Design and Implementation of Smooth Renewable Power in Cloud Data Centers

Xinxin Liu, Yu Hua, Xue Liu[†], Ling Yang, Yuanyuan Sun Huazhong University of Science and Technology [†] McGill University

Abstract—The renewable power has been widely used in modern cloud data centers, which also produce large electricity bills and the negative impacts on environments. However, frequent fluctuation and intermittency of renewable power often cause the challenges in terms of the stability of both electricity grid and data centers, as well as decreasing the utilization of renewable power. Existing schemes fail to alleviate the renewable power fluctuation, which is caused by the essential properties of renewable power. In order to address this problem, we propose an efficient and easy-to-use smooth renewable power-aware scheme, called Smoother, which consists of Flexible Smoothing (*FS*) and Active Delay (*AD*). First, in order to smooth the fluctuation of renewable power, FS carries out the optimized charge/discharge operation via computing the minimum variance of the renewable power that is supplied to data centers per interval. Second, AD improves the utilization of renewable power via actively adjusting the execution time of deferrable workloads. Extensive experimental results via examining the traces of real-world data centers demonstrate that Smoother significantly reduces the negative impact of renewable power fluctuations on data centers and improves the utilization of renewable power by 250.88% on average. We have released the source codes for public use.

1 INTRODUCTION

I N cloud environment, large electricity bills have become almost unaffordable with the rapid growth of computation capability and the scale of cloud data centers [1]. In the meantime, the negative impacts of the high consumption of brown energy (i.e., carbon-intensive fuels) have caused significant environmental concerns [2]. In order to address these problems, energy-efficient management schemes have been proposed [3], [4], [5]. The economical and environment-friendly renewable energy has been used to meet the needs of data centers and receives many attentions [6], [7], [8].

However, due to the fluctuation and intermittency properties of renewable power [2], using renewable power to supply data centers needs to address two main challenges:

Challenge (1): *High Fluctuation.* High fluctuation and intermittency of renewable power pose the risks to data centers and often lead to the instability of both grid [9] and clusters [10]. For example, when lots of renewable power penetrates the system, frequent fluctuation in renewable power generation can generally degrade system frequency stabilization, resulting in higher maximum rate-of-change-of-frequency (ROCOF) [11], which is unsafe and unreliable for systems. Furthermore, frequent fluctuation of renewable power also increases the management overheads for renewable energy, such as the overhead of frequent load migration between grid-powered/renewable-energy-powered clusters [12] and the requirement for a large battery capacity [13].

Challenge (2): *Low Power Utilization.* Based on the analysis of real-world traces, the generation of renewable power and the workload power demand in data centers are different and possibly lead to the imbalance between demand and supply. This imbalance fails to make full use of

renewable power in data centers and reduces the utilization of renewable power [14].

Unfortunately, existing solutions fail to address the above challenges. Multigreen [7] proposes a costminimizing control algorithm that uses the generated renewable power as much as possible, which incurs the fluctuation of renewable power to be used without considering the renewable energy in battery, and overlooking the effects of the fluctuation on the stability of grid and data centers. Moreover, the cost minimizing online scheme [13] stores all the renewable power into the battery and selectively charge the battery with the grid. However, storing all the wind energy of a wind farm with an installed capacity of 12MW requires a large battery capacity and high battery charging/discharging rate. iSwitch [12] allows the grid and the renewable energy circuit to be independent of each other, thus avoiding the impact of wind power fluctuation on the stability of the grid [9]. Depending on the amount of renewable energy, iSwitch migrates some virtual machines to grid-powered clusters, or vice versa. However, for the servers powered by renewable energy sources, the fluctuation of renewable power still decreases the stability of clusters. Moreover, this work fails to use the energy buffering function of batteries, and the servers hence need to constantly convert power sources due to the fluctuation of wind power, which introduces high operation overheads of virtual machine migration to clusters. Furthermore, workload-based scheduling algorithms [5], [14], [15] allow the workloads to match renewable power and postpone workloads until the renewable power is sufficient or the electricity price is low before the soft-deadline of batch jobs. But these scheduling algorithms fail to consider the impact of renewable power fluctuation on the stability of the grid and data centers.

Our Solution. Unlike existing schemes, we propose an

This is an extended version of our paper published in the Proceedings of the 39th IEEE International Conference on Distributed Computing Systems (ICDCS), 2019, pages: 249-260.

efficient smooth renewable power-aware scheme, called Smoother [16], which contains two main components, i.e., Flexible Smoothing (FS) and Active Delay (AD). Flexible Smoothing in Smoother smooths the frequent fluctuation of renewable power and provides a relatively stable supply using energy storage devices (ESDs) of data centers with finite battery capacity and charging/discharging rate. We formulate the problem of Flexible Smoothing in Smoother into a constrained nonlinear programming problem which determines the optimized charge/discharge scheme for renewable energy. Moreover, Active Delay in Smoother improves the utilization of the smoothed renewable power. Active Delay in Smoother adjusts the execution time of deferrable workloads, such as batch work, off-line tasks and low-priority jobs to match the generation of renewable energy, thus improving the utilization of renewable power.

Contributions. Our paper contains the following contributions:

1) *Cost-efficient Synergized Design.* We propose an easyto-use and efficient smooth renewable power-aware scheme, called Smoother, which consists of Flexible Smoothing (*FS*) and Active Delay (*AD*). In order to smooth the fluctuation of renewable power and provide a relatively stable supply to data centers (challenge 1), Flexible Smoothing in Smoother carries out the optimized charge/discharge operation via computing the minimum variance of the renewable power that is supplied to data centers per interval. In order to improve the utilization of the smoothed renewable power (challenge 2), Active Delay in Smoother matches deferrable workloads of data centers with the generation of renewable power.

2) *Practical Mitigation Scheme.* Smoother mitigates the impact of renewable power fluctuation on the stability of both grid and data centers using cost-efficient ESD management. We explore and exploit finite battery capacity, as well as the limited charging/discharging rate, which meet the demands of real-world situations.

3) *Suitable for A Variety of Renewable Energy.* Smoother can be used for a variety of renewable power sources, such as wind power and solar power via executing similar operations.

4) *Real-World Traces and Open-Source Code*. Real-world traces from workloads and renewable power generation are used in our extensive experiments to show the efficiency of Smoother. Evaluation results demonstrate that Smoother significantly reduces the negative impact of renewable power fluctuations on data centers, and improves the utilization of renewable power by 250.88% on average. We have released the source codes for public use in GitHub.

The rest of this paper is organized as follows. We present the background of data center energy management in Section 2. In Section 3, we present Smoother design. The experimental results are shown in Section 4. Section 5 shows the related work. Finally, we conclude this paper in Section 6.

2 BACKGROUND

In this Section, we present the backgrounds of power systems of data centers. First, we give a brief description of data centers' power infrastructure, and then discuss the characteristics of renewable energy supply and power consumption in data centers.

2.1 Power Infrastructure in Data Centers

As energy consumers, data centers have two main components of power consumption: IT equipments and cooling devices [17]. These IT equipments contain all servers supporting data computation and storage, as well as networking devices for data communications. In addition, cooling devices need to be installed in the machine room, which extracts heat released from IT equipments and decreases the room temperature. According to the Report to Congress on Server and Data Center Energy Efficiency [17], the cooling system consumes a significant amount of energy, which is close to 30% of total power consumption in data centers. In general, there is a micro power grid that integrates all power supplies, such as the electric grid, diesel generator and renewable energy generators, which allow power infrastructure to generate and distribute power for IT equipments and cooling devices. In order to guarantee the availability, data centers need to rely on brown energy resources, including the grid and diesel powers. However, when there are sufficient renewable sources, it is more sustainable to first consider to leverage renewable energy, rather than the grid or diesel generators, to save energy consumption. Although the renewable power is superior to brown power, some disadvantages, like time-variant and non-dispatchable, need to be considered for optimization.

2.2 Renewable Energy Sources

Owing to the characteristics of the brown energy, more renewable energy has been used in the power consumption of data centers to lower their operating costs as well as alleviate their impacts upon environments.

Wind Power: Wind resource characteristics and turbine properties become the main factors to evaluate the quantity of power generated from wind energy sources [18]. There are two features about wind power [19]. First, the power converted from the wind resource is a fraction, which largely depends on the wind speed variance and is defined as power coefficient. Second, different types of wind turbines have different power coefficient curves due to different cutin speeds and/or rated output power/speeds. Thus, the output power $P_{wind}(\nu)$ of a turbine type, with respect to wind speed ν , can be expressed by a piecewise function:

$$P_{wind}(\nu) = \begin{cases} 0 & \nu \leq V^{in} \\ G(\nu) & V^{in} < \nu \leq V^{rate} \\ P^{rate} & V^{rate} < \nu \leq V^{out} \\ 0 & \nu > V^{out} \end{cases}$$
(1)

where V^{in} and V^{out} are the cut-in and cut-out speeds, respectively. V^{rate} is the rated speed, P^{rate} is the rated power, and $G(\nu)$ is defined as the power curve between the cut-in and rated speeds. We use an example to illustrate the concept of the cut-in, rated, and cut-out speeds. For example, when the wind speed is 3 m/s (the cut-in speed), the wind turbine starts to generate usable power. Then the output power $G(\nu)$ increases with the growth of wind speed until the wind speed reaches 14 m/s (the rated speed). The corresponding output power is 800 KW (the rated power). When the wind speed increases more than the rated speed, the output power maintains at the rated power. When the wind speed reaches 25 m/s (the cut-out speed), the wind turbine has to be shut down to prevent damages.

The optimal wind power generation scheme [18] tries to leverage multiple curve fitting methods for $G(\nu)$. Compared with polynomial regression and exponential fitting, Gaussian Regression performs the best simulation of the relations between power coefficient and wind speed, i.e.,

$$G(\nu) = a_1 e^{-\frac{(\nu-b_1)^2}{c_1^2}} + \dots + a_n e^{-\frac{(\nu-b_n)^2}{c_n^2}} \quad 1 \le n \le 5 \quad (2)$$

where a_i , b_i and c_i are parameters to be fitted according to the real-world turbine type.

Solar Power: In addition to wind power, solar resource is another renewable power which is widely used to supply data centers. There are two critical factors, including solar radiation intensity and features of photovoltaic (PV) solar cells, to determine the amount of power generated from solar resource. One factor, i.e. solar radiation intensity, depends on weather conditions. By analyzing the solar traces in National Solar Radiation Data Base [20], the irradiance follows a clear daily periodical distribution. The availability of solar energy is merely during daytime and the radiation intensity is zero during night. Another factor related with solar power is the features of PV solar cells, which can convert sunlight directly into electricity. To boost the output of PV cells, they are connected together to form larger units known as panels [21]. We estimate the output power converted from solar irradiation [22] by

$$P_{solar} = A \cdot r \cdot H \cdot PR \tag{3}$$

where A is the total solar panel area (m^2) ; r indicates the ratio of the solar panel yield which is defined as electrical power of solar panel divided by the panel area; H is the average solar radiation on titled panels; PR (performance ratio) is the coefficient of solar energy losses, which evaluates the quality of photovoltaic installation.

In addition, Solar-TK [23], [24] and PVlib [25] are stateof-the-art schemes for modeling and forecasting solar power. In our evaluation, we predict solar power using both a model based on Equation (3) (e.g., Fig. 9, Fig. 14 and Fig. 17) and Solar-TK [24] (e.g., Fig. 15, Fig. 16 and Fig. 18), showing that different models don't have impacts on our proposed Smoother.

2.3 The Power Consumption

For a data center, the main power consumption comes from IT and cooling equipments. The power usage effectiveness (PUE), denoted by R_{pue} , represents the power ratio of two components, which is equal to the ratio of the data center's total power usage $P_{datacenter}$ to the power usage of IT equipments P_{IT} . Thus, the total power consumption of the data center at interval *t* can be estimated as [2]:

$$P_{datacenter}(t) = P_{IT}(t) * R_{pue} \tag{4}$$

where the total energy use of IT equipment P_{IT} is defined as the combined energy use of servers (including data processing servers and data storage servers) and networking devices [17]:

$$P_{IT}(t) = P_{server}(t) + P_{network}(t)$$
(5)

The power consumption of networking equipment is approximately less than 10% of the total peak power of all servers, which usually can be estimated as a constant.

We assume that there are N machines assembled at a data center and all machines have similar hardware configurations, i.e., each machine consumes the same power at the same central processing unit (CPU) utilization. The power consumptions of all servers are the sum of all machines' power, and the power consumed by individual machine is linearly scaled by CPU utilization shown in [2]:

$$P_{server}(t) = p^{idle} + (p^{full} - p^{idle}) * \mu$$
(6)

$$\mu(t) = \frac{1}{N} \sum_{i=1}^{N} \mu_i(t) \quad \{i = 1, 2, ..., N\}$$
(7)

where p^{idle} and p^{full} are the powers used by all machines at idle and fully utilized states, respectively. μ_i represents the machine i's CPU utilization, and μ is the average CPU utilization of all machines.

Evaluating the Effectiveness of Model-Based Power Characterization [26] demonstrates various models for power characterization. The results and the in-depth analysis in the paper show that for modern platforms, the modeling technique suffers from a basic error that can't be overcome by adding complexity. Considering that workloads with CPU utilization are easy to obtain and portable across heterogeneous platforms, in this paper, we use a widelyused power model [27], [28], [29], [30], i.e., Equations (4) – (7), to convert the CPU utilization into power traces.

3 THE DESIGN OF SMOOTHER

3.1 Design Goals

In order to address the problems of high fluctuation and low power utilization of renewable power, our paper has two design goals:

1) Alleviate fluctuation in renewable energy supply, thereby mitigating the impact of renewable power fluctuation on the stability and safety of both grid and data centers as well as reducing the overhead of energy switching in data centers (e.g., the overhead of virtual machine migration between grid-powered clusters and renewable-energy-powered clusters).

2) Improve the utilization of renewable power by alleviating the imbalance between supply and demand, making full use of renewable power in data centers.

The idea behind Smoother is to alleviate frequent fluctuation of renewable power and improve the utilization of renewable power. Specifically, we divide the renewable power supply into three regions and put forward the corresponding processing schemes for different regions. Smoother consists of two main components, including Flexible Smoothing (*FS*) and Active Delay (*AD*). Flexible Smoothing in Smoother provides a relatively stable renewable power supply by using energy storage devices, thus offering the smoothing effect. Active Delay in Smoother matches the deferrable workloads with renewable power generation and thus improve the utilization of renewable power.



Fig. 1. The normalized variance of capacity factor for solar power for every two hours in June 2015, latitude 42° and longitude -72° [24].



Fig. 2. Differentiated regions in solar power trace.

3.2 Regions Division in Power Generation

3.2.1 Regions in Solar and Wind Power Traces

Based on the degree of renewable power fluctuation, we divide renewable power traces into different regions. Optimal Harvesting Wind Power [18] introduces the concept of the capacity factor, which is the ratio of the actual output power P(t) to the rated output power P^{rate} . The higher the capacity factor is, the more electricity is actually generated. In this paper, we use the capacity factor variance σ_{power}^2 during an interval [0,T] to represent the fluctuation of renewable power supply. We calculate the capacity factor variance σ_{power}^2 :

$$\sigma_{power}^{2} = \frac{1}{T} \sum_{t=0}^{T} (\frac{P(t)}{P^{rate}} - \mu_{cf})^{2}$$
(8)

where μ_{cf} is the average capacity factor during an interval [0, T]:

$$\mu_{cf} = \frac{1}{T} \sum_{t=0}^{T} \frac{P(t)}{P^{rate}}$$
(9)

Fig. 1 and 3 show the normalized variances of capacity factor, i.e., the ratio of the standard deviation to the mean, for solar and wind power within a month. We observe that the wind power fluctuates more frequently than solar power. Therefore, we divide solar and wind power traces into two and three regions respectively.

Two Regions in Solar Power Traces. By analyzing the solar traces in National Solar Radiation Data Base [20], the solar power follows a clear daily periodical distribution, which means the solar energy is merely available during daytime. Fig. 2 shows a real-world example of solar power trace. we divide solar power traces into two regions.

In Region-I, compared with other regions, the solar power supply is relatively stable and unnecessary to be processed by Flexible Smoothing in Smoother. Specifically, Region-I consists of two situations according to the daily



Fig. 3. The normalized variance of capacity factor for wind power in each hour in May 2011, California, USA [31].



Fig. 4. Differentiated regions in wind power trace.

periodical distribution of solar power: 1) Solar power is unavailable during night because the radiation intensity is zero. 2) The solar power is stable when the solar radiation remains constant during daytime.

In Region-II, the fluctuation of solar power often results in the instability of power supplies in data centers, as discussed in Section 1. We hence consider how to decrease the degree of solar power fluctuation during this region. Specifically, the fluctuation of solar power is relatively moderate compared with that of wind power. It is hence feasible to use the uninterrupted power supply (UPS) battery to complement this fluctuation. We adopt Flexible Smoothing in Smoother to provide a relatively stable supply of renewable energy, as described in Section 3.3.

Three Regions in Wind Power Traces. Fig. 4 shows a real-world example of wind power trace, which comes from National Renewable Energy Laboratory (NREL) [31]. Based on the degree of wind power fluctuation, we divide wind power traces into three regions.

Specifically, Region-I consists of two situations of wind power generation according to Equation (1): 1) Wind power is nearly unavailable when the wind speed is less than the cut-in speed or larger than the cut-out speed. 2) The wind power is stable at the designated rated power of wind turbines when the wind speed is between the rated speed and the cut-out speed. In Region-I, the wind power supply is stable and unnecessary to be processed by Flexible Smoothing.

In Region-II, the wind power fluctuates frequently. In order to obtain a suitable trade-off between the smoothing effect and the required maximum rate of charging/discharging battery (and the resulting battery capacity), we further divide Region-II into two sub-regions, i.e., Region-II-1 and Region-II-2, based on the degree of wind power fluctuation.

In Region-II-1, the fluctuation of renewable power is relatively moderate compared with that in Region-II-2. We



Fig. 5. The cumulative distribution function of capacity factor variance in May 2011, California [31].

process the wind power supply in Region-II-1 like the solar power supply in Region-II, and execute Flexible Smoothing in Region-II-1 to provide a relatively stable supply of wind power.

In Region-II-2, renewable power fluctuates too frequently. In order to alleviate the fluctuations, Region-II-2 needs to carry out higher battery charging/discharging rate and battery capacity than other regions. Considering the tradeoff between the smoothing effect and the required maximum rate of charging/discharging battery (and the resulting battery capacity), we do not execute Flexible Smoothing in Smoother in Region-II-2.

In summary, we set the proportion of Region-II-2 to 0%-5% of the total regions of renewable power trace according to the actual situation. Active Delay in Smoother in Section 3.4 is used to address the mismatch between the smoothed renewable power and workload demand, thus increasing the utilization of renewable power.

3.2.2 Distinguishing Different Regions

We distinguish different regions by defining the thresholds of the variance of capacity factors. A specific threshold can be determined via the supply history of renewable energy and the required maximum rate of charging/discharging battery.

We use an example to illustrate how to use the thresholds of the capacity factor variance to distinguish Region-I, Region-II-1 and Region-II-2. The effects of different thresholds are discussed in the next subsection.

In order to illustrate how to use the thresholds of the variance to distinguish different regions, we leverage the wind power data every five minutes in May 2011, California, USA from NREL [31] and then calculate the variance of capacity factor in each hour. Figure 5 shows the cumulative distribution function (CDF) of the above hourly values of variance in May 2011. In Figure 5, the probability reaches 95% when the variance of capacity factor is smaller than $2 * 10^{11}$ (the threshold between Region-II-2 and Region-II-1). The probability is 30% when the variance of capacity factor is smaller than $4.45 * 10^8$ (the threshold between Region-II-1 and Region-I). When the variance of capacity factor is larger than $2 * 10^{11}$, we observe the CDF curve increases slowly in Figure 5, meaning that this region accounts for a low proportion of the entire region. At the same time, as shown in Fig. 6, the required maximum rate of charging/discharging battery (and the resulting battery capacity) grows rapidly when the CDF of the variance increases, which is used



Fig. 6. The effects of different variance thresholds for distinguishing Region-II-1 and Region-II-2 on the optimization results and battery charging/discharging rate requirements. W/ Smooth and W/O Smooth means with and without performing Flexible Smoothing in Smoother respectively. Battery MaxVol represents the required maximum battery charging/discharging rate for Flexible Smoothing, which also represents the trend of the required battery capacity.

to distinguish Region-II-1 and Region-II-2 (For example, a CDF value of 0.95 means Region-II-2 accounts for 5% of the total regions). At the same time, the "energy switching times" becomes smaller as the CDF value increases. Since there is "switching times" in Fig. 6, we introduce the concept of "energy switching times", which is interpreted as the virtual machine migration times between grid-powered clusters and renewable-energy-powered clusters. Frequent fluctuations of renewable energy usually mean frequent load migration, which increases operation overheads. Thus we use the metric "energy switching times" in the power supply between the grid and renewable power to represent the impact of renewable power fluctuation on data centers. For the difference between Region-II-1 and Region-I, a smaller probability of Region-I means a larger number of charging/discharging operations in Flexible Smoothing in Smoother. However, frequent charging and discharging operations exacerbate battery lifetime and increase energy loss [32]. Therefore, in a real-world system, we need to trade off the battery consumption and energy switching overhead.

3.3 Flexible Smoothing in Smoother

The main idea of Flexible Smoothing in Smoother is to alleviate the frequent fluctuation of renewable power in Region-II-1 and achieve a relatively stable supply of renewable energy to data centers. Renewable energy generation can be estimated by Equations (1) - (3) in Section 2.2 given meteorological data [18], [22]. On the other hand, the future generation of renewable energy can be predicted by lots of methods [23], [24], [25], [33], [34], [35], [36], [37]. For example, LSSVM-GSA model for wind power prediction achieves less than 15% absolute error within 48 hours for 100% data points, and less than 10% absolute error for 97.92% data points [36]. The SUNY solar forecast model reports 13.4% – 17.4% MAPE and 21% – 25.3% rRMSE for 3-hours ahead solar forecasts [33]. Note that the prediction mechanism is out of the scope of this paper, and we leverage a widely-used method based on Equations (1) - (3) in Section 2.2 and Solar-TK [24] to predict renewable power.

At the beginning of each interval (e.g., one hour), Flexible Smoothing in Smoother determines the optimized battery charging/discharging scheme in the incoming interval. Through the implementation of the optimized battery



Fig. 7. The smoothed wind power (W/ FS) versus the original wind power (W/O FS). The region in Region-II-1 becomes stable like Region-I after using Flexible Smoothing in Smoother.

charging/discharging operation, the data center can obtain a relatively smooth renewable energy supply.

Suppose that there are *m* time points (e.g., m = 12) in the interval [1, m]. $U = [u_1 u_2 ... u_m]^T$ represents the amount of the generated renewable energy at each time point (e.g., five minutes). $S = [s_1 s_2 ... s_m]^T$ represents the battery charging/discharging amount at each time point. A positive s_i indicates that the battery is discharged with the amount of $|s_i|$ at time point *i*, and a negative s_i indicates the battery is charged with the amount of $|s_i|$ at time point *i*. We use *A* to represent the final amount of renewable energy supplied to the data center after the battery charging/discharging operation at each time point in the interval [1, m]:

$$A = U + S = \begin{bmatrix} a_1 \\ a_2 \\ \dots \\ a_m \end{bmatrix} = \begin{bmatrix} u_1 + s_1 \\ u_2 + s_2 \\ \dots \\ u_m + s_m \end{bmatrix}$$
(10)

We formulate the problem of Flexible Smoothing in Smoother into the constrained nonlinear programming problem:

min
$$\sigma_A = \sqrt{\frac{1}{m} \sum_{i=1}^m (a_i - \bar{a})^2}$$
 (11)

subject to:

$$\forall i \in [1, m], \begin{cases} 0 \le |s_i| \le 0.8M, \quad s_i \ge 0\\ 0 \le |s_i| \le u_i, \quad s_i < 0 \end{cases}$$
(12)

$$\forall i \in [1, m], 0.2M \le |\sum_{t=1}^{i} s_t| \le M$$
 (13)

where M is the battery capacity, and \bar{a} is the average value of a_i ($i \in [1, m]$). We assume that each time point is five minutes, and m is equal to 12 since the decision is computed each hour. Equation (12) means that at each time point, the amount of charging battery can not exceed that of generating renewable energy at that time. The amount of discharging battery can not exceed 80% of the battery capacity to avoid the damage of full discharge on the battery [38]. Equation (13) indicates that the accumulated electric quantity of the battery can not exceed the battery capacity or be less than 20% of the battery capacity during the interval [1, m]. We use MATLAB to solve this nonlinear programming problem with negligible computational overhead.

It is worth noting that the rate limits of charging/discharging battery are implicitly considered in our



Fig. 8. The imbalance between workload demand and renewable power supply. The renewable energy in the green area can not be used by data centers.

model proposed above. In our implementation in Section 4, the battery capacity is set to sustain half interval (e.g., half an hour) of operations at the maximum rate of charging/discharging battery, while the decision is computed each interval (e.g., one hour). Actually, the larger battery capacity will yield the better smoothing effect.

Based on the results of Equations (10) - (13), we execute the battery charging/discharging operations on the UPS battery. Fig. 7 shows the smoothed wind power versus the original wind power, where *W*/*FS* means "*with* Flexible Smoothing", which represents the final wind power supply after performing Flexible Smoothing in Smoother. *W*/*O FS* means "*without* Flexible Smoothing", representing the original output power of wind turbines. After performing Flexible Smoothing in Smoother in Region-II-1, the fluctuation of the renewable energy supply becomes moderate like Region-I, as shown in the red dotted circle in Fig. 7.

3.4 Active Delay in Smoother

The main idea of Active Delay in Smoother is to match the deferrable workloads with the smoothed renewable power generation under the condition of meeting the soft-deadline of deferrable workloads, thus improving the utilization of renewable power.

Random requests from users often result in the fluctuation of workload power demands in data centers. The fluctuations of workload power demand and the smoothed renewable power supply are different and thus possibly lead to the imbalance between supply and demand, which fails to make full use of renewable energy in data centers. As shown in Fig. 8, the workload power demand and the renewable power supply change over time. When the renewable power supply is larger than the workload power demand as shown in the green area, the excess renewable energy can not be used by data centers.

Real-time workloads, such as Web access [39], need to be executed immediately once a request is issued. However, for the deferrable workloads such as batch jobs, their characteristics provide opportunities to increase the utilization of the renewable energy. When the renewable power supply is insufficient, the workloads can be deferred to a future time slot. In many workload schedulers, the deadline and the expected running time for each job can be provided by users or estimated using historical statistics [14]. Therefore, we assume to obtain the above information from users or historical statistics.

We present the details of Active Delay in Smoother in Algorithm 1. Note that when using Fixed pricing (FP) or

Tiered Pricing (TP) as electricity Pricing, increasing renewable power utilization can reduce electricity bills. At this point, the main idea of Active Delay in Smoother is to schedule jobs for increasing renewable power utilization. However, Time-of-Use pricing (ToU) has three periods (i.e., on-peak, mid-peak and off-peak) with different grid prices, and scheduling jobs only for the highest renewable energy use may incur higher grid costs. Therefore, when using ToU, Active Delay in Smoother is to schedule jobs for reducing grid costs.

Specifically, in Algorithm 1, the requestJob (Line 1) is a workload request queue which contains the original requests for each job. The queueJob (Line 2) contains the ordered jobs (requests) that will be scheduled in sequence into the following time. For every small time slot (e.g., one minute), we first calculate the power quantity that needs to be consumed by each job in requestJob (Line 6). More details are described in Section 2.3. Then, each job is inserted into the queueJob in the ascending order of its slack-time, which is defined as deadline minus the sum of running time and current time (Lines 7-8). We schedule each job in this queue in sequence. For each job in the queueJob, we determine whether the slack-time of the current job is larger than 0 (Lines 11–12). Slack-time > 0 means that the current job is a non-real-time job and can be scheduled into the following time. If so, we traverse the slack-time and choose the time with the lowest grid costs (using ToU, Lines 13-18) or with the highest renewable energy use (using FP or TP, Lines 20-25) as the real execution time for the current job. It is worth noting that the scheduling of the current job is based on the scheduling results of previous jobs (i.e., Active Delay in Smoother maintains renewable power for previous jobs) in queueJob, and there are no conflicts of the optimal execution time among jobs. Line 19 (or Line 26) updates the remaining renewable power after the scheduling (i.e., maintaining resources for the current job). But if the slacktime of the current job is not larger than 0, this job needs to be carried out immediately (Lines 28–30).

4 **PERFORMANCE EVALUATION**

Smoother consists of Flexible Smoothing (FS) and Active Delay (AD). Flexible Smoothing in Smoother is the design focus of this paper, which aims to alleviate fluctuations of renewable power, thereby reducing operation overheads of data centers (e.g., the overhead of virtual machine migration between grid-powered clusters and renewable-energypowered clusters). Active Delay in Smoother improves the utilization of renewable power. We examine the performance of Smoother in terms of multiple metrics, including energy switching times between the grid and renewable power in the power supply via real-world workloads and renewable power traces, which represent the impact of renewable power fluctuation on data centers, renewable energy utilization and grid costs

4.1 Experimental Setup

Data Center Configurations: We assume that a data center is equipped with 11,000 servers and each server has the same processing power [2], i.e., identical energy consumption for Algorithm 1 Active Delay in Smoother.

8
Input: Data center workload requests and the smoothed
renewable power supply
Output: The execution time of each job and the optimal use
of the grid and renewable power
1: queue <job> requestJob;</job>
2: priority_queue <job> queueJob;</job>
3: for each small time slot do
4: if requestJob.size() > 0 then
5: for job: requestJob do
6: calWorkloadPower(job);

- calWorkloadPower(job);
- 7: slackTime = job.deadline - job.runTime - currentTime; queueJob.push(job, slackTime); 8:
- requestJob.pop(); 9:

10: end for

for job: queuelob do 11.

· · ·	ior jour queuejou no			
12:	if job.slackTime > 0 then			
13:	if electricityPricing $==$ ToU then			
14:	for time: slackTime do			
15:	costs = getGridCostsWhenStartAt-			
	Time(job,time);			
16:	end for			
17:	startTime = getStartTimeOfJobForMini-			
	mumGridCosts(job);			
18:	executeJob(job, startTime);			
19:	updateRemainRPower(job, startTime);			
20:	else			
21:	for time: slackTime do			
22:	power = getRenewablePowerUsedWhen-			
	StartAtTime(job, time);			
23:	end for			
24:	startTime = getStartTimeOfJobForMaxi-			
	mumRenewablePower(job);			
25:	executeJob(job, startTime);			
26:	updateRemainRPower(job, startTime);			
27:	end if			
28:	else			
29:	executeJob(job, current'lime);			
30:	updateRemainRPower(job, currentTime);			
31:	end if			
32:	queueJob.pop();			
33:	end for			
34:	end if			
35:	end for			

executing the same jobs. The peak power of each server is 186W and the idle power is 62W [12]. Each server runs the same workloads and has similar CPU utilization, like existing experimental configurations [12].

Workload Traces: Our evaluation uses three typical realworld workloads, including Web workloads [39], Google cluster-data [40] and batch workloads [41]. Specifically, Google trace provides real-world power consumption from a 12,500-machine cell over about a month-long period in May 2011 [40]. According to Equations (4) - (7) in Section 2.3, we convert the CPU utilization into power trace. The time-sensitive Web workloads are collected from the real-world logs provided by the Internet Traffic Archive [39]. We generate the number of requests per minute from the

request logs and convert this number into CPU utilization via a linear analog [12], which sets the utilization as 100% when the request rate is the maximum and 0% when the request rate is the minimum. As shown in Table 1, we select five Web workload traces with different average CPU utilizations during one week. The third workloads are batch workloads. The logs of Real Parallel Workloads from Production Systems [41] provide system information, such as average CPU utilization, work arrival time, execution time, deadline and the number of the required servers. For each job, the energy consumption requirements can be calculated by the CPU utilization, the number of servers used, etc. In our simulation, we choose four batch workload traces with different CPU utilizations as shown in Table 2.

TABLE 1 Five Web workload traces with different average CPU utilizations [39].

Web	Web Description	
Calgary	CS departmental Web server	3.63%
U of S	University Web server	7.21%
NASA	Kennedy Space Center Web server	28.89%
Clark	ClarkNet Web server	35.78%
UCB	UC Berkeley IP Web server	46.04%

TABLE 2 Batch workload traces with different CPU utilizations [41].

Batch Workload Traces	Avg. CPU utilization
LLNL_Thunder	86.7%
LANL_CM5	74.4%
HPC2N	60.1%
Sandia Ross	49.9%

TABLE 3 Renewable power traces with different volatility [42] [20].

Renewable Power Traces	Site ID (Wind)	Date (Solar)
Low volatility	CA(9122)	May 8
	OR((24258)	May 17
	WA(29359)	May 22
	TX(10)	May 9
High volatility	CO(11005)	May 18
	WY(16419)	May 23

Renewable Power Traces: Wind power traces come from Wind Data Resources [42] of the National Renewable Energy Laboratory (NREL) [31]. Solar power traces come from National Solar Radiation Data Base [20]. As shown in Table 3, we select two groups of renewable power traces that are differentiated by volatility intensity. The low volatility traces have relatively stable and smooth generation, while the high volatility traces have output rate with high volatility. In addition, we use Solar-TK [23], [24] to generate a solar power trace in June 2015, latitude 42° and longitude -72°.

4.2 Results and Analysis

We use energy switching times between the grid and wind power in the power supply of server clusters to represent the impact of renewable energy fluctuations on a data center. Smoother is compared with the standard battery storage scheme where the data center first uses renewable power as



(a) Solar power traces on May 9, 17, 18 and 22, 2011, California, USA [20].



Fig. 9. The comparison of energy switching times with different solar fluctuation degrees between "W/O FS" scheme and "W/ FS" scheme.

much as possible without the battery, and then the battery stores the remaining renewable power and discharges when the renewable power is insufficient. This battery storage solution is efficient and adopted by many works such as Multigreen [7]. We use *Comp* to represent the standard battery storage scheme. *W*/ means "*with*", and *W*/*O* means "*without*".

4.2.1 Evaluation of Flexible Smoothing in Smoother

We present the performance evaluation of Flexible Smoothing in Smoother via the metric of energy switching times. As described in Section 3.2, we set the capacity factor variance whose corresponding CDF value is 0.95 as the threshold to distinguish Region-II-1 and Region-II-2. We execute Flexible Smoothing in Smoother in Region-II-1. Flexible Smoothing in Smoother is able to achieve a smooth and stable renewable energy supply by pre-calculating the optimized charge/discharge strategy of the battery.

In Fig. 9, we obtain four-day solar power with different fluctuation degrees in May 2011, California, USA [20], and compare the energy switching times after carrying out Flexible Smoothing in Smoother with the initial energy switching times within four days. From Fig. 9(a), we observe that solar power fluctuates most frequently on May 18. By using Flexible Smoothing in Smoother, the energy switching times significantly decrease on May 18, but the decreased degrees become smaller on other days, as shown in Fig. 9(b). The main reason is that Flexible Smoothing in Smoother aims to alleviate the fluctuation degree of renewable power. Hence if the renewable power is smooth, Flexible Smoothing in Smoother plays the less important role.

In the similar way, we use five Web workload traces in Table 1 and two groups of wind power traces in Table 3 respectively for comparisons in terms of energy switching times. When the total installed wind turbine capacity is 976KW, the results are shown in Fig. 10 and Fig. 11. When the total installed wind turbine capacity is 1525KW, the



Fig. 10. The comparison of energy switching times between "W/ *Comp*" scheme and "W/ *FS*" scheme with different workloads (The rated output power of wind energy is 976KW).



Fig. 11. The comparison of energy switching times between "W/ Comp" scheme and "W/ FS" scheme with different wind power traces (The rated output power of wind energy is 976KW).

results are shown in Fig. 12 and Fig. 13. We observe that with different workloads and renewable power traces, Flexible Smoothing in Smoother can effectively reduce the energy switching times compared with the general battery storage solution, thus significantly reducing the negative impact of renewable power fluctuation on data centers. Fig. 11 and Fig. 13 also demonstrate that Flexible Smoothing in Smoother brings more remarkable effect when processing renewable traces with high volatility.

4.2.2 Evaluation of Active Delay in Smoother

Flexible Smoothing in Smoother aims to alleviate fluctuations of renewable power, while Active Delay in Smoother aims to improve the utilization of renewable power. We examine the performance of Active Delay in Smoother in the metric of renewable power utilization and grid costs.

Table 2 shows batch workload traces with different CPU utilizations. As shown in Fig. 14, with different workloads under different renewable energy supplies, the utilization of renewable power increases by an average of 250.88% by using Active Delay in Smoother. Whether or not the renewable power is sufficient, all the workloads increase the utilization of renewable power after performing Active Delay in Smoother. Fig. 15 uses one-month solar power trace in June 2015 [24] and the Sandia Ross batch workload trace. Active Delay in Smoother increases the utilization of renewable power by an average of 241.09%. Moreover, we use AD_ToU to represent the Active Delay algorithm with Time-of-Use pricing. Since the time of the solar power traces is June 2015, we adopt the corresponding historical ToU prices of Ontario Hydro [43], i.e., 12.2 cents/kWh for midpeak period (7:00 AM - 11:00 AM and 5:00 PM -7:00 PM on weekdays), 16.1 cents/kWh for on-peak period (11:00 AM - 5:00 PM on weekdays), and 8 cents/kWh for



Fig. 12. The comparison of energy switching times between "W/ *Comp*" scheme and "W/ *FS*" scheme with different workloads (The rated output power of wind energy is 1525KW).



Fig. 13. The comparison of energy switching times between "W/ *Comp*" scheme and "W/ *FS*" scheme with different wind power traces (The rated output power of wind energy is 1525KW).

off-peak period (7 : 00 PM – 7 : 00 AM on weekdays and all day on weekends and holidays). As shown in Fig. 16, the grid costs decrease by an average of 7.66% when using Active Delay in Smoother with ToU pricing. We argue that Active Delay in Smoother can significantly improve the utilization of renewable energy, reduce the use of brown energy and decrease the electricity costs of data centers and carbon emissions.

We also compare energy switching times between "W/O *FS* and W/ *AD*" and "W/ *FS* and W/ *AD*" schemes. As shown in Fig. 17 and 18, the "W/ *FS* and W/ *AD*" scheme significantly reduces the energy switching times by an average of 25% and 24.74% respectively, thus mitigating the negative impact of renewable energy fluctuations on data centers.

In summary, by using Flexible Smoothing and Active Delay, our proposed Smoother significantly reduces the negative impact of renewable power fluctuation on data centers and increases the utilization of renewable power, thus improving the stability of the grid and data centers and reducing the cost of data centers.

5 RELATED WORK

In order to meet the energy needs of data centers, existing schemes have been proposed by exploiting the economical and environment-friendly renewable powers in terms of leveraging energy storage devices and workload scheduling.

Energy Storage: Existing schemes generally leverage energy storage devices to reduce electricity costs [7], [13], [27], [44], [45], [46], [47], mitigate peak power [48], minimize carbon emissions of the grid [49] and eliminate the fluctuation of renewable energy [13], [44], [49]. The optimal control of



Fig. 14. The comparison of renewable power utilization between "W/ *FS* and W/O *AD*" scheme and "W/ *FS* and W/ *AD*" scheme with different workloads and renewable power traces [20], [31].



Fig. 15. The comparison of renewable power utilization between "W/ FS and W/O AD" scheme and "W/ FS and W/ AD" scheme with the Sandia Ross workload and one-month solar power traces [24].

end-user energy storage [46] stores energy at much lower prices in a battery that discharges when the energy prices are high to meet the demand for cost savings. Multigreen [7] proposes a cost-minimizing control algorithm to determine the amount of energy drawn from a long-term grid market, a real-time grid market and energy storage devices which store excess renewable power. Optimal power cost management [47] designs an algorithm for exploiting the UPS unit and delay-tolerance of workloads with time-varying power prices to reduce the electricity bill of a data center. The comprehensive understanding of operation cost reduction [45] conducts a quantitative analysis on normalized electricity price in the battery storage and thermal energy storage for internet data centers, and concludes that the cost of the energy storage devices are largely affected by the storage capacity and the location of data centers. The cost minimizing online scheme [13] stores the wind energy of a wind farm into batteries to eliminate the fluctuation of renewable power. However, storing all the wind energy needs a large battery capacity. In order to address the problem of peak power, EBuff [48] proposes a peak reduction algorithm leveraging energy storage. A low-complexity algorithm [27] schedules heterogeneous workloads with UPS system to minimize the energy cost of an internet data center. The batteries are charged when the renewable power is more than the power consumption or the electricity price is low. However, this work overlooks the fluctuation of the renewable power without using the batteries, which possibly incurs the instability of grid and data centers. Emissionaware Energy Storage Scheduling [49] leverages distributed energy storage and aims to minimize carbon emissions of the grid. GreenFlowing [44] is a scheduling scheme that leverages different types of ESDs to reduce the electricity cost for a cloud data center. These schemes with energy storage [44], [49] typically store all the available renewable



Fig. 16. The comparison of grid costs between "W/ FS and W/O AD_ToU" scheme and "W/ FS and W/ AD_ToU" scheme with the Sandia Ross workload and one-month solar power traces [24].



Fig. 17. The comparison of energy switching times between "W/O FS and W/ AD" scheme and "W/ FS and W/ AD" scheme with different workloads and renewable power traces [20], [31].

sources in batteries for later use, which requires a large storage capacity and high battery charging/discharging rate. Unlike them, Smoother alleviates the fluctuation of renewable energy supply in a complementary manner with finite battery capacity and charging/discharging rate.

Renewable Energy Generation: Optimal Harvesting Wind Power [18] and PV Solar Energy Calculation [22] show the computation of estimating the renewable energy generation with meteorological data. On the other hand, the future generation of renewable energy can be predicted by lots of methods [23], [24], [25], [33], [34], [35], [36], [37]. For example, LSSVM-GSA model for wind power prediction achieves less than 15% absolute error within 48 hours for 100% data points, and less than 10% absolute error for 97.92% data points [36]. The SUNY solar forecast model reports 13.4% - 17.4% MAPE and 21% - 25.3% rRMSE for 3-hours ahead solar forecasts [33]. Note that the prediction mechanism is out of the scope of this paper, and the above prediction methods can be integrated into our system. In this paper, we leverage a widely-used method based on Equations (1) – (3) in Section 2 and Solar-TK [24] to predict renewable power.

Workload Scheduling: In order to address the problem of the mismatch between data center workloads and green energy supplies, existing workload scheduling schemes have been proposed [5], [14], [15]. Their basic idea is to postpone the deferrable workloads until the renewable power is sufficient or the electricity price is low before the softdeadline of workloads. In addition, some schemes leverage geographical load balancing among distributed data centers to improve the utilization of renewable power [8]. However, these algorithms fail to consider the impact of renewable power fluctuation on the stability of the grid and data centers. A self-adaptive approach to managing applications and harnessing renewable energy [50] proposes



Fig. 18. The comparison of energy switching times between "W/O FS and W/ AD" scheme and "W/ FS and W/ AD" scheme with the Sandia Ross workload and one-month solar power traces [24].

a deferring algorithm for batch workloads and a brownoutbased algorithm for interactive workloads, which improve the renewable energy usage and reduce the carbon footprint of the data centers. However, the self-adaptive approach currently fails to deal with the renewable power fluctuation but includes a battery model as future work.

Compared with the conference version [16], the improved journal version includes the new description about solar power, the division of regions in the solar power traces, the adaptive Flexible Smoothing scheme, the improved Active Delay algorithm with ToU pricing, as well as many new evaluation results using real-world solar power traces.

6 CONCLUSION

Providing a smooth and stable supply of renewable power and improving the utilization of renewable power are important in modern data centers. We propose a smooth renewable power-aware scheme, called Smoother, which consists of Flexible Smoothing (FS) and Active Delay (AD). The novelty behind Smoother is that we emphasize the impact of frequent fluctuation of renewable power on the stability of the grid and data centers, as well as improving the utilization of renewable power. The trace-driven evaluation demonstrates that our proposed Smoother offers a smooth and stable supply of renewable power for data centers, reducing the energy switching times by 25% on average and improving the utilization of renewable power by an average of 250.88%. Smoother is able to improve system performance and the stability of the grid and data center systems, while reducing the costs of data centers. We have released the source code of Smoother for public use at https://github.com/csXinxinLiu/Smoother.

REFERENCES

- A. Qureshi, "Power-demand routing in massive geo-distributed systems," Ph.D. dissertation, Massachusetts Institute of Technology, 2010.
- [2] F. Kong and X. Liu, "A survey on green-energy-aware power management for datacenters," ACM Computing Surveys (CSUR), vol. 47, no. 2, p. 30, 2015.
- [3] M. H. K. Tushar, A. W. Zeineddine, and C. Assi, "Demand-side management by regulating charging and discharging of the ev, ess, and utilizing renewable energy," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 1, pp. 117–126, 2018.
- [4] T. Liu, H. Yu, H. Guo, Y. Qin, and Y. Zou, "Online energy management for multimode plug-in hybrid electric vehicles," *IEEE Transactions on Industrial Informatics*, 2018.
- [5] Q. Sun, S. Ren, C. Wu, and Z. Li, "An online incentive mechanism for emergency demand response in geo-distributed colocation data centers," in *e-Energy*, 2016.

- [6] D. van der Meer, G. R. C. Mouli, G. M.-E. Mouli, L. R. Elizondo, and P. Bauer, "Energy management system with pv power forecast to optimally charge evs at the workplace," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 1, pp. 311–320, 2018.
- [7] W. Deng, F. Liu, H. Jin, C. Wu, and X. Liu, "Multigreen: Costminimizing multi-source datacenter power supply with online control," in *e-Energy*, 2013.
- [8] Y. Zhang, Y. Wang, and X. Wang, "Greenware: Greening cloudscale data centers to maximize the use of renewable energy," in Distributed Systems Platforms and Open Distributed Processing, 2011.
- [9] J. P. Lopes, N. Hatziargyriou, J. Mutale, P. Djapic, and N. Jenkins, "Integrating distributed generation into electric power systems: A review of drivers, challenges and opportunities," *Electric power* systems research, vol. 77, no. 9, pp. 1189–1203, 2007.
- [10] L. Xie, P. M. Carvalho, L. A. Ferreira, J. Liu, B. H. Krogh, N. Popli, and M. D. Ilic, "Wind integration in power systems: Operational challenges and possible solutions," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 214–232, 2011.
- [11] Y. Wu, Z. Wei, J. Weng, X. Li, and R. H. Deng, "Resonance attacks on load frequency control of smart grids," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4490–4502, 2018.
- [12] C. Li, A. Qouneh, and T. Li, "iswitch: coordinating and optimizing renewable energy powered server clusters," in *ISCA*, 2012.
- [13] C.-K. Chau, G. Zhang, and M. Chen, "Cost minimizing online algorithms for energy storage management with worst-case guarantee," *IEEE Transactions on Smart Grid*, vol. 7, no. 6, pp. 2691–2702, 2016.
- [14] Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Greenhadoop: leveraging green energy in data-processing frameworks," in *Eurosys*, 2012.
- [15] A. Krioukov, S. Alspaugh, P. Mohan, S. Dawson-Haggerty, D. E. Culler, and R. H. Katz, "Design and evaluation of an energy agile computing cluster," *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2012-13, 2012.*
- [16] X. Liu, Y. Hua, X. Liu, L. Yang, and Y. Sun, "Smoother: a smooth renewable power-aware middleware," in ICDCS, 2019.
- [17] R. Brown et al., "Report to congress on server and data center energy efficiency: Public law 109-431," Lawrence Berkeley National Laboratory, 2008.
- [18] F. Kong, C. Dong, X. Liu, and H. Zeng, "Quantity versus quality: Optimal harvesting wind power for the smart grid," *Proceedings of the IEEE*, vol. 102, no. 11, pp. 1762–1776, 2014.
- [19] "Enercon," http://www.enercon.de/home, 2015.
- [20] "National Solar Radiation Data Base," http://rredc.nrel.gov/ solar/old_data/nsrdb/.
- [21] "Solar Photovoltaic Technology Basics," http://energy.gov/eere/ energybasics/articles/solar-photovoltaic-technology-basics.
- [22] "PV solar energy calculation," http://photovoltaic-software.com/ PV-solar-energy-calculation.php.
- [23] N. Bashir, D. Chen, D. Irwin, and P. Shenoy, "Solar-tk: A datadriven toolkit for solar pv performance modeling and forecasting," in *Proceedings of the 16th International Conference on Mobile Ad-hoc and Smart Systems (MASS19)*, 2019.
- [24] "The source code of Solar-TK." https://github.com/ sustainablecomputinglab/solar-tk.
- [25] W. F. Holmgren, C. W. Hansen, and M. A. Mikofski, "pvlib python: A python package for modeling solar energy systems," *Journal of Open Source Software*, vol. 3, no. 29, p. 884, 2018.
- [26] J. C. McCullough, Y. Agarwal, J. Chandrashekar, S. Kuppuswamy, A. C. Snoeren, and R. K. Gupta, "Evaluating the effectiveness of model-based power characterization," in USENIX ATC, 2011.
- [27] K. Haghshenas, S. Taheri, M. Goudarzi, and S. Mohammadi, "Infrastructure aware heterogeneous-workloads scheduling for data center energy cost minimization," *IEEE Transactions on Cloud Computing*, 2020.
- [28] W. Wu, W. Wang, X. Fang, L. Junzhou, and A. V. Vasilakos, "Electricity price-aware consolidation algorithms for time-sensitive vm services in cloud systems," *IEEE Transactions on Services Computing*, 2019.
- [29] T. Van Damme, C. De Persis, and P. Tesi, "Optimized thermalaware job scheduling and control of data centers," *IEEE Transactions on Control Systems Technology*, vol. 27, no. 2, pp. 760–771, 2018.
- [30] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 732–794, 2015.
- [31] "National Renewable Ênergy Laboratory," https://www.nrel. gov/.

- [32] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?" in SIGMETRICS, 2012.
- [33] R. Perez, J. Schlemmer, S. Kivalov, J. Dise, P. Keelin, M. Grammatico, T. Hoff, and A. Tuohy, "A new version of the suny solar forecast model: A scalable approach to site-specific model training," in PVSC, 2018.
- [34] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser, "Renewable and cooling aware workload management for sustainable data centers," in SIGMETRICS, 2012.
- [35] N. Sharma, P. Sharma, D. Irwin, and P. Shenoy, "Predicting solar generation from weather forecasts using machine learning," in SmartGridComm, 2011.
- [36] X. Yuan, C. Chen, Y. Yuan, Y. Huang, and Q. Tan, "Short-term wind power prediction based on lssvm-gsa model," Energy Conversion and Management, vol. 101, pp. 393-401, 2015.
- [37] M. Rana, I. Koprinska, and V. G. Agelidis, "Forecasting solar power generated by grid connected pv systems using ensembles of neural networks," in *IJCNN*, 2015.
- [38] "Northern Arizona Wind & Sun. Deep cycle battery," https:// www.solar-electric.com/deep-cycle-battery-faq.html. [39] "The Internet Traffic Archive," ftp://ita.ee.lbl.gov/html/contrib. [40] "Google cluster-data," https://github.com/google/cluster-data.

- [41] "Logs of Real Parallel Workloads from Production Systems." http: /www.cs.huji.ac.il/labs/parallel/workload/logs.html.
- [42] "Wind Data Resources of the National Renewable Energy Laboratory." https://www.nrel.gov/wind/data_resources.html.
- [43] "Ontario Hydro." http://www.ontario-hydro.com/ historical-rpp-rates.
- [44] 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 ISCC, 2019.
- [45] H. Zhou, J. Yao, H. Guan, and X. Liu, "Comprehensive understanding of operation cost reduction using energy storage for idcs," in INFOCOM, 2015.
- [46] P. M. van de Ven, N. Hegde, L. Massoulié, and T. Salonidis, "Optimal control of end-user energy storage," IEEE Transactions on Smart Grid, vol. 4, no. 2, pp. 789-797, 2013.
- [47] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam, "Optimal power cost management using stored energy in data centers," in SIGMETRICS, 2011.
- [48] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and limitations of tapping into stored energy for datacenters," in ISCA. 2011.
- [49] R. Jha, S. Lee, S. Iyengar, M. H. Hajiesmaili, D. Irwin, and P. Shenoy, "Emission-aware energy storage scheduling for a greener grid," 2020. [Online]. Available: https://doi.org/10.1145/ 3396851.3397755
- [50] M. Xu, A. N. Toosi, and R. Buyya, "A self-adaptive approach for managing applications and harnessing renewable energy for sustainable cloud computing," IEEE Transactions on Sustainable Computing, 2020.



Yu Hua is a professor at Huazhong University of Science and Technology. He was a postdoc research associate at McGill University in 2009 and a postdoc research fellow at University of Nebraska-Lincoln in 2010-2011. He obtained his B.E and Ph.D degrees from Wuhan University, respectively in 2001 and 2005. His research interests include file systems, cloud storage systems, non-volatile memory, big data analytics, machine learning, etc. His work has been published in multiple top-tier conferences and jour-

nals in the area, including OSDI, FAST, MICRO, USENIX ATC, ACM SoCC, SC, HPDC, ICDCS, IPDPS, MSST. He serves as PC chairs, committee members for multiple international conferences, including ASPLOS, SOSP, SC, USENIX ATC, ICS, RTSS, ICDCS, INFOCOM, IPDPS, MSST, DATE. He is a distinguished member of CCF, a senior member of ACM and IEEE, and a member of USENIX. He serves as the Distinguished Speaker of ACM and CCF.



Xue Liu is a professor in the School of Computer Science and a William Dawson Scholar at McGill University. He obtained his Ph.D. in Computer Science from The University of Illinois at Urbana-Champaign and his B.S. degree in Mathematics and M.S. degree in Automatic Control both from Tsinghua University, China. He has published over 250 research papers in major highlyreputable international academic journals and conference proceedings. He has also authored 3 books/monographs and several book chapters.

Dr. Liu has served on the organizing committees of 40 major international conferences and workshops. Dr. Liu is an editor/associate editor of IEEE/ACM Transactions on Networking (ToN), ACM Transactions on Cyber-Physical Systems (TCPS), IEEE Transactions on Vehicular Technology (TVT), and IEEE Communications Surveys and Tutorial (COMST). Dr. Liu is a member of the Tau Beta Pi Engineering Honorary Society, a lifetime member of the ACM, a member of the IEEE, and USENIX.



Ling Yang received the B.E degree in software engineering from China University of Geosciences (CUG), China, in 2014, and received the master degree in computer system architecture from Huazhong University of Science and Technology (HUST), China, in 2017. She has published several papers in major conferences and journal including MSST, ICDCS and TPDS.



Xinxin Liu received the B.E degree in software engineering from Central China Normal University (CCNU), China, in 2016. She is currently a PhD student majoring in computer system architecture at Huazhong University of Science and Technology (HUST), China. Her current research interests include power-aware architecture, non-volatile storage system and RDMA performance. Her work has been published in ICD-CS.



Yuanyuan Sun received the B.E and Ph.D degrees in computer science and technology from Huazhong University of Science and Technology (HUST), China, respectively in 2014 and 2019 and joins Alibaba Damo Academy. Her current research interests include cloud storage systems, semantic hashing, metadata management, index structures and big data analytics. She has published several papers in major conferences and journals including USENIX ATC, ACM SoC-C, MSST, ICDCS, IPDPS, IWQoS and TPDS.

She serves as program committee member in MSST 2020.