consumption. Our system kernel can be configured with the *Userspace* frequency scaling governor, which provides 8 – 13 CPU frequencies, and dynamically adapts to a proper CPU frequency according to the current renewable energy and UPS state during a grid power outage. We set the CPU frequency of all servers in a rack from f_i to f_j. The reduced power can therefore be expressed as $P_{Th}(f_i, f_j)$, and we have the updated power demand $P_D^{'}$ of the rack: $P_D{}' = P_D - P_{Th}(f_i, f_j)$ (P_D is the power demand before throttling). We calculate power reduction $P_{Th}(f_i, f_j)$ for different CPU frequencies for each application.

Migration. Given insufficient renewable power and UPS power, an alternative is to migrate workloads to racks where the power supply is sufficient. After migration, some servers are shut down to reduce power consumption. We use virtual machine (VM) migration to achieve the best workload survival, which has become the handiest approach to move the workload to other servers and to guarantee the Service-Level Agreement (SLA). We use n to denote the number of servers of a rack and m is the number of active servers after migration. That is, $n - m$ servers $(S_0, S_1, \ldots, S_{n-m-1})$ can be shut down and the reduced power can be expressed as $P_{Mig}(S_0, S_1, \ldots, S_{n-m-1})$. We obtain the updated power demand of the rack as follows: ${P_D}' = P_D - P_{Mig}(S_0, S_1, \ldots, S_{n-m-1})$.

Sleep and Hibernation. Sleep mode is a power saving state as everything in the system is put into a low-power state except memory [32]. Hibernation mode offers the greatest energy savings by pushing the application state to local persistent storage and shutting down the server. Many workloads which are not real-time and run for a long time, such as scientific computation, can tolerate some delays. Such loose time demands provide the flexibility that allows the computation workloads to be unavailable during a rare and short outage. When all power sources are insufficient or unavailable, servers can be set to sleep or hibernation to reduce power consumption.

Although migration can greatly reduce power consumption, there are three situations where migration may not be considered as the first choice: (1) there must be headroom elsewhere in the rack to accommodate the migrated load and meet the application's SLA, which may not happen if the rack is highly utilized; (2) large application state and frequent memory writing can result in long migration times; (3) when grid power is active again, we need to migrate the workload back. If UPS and throttling can handle the grid power outages while meeting the application's SLA, the additional migration cost can be avoided. The agility of throttling makes it more attractive than migration.

We present rack-level renewable-energy-aware dynamic power management to dynamically control power consumption of a datacenter by adopting different power capping techniques. It monitors the renewable power, the status of the UPSs, and the resource usage at a coarse-grained time interval (e.g., 1 minute). According to the state, it selects proper operations to handle grid power outages. The goal is to guarantee the availability and normal performance of applications as much as possible during different durations of grid power outages.

Renewable-energy-aware dynamic power management divides the time-varying renewable energy into five levels: *max*, *large*, *medium*, *low* and *min* (which includes unavailable). We propose five corresponding heuristic techniques, as shown in Table 2. The *max* level of renewable energy can meet power demands of all servers in the rack and the excess energy can charge the UPS. In this case, our power management can guarantee the maximum performance of an application, corresponding to the *Normal* state. The *large* level of renewable energy is insufficient, but can still meet the power demand of applications with SLA in a proper throttling power state (*Throt*). In order to slow down the power consumption from UPSs for unknown-duration grid power outages, our power management selects a set of servers and sets proper CPU frequency for them to control their power consumption from the *large* level of renewable energy rather than from UPSs.

TABLE 2
Heuristic Handling for Different Renewable Energy Status

RE Level	Heuristic	Description
max	Normal	only use renewable energy to supply power.
large	Throt	throttling.
medium	UPS+Mig	drain UPS first before migration
low	UPS&Mig+Throt	drain UPS while migration; after UPS is depleted, then throttling.
min (including unavailable)	UPS&Mig&ThrotG	UPS discharge while migration, while **gradually** throttling.

Assuming that there is a headroom in a rack that allows to migrate workloads to fewer servers with little performance impact, the *medium* level of renewable energy can meet the power demand after migration. For example, there are 10 servers and 30 percent of headroom in the rack. We can seamlessly migrate the workloads of 3 servers to the other 7 servers. Renewable-energy-aware dynamic power management selects the *UPS+Mig* technique to handle the grid power outages for the *medium* level. In this case, UPSs are first activated and cooperate with renewable energy to supply power. The power management defers the migration operation until the UPS capacity reaches the residual capacity needed for migration (~4 minutes for migrating 10 GB application state). The number of servers from which the workload is to be migrated depends on how many servers need to be shut down within a limited power budget.

Renewable-energy-aware dynamic power management employs the *UPS&Mig+Throt* technique for the *low* renewable energy level which even cannot meet the power demand after migration. It first activates UPSs to supplement renewable energy and migration is performed at the same time to reduce overall power consumption. When UPSs are depleted, throttling is employed to control the power demand within the current power budget. The power management delays the throttling to minimize the performance degradation during grid power outages.

When renewable energy is in the *min* level (including unavailable) which cannot meet the power demand after migration and the lowest throttling state, the UPS is the main power source. We need to not only consider the performance of applications but also prolong the UPS's runtime as much as possible because of the unknown grid power outage duration. Our power management therefore employs the *UPS&Mig&ThrotG* technique to sustain power. First, UPSs are activated and workloads are migrated to fewer servers to reduce power requirement. Meanwhile, it starts by throttling at full performance mode (assuming grid outage will be short), then gradually transitions to lower power modes which can slow down the discharging of the UPS, and finally (when UPS batteries are depleted) uses the sleep or hibernate techniques. *UPS&Mig&ThrotG* guarantees high performance for short outages and high availability for long outages as much as possible.

It is worth noting that, if renewable energy is in the medium, low or min level during a power outage, some

TABLE 3
Benchmarks Description

WorkLoad	Memory Usage	Performance Metric
SPECjbb	10 GB	ops/sec
Memcached	10 GB	queries/sec
Web-Search	12 GB	queries/second
mcf in SPEC CPU2017	8 GB	completion time

virtual machines should be migrated. The job scheduler is aware of this situation and distributes new jobs in the job queue to the machines that do not need to migrate jobs.

4 Experimental Setup

We use 10 identical servers with 32-core 3.2 GHz Intel(R) Xeon(R) E5-2630 v3 processors, 64 GB DRAM, and 1 Gbps Ethernet interface in a rack and run our applications hosted on the Linux OS. The power consumption of each server is measured by a power meter. The server idle power is around 100 W and the peak power that we have measured is 224 W. The dynamic power consumption can be modulated with 13 CPU frequencies (from 1.2 to 2.4 GHz).

We randomly choose one of the solar and wind power production traces of one-day from NREL [28] in our experimental platform. In our setup, we provision 2,240 W of rated solar energy power and 1,478 W (2240×0.66) of UPS with 68-minute runtime (obtained by the Least UPS Capacity Determining technique for grid power outages). We use *cpufreq* (the Linux interface) [33] to implement throttling, which allows the system to set to a given CPU frequency, and leverage the *KVM live migration* technique [34] for migration. Moreover, we use the OS commands in Linux to implement the sleep and hibernation power saving modes.

To evaluate the efficacy of our COPA approach, we experiment with three relevant scenarios during grid power outages: (1) different availability levels of renewable energy: *max*, *large*, *medium*, *low*, and *min*; (2) different durations of grid power outages: 5 minutes, 20 minutes, 60 minutes, 120 minutes, and 180 minutes; (3) four representative workloads that have different characteristics and different SLA demands, as shown in Table 3. Interactive applications include SPECjbb [30], an in-memory key-value store Memcached benchmark [35], and Web-Search from CloudSuite [36]. We use *mcf* from SPEC CPU2017 in our experiments as a representative for a computation-intensive workload.

Because COPA power management (e.g., throttling) operates on individual servers and does not consider all servers as a whole for scheduling, the COPA approach can easily scale out to thousands of servers.

5 Results and Analysis

5.1 Cost Analysis

UPS Cost. Since the power backup infrastructure is rarely used, we ignore the operational expenditure such as the cost for management and energy loss, and only focus on the up-front procurement cost and the amortized future replacement cost, called capital expenditure (cap-ex). UPS units come with rated runtime and power capacity. We express cap-ex as amortized $/year, using a linear depreciation model. The cost values of backup infrastructure are depreciated based on their lifetime, i.e., 4 years for UPS lifetime. According to [13], the UPS cap-ex, C_{ups} ($/year) depends on the rated power capacity (P_{ups} in KW) and runtime (RL_{ups} in h), and can be expressed as:

$$C_{ups} = PC_{ups} \times P_{ups} + EC_{ups} \times P_{ups} \times (RL_{ups} - RL_{free})), \tag{6}$$

where PC_{ups} ($/KW/year) and EC_{ups} ($/KWh/year) represent power cost per year and energy cost per year, respectively. PC_{ups} and EC_{ups} reach 50$/KW/year and 50 $/KWh/year, respectively [37]. RL_{free} refers to the free runtime expected with rated power, and its value is 2-minute [13]. In our experiments, the UPS with 0.66 of the peak power and 68-minute runtime costs 1.39 M$/year in a solar-powered 20 MW datacenter. The UPS with 0.6 of the peak power and 20-minute runtime costs 0.78 M$/year in a wind-powered 20 MW datacenter.

Cost for Cooperative UPS/Renewable Power Supply. Cooperative UPS/renewable power supply needs additional control switches to make renewable energy collaborate with UPS in a rack (details in Section 3.3). Assuming that the number of servers in a datacenter is S, we need S switches. The per switch cost is C_s and lifetime is T_{life} years. The cost for cooperative UPS/renewable power supply, C_{cps} ($/year) can therefore be expressed as:

$$C_{cps} = S \times C_s / T_{life}. \tag{7}$$

A control switch can be implemented by a programmable chip, which costs as low as 0.14$ [38] and have 10 years of lifetime [39]. Assuming that there are 90,000 servers with 20 MW power demand of a datacenter (the peak power of each server is 224 W in our experimental platform), the cost of Cooperative UPS/renewable Power Supply for the SEC implementation equals 1,260$/year. The PEC also needs control switches to select multiple batteries, for which the cost is similar as for SEC.

Renewable Energy Cost. In green datacenters, renewable energy is used all the time but we only consider the amortized cost during the grid power outages rather than the direct construction cost because we focus on power back-up solutions. The renewable energy cost for COPA C_{re} ($/year) can be expressed as:

$$C_{re} = C_{perKWh} \times P_{re} \times duration, \tag{8}$$

where C_{perKWh} ($/KWh) represents the cost per KWh of renewable energy, P_{re} (KW) is the rated renewable power, and the $duration$ is the total length of all grid power outages per year. Taking solar and wind energy as examples, the solar and wind energy cost C_{perKWh} is as low as 0.12$/KWh [40] and 0.1$/KWh[41], [42], respectively. The $duration$ has 638 minutes/year in the worst case (240 minutes $\times$ 2.66/ year) for datacenters [9]. For a 20 MW datacenter , the solar energy cost and wind energy cost for COPA are only 24,000 $/year and 20,000$/year, respectively.

Now we investigate the total cost of COPA, and compare it with that of the traditional power backup infrastructure, as shown in Table 4. We see that COPA reduces 47 percent

TABLE 4
Cost Comparison with [14],[43]; M$ - Million Dollars

	COPA		OPB
	Solar energy	Wind energy	
Peak power (MW)	20	20	20
DG cost (M$/year)	0	0	1.66
UPS cost ((M$/year))	1.39	0.78	1.02
RE cost (M$/year)	0.024	0.02	0
CPS cost(M$/year)	0.0013	0.0013	0
All cost(M$/year)	**1.42**	**0.79**	2.68

CPS - Cooperative UPS/renewable Power Supply. RE - renewable energy.

of power backup cost for a solar-powered datacenter and 70 percent for a wind powered datacenter. In detail, COPA can save 1.26 M$ (million dollars) and 1.89 M$ per year for 20 MW solar and wind-powered datacenter, respectively.

5.2 Effectiveness of COPA

We first evaluate the performance of SPECjbb by using COPA during different durations of grid power outages and compare with the traditional power backup approach. We use the average operation/second (ops) of servers in the rack as our performance metric for SPECjbb. Fig. 11 shows the aggregated power of the 10 green-provisioned servers running SPECjbb under different levels of renewable energy. We see high variation of solar energy production over time. In this paper, we assume that the performance loss within 10 percent compared to normal performance can still meet the SLA of SPECjbb, corresponding to ≥2 GHz CPU frequency and 2,040 W power. Meanwhile, there is 30 percent headroom in the rack. In other words, the workloads of 3 servers can be migrated to the other 7 servers without performance degradation. On the other hand, more workload migration will cause larger performance loss. After migration, the 3 servers can be powered down to reduce power consumption. Subsequently, we compute the values of P_{Mig} and $P_{LowestTh}$ a priori.

We compute the thresholds of the five levels of renewable energy for SPECjbb: 2,240 W (max), 2,240 W - 2,040 W (large), 2,040 W – 1,568 W (medium), 1,568 W – 840 W (low), and 840 W – 0 W (min), as shown in Fig. 11. We have evaluated the performance for all the cases of time-varying renewable energy over different grid power outage durations. We only present the performance and availability impact on applications for five representative levels of renewable power: 2240 W, 2150 W, 1800 W, 1250 W, 450 W and we assume at this moment a grid power failure occurs.

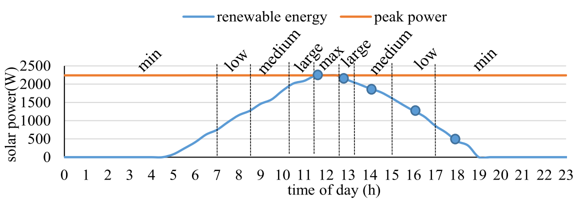

Fig. 11. The five levels of availability of renewable energy for SPECjbb over time. The five points in different solar power levels indicate the moment when a unpredicted grid power outage occurs (e.g., a outage may be occurs in the *max* level or *large* level).

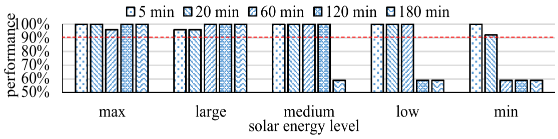

Fig. 12. Performance of SPECjbb for COPA normalized to MaxPerf (with the original power backup) for five solar energy levels. The red line is drawn as a reference for acceptable performance.

The renewable-energy-aware dynamic power management of COPA employs different techniques for different levels.

Impact of Renewable Energy Levels on Performance. We observe the performance of a program from the moment when the renewable energy is in a certain level and a grid power outage occurs at the same time to the end of the outage. Fig. 12 shows how the solar energy levels affect the performance (operations per second) of SPECjbb when using COPA. The X axis represents the five levels of the solar energy. The performance corresponding to the *max* level represents the one observed from the moment when the solar energy is in the *max*level and at this time a grid power outage occurs to the end of the outage. The same applies to other levels of renewable energy. We analyze the performance impacts for five levels of the solar energy next.

Max Level. From Fig. 11, we see that the solar energy can be in the *max* level for about 40 minutes after a grid power outage occurs. For the 5-minute and 20-minute power outages, SPECjbb can run with the maximum (full) performance, as shown in Fig. 12. However, Fig. 12 shows that SPECjbb experiences a little bit performance degradation for the 60-minute grid power outage. The reason is that the solar energy enters its *large* level and it can not support SPECjbb to run with the full performance after 40 minutes. In this level, the renewable-energy-aware dynamic power management in COPA employs the corresponding tuning technique to control the power consumption by *Throt* without using UPS, which reduces performance by 4 percent. The *large* level of solar energy lasts for about 50 minutes (from 40 to 90 minutes after the grid power fails). Then the solar energy enters the *medium* level and the power management technique employs the *UPS+Mig* to sustain operation which makes SPECjbb run at full performance. In summary, COPA makes a datacenter meet the SLA of applications for at least 3 hours if the grid power outage occurs at the moment when the solar energy is in its *max* level, as shown in Fig. 12.

Large Level. If the solar energy enters its *large* level at the moment when a grid power outage occurs, the performance is still able to meet the SLA of SPECjbb for 3 hours. The *Throt* is first used for 50 minutes with 96 percent of the full performance of SPECjbb. Then renewable-energy-aware dynamic power management employs *UPS+Mig* to sustain power with the full performance for 100 minutes, which the solar energy already enters its *medium* level.

Medium Level. If the solar energy is in its *medium* level and at a certain moment the grid power fails, COPA can meet the SLA of SPECjbb for 160 minutes by using *UPS+Mig* and *UPS&Mig+Throt*. Since the UPS power is enough for the *medium* level, migration is executed at the beginning of the *low* level for *UPS&Mig+Throt*. Due to the 30 percent of headroom in the rack, we can only migrate the workloads of up

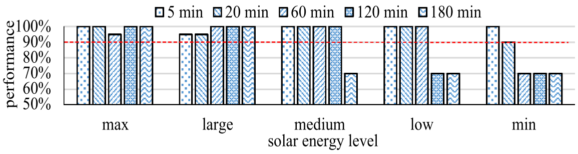

Fig. 13. Performance of Memcached for COPA normalized to MaxPerf (with the original power backup) for five solar energy levels. The red line is drawn as a reference for acceptable performance.

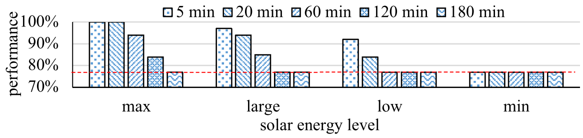

Fig. 15. Performance of mcf for COPA normalized to MaxPerf (with the original power backup) for five solar energy levels and five moments of the outage. The red line is drawn as a reference for acceptable performance.

to 3 servers to the other 7 servers, which causes little performance impact and it takes about 4 minutes for SPECjbb. When the UPSs are depleted, COPA reduces the performance by 41 percent with the lowest throttling state. We can see that COPA prolongs the supply time of UPS by $2.35\times$ (from 68 to 160 minutes).

Low Level. The *low* level of solar energy does not meet the power demand after migration. In this case, COPA achieves full performance by using *UPS&Mig+Throt* for 60 minutes. Subsequently, throttling is used to reduce power consumption with a 59 percent performance degradation.

Min Level. If the solar energy is at its *min* (minimal or unavailable) level and a grid power outage occurs, renewable-energy-aware dynamic power management uses *Mig&UPS&ThrotG* with gradually lower CPU frequency to provide high performance for short grid outages and high availability for long outages but with lower performance. As such, our power management can achieve the trade-off of performance and availability by using throttling.

5.3 Impact of Application Characteristics

For the rest of the evaluation, we consider the COPA in the solar powered datacenter (configuring UPS with 66 percent demand power and 68 minute runtime) and compare the performance and availability impact for applications with diverse characteristics.

Memcached. Due to high memory-related CPU stalls which are better for throttling-based techniques [44], we find that Memcached achieves better performance through throttling than SPECjbb for the *low* and *min* levels of renewable energy. For the *low* level of the renewable energy, the performance achieves the full performance for 60 minutes after a grid power outage occurs, then degrades to 70 percent for lowest throttling state, which is 11 percent higher than for SPECjbb, as shown in Fig. 13. As for the *min* level, our renewable-energy-aware dynamic power management employs gradual throttling to achieve different performance levels (from high to low). The performance of Memcached is 10 percent higher than for SPECjbb for the same level on average.

Web-Search. The performance of Web-search is insensitive to CPU frequency. For example, the performance of Web-search can still achieve 93 percent of the full performance

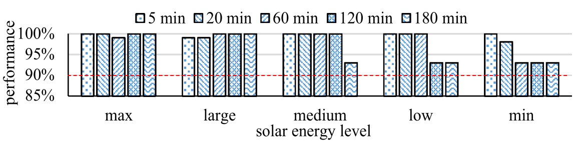

Fig. 14. Performance of Web-search for COPA normalized to MaxPerf (with the original power backup) for five solar energy levels. The red line is drawn as a reference for acceptable performance.

and meets the SLA at the lowest CPU frequency. This is beneficial to control power while achieving high performance through throttling. Web-search shows better performance than Memcached during grid power outages, as shown in Fig. 14. For the *low* level of the renewable energy, the performance only degrades to 93 percent of the full performance after 180 minutes, 23 percent higher than Memcached. As for the *min* level, the performance of different throttling state for Web-search can achieve at least 93 percent of the full performance until the UPSs are depleted.

We take *mcf* from SPEC CPU2017 as a representative of scientific applications. In Fig. 15, the renewable energy enters the *max* level at the moment when a grid power outage occurs, it can achieve full performance for 30 minutes and degrades to 94 percent, further 84, 77 percent as the renewable energy gradually falls. For the *large* and *low* renewable energy levels, proper throttling state is set based on the renewable power state without using UPSs. When renewable energy enters its *min* level, UPSs are activated and the lowest CPU frequency is set to prolong the UPS's runtime with a 23 percent performance degradation. The availability of *mcf* is still at least 3 hours for a grid power outage.

Key Insights. (1) For applications with low sensitivity of CPU frequency, throttling is a good choice to control power while meeting the application's SLA. (2) For scientific applications that can tolerate delays, we can only use throttling to the lowest CPU frequency to prolong the UPS supply time without migration.

5.4 Availability Analysis of Datacenters

High availability of datacenters is the guarantee of high service quality. The availability of a datacenter per year (A_{dc}) can be expressed as: $A_{dc} = (1 - T_{dt}/T_{year}) \times 100\%$, where the T_{dt} represents the actual power-down time of a datacenter per year and T_{year} is the time of one year ($24 \times 365 = 8760 hours$). We can see that COPA can sustain power for at least 3 hours. In most (over 90 percent) cases , A_{dc} can achieve 100 percent. In the worst case (i.e., 240 minutes for each grid power outage), A_{dc} can also achieve at least 99.97 percent.

6 Related Work

Datacenter Power Backup. Due to the high cost of power infrastructure in datacenters, several efforts on under-provisioning power infrastructures have been proposed. With power under-provisioning, backup power works not only during grid power outages, but also during the procedure for shaving the peak power and demand response [45], [46]. Wang et al. [10] explore the pros and cons of placing different energy storage devices in different layers of the power hierarchy. Narayanan et al. [47] perform a detailed study of the pros and cons when dual-purposing (outages and demand response) UPSs are warranted. These works do not focus on the configuration exploration of power backup for grid power outages in green datacenters.

Wang et al. [13] present cost-performance-availability trade-offs for different levels of backup under-provisioning for applications with diverse reliance on the backup infrastructure. However, this work only gives different rough power backup configurations without considering renewable energy for green datacenters. In contrast, COPA proposes a careful power backup configuration exploration by leveraging the availability characteristics of renewable energy as well as the grid power outage characteristics to significantly reduce the power backup cost, while minimizing the performance and availability impacts.

Using Renewable Energy in Datacenters. With the wide use of renewable energy in datacenters, many approaches have been proposed to under-provision the grid power infrastructure to save grid power investments [15], [48], [49], [50] and optimize renewable resource utilization to improve performance per Watt [51], [52], [53], [54], [55], [56]. To our best knowledge, while much prior work has focused on renewable energy, the cost-benefit power backup approaches by using renewable energy have not yet been explored so far. COPA under-provisions the power backup infrastructure reasonably.

Power Management. Throttling such as dynamic voltage/frequency scaling (DVFS) and migration are effective power management techniques to reduce power consumption [13], [15], [57]. However, a single power capping technique or a fixed technique combination is not enough for handling the time-varying renewable energy and unplanned (unknown-length) grid power outages. Raghavendra et al. [58] propose a hierarchy of coordinated controllers for peak and average power management across hardware and software for complex enterprise environments to minimize the performance impact. This work however does not consider power backup for grid power outages. Given the time-varying renewable energy and limited power backup capacity during unpredicted outages, COPA employs different power capping techniques for different renewable power levels in green datacenters, achieving high performance for short outages and high availability for long outages. Liu et al. [59] studied how to select a power source for normal operation rather than for backing up power in a datacenter. Recently, security has become an important concern in datacenter power management. Hou et al. [60], [61] investigate datacenter power management in the context of denial-of-service (DoS) attacks.

7 Conclusion

This paper proposes COPA, a highly cost-effective power backup approach by leveraging the availability characteristics of renewable energy as well as the grid power outage characteristics. COPA consists of three techniques: least UPS capacity planning, cooperative UPS/renewable power supply, and renewable-energy-aware dynamic power management. Least UPS capacity planning explicitly reduces the power backup cost by reconfiguring the UPSs while cooperative UPS/renewable power supply and renewable-energy-aware dynamic power management implicitly decrease the cost by prolonging the supply time of UPS batteries as long as possible during grid power outages. As a result, COPA significantly reduces the power backup cost and substantially improves performance and availability, yielding a much improved trade-off between cost, performance, and availability for green datacenters compared to traditional power backup solutions.

Acknowledgments

The authors would also like to thank the anonymous reviewers for their valuable comments. This work is supported by the national key research and development program under Grant No. 2016YFB1000204, outstanding technical talent program of CAS, NSFC under (Grant No. 61672511, 61702495, and 61802384), Postdoctoral Science Foundation of China under No. 2017M622830, and NSF of Guangdong province under Grant No. 2017A030310350. Lieven Eeckhout is partly supported by a Chinese Academy of Sciences (CAS) visiting professorship for senior international scientists.

References

[1] A. Shehabi *et al*., "United statesdata center energy usage report, lawrence berkeley national labo-ratory, berkeley, california," Lawrence Berkeley National Laboratory, LBNL-1005775, vol. 4, 2016.
[2] N. Rasmussen, "Allocating data center energy costs and carbon to it users," *APC White Paper*, 2009.
[3] P. Delforge, "America's data centers consuming and wasting growing amounts of energy," 2015. [Online]. Available: https://www.nrdc.org/resources/americas-data-centers-consuming-and-wasting-growing-amounts-energy
[4] 2014. [Online]. Available: https://www.apple.com/environment
[5] 2017. [Online]. Available: https://environment.google/projects/announcement-100/
[6] 2014. [Online]. Available: https://sustainability.fb.com/clean-and-renewable-energy/
[7] G. Cook, "How clean is your cloud," *Catalysing an Energy Revolution*, 2012, Art. no. 11.
[8] G. Cook *et al*., "Clicking clean: Who is winning the race to build a green internet?" Greenpeace Inc., Washington, DC, pp. 1–102, 2017.
[9] P. Institute, "National survey of datacenter outages," 2010. [Online]. Available: http://www.inquirere.com/wp-content/uploads/2010/11/National-Survey-on-Data-Center-Outages.pdf
[10] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: What, where, and how much?" *ACM SIGMETRICS Performance Eval. Rev*., vol. 40, no. 1, 2012, pp. 187–198.
[11] M. A. Bell, "Use best practices to design data center facilities," *Gartner Res.*, vol. 22, no. G00127434, Apr. 2005.
[12] Z. S. Yang Jun, "Analysis of the data center investment income, china united network communications group co.ltd.," *Desig. Techn. Post Telecommun*., vol. 11, pp. 84–87, 2015.
[13] D. Wang, S. Govindan, A. Sivasubramaniam, A. Kansal, J. Liu, and B. Khessib, "Underprovisioning backup power infrastructure for datacenters," in *Proc. Int. Conf. Architectural Support Program. Languages Operating Syst*., 2014, pp. 177–192. [Online]. Available: http://doi.acm.org/10.1145/2541940.2541966
[14] L. A. Barroso, J. Clidaras, and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis Lectures Comput. Architecture*, vol. 8, no. 3, pp. 1–154, 2013.
[15] X. Zhou, Q. Cao, H. Jiang, and C. Xie, "Underprovisioning the grid power infrastructure for green datacenters," in *Proc. 29th ACM Int. Conf. Supercomputing*, 2015, pp. 229–240.
[16] P. Costello and R. Rathi, "Data center energy efficiency, renewable energy and carbon offset investment best practices," *RealEnergyWriters.com*, , pp. 1–92, Jan., 2012.
[17] F. Kong and X. Liu, "A survey on green-energy-aware power management for datacenters," *ACM Comput. Surveys*, vol. 47, no. 2, 2015, Art. no. 30.
[18] Q. Deng, D. Meisner, A. Bhattacharjee, T. F. Wenisch, and R. Bianchini, "Coscale: Coordinating CPU and memory system DVFS in server systems," in *Proc. 45th Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2012, pp. 143–154.

[19] A. Gandhi, M. Harchol-Balter, R. Das, and C. Lefurgy, "Optimal power allocation in server farms," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 37, no. 1, 2009, pp. 157–168.
[20] P. Ranganathan, P. Leech, D. Irwin, and J. Chase, "Ensemble-level power management for dense blade servers," *ACM SIGARCH Comput. Architecture News*, vol. 34, no. 2, 2006, pp. 66–77.
[21] 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.*, 2009, pp. 317–330.
[22] J. Stoess, C. Lang, and F. Bellosa, "Energy management for hypervisor-based virtual machines," in *Proc. USENIX Annu. Tech. Conf.*, 2007, pp. 1–14.
[23] K. Fehrenbacher, "The era of the 100 MW data center," 2012. [Online]. Available: https://gigaom.com/2012/01/31/the-era-of-the-100-mw-data-center/
[24] J. Hidalgo, "A rare look inside apples expanding reno data center," 2018. [Online]. Available: https://www.rgj.com/story/money/business/2018/01/17/sneak-peek-rare-look-inside-apple-reno-data-center/1038984001/
[25] "Facebook rack-level ups for improved efficiency," [Online]. Available: http://www.datacenterknowledge.com/archives/2011/04/07/
[26] "Microsoft reveals its specialty servers,racks," [Online]. Available: http://www.datacenterknowledge.com/archives/2011/04/25/microsoft-reveals-its-specialty-servers-racks/
[27] V. Kontorinis *et al.*, "Battery provisioning and associated costs for data center power capping," Univ. California, San Diego, Tech. Rep., CS2012-0985, 2012.
[28] "Measurement and instrumentation data center," 2013. [Online]. Available: http://www.nrel.gov/midc/
[29] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," *ACM SIGARCH Comput. Architecture News*, vol. 35, no. 2, 2007, pp. 13–23.
[30] "Specjbb 2015: Java business benchmark," 2015. [Online]. Available: http://www.spec.org/jbb2015
[31] W. Peukert, "Introdution to Peukert's law," 2015. [Online]. Available: http://vlab.amrita.edu
[32] D. Meisner and T. F. Wenisch, "Dreamweaver: Architectural support for deep sleep," *ACM SIGPLAN Notices*, vol. 47, no. 4, pp. 313–324, 2012.
[33] "CPU frequency scaling," 2019. [Online]. Available: https://wiki.archlinux.org/index.php/CPU_frequency_scaling
[34] "KVM live migration," 2018. [Online]. Available: http://www.linux-kvm.org/page/Migration
[35] "Memcached: Memory-object caching system," 2018. [Online]. Available: http://www.memcached.org.
[36] M. Ferdman *et al.*, "Clearing the clouds: A study of emerging scale-out workloads on modern hardware," *ACM SIGPLAN Notices*, vol. 47, pp. 37–48, 2012.
[37] APC, "Infrastruxure total cost of ownership," 2013. [Online]. Available: http://www.apc.com/in/en/tools/isx/tco/index.cfm
[38] "The cost of electronic components," 2013. [Online]. Available: http://www.lcsc.com
[39] NASA, "Programmable logic devices (pld) handbook," 2013. [Online]. Available: https://standards.nasa.gov/sites/standards.nasa.gov/files/standards/NASA/PUBLISHED/Baseline1/NASA-HDBK-4008.pdf
[40] "The cost of renewable energy," 2001. [Online]. Available: http://www.china-nengyuan.com/news
[41] G. Zhang, "The analysis on the cost of per kwh of solar power and wind power," *China Electric Power*, vol. 6, pp. 69–73, 2016.
[42] "The cost analysis of wind energy," 2017. [Online]. Available: http://www.china5e.com
[43] N. Rasmussen, "Determining total cost of ownership for data center and network room infrastructure," *White Paper*, vol. 6, 2011.
[44] A. Weissel and F. Bellosa, "Process cruise control: Event-driven clock scaling for dynamic power management," in *Proc. Int. Conf. Compilers Architecture Synthesis Embedded Syst.*, 2002, pp. 238–246.
[45] V. Kontorinis *et al.*, "Managing distributed ups energy for effective power capping in data centers," in *Proc. 39th Annu. Int. Symp. Comput. Architecture*, 2012, pp. 488–499.
[46] Y. Li *et al.*, "Sizecap: Efficiently handling power surges in fuel cell powered data centers," in *Proc. IEEE Int. Symp. High Performance Comput. Architecture*, 2016, pp. 444–456.
[47] I. Narayanan, D. Wang, A. Al Mamun, A. Sivasubramaniam, and H. K. Fathy, "Should we dual-purpose energy storage in datacenters for power backup and demand response?" in *Proc. 6th USENIX Conf. Power-Aware Comput. Syst.*, 2014, pp. 7–7.
[48] J. Pierson *et al.*, "DATAZERO: Datacenter with zero emission and robust management using renewable energy," *IEEE Access*, vol. 7, pp. 103 209–103 230, 2019. [Online]. Available: https://doi.org/10.1109/ACCESS.2019.2930368
[49] L. Grange, G. D. Costa, and P. Stolf, "Green IT scheduling for data center powered with renewable energy," *Future Generation Comput. Syst.*, vol. 86, pp. 99–120, 2018. [Online]. Available: https://doi.org/10.1016/j.future.2018.03.049
[50] P. Luo, X. Wang, H. Jin, Y. Li, and X. Yang, "Load management for multiple datacenters towards demand response in the smart grid integrating renewable energy," in *Proc. 2nd Int. Conf. Comput. Sci. Artif. Intell., 10th Int. Conf. Inf. Multimedia Technol.*, 2018, pp. 140–144. [Online]. Available: https://doi.org/10.1145/3297156.3297197
[51] 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 Performance Comput. Architecture*, 2011, pp. 205–216.
[52] C. Li, A. Qouneh, and T. Li, "iSwitch: coordinating and optimizing renewable energy powered server clusters," in *Proc. 39th Annu. Int. Symp. Comput. Architecture*, 2012, pp. 512–523.
[53] C. Li *et al.*, "Towards automated provisioning and emergency handling in renewable energy powered datacenters," *J. Comput. Sci. Technol.*, vol. 29, no. 4, pp. 618–630, 2014.
[54] C. Li, R. Wang, D. Qian, and T. Li, "Managing server clusters on renewable energy mix," *ACM Trans. Autonomous Adaptive Syst.*, vol. 11, no. 1, 2016, Art. no. 1.
[55] E. Sheme, S. Lafond, D. Minarolli, E. K. Meçe, and S. Holmbacka, "Battery size impact in green coverage of datacenters powered by renewable energy: A latitude comparison," in *Proc. 6th Int. Conf. Emerging Internet Data Web Technol.*, 2018, pp. 548–559. [Online]. Available: https://doi.org/10.1007/978-3-319-75928-9_49
[56] Y. Lu, Ran, P. Wang, Y. Cao, J. Hao, and K. Zhu, "Energy-efficient task scheduling for data centers with unstable renewable energy: A robust optimization approach," in *Proc. IEEE Int. Conf. Internet Things, IEEE Green Comput. Commun., IEEE Cyber Physical Social Comput., IEEE Smart Data*, 2018, pp. 455–462. [Online]. Available: https://doi.org/10.1109/Cybermatics_2018.2018.00101
[57] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters," in *Proc. 17th Int. Conf. Architectural Support Program. Languages Operating Syst.*, 2012, pp. 75–86.
[58] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu, "No power struggles: Coordinated multi-level power management for the data center," *ACM SIGARCH Comput. Architecture News*, vol. 36, 2008, pp. 48–59.
[59] L. Liu, H. Sun, C. Li, Y. Hu, N. Zheng, and T. Li, "Towards an adaptive multi-power-source datacenter," in *Proc. Int. Conf. Supercomputing*, 2016, pp. 11:1–11:11. [Online]. Available: https://doi.org/10.1145/2925426.2926276
[60] 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.*, 2019, pp. 13:1–13:10. [Online]. Available: https://doi.org/10.1145/3337821.3337856
[61] 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. 36th IEEE Int. Conf. Comput. Design*, 2018, pp. 26–34. [Online]. Available: https://doi.org/10.1109/ICCD.2018.00015

Yan Yin is currently working toward the PhD degree from the Department of Computer Science and Technology, University of Science and Technology of China (USTC), Hefei, China. Her research interests include computer architecture, virtualization technologies, cluster computing, and deep neural networks.

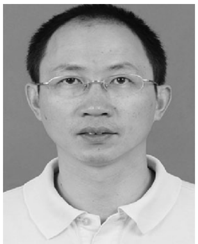

Junmin Wu received the PhD degree in computer science and engineering from the University of Science and Technology of China (USTC), in 2005. He is an associate professor with the Department of Computer Science and technology of USTC, Hefei, and assistant dean of the Suzhou Institute for Advanced Study of USTC, Suzhou. His research interests include computer architecture, virtualization technology, cluster computing, multi-core computing and simulation.

Xu Zhou received the PhD degree in computer science and technology from the Huazhong University of Science and Technology, in 2016. His research interests span the areas of computer systems, computer architecture, energy-efficient systems design, and sustainable computing.

Amer Qouneh received the PhD degrees in computer engineering from the University of Florida, in 2014. He is currently an assistant professor with the Electrical and Computer Engineering Department, Western New England University, Springfield, Massachusetts. He was with the Royal Scientific Society, Jordan, from 1993 to 2000. His current research interests include energy efficiency and power management in data centers and servers, heterogeneous architectures, and power aware scheduling.

Lieven Eeckhout received the PhD degree in computer science and engineering from Ghent University, in 2002. He is currently a full professor at Ghent University, Belgium. He is the recipient of the 2017 ACM SIGARCH Maurice Wilkes Award, the 2017 ACM SIGPLAN OOPSLA Most Influential Paper Award, and was elevated to IEEE fellow in 2018. He served as the editor-in-chief of IEEE Micro (2015–2018) and serves as Program Chair for ISCA 2020. His research interests include the area of computer architecture, with a specific interest in performance analysis, evaluation and modeling, as well as dynamic resource management.

Tao Li received the PhD degree in computer engineering from the University of Texas at Austin. He is an associate professor with the Department of Electrical and Computer Engineering, University of Florida. He received 2009 National Science Foundation Faculty Early CAREER Award, 2008, 2007, 2006 IBM Faculty Awards, 2008 Microsoft Research Safe and Scalable Multi-core Computing Award and 2006 Microsoft Research Trustworthy Computing Curriculum Award. He co-authored a paper that won the Best Paper Award in HPCA 2011 and three papers that were nominated for the Best Paper Awards in DSN 2011, MICRO 2008 and MASCOTS 2006. He is one of the College of Engineering winners, University of Florida Doctor Dissertation Advisor/Mentoring Award for 2013-2014 and 2011-2012. He is a fellow of the IEEE. His research interests include computer architecture, microprocessor/memory/storage sys temdesign, virtualization technologies, energy-efficient/sustainable/dependable data center, cloud/big data computing platforms, the impacts of emerging technologies/applications on computing, and evaluation of computer systems.

Zhibin Yu received the PhD degree in computer science from the Huazhong University of Science and Technology (HUST), in 2008. He visited the Laboratory of Computer Architecture (LCA) of ECE of the University of Texas at Austin for one year and he worked in Ghent University as a postdoctoral researcher for half of a year. Now he is a professor in the SIAT. He won the outstanding technical talent program of Chinese Academy of Science (CAS) in 2014 and the 'peacock talent' program of Shenzhen City in 2013. He also won the first award in teaching contest of HUST young lectures in 2005 and the second award in teaching quality assessment of HUST in 2003. He is a member of the IEEE and ACM. He serves for ISCA 2013, MICRO 2014, HPCA 2015, 2018, and ICS2018. His research interests include micro-architecture simulation, computer architecture, workload characterization and generation, performance evaluation, multi-core architecture, GPGPU architecture, virtualization technologies and so forth.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/csdl.