# Energy-efficient Management of Data Centers using a Renewable-aware Scheduler


Xiaopu Peng
*Department of Mathematics & Computing*
*Lander University*
Greenwood, SC, USA
xzp0007@auburn.edu

Tathagata bhattacharya
*Department of Computer Science*
*Auburn University at Montgomery*
Montgomery, AL, USA
tzb0063@auburn.edu

Jianzhou Mao
*Computer Science & Software Engineering*
*Auburn University*
Auburn, AL, USA
jzm0122@auburn.edu

Ting Cao
*Department of Computer Science*
*Truman State University*
Kirksville, MO, USA
tzc0028@auburn.edu

Chao Jiang
*Computer Science & Software Engineering*
*Auburn University*
Auburn, AL, USA
czj0043@auburn.edu

Xiao Qin
*Computer Science & Software Engineering*
*Auburn University*
Auburn, AL, USA
xqin@auburn.edu



2022 IEEE International Conference on Networking, Architecture and Storage (NAS) | 978-1-6654-5408-7/22/$31.00 ©2022 IEEE | DOI: 10.1109/NAS55553.2022.9925479



***Abstract*—Leveraging on-site renewable sources like solar and wind provides ample opportunities on developing environmental friendly and energy-efficient data centers. Evidence shows that renewable-aware job schedulers conserve energy by adjusting the arrangement of non-urgent workload according to renewable energy states. We propose an energy management system with a renewable-aware scheduler called REDUX3, which offers a smart way of managing the energy supply of data centers powered by the grid and renewable energy. Due to the intermittent nature of renewable energy resources, REDUX3 judiciously back-fills workload when renewable energy is sufficient, and defer workload to the next time slot if renewable energy is at outage state. As an integrated and smarter update from our previous work [1], [2] and [3], REDUX3 also 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 fluctuating level. Compared with the existing strategies, REDUX3 demonstrates a prominent capacity of boosting renewable energy utilization.**



***Index Terms*—Renewable energy, Data centers job scheduling, Renewable-aware scheduling, Resource management, Data center energy efficiency**


## I. Introduction

Energy efficiency is one of the key design principles of modern data centers. Inspired by the observations of renewable energy utilization and distinct pattern of job power profiles, we hypothesize that a renewable-aware scheduling mechanism running prior to a dynamic power manager has a potential to enhance the energy efficiency of data centers.

In our prior project [1], [2] and [3], we have demonstrated that the REDUX system save electricity bills by judiciously managing both inexhaustible on-site renewable energy and inexpensive stored energy from power grid for a typical data center. The overarching goal of REDUX is to curtail energy cost while maintaining high performance in data centers to address two challenging issues: (1) the intermittent feature of renewable energy and (2) dynamically changing grid electricity price. REDUX is responsible for orchestrating daily computing workloads in data centers while cutting back energy costs through improved energy resource utilization.

We anchor this study in a renewable-energy-aware job scheduler to boost the energy-efficiency performance of REDUX. We refer to a system powered by the smart job scheduler as *REDUX3*, in which workloads are scheduled and dispatched to computing nodes in a data center. The rationale behind REDUX3 is sparked on a key observation of jobs from the perspective of users - each job has a distinct power consumption profile [4]. It is feasible to schedule jobs by considering expected power profiles accompanied by the status of renewable energy (available or outage) at the beginning of each time period. Our renewable-aware job scheduling mechanism enables REDUX3 to either (1) dispatch hefty load to the computing nodes with available cheap and green renewable energy or (2) defer non-urgent jobs when grid energy price is high amid insufficient energy stored in batteries. Compared with the legacy REDUX system, REDUX3 is adept at scheduling jobs in a holistic way to optimize energy efficiency of data centers fueled by renewable energy.

There have been practical and promising studies into the topic of job scheduling customized for energy-efficient and parallel computing systems [4]. The prior evidence both demonstrates praiseworthy power-aware scheduling policies, which help in significantly shaving electricity cost by the virtue of window-based scheduling coupled with dynamic electricity pricing. Rather than allocating jobs one by one from a wait queue, these scheduling policies dispatch a ”window” of jobs at a time. Jobs placed into the window are chosen to maintain fairness through job priorities. Complementing these scheduling idea on large-scale parallel computing jobs, this work is anchored to a renewable-energy-aware scheduling design preparing for the next-step energy management in data centers. By integrating this scheduling mechanism into the REDUX3 system, we expect that REDUX3 has an edge over its predecessor REDUX system [1], [2] and [3] with respective to


978-1-6654-5408-7/22/$31.00 ©2022 IEEE
Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

saving energy cost. Our novel mechanism not only schedules and dispatches jobs on the basis of job-power profiles and renewable-energy status, but also triggers negligible impacts on system utilization and fairness. Preferentially, the proposed mechanism dispatches intensive workload amid fluent and stable renewable energy, hold non-urgent jobs for the next time slot during the outage of renewable energy.

In this study, we propose a workload-energy management system - *REDUX3* - to orchestrate renewable energy integrated with a power grid and a distributed UPS system. REDUX3 first examines jobs according to their power profiles with a help of renewable-aware scheduler to lay out a foundation for the subsequent executions in energy management; Then adept at smoothing power supplies with the supply levels of renewable energy resources; Utilize per-server distributed batteries as an economical and secure battery backup; And finally, makes judicious decisions on UPS charging and discharging with vital information on renewable-energy supply levels and time-dependent grid power price.

This paper makes three main contributions:

- *Unified Framework Design.* We design REDUX3 - a unified framework incorporating all the aspects of energy sources and workload, which takes full advantages of pre-execution renewable-aware workload scheduling, as a deep extension of our previous systems [1] [2] [3].
- *Renewable-aware Job Scheduling* We devise the renewable-aware job scheduling with a dynamic scaling scheduling window as a vital part of the workload management. The scheduling policy leverages low-cost renewable energy in accordance with power profiles, priorities, and renewable energy states, and back-fill workload to idle nodes or defers the executions of non-urgent workload to further cut back the overall electricity cost.
- *Optimization.* We evaluate REDUX3 using the real-world workload traces and green energy data. We demonstrate that REDUX paves a way toward constructing modern data centers that are economically and environmentally friendly.

## II. Background and Related Work

The topics related to energy efficiency of data centers is broadly and diversely discussed. Table I summarizes the major differences between our proposed REDUX3 and the existing power management techniques for 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 [10]. Chun *et al.* investigated a hybrid datacenter architecture mixing low power systems and high performance ones [11]. Dabbagh *et al.* developed an integrated energy-aware resource provisioning framework to optimize energy efficiency of data centers [12]. Virtualization techniques have been applied to reduce the power consumption by consolidating multiple virtual machines into one physical server [13] [14] [15]. For example, Tseng *et al.* designed a virtual machine management strategy to predict resource utilization for upcoming videos, thereby turning off idle servers to conseve energy [16]. Another thread of techniques that minimize the electricity cost of data centers is exploiting the temporal and regional diversity of electricity prices (see, for example, [17]).

**Renewable Energy in Data Centers.** There are pressing demands from cloud service providers for cost-effective renewable energy to power data centers [18]. An increasing number of prior studies focused on making green data centers through the integration of renewable power sources [6]. 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 [19]. To integrate renewable energy, Guo *et al.* incorporated geographical load balancing, opportunistic workload scheduling, and thermal storage management in data centers [20].

**Distributed UPS Systems.** Uninterruptable power supplies (i.e., UPS devices) offer efficient peak power shaving in data centers [5]. Urgaonkar *et al.* proposed the usage of UPS units to offer ample opportunities of cost reduction in data centers [21], by first prove the feasibility of leveraging energy storage device (e.g., lead-acid batteries) in data centers to reduce peak power cost. Govindan *et al.* designed a UPS-based energy buffer to reduce energy costs in data centers [22].

**Renewable-Aware Scheduling** Curtailing energy consumption has become a necessary objective of modern job schedulers running on large-scale computing system housed in data centers. For example, Zhou *et al.* [23] proposed a power-aware job scheduling approach incorporating variable energy prices and job power profiles. The scheduler relies on a power budget for energy savings, thereby degrading system utilization slightly during on-peak electricity price period. Yang *et al.* [4] intended to provide a generic job-power-aware scheduling mechanism to minimize electricity bills without impacting system utilization during both on-peak and off-peak pricing periods. Wallace *et al.* [7] proposed a data-driven scheduling approach to controlling power consumption of an entire system under any user defined budget. Though renewable-energy-aware schedulers receive attention in the research community, there is the lack of a unified high-level framework design. For example, Goiri *et al.* [9] proposed GreenSlot, a scheduler for parallel batch jobs in a data center powered by a photovoltaic solar array where grid power is only treated as a backup energy source. Grange *et al.* [8] again implemented the "GreenVarPrice" version of the GreenSlot scheduler [19], which takes into account the variation of grid energy price.

## III. Framework Design

### A. Overview

The *REDUX3* system is centered around two key components, namely, (1) renewable-aware workload scheduler and (2) REDUX workload-power manager [1], [2] and [3].

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

TABLE I
COMPARISONS BETWEEN REDUX3 AND THE EXISTING ENERGY MANAGEMENT TECHNIQUES

| | Renewable-aware Workload Scheduling | Workload Scheduling | Renewable Energy | Grid Price Fluctuations | UPS Enabled | Unified Framework |
|---|---|---|---|---|---|---|
| ReUPS [5] | × | × | √ | × | √ | √ |
| GreenSwitch [6] | × | × | √ | √ | √ | × |
| CQSim [4] [7] | × | √ | × | × | × | × |
| Green IT [8] | √ | √ | × | √ | × | × |
| GreenSlot [9] | √ | √ | √ | √ | × | × |
| **REDUX3** | √ | √ | √ | √ | √ | √ |

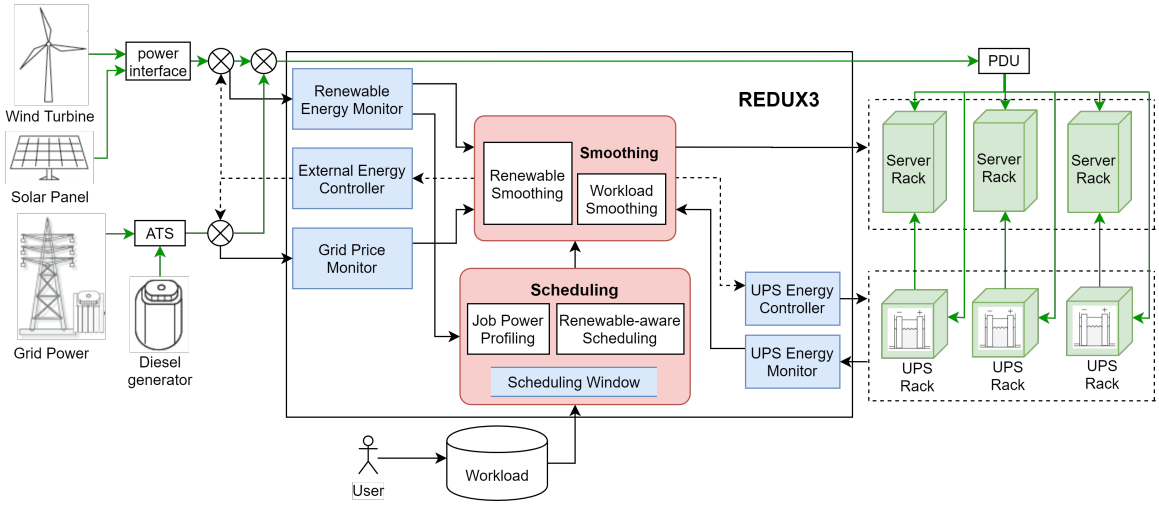



Fig. 1. The framework of the *REDUX3* system.

REDUX3 is a cross-layer power optimization system, which resides in both the front end as a workload scheduler, and the back end as a data-center energy resource manager. REDUX3 is adept at tuning power and managing workload in data centers. The scheduler embedded in REDUX3 is running in the batching mode, which is popular in the realm of big-data applications running in data centers.

Fig. 1 plots an overview of the REDUX3 framework design. The system inside the large box persistently communicates with the five external entities, namely, (1) renewable energy supplies, (2) power grid with diesel generator as a backup, (3) distributed UPS units, (4) server racks, and (5) workload submitted by users. Renewable energy (e.g., wind and solar), grid power, and diesel generators are seamlessly integrated into the system. On the left-hand side of Fig. 1, there are a green colored power circuit, a black colored monitoring data flows, and a dashed arrows of control signal flows. Through these flows, the monitoring data will feed both to the REDUX3 and the renewable aware workload scheduler, and the control signal flow will send to all external entities by the REDUX3 system except user workload. On the right-hand side of the framework, a distributed UPS system is attached to the server racks to improve efficiency, scalability, and reliability over a traditional centralized counterpart. The UPS system aims to buffer energy under light workload or power at a cheap price, and to consume stored energy in case of power spikes. Inspired by a prior work [24], we derive the required battery capacity based on the amount of power demand from smoothing processes and the corresponding energy stored in a battery for a given daily power profile. At the bottom of Fig. 1, jobs are submitted by users and queued in a waiting list before being fed into the renewable-aware scheduling system.

All the modules of the REDUX3 system falls into two categories. The first category includes control and monitor modules (see all light-blue colored modules in Fig. 1); the second category groups the modules forming the two core subsystems (see the modules in light-red colored in Fig. 1). There are three monitoring modules and two controller modules for the external entities: (1) the renewable energy monitor, the grid price monitor and UPS monitor which collect trace data and update the status of all the external entities; and (2) the external energy controller and the UPS energy controller which manages the system according to received control signals from REDUX3. Seamlessly integrating the three monitoring modules, the system keeps track of energy supply trace data from on-site renewable devices or off-site grid networks, and energy storage amount in UPS racks. On the other hand, the monitoring modules facilitate the the two core subsystems with updated workload and power status. The controllers deliver control signals to the external entities, thereby tuning

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

renewable and grid energy supply and determine recharging or discharging of UPS devices.

### B. The Renewable-aware Workload Scheduler

In the REDUX3 system, the renewable-aware workload scheduler keeps track of renewable energy status and workload power profile data in each time slot. The following three key components are seamlessly integrated in the scheduler.

*The scheduling window module.* This module is responsible for filling up a windowed list by picking jobs from the job waiting queue submitted during the time slot, thereby striving to maintain fairness among the jobs. In addition, the scheduling window module updates the states of the job list after the smoothing mechanism.

*The renewable-aware scheduling module.* The scheduling module first retrieves state of renewable energy from the monitor, and then judiciously makes final scheduling decides according to the renewable-aware scheduling policy. Basically, the scheduling module (1) back-fill workloads of next time slot from executing jobs if existed when renewable energy is abundant, or (2) defer non-urgent jobs to the next time slot at outage time slot of renewable energy.

*Job power profiling module.* The profiling module, managing power profile information, assigns a power profile to each job. After assigning profile information, this module records the total energy demand from the jobs residing in the scheduled window.

## IV. Methodology

We advocate for a modularized design approach to constructing REDUX3, in which a high-level coordinating algorithm orchestrates the monitoring, scheduling and smoothing modules (see Section IV-A). Next, Section IV-B sheds light on the design issues of the renewable-aware job scheduler.

### A. The High-Level Control Algorithm

Algorithm 1 depicts the pseudo-code of the high-level algorithm in REDUX3. The inputs include (1) energy supply values, namely, wind energy $E^W(t)$, solar energy $E^S(t)$, and energy storage $E^{UPS}(t)$ in UPS device, (2) the energy demand from workload $W(t)$, and (3) grid energy price $P^G(t)$ as variable, wind energy price $P^W$, solar energy price $P^S$, and UPS energy price $P^{UPS}$ as constants. And we denote $E^G(t)$ as the energy draw from grid network after $E^{RE}(t)$ and $E^{UPS}(t)$ are decided during the time slot.

The high-level controller kicks off the power management task with determine the renewable energy state $S_{RE}$ (see $S_{RE} = get_renewable_energy_state()$; in line 3), which is categorized as, namely, *stable* (i.e., $STA$), *fluctuate* (i.e., $FLU$), and *outage* (i.e., $OTA$).

- *The Outage Case*. If the renewable supply is at the minimum level, the renewable state will be defined as $OTA$.
- *The Stable Case.* If the renewable-energy supply exceeds the maximum level, the renewable state will be referred to as $STA$. For the other cases, we compare the the renewable-energy supply with its previous exponential windowed average.
- *The Fluctuate Case.* 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 $FLU$.

**Algorithm 1** The high-level control algorithm of REDUX3.

**Input:**
Renewable energy $E^W(t)$ and $E^S(t)$
UPS energy storage $E^{UPS}(t)$
Submitted workload $W(t)$
Price of all energy sources $P^W, P^S, P^{UPS}$ and $P^G(t)$
**Output:**
Energy Cost $C(t)$ of time slot $t$

1: $t = 0$
2: **while** $t <= T$ **do**
3: $S_{RE} = get_renewable_energy_state()$;
4: $renewable - aware_job_scheduling()$;
5: **if** $S_{RE} = FLU$ **then**
6: $ren_supply_smooth()$;
7: **end if**
8: $ups_storage_update()$;
9: $E^G(t) = W(t) - E^{RE}(t) - E^{UPS}(t)$;
10: $C(t) = P^G(t) \cdot E^G(t) + P^{ren}(t) \cdot E^{ren}(t) + P^{UPS} \cdot E^{UPS}(t)$;
11: $t++$;
12: **end while**
13: **return** REDUX cost

Taking both when $S_{RE} = STA$ and $S_{RE} = FLU$ as available states of renewable energy, the renewable-aware job scheduling procedure (see $renewable - aware_job_scheduling()$; in Line 4) then take place to apply the renewable-aware algorithm on a windowed priority queue, followed by transforming the job workload into energy demands derived from job power profiles. After determining the state of the current renewable supply, REDUX3 further manages energy supply in accordance to the renewable-energy state. More specifically, if the renewable energy is in the fluctuation state, the renewable-energy-supply smoothing policy is invoked in Line 6. The smoothing procedure is governed by the current renewable supply data, a stable supply level that is dynamically updated, and a UPS energy storage condition. The functionality of $ups_storage_control()$ procedure (see Line 8) is two-fold, namely, (1) making decisions on UPS discharging and charging and (2) updating the UPS status information. In responsible for protecting UPS units against frequent charging and discharging to maintain reliability of the distributed UPS system, this procedure stipulates a minimum UPS energy level, under which the UPS units should be protected. With $W(t)$ denotes workload demand after renewable-aware scheduling and in terms of energy amount, REDUX3 computes the energy supply obtained from the grid network (see Line 9). The function in Line 10

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

calculates the total energy cost occurred in time slot $t$.

### B. Renewable-aware Job Scheduler

In this subsection we brush up the description of renewable-aware job scheduling. Inspired by [4] and [7], our design embraces three key mechanisms, (1) a scheduling window, (2) a scheduling algorithm, and (3) a job power profiling module.

*1) Scheduling Window:* Maintaining fairness among jobs with various priorities and computing recourse requirements is a well-received concern amid the development of job schedulers. Demonstrated as a gray box in Fig. 2, we advocate for a window-based scheduling mechanism to facilitate the renewable-energy-aware scheduling algorithm without breaking fairness. Subsequently, the job is fed into the scheduling window if resources are permitted. Any job that has already entered the scheduling window will never be kicked back to the waiting queue until being executed and marked as "finished".

Another key design of the scheduling window lies in the renewable-aware dynamic scaling mechanism. Since any submitted job requests a number of processors during its execution, we derive and measure the size of the scheduling window as the total number of processors to be allocated by the data center. More specifically, renewable-aware and dynamic scaling is performed in the following fashion: (1) enlarging the scheduling window if renewable energy is sufficient or stable, or (2) shrinking the scheduling window when renewable energy is in the outage state. Guided by our empirical studies, we set this dynamic scaling ratio to $0.2$ for the enlarging and shrinking process. The upper part of Fig. 2 demonstrates that the renewable energy state from renewable energy monitor dynamically scales the scheduling-window size.

*2) Scheduling Algorithm:* Our renewable-aware scheduling algorithm selects jobs from the scheduling window to optimize utilization of renewable energy. At the beginning of each time slot, the algorithm fills up the dynamically scaled scheduling window (either enlarged or reduced) with jobs from the job waiting queue in a FCFS order. Next, either back-filling or deferring of workload is determined according to the following two cases:

- **Case 1.** If the renewable energy is abundant accompanied by an enlarged scheduling window, extra workload form existing jobs in the next time slot are back-filled into the scheduling window and dispatched to idle processors (see the LHS of Fig. 2). We avoid back-filling any new job from the job waiting queue because it is imperative to maintain an on-line processing status for the scheduling step, which is subject to a future study.
- **Case 2.** With a narrow-scaled scheduling window coupled with the outage state of renewable energy, the scheduling algorithm switches to deferring low priority jobs to the next time slot (see the RHS of Fig. 2). Depending on the aggressiveness of the energy-saving policy, the REDUX3 system may either prevent all low priority jobs from entering the scheduling window in the first step or attempt to fill up the scheduling window with all high priority jobs first, then low-priority jobs.

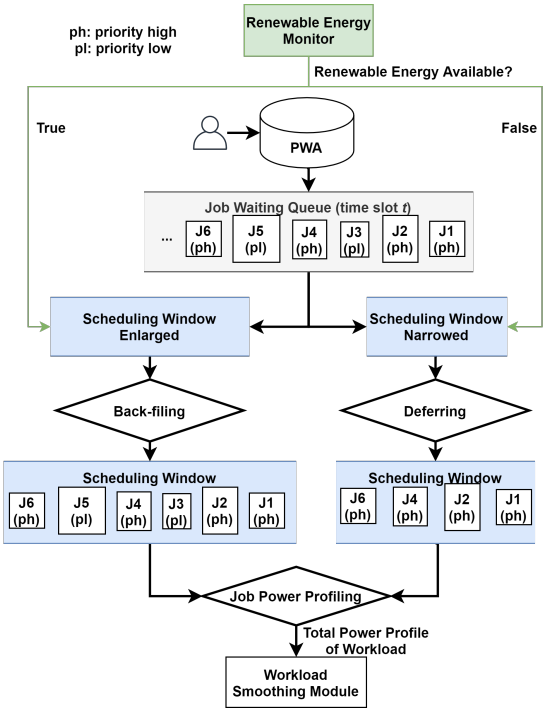



Fig. 2. Overview of Renewable-aware Job Scheduling

Last but not least, as depicted in Fig. 2, each box with job number inside the job waiting queue represents one hour of workload of that job. This key design is driven by three reasons. (1) Any job's execution time spans over one time slot can be treated as multiple one-time-slot workloads to be fed into the job waiting queue at the beginning of each time slot. (2) Parallel computing among multiple nodes provides possibility of allocating a job's workload that belongs to the next time slot to the current time slot. (3) This design can eschew overflow of any unfinished job when an enlarged scheduling window of the current time slot is scaled to a narrowed scheduling window in the next slot.

*3) Job Power Profiling:* To make an intelligent job scheduling decision, we ought to profile and model power consumption requirements of submitted jobs. Estimating power consumption of jobs running on clusters in a data center is nontrivial in that the running time and energy usage of a job are reliant on the allocated resources (e.g., the number of computing nodes and main memory capacity). For each job to be submitted to a data center, we profile the energy efficiency and performance of the job on a cluster. Such profiling data allow us to build a model to project the execution time and power consumption of the job given any assigned resources. Upon a job submission, such profiling and modeling data should be submitted along with jobs. In doing so, the REDUX3 system is capable of speculating each job's execution time and energy demands to make wise scheduling decisions. In our experiments, the profiling data are furnished in job traces

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

containing resource requirements (e.g., CPU time and energy demand).

Motivated by the findings in [4] and [7], we implement a job power-profile function leveraging a categorized normal distribution to estimate power consumption of a job from its power profile. More specifically, the power demand of the job is initially derived by the job's requested computing resources, followed by the categorized normal distribution method to finalized the job's power requirement value. The detail settings of the job power-profile function can be found in Section V-B2.

## V. Evaluation

We systematically evaluate the design of REDUX3 by developing a simulated data center processing real-world workload.

### A. ReCQSim: Renewable-aware Scheduling Simulation

We implement our renewable-aware scheduling method in *ReCQSim*, which follows the scheduling framework of the existing trace-based event driven scheduling simulator named *CQSim* [4] [25]. ReCQSim is a renewable-aware workload scheduling simulator that takes renewable energy supply states into consideration during a scheduling process.

ReCQSim detects the state of renewable energy with help of the modules from REDUX3 at the beginning of each time slot. With the mechanism elaborated in Section IV-B, the renewable-aware scheduler makes a decision on enlarging or shrinking the scheduling window based on the status of renewable energy supply. Next, the job power profile module computes workload of all the jobs in the scheduling window; Finally, the workloads are measured in terms of $kWh$ with each job's power profile. The workload data coupled with all the job information are then fed into REDUX3. More specifically, REDUX3 allocates the sorted jobs in the scheduling window into the available computing nodes of the data center. Then, REDUX3 acts as a management system orchestrating grid power, UPS units and renewable energy for computing servers according to workload requirements in data centers. After each simulation round is accomplished, the scheduling window is updated with the current status of each job.

### B. Experiment Configurations

We elaborate the experimental configurations from the perspectives of (1) job traces (Section V-B1), (2) job power profile (Section V-B1), and (3) energy supply (see Section V-B3).

TABLE II
Comparisons between REDUX and the existing power management techniques for data centers.

| Log Name | From | To | Mo | CPUs | Jobs |
|---|---|---|---|---|---|
| UniLu Gaia | May14 | Aug14 | 3 | 2,004 | 51,897 |
| MetaCentrum2 [26] | Jan13 | Apr15 | 28 | 8,412 | 5,731,100 |
| CIEMAT Euler [27] | Nov08 | Dec17 | 110 | 1920 | 9263012 |
| KIT FH2 | Jun16 | Jan18 | 19 | 24048 | 114,355 |

*1) **Job Trace and Size of Data Center**:* We replay real-world workload traces archived in a well established online repository [28]. To gauge the performance of the renewable-aware scheduling algorithms, we use five key parameters in each trace file, namely, job ID, submit time, requested time, requested number of processors, and the priority level number. In order to demonstrate the effectiveness of our design on a broad range of workload settings and traces, we exterminate the following log files from *PWA* [28]: *UniLu Gaia*, *MetaCentrum2* [26], *CIEMAT Euler* [27], and *KIT FH2*. We choose to replay these traces, because the traces encompass at least three-month workload data, which are aligned with the duration of the power supply data set. We also expand the trace data to resemble the scale of modern data centers simulated in REDUX3.

*2) Job Power Profile:* It's nontrivial to assess the system power profiles of job traces for parallel systems. Because the *PWA* trace excludes the power information connected to the tested workloads, we implement a job power-profiling module (see also Section IV-B3) to render power profiles for jobs within time slot $t$. The input of the power-profiling module includes the number of requested processors as well as expected running time during slot $t$ of each job. Following [4] and [7], we set the power demand of a computing node to a value anywhere between 20 to 60 Watts within a normal distribution. Next, the power profile of each job running in slot $t$ is multiplied by the power dissipation of all the allocated processors and the job running time. Finally, the profiling module accumulates the power demands of all the jobs running in slot $t$, thereby deriving power demand measured in $kWh$.

*3) Energy Supply:* We test the renewable-aware scheduler and the REDUX3 system using the electricity price of a real-world power grid in the New York state during a 3-month period from June 1 to August 31, 2016. We select this time period, because the solar and wind data were also collected and available for our simulation study. Fig. **??** reveals that The power grid price dynamically changes during the three-month time interval, with the maximum and minimum prices in the trace are \$1.82 and \$0.26 $/Kwh$, respectively.

We retrieve the solar supply of the New York state the National Renewable Energy Laboratory's (NREL) database. To align the time window with the that of the above electricity price, we focus on the solar trend from June 1 to August 31, 2008. The solar and wind power price levels are fixed at \$0.09/Kwh and \$0.07/Kwh, respectively.

Inspired by [24], we assert that an array of distributed UPS devices with 12V lithium iron phosphate (LFP) batteries are installed. Attaching to each server provisioned at 60Ah, the batteries per rack have a maximum capacity fitting within server-size constraints. Because a distributed UPS system has the distinctive advantage of scalability, this setting can be configured according to servers' maximum energy demands. The energy cost paid to recharge/discharge batteries is calculated as \$0.07/kWh.

### C. Results and Analysis

In this section we present experimental results to demonstrate the benefits of applying the renewable-aware scheduler

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

to various energy-efficiency management systems including our baseline REDUX design.

*1) Energy Cost Saving:* We first conduct the energy-cost-efficiency experiments driven by the representative workload trace $SDSC-SP2$. We compare the discrepancy of workload and accumulated energy cost between the renewable-aware REDUX3 and the baseline REDUX system. Fig. 3 depicts the workload difference between REDUX3 and REDUX systems during each time slot. A clear observation is that the $kWh$ value of workloads scheduled by the renewable-aware design are more fluctuate than the workloads from the baseline REDUX system. This pattern is expected because parallel workloads can be either back-filled into a scheduling window amid sufficient renewable energy or deferred to the next time slot.

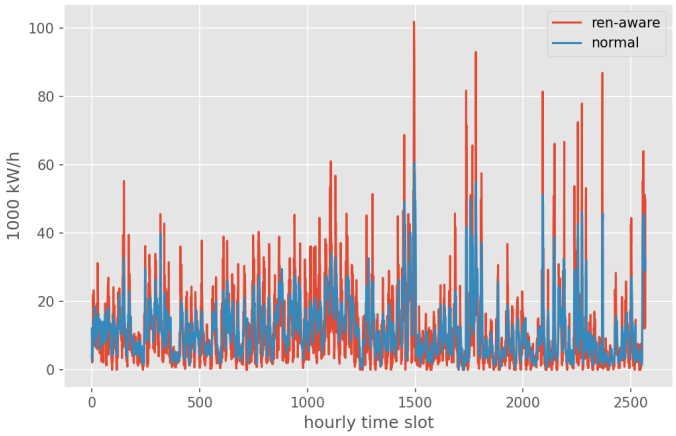



Fig. 3. The difference of workload in *1000 kWh* between renewable-aware and normal workload scheduler across time slots

Table III summarizes three sample differences when renewable energy is available and experiencing an outage. It is obvious that the values (in $1000kWh$) of workload from renewable-aware scheduling is substantially higher when renewable energy is available, and greatly lower when renewable energy is at the outage state.

TABLE III
WORKLOAD (IN $1000kWh$) COMPARISONS OF RENEWABLE-AWARE REDUX3 AND ORIGINAL REDUX SYSTEM IN REPRESENTATIVE TIME SLOTS

| Time Slot | 2370 | 1501 | 2092 | 1736 | 1495 | 2368 |
|---|---|---|---|---|---|---|
| Renewable Energy Available? | Yes | Yes | Yes | No | No | No |
| ren-aware | 5.51 | 2.51 | 9.15 | 81.69 | 96.09 | 86.87 |
| normal | 45.73 | 46.92 | 50.45 | 39.90 | 53.88 | 44.40 |
| difference | 44.55 | 44.42 | 41.3 | -41.79 | -42.21 | -42.48 |

Fig. 4 demonstrates the accumulated total energy cost (in million $) of the REDUX3 and REDUX systems. We observe that the energy-cost gap between the two systems is widening along with the timeline. Intuitively, the renewable-aware REDUX3 has an edge over REDUX in terms of energy efficiency, offering cost savings at the rate of almost 8% at the end of the timeline. It is evident that the renewable-aware scheduler optimize the energy efficiency of REDUX3 during each time slot by (1) providing higher workload when inexpensive renewable energy is available and (2) lowering workload when expensive grid or UPS energy is the only choice. The detailed value of accumulated energy cost and relative saving rate against the REDUX system at the end of each month can be found in Table IV.

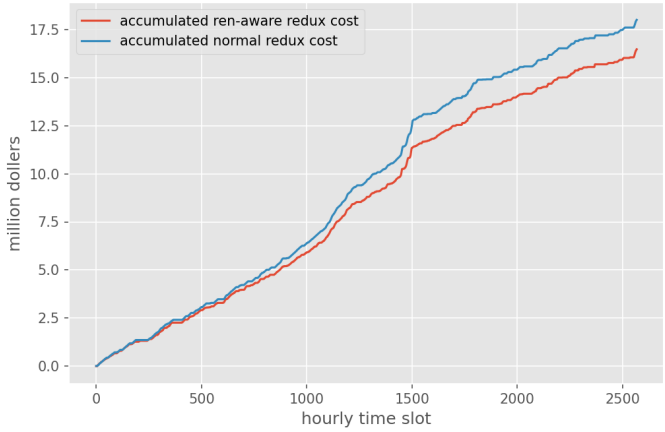



Fig. 4. The accumulated energy cost of renewable-aware REDUX3 and original REDUX system

| Time Slot | 1st | Cost Saving | 2nd | Cost Saving | 3rd | Cost Saving |
|---|---|---|---|---|---|---|
| Ren-aware | 4.82 | 0.36 | 12.57 | 1.33 | 16.23 | 1.74 |
| Original | 5.18 | - | 13.9 | - | 17.97 | - |

TABLE IV
CUMULATIVE ENERGY-COST (IN MILLION DOLLARS) COMPARISONS AND SAVINGS BETWEEN RENEWABLE-AWARE REDUX3 AND ORIGINAL REDUX SYSTEM AT THE END OF EACH MONTH

*2) Impacts on Energy Cost:* In this subsection we explore each of the following impacts of the renewable-aware REDUX3 design on energy cost.

**Impacts of Workload Conditions** To assess the effect of workload traces on energy efficiency, we conduct the same energy-cost experiments driven by the three other traces extracted from the Parallel Workload Archive, namely $SDSCBLUE$, $SandiaRoss$, $UniLuGaia$. We pick these traces, because the traces are comprised of user run-time estimates data - an input of the power-profile module. Table V lists key features of these workload data. Fig. 5 demonstrates that the renewable-aware REDUX3 system always prevails the original REDUX system on accumulated energy cost at time slots 1000, 2000, and 2600. The results indicate that REDUX3 boosts the energy efficiency across the tested workload traces. In particular, REDUX3 cuts back the energy cost by an average around 25% under the $UniLuGaia$ workload.

## VI. CONCLUSIONS

In this paper, we presented a renewable-aware scheduler based data center energy efficiency system called *REDUX3* with on-site renewable sources coupled with a distributed UPS

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.

| Data Trace | Months | CPU | Job executed | Utilization Rate % |
|---|---|---|---|---|
| SDSC SP2 | 24 | 128 | 59,715 | 83.4 |
| SDSC Blue | 32 | 1152 | 243,306 | 76.7 |
| Sandia Ross | 37 | 1524 | 57,882 | 49.9 |
| UniLu Gaia | 3 | 2004 | 51,987 | 84 |

TABLE V
GENERAL FEATURES OF WORKLOAD TRACES

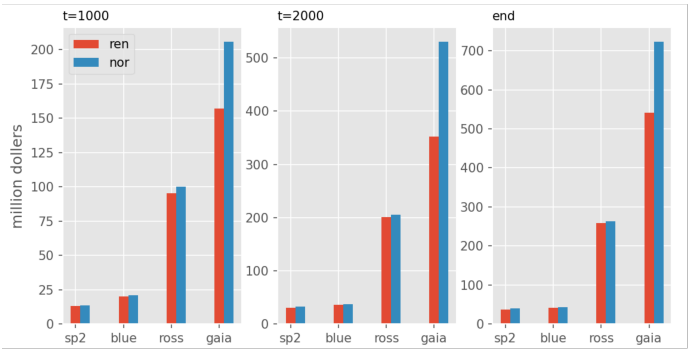



Fig. 5. Comparison of Different Workload Traces

system. The renewable-aware scheduler judiciously cope with the intermittent nature of renewable energy integrated with power grids. One promising feature of REDUX3 is to make data centers work in full capacity when inexpensive renewable energy is abundant and scale down workload when grid or UPS energy is the only choice. Compared with the prior solutions, REDUX3 conserves total energy cost while boosting renewable-energy utilization. We show that REDUX3 is an underpinning technique to build energy-efficient data centers.

## REFERENCES


[1] X. Peng, C. Kauten, C. Zhang, J. Mao, T. Heckwolf, T. K. Al Tekreeti, and X. Qin, "Redux: Managing renewable energy in data centers using distributed ups systems," in *2018 IEEE International Conference on Smart Cloud (SmartCloud)*. IEEE, 2018, pp. 46–53.

[2] X. Peng and X. Qin, "Energy efficient data centers powered by on-site renewable energy and ups devices," in *2020 11th International Green and Sustainable Computing Workshops (IGSC)*. IEEE, 2020, pp. 1–3.

[3] X. Peng, T. Bhattacharya, T. Cao, J. Mao, T. Tekreeti, and X. Qin, "Exploiting renewable energy and ups systems to reduce power consumption in data centers," *Big Data Research*, vol. 27, p. 100306, 2022.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] L. Grange, G. Da Costa, and P. Stolf, "Green it scheduling for data center powered with renewable energy," *Future Generation Computer Systems*, vol. 86, pp. 99–120, 2018.

[9] Í. Goiri, M. E. Haque, K. Le, R. Beauchea, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Matching renewable energy supply and demand in green datacenters," *Ad Hoc Networks*, vol. 25, pp. 520–534, 2015.

[10] J. G. Koomey, "Worldwide electricity used in data centers," *Environmental research letters*, vol. 3, no. 3, p. 034008, 2008.

[11] 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.

[12] 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.

[13] 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.

[14] J. Mao, X. Peng, T. Cao, T. Bhattacharya, and X. Qin, "A frequency-aware management strategy for virtual machines in dvfs-enabled clouds," *Sustainable Computing: Informatics and Systems*, vol. 33, p. 100643, 2022.

[15] T. Bhattacharya, X. Peng, J. Mao, C. Zhang, T. Takreeti, Y. Wang, T. Cao, and X. Qin, "Performance modeling for i/o-intensive applications on virtual machines," *Concurrency and Computation: Practice and Experience*, vol. 34, no. 10, p. e6823, 2022.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] Z. Zhou, Z. Lan, W. Tang, and N. Desai, "Reducing energy costs for ibm blue gene/p via power-aware job scheduling," in *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer, 2013, pp. 96–115.

[24] 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.

[25] Y. Fan, Z. Lan, P. Rich, W. E. Allcock, M. E. Papka, B. Austin, and D. Paul, "Scheduling beyond cpus for hpc," in *Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing*, 2019, pp. 97–108.

[26] D. Klusáček, Š. Tóth, and G. Podolníková, "Real-life experience with major reconfiguration of job scheduling system," in *Job Scheduling Strategies for Parallel Processing*. Springer, 2015, pp. 83–101.

[27] M. Rodríguez-Pascual, A. J. Rubio-Montero, J. A. Moríñigo, and R. Mayo-García, "Execution data logs of a supercomputer workload over its extended lifetime," *Data in brief*, vol. 28, p. 105006, 2020.

[28] D. Feitelson, "Parallel workloads archive." [Online]. Available: https://www.cs.huji.ac.il/labs/parallel/workload/


Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 23:18:33 UTC from IEEE Xplore. Restrictions apply.