

Smoother: A Smooth Renewable Power-Aware Middleware

Xinxin Liu, Yu Hua*, Xue Liu[†], Ling Yang, Yuanyuan Sun

Wuhan National Laboratory for Optoelectronics

School of Computer, Huazhong University of Science and Technology, China

[†] McGill University, Canada

*Corresponding author: Yu Hua (csyhua@hust.edu.cn)

Abstract—The large electricity bills and the negative impacts on environments accelerate the use of renewable power to supply systems. However, frequent fluctuation and intermittency of renewable power often cause the challenges in terms of the stability of both electricity grid and systems, as well as decrease 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 middleware, called Smoother, which consists of Flexible Smoothing (FS) and Active Delay (AD). First, in order to smooth the fluctuation of renewable power, Flexible Smoothing carries out the optimized charge/discharge operation via computing the minimum variance of the renewable power that is supplied to systems per interval. Second, Active Delay 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 systems demonstrate that Smoother significantly reduces the negative impact of renewable power fluctuations on systems and improves the utilization of renewable power by 169.85% on average. We have released the source codes for public use.

Keywords-renewable power; middleware; energy management;

I. INTRODUCTION

Today, with the rapid growth of computation capability and the scale of systems, the large electricity bills have become almost unaffordable [1]. At the same time, 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]–[9]. In the meantime, the economical and environment-friendly renewable energy has been used to meet the needs of systems and receives many attentions [10]–[14].

However, renewable power has the essential properties of fluctuation and intermittency [2]. Using renewable power to supply systems needs to address two main challenges:

Challenge (1): High Fluctuation. Frequent fluctuation of renewable power poses the risks to systems and often leads to the instability of both grid [15] and clusters [16]. 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), which is unsafe and unreliable for systems. Furthermore, high fluctuation and intermittency of renewable power also increase the management overheads for renewable energy, such as the overhead of frequent load migration between grid-powered clusters and renewable-energy-powered clusters [17] and the requirement for a large battery capacity [18].

Challenge (2): Low Power Utilization. The renewable power generation and the workload power demand in systems are different and possibly lead to the imbalance between supply and demand, based on the analysis of real-world traces as shown in Fig. 6. This imbalance fails to make full use of renewable power in systems and reduces the utilization of renewable power [19].

Existing solutions fail to address the above challenges. Multigreen [13] proposes a cost-minimizing 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 systems. Moreover, the cost minimizing online algorithms [18] store 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 [17] 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 [15]. 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 [4], [19], [20] 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 systems.

Our Solution. Unlike existing schemes, we propose an efficient smooth renewable power-aware middleware, called Smoother, which consists of 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) with finite battery capacity and charging/discharging rate. We formulate the problem of Flexible Smoothing into a constrained nonlinear programming problem which determines the optimized charge/discharge scheme for renewable energy. Moreover, Active Delay improves the utilization of the smoothed renewable power. Active Delay 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 efficient and easy-to-use smooth renewable power-aware middleware, called Smoother, which contains two main components, i.e., Flexible Smoothing (*FS*) and Active Delay (*AD*). In order to smooth the fluctuation of renewable power and provides a relatively stable supply to systems (challenge 1), Flexible Smoothing (*FS*) carries out the optimized charge/discharge operation via computing the minimum variance of the renewable power that is supplied to systems per interval. In order to improve the utilization of the smoothed renewable power (challenge 2), Active Delay (*AD*) matches deferrable workloads of systems 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 systems using cost-efficient ESD management. We explore and exploit finite battery capacity, as well as the limited charging/discharging rate, which meet the demand of real-world situations.

3) **Suitable for A Variety of Renewable Energy.** Smoother can be used for a variety of renewable power sources, while executing similar operations. In this paper, without loss of generality, we only discuss the case study of wind power due to space limitation.

4) **Real-World Traces and Open-Source Code.** We have simulated Smoother by using MATLAB and released the source codes for public use in GitHub. 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 systems, and improves the utilization of

renewable power by 169.85% on average.

The rest of this paper is organized as follows. In Section II, we present the background of cluster energy management. In Section III, we present Smoother design. The experimental results are shown in Section IV. Section V shows the related work. Finally, we conclude this paper in Section VI.

II. BACKGROUND

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

A. Power Infrastructure and Consumption in Systems

As energy consumers, systems require power infrastructure to generate and distribute power. In general, there is a micro power grid that integrates all power supplies, such as the electric grid, diesel generator and renewable energy generators. In order to guarantee the availability, systems 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.

An accurate and useful linear power model is widely used to estimate servers' power consumption [2]:

$$P_{server}(\mu) = p^{idle} + (p^{full} - p^{idle}) * \mu \quad (1)$$

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

B. Wind Energy Source

Wind resource characteristics and turbine properties become the main factors to evaluate the quantity of power generated from wind energy sources [21]. There are two features about wind power [22]. 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 cut-in 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:

$$P_{wind}(\nu) = \begin{cases} 0 & \nu \leq V^{in} \\ G(\nu) & V^{in} < \nu \leq V^{rate} \\ P^{rate} & V^{rate} < \nu \leq V^{out} \\ 0 & \nu > V^{out} \end{cases} \quad (2)$$

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




Figure 1. Differentiated regions in wind power trace.

power, and $G(\nu)$ is defined as the power curve between the cut-in and rated speeds.

The optimal wind power generation scheme [21] tries to leverage multiple curve fitting methods for $G(\nu)$. Compared with polynomial regression and exponential fitting, Gaussian Regression performs the simulation of the relations between power coefficient and wind speed, i.e.,

$$G(\nu) = a_1 e^{-\frac{(\nu-b_1)^2}{c_1^2}} + \dots + a_n e^{-\frac{(\nu-b_n)^2}{c_n^2}} \quad 1 \leq n \leq 5 \quad (3)$$

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

III. THE DESIGN OF SMOOTHER

A. Design Goals

In order to address the problems of high fluctuation and low power utilization of renewable power, this 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 systems as well as reducing the overhead of energy switching in systems (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 systems.

The idea behind Smoother middleware 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 middleware consists of two main components, including Flexible Smoothing (*FS*) and Active Delay (*AD*). Flexible Smoothing (*FS*) in Smoother provides a relatively stable renewable power supply by using energy storage devices, thus offering the smoothing effect. Active Delay (*AD*) in Smoother matches the deferrable workloads with renewable power generation and thus improve the utilization of renewable power.

B. Three Regions in Power Generation

Fig. 1 shows a real-world example of wind power trace, which comes from National Renewable Energy Laboratory

(NREL) [23]. Based on the degree of wind power fluctuation, we divide wind power trace into three regions.

In Region-I, compared with other regions, the wind power supply is relatively stable and unnecessary to be processed by Flexible Smoothing in Smoother. Specifically, Region-I consists of two situations of wind power generation according to Equation (2): 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-II, the frequent fluctuation of wind power often results in the instability of power supplies in systems, as discussed in Section I. We hence consider how to decrease the degree of wind power fluctuation during this region. 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. It is hence feasible to use the 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 III-C.

In Region-II-2, renewable power fluctuates too frequently. In order to alleviate the fluctuations, Region-II-2 needs to carry out much higher battery charging/discharging rate and battery capacity than other regions. Considering the trade-off 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 practical applications, we set the proportion of Region-II-2 to 0.05%–5% of the total regions of wind power trace according to the actual situation. Active Delay in Smoother in Section III-D is used to address the mismatch between the smoothed renewable power and workload demand, thus increasing the utilization of renewable power.

As shown in Section II-B, the power generated by the wind generator varies with wind speeds. Equations (2) and (3) calculate the electricity converted from wind resource, which is tightly correlated with the wind speed and the properties of wind turbines. Optimal Harvesting Wind Power [21] 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_{wind}^2 during an interval $[0, T]$ to represent the fluctuation of wind power supply. We calculate




Figure 2. The cumulative distribution function of capacity factor variance in May 2011, California [23].

the capacity factor variance σ_{wind}^2 :

$$\sigma_{wind}^2 = \frac{1}{T} \sum_{t=0}^T \left(\frac{P_{wind}(\nu_t)}{P_{rate}} - \mu_{cf} \right)^2 \quad (4)$$

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

$$\mu_{cf} = \frac{1}{T} \sum_{t=0}^T \frac{P_{wind}(\nu_t)}{P_{rate}} \quad (5)$$

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 wind energy and the required maximum rate of charging/discharging battery. As shown in Fig. 5, 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 to distinguish Region-II-1 and Region-II-2. Since there is “switching times” in Fig. 5, 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, like iSwitch [17]. 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 wind power to represent the impact of wind power fluctuation on systems. 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. However, frequent charging and discharging operations exacerbate battery lifetime and increase energy loss [24]. Therefore, in a real-world system, we need to trade off the battery consumption and energy switching overhead.

C. Flexible Smoothing in Smoother

The main idea of Flexible Smoothing is to alleviate the frequent fluctuation of renewable power in Region-II-1 and achieve a relatively stable supply of renewable energy to systems. Fig. 3 shows the ideal results of Flexible Smoothing. Raw power represents the original renewable power. Ideal power represents the renewable power with Flexible Smoothing in Smoother that is ultimately supplied




Figure 3. The case with Flexible Smoothing in Smoother. Raw power represents the original renewable power. Ideal power represents the renewable power with Flexible Smoothing in Smoother that is ultimately supplied to systems.

to systems. In Fig. 3, Flexible Smoothing in Smoother transfers the renewable power in the orange area to the light blue area and achieves a stable supply (the dark blue line).

Renewable energy generation can be estimated by Equations (2) and (3) given meteorological data [21], [25]. On the other hand, the future generation of renewable energy can be accurately predicted by machine learning methods [26]–[29]. For example, LSSVMCGSA model achieves about 5%–10% prediction error within 48 hours [28]. Note that the prediction mechanism is out of the scope of this paper, and we leverage a widely-used method to predict renewable power based on Equations (2) and (3) in Section II-B [21].

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 charging/discharging operation, the system can obtain a relatively smooth and stable renewable energy supply.

Suppose that there are m time points (e.g., $m = 12$) in the interval $[1, m]$. $U = [u_1 u_2 \dots u_m]^T$ represents the amount of the generated renewable energy at each time point (e.g., five minutes). $S = [s_1 s_2 \dots s_m]^T$ 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 system after the battery charging/discharging operation at each time point in the interval $[1, m]$:

$$A = U + S = \begin{bmatrix} a_1 \\ a_2 \\ \dots \\ a_m \end{bmatrix} = \begin{bmatrix} u_1 + s_1 \\ u_2 + s_2 \\ \dots \\ u_m + s_m \end{bmatrix} \quad (6)$$

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

lem:

$$\min \sigma_A = \sqrt{\frac{1}{m} \sum_{i=1}^m (a_i - \bar{a})^2} \quad (7)$$

$$\text{subject to:} \quad \forall i \in [1, m], \begin{cases} 0 \leq |s_i| \leq 0.9M, & s_i \geq 0 \\ 0 \leq |s_i| \leq u_i, & s_i < 0 \end{cases} \quad (8)$$

$$\forall i \in [1, m], 0.1M \leq \left| \sum_{t=1}^i s_t \right| \leq M \quad (9)$$

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 based on the interval of wind power information [23], and m is equal to 12 since the decision is computed each hour. Equation (8) 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 90% of the battery capacity to avoid the damage of full discharge on the battery [30]. Equation (9) indicates that the accumulated electric quantity of the battery can not exceed the battery capacity or be less than 10% of the battery capacity during the interval $[1, m]$.

It is worth noting that the rate limits of charging/discharging battery are implicitly considered in our model proposed above. In our implementation in Section IV, the battery capacity is set to sustain one time point (five minutes) of operations at the maximum rate of charging/discharging battery, while the decision is computed each interval (one hour). Actually, the larger battery capacity (e.g., which can sustain thirty minutes of operations at the maximum rate of charging/discharging battery) will yield the better smoothing effect.

Based on the results of Equations (6) – (9), we execute the battery charging/discharging operations on the UPS battery. Fig. 4 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.

We carry out Flexible Smoothing in Smoother at Region-II-1, as shown in the red dotted circle in Fig. 4. As described in Section III-B, considering the trade-off between the smoothing effect and the battery overhead, we do not perform Flexible Smoothing in Region-II-2.

In order to illustrate the effects of different thresholds for distinguishing Region-II-1 and Region-II-2, we examine a set of data as shown in Fig. 5. *W/ Smooth* and *W/O Smooth* respectively represent the energy switching times between the grid and wind power in the power supply of server clusters with and without performing Flexible




Figure 4. The smoothed wind power versus the original wind power. The region in red dotted circle in Region-II-1 becomes stable after using Flexible Smoothing in Smoother.

Smoothing in Smoother. Battery MaxVol represents the required maximum battery charging/discharging rate for Flexible Smoothing in Smoother. As mentioned above, the battery capacity is set to sustain one time point (five minutes) of operations at the maximum battery charging/discharging rate in our implementation. Therefore, Battery MaxVol also represents the trend of the required battery capacity. The CDF values of the capacity factor variance (on the x-axis) are used to distinguish Region-II-1 and Region-II-2. For example, when the CDF value is 0.95, which means region-II-2 accounts for 5% of the total regions, the corresponding capacity factor variance is $2 * 10^{11}$ (as shown in Fig. 2). We regard this variance value as the upper boundary of Region-II-1 (also the lower bound of Region-II-2). The increase of CDF value in Fig. 5 means that more regions are classified as Region-II-1. Fig. 5 shows the comparisons between the initial energy switching times and the energy switching times after executing Flexible Smoothing in Smoother as well as the required maximum rate of charging/discharging battery based on different CDF values. As the CDF value increases, the energy switching times becomes smaller while the corresponding required maximum rate of charging/discharging battery (and the resulting battery capacity) becomes larger. We hence need to obtain a suitable trade-off between the energy switching times and the battery capacity. In this paper, we set the ratio of Region-II-2 to be 5%.

Finally, 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 Fig. 4.

D. 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 systems. 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 systems. As shown in Fig. 6, the




Figure 5. 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. The energy switching times between the grid and wind power in the power supply of server clusters are on the left y-axis. The required maximum battery charging/discharging rate for Flexible Smoothing is on the right y-axis, which also represents the trend of the required battery capacity. The CDF values of the capacity factor variance for distinguishing Region-II-1 and Region-II-2 are on the x-axis.

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 systems.

Real-time workloads, such as Web access [31], 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 Fig. 7, there are three deferrable jobs, i.e., $J1$, $J2$ and $J3$. The dark grey regions with oblique lines under black solid line represent the power demand of deferrable jobs. The transparent regions under black dotted line represent the power of renewable energy. The light grey regions with oblique lines represent the overlapping regions of the renewable power and demands. Fig. 7 (a) shows the original case without Active Delay in Smoother, where a system can not use renewable energy to supply $J2$, leading to a low utilization of renewable power. Fig. 7 (b) shows the use of renewable energy for each job via Active Delay in Smoother, where the renewable power utilization increases through delaying the execution of $J2$ to the time with the most renewable energy use before the soft-deadline (red dotted line in Fig. 7 (b)).

In many workload schedulers, the deadline and the expected running time for each job can be provided by users or be estimated using historical statistics [19]. Therefore, we assume we get the above information from users or historical statistics.

We present the details of Active Delay in Smoother in Algorithm 1. Specifically, 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




Figure 6. The imbalance between workload demand and renewable power supply. The renewable energy in the green area can not be used by systems.




Figure 7. Active Delay diagram.

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 II-A. 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 compute the amount of renewable power that the current job will use to facilitate the execution within the slack-time. Then we choose the time with the highest renewable energy use as the real execution time for the current job (lines 13–17). 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 18 updates the remaining renewable power after the scheduling (i.e., maintaining resources for the current job). But if the slack-time of the current job is not larger than 0, this job needs to be carried out immediately (lines 19–21).

IV. PERFORMANCE EVALUATION

Smoother middleware 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 systems (e.g., the overhead of virtual machine migration between grid-powered clusters

Algorithm 1 Active Delay in Smoother.

Input: System 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;
2: priority_queue<Job> queueJob;
3: for each small time slot do
4:   if requestJob.size() > 0 then
5:     for job: requestJob do
6:       calWorkloadPower(job);
7:       slackTime = job.deadline - job.runTime -
         currentTime;
8:       queueJob.push(job, slackTime);
9:       requestJob.pop();
10:    end for
11:    for job: queueJob do
12:      if job.slackTime > 0 then
13:        for time: slackTime do
14:          power = getRenewablePowerUsedWhen-
            StartAtTime(job, time);
15:        end for
16:        startTime = getStartTimeOfJobFor-
            MaximizeRenewablePower(job);
17:        executeJob(job, startTime);
18:        updateRemainRPower(job, startTime);
19:      else
20:        executeJob(job, currentTime);
21:        updateRemainRPower(job, currentTime);
22:      end if
23:      queueJob.pop();
24:    end for
25:  end if
26: end for
```

and renewable-energy-powered 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 systems, and renewable energy utilization.

A. Experimental Setup

System Configurations: We assume that a system is equipped with 11,000 servers and each server has the same processing power [2], i.e., identical energy consumption for executing the same jobs. The peak power of each server is 186W and the idle power is 62W [17]. Each server runs the same workloads and has similar CPU utilization, like existing experimental configurations [17].




Figure 8. Power consumption of Google cluster.

Workload Traces: Our evaluation uses three typical real-world workloads, including Google cluster-data [32], Web workloads [31] and batch workloads [33]. Specifically, Google trace provides real-world power consumption from a 12,500-machine cell over about a month-long period in May 2011 [32]. According to Equation (1) in Section II, we convert the CPU utilization into power trace, as shown in Fig. 8. The time-sensitive Web workloads are collected from the real-world logs provided by the Internet Traffic Archive [31]. We generate the number of requests per minute from the request logs and convert this number into CPU utilization via a linear analog [17], 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 I, 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 [33] 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 II.

Table I
FIVE WEB WORKLOAD TRACES WITH DIFFERENT AVERAGE CPU UTILIZATIONS [31].

Web	Description	Avg. CPU utilization
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 II
BATCH WORKLOAD TRACES WITH DIFFERENT CPU UTILIZATIONS [33].

Batch Workload Traces	Avg. CPU utilization
LLNL Thunder	86.7%
LANL CM5	74.4%
HPC2N	60.1%
Sandia Ross	49.9%

Renewable Power Traces: Wind power traces come from Wind Data Resources [34] of the National Renewable Energy Laboratory (NREL). We select two groups of traces which are differentiated by volatility intensity, including low and high volatility traces. The low volatility traces have relatively stable and smooth generation, while the high volatility traces have output rate with high volatility. Table III shows the two groups of traces and each group has three traces with different capacity factors.

Table III
WIND POWER TRACES WITH DIFFERENT VOLATILITY INTENSITIES [34].


Wind Power Traces	Site ID	Capacity Factor
Low volatility	CA(9122)	17.9%
	OR((24258)	19.0%
	WA(29359)	17.9%
High volatility	TX(10)	32.4%
	CO(11005)	29.9%
	WY(16419)	29.6%

B. 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 system. Smoother is compared with the efficient battery storage solution where the system first uses renewable power as 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 [13]. We use *Comp* to represent the efficient battery storage solution. *W/* means “with”, and *W/O* means “without”.

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 III-C, we set $2 * 10^{11}$ whose corresponding CDF value is 0.95 as the threshold of capacity factor variance to distinguish Region-II-1 and Region-II-2. We execute Flexible Smoothing 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 wind power with different fluctuation degrees on May 2, 14, 18 and 23, 2011, California, USA [23], and compare the energy switching times after carrying out Flexible Smoothing (*FS*) with the initial energy switching times within four days. From Fig. 9(a), we observe that wind power fluctuates most frequently on May 18, and becomes smooth on May 2. Therefore, by using Flexible Smoothing, the energy switching times significantly decrease on May 18, but the decrease degree becomes smaller on May 2, as shown in Fig. 9(b). The main reason is that Flexible



(a) Wind power traces on May 2, 14, 18 and 23, 2011, California, USA [23].



(b) Energy switching times on May 2, 14, 18 and 23.

Figure 9. The comparison of energy switching times with different wind fluctuation degrees between “W/O *FS*” scheme and “W/ *FS*” scheme.

Smoothing in Smoother aims to alleviate the fluctuation degree of renewable power. Hence if the wind power is stable, Flexible Smoothing plays the less important role.

In the similar way, we use five Web workload traces in Table I and two groups of wind power traces in Table III respectively for comparisons in terms of energy switching times. When the total installed wind turbine capacity is set to 976KW, the results are shown in Fig. 10 and Fig. 11. When the total installed wind turbine capacity is set to 1525KW, the results are shown in Fig. 12 and Fig. 13.

We observe that with different workload 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 systems. Fig. 11 and Fig. 13 also demonstrate that Flexible Smoothing in Smoother brings more remarkable effect when processing renewable traces with high volatility.

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.

Fig. 14 and Fig. 15 show the use of renewable energy when being equipped with sufficient renewable power and insufficient renewable power, respectively. The blue area represents the amount of renewable power supply, the green line indicates initial power consumption of workload, and the red line indicates the adjusted power consumption of workload when Active Delay in Smoother is used. We observe whether the renewable power is sufficient or not,




Figure 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).




Figure 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).




Figure 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).




Figure 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).

Active Delay in Smoother makes full use of renewable power.

Table II shows batch workload traces with different CPU utilizations. As shown in Fig. 16, with different workloads under different renewable energy supplies, the utilization of renewable power increases by an average of 169.85% by using Active Delay in Smoother. All the workloads increase the utilization of renewable power after performing Active Delay in Smoother. Especially in the case of HPC2N with a low supply of renewable energy, the utilization increases significantly from 0.19 to 0.81. When the renewable power is sufficient, the supply of renewable power is larger than the total energy required for the workload. Therefore, the utilization of renewable power is low relatively. However, it does not mean Active Delay in Smoother is more suitable for low supply of renewable power. In fact, Active Delay in Smoother can be used for each case. We argue that Active Delay in Smoother can significantly improve the utilization of renewable energy, and reduce the use of brown energy, thus decreasing the electricity costs of systems 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, the “W/ FS and W/ AD” scheme significantly reduces the energy switching times by more than 25%, thus mitigating the negative impact of renewable energy fluctuations on systems.

In summary, by using Flexible Smoothing and Active

Delay, our proposed Smoother significantly reduces the negative impact of renewable power fluctuation on systems and increases the utilization of renewable power, thus improving the stability of the grid and systems and reducing the cost of systems.

V. RELATED WORK

In order to meet the energy needs of systems, 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 [13], [18], [35], mitigate peak power [36], minimize the amount of user information leaked to the grid [37], [38] and eliminate the fluctuation of renewable energy [18]. The optimal control of end-user energy storage [35] stores energy at much lower prices in a battery that discharges when the energy prices are high to satisfy the demand for cost savings. Multigreen [13] 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. The cost minimizing online algorithms [18] store all 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 [36] proposes a peak reduction algorithm leveraging




Figure 14. The case with sufficient renewable power.




Figure 15. The case with insufficient renewable power.

energy storage. There also exist schemes to minimize the fluctuation for the demand with energy storage for privacy issues [37], [38]. Unlike them, the design goal of Smoother is to alleviate the fluctuation in renewable energy supply.

Renewable Energy Generation: Optimal Harvesting Wind Power [21] shows the computation of estimating the wind energy generation with meteorological data. Study of Forecasting Renewable Energies [39] proposes a methodology based on linear predictive coding and digital image processing principles to estimate both wind speed and solar radiation. Wind Profile Prediction [40] leverages the Mc-FCRBF network for predicting the wind speed and wind direction. On the other hand, the future generation of renewable energy can be accurately predicted by machine learning methods [26]–[29]. For example, LSSVMCGSA model achieves about 5%–10% prediction error within 48 hours [28]. 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 to predict renewable power based on Equations (2) and (3) in Section II-B [21].

Workload Scheduling: In order to address the problem of the mismatch between system workloads and green energy supplies, existing workload scheduling schemes have been proposed [4], [19], [20]. Their basic idea is to postpone the deferrable workloads until the renewable power is sufficient or the electricity price is low before the soft-deadline of workloads. In addition, some schemes leverage geographical load balancing among distributed systems to




Figure 16. 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.




Figure 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.

improve the utilization of renewable power [14]. However, these algorithms fail to consider the impact of renewable power fluctuation on the stability of the grid and systems.

VI. CONCLUSION

It is an important problem of providing a smooth and stable supply of renewable power and improving the utilization of renewable power in systems. In order to address the two problems, we propose a smooth renewable power-aware middleware, called Smoother, which consists of Flexible Smoothing (FS) and Active Delay (AD). The novelty behind Smoother middleware is that we emphasize the impact of frequent fluctuation of renewable power on the stability of the grid and systems, as well as improve the utilization of renewable power. The trace-driven evaluation demonstrates that our proposed Smoother middleware offers a smooth and stable supply of renewable power for systems, and improves the utilization of renewable power by an average of 169.85%. Smoother is able to improve system performance and the stability of the grid and systems, while meantime reducing the costs of systems. We have released the source code of Smoother for public use at <https://github.com/csXinxinLiu/Smoother>.

ACKNOWLEDGMENT

This work was supported by National Natural Science Foundation of China (NSFC) under Grant 61772212.

REFERENCES

- [1] A. Qureshi, "Power-demand routing in massive geodistributed 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] J. Wu, B. Cheng, and M. Wang, "Energy minimization for quality-constrained video with multipath tcp over heterogeneous wireless networks," in *ICDCS*, 2016.
- [4] 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.
- [5] D. Cheng, P. Lama, C. Jiang, and X. Zhou, "Towards energy efficiency in heterogeneous hadoop clusters by adaptive task assignment," in *ICDCS*, 2015.
- [6] R. Zhou, Z. Li, and C. Wu, "An online procurement auction for power demand response in storage-assisted smart grids," in *INFOCOM*, 2015.
- [7] S. Nikolettseas, T. P. Raptis, and C. Raptopoulos, "Low radiation efficient wireless energy transfer in wireless distributed systems," in *ICDCS*, 2015.
- [8] R. Zhou, Z. Li, C. Wu, and M. Chen, "Demand response in smart grids: A randomized auction approach," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 12, pp. 2540–2553, 2015.
- [9] X. Mei, X. Chu, H. Liu, Y.-W. Leung, and Z. Li, "Energy efficient real-time task scheduling on cpu-gpu hybrid clusters," in *INFOCOM*, 2017.
- [10] Y. Zou, X. Lin, D. Aliprantis, and M. Chen, "Robust multi-stage power grid operations with energy storage," in *INFOCOM*, 2018.
- [11] G. T. R. Das *et al.*, "Fuzzy set theory applications for facts devices in grid connected renewable power systems," in *ICDCS*, 2016.
- [12] Z. Zhou, F. Liu, and Z. Li, "Bilateral electricity trade between smart grids and green datacenters: Pricing models and performance evaluation," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 12, pp. 3993–4007, 2016.
- [13] W. Deng, F. Liu, H. Jin, C. Wu, and X. Liu, "Multigreen: Cost-minimizing multi-source datacenter power supply with online control," in *e-Energy*, 2013.
- [14] Y. Zhang, Y. Wang, and X. Wang, "Greenware: Greening cloud-scale data centers to maximize the use of renewable energy," in *Distributed Systems Platforms and Open Distributed Processing*, 2011.
- [15] J. P. Lopes, N. Hatziaargyriou, 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.
- [16] 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.
- [17] C. Li, A. Qouneh, and T. Li, "iswitch: coordinating and optimizing renewable energy powered server clusters," in *ISCA*, 2012.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] "Enercon," <http://www.enercon.de/home>, 2015.
- [23] "National Renewable Energy Laboratory," <http://www.nrel.gov/irred/>.
- [24] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?" in *SIGMETRICS*, 2012.
- [25] "PV solar energy calculation," <http://photovoltaic-software.com/PV-solar-energy-calculation.php>.
- [26] 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.
- [27] N. Sharma, P. Sharma, D. Irwin, and P. Shenoy, "Predicting solar generation from weather forecasts using machine learning," in *SmartGridComm*, 2011.
- [28] 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.
- [29] 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.
- [30] "Northern Arizona Wind & Sun. Deep cycle battery," <https://www.solar-electric.com/deep-cycle-battery-faq.html>.
- [31] "The Internet Traffic Archive," <http://ita.ee.lbl.gov/>.
- [32] "Google cluster-data," <https://github.com/google/cluster-data>.

- [33] “Logs of Real Parallel Workloads from Production Systems.” <http://www.cs.huji.ac.il/labs/parallel/workload/logs.html>.
- [34] “Wind Data Resources of the National Renewable Energy Laboratory.” https://www.nrel.gov/wind/data_resources.html.
- [35] 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.
- [36] S. Govindan, A. Sivasubramaniam, and B. Uргаonkar, “Benefits and limitations of tapping into stored energy for datacenters,” in *ISCA*, 2011.
- [37] F. Laforet, E. Buchmann, and K. Böhm, “Towards provable privacy guarantees using rechargeable energy-storage devices.” in *e-Energy*, 2016.
- [38] S. Li, A. Khisti, and A. Mahajan, “Privacy-optimal strategies for smart metering systems with a rechargeable battery,” in *ACC*, 2016.
- [39] A. A. Moghaddam and A. Seifi, “Study of forecasting renewable energies in smart grids using linear predictive filters and neural networks,” *IET Renewable Power Generation*, vol. 5, no. 6, pp. 470–480, 2011.
- [40] E. Sathish, M. Sivachitra, R. Savitha, and S. Vijayachitra, “Wind profile prediction using a meta-cognitive fully complex-valued neural network,” in *ICoAC*, 2012.