Managing Battery Aging for High Energy Availability in Green Datacenters

Longjun Liu, *Member, IEEE*, Hongbin Sun, *Member, IEEE*, Chao Li, *Member, IEEE*, Tao Li, Jingmin Xin, *Senior Member, IEEE*, and Nanning Zheng, *Fellow, IEEE*

Abstract—Energy storage devices (ESD), such as UPS batteries, have been repurposed in datacenter as a promising tuning knob for peak power shaving and power cost reducing. However, batteries progressively aging due to irregular usage patterns, which result in less effective capacity and even pose serious threat to server availability. Nevertheless, prior proposals largely ignore the aging issues of battery which may lead to low energy availability for datacenter servers. To fill this critical void, we thoroughly investigate battery aging on a heavily instrumented prototype system over an observation period of ten months. We propose Battery Anti-Aging Treatment Plus (*BAAT-P*), a novel power delivery architecture included aging management algorithms from the perspective of computing system to hide, reduce, mitigate and plan the battery aging effects for high energy availability in datacenter. Our techniques exploit diverse battery aging mechanisms and dynamic aging management algorithms to provide system-level availability guarantee for datacenter. We evaluate the *BAAT-P* design with a real prototype. Compared with a battery powered datacenter without aging management policies, the results show that *BAAT-P* can extend battery lifetime by 72 percent, reduce battery cost by 33 percent and effectively improve energy availability for datacenter servers while maintaining workload performance for the performance critical workloads.

Index Terms—Green datacenters, battery aging management, power management, availability

1 Introduction

We can see that the growing adoption of massive distributed batteries at the server/rack level could fundamentally transform the way we manage datacenter power and energy [1], [2], [3]. For example, distributed batteries are important energy buffers in several topical studies that exploit renewable energy resources for capping datacenter carbon footprint [5]. They are also the key enablers that allow datacenters to smooth out load power peaks to greatly reduce total cost [4], [6], [7]. In general, these critical energy storage devices can help today's power-constrained datacenter servers maintain a continuous balance between power supply and demand, thereby protecting sensitive loads from possible power disturbances. As a result, the world-wide installed renewable energy battery is projected to increase by 22 GW in the next decade [8], which almost matches the global server power demand today [9].

Battery aging issue starts to become the major challenges for energy availability of datacenter with the dramatic increase of their utilization frequency in datacenter, especially for the renewable energy powered datacenter in which batteries has become necessary energy storage devices. Most of the prior power management works focus on managing server power demand [5], [6], [7], [10], [11], [12], [13], whereas the battery aging effect in datacenters has been largely ignored. However, batteries progressively and inevitably age. Without smart control, the aged batteries often lose their effective energy storage capacity and fail to power servers in emergency events.

Previous battery aging and wear-out research in electrochemistry research community focus on improving battery's internal design [14], [15], [16], [17], [18], [19]. When the manufactured batteries are massively deployed in datacenters, new aging management strategies at computer architecture and system levels are highly desired for reducing battery aging and improving datacenter energy availability.

This paper exploit diverse battery aging mechanisms and aging related server availability issues from the perspective of computing system to mitigate the threat of battery aging for green datacenters. We investigate emerging green datacenters that rely on energy storage devices (batteries) to jointly store green energy and shave load power spikes. In contrast to conventional batteries mainly deployed for handling power outages (rarely used), batteries in green datacenters often incur cyclic usage, i.e., they are charged and discharged in a much more frequent and irregular manner. In addition, distributed battery units also incur significant aging variations due to imperfect manufacturing process and different server load power behaviors. To manage battery aging in such a dynamic, complex environment, it is critical to identify key factors that correlate with battery

- *L. Liu, H. Sun, J. Xin and N. Zheng are with the School of Electrical and Information Engineering, Xi'an Jiaotong University, Xi'an 710049, China. E-mail: {liulongjun, hsun, jxin, nnzheng}@xjtu.edu.cn.*
- *C. Li is with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China. E-mail: lichao@cs.sjtu.edu.cn.*
- *T. Li is with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611. E-mail: taoli@ece.ufl.edu.*

Manuscript received 19 Jan. 2017; revised 16 May 2017; accepted 2 June 2017. Date of publication 7 June 2017; date of current version 15 Nov. 2017. (Corresponding author: Hongbin Sun.)
Recommended for acceptance by M. Kandemir.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TPDS.2017.2712778

aging and develop a holistic management strategy that can capture various anti-aging opportunities.

To this end, we conduct a thorough battery aging analysis based on our one-year deployment experience on a state-of-the-art green datacenter prototype. Our scaled-down system combines a Xeon-based server cluster, solar panels, a professionally assembled battery array, and a software management console built from scratch. For over ten months, we continuously monitor detailed battery usage events, investigate various factors that may induce battery aging and server availability issues. We propose Battery Anti-Aging Treatment Plus (*BAAT-P*), a new power delivery architecture and aging management framework from the perspective of computing system to avoid the degradation of energy availability caused by battery aging. The novelty of *BAAT-P* includes three aspects: 1) it can sense the synergistic effect of multiple aging factors to assess battery aging; 2) It leverages hybrid ESD (Energy Storage Device) technology, comprised by battery and super-capacitor, to efficiently mitigate battery aging processes in green datacenters; 3) At the computing architecture/system level, it is able to further leverage a multifaceted algorithms to mitigate battery aging and improve the energy availability of datacenter while maintaining workload performance.

BAAT-P further offers four key benefits to manage battery aging for high energy availability of green datacenter. First, it hides battery aging variation within datacenter. When datacenter adds new workloads or consolidates existing workloads, *BAAT-P* can intelligently identify battery units that wear out much faster than others. It can balance the aging effect across the battery units through an aging-driven load management policy. Second, it can slow down the battery aging process in cases that a battery is prone to wear out, e.g., when it is frequently used under low state of charge (SoC). *BAAT-P* leverages workload migration and power capping mechanisms at critical points to avoid aggressively discharging the batteries. Third, it further leverages emerging hybrid energy storage technology to efficiently manage hybrid energy flow for mitigating battery aging rate while maintaining workload performance for critical workloads. Forth, *BAAT-P* also can gainfully plan the battery aging speed to synchronize battery life with the end-of-life of IT systems. It proactively predicts battery lifetime and trades off unnecessary battery service life.

The challenges for *BAAT-P* to dynamically and intelligently manage battery aging primarily stem from the following four aspects: 1) how to assess various battery aging degrees; 2) how to select a battery whose aging should be hidden; 3) how to balance battery aging and workload performance; 4) how to effectively control and manage hybrid ESD energy flow to mitigate battery aging; 5) how to fully unleash the performance potential of green datacenter via planned battery aging. This paper can effectively solve these obstacles and improves the energy for datacenters.

2 Related Work

To the best of our knowledge, this is the first extensive analysis of battery aging and availability studies for datacenters. The relevant prior works are summarized as bellow.

1) *Studies of battery aging and battery lifetime:* There have been many studies on battery aging and lifetime, which can be categorized into: (1) Battery system model and lifetime evaluation [14], [15], [20], [21]. The main focus is to leverage mathematical formula to model battery charging and discharging behavior. (2) Battery aging mechanism studies [16], [17]. Among those, [16] presents an overview of battery aging mechanisms and lifetime evaluation for lead-acid battery. [17] focus on the sulfation aging issues in battery. (3) Battery lifetime prediction model [18], [19], [22]. The most previous works studied the battery aging issues for various application, but none of prior work addresses the battery aging and availability issues at datacenter level.
2) *Battery provisioning in datacenters:* Deploying batteries as energy buffers in datacenters to reduce power cost and improve power quality have received increasing attention. Recently, considerable proposals investigate battery provisioning and management in both conventional [6], [7], [10], [11], [23] and emerging renewable energy powered datacenters [4], [5], [12], [13], [24], [25], [26], [27]. Among those, [6], [7], [10], [11] primarily emphasize battery provisioning topologies (e.g., centralized, distributed or hierarchical deployment), and battery usage manner (aggressively leverage battery to shave peak power demands and stores energy during low load actively periods). While many recent studies [4], [5], [12], [13] also employ battery to buffer renewable power in datacenters. However, there has been no work that explores battery aging issues in the context of green datacenters. Battery aging issues can easily cause energy availability degradation with the growing proportion of massive batteries are deployed in datacenter.
3) *Energy aware/availability studies for computing system:* Several energy availability related proposals [28], [29], [30], [31], [32], [33] have been proposed for computing system to improve their availability. For example, [29] models the energy-aware scheduling problem and proposes two heuristic algorithms to find optimal job scheduling and minimize the energy consumption for datacenter. [33] investigates the realistic threat of power attacks for datacenter, an effective management policies to manage the power attacks and effectively improve the availability of datacenters. Our work distinguishes itself from previous studies in four aspects: (1) It exposes the emerging battery aging issues at the datacenter level. (2) It comprehensively analyzes battery aging mechanisms availability issues from the perspective of architecture and system designers. (3) It proposes a novel aging management framework that can jointly balance, slow down, mitigate and plan the battery aging process for high energy availability of datacenter. (4) It builds a real prototype to evaluate the battery aging behaviors and the efficiency of various aging management algorithms.

3 Background and Motivation

3.1 The Usage of Battery in Current Datacenters and Energy Availability Challenges

A large number of recent efforts [10], [11], [12], [33], [34], [35], [36] has repurpose energy storage devices (e.g., UPS

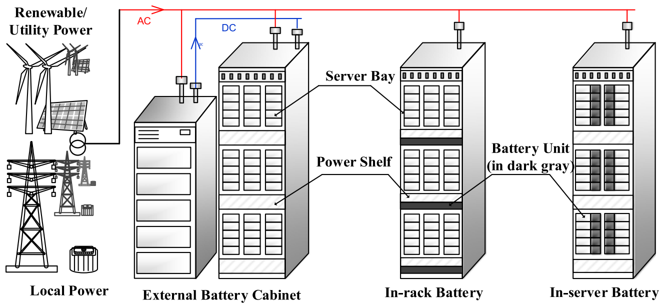

Fig. 1. Massive battery deployment in green datacenter and differ-ent methods for integrating batteries with computer servers.

batteries) to under-provisioning the power infrastructure in traditional datacenter for capital expenditure savings, and to smoothen power supply in renewable energy datacenter for operating cost savings. The deployment of batteries has been shifted from the traditional datacenter level to the server rack level, as shown in Fig. 1. One of the primary reasons is that such a distributed battery system reduces power conversion loss and provides a way to manage power in a fine-grained manner. For example, Facebook proposes both external battery and in-rack battery designs in their Open Rack project [3]. Microsoft's in-rack battery design uses at least two sets of battery packs to provide short-term backup [2]. To improve datacenter power usage effectiveness (PUE), Google has also tested in-server lead-acid batteries [1]. Several researchers at Hitachi have compared different battery integration methods in [37].

In recent years, datacenters are aggressively exploiting battery to reduce cost and improve sustainability. For example, batteries allow datacenters to shift load power, referred to as demand response [38]. By reshaping server power demand, primarily shaving power peaks, there are considerable opportunities for reducing operational cost (deferring loads to non-peak tariffs) and capital cost (oversubscribing power capacity).

In addition, datacenters are forced to integrate eco-friendly source of power (solar/wind energy) due to the ever-growing concern on greenhouse gas emission and global climate change. Unlike the stable supply of grid power, renewable energy generation is intermittent and it is not suitable to be directly consumed by datacenter servers. To avoid power brownout, a mass of inexpensive lead-acid batteries are cyclically used to smooth the variable renewable power supply [4], [5], [6], [7], [25], [27].

As new architecture emerges and massive usage pattern evolves, batteries have become the Achilles's heel in green datacenters. In Table-1, we briefly summarize different

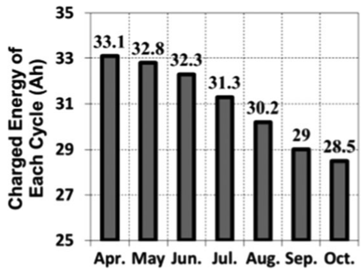

Fig. 3. Measured battery capacity drop due to aging over 6 months.

battery usage scenarios. In contrast to conventional battery systems used in emergency (rarely occurred and only lasts a few minutes [39]), emerging battery provisioning schemes result in a much higher charging/dis- charging rate, which hastens battery aging speed even cause serious threats to server energy reliability when the aged batteries are used to power datacenter servers. Thus, the most challenges of energy availability caused by battery aging are how to understand the battery aging mechanisms, identify the aged battery, and intelligently manage the aging issues in emerging green datacenter.

3.2 Aging Mechanism and the Impact for Availability

Based on our battery aging experiment prototype, we perform various battery aging test to present its impact on energy availability. As shown in Fig. 2, over a continuous operation of six months, the terminal voltage of a battery (fully charged) in our system can be decreased by approximately 9 percent. Although the average server loading is relatively the same, the voltage dropping rate increases as battery ages (0.1V/Month from Apr. to Jun. and about 0.3V/Month from Jul. to Sep.). Low terminal voltage often triggers server power blackout as the under-voltage battery cannot sustain high-current drawn to power server system [40]. In addition, the effectively stored energy in each charging cycle has also dropped by 14 percent under aggressive usage (Fig. 3). Typically, a battery unit is considered at end-of-life (i.e., not suitable for backup purpose in mission critical systems) when it fails to deliver 80 percent of its initial capacity [41]. If used as green energy buffer, such an aged battery can cause degraded energy efficiency. Our historical record shows that after six months the round-trip efficiency has also decreased by 8 percent, as shown in Fig. 4. As a result, unpredictable battery energy capacity reducing and energy efficiency decreasing caused by aging are more prone to lead to energy availability degradation during server powered by batteries.

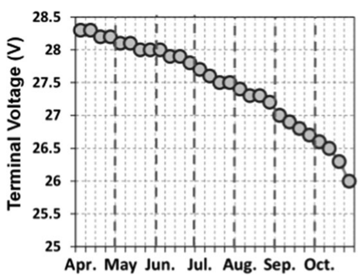

Fig. 2. Measured battery voltage drop due to aging over 6 months.

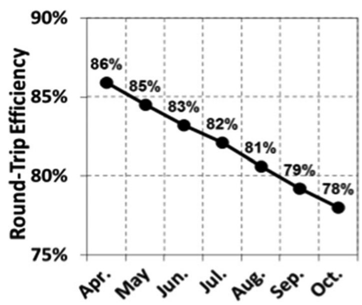

Fig. 4. Measured energy efficiency reducing as aging.

Battery aging mechanism refers to the processes of gradual deterioration of inner materials and the irreversible chemical reactions within the battery [15]. This study primarily focuses on lead-acid battery aging, which account for over 97 percent of industry batteries [42]. They are widely deployed in datacenters due to their maturity, low cost, and easy maintenance.

1) *Grids Corrosion*. The high positive potential at the positive electrode can result in lead grid corrosion. The aging processes cause the cross-section of the grid to decrease and grid resistance to increase. Consequently, the battery voltage drops and the maximum energy that can be stored becomes lower. The corrosion rate depends on the acid density, and electrode polarization [43].
2) *Active Mass Degradation/Shedding*. It contains many complex aging processes in the positive active mass (PAM) and the negative active mass (NAM). It leads to a change in the microstructure (e.g. active mass softening, recrystallization, loss of surface of active mass [44]). The AM shedding represents the active mass is permanently removed from the electrode. The AM shedding are usually accelerated by a high energy throughput, very low states of charge (SOC) and fast temperature changes.
3) *Irreversible Formation of Lead Sulfate*. It is also known as sulphation. When the electrodes are discharged, the active masses ($PbO2$ and Pb) are transformed into $PbSO4$ [16]. However, if a battery is not recharged timely, sulphate crystals can grow almost linearly with the solubility of sulphate ions and temperature. Therefore, the active masses are irreversibly converted to $PbSO4$ and no longer participate in electrochemical reactions [45]. The process is accelerated if battery keeps running at low SOC.
4) *Loss of Water*. In a valve-regulated lead-acid (VRLA) battery, water can gradually diminish due to extensive gassing. Moreover, water cannot be re-filled and the aging process is called drying out [46]. Overcharging and high battery temperature can affect the loss rate of water.
5) *Electrolyte Stratification*. The vertical distribution of the electrolyte density in battery is different, which leads to a preferred discharge at the bottom and preferred charge at the top of the electrolyte. The heterogeneous distribution accelerates sulphation at the bottom of the electrodes. The stratification aging occurs on the battery that is rarely fully recharged and the cells are deeply discharged with very low current [47]. It reduces the available capacity of a battery.

In general, the aging of battery is a synergistic effect of the above factors. Without careful management, cyclically used batteries can age quickly and further affect the energy availability of datacenter servers. They are often the hidden cause of undesired load shedding (for green datacenters that have stringent power budget). As a result, it is wise to model the battery aging factors and proactively manage the aging process from the view of computer system for the green datacenters.

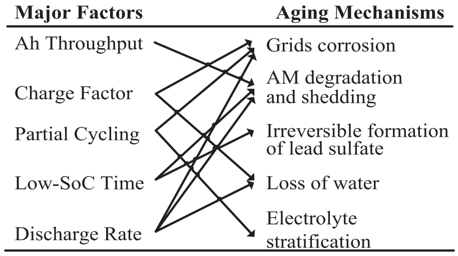

Fig. 5. Five key factors correlated with battery aging mechanisms.

4 Modeling and Quantitative Analysis of Lead-Acid Battery Aging Factors

Battery operating conditions (different voltage, current and temperature) largely determine the rate of aging processes. If we pose battery in different operating conditions and use it for a long term, it will lead to distinct aging processes. In this section, we calibrate the impact of different operating conditions using five metrics. Fig. 5 shows the correlation between these metrics and different aging mechanisms.

4.1 Normalized Ah Throughput (NAT)

The *Ah throughput* is defined as the ratio between the cumulative ampere-hour (Ah) output of a battery and the nominal total discharge capacity [48],

$$NAT = Q_{AT} \Big/ CAP_{nom} = \frac{\int_{t_0}^{T} I_{bat} dt}{CAP_{nom}} \tag{1}$$

In Eq. (1), Q_{AT} is the cumulative ampere-hour (Ah) output of battery from time t_0 to T. CAP_{nom} is the nominal lifelong output of the battery. It has been shown that the aggregated electric charge that can be cycled from a battery (before it wears out) is almost constant [48], [49]. Therefore, prior work has used *Ah* throughput for predicting battery lifetime under different charge/discharge conditions [49]. This factor can also be used to distinguish between backup battery operations (low *NAT*) and full cycling operations (high *NAT*). A high *NAT* value increases active mass degradation and shedding [49].

4.2 Charge Factor (CF)

The charge factor is expressed as the ratio of cumulative Ah throughput between battery charging and discharging,

$$CF = \frac{Ah_{charge}}{Ah_{discharge}} = \frac{\int_{t_0}^{T} I_{charge} dt}{\int_{t_0}^{T} I_{discharge} dt} \tag{2}$$

The charge factor indirectly indicates the operating conditions of a battery (e.g., partial discharging or float charging). Typically the charge ratio is between 1~1.3 [46]. In normal partial cycling conditions, the charge factor is close to 1. If the battery frequently receives float charge, its charge factor can increase dramatically. When the charge factor is too low, sulphation and stratification may become the major causes of fast aging. If the charge factor is above its normal range, the following aging mechanisms may be accelerated: active mass shedding, water loss, and corrosion.

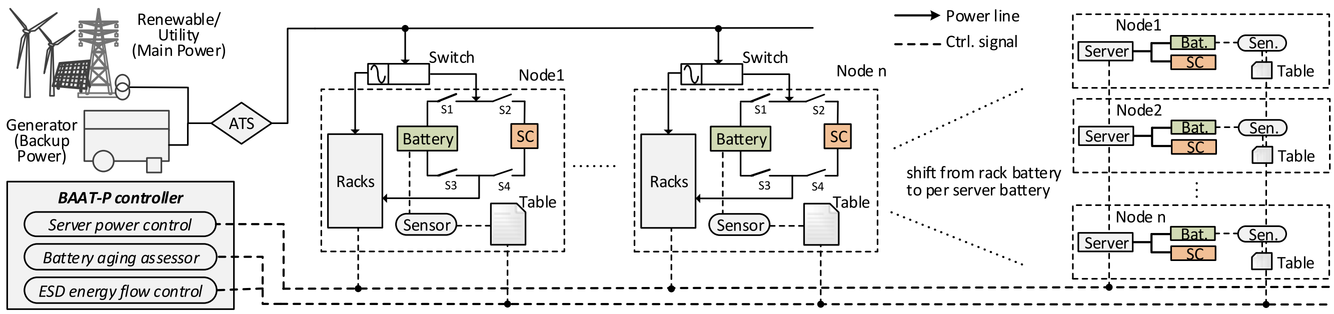

Fig. 6. *BAAT-P* battery aging management architecture. It demonstrates two types of architecture: per-rack integration and per-server integration of ESD.

4.3 Partial Cycling (PC)

PC reflects how a battery is used. The state of charge (SoC) of battery is divided into four ranges: A (100-80 percent), B (79-60 percent), C (59-40 percent) and D (39-0 percent). The probability of the Ah output being in range *X* is given by,

$$PC_X = \frac{\int_{t_0}^{T} I_{SoC\_X} dt}{\int_{t_0}^{T} I_{all} dt} \times 100\% \tag{3}$$

In Eq. (3), the numerator is the cumulative *Ah* output during the time the battery falls into SoC range *X*. The partial cycling value is then calculated by weighting functions as below:

$$\mathrm{PC} = (\mathrm{PC}_{(\mathrm{A})} \times 1 + \mathrm{PC}_{(\mathrm{B})} \times 2 + \mathrm{PC}_{(\mathrm{C})} \times 3 + \mathrm{PC}_{(\mathrm{D})} \times 4)/4 \tag{4}$$

In Eq. (4), the linear weighting factors reflect that the cycling at high SoC has less damaging for battery aging than the cycling at low SoC [43]. The higher value of *PC* will accelerate the battery aging such as corrosion and electrolyte stratification.

4.4 Deep Discharge Time (DDT)

This factor expresses the percentage (percent) of the entire period (from t_0 to *T*) within which the battery operates below 40 percent SoC. It can be calculated as,

$$T_{DD} = \frac{\int_{t_0}^{T} H(39\% - SoC) dt}{\int_{t_0}^{T} dt} \times 100\% \tag{5}$$

In Eq. (5), *H* is the Heaviside Step Function ($H(x) = 0$, $if\, x < 0$; $H(x) = 1$, if $\mathrm{x} \geq 0$). Staying at a low SoC accelerates irreversible sulphation. In contrast to *PC*, this factor is based only on the time and not on the *Ah* throughput

4.5 Discharge Rate (DR)

Low discharge rate (Ampere) has little impact on battery aging processes. But high discharge rate during low SoC duration can significantly accelerate aging. For example, the high discharge rate which exceeds the tolerable range can cause increased battery temperature. Taking the battery lifetime at 20°C as a baseline, a 10 °C temperature increase will result in a obvious reduction of the lifetime by 50 percent [45].

When using battery to power server system, the aging factors are related to the utilization of server. Higher utilization of servers will cause more dynamic power consumption and larger battery discharging current, which proportionally increases the value of aging factors of *NAT*, *CF*, *PC* and *DR*.

5 Aging Management Architecture and Algorithms for High Energy Availability

5.1 Aging Management Architecture

Fig. 6 depicts the architecture diagram of *BAAT-P*. The design mainly contains a sensor-table based power monitoring system, a *BAAT-P* controller and hybrid energy buffers: battery and supercapacitor (SC). The *BAAT-P* system can be integrated in green datacenter with emerging distributed energy storage [1], [2], [3] and it also applies to two types of distributed energy storage architectures: 1) rack level architecture: several racks share a pool of batteries (akin to Facebook's Open Rack design [3]), and 2) server level architecture: each server is equipped with a separate battery (similar to that in Google datacenters [1]).

The renewable/utility power can be tapped into each ESD node, as shown in Fig. 6. Each hybrid ESD node includes batteries and super-capacitors (SCs). The switches S1 and S2 control the charging operation of batteries and SCs respectively, while both S3 and S4 control the discharge operation of battery and SC to powers racks/servers. The *BAAT-P* controller contains *ESD energy flow control* algorithm that decides how to control the four switches ON/OFF for executing charging/discharging. To monitor the aging processes, each group of batteries are equipped with a sensor and a power table which records the battery utilization history logs as shown in Table 2. These log data are collected from the corresponding sensor of each battery and are sent to *BAAT-P* controller. The *Battery aging assessor module* in *BAAT-P* controller then calculates various metrics (*NAT*, *CF*, *PC*, *DDT*, and *DR*) for evaluating the aging process of each battery unit. The *Server power control* module in *BAAT-P* controller has the knowledge of the power information of each server through datacenter IT infrastructure such as IPDU or other power meters. By a software driver within each server, it can control the switches ON/OFF and can tune the power state of each server, such as Dynamic voltage and frequency scaling (DVFS), CPU clock throttling and virtual machine (VM) migration/consolidation.

With the architecture and monitored logs, *BAAT-P* controller can intelligently manage various battery aging issues and effectively improve the energy availability of datacenter servers. We propose four management schemes integrated with the *BAAT-P* controller: (1) load scheduling policy for hiding aging, (2) power capping strategy for slowing down

TABLE 1
Different Battery Usage Scenarios in Datacenter

Usage Objective	Usage Frequency	Aging Speed	Impact for Server Availability
Power Backup	Rarely	Slow	Light
Demand Response	Occasionally	Medium	Medium
Power Smoothing	Cyclically	Fast	Severe

TABLE 2
The Battery Data From Sensors

Variables	Description
Current	Charging and discharging current of battery
Voltage	Discharging voltage used for calculating SoC
Temperature	Battery surface temperature
Time	Total working time of a battery

aging, (3) hybrid energy storage scheduling policy for mitigating aging, and (4) algorithm for planning the aging speed. Among those, the aging-hiding scheduling is a battery aging-aware workload placement and consolidation policy, which can be added to existing datacenter workload management schemes for better addressing battery aging issues across datacenters. The aging slowing down strategy can reduce battery aging rate and avoid server availability degradation. The aging mitigating policy is a hybrid energy storage sources scheduling scheme which contributes to maintain workload performance for the performance critical workload. The planned aging is used when datacenter lifetime and battery aging rate are discrepant and require synchronization.

5.2 Aging Management Algorithms

5.2.1 Hiding Aging: Aging-Driven Workload Scheduling

1) *Technique Background*

In a distributed energy storage system, different commodity battery nodes can experience significant aging variation. The reasons of aging variation are: (1) current battery manufactures techniques lead to the deviations of actual aging time from their nominal specification, and (2) different power demand and supply of each server lead to varied charging/discharging behaviors on each battery node. If datacenter operators neglect the aging variation issue, they have to replace batteries that undergo faster aging irregularly, which unavoidably increases battery maintenance and replacement cost. Moreover, in case that some critical workloads are running on a "prone-to-wear-out" battery node but its unusual fast aging rate is ignored, the server energy availability is severely threatened as a battery node that ages faster can more easily cause unexpected server downtime. Therefore, to improve datacenter availability, it is wise to carefully balance battery aging variation issues.

We propose to hide the effects of battery aging across the datacenter. In detail, we schedule the workloads on different server nodes (associated with different battery units) in an aging-driven manner. We want the scheduling to be such that the aging slowest battery node can age faster, while the fast-aging battery node ages slower. As a result, the aging process of some of the worst battery units is hidden and the reliability of the node is improved.

2) *Implementation Details*

The battery aging-aware scheduling is performed when datacenter operators deploy new applications or perform workload consolidation. To effectively accomplish the scheduling, we consider the load power demand and the battery aging conditions in a coordinated manner.

a) *Load Power Demand Profiling*

Providing detailed and accurate workload power profiling information can help us place the workloads to the most appropriate battery nodes. Many datacenter applications can provide coarse granularity power profile, e.g., long time running services (web searching, memcached, etc.) and periodic/repetitive workloads (e.g., web crawling) [50]. In our study, using the power profiling information contributes to better estimation of its impact on battery aging.

b) *Battery Aging Consideration*

As mentioned in the section above, *BAAT-P* is able to assess battery aging process via five metrics (*NAT*, *CF*, *PC*, *DDT*, and *DR*) and we can calculate each aging quantization value based on the *BAAT-P* runtime logs. *BAAT-P* mainly relies on Ah-throughput (*NAT*), charge factor (*CF*) and partial cycling (*PC*) to determine workload allocation for hiding the aging effect. For example, a very high value of *Ah-throughput* indicates faster aging, since normally there is a fixed number of electric charge that can be cycled from a battery before it needs replacement [48], [49]. Meanwhile, a low *CF* value implies that the battery has more discharging events than charging to their full capacity. A lower *PC* indicates that a battery mostly stays at a very high depth of discharge (DoD).

Based on the implications of these three metrics, we can decide how to allocate battery nodes for the given workloads. Intuitively, we should place more loads on the aging slowest battery node, but the challenge is how to find the most suitable battery node. Base on different aging mechanisms, our solution is to combine the power demand profiling and the weighted value of the three metrics to dispatch workloads.

To find the optimal battery node to place new workloads, we roughly classify the power & energy profile of green datacenters into four scenarios as shown in Table 3. The power demand is treated as "Large" if the load power consumption exceeds 50 percent of the peak power. Otherwise, we define the power demand as "Small". Similarly, we classify the energy demand as either "More" or "Less" based on the duration of large power demand. The power and energy demands imply the load running length and the total energy request.

Different metrics have different sensitivities to the power and energy demand. For example, "Large" power demand reduces the value of $\triangle CF$ and $\triangle PC$, whereas the $\triangle NAT$ is more likely to decrease when the energy request becomes "Less". We use "High", "Low" and "Medium" to

TABLE 3
Different Battery Usage Scenarios in Datacenters

Power	Energy	ΔNAT	ΔCF	ΔPC
Large	Less	Medium	High	High
Large	More	High	High	High
Small	More	High	Low	Medium
Small	Less	Low	Low	Low

```
//Description: Hiding aging algorithm of BAAT-P;
//Input: Battery aging factors, workload power profiling results;
//Output: The battery nodes for workload adding or consolidation;
Workload power and energy demand profiling;
Calculate the Weighted_aging of all the battery nodes in datacenter;
Rank the aging conditions of all the battery nodes in datacenter;
IF ( Add new workload or task )
  do {
      Place the new workload on the battery nodes
            with minimal weighted value of △NAT, △CF, △PC ;
      Mark the battery node has been used;
  } While (All the new workloads are placed on battery nodes);
ELSE IF (Perform workload consolidation)
      {
        Virtual Machines migrate to the battery nodes
            with minimal weighted value of △NAT, △CF, △PC ;
        Shut down the aging fastest nodes
                      //The workload in this node has been migrated;
      }
END
```

Fig. 7. The flow chart of the aging hiding algorithm for *BAAT-P*.

respectively denote the impact of load power/energy demand on the three metrics as shown in Table 3. We set weighting factors for different metrics and calculate a weighted aging value as:

$$Weighted\_aging = a \times \Delta CF \ + b \times \Delta PC \ + c \times \Delta NAT \quad (6)$$

In Eq. (6), a, b and c are all weighting factors. The value of these factors is set as 50 percent in the "High" scenario, 30 percent in the "Middle" scenario and 20 percent in the "Low" scenario respectively. Our extensive training and experiments show that these weighting factors are fairly effective in our battery aging evaluation system. Note, the parameters above are related to the capacity of battery in different datacenters. A large value of the weighted aging indicates the fast aging pace. We can rank the weighted aging value of all the battery nodes in datacenters for the load placement, which is triggered when adding new jobs or performing workload consolidation, as shown in Fig. 7.

In summary, the battery aging-driven workload scheduling can balance the battery aging processes and improve the availability of battery nodes with faster aging rate in datacenters. As the inaccurate power profiling may lead to wrong load placement or consolidation sometimes, we can further leverage battery aging slowing down technique to remedy it.

5.2.2 Slowing Down Aging: Server Level Control

1) Technique Background. It is dangerous to discharge battery with high discharge rate during low SoC state. Doing so not only accelerates battery aging but also puts the server at the risk of low availability. Due to the intermittency of renewable energy supply and the imbalanced workload power demand of each server, some battery nodes may always stay at low SoC but experience high discharging rate. To further improve the server availability of low SoC battery nodes, we need to slow down the battery aging at appropriate timestamp.

2) Implementation Details. Our algorithm periodically checks two metrics: accumulated deep discharge time

```
//Description: Aging slowdown algorithm of BAAT-P;
//Input: Battery aging factors;
//Output: The battery nodes for workload migration and workload performance scaling;
Monitoring Deep Discharge Time ( △DDT)
     and Discharge Rate ( △DR) of each battery node during low SoC district;
IF ( △DDT > T_threshold && △DR > P_threshold) //the battery node reaches to the threshold
  IF (the Virtual Machines in the battery node can be migrated)
    {
       Calculate the Weighted_aging of all the battery nodes in datacenter;
       Rank the aging conditions of all the battery nodes in datacenter;
       Virtual Machines migrate to the battery nodes
              with minimal weighted value of △NAT, △CF, △PC ;
    }
  ELSE
       {
           Power throttling for the server nodes (e.g., DVFS);
       }
  END
ELSE
    Keep monitoring each battery node;
END
```

Fig. 8. The flow chart of aging slowdown algorithm for *BAAT-P*.

($\triangle DDT$) and high discharge rate ($\triangle DR$). If a battery always exhibits low SoC, it indicates the server loads allocated to the battery node are too heavy. In this case, if there is a peak power demand from the server and the primary power source (intermittent renewable power or utility power) lacks enough power budgets, battery voltage may reduce to the cut-out line and lead to server downtime. Therefore, we set thresholds $T_{threshold}$ and $P_{threshold}$ for $\triangle DDT$ and $\triangle DR$ to avoid unplanned battery cut out. $P_{threshold}$ is the maximal discharge current that can sustain discharge for 2 minutes ($T_{threshold}$). The two threshold parameters are related to the capacity of battery.

When the SoC of battery drops below 30 percent (low capacity for battery in general), *BAAT-P* periodically checks *DDT* and *DR* to see if they reach a preset threshold. It leverages virtual machine (VM) migration or performance scaling (DVFS) to avoid battery cut-off. It first checks the workload running on the battery node to see whether VM migration can be performed. If so, *BAAT-P* selects a target battery node, which has a minimal weighted aging value of *NAT*, *CF* and *PC* similar to the aging hiding technique. If the VM cannot be migrated due to resource constrains (e.g., CPU/Memory/Disk) elsewhere in the datacenter, we perform DVFS on servers to reduce power demand and promote the chances of battery charging to a higher SoC when the intermittent power supply becomes sufficient again. As DVFS may cause degraded performance, we preferentially use VM migration to reduce performance penalty. The details of our slowdown aging policy are depicted in Fig. 8.

5.2.3 Mitigating Aging: Hybrid ESD Energy Flow Control

1) Technique Background. Both of the battery aging management algorithms above can effectively mitigate battery aging processes and improve server availability, but they may lead to workload performance degradation if servers are frequently performed performance scaling policies such as DVFS, CPU clock gating and VM migration/consolidation) [65]. Therefore, we further leverage hybrid ESD energy flow control algorithm to mitigate battery aging rate while maintaining workload performance. In this paper, the

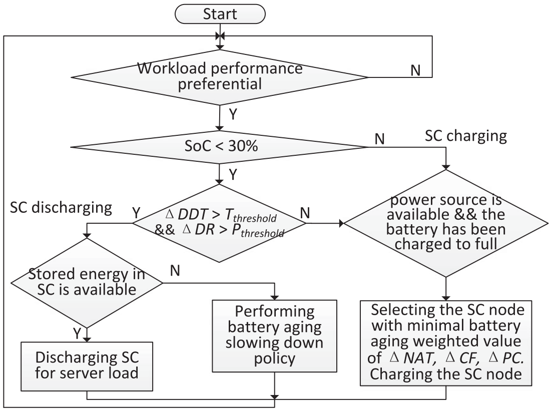

Fig. 9. The flow chart of aging mitigating algorithm for *BAAT-P.*

supercapacitors (SCs) and lead-acid batteries constitute the hybrid ESD and are deployed in the *BAAT-P* aging management architecture, as mentioned in Fig. 6 above.

2) Implementation Details

Many recent effects [11], [12], [26], [27], [51], [52] employ hybrid energy storage device (e.g., lead-acid battery, SCs, Li-on battery, etc.) to shave datacenter peak power. We leverage SCs to handle battery aging issues because SCs manifest distinct merits [53], [54]: (1) SCs commonly have much longer life cycle than batteries, (2) SCs has higher charging/discharging energy efficiency, (3) SCs has slower aging rate than battery because battery stores energy electrochemically while there is no chemical reaction in SCs. However, current cost of SCs is still high than batteries for large-scale and exclusive deployment in datacenter [53], [55], [68]. The self-discharge property of SCs may reduce the energy availability and leads to workload performance degradation. With frequent and timely charging for SCs, the self-discharge effect can be alleviated. As a result, we employ a fraction of SCs as hybrid ESD energy deployed in *BAAT-P* system to mitigate battery aging rate while improve workload performance for the performance critical workloads. We propose an efficient hybrid ESD energy flow aging management algorithm to handle the battery aging issue, as shown in Fig. 9. For the performance preferential workloads, the aging management algorithm periodically checks two battery aging factors *DDT* and *DR* when the SoC of battery drops below 30 percent. If the $\triangle DDT$ and $\triangle DR$ reach the preset threshold, it indicates the aging rate of the battery node is accelerated and may threaten server availability. At the moment, if the stored energy in SCs is available, the aging mitigating algorithm will discharge SCs to power the critical server nodes and stop using battery to power the nodes. Otherwise, it will perform battery aging slowing down policies, as shown in Fig. 9. When power source is available and battery has been charged to full, *BAAT-P* first find the hybrid energy storage unit in which the battery node has maximal aging weighted value, then SC in this unit will be timely charged for effectively mitigating battery aging rates in next discharging cycle.

The most benefits of battery aging mitigating algorithm is that it preferentially performs hybrid ESD energy sources scheduling policy rather than workload performance scaling scheme. The solution can effectively mitigate battery aging rate while maintaining the workload performance for the performance critical workload nodes. To our best knowledge, this is the first work that control hybrid ESD energy flow to handle battery aging issue in datacenter.

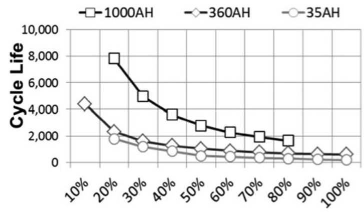

Fig. 10. Battery cycle life under varying depth of discharge (DoD).

5.2.4 Planned Aging: Aging Rate Management

The speed of battery aging and datacenter infrastructure aging are different. The average aging rate of battery is usually faster than servers and other IT equipment. Consequently, datacenter operators will end up with discarding servers or batteries before their expected end-of-life. In this case, significant performance may be wasted if slowing down the battery aging rate is excessively emphasized. If we know when the batteries will be discarded, we can use *BAAT-P* to "shift" some performance from the unused portion of the battery's lifetime to the used time, which refers to planned aging.

In Fig. 10 we show the cycle life data from different manufacturers (Hoppecke, Trojan and UPG Battery). It shows that the battery cycle life decreases by 50 percent if it is frequently discharged at a DoD above 50 percent. If we want to discard the battery after a certain life cycle, our goal is to apply planned aging techniques so that the battery is aggressively used before being discarded. Since different battery DoD implies different lifecycles, we can regulate the battery DoD to plan battery aging rate. We leverage the battery *Ah-throughput* capacity and the expected battery life cycles to calculate the needed DoD, shown below:

$$DoD_{goal} = (C_{total} - C_{used})/Cycle_{plan} \times 100\% \tag{7}$$

In Eq. (7), DoD_{goal} is the battery DoD for the planned aging rate. C_{total} is the nominal value of battery total Ah-throughput specified by the manufacturer, C_{used} is the past Ah-throughput that has been discharged; $Cycle_{plan}$ is the planned cycle which can be estimated base on the battery usage log in datacenter.

The planned aging technique can be built on the slowdown aging technique. We implement planned aging by replacing the low SoC value in slowdown aging technique with $(1 - DoD_{goal})$. The planned aging technique can effectively help us synchronize the battery aging rate with the end-of-life of datacenter infrastructures.

According to the four algorithms above, *BAAT-P* can jointly handling battery aging issues in datacenter. In brief, for the workload placement, battery aging-aware scheduling can hide aging; For the low SoC battery nodes, by keep checking the *DDT* and *DR* of battery, battery aging rate can be slowed down; For the performance critical workloads, leveraging hybrid energy flow control policy can effectively mitigate aging; For synchronizing battery aging with the aging of servers, battery aging planning algorithms can actively manage the aging rate. All the four aging

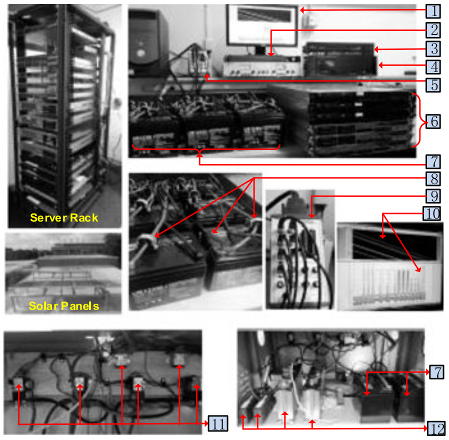

Hybrid energy storage device array	
(7)	Twelve 12V 35Ah sealed lead-acid batteries (420Wh)
(12)	Super-capacitors (SCs) , (Maxwell, 16V, 600F) (42Wh)
Sensor devices	
(5) (9)	Shielded connector block (NI BNC-2110[56]) for sensor signal transmitting from front sensors to sensor data acquisition card (NI PCI-6221[57])
(8)	Front sensors for batteries for collecting battery voltage, current and temperature [58]
Control server	
(2)	A customized low power intel i7 server for collecting monitored batteries/servers information and sending control commands to computing servers
Power Switcher	
(3)	IPDU[59], power monitor for computing servers
(4)	Power switch controller included some PLC, relays and DC-AC inverter to switch the power sources among utility, renewable power or battery power
(11)	Power relays control the switching between battery and SCs
Computing nodes	
(6)	Three IBM X series 330 and three HP ProLiant servers
Display	
(1) (10)	Displayed information includes data captured by sensors, system log trace, and various aging metrics calculated by the control server in real time

Fig. 11. A full-system implementation of *BAAT-P* prototype and six major function modules of the prototype system.

management algorithms can be packaged as specific APIs integrated in existing power scheduler. The aging related APIs can be invoked to handle battery aging issues similarly to other power management APIs in the data centers.

6 Evaluation Methodology

6.1 Prototype System Set Up and Configuration

We have built a heavily instrumented hybrid energy storage system (include battery and SCs). It enables us to analyze various battery aging issues and evaluate the proposed techniques. We synergistically integrate our energy storage system with a scaled-down green datacenter prototype also built from scratch. As shown in Fig. 11, our prototype system consists of a pack of batteries, SCs, sensors, servers, power switches and power meters. Specifically, we classify these hardware components into ten modules based on their functionality.

1) *Hybrid ESD Module.* Our system employs emerging distributed energy storage architecture. Each server is equipped with individual battery and SC unit (the initial capacity ratio of SC and batteries is set as 4:6 for evaluation). Multiple new sealed lead-acid batteries/super-capacitors are partitioned in our experiment to conduct comparable experiments with different aging management policies, and we can obtain the battery usage log of each power management policy respectively.
2) *Sensor Module.* The sensor device module includes several front-end sensor devices, which are used to measure the voltage, current and temperature of each battery. A data acquisition card [57] plugged-in server motherboard via PCI-E interface is used to collect the data from sensors to hard disk. The sensor data can be viewed by LabVIEW [60] in real time.
3) *Control Module.* Our control server is a customized i7 low power server, which runs our *BAAT-P* algorithms. The control server can collect the sensor data and calculate different metrics to access the aging process, monitor computing server power consumption information via IPDU [58]. Various aging management policies can be integrated into the control server to monitor and control both computing servers and batteries.
4) *Power Module.* The power switch module mainly contains power switch and power conversion equipment such as IPDU, PLC, relays, battery charger and DC-AC inverters. The power switcher can dynamically switch the power sources among utility, battery power and renewable energy (we tap into one solar power line from the PV panel on the roof of the building to our prototype system) to power servers and it also can switch the utility or renewable power to charge batteries and SCs. It also controls the power switch between battery and SCs for charging and discharging. The switch module is controlled by control server via SNMP commands over the Ethernet.
5) *Compute Nodes and Display Module.* We use three IBM servers and three HP servers to run workloads for the aging experiments. The display module can visualize the data captured by sensors via LabVIEW and the aging impact factors calculated by control server in real time.

6.2 Workload Deployment

We deploy six datacenter workloads. Three of them come from *Hibench* [61]: *Nutch Indexing (ID); K-Means Clustering (KM)*, and *Word Count (WC)*. These workloads represents today's large-scale search indexing application, machine learning application, and MapReduce jobs. We also select three popular cloud workloads from the *CloudSuite* [62]: *Software Testing (ST)*, *Web Serving (WS)*, and *Data Analytic (DA)*. For example, *Software Testing* is a resource-hungry and time-consuming application that allows us to stress our servers and distributed batteries.

We deploy Xen 4.1.2 hypervisor as the virtual machine monitor (VMM) in our system prototype. All the workloads are hosted in virtual machines (VMs) and the workloads can be easily managed by performing VM spawning, pausing and migration among server nodes. Through software driver, we can dynamically set the frequency of processors.

TABLE 4
The Comparison of Four Aging Management Schemes

Schemes	Method Description
e-Buff	Aggressively use battery as the green energy buffer to manage supply/load power variability
BAAT-h	Only use aging-aware VM migration technique to hide battery aging variation
BAAT-s	Only use aging-aware CPU frequency throttling to slow down battery aging
BAAT-P	Coordinate aging hiding, slowing down and mitigating techniques to dynamically manage battery aging

The operating condition of our system is affected by the available solar energy. Normally we turn on the first server at 8:30 AM and all servers are shut down usually after 6:30 PM. When solar power budget is temporarily unavailable, our system can make checkpoint and all VM states are saved. Our controller can precisely control battery charger so that the stored energy reflects the actual solar power supply on our prototype.

7 Evaluation Results

In this section, we evaluate the impact of various power management schemes on battery aging. We first compare the *BAAT-P* with three baseline aging management algorithms as shown in Table 4. Among those, *e-Buff* represents the power design approaches similar to previous work [4], [7], which aggressively employs battery energy to manage power mismatch between supply and demand. *BAAT-s* and *BAAT-h* are two simplified versions of our *BAAT-P* aging management scheme. *BAAT-s* only focuses on slowing down the battery aging processes, while *BAAT-h* mainly emphasizes battery aging hiding.

In the following evaluation section, we first employ pure lead-acid battery to evaluate the battery lifetime, battery aging test and the effect of battery planned aging for datacenter. Then, we discuss the server availability, workload performance improvement and cost variation when using the hybrid ESD (battery and SC) scheme in the *BAAT-P* prototype system. Notes, when comparing the experiment results of different aging management policies, we select a worst battery node that has the most *Ah-throughput* in each aging management scheme.

7.1 Prototype System Running Profiling

We first profile the system runtime of our prototype across different solar generation scenarios (as shown in Fig. 12) by analyzing the generated logs and collected system runtime traces. The six servers in the prototype are respectively powered by six groups of battery. The solar energy is used to charge each battery. As expected, due to the intermittent solar power budget and different server power demands, the usage frequency of the six battery pack varies significantly, which leads to different battery aging rates.

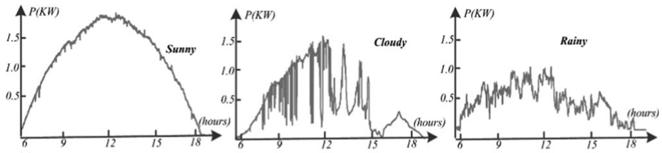

(a) Sunny day solar generation log (b) Cloudy day solar generation log (c) Rainy day solar generation log

Fig. 12. Solar generation scenarios in different weather conditions.

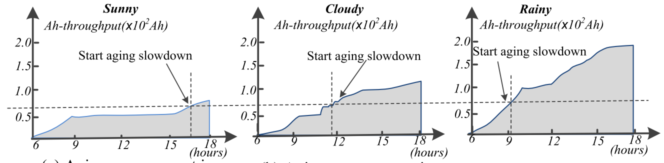

(a) Aging management in sunny day (b) Aging management in cloudy day (c) Aging management in rainy day

Fig. 13. The variation of *Ah-throughput* in different weather conditions.

To quantify the aging processes, we present the variation of several aging-related metrics of one battery node. We present the values of *NAT*, *CF*, and *PC* under three typical weather conditions. The total energy budget for the Sunny, Cloudy, and Rainy day is 8 kWh, 6 kWh and 3 kWh, respectively.

It is clear that the battery nodes yield less *Ah-throughput* in sunny day than the other two weather conditions, as shown in Fig. 13. This is because the solar energy can afford most server power demands and the batteries are rarely used. The shadow regions of the figure also denote the battery lifetime distribution under different weather conditions when using the accumulated *Ah-throughput* model to present the anticipant battery lifetime. In contrast, the *Ah-throughput* of cloudy and rainy day obviously increases. As shown in Fig. 14, the *CF* in sunny day is higher than cloudy and rainy days, which implies that the battery node has been recharged more frequently. The *PC* value in sunny day reflects the fact that battery node stays in high SoC region. We can see that the battery node stays high SoC at most time in sunny day. Therefore, on the cloudy and rainy days, the battery node has more aging decay by exhibiting high *Ah-throughput*, low *CF* and low *PC*.

With the quantized aging process, our system has the knowledge of all the battery aging processes under different weather conditions. Therefore, the aging-aware operations can be timely performed. Take *Ah-throughput* for example, we start to slow down battery aging when the accumulative Ah-throughput of the battery reaches to the pre-defined

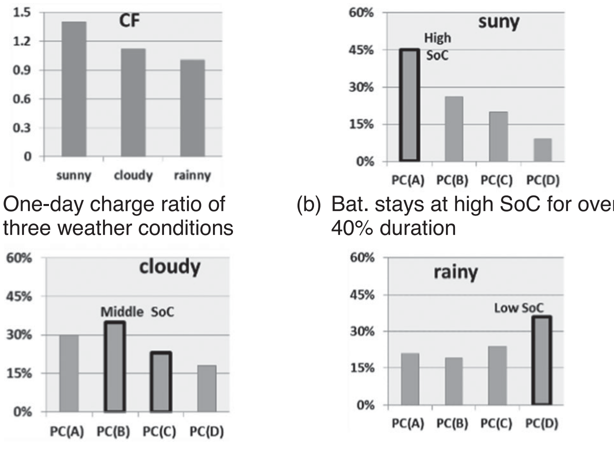

(a) One-day charge ratio of three weather conditions
(b) Bat. stays at high SoC for over 40% duration
(c) Bat. stays at medium SoC for 60% duration
(d) *Bat. stays at low SoC for over 30% duration*

Fig. 14. The comparison of of *CF* and *SoC* in three weather conditions. Note: PC(A): $SoC > 80\%$; PC(B): $80\% >= SoC > 60\%$, PC(C): $60\% >= SoC > 40\%$; PC(D): $SoC =< 40\%$.

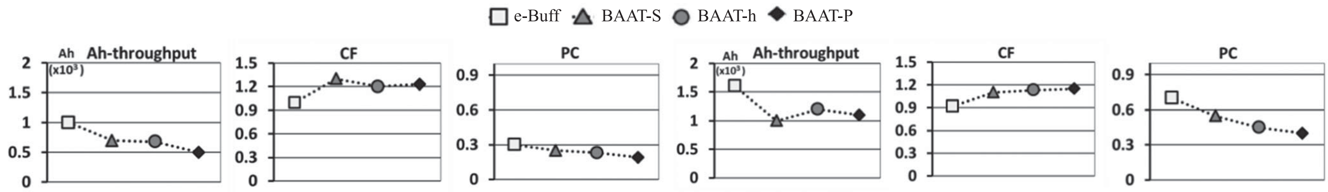

(a) Battery aging experiment in a sunny day

(b) Battery aging experiment in a cloudy day

Fig. 15. The aging metrics comparison of four aging management algorithms in four power management schemes.

threshold. The slowdown time varies in different weathers, as marked in Fig. 13. Based on the logs of different aging metrics, our *BAAT-P* aging management framework can combine load power profiling information and weighted aging ranking to effectively hide and slow down battery aging.

7.2 Aging Test and Experiments Design

In this section, we design an aging experiment that can be launched iteratively. By executing the battery aging test for ten months, we further present the effectiveness of different aging management schemes on battery aging experiments. Specially, we compare our aging management algorithm *BAAT-P* with three baseline algorithms.

We start by using the new batteries to conduct the aging experiments after completing the prototype system. We deploy and iteratively run the workloads hosted in virtual machines on our computing server nodes. The workload is aggressively run to accelerate battery aging. We record the running logs on our control server node. The logs contains the workload power demands of the six computing server nodes, the one day solar power generation trace and aging metrics information (*NAT, CF, PC, DDT,* and *DR*) of six battery nodes.

After ten months, our system collects various log data during runtime. By analyzing the log data, we are able to compare the effectiveness of the four aging management algorithms. Figs. 15a and 15b respectively show the impact of the four aging management algorithms on the battery aging test by battery aging metrics. (1) Battery ages faster in the harsh usage conditions, for example, the Ah-throughput of *e-Buff* algorithm in cloudy is increased by 35 percent than the sunny day on average. (2) By leveraging our *BAAT-P* battery aging aware power management algorithm, the battery aging progresses can be effectively mitigated on average.

The *e-Buff* power management scheme always aggressively uses battery to bridge the gap between server power demands and solar power budget, which obviously accelerates the battery Ah throughput (1.3 X more than *BAAT-P* on average), especially when it is cloudy and the battery stays old aging states (2.1X). The *BAAT-s* power management scheme is aware of the aging issues and leverages power capping technique (DFVS) to slow down the battery aging progresses, but the simplified aging management policy only perform DVFS operation for some computing server when the battery power cannot afford the mismatch of solar budget and server power demands, which is a passive solution and leads to workload performance degradation (detailed in next Section 7.5). The *BAAT-h* aging management scheme employs VM migration mechanism to alleviate one battery node aging progress. But it lacks the holistic battery node aging information (e.g., weighted battery aging metrics) for VM migration, Therefore, the VM migration is unaware the aging state of other battery nodes, which make the migration become random and low efficiency. Our *BAAT-P* algorithm is a holistic battery aging-aware policy which calculates and ranks the weighted aging metrics of all battery nodes. It also refers the workload power profiling information to manage aging. Based on the techniques, *BAAT-P* can dynamically slow down, mitigate and hide the battery aging effect across all the battery nodes in datacenters, which balances the aging effect and prevents severe aging of some battery nodes.

By comparing and analyzing their aging metrics logs, we can see that the *BAAT-P* can: (1) Effectively reduce the total Ah-throughput and avoid the worst battery node which is aggressively used. (2) By workload power capping, the worst battery node can obtain more solar charging chances and has higher CF. (3) By timely load scheduling, the power pressure of worst battery node is mitigated and its PC value is increased. Therefore, based on the aging information, our *BAAT-P* framework can effectively manage the battery aging in datacenter. By weighting the three aging metrics (using Eq-6 with same weighting factors) in worse case condition (cloudy and old battery), we conclude that the *BAAT-P* can reduce battery aging speed in the worst case by 41 percent.

7.3 Impact on Battery Lifetime

We find that the renewable energy availability and load power demand greatly affect the battery lifetime as the battery is cyclically and frequently used. We further evaluate the battery lifetime impacts under different solar energy potentials and server capacities (Figs. 16 and 17).

We first consider geographic locations that have different solar energy availabilities (represented by sunshine fraction, the percentage of time when sunshine is recorded [63]). It is clear that battery lifetime increases when the availability of solar energy grows. This is primarily because batteries do not need to be frequently discharged when the solar

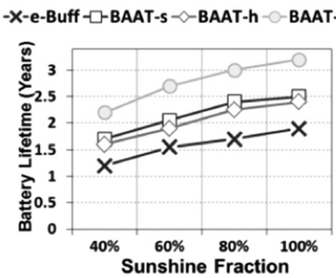

Fig. 16. Battery lifetime under different green energy supplies.

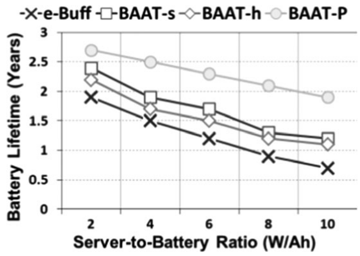

Fig. 17. Battery lifetime under different capacity (server to battery).

power output is high. In fact, sometimes the entire load can be directly powered by solar panels without using the stored green energy. *BAAT-P* could extend battery life by 72 percent on average, compared to *e-Buff*. The lifetime improvement for *BAAT-s* and *BAAT-h* is 40 and 32 percent, respectively. Our results show that aging slowdown policy has a larger impact on battery lifetime, compared to the aging balancing policy.

On the other hand, by varying the loading placed on batteries (represented by server-to-battery capacity ratio), our results demonstrate three key findings: (1) A heavy server-to-battery ratio accelerates aging. As we increase the server-to-battery capacity ratio from 2W/Ah to 10W/Ah, the average battery lifetime decreases by 33 percent (Fig. 17). This is mainly because heavy server loading is more likely to create power spikes, which launch aging by deep battery discharging and high discharge rate. (2) The optimization effectiveness of *BAAT-P* on battery lifetime becomes greater when the server system is heavily power-constrained. Although the battery life decreases when adding servers, we observe that the performance improvement of *BAAT-P* (compared to *e-Buff*) grows from 39 percent to 1.4X. This indicates the benefits of our battery aging management scheme actually increase when a green datacenter has to frequently use energy storage systems to handle power shortfall. (3) Excessively increasing battery capacity to reduce server-to-battery ratio may not be wise. In Fig. 15, doubling the installed battery can cut the battery-to-server ratio by half, but may result in less than 30 percent lifetime improvement. This is because the aging process is not linearly correlated with the reduction in server loads. One should carefully plan the battery capacity for aging management.

7.4 Impact on Server Availability

To improve the availability of datacenter, one must carefully manage battery aging and improve the availability of each server. This section focuses on the availability evaluation of server system from the perspective of various battery aging tests. The key aging factor that directly correlates with server availability is deep discharge time (DDT). Prior work

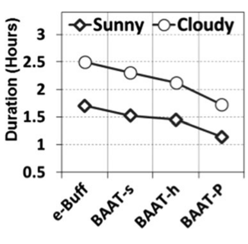

Fig. 18. Low-SoC duration. *BAAT-P* reduces deep discharge duration.

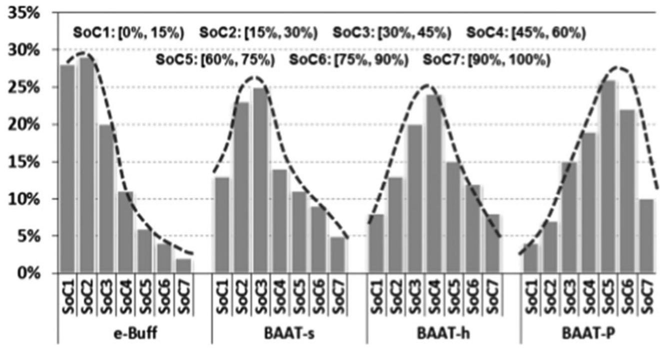

Fig. 19. Distribution of theSoC under different battery aging management schemes.

has shown that datacenter must leave 2 minutes of reserve capacity in UPS battery for high availability [64]. A low SoC means less reserved energy, which is dangerous when servers plan to draw large amount of power from the battery.

We collect and calculate the low-SoC duration of different aging management schemes from the experiment logs, as shown in Fig. 18. As can be seen, the *e-Buff* scheme can easily make some batteries enter low-SoC state for a long time. It potentially increases the chance of power budget violation and causes single point of failure (SPOF) when a battery happens to run out of power upon load power spikes. In contrast, *BAAT-P* can dynamically slow down and balance the battery SoC across all the battery nodes. Moreover, the hybrid ESD energy flow control policy in *BAAT-P* can effectively leverage supercapacitor to mitigate the battery aging rate, which can effectively eliminate SPOF. The results show that *BAAT-P* could increase battery availability by 55 percent on average based on the statistics of low-SoC duration of the worst-case battery node. In Fig. 19 we further evaluate the distribution of deep discharging over 6 months. It is clear that *e-Buff* tends to create low-SoC batteries, whereas *BAAT-P* can shift the most likely SoC region towards 90-100 percent. Therefore, *BAAT-P* increases the resiliency and emergency handling capability.

7.5 Workload Performance Discussion

Battery hiding and slowing down aging management schemes may cause workload performance degradation because both of the two aging management algorithms always protect battery from aging by scaling workload performance. In comparison, the battery aging mitigating policy can effectively avoiding workload performance degradation by leveraging hybrid ESD energy source scheduling rather than workload performance scaling. In this section, we first compare the impact of different battery aging management schemes on workload performance. Then we present the performance varies with the capacity increasing of supercapacitors (SC).

The experimental method is the same as Section 7.2. First, we evaluate the number of completed tasks (i.e., throughput) for the four aging management policies under fixed capacity (30 Ah) of hybrid energy storage devices (battery and ultra-capacitor), as shown Fig. 20. We can see that the six workloads have different task completion rate under the same energy budget. The average number of completed

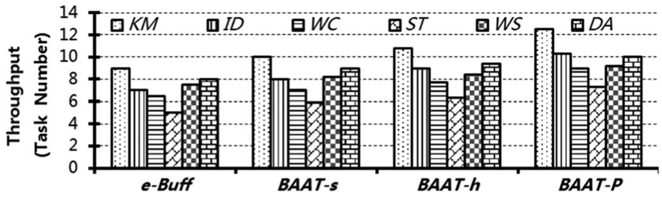

Fig. 20. Throughputs comparison of different policies for six workloads.

task of *e-Buff* policy is 7. In contrast, with improved policies (*BAAT-s*, *BAAT-h* and *BAAT-p*), the throughput gradually increases, because the optimized algorithms can effectively improve the energy availability which increases the throughputs of tasks. We further evaluate the total throughput of the four aging management schemes in one day (Fig. 21). Intuitively, the *e-Buff* algorithm can yield the best performance as it ignores battery aging issue and aggressively uses battery to satisfy workload performance. However, when the solar budget is inadequate and the capacity of battery reduces to the cut-off level, the server has to be shut down. During server downtime, the throughput is zero for *e-Buff*. The *BAAT-s* aging management policy always leverages power capping mechanisms to alleviate the battery aging and avoid aggressive battery usage. However, it reduces the CPU computing speed and leads to workload throughput degradation. The purpose of workload migration in *BAAT-h* algorithm is to avoid battery aging acceleration. However, as mentioned earlier, *BAAT-h* lacks the holistic information of battery aging and power demands, and its low-efficiency migration may causes severe performance overhead (e.g., frequent VM stop and restart). Based on the profiling information and the calculated weighted aging metrics of all the battery nodes, our *BAAT-P* algorithm can dynamically slow down battery aging and schedule workloads according. Compared to *e-Buff*, *BAAT-P* can improve the performance by 35 percent on average (Fig. 21).

We further evaluate the workload performance improvement when provisioning different capacity ratios of SC in hybrid ESD buffer. By keeping the constant total capacity of ESD, we change the capacity ratio between SCs and batteries. We adjust the *Depth-of-Discharge* (DoD) of energy buffers to generate different available capacity of battery and SC, The power switches in *BAAT-P* can control the energy budget by disabling the utilization of battery/SC once they hit the DoD threshold.

By iteratively running the workloads with *BAAT-P* power scheme, we respectively obtain the average per-formance improvement in different ESD ratios, as shown in Fig. 22. Note that the result is normalized to the ratio of 2:8.

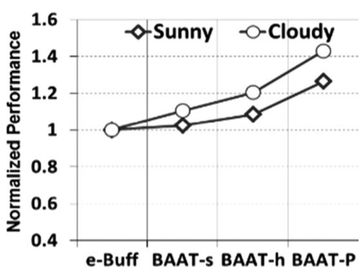

Fig. 21. Performance improvement of different policies.

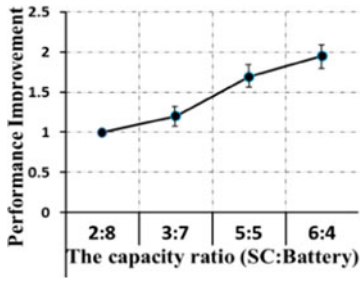

Fig. 22. Performance impact of different capacity ratio of SC.

It obviously shows that the more proportion of SCs can obtain better performance improvement. However, the relationship between performance and SC capacity ratio is nonlinear. The results contribute to the rightsizing of an appropriate proportion of SC in the hybrid ESD energy buffer for the optimization among the cost of SC, battery aging rate, server availability and workload performance, etc.

7.6 Benefits of Planned Aging for Availability

Another distinctive feature of *BAAT-P* is that it is capable of planning the battery aging rate for better utilizing the energy storage resources. Typically, the lifetime range of a lead-acid battery and a datacenter is 3∼10 years [66] and 10∼15 years [67], respectively. Without aging planning, it is highly likely that datacenter operators have to discard the latest replaced batteries before they wear out.

BAAT-P can adjust battery DoD to modulate battery service life and synchronize it with the end-of-life of datacenter infrastructures to gain more performance benefits. As shown in Fig. 23, the performance improvement is not linear with DoD variation. When the DoD increases from 40 to 60 percent, the performance improvement is more visible than when the DoD increases from 70 to 90 percent. This is because compared with a normal DoD range of 20∼40 percent, the latter will yield too low battery SoC, which leads to reduced battery lifetime.

7.7 Cost Benefits and Discussion for BAAT-P

Efficient aging management scheme can greatly increase battery lifetime and return on investment (ROI) of battery due to the reduced battery depreciation cost. By varying the threshold of our aging slowdown optimization algorithm, we observe that the cost benefits changes. Increasing the threshold allow batteries to offload more burden, thereby increasing their lifetime and reducing cost. Compared to *e-Buff*, *BAAT-P* can achieve 33 percent cost reduction, as shown in Fig. 24. Note that aggressively applying the aging slowdown algorithm is not wise since it may cause unnecessary performance degradation.

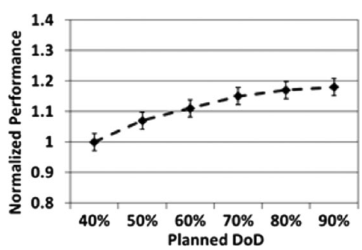

Fig. 23. Performance impact of planned DoD.

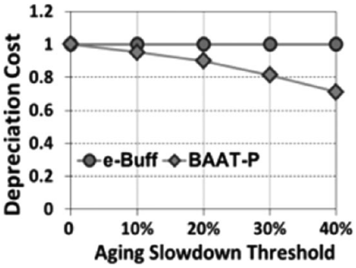

Fig. 24. BAAT-P reduces annual depreciation cost.

Another benefits of *BAAT-P* aging management scheme is that it allows existing green datacenters to expand (scale-out) without increasing the total cost of ownership (TCO). This is because the cost savings due to improved battery lifetime can actually be used to purchase more servers. Fig. 25 simulates the number of server that is allowed to be added to datacenter without increasing the TCO. The result is closely related with the sunshine fraction since the actual server that can be installed depends on the available solar power budget. In geographic locations that have abundant solar energy, one can add up to 17 percent more servers. Note that the server expansion ratio does not linearly grow when server number increases. The main reason is that the battery lifetime decreases as server number grows, which adds depreciation cost.

At last, we present the ESD cost related simulated results based on our prototype, as shown in Fig. 26. The cost mainly includes initial cost of ESD device, replacement cost, maintenance cost and server downtime cost caused by battery aging. At the beginning, the initial cost of *BAAT-P* is higher than *e-Buff* (about 3X), because *BAAT-P* is comprised by battery and expensive SC, but *e-Buff* only contains homogenous battery. However, the interesting observation is that the cost growth of *BAAT-P* is very slow but the cost of *e-buff* increases exponentially. This is because the SC in *BAAT-P* has longer lifecycle which reduces the maintenance cost and replacement cost than the pure battery system of *e-Buff*. The lifetime of lead-acid battery is about 1 to 3 year base on the different usage frequency. Therefore, battery has to be replaced every two years on average which increases the replacement and maintenance cost. Moreover, *BAAT-P* has higher efficiency for handling battery aging and can maintain better workload performance, both of which can further reduce the datacenter downtime cost. As a result, due to the expensive initial cost, the cost of *BAAT-P* scheme is higher than the *e-Buff* scheme in the first five years, but the cost of *e-buff* gradually exceeds the *BAAT-P* during the last five years.

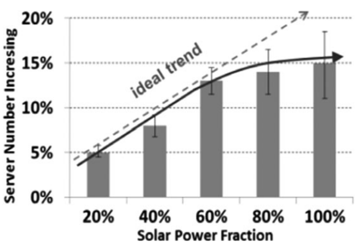

Fig. 25. Datacenter can economically trade off battery life for server capacity.

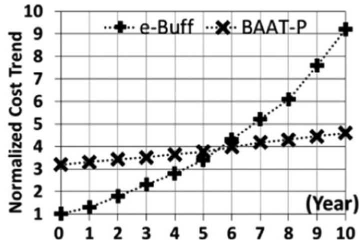

Fig. 26. The comparison of predicted cost trend in 10 years.

8 Conclusion

In this study we explores battery aging issues on a scaled-down prototype over ten months and proposes a novel aging management framework to dynamically manage battery aging in emerging green datacenters, in which massive distributed battery systems are increasingly being deployed for the purpose of power and energy management.

By measurement and modeling battery aging issue in datacenter, we propose Battery Anti-Aging Treatment Plus (*BAAT-P*), a novel aging management framework which can jointly hide, slow down, mitigate and plan battery aging for high energy availability of datacenters. It leverages quantified battery aging metrics abstracted from runtime performance statistics to efficiently handle battery aging at the computer architecture and system levels. We conduct detailed experiments on a real system prototype build from scratch. Our results show that *BAAT-P* can leads to 41 percent battery aging rate reduction even in the worst case and can improve battery lifetime by 72 percent on average. Meanwhile, *BAAT-P* could reduce the performance overhead caused by inefficient battery management, thereby improving workload performance by 35 percent. Our design allows datacenter to reduce 33 percent battery annual depreciation cost.

Acknowledgments

This work is supported in part by NSF grants 1423090, 1320100, 1117261, by NSFC grant 61602368, 61502302, by the National Key Research and Development Program of China: 2016YFB0200501 and by China Postdoctoral Science Foundation 168236.

References

[1] (2009). [Online]. Available: http://www.cnet.com/news/google-uncloaks-once-secret-server-10209580/

[2] (2011). [Online]. Available: http://www.datacenterknowledge.com/archives/2011/04/25/microsoft-reveals-its-specialty-servers-racks/

[3] Facebook, Hacking Conventional Computing Infrastructure, 2011. [Online]. Available: http://opencompute.org/, 2011.

[4] D. Meisner and T. F. Wenisch, “Peak power modeling for data center servers with switched-mode power supplies,” in *Proc. Int. Conf. Low Power Electron. Des.*, 2010, pp. 319–324.

[5] I. Goiri, W. Katsak, K. Le, T. Nguyen, and R. Bianchini, “Parasol and GreenSwitch: Managing datacenters powered by renewable energy,” in *Proc. Int. Conf. Architectural Support Programming Languages Operating Syst.*, 2013, pp. 51–64.

[6] V. Kontorinis, et al., “Managing distributed UPS energy for effective power capping in data centers,” in *Proc. Int. Symp. Comput. Archit.*, 2012, pp. 488–499.

[7] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, “Benefits and limitations of tapping into stored energy for datacenters,” in *Proc. Int. Symp. Comput. Archit.*, 2011, pp. 341–352.

[8] Navigant Consulting, Inc. "Energy storage systems for solar and wind power integration will total nearly 22 Gigawatts of installed capacity from 2013 to 2023," *Forecasts Navigant Res.*, 2013.

[9] P. Jones, "DCD industry census 2013: Data center power," *Datacenter Dynamics Industry Census*, 2014. http://www.datacenterdynamics.com/.

[10] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Khessib, "Ag-gressive datacenter power provisioning with batteries," *ACM Trans. Comput. Syst.*, vol. 31, no. 1, 2013, Art. no. 2.

[11] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: What, where, and how much," in *SIGMETRICS Performance Eval. Rev.*, vol. 40, pp. 187–198, 2012.

[12] L. Liu, et al., "Leveraging heterogeneous power for improving datacenter efficiency and resiliency," *IEEE Comput. Archit. Lett.*, vol. 14, no. 1, pp. 41–45, Jan.-Jun. 2015.

[13] N. Sharma, S. Barker, D. Irwin, and P. Shenoy, "Blink: Managing server clusters on intermittent power," in *Proc. Int. Conf. Archit. Support Programming Languages Operating Syst.*, 2011, pp. 185–198.

[14] S. Buller, M. Thele, E. Karden, and R. W. Doncker, "Impedance-based non-linear dynamic battery modeling for automotive applications," *J. Power Sources*, vol. 113, no. 2, pp. 422–430, 2003.

[15] Y. Sun, H. Jou, and J. Wu "Aging estimation method for lead-acid battery," *IEEE Trans. Energy Conversion*, vol. 26, no. 1, pp. 264–271, Mar. 2011.

[16] P. Ruetschi, "Aging mechanisms and service life of lead–acid batteries," *J. Power Sources*, vol. 127, no. 1, pp. 33–44, 2004.

[17] H. Catherino, F. Feres, and F. Trinidad, "Sulfation in lead–acid batteries," *J. Power Sources*, vol. 129, no. 1, pp. 113–120, 2004.

[18] D. Sauer and H. Wenzl, "Comparison of different approaches for lifetime prediction of electrochemical systems using lead-acid batteries as example," *J. Power Sources*, vol. 176, no. 2, pp. 534–546, 2008.

[19] D. Pavlov, G. Petkova, M. Dimitrov, M. Shiomi, and M. Tsubota, "Influence of fast charge on the life cycle of positive lead–acid battery plates," *J. Power Sources*, vol. 87, no. 1, pp. 39–56, 2000.

[20] R. A. Dufo-López, J. M. Lujano-Rojas, and J. L. Bernal-Agustín, "Comparison of different lead–acid battery lifetime prediction models for use in simulation of stand-alone photovoltaic systems," *Appl. Energy*, vol. 115, no. 4, pp. 242–253, 2014.

[21] L. Cloth, M. Jongerden, and B. Haverkort, "Computing battery lifetime distributions," in *Proc. Int. Conf. Dependable Syst. Netw.*, 2007, pp. 780–789.

[22] T. Panigrahi, et al., "Battery life estimation of mobile embedded systems," in *Proc. Int. Conf. VLSI Des.*, 2001, Art. no. 57.

[23] C. Li, et al., "Oasis: Enabling datacenter to scale out economically and sustainably," in *Proc. Int. Symp. Microarchitecture*, Dec. 2013, pp. 322–333.

[24] S. McCluer, "Battery technology for data centers and network rooms: lead-acid battery options," APC White Paper #30.

[25] C. Li, et al., "Towards sustainable in-situ server systems in the big data era," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2015, pp. 14–26.

[26] L. Liu, et al., "HEB: Deploying and managing hybrid energy buffers for improving datacenter efficiency and economy," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2015, pp. 463–475.

[27] L. Liu, H. Sun, C. Li, Y. Hu, N. Zheng, and T. Li, "Towards an adaptive multi-power-source datacenter," in *Proc. Int. Conf. Supercomputing*, Jun. 2016, Art. no. 11.

[28] X. Li, Y. Wang, H. Chen, X. Chu, Y. Wu, and Y. Qi, "Reliable and energy-efficient routing for static wireless ad hoc networks with unreliable links," *IEEE Trans. Parallel Distrib. Syst.*, vol. 20, no. 10, pp. 1408–1421, Oct. 2009.

[29] L. Mashayekhy, M. Nejad, D. Grosu, Q. Zhang, and W. Shi, "Energy-aware scheduling of MapReduce jobs for big data applications," *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, no. 10, pp. 2720–2733, Oct. 2015.

[30] H. Nguyen, J. Song, and Z. Han, "Distributed demand side management with energy storage in smart grid," *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, no. 12, pp. 3346–3357, Dec. 2015.

[31] Y. Zhang, J. Liu, and M. Kandemir, "Software-directed data access scheduling for reducing disk energy consumption," in *Proc. IEEE Int. Conf. Distrib. Comput. Syst.*, 2012, pp. 281–282.

[32] J. Lee, B. Yun, and G. Kang, "Reducing peak power consumption inmulti-core systems without violatingreal-time constraints," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 4, pp. 1024–1033, Apr. 2014.

[33] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, "Power attack defense: Securing battery-backed data centers," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2016, pp. 493–505.

[34] W. Zheng, K. Ma, and X. Wang, "TE-shave: Reducing data center capital and operating expenses with thermal energy storage," *IEEE Trans. Comput.*, vol. 64, no. 11. pp. 3278–3292, Nov. 2015.

[35] Y. Guo and Y. Fang, "Electricity cost saving strategy in data centers by using energy storage," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no 6, pp. 1149–1160, Jun. 2013.

[36] L. Liu, et al., "RE-UPS: An adaptive distributed energy storage system for dynamically managing solar energy in green datacenters," *J. Supercomputing*, vol. 72, no. 1, pp. 295–316, 2015.

[37] Y. Kuroda, A. Akai, T. Kato, and Y. Kudo, "High-efficiency power supply system for server machines in data center," in *Proc. Int. Conf. High Performance Comput. Simul.*, 2013, pp. 172–177.

[38] L. Narayanan, D. Wang, A. Mamun, A. Sivasubramaniam, and H. Fathy, "Should we dual-purpose energy storage in datacenters for power backup and demand response?" in *Proc. Workshop Power-Aware Comput. Syst.*, 2014, pp. 1–5.

[39] APC White Paper: Comparing UPS System Design Configurations, 2008.

[40] H. Kim and K. Shin, "Scheduling of battery charge, discharge, and rest," in *Proc. Real-Time Syst. Symp.*, 2009, pp. 13–22.

[41] M. Nispel, "Important considerations when reducing the runtimes of VRLA UPS batteries," 2011. [Online]. Available: http://www.battcon.com/.

[42] R. Lizotte, Data Center VRLA Battery End-of-Life Recycling Procedures, APC White Paper #36.

[43] J. Garche, "Corrosion of lead and lead alloys: Influence of the active mass and of the polarization conditions," *J. Power Sources*, vol. 53, no. 1, pp. 85–92, 1995.

[44] K. Constanti, A. Hollenkamp, M. Koop, and K. McGregor, "Physical change in positive-plate material - an underrated contributor to premature capacity loss," *J. Power Sources*, vol. 55, no. 2, pp. 269–275, 1995.

[45] A. Jossen, J. Garche, and U. Sauer, "Operation conditions of batteries in PV applications," *Solar Energy*, vol. 76, no. 6, pp. 759–769, 2004.

[46] V. Svoboda, et al., "Operating conditions of batteries in off-grid renewable energy systems," *Solar Energy*, vol. 81, no. 11, pp. 1409–1425, 2007.

[47] U. Sauer, et al., "Analysis of the performance parameters of lead/acid batteries in photovoltaic systems," *J. Power Sources*, vol. 64, no. 1, pp. 197–201, 1997.

[48] K. Yabuta, T. Matsushita, and T. Tsujikawa, "Examination of the cycle life of valve regulated lead acid batteries," in *Proc. Int. Telecommun. Energy Conf.*, 2007, pp. 97–101.

[49] H. Bindner, T. Cronin, P. Lundsager, J. Manwell, U. Abdulwahid, and I. Gould, "Lifetime modelling of lead acid batteries," *J. Power Sources*, vol. 53, no. 1, pp. 111–1425, 2005.

[50] D. Wang, C. Ren, and A. Sivasubramaniam, "Virtualizing power distribution in datacenters," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2013, pp. 595–606.

[51] Y. Li, et al., "SizeCap: Coordinating energy storage sizing and power capping for fuel cell powered Da-ta centers," in *Proc. Int. Symp. High-Performance Comput. Archit.*, 2016.

[52] W. Zheng, K. Ma, and X. Wang, "Hybrid energy storage with supercapacitor for cost-efficient data center power shaving and capping," *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 4, pp. 1105–1118, Apr. 2017.

[53] H. Chen, T. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding, "Progress in electrical energy storage system: A critical review," *Natural Sci.*, vol. 19, no. 3, pp. 291–312, 2009.

[54] D. Rastler, "Electricity ENERGY STORAGE TECHNOLOGY OPTIONS: A White Paper Primer on Applications, Costs, and Benefits," Electric Power Research Institute, 2010.

[55] S. Schoenung and W. Hassenzahl, "Long-Term versus short-term energy storage technologies analysis: A life-cycle cost study," Sandia National Laboratories, New Mexico; Livermore, CL, USA, 2003.

[56] (2014). [Online]. Available: http://sine.ni.com/nips/cds/view/p/lang/en/nid/1865

[57] (2014). [Online]. Available: http://sine.ni.com/nips/cds/view/p/lang/en/nid/14132

[58] (2014). [Online]. Available: http://www.u-rd.com/technical/tech2.html

[59] (2014). [Online]. Available: http://www.wti.com/c-34-high-amp-c19-switched-pdu.aspx

[60] (2014). [Online]. Available: http://www.ni.com/labview/

[61] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang, "The HiBench benchmark suite: Characterization of the MapReduce-based data analysis," in *Proc. Data Eng. Workshops Int. Conf. Data Eng.*, 2010, pp. 209–228.
[62] The CloudSuite 2.0. (2014). [Online]. Available: http://parsa.epfl.ch/cloudsuite
[63] L. Wang and W. Chow, "Solar radiation model," *Appl. Energy*, vol. 69, no. 1, pp. 191–224, 2011.
[64] S. Govindan, D. Wang, L. Chen, A. Sivasubramaniam, and B. Urgaonkar, "Towards Realizing a Low Cost and Highly Available Datacenter Power Infrastructure," in *Proc. Workshop Power-Aware Comput. Syst.*, 2011, Art. no. 7.
[65] L. Liu, C. Li, H. Sun, Y. Hu, J. Gu, and T. Li, "BAAT: Towards dynamically managing battery aging in green datacenters," in *Proc. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2015, pp. 307–318.
[66] D. Baert and A. Vervaet, "Lead-acid battery model for the derivation of peukert's law," *Electrochimica Acta*, vol. 44, pp. 3491–3504, 1999.
[67] Lifecycle Costing for Data Centers: Determining the True Costs of Data Center Cooling, White Paper, Emerson Network Power, 2003.
[68] B. Gregory, "Ultra-Capacitor Sizing and Packaging for Cost Effective," in White Paper, IOXUS Inc., 2013.

Longjun Liu received the PhD degree in control engineering from Xi'an Jiaotong University, Xi'an, China, in 2015. He was a visiting PhD student in the Department of Electrical and Computer Engineering, University of Florida, Gainesville, Florida, from 2012 to 2014. Currently, he is an assistant professor with Xi'an Jiaotong University. His research interests include computer architecture, power-efficient computing systems, and energy management for green dataenters. He is a member of the IEEE.

Hongbin Sun (M'11) received BS and PhD degrees in electrical engineering from Xi'an Jiaotong University, Xi'an, China, in 2003 and 2009, respectively. He was a visiting PhD student in electrical, computer and systems engineering department at Rensselaer Polytechnic Institute, Troy, New York, from 2007 to 2008. Currently he is an associate professor with Xi'an Jiaotong University. His research interests include memory hierarchy design and VLSI architecture for digital video processing. He is a member of the IEEE.

Chao Li received the PhD degree from the University of Florida, in 2014. He is an assistant professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University. His research focuses on computer architecture, energy-efficient systems, and green computing. He is a member of the IEEE.

Tao Li received the PhD degree in computer engineering from the University of Texas, Austin, Texas. He is a professor in the Department of Electrical and Computer Engineering, University of Florida, Florida. His current research interests include computer architecture, microprocessor, memory storage system design, virtualization technologies, energy efficient/sustainable dependable datacenters, cloud/big data computing platforms, the impacts of emerging technologies/applications on computing, and evaluation of computer systems.

Jingmin Xin (S'92–M'96–SM'06) received the BE degree in information and control engineering from Xi'an Jiaotong University, Xi'an, China, in 1988 and the MS and PhD degrees in electrical engineering from Keio University, Yokohama, Japan, in 1993 and 1996, respectively. Since 2007, he has been a Professorat Xi'an Jiaotong University. His research interests include in the areas of adaptive filtering, statistical and array signal processing, system identification, and pattern recognition. He is a senior member of the IEEE.

Nanning Zheng (SM'93–F'06) received the graduate degree from the Department of Electrical Engineering, Xi'an Jiaotong University, Xi'an, China, in 1975, and received the MS degree in information and control engineering from Xi'an Jiaotong University, in 1981 and the PhD degree in electrical engineering from Keio University, Yokohama, Japan, in 1985. He jointed Xi'an Jiaotong University in 1975, and is currently a professor and the director of the Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University. His research interests include computer vision, pattern recognition and image processing, and hardware implementation of intelligent systems. He is a fellow of the IEEE.