

# Exploiting Renewable Energy and UPS Systems to Reduce Power Consumption in Data Centers

Xiaopu Peng (✉ [xzp0007@auburn.edu](mailto:xzp0007@auburn.edu))

Auburn University <https://orcid.org/0000-0003-4995-7746>

Tathagata Bhattacharya

Auburn University

Ting Cao

Auburn University

Jianzhou Mao

Auburn University

Taha Tekreeti

Auburn University

Xiao Qin

Auburn University

---

## Research

**Keywords:** Renewable energy, uninterruptible power supply (UPS), distributed UPS systems, resource management, energy cost, data centers.

**Posted Date:** July 9th, 2021

**DOI:** <https://doi.org/10.21203/rs.3.rs-668107/v1>

**License:**  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

**Version of Record:** A version of this preprint was published at Big Data Research on December 1st, 2021. See the published version at <https://doi.org/10.1016/j.bdr.2021.100306>.

# Exploiting Renewable Energy and UPS Systems to Reduce Power Consumption in Data Centers

Xiaopu Peng, (Corresponding Author xzp0007@auburn.edu)

Tathagata Bhattacharya, Ting Cao, Jianzhou Mao, Taha Tekreeti, and Xiao Qin, *Senior Member IEEE*  
*Department of Computer Science and Software Engineering Auburn University, Auburn, AL.*



**Abstract**—To develop environmental friendly and energy-efficient data centers, it is prudent to leverage on-site renewable sources like solar and wind. Data centers deploy distributed UPS systems to improve efficiency, scalability, and reliability of UPS systems, thereby handling the intermittent nature of renewable energy. We propose a renewable-energy manager called *REDUX* to (1) offer a smart way of managing energy supply of data centers powered by grid and renewable energy and (2) maintain a desirable balance between energy cost and system performance. To achieve this overarching objective, *REDUX* judiciously orchestrates distribute UPS devices (i.e., recharge or discharge) to allocate energy resources when (1) grid price is at low or high states or (2) renewable energy generation is at a low or fluctuate level. *REDUX* not only guarantees the stable operation of daily workload conditions, but also cuts back the energy cost of data centers by improving power resource utilization. Compared with the existing strategies, *REDUX* demonstrates a prominent capability of mitigating average peak workload and boosting renewable-energy utilization.

**Index Terms**—Renewable energy, uninterruptible power supply (UPS), distributed UPS systems, resource management, energy cost, data centers.

## 1 INTRODUCTION

Existing power management techniques for data centers either overlook the usage of uninterruptible power supply devices (i.e., UPS devices) as an energy storage or source, or pay no attention to the nature price fluctuation of grid power. In this study, we develop an energy resource manager called *REDUX* to cost-effectively allocate energy resources by incorporating a distributed UPS system tailored for renewable energy like solar and wind.

*REDUX* aims to minimize energy cost while maintain required data center performance by deploying distributed UPS devices in data centers to address two challenging issues, namely, 1) the intermittency feature of renewable energy and 2) dynamically changing grid electricity price. *REDUX* is responsible for maintaining stable operation of daily computing workload in data centers while cutting back energy costs by improving overall energy resource utilization. Four emerging trends below strongly motivate us to contrive the *REDUX* system.

- Energy cost of large-scale data centers is skyrocketing.
  - Consuming renewable energy in data centers brings economical and environmental benefits.
- 
- *X. Qin is with Department of Computer Science and Software Engineering, Shelby Center for Engineering Technology, Samuel Ginn College of Engineering, Auburn University, AL 36849-5347. E-mail: xqin@auburn.edu.*

- There are pressing demands for data centers powered by cost-effective energy.
- Distributed uninterruptible power supplies systems are increasingly popular in data centers.

In the past decades, there has been a rapid growth of data centers built for a wide range of cloud-based computing services [1]. However, massive amounts of power consumption of large-scale data centers has become a serious challenge to data center designers and operators worldwide [2]. It is evident that data centers running around the world are at a risk of doubling their energy consumption every five years [3]. To make it worse, a significant portion of electricity is generated through carbon-intensive approaches (a.k.a., brown energy), and the environmental impacts caused by these data centers are under a continuous pressure from the media and society.

In contrast to brown energy, renewable energy harvested from sources such as wind turbines and solar panels exhibits an assortment of benefits [4]. Most renewable energy sources are clean and produce few global warming emissions while reduces energy costs by offering affordable electricity. As such, a growing number of data centers have kicked off various initiatives to integrate on-site renewable sources into power supplies. For instance, *Apple* has recently constructed a massive 100-acre solar farm adjacent to the *iCloud* data center in North Carolina. This solar farm yields 84 million kWh of clean renewable energy annually.

Designing renewable-energy powered data centers has promising benefits beyond low carbon footprint [5]. The first benefit is that compared with grid energy supply, renewable-energy price tends to remain reasonably low and stable after on-site installations are completed [6]. Running 24/7, data centers lead to grid energy bills with higher unit prices at on-peak period than off-peak period. Although non-urgent workloads may be deferred to off-peak time, it would be profitable and reliable for the data centers to construct its own on-site renewable facility in the long run. Renewable energy supplies, being highly modular, can quickly scale up in power capacity to accommodate load growth. Another benefit is that on-site renewable energy facilities can be constructed in a shorter time interval compared with conventional power plants [7].

To prevent costly down time incurred by power budget violations, Govindan *et al.* employed a centralized UPS scheme to furnish energy during a peak demand period [13]. The UPS-based technique effectively hides the extra power from power grids by completely shaving high-magnitude power spikes in a short time window (e.g., 1-2 hours). When it comes to long spike intervals (e.g., 8-10 hour), distributed and per-

Table I. Comparisons between REDUX and the existing power management techniques for data centers.

Power Management	UPS Enabled	Renewable Energy	Management Regions	Grid Price Fluctuations
ReUPS [8]	✓	✓	✓	×
Greenworks [9]	✓	✓	×	×
iSwitch [7]	×	✓	✓	×
EcoPower [10]	×	✓	×	✓
GreenSlot [11]	×	✓	×	✓
GreenSwitch [12]	×	✓	×	✓
<b>REDUX (This Study)</b>	✓	✓	✓	✓

server UPS devices represent an economical solution by the virtue of battery backup. In contrast to the conventional UPS design, distributed UPS batteries allow data-center operators to readily shave peak power with stored energy [14]. The advantages of distributed UPS systems include good efficiency, high scalability, and adequate reliability. A recent study showed that a hybrid distributed UPS at the power distribution unit (a.k.a., PDU) and server levels is a precursor architecture. As an example, a novel distributed battery control system was deployed in a Google state-of-the-art data center [15].

In this study, we propose a workload-energy management system - REDUX - to allocate renewable energy integrated with a power grid. REDUX is adept at smoothing power supplies with the three supply levels of renewable energy resources, namely, high, fluctuate, and low. By naturally scaling with the data center size and eliminating a potential single point of failure, per-server distributed batteries become an economical and secure solution for battery backup. REDUX also makes judicious decisions on UPS charging and discharging with vital information on renewable-energy supply levels and time-dependent grid power price. For instance, REDUX makes full use of UPS devices to render energy resources when renewable energy supply is inadequate or under a fluctuate condition.

One prominent feature of REDUX is that it orchestrates a desirable balance between renewable-energy utilization and data center performance. Importantly, REDUX conserves the energy cost of data centers through power resource management while offering reliable computing operations. Compared with existing strategies, REDUX demonstrates an exemplary capability of mitigating average peak workload and boosting renewable-energy utilization.

This paper makes three main contributions:

- *Design.* We design REDUX, a workload-energy management system governing UPS units and renewable energy for computing servers according to workload requirements in data centers. Taking advantage of a hierarchical power management framework design, REDUX enables a multi-source driven multi-objective power management coordination, thereby maximize the benefits of the hybrid energy system without heavily relying on grid, renewable or UPS power supply.
- *Characterization.* After characterizing the states of renewable energy and workload conditions, we craft the state-transition diagrams to illustrate transitions among the states. Our system governs energy draw either from cost-effective grid power or UPS devices, thereby saving energy cost in data centers while maintaining acceptable performance. Depending on the characterized energy and workload states, REDUX either takes full advantage of low-cost energy or defers the execution of non-urgent workload to significantly cut back the overall electricity

cost in a long run.

- *Optimization.* We evaluate REDUX using real-world workload traces and green energy data. We show that compared with the prior solutions, REDUX conserves the total energy cost anywhere between 10% to 25%. We demonstrate that REDUX paves the way for constructing modern data centers that are economically and environmentally friendly.

The rest of the paper is organized as follows. Section 2 surveys the related work of energy-efficient computing, renewable energy in data centers, and distributed UPS systems. The framework and models of REDUX are detailed in Section 3. Section 5 articulates an array of heuristic algorithms at the heart of REDUX. Section 6 describes the experimental settings and results. Finally, Section 7 concludes this paper.

## 2 BACKGROUND AND RELATED WORK

Let us classify the related studies into four core areas, namely, energy-efficient computing, conserving energy in data centers, renewable energy in data centers, and distributed UPS systems. Table I summarizes the major differences between our proposed REDUX and the existing power management techniques for data centers.

**Energy-Efficient Computing.** A large body of early studies focused on reducing power consumption of a single server by applying the dynamic voltage and frequency scaling technique (i.e., DVFS) [16], low-power chipsets [17], and advanced cooling techniques [18]. Emerging energy-management schemes aim to optimize energy efficiency of servers equipped with multi-core processors [19], GPUs [20], and smart memory cubes [21]. In contrast to the above energy-efficient computing strategies, our REDUX pays attention to reducing energy cost of large-scale data centers.

**Conserving Energy in Data Centers.** In recent years, there has been a pressing need to minimize the total power consumption of data centers [22] [10]. Chun *et al.* investigated a hybrid datacenter architecture mixing low power systems and high performance ones [23]. Dabbagh *et al.* developed an integrated energy-aware resource provisioning framework to optimize energy efficiency of data centers [24]. Shuja *et al.* investigated practical approaches to maximizing quality of services while minimizing energy consumption of data center resources.

Virtualization techniques have been applied to reduce the power consumption by consolidating multiple virtual machines into one physical server [25]. For example, Tseng *et al.* designed a virtual machine management strategy to predict resource utilization for upcoming videos, thereby turning off idle servers to conserve energy [26]. Another thread of techniques that minimize the electricity cost of data centers is exploiting the temporal and regional diversity of electricity prices (see, for

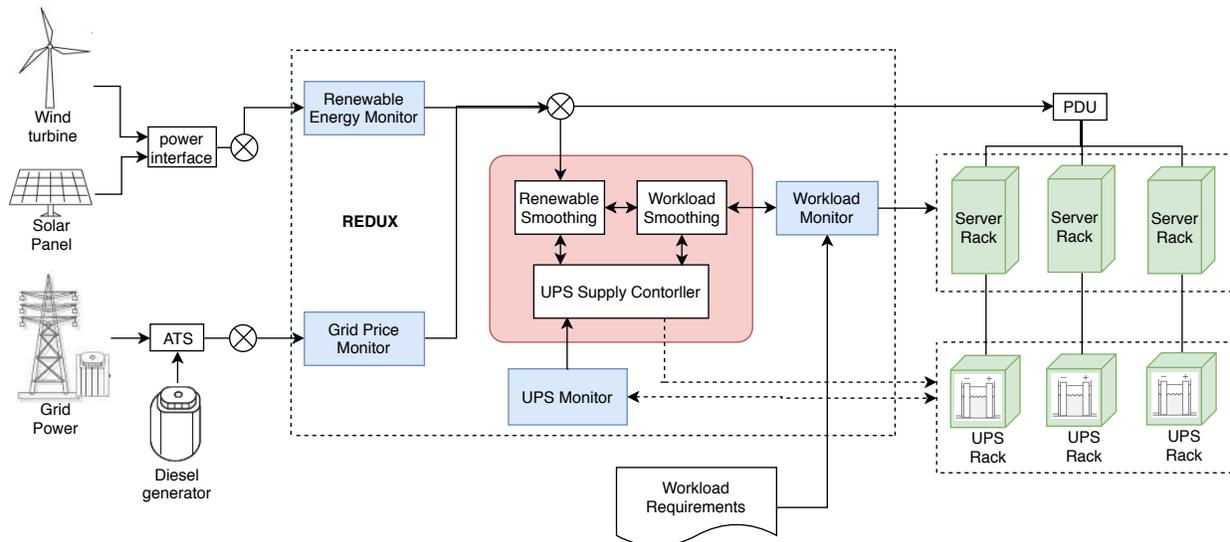


Figure 1. The framework of the REDUX system. REDUX orchestrates renewable energy, UPS units, and server racks according to workload requirements.

example, [27]). The pragmatic idea of these schemes is to dynamically allocate resources to service workload in data centers with a cheap electricity price. Unlike the existing energy-saving techniques, our proposed REDUX conserves energy by taking a full advantage of renewable energy and UPS devices.

**Renewable Energy in Data Centers.** There are pressing demands from cloud service providers for cost-effective renewable energy to power data centers [9]. An increasing number of prior studies focused on making green data centers through the integration of renewable power sources [28] [12]. A handful of resource management techniques were developed to address the shortfall intermittency of renewable power. For example, Goiri *et al.* designed a parallel batch job scheduler to match the computational workload with solar energy supplies in a datacenter [11]. To integrate renewable energy, Guo *et al.* incorporated geographical load balancing, opportunistic workload scheduling, and thermal storage management in data centers [29].

**Distributed UPS Systems.** Uninterruptable power supplies (i.e., UPS devices) offer efficient peak power shaving in data centers [15] [14] [8]. For example, Google data centers employ server-level UPS units, after which a battery is attached on each server [30]. Urgaonkar *et al.* proposed the usage of UPS units to offer ample opportunities of cost reduction in data centers [31], by first prove the feasibility of leveraging energy storage device (e.g., lead-acid batteries) in data centers to reduce peak power cost. Aksanli *et al.* applied distributed UPS devices to store energy during low load activity periods while using UPS stored energy during power spikes [15]. Kontorinis *et al.* explored the total cost of operation of the distributed UPS system in datacenters and proposed using local distributed UPS systems to shave the datacenter’s peak power [14]. Govindan *et al.* designed a UPS-based energy buffer to reduce energy costs in data centers [30]. To better utilize the temporal variation of electricity prices, Wang *et al.* deployed UPS systems to shift demand peak away from high-price periods in data centers [32]. Unlike the aforementioned techniques, REDUX stores energy in UPS units during intervals when renewable energy is stable and incessant or power grid prices are low. For the rest of this paper, we use terms UPS or battery interchangeably.

### 3 FRAMEWORK DESIGN

We start this section by depicting the REDUX system framework (see Section 3.1), which offers a high-level design from the perspective of power management in data centers. Next, we shed light on the responsibilities of the core modules in REDUX (see Section 3.3).

#### 3.1 The REDUX Platform

REDUX is a management system orchestrating UPS units and renewable energy for computing servers according to workload requirements in data centers. In what follows, we outline the correlations between REDUX and the external entities such as renewable and grid energy supply, followed by the descriptions of the primary modules in REDUX.

Fig. 1 plots the framework design of the REDUX system, where renewable energy (e.g., wind and solar), grid power, and diesel generators are seamlessly integrated (see the left-hand side of Fig. 1). REDUX employs a distributed UPS system, where a UPS device is attached to each server rack (see the right-hand side of Fig. 1).

REDUX persistently communicates with the five external entities, namely, renewable energy supplies, a power grid with diesel generator as backup, distributed UPS units, server racks and workload requirements. A wide range of energy supplies may include, but not limited to wind or solar energy, electrical grid and diesel generators. To alleviate the counter-varying load effect from the perspective of power grids, smart on-site diesel generators are installed in data centers as an energy buffering mechanism. Such an energy buffer aims at smoothing dynamically changing renewable energy to minimize the adverse effect of counter-varying load.

REDUX consists of two layers : an interface layer (see all light-blue colored modules in Fig. 1) and a backbone layer (see the modules inside the light-red colored box in Fig. 1), which are detailed in the following two subsections.

#### 3.2 The Interface of REDUX

The interface layer is comprised of four modules, including a renewable energy monitor, a grid price monitor, a UPS monitor, and a workload monitor. Seamlessly integrating these modules,

REDUX keep track of (1) the necessary energy supply trace data from on-site renewable device or off-site grid networks, (2) an amount of workload to initiated in the current time slot or deferred from the previous time slot, and (3) energy storage amount in each UPS rack. On the other hand, we also need these monitors to update the necessary states of the specific data. For example, the renewable energy monitor and the grid price monitor will determine the state of renewable energy supply state and grid price state, respectively. Then the backbone layer can continue on the energy management process; The workload monitor will also update the workload states before the workload smoothing method take charge of workload arrangement tasks.

Furthermore, we introduce a primitive prediction mechanism in all the modules except for the UPS monitor. By considering the fluctuating patterns of renewable energy supply or grid power price, and the status of previous workload conditions, this mechanism helps the backbone layer of REDUX to achieve rational and improved decisions on both energy supply and workload management. Although we make no assumption on statistical distribution of workloads, this prediction method on workloads can help the core module in backbone layer by smoothing the workloads that are dramatically changing between subsequent time slots. Please refer to Section 6.1 for details on this mechanism.

In a nutshell, these monitors collect trace data from all the external entities, which feed data into the modules residing in the backbone layer.

### 3.3 The Backbone of REDUX

The three core modules in the backbone layer communicate with five external entities (i.e., renewable energy supplies, power grid, UPS units, server racks, workload requirements) through the four interface modules in the interface layer.

The backbone layer of REDUX contains three modules, namely the Renewable Smoothing module, the Workload Smoothing module, and the UPS storage controller module. The algorithms of these three modules are presented in Section 5. The functionalities of the three core modules of REDUX are summarized as follows:

- **Renewable Smoothing.** The renewable smoothing module (see Algorithm 2 in Section 5.) is responsible for effectively dealing with cases where renewable energy fluctuates due to an environmental status. This module copes with the fluctuating renewable energy through a distributed UPS system.
- **Workload Smoothing** The workload smoothing module is focused on treating overload conditions, in which workload exceeds the underlying data center’s computing and energy capacity, or cases when deferral is a better choice, depends on prediction. This module is invoked in Steps 3 and 10 of Algorithm 1 in Section 5.
- **UPS storage controller** The responsibility of the UPS storage controller module is two-fold. First, this module keeps track of the status of all UPS units in a data center. Second, the module is in control of making battery charging and discharging decisions. Refer to Algorithm 3 in Section 5 for the design of the UPS-usage controller module.

## 4 PROBLEM STATEMENT

In this section, we formally formulate the problem statement of this study by building a time-slotted model. To facilitate the

presentation of the model, the notations used throughout this section are summarized in Table II.

Table II. Symbol and Annotation

Symbol	Annotation
$E(t)$	total energy needed in time slot $t$
$E^W(t), E^S(t)$	energy supply of wind and solar in time slot $t$ , respectively
$s(t)$	the solar irradiance trace during time slot $t$
$v(t)$	the average wind speed trace during time slot $t$
$E^G(t)$	energy supply needed from power grid in time slot $t$
$E^{UPS}(t)$	net energy supply from UPS device in time slot $t$
$E_{storage}^{UPS}(t)$	currently stored energy in UPS device
$E_{min}^{UPS}, E_{max}^{UPS}$	lower and higher bound of UPS device storage capability
$P^W, P^S$	price of wind and solar energy, respectively
$P^G(t)$	price of grid energy at time slot $t$ , respectively
$W(t)$	the total workload need to be executed during time slot $t$
$w(t)$	the workload generated by users in time slot $t$
$w_d(t)$	the remaining and undeferable workload from time slot $t$
$C(t)$	total cost per unit of energy in time slot $t$

### 4.1 Modeling Energy Sources

In this subsection, we model multifaceted energy sources including renewable energy, grid power, and UPS systems. Given energy supply sources, we denote the total energy supply of the current time slot as  $E(t)$ . In a general scenario,  $E(t)$  is derived from (1) grid energy drawn from off-site power grid  $E^G(t)$ , (2) renewable energy generated from on-site facilities  $E^{ren}(t)$ , and (3) energy previously stored in UPS devices  $E^{UPS}(t)$ . Thus, total energy supply  $E(t)$  is expressed as

$$E(t) = E^G(t) + E^{ren}(t) + E^{UPS}(t)$$

In what follows, we construct the models of these three energy sources.

#### 4.1.1 Grid Energy Model

The grid energy cost of a datacenter is determined by the product of two factors, namely, (1) the grid energy price and (2) the amount of grid energy consumed by the datacenter. Recall that the main goal of REDUX system is to aggressively utilize renewable energy by employing UPS devices with high reliable condition. The amount of grid energy utilized in the current time slot,  $E^G(t)$ , will then reliant on (1) workload status of the current time slot, (2) price level of grid energy, and (3) the available amount of renewable energy combined with energy stored in UPS devices. Let  $P(t)$  be the grid energy price during time slot  $t$ , which is a random variable determined by the electricity market [10]. We assume that the time slot (e.g., one hour) is small enough so that  $P(t)$  remains stable during each time slot, unlike the system developed by Deng *et al.* [10]. We also assume the datacenter is supplied by a power grid source with infinite availability compares to the energy demand of the datacenter, meaning that there is no maximum bound to  $E^G(t)$ . Hence, the necessary off-site grid energy cost in time slot  $t$  can simply written as  $P(t)E^G(t)$ .

### 4.1.2 Renewable Energy Model

In order to model renewable energy  $E^{ren}(t)$  in time slot  $t$ , we consider solar and wind energy - two popular renewable energy sources. Let  $E^W(t)$  and  $E^S(t)$  be the wind and solar energy amount generated in time slot  $t$ . We express  $E^{ren}(t)$  as

$$E^{ren}(t) = E^W(t) + E^S(t).$$

It is reasonable to assume solar irradiance and wind speed remain stable within any time slot if the time window is small (e.g., 1 hour). Following renewable energy models developed by Richard, Komp, and Gipe [33] [34], we construct the solar energy model as:

$$E^S(t) = \alpha \cdot As(t) \cdot l,$$

where  $s(t)$  is the solar irradiance trace during time slot  $t$ ,  $A$  is the total active irradiance area of all solar panels,  $\alpha$  represents the solar-to-electricity conversion efficiency rate, and  $l$  indicates the duration of a time slot. It is mindful that the total active irradiance area  $A$  is unlikely to change during the investigated time horizon; such an assumption is validated in the prior studies [33] [34].

The wind energy model is written as:

$$E^W(t) = \beta \cdot (1/2)Bpv^3(t) \cdot l,$$

where  $v(t)$  is the average wind speed trace during time slot  $t$ ,  $B$  is the total rotor area of all wind turbines,  $\beta$  represents the wind-to-electricity conversion efficiency rate,  $\rho$  is the air density, and  $l$  is again the duration of a time slot. Similar to active irradiance area  $A$ , rotor area of all wind turbines  $B$  remains unchanged during the course of the studied time horizon.

It is worth to notice that in both of these two models, the only changing variable according to different time slot  $t$  are only solar irradiance  $s(t)$  and average wind speed  $v(t)$ . Which means all other variables will act as constants as long as we focus our research in an appropriate the time horizon, or determined by experiment settings such as the duration of time slot  $l$ .

After modeling the above two renewable energy sources, we denote the unit prices of solar and wind energy as  $P^S$  and  $P^W$ , respectively. Unlike the grid power price, the price of renewable energy includes one-time capital expense and long-term operational costs. Renewable energy price is calculated by averaging the total cost of installment and operation over the total amount of energy generated during an entire available period [35]. Modeling and deriving renewable energy prices is beyond the scope of our research. Therefore, we set these prices as a constant in our experiment configurations. In the not-too-distant future, we will delve in the impact of dynamically changing renewable energy price on our system performance.

### 4.1.3 UPS Energy Model

Traditional UPS devices merely act as transitional fail-over mechanisms between utility and captive sources such as diesel generators. We pay particular heed to a powerful distributed UPS system that not only powers up a data center when grid power is interrupted, but also is capable of storing a noticeable amount of energy when grid energy price is cheap or renewable energy is abundant. Such a distributed UPS system provides power supply when grid energy price is high or renewable supply becomes unreliable.

As shown in our framework design, batteries can be recharged by either off-site power grid or on-site renewable energy generators. Let  $E^{UPS}(t)$  denote the net energy supply by UPS devices, with positive value as discharge and negative value as recharge.

Throughout the lifespan of battery there are two affecting factors determine battery failure rates, namely, (1) the depth of discharge (DOD) and (2) discharge rates. DOD is gauged as a percentage of energy storage; a discharge rate is represented as discharge current specified as a C-rate. We configure the DOD lower bound of UPS devices (e.g., 25%) to minimize the battery failure rate. Such a DOD lower bound is a secured level to boost battery reliability by tackling the running-down issue induced by power outages. We setup the discharge and recharge current of UPS batteries as a reasonable constant.

To stay focused on this study, we also follow a reasonable assumption imposed by Uргаonkar *et al.* [31], stating that (1) batteries lost no power either when recharging and discharging, and (2) the stored energy will not decrease except discharged (i.e., batteries are not "leaky"). The batteries have a finite and unchanged capability of energy storage,  $E_{max}^{UPS}$ . On the other hand, it is required to maintain a minimum energy storage amount in UPS devices for the reliability purpose; in case of both grid and renewable outage, UPS devices must provide sufficient energy until a secondary power source becomes available. Given a lower bound  $E_{min}^{UPS}$  of energy levels in a UPS device's batteries, we have the following expression where UPS energy  $E^{UPS}(t)$  is in a range between the upper bound  $E_{max}^{UPS}$  and lower bound  $E_{min}^{UPS}$ . Last but not least, there is one noticeable idea for the setting of  $E_{min}^{UPS}$ . To maintain the function as backup power, we could set  $E_{min}^{UPS}$  to a certain portion to the  $E_{max}^{UPS}$  (20

$$E_{min}^{UPS} \leq E^{UPS}(t) \leq E_{max}^{UPS}.$$

Furthermore, for the purpose of extending a healthy lifetime of batteries, in each time slot we restrict the use of batteries in a way that one can either recharge, discharge the batteries, or neither; the batteries can not be recharged and discharged at the same time (see, for example, a recent study by Deng *et al.* [10]). We define  $E_t^r$  and  $E_t^d$  respectively as energy rechargeable to or dischargeable from batteries in time slot  $t$ . The following constrains should be met:

$$0 \leq E_t^r \leq E_{max}^{UPS} - E^{UPS}(t), \quad (1)$$

$$0 \leq E_t^d \leq E^{UPS}(t) - E_{min}^{UPS}. \quad (2)$$

Eq. 1 indicates that the possible amount of recharged energy is upper bounded by the difference between the capability and the energy storage level in batteries during time slot  $t$ . Similarly, Eq. 2 signifies that the energy discharge level is lower bounded by the difference between the minimum safety level and the energy storage level in batteries during the time slot.

As the cost of recharged energy is incorporated in the grid and renewable power cost, the battery power cost is simplified to a one-time capital expenditure only. Suppose a battery that is worth  $Q$  dollars has a finite non-negative integer  $N$  of recharge/discharge times, then we model the cost of UPS devices as  $P^{UPS} = Q/N$  during the device lifespan. Though such cost becomes negligible compared to grid or renewable energy cost, we do consider the UPS cost in our experiments.

## 4.2 Modeling Workloads

Prior to the development of a workload model, we first measure the total computational workload  $w(t)$  generated from datacenter users in time slot  $t$ . Since job-power profiles are missing in the job traces, we apply the power-profile function to derive workload  $w(t)$  of each job during time slot  $t$ . Power demand of the individual job is then calculated by the job's allocated number of processors and running time during time slot  $t$ , followed by measuring power demands in terms of  $kW/h$ . The details of the job power profile function can be found in section 6.1.

Workload  $w(t)$  is now gauged in terms of power demand as  $kW/h$ . All these workloads vary in a wide range, taking values from a non-negative finite set, with a maximum value denoted as  $w_{max}$ . Thus, we have

$$0 \leq w(t) \leq w_{max}.$$

To simplify the design while meeting our experimental needs, there is no necessity to keep track of the underlying probability distribution or statistical characterization of  $w(t)$ . Nevertheless, we conclude that the overall workload generated during a time horizon  $T$  should match the total energy consumed by a data center. We formally express this trend as:

$$\sum_{t=1}^T E(t) = \sum_{t=1}^T w(t). \quad (3)$$

Because the quality of service (QoS) level offered a data center highly depends on whether the data center can handle the workloads in a timely manner, it is nature to build a mechanism to deal with cases in which heavy workload in a certain time slot exceeds the data center's computing capability. To cope with such heavy workload conditions, REDUX smartly defers computing loads to the next subsequent time slot. It is worth noting that such a computing deferral strategy is applicable to the other three scenarios, namely, (1) high grid power price and (2) outage of renewable energy.

Following the assumption imposed by Goiri *et al.* [12] and Urgaonkar *et al.* [31], it is arguably true that a data center's energy demand is either satisfied in the current slot or delayed to the next time slot due to the above three scenarios. In other words, our scheduling policy makes jobs wait for a limited amount of time (e.g., one time slot) to acquire an improved QoS performance (see also a similar design philosophy by Zaharia *et al.* [36]). Now, we extend Eq. 3 to model workload deferral cases below.

$$\forall t \in T : \sum_{t'=1}^{t'=t} E(t') \leq \sum_{t'=1}^{t'=t} w(t'), \quad (4)$$

Eq. 4 implies the accumulated energy demand from workloads may greater or equal than the total energy supply. Eq. 4 is reasonable, because in order to maintain a relatively high QoS level for any time slot  $t'$ , some heavy computing loads are deferred to the next time slot by the delay scheduling method.

In delay scheduling, the system strives to finish deferred workloads in time slot  $t+1$  by giving a high priority to the deferred workloads. In doing so, deferred tasks are unlikely to be further delayed to avoid the starvation problem. Let  $w_d(t)$  be delayed workload in time slot  $t$ . We set deferred workload as undeferrable in slot  $t+1$ , placing the delayed workload at the beginning of an FIFO queue in the next time slot. We denote

$W(t)$  as the total amount workload to be executed by the data center in time slot  $t$ . Hence, we express  $W(t)$  as

$$W(t) = w(t) + w_d(t-1).$$

In a nutshell, the priority of delayed  $w_b(t-1)$  is higher than that of the newly submitted workload  $w(t)$ . Ideally, jobs in high-priority  $w_b(t-1)$  should be accomplished before executing the first job in  $w(t)$ .

Recall that  $w(t)$  as well as  $W(t)$  are measured in terms of energy consumption in the unit of  $kW/h$ ;  $E(t)$  is defined as the total energy supply consumed by a data center handling all the workloads during time slot  $t$ . Given time slot  $t$ , we express the relationship between consumption  $W(t)$  and supply  $E(t)$  as

$$W(t) = E(t).$$

## 4.3 Modeling Total Cost

Now, we are positioned at construct a total-cost-optimization model that incorporates with all the constrains presented in Sections 4.1 and 4.2.

We define  $C(t)$  as the total cost of energy supply in time slot  $t$ . Cost  $C(t)$  is derived from grid-power cost, renewable energy cost, and the UPS-device cost. Thus, cost  $C(t)$  can be calculated as

$$C(t) = P^G(t) \cdot E^G(t) + P^{ren}(t) \cdot E^{ren}(t) + P^{UPS} \cdot E^{UPS}(t).$$

The overall optimization goal of our system is to minimize the cumulative  $C(t)$  through time horizon  $T$  under all the aforementioned constrains. We formally describe the goal as follows

$$\text{Min} \sum_{t=1}^T [P^G(t) \cdot E^G(t) + P^W(t) \cdot E^W(t) + P^{UPS} \cdot E^{UPS}(t)]$$

$$\text{s.t. } W(t) = E(t)$$

$$E^G(t) + E^{ren}(t) + E^{UPS}(t) = E(t)$$

$$W(t) = w(t) + w_b(t-1)$$

$$0 \leq w(t) \leq w_{max}$$

$$0 \leq E^G(t) \leq E_{max}^G$$

$$0 \leq E^{ren}(t) \leq E_{max}^{ren}$$

$$E_{min}^{UPS} \leq E^{UPS}(t) \leq E_{max}^{UPS}$$

It is clear that the above optimization goal has trade-off between power cost, renewable energy reflected in the constraints. By solving the preceding problem, we can obtain the corresponding power management and load scheduling strategies to optimize the system.

- Power Management Strategy: For a data-center with on-site renewable energy facilities and fluctuating grid power price, REDUX can determine the amount of power supplied by the grid at time slot  $t$   $E^G(t)$ , and the amount of power recharged to the UPS devices  $E^{UPS}(t)$ .
- Load Scheduling Strategy: The workload smoothing module can smoothing the workload generated by users and decide whether to dispatch more or less the workload  $W(t)$  to the data-center.

## 5 HEURISTIC ALGORITHMS

We first present a high-level control algorithm in Section 5.1. Then, Sections 5.2 and 5.3 shed light on the design of renewable smoothing and UPS storage control algorithms implemented in the backbone layer (see also Section 3.3) of REDUX.

## 5.1 The High-Level Control Algorithm

We take a hierarchical design approach to the development of REDUX, in which a high-level control algorithm coordinates all the three core modules in the backbone layer (see the backbone layer in Section 3.3). This subsection is dedicated to the design of the high-level controller, whereas the subsequent subsections articulate the algorithms of the core modules, namely, renewable energy smoothing and UPS-storage controller (see also Sections 5.2 and 5.3).

---

### Algorithm 1 The high-level control algorithm in REDUX.

---

**Input:**

Renewable energy trace data  $E^W(t)$  and  $E^S(t)$   
 UPS energy storage  $E^{UPS}(t)$  from previous time slot  
 Generated workload  $W(t)$  and deferred workload  $w_d(t)$   
 Price of all energy source  $P^W, P^S, P^{UPS}$  and  $P^G(t)$

**Output:**

REDUX Cost  $C(t)$  of All time slots

```

1:  $t = 0$ 
2: while  $t \leq T$  do
3:    $workload\_shaving()$ ;
4:    $workload\_smoothing()$ ;
5:   determine renewable energy supply state  $S_{RE} = STA, FLU$  or  $OTA$ ;
6:   if  $S_{RE} = FLU$  then
7:      $ren\_supply\_smooth()$ ;
8:   end if
9:   update over-peak workload threshold then determine workload over-peak states  $S_w = OVR$  or  $OFFP$ ;
10:  if ( $S_w = OVR$ ) then
11:     $workload\_shaving()$ ;
12:  end if
13:   $ups\_storage\_control()$ ;
14:   $E_t^G = W_t - E_t^{RE} - E_t^{UPS}$ ;
15:   $calculate\_energy\_cost()$ ;
16:   $t++$ 
17: end while
18: return REDUX cost

```

---

Algorithm 1 depicts the pseudo-code of the high-level algorithm in REDUX, where the inputs include (1) all energy supply data, namely, wind energy  $E^W(t)$ , solar energy  $E^S(t)$  and energy storage amount in UPS device  $E^{UPS}(t)$ , (2) workload data  $W(t)$  generated by users and previous deferred workload  $w_d(t)$  and (3) grid energy price  $P^G(t)$  as variable and  $P^W, P^S, P^{UPS}$  as constants.

The high-level controller kicks off the power management task with the workload-shaving procedure (see  $workload\_shaving()$ ; in Line 3). The purpose of this first workload shaving is the overloading of a data center with respect to power supplies and computing capacity. More specifically, if the submitted workload is below the overloading threshold, the pending jobs will immediately be processed in the current time slot. Otherwise, these jobs will be delay scheduled by the high-level controller. Jobs resulting in overload should be scheduled and processed in the subsequent time slots. It is noteworthy that scheduling such jobs under the overloading condition imposes significant impact on the data center's overall performance metrics such as throughput and response time. The investigation of scheduling policies in the overloading case is out of scope of this study. We intend to address this concern as a future research direction.

The procedure of workload smoothing in Line 4 (i.e.,  $workload\_smoothing()$ ) aims to mitigate possible dramatic workload changes between subsequent time slots. Prior to kicking in this procedure, an exponential windowed average result of a previous workload should be maintained as a threshold. If the current workload is greater than the threshold (i.e., exponential window average), the workload smoothing method will simply update the current workload level with a smoothed average result. If the current workload is smaller than the threshold, the procedure will increase the current workload by handling deferred workload if there is any. In this case, the current workload level is increased until the level hits the threshold.

In the high-level control algorithm 1, the next task after workload smoothing is to determine the renewable-energy state  $S_{RE}$  (see Line 5), which has three candidate cases, namely, stable (i.e.,  $STA$ ), fluctuate (i.e.,  $FLU$ ) and outage (i.e.,  $OTA$ ). To categorize these states, if the renewable supply is at the minimum level, the renewable states will be defined as  $OTA$  (see the purple dots in Fig. 2). If the renewable-energy supply is beyond the maximum level, the renewable states will be referred to as  $STA$  (see the gray dots Fig. 2). For the other cases, we compare the renewable-energy supplies with their previous exponential windowed average. If the change in the renewable supply from the previous time slots is in a "stable" interval, we envision this current renewable supply as  $STA$ ; otherwise, the state should be categorized as  $FLU$  (values unannotated in Fig. 2).

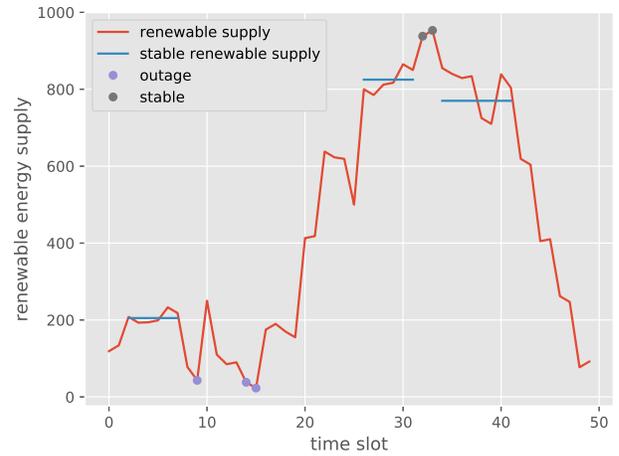


Figure 2. An example for the categorization of renewable energy supplies. The unannotated values are defined as fluctuate states.

After REDUX determines the state of the current renewable supply, REDUX further manages workload in accordance to the state. More specifically, if the renewable energy is in the fluctuation state, the renewable-energy-supply smoothing policy is invoked in Line 7. The smoothing procedure is governed by the current renewable supply data, a dynamically updating stable supply level, and UPS energy storage condition. Please refer to Section 5.2 for the details of this procedure.

After the fluctuating renewable supply is smoothed, the renewable and UPS energy supplies are finalized. To further improve the performance of REDUX with these measures in place, REDUX dynamically updates an over-peak workload threshold. Such an update is accomplished by providing the information pertaining to the current workload amount, re-

newable and UPS energy supply and grid energy price. From this threshold, REDUX can determine whether or not a second workload shaving is necessary by the workload's over-peak state  $S_W$ , either over-peak (i.e.,  $OVR$ ) or off-peak (i.e.,  $OFF$ ) (see Line 9). If the workload exceeds the over-peak threshold, then the workload-shaving procedure is kicked in to defer the extra load to the next time slot (see Line 11). Again, such second workload shaving makes a positive impact on the data center's overall performance as the first shaving.

The `ups_storage_control()` procedure (see Line 13) is responsible for protecting UPS units against frequent charging and discharging, thereby prolonging lifetime of the distributed UPS system. This procedure determines a UPS energy level, under which the UPS units should be protected. Such levels largely depend on current workload conditions, grid energy price as well as renewable energy capacity. The functionality of `ups_storage_control()` is two-fold: (1) updating the UPS status information and (2) making UPS discharging and charging decisions.

Recall that vital factors include workload  $W_t$ , renewable energy supply  $E_t^{RE}$  and net energy supply from UPS devices  $E_t^{UPS}$ . After all these factors are determined, REDUX computes the energy supply obtained from the grid network (see Line 14). The `calculate_energy_cost()` function in Line 15 calculates the total energy cost occurred in time slot  $t$ . This function also enables us to quantitatively compare REDUX against other exiting power management strategies (see also Section 6).

## 5.2 Smoothing Renewable Energy Supply

Recall that Line 7 in Algorithm 1 is a smoothing operation for renewable energy supplies. This operation is dedicated to smooth renewable energy output for data centers.

---

### Algorithm 2 Smoothing Renewable Energy Supplies

---

**Input:**

grid price  $P^G(t)$ , and state  $S_{GP}$   
renewable price  $P^{RE}$ , renewable supply  $E^{RE}$   
updated stable renewable supply  $E^{STA}$

**Output:**

updated renewable supply  $\bar{E}^{RE}$   
updated stable renewable supply  $\bar{E}^{STA}$

```

1: if  $E^{RE} < E^{STA}$  then
2:   if  $S_{GP} = HIGH$  then
3:     Discharge UPS units until stable renewable supply;
4:   else
5:     Use the power grid for smoothing;
6:     if  $P^G(t) < P^{RE}$  then
7:       Recharge the UPS units;
8:     end if
9:   end if
10: else
11:   Use renewable energy to power servers;
12:   Recharge the UPS units using renewable energy;
13: end if

```

---

Algorithm 2 outlines the pseudo-code of the smoothing operation. First of all, to simulate a stable threshold for the following steps, we again use the exponential windowed average method to figure out the result of stable renewable supply level. Then, if the renewable energy supply is lower than the stable threshold (i.e., see  $E^{RE} < E^{STA}$  in Line 1), the renewable energy output fluctuates due to renewable

resource variations. In such a fluctuation case, energy should be supplemented by the underlying UPS systems (see Lines 2-3) or the power grid (see Lines 5-8). Specifically, when the price of the electricity grid is high, the UPS units discharge and supply energy to the servers (see Line 3). In case the power grid offers cheap energy, REDUX fully utilizes the grid to smooth the renewable energy supply (see Line 5). It is economically wise to recharge the UPS system using the power grid at cheap price (see Line 7). The servers will be powered by cost-effective renewable energy if the renewable energy generation is sufficient (see Line 11). Further, when the renewable resources are abundant, the UPS devices are recharged by the renewable energy (see Line 12).

## 5.3 Updating UPS Energy Supply

Now we delineate an algorithm for UPS energy supply and storage control. This algorithm is invoked in Step 13 in Algorithm 1 (see Section 5.1. The UPS storage control algorithm keeps track of the status of the UPS energy supply and storage (i.e.,  $E^{UPS}(t)$  and  $E_{storage}^{UPS}(t)$ ).

---

### Algorithm 3 Updating UPS Energy Supply

---

**Input:**

Energy demand from workload amount in time slot  $t$ ,  $W(t)$   
updated over-peak workload,  $W_{OVR}$   
Renewable energy states,  $S_{RE} = \{STA, FLU, OTA\}$   
Grid price states,  $S_{GP} = \{HIGH, LOW\}$   
UPS energy storage,  $E_{storage}^{UPS}(t)$

**Output:**

net UPS energy supply,  $E^{UPS}(t)$   
updated UPS energy storage,  $\bar{E}_{storage}^{UPS}(t)$

```

1: if  $S_{RE} = OTA$  then
2:   if  $W_t \geq W_{OVR}$  and  $S_{GP} = HIGH$  then
3:     Discharge UPS units;
4:   end if
5:   if  $W_t < W_{OVR}$  and  $S_{GP} = LOW$  then
6:     recharge UPS units until full
7:   end if
8: end if
9: if  $S_{RE} = FLU$  then
10:  if  $W_t < W_{OVR}$  and  $S_{GP} = LOW$  then
11:    Recharge the UPS units;
12:    Update the UPS supply with  $E^{UPS}$ ;
13:  end if
14: end if
15: if  $S_{RE} = STA$  then
16:  if  $S_{GP} = LOW$  or ( $W_t < W_{OVR}$  and  $S_{GP} = HIGH$ ) then
17:    Recharge the UPS units while renewable energy is accessible;
18:  end if
19: end if

```

---

Algorithm 3 illustrates the procedure of updating UPS energy supply and storage information under three distinctive renewable-energy conditions, namely, the outage (see Lines 1-8), fluctuation (see Lines 9-14), and stable state (see Lines 15-19).

If a renewable energy outage occurs, energy must be supplied by either the power grid or the distributed UPS system. REDUX chooses to discharge the UPS units if the current workload is high and the power price from the grid is expensive

(see Lines 2-3); the UPS devices are recharged when the grid price becomes cheap (see Lines 5-6).

When the renewable energy resource is in the fluctuation state and the current workload is below the peak level, the UPS units are recharged by cheap power grids (see Lines 10-11).

In case of abundant renewable energy resources (i.e.,  $S_{RE} = STA$ ), REDUX decides to recharge the UPS units under one of the two conditions: (1) the power price of the grid is fairly low or (2) current workload is in the off-peak window and the power grid price is high (see Lines 16-17).

## 6 RESULTS AND ANALYSIS

In this section, we evaluate the performance of REDUX driven by real-world renewable energy and power grid price data. We compare REDUX with a baseline solution called *noREDUX* as well as two existing power management techniques - *GreenSlot* [11] and *ReUPS* [8]. We measure the accumulated energy cost of a simulated data center governed by REDUX, *noREDUX*, *GreenSlot*, and *ReUPS* under various conditions.

Before presenting experimental results, let us discuss the experimental settings in Section 6.1.

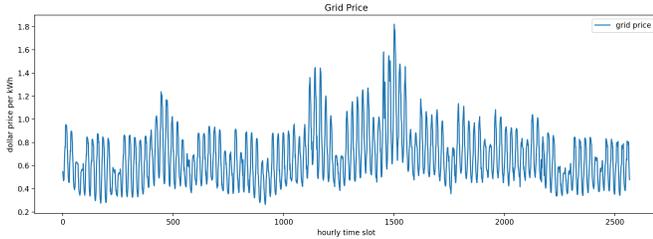


Figure 3. The electricity price of a power grid in the New York state during the period of three month ranging from June 1 to August 31, 2006. Unit:  $\times 100\$/MWh$

### 6.1 Experimental Settings

#### 6.1.1 Power Demands

We conduct extensive experiments using real-world workloads recorded in the Parallel Workloads Archive (a.k.a., PWA) [37] [38] – a popular repository of parallel computing workloads. Traces archived in PWA are job logs collected from large-scale parallel computing systems. Each job information includes job ID, running time, the number of allocated processors, used memory resources and the like. The specific trace evaluated in this study is LLNL Atlas, which contains 8-month (Nov 2006 to Jun 2007) accounting records from a large Atlas Linux cluster. The Atlas cluster is comprised of 1152 nodes in 32 server racks (36 nodes per rack), each of these nodes contains eight AMD Opteron processors clocked at 2.4 GHz and 16 GB of memory. The total number of processors in the cluster is 9216.

All the job trace data acquired from *PWA* represent the behavior of a single parallel system, which is relatively small compared with large-scale computing platforms housed in data centers. We expand the traces along the following two directions to address this concern. First, we scale up the size of the LLNL Atlas platform by a factor of 5 to mimic a large system that is comprised of 46,080 processors ( $5 \times 9216$ ) with a homogeneous configuration. Second, we scale up by a factor of five the number of required processors for each job to resemble the size of a data center.

The time slot length is set to one hour in our experiment settings to (1) accommodate large batching jobs and (2) match the measurement unit of energy consumption as *kWh*. Therefore, if a large job started during any time slot requests an execution time exceeding one hour, the REDUX system will label this job as “paused” at the end of this time slot. Here “paused” is not a sign of deferred workload; rather, “paused” signifies that there is no time gap between two consecutive time slots. More importantly, “paused” serves as a “stop sign” of accumulated current workload and a starting point of counting workloads in the next time slot.

We implement a job power profiling module to render a power profile for each job executed during time slot  $t$ . The power profiling module takes the number of allocated processors and running time during slot  $t$  of each job as inputs. Following a guideline documented in [39] and [40], we set the power demand of a node to a value anywhere between 20 to 60 *Watt* using a normal distribution; This setting will lead to the following two deductions: (1) the profiling module then assigns the power demand of each processor in a range between 2.5 and 7.5 $W$  following a normal distribution; and more importantly (2) as the number of nodes of the entire data center was set as 46080, the maximum execution capability of the whole data center in terms of energy consumption can then be calculated as  $46080 \times 60 Watt \times 3600 Secs = 9.95 \times 10^9 J = 2764.8 kWh/h$ . The power profile of each job running in slot  $t$  is then multiplied by the power dissipation of all the allocated processors and the job running time. The final step of the profiling module is to accumulate all the jobs’ power demands during slot  $t$ , thereby converting job trace data into power demand measured in *kWh* in slot  $t$ .

#### 6.1.2 Power Supply

In our experiments, we test REDUX and its competitors using the electricity price of a real-world power grid in New York state during a period of three months from June 1 to August 31, 2006. We pay attention to this time period, because the solar data was collected during this period. Fig. 3 shows that the power grid price dynamically changes during the three-month time interval, where three noticeable price peaks occur at around 480, 1200, and 1500 hours. The maximum and minimum prices in the trace are 1.82 and 0.26  $\$/Kwh$ , respectively.

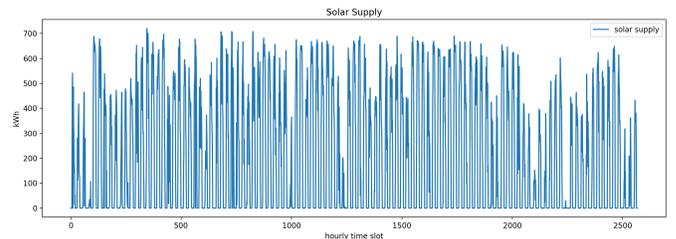


Figure 4. The solar supply of New York state during the three-months period ranging from June 1 to August 31, 2006.

From the National Renewable Energy Laboratory’s (NREL) database, we retrieve the solar supply of New York state during the three-month period also from June 1 to August 31, 2006. The solar supply data measured in *KWh* is plotted in Fig. 4. Follow our previous problem statement design, the solar and wind power price is set as constant at  $\$0.09/Kwh$  and  $\$0.07/Kwh$ , respectively.

Last but not least, the distributed UPS system - 6 per rack (192 per cluster computing system) - is scaled up by a factor of 10 for the entire large data center. And the energy cost for recharging and discharging batteries is calculated as \$0.07/kWh from the UPS model we use.

### 6.2 Predicting from Smoothed Workload

Recall that in this research we introduce a simple yet effective prediction model, where an exponential windowed average is continuously updated. These predicted results serve as a threshold or a baseline to guide the relative smoothing method carried out in both workload smoothing and renewable energy smoothing. With each updated threshold on hand, the REDUX will then determine whether shave and defer the extra workload from the current time slot to the next. Fig. 5 depicts an overall view of the smoothing method, which makes predictions on workload prior to any workload shaving. Intuitively, workload distribution after the exponential windowed average (see the red dots in Fig. 5) is more convergent to the total average. When we sort and plot both these data traces, this averaging method offers a flat sloped curve (i.e., see the red curve). We draw similar observations when we apply this prediction method to grid price.

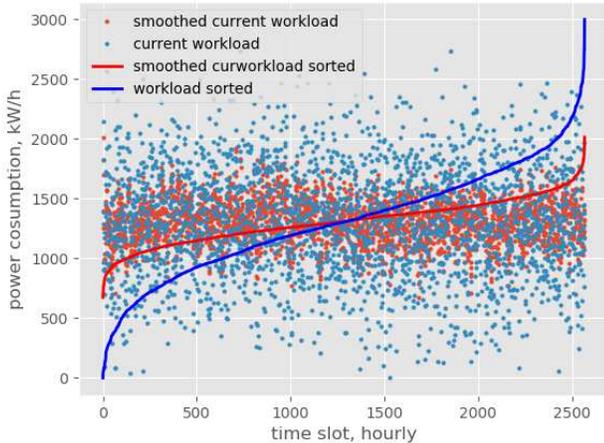


Figure 5. Total cost of REDUX during specific time slot maybe more than noREDUX

### 6.3 REDUX vs noREDUX

In the first experiments, we compare REDUX with noREDUX - a baseline power manager where renewable energy resources and UPS units are not incorporated. More specifically, noREDUX is a scheme where renewable energy is just passively utilized as it generated, and the energy stored in UPS devices (if there exists) is not actively utilized.

Fig. 6 shows a few time slots are outliers, where total energy cost for a REDUX-enabled data center is higher than a noREDUX data center. This may be because the deferred workloads from previous time slot has to be processed while noREDUX system do not have this mechanism. Nevertheless, REDUX judiciously cuts back, in a majority number of cases, the total cost of a data center.

Fig. 7 again reveals the superiority of the REDUX system against noREDUX systems by an accumulative metric. Regardless of REDUX or noREDUX, the cumulative energy cost

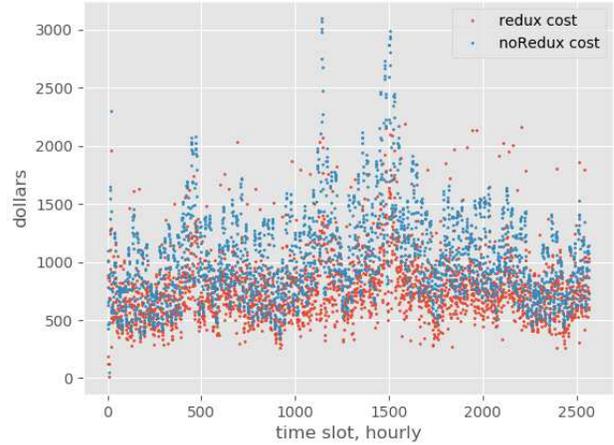


Figure 6. Total cost of a data center with REDUX during most time slot are lower than noREDUX

linearly increases with time. The results show that REDUX significantly cut back the energy of the data center governed by noREDUX. For example, REDUX is able to reduce the energy cost of noREDUX by 24.5% by the end of the first month; such cost savings becomes 27.9% and 33.2% by the end of the second month and third month, respectively.

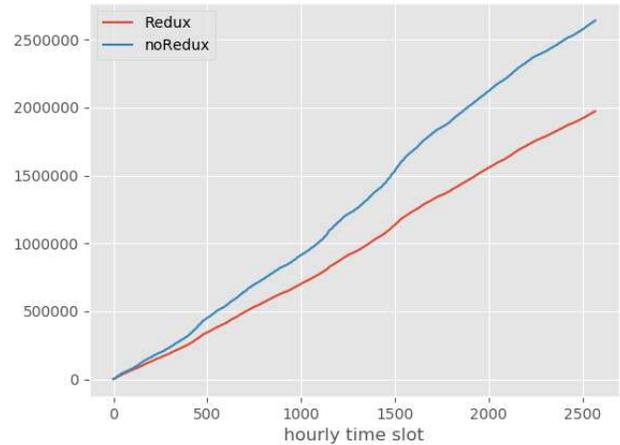


Figure 7. REDUX significantly reduces the energy cost of a data center managed by noREDUX - a traditional scheme.

The energy-cost savings are expected because REDUX orchestrates renewable energy resources and UPS units to conserve energy cost when the electricity price of the power grid is expensive. By defer some workload to the following time slot while still maintain a desirable QoS, a data center with REDUX can then smartly control the arrangement of workload and distribution of energy resources. For example, there are multiple sharp peaks (e.g., see 480, 1200, and 1500 hours in Fig. 3) in the grid price trace data. REDUX makes judicious decisions to power the data center using either renewable energy or the distributed UPS system rather than the expensive power grid.

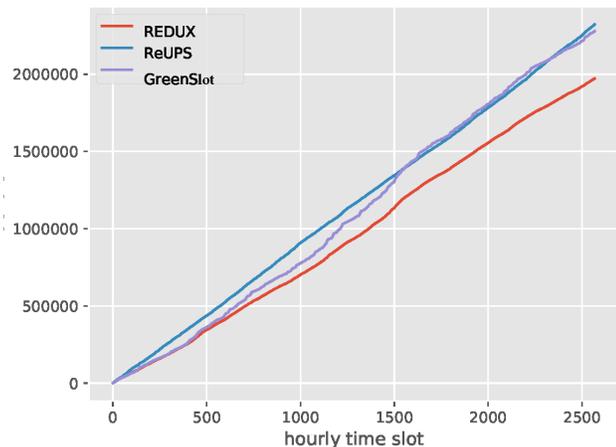


Figure 8. REDUX reduces the accumulated energy cost of a data center managed by other schemes

#### 6.4 Improving GreenSlot and ReUPS

Now we compare REDUX with two existing schemes - GreenSlot and ReUPS. Fig. 8 shows the cumulative energy cost of data centers managed by REDUX, GreenSlot, and ReUPS, respectively. Like the trend plotted in Fig. 7, the energy cost of the three management schemes grows with time.

More importantly, Fig. 8 illustrates that REDUX noticeably outperforms GreenSlot and ReUPS in terms of energy cost. For instance, by the end of the first month, REDUX slashes the energy cost of GreenSlot and ReUPS by 10.3% and 26.5%, respectively; such cost savings becomes 16.8% and 14.5% by the end of second month, and 15.5% and 17.7% by the end of the third month, respectively. Table III summarizes the cumulative energy cost of the three management strategies by the end of the first, second, and third month, respectively. The energy-cost saving rates of REDUX over GreenSlot and ReUPS can also be found in Table III.

REDUX is superior to GreenSlot, because REDUX makes use of the distributed UPS system to smooth renewable energy supply, which helps REDUX to smarter control the energy resources distribution over time slots. The improvement of REDUX over ReUPS is ascribed to the awareness of changing power grid price. REDUX makes full advantage of renewable energy coupled with the UPS devices to offer energy supply when the power grid price becomes high (see, for example, multiple sharp peaks in Fig. 3). When the power grid price is surging, REDUX switches power supply from the grid to either renewable energy resources or the UPS system, depending on current workload conditions.

#### 6.5 Impacts of UPS Capability

In the following experiments, we study the impacts of UPS capacity on energy efficiency of the REDUX system.

Fig. 9 unveils the impacts of UPS storage capability on accumulated total cost. In this group of experiments, we test three UPS capacity levels while keeping the other parameters unchanged. Though the improvement made by large UPS capability is marginal, the results still shows the advantage of large UPS devices on total cost savings in the long run. We conclude that given UPS devices with the high capability, RE-

DUX is offered ample opportunities to boost energy efficiency by storing low-cost energy for future usage.

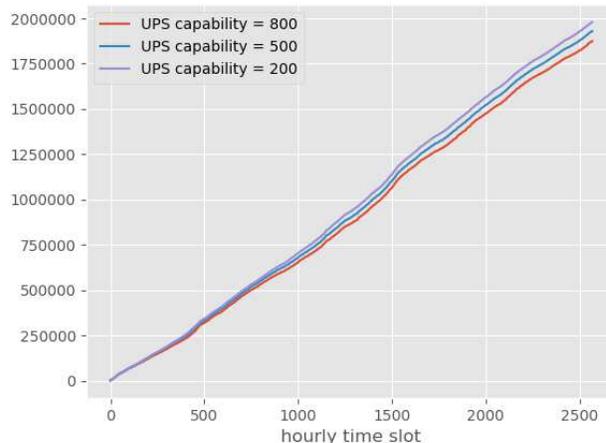


Figure 9. Total cost drops with an increased UPS capability.

## 7 CONCLUSIONS AND FUTURE WORK

In this paper, we presented a resource manager called *REDUX* to on-site renewable sources like solar coupled with a distributed UPS system to reduce energy cost of data centers. REDUX orchestrates UPS units to judiciously cope with the intermittency nature of renewable energy integrated with power grids. REDUX makes critical decisions on UPS charging and discharging with sufficient information on renewable-energy supply levels and time-dependent grid power prices. More specifically, REDUX makes prudent discharge of batteries to offer energy resources when renewable energy production is low or in the fluctuate condition. REDUX makes judicious decisions to charge batteries when renewable energy levels are high coupled with low loads.

One salient feature of REDUX is to effectively slash energy costs of data centers powered by both renewable and conventional energy. Compared with the prior solutions, REDUX delivers a prominent capability of mitigating average peak workload and boosting renewable-energy utilization. The experimental results demonstrate that REDUX paves the way for constructing modern data centers that are economically and environmentally friendly.

There are two future research directions. First, we will unravel a plan to explore scheduling policies during peak load. We expect that the scheduling policies undoubtedly have a tremendous impact on energy efficiency and performance of data centers. Second, the smoothing algorithm implemented in REDUX is a pilot study towards optimizing renewable energy. To extend the smoothing algorithm, we intend to propose a predictive scheme to make judicious smoothing decisions on renewable energy supplies.

#### ABBREVIATIONS

The following abbreviations are used in this manuscript:

Uninterruptible Power Supply	UPS
Power Distribution Unit	PDU
Dynamic Voltage and Frequency Scaling	DVFS
Depth of Discharge	DOD
Parallel Workloads Archive	PWA

Table III. Cumulative energy-cost comparisons among REDUX, GreenSlot, and ReUPS the end of the first, second, and third month.

Power Management	1st Month	1st Month Reduction Rate	2nd Month	2nd Month Reduction Rate	3rd Month	3rd Month Reduction Rate
REDUX (This Study)	\$604,821.44	N/A	\$1,330,663.89	N/A	\$1,972,270.34	N/A
GreenSlot [12]	\$765,215.58	10.3%	\$1,524,093.47	16.8%	\$2,321,717.9	15.5%
ReUPS [8]	\$667,445.95	26.5%	\$1,555,229.6	14.5%	\$2,278,287.49	17.7%

## DECLARATIONS

### Ethics Approval and Consent to Participate

Not Applicable

### Consent for publication

Not Applicable

### Availability of data and material

Solar Energy: New York state solar energy from National Renewable Energy Laboratory's (NREL) database, June 1 to August 31, 2006. Workload: Parallel Workloads Archive - LLNL Atlas Grid Power Price: New York State electricity price from nationalgridus.com database, June 1 to August 31, 2006.

### Competing interests

The authors whose names are listed immediately below certify that they have NO affiliations with or involvement in any organization or entity with any financial interest (such as honoraria; educational grants; participation in speakers' bureaus; membership, employment, consultancies, stock ownership, or other equity interest; and expert testimony or patent-licensing arrangements), or non-financial interest (such as personal or professional relationships, affiliations, knowledge or beliefs) in the subject matter or materials discussed in this manuscript.

### Funding

There is no funding related information.

### Authors' contributions

**Xiaopu Peng** The main author. Draft and revised the whole paper with original main idea. Conducted experiments and finalized the conclusion. **Tathagata Bhattacharya** Helped with collecting datasets and references **Ting Cao** Helped draft and revise the background and introduction sections **Jianzhou Mao** Helped draft and revise the framework and algorithms sections **Taha Al Tekreeti** Helped with draft and revise the experiments section **Xiao Qin** Helped polish the writing and guided the main idea and experimental part of the paper

### Acknowledgments

Xiao Qin's work is supported by the U.S. National Science Foundation under Grants IIS-1618669, CCF-0845257 (CAREER), CNS-0917137, CNS-0757778, CCF-0742187, CNS-0831502, CNS-0855251, and OCI-0753305. Xiao Qin's study is also supported by the 111 Project under grant No. B07038.

## REFERENCES

- [1] E. Baccarelli, N. Cordeschi, A. Mei, M. Panella, M. Shojafar, and J. Stefa, "Energy-efficient dynamic traffic offloading and re-configuration of networked data centers for big data stream mobile computing: review, challenges, and a case study," *IEEE Network*, vol. 30, no. 2, pp. 54–61, 2016.
- [2] 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, 2016.
- [3] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United states data center energy usage report," 2016.
- [4] L. Olatomiwa, S. Mekhilef, M. Ismail, and M. Moghavvemi, "Energy management strategies in hybrid renewable energy systems: A review," *Renewable and Sustainable Energy Reviews*, vol. 62, pp. 821–835, 2016.
- [5] L. Xie, Y. Gu, X. Zhu, and M. G. Genton, "Short-term spatio-temporal wind power forecast in robust look-ahead power system dispatch," *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 511–520, 2014.
- [6] J. M. Carrasco, L. G. Franquelo, J. T. Bialasiewicz, E. Galván, R. C. PortilloGuisado, M. M. Prats, J. I. León, and N. Moreno-Alfonso, "Power-electronic systems for the grid integration of renewable energy sources: A survey," *IEEE Transactions on industrial electronics*, vol. 53, no. 4, pp. 1002–1016, 2006.
- [7] C. Li, A. Qouneh, and T. Li, "iswitch: coordinating and optimizing renewable energy powered server clusters," in *Computer Architecture (ISCA), 2012 39th Annual International Symposium on*. IEEE, 2012, pp. 512–523.
- [8] L. Liu, H. Sun, C. Li, Y. Hu, J. Xin, N. Zheng, and T. Li, "Re-ups: an adaptive distributed energy storage system for dynamically managing solar energy in green datacenters," *The Journal of Supercomputing*, vol. 72, no. 1, pp. 295–316, 2016.
- [9] C. Li, R. Wang, T. Li, D. Qian, and J. Yuan, "Managing green datacenters powered by hybrid renewable energy systems." in *ICAC*, 2014, pp. 261–272.
- [10] X. Deng, D. Wu, J. Shen, and J. He, "Eco-aware online power management and load scheduling for green cloud datacenters," *IEEE Systems Journal*, vol. 10, no. 1, pp. 78–87, 2016.
- [11] Í. Goiri, R. Beauchea, K. Le, T. D. Nguyen, M. E. Haque, J. Guitart, J. Torres, and R. Bianchini, "Greenslot: scheduling energy consumption in green datacenters," in *High Performance Computing, Networking, Storage and Analysis (SC), 2011 International Conference for*. IEEE, 2011, pp. 1–11.
- [12] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, "Parasol and greenswitch: Managing datacenters powered by renewable energy," in *ACM SIGARCH Computer Architecture News*, vol. 41, no. 1. ACM, 2013, pp. 51–64.
- [13] S. Govindan, D. Wang, L. Chen, A. Sivasubramaniam, and B. Urgaonkar, "Towards realizing a low cost and highly available datacenter power infrastructure," in *Proceedings of the 4th Workshop on Power-Aware Computing and Systems*. ACM, 2011, p. 7.
- [14] V. Kontorinis, L. E. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, D. M. Tullsen, and T. S. Rosing, "Managing distributed ups energy for effective power capping in data centers," in *Computer Architecture (ISCA), 2012 39th Annual International Symposium on*. IEEE, 2012, pp. 488–499.
- [15] B. Aksanli, T. Rosing, and E. Pettis, "Distributed battery control for peak power shaving in datacenters," in *Green Computing Conference (IGCC), 2013 International*. IEEE, 2013, pp. 1–8.
- [16] D. Cheng, X. Zhou, P. Lama, M. Ji, and C. Jiang, "Energy efficiency aware task assignment with dvfs in heterogeneous hadoop clusters," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 1, pp. 70–82, 2018.
- [17] M. Geveler, D. Ribbrock, D. Donner, H. Ruelmann, C. Höppke, D. Schneider, D. Tomaschewski, and S. Turek, "The icarus white

- paper: a scalable, energy-efficient, solar-powered hpc center based on low power gpus," in *European Conference on Parallel Processing*. Springer, 2016, pp. 737–749.
- [18] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renewable and Sustainable Energy Reviews*, vol. 31, pp. 622–638, 2014.
- [19] V. Hanumaiah and S. Vrudhula, "Energy-efficient operation of multicore processors by dvfs, task migration, and active cooling," *IEEE Transactions on Computers*, vol. 63, no. 2, pp. 349–360, 2014.
- [20] G. Falcao, L. A. Alexandre, J. Marques, X. Frazão, and J. Maria, "On the evaluation of energy-efficient deep learning using stacked autoencoders on mobile gpus," in *Parallel, Distributed and Network-based Processing (PDP), 2017 25th Euromicro International Conference on*. IEEE, 2017, pp. 270–273.
- [21] E. Azarkhish, D. Rossi, I. Loi, and L. Benini, "Neurostream: Scalable and energy efficient deep learning with smart memory cubes," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 2, pp. 420–434, 2018.
- [22] J. G. Koomey, "Worldwide electricity used in data centers," *Environmental research letters*, vol. 3, no. 3, p. 034008, 2008.
- [23] B.-G. Chun, G. Iannaccone, G. Iannaccone, R. Katz, G. Lee, and L. Niccolini, "An energy case for hybrid datacenters," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 1, pp. 76–80, 2010.
- [24] M. Dabbagh, B. Hamdaoui, M. Guizani, and A. Rayes, "Energy-efficient resource allocation and provisioning framework for cloud data centers," *IEEE Transactions on Network and Service Management*, vol. 12, no. 3, pp. 377–391, 2015.
- [25] T. H. Nguyen, M. Di Francesco, and A. Yla-Jaaski, "Virtual machine consolidation with multiple usage prediction for energy-efficient cloud data centers," *IEEE Transactions on Services Computing*, 2017.
- [26] H.-W. Tseng, T.-T. Yang, K.-C. Yang, and P.-S. Chen, "An energy efficient vm management scheme with power-law characteristic in video streaming data centers," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 2, pp. 297–311, 2018.
- [27] J. Li, Z. Li, K. Ren, and X. Liu, "Towards optimal electric demand management for internet data centers," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 183–192, 2012.
- [28] M. Ghamkhari and H. Mohsenian-Rad, "Energy and performance management of green data centers: A profit maximization approach," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 1017–1025, 2013.
- [29] Y. Guo, Y. Gong, Y. Fang, P. P. Khargonekar, and X. Geng, "Energy and network aware workload management for sustainable data centers with thermal storage," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 8, pp. 2030–2042, 2014.
- [30] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and limitations of tapping into stored energy for datacenters," in *ACM SIGARCH Computer Architecture News*, vol. 39, no. 3. ACM, 2011, pp. 341–352.
- [31] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam, "Optimal power cost management using stored energy in data centers," in *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*. ACM, 2011, pp. 221–232.
- [32] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?" in *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, no. 1. ACM, 2012, pp. 187–198.
- [33] J. Richard and J. Komp, "Practical photovoltaics electricity from solar cells," 2001.
- [34] P. Gipe, "Wind power: renewable energy for home, farm, and business, chelsea green pub," 2005.
- [35] L. A. Barroso, J. Clidaras, and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 8, no. 3, pp. 1–154, 2013.
- [36] M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, and I. Stoica, "Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling," in *Proceedings of the 5th European conference on Computer systems*. ACM, 2010, pp. 265–278.
- [37] D. Feitelson, "Parallel workloads archive." [Online]. Available: <https://www.cs.huji.ac.il/labs/parallel/workload/>
- [38] D. G. Feitelson, D. Tsafir, and D. Krakov, "Experience with using the parallel workloads archive," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2967–2982, 2014.
- [39] X. Yang, Z. Zhou, S. Wallace, Z. Lan, W. Tang, S. Coghlan, and M. E. Papka, "Integrating dynamic pricing of electricity into energy aware scheduling for hpc systems," in *SC'13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE, 2013, pp. 1–11.
- [40] S. Wallace, X. Yang, V. Vishwanath, W. E. Allcock, S. Coghlan, M. E. Papka, and Z. Lan, "A data driven scheduling approach for power management on hpc systems," in *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2016, pp. 656–666.