

BAAT: Towards Dynamically Managing Battery Aging in Green Datacenters

Longjun Liu^{1,3}, Chao Li², Hongbin Sun^{1*}, Yang Hu³, Juncheng Gu³ and Tao Li³

¹ School of Electrical and Information Engineering, Xi'an Jiaotong University, China
longjun.liu@stu.xjtu.edu.cn
hsun@mail.xjtu.edu.cn

² Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China
lichao@cs.sjtu.edu.cn

³ Department of Electrical and Computer Engineering, University of Florida, USA
{huyang.ece, gujuncheng}@ufl.edu, taoli@ece.ufl.edu

Abstract— Energy storage devices (batteries) have shown great promise in eliminating supply/demand power mismatch and reducing energy/power cost in green datacenters. These important components progressively age 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 issue of batteries or simply use ad-hoc discharge capping to extend their lifetime. To fill this critical void, we thoroughly investigate battery aging on a heavily instrumented prototype over an observation period of six months. We propose battery anti-aging treatment (BAAT), a novel framework for hiding, reducing, and planning the battery aging effects. We show that BAAT can extend battery lifetime by 69%. It enables datacenters to maximally utilize energy storage resources to enhance availability and boost performance. Moreover, it reduces 26% battery cost and allows datacenters to economically scale in the big data era.

Keywords— Battery Aging, Analysis, Green Datacenters, Power Management, System Prototype

I. INTRODUCTION

The growing adoption of massive distributed batteries at the server rack level could fundamentally transform the way we manage datacenter power and energy [1-3]. For example, distributed batteries are important energy buffers in several topical studies that exploit renewable energy resources for capping datacenter carbon footprint [4, 5]. They are also the key enablers that allow datacenters to smooth out load power peaks to greatly reduce total cost [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].

With the dramatic increase in their quantity, size, and utilization frequency in datacenters, batteries start to become the major design bottleneck. Most of the prior works in this context focus on managing server power demand [4-7, 10-13], whereas the battery aging effect in datacenters has been largely ignored. However, batteries progressively and inevitably age. Without smart control, datacenter operators may incur very high maintenance and replacement cost. Even worse, aged batteries often lose their effective energy storage capacity and may fail to handle emergency events, which poses serious threat to server availability. Therefore, if the aging issues are overlooked, the benefits of carbon emission capping and cost reduction brought by batteries can be significantly compromised.

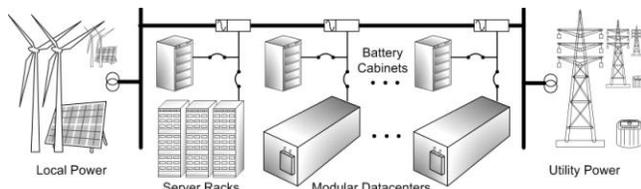


Fig. 1. Distributed energy storage devices in emerging green datacenters

In the past, battery aging or wear-out issues have been well studied by the electrochemistry research community with the aim of improving a battery's internal design [14-19]. When batteries are massively deployed in datacenters, new aging-aware power management strategies at computer architecture and system levels are highly desired.

One goal of this paper is to establish an understanding of the battery aging mechanism from a system designer's perspective. As shown in Figure 1, we investigate emerging green datacenters that rely on energy storage devices 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 aging and develop a holistic management strategy that can capture various anti-aging opportunities.

In this study 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 six months, we continuously monitor detailed battery usage events and investigate a diverse set of factors that may induce aging.

We propose *battery anti-aging treatment (BAAT)*, an aging-aware power management framework for better coordinating massive amount of batteries in green datacenters. The novelty of BAAT is two-folded: 1) it is able to sense the synergistic effect of multiple aging factors to assess battery aging; 2) it can further leverage a multifaceted approach to manage battery aging at the system/architecture level.

Specifically, BAAT offers three key benefits. First, it hides battery aging variation within datacenter. When datacenter adds new workloads or consolidates existing workloads, BAAT can

*Corresponding author: Hongbin Sun.

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 leverages workload migration and power capping mechanisms at critical points to avoid aggressively discharging batteries. Third, BAAT can further 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 for better datacenter productivity.

To the best of our knowledge, this is the first extensive study of battery aging in emerging green datacenters. The challenges for BAAT 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 fully unleash the performance potential of green datacenter via planned battery aging. This paper solves these obstacles and makes the following contributions:

- We investigate a diverse set of battery aging factors and formulate five metrics (i.e. *Ah throughput*, *charge factor*, *partial charging*, *deep discharge time*, and *discharge rate*). These metrics provide an effective way to dynamically evaluate battery aging at the system level.
- We propose BAAT, the first battery aging management framework tailored for green datacenters. BAAT tackles battery aging, an important, yet so-far overlooked issues in green datacenters. It features a multifaceted approach to managing battery aging.
- We build up a prototype and describe the details of our implementation. We analyze battery aging based on real deployment experience over six months. We evaluate the effectiveness of BAAT and show that it can improve battery lifetime by 69%, reduce battery depreciation cost by 26%, and boost performance by 28%.

The rest of this paper is organized as follows. Section 2 introduces distributed batteries in green datacenters and their aging mechanisms. Section 3 formulates the battery aging impact factors. Section 4 elaborates BAAT architecture and aging management schemes. Section 5 describes experimental methodology. Section 6 presents evaluation results. Section 7 discusses related work and Section 8 concludes this paper.

II. BACKGROUND

In this section, we first introduce battery usage scenarios in the state-of-the-art datacenter design. We then provide an overview of the battery aging mechanisms.

A. Battery Deployment in Datacenters

Many recent designs start to shift battery system from the datacenter level to the server rack level, as shown in Figure 2. 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

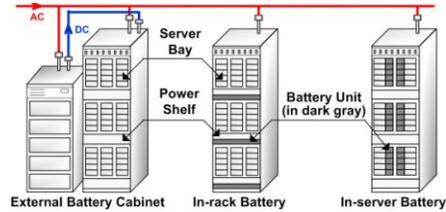


Fig 2. Different methods for integrating battery with computer servers

TABLE I. DIFFERENT BATTERY USAGE SCENARIOS IN DATACENTERS

Usage Objective	Usage Frequency	Aging Speed	Aging Variation
Power Backup	Rarely	Light	Small
Demand Response	Occasionally	Medium	Medium
Power Smoothing	Cyclically	Severe	Large

batteries [1]. Several researchers at Hitachi have compared different battery integration methods in [20].

Meanwhile, datacenters are aggressively exploiting battery for reducing cost and improving sustainability. This is different from conventional designs that mainly use battery to provide power backup upon a grid outage. For example, batteries allow datacenters to shift load power, referred to as demand response [21]. By re-shaping server power demand, primarily shaving power peaks, there are considerable opportunities for reducing operational cost (deferring loads to non-peak tariffs) and capital cost (under-provision the power infrastructure).

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, inexpensive lead-acid batteries are used to smooth the variable renewable power supply [4, 5].

As new architecture emerges and usage pattern evolves, batteries have become the Achilles’s heel in green datacenters. In Table-1 we briefly summarize different battery usage scenarios. In contrast to conventional battery systems used in emergency (rarely occurred and only lasts a few minutes [22]), emerging battery provisioning schemes result in a much higher charging/discharging rate that hastens battery aging. Thus, it is necessary to understand battery aging mechanisms and intelligently manage the battery aging issue in datacenters.

B. Battery Aging Mechanisms

Battery aging refers to the processes of gradual deterioration of inner materials and the irreversible chemical reactions within the battery [16]. Different types of batteries (e.g., lithium-ion, nickel-cadmium, and lead-acid) have different aging mechanisms. This study primarily focuses on lead-acid batteries, which account for over 97% of industry batteries [23]. 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, electrode polarization and battery temperature [24].

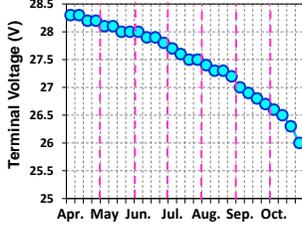


Fig 3. Measured battery voltage drop due to aging over 6 months

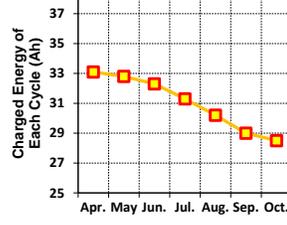


Fig 4. Measured battery capacity drop due to aging over 6 months

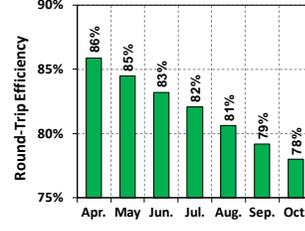


Fig 5. Measured energy efficiency degradation due to aging over 6 months

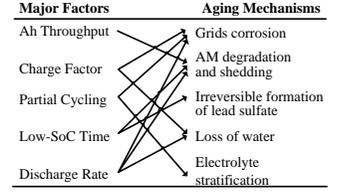


Fig 6. Key factors that correlate with different battery aging mechanisms

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, etc. [25]). The AM shedding represents the active mass is permanently removed from the electrode. The AM degradation/shedding are usually accelerated by a high Ah 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 (PbO_2 and Pb) are transformed into $PbSO_4$ [16]. However, if a battery is not recharged timely, sulphate crystals can grow almost linearly with the solubility of sulphate ions and temperature. In this case, the active masses are irreversibly converted to $PbSO_4$ and can no longer participate in electrochemical reactions [26]. This process is accelerated if a 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. In a VRLA battery, water cannot be re-filled and the aging process is called drying out [27]. Over-charging 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 occur on the battery that is rarely fully recharged and the cells are deeply discharged with very low current [28]. 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. They are often the hidden cause of undesired load shedding (for green datacenters that have stringent power budget) and unnecessary renewable energy feedback (sold at an unprofitable energy price to the utility grid, e.g., 40% of the whole sale price [5]).

As shown in Figure 3, over a continuous operation of six months, the terminal voltage of a battery (fully charged) in our system can be decreased by approximately 9%. 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 emergency handling as the under-voltage battery cannot sustain high-current drawn and may be disconnected from the system [29]. In addition, the effectively stored energy in each charging cycle has also dropped by 14% under aggressive usage (Figure 4). 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% of its

initial capacity [30]. 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 decreased by 8%, as shown in Figure 5.

III. MODELING AND QUANTIFYING BATTERY AGING

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. Figure 6 shows the correlation between these metrics and different aging mechanisms. These metrics quantitatively reflect battery aging progresses and allow us to devise power management policies for handling various battery aging issues at the computer architecture level.

A. 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 [27]:

$$NAT = \frac{Q_{AT}}{CAP_{nom}} = \frac{\int_{t_0}^T I_{bat} dt}{CAP_{nom}} \quad (1)$$

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 life-long 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 [31, 32]. Therefore, prior work has used Ah throughput for predicting battery lifetime under different charge/discharge conditions [4]. 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 [32].

B. Charge Factor (CF)

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

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

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 [27]. 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.

C. 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%), B (79%-60%), C (59%-40%) and D (39%-0%). The probability of the Ah output being in range X is given by:

$$PC_X = \frac{\int_{t_0}^T I_{SoC_X} dt}{\int_{t_0}^T I_{all} dt} \times 100\% \quad (3)$$

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:

$$PC = (PC_{(A)} \times 1 + PC_{(B)} \times 2 + PC_{(C)} \times 3 + PC_{(D)} \times 4) / 4 \quad (4)$$

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 [28]. The higher value of PC will accelerate the battery aging such as corrosion and electrolyte stratification.

D. Deep Discharge Time (DDT)

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

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

In Eq-5, H is the Heaviside Step Function ($H(x)=0$, if $x < 0$; $H(x)=1$, if $x \geq 0$). 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.

E. 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 reduction of the lifetime by 50% [26].

IV. HIDING, SLOWING DOWN, AND PLANNING BATTERY AGING IN A GREEN DATACENTER

This section discusses how to alleviate the deleterious effect of battery aging. We first present the detailed design of our battery anti-aging treatment (BAAT) framework. We then elaborate how to hide aging, slow down aging, and plan the aging speed of battery in green datacenters.

A. BAAT Framework Overview

Figure 7 depicts the schematic diagram of BAAT. Our framework contains a sensor-table based power monitoring architecture and a workload scheduler, which are built on the emerging distributed energy storage architecture [1-3]. BAAT supports two types of distributed energy storage architectures: 1) each server is equipped with a separate battery (similar to that in Google datacenters [1]) and 2) several racks share a pool of batteries (akin to Facebook's Open Rack design [3]).

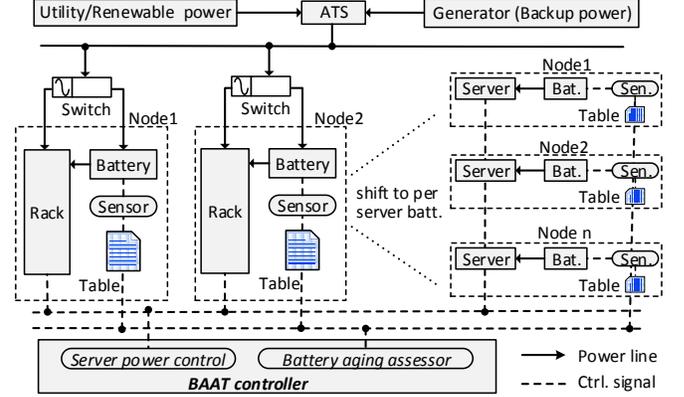


Fig 7. BAAT battery aging management framework. It demonstrates two types of architecture: per-rack integration and per-server integration

Each group of batteries has a power table which records the battery utilization history logs as shown in Table 2. These log data are collected from corresponding sensor of each battery and are sent to BAAT controller. The BAAT controller then calculates various metrics (NAT, CF, PC, DDT, and DR) for evaluating the aging process. The BAAT controller also has the knowledge of the power information of each server through datacenter IT infrastructure such as IPDU or other power meters. It can control the power state of each server (e.g., DVFS, clock throttling, and VM migration and consolidation) by a software driver within each server.

With the architecture and monitored logs, BAAT controller can dynamically and intelligently manage various battery aging issues. We propose three management schemes integrated with the BAAT controller: (1) *load scheduling policy for hiding aging*, (2) *power capping strategy for slowing down aging*, and (3) *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 effectively mitigate battery aging rate and avoid server availability degradation. The planned aging is used when datacenter lifetime and battery aging rate are discrepant and require synchronization.

TABLE II. 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

B. Hiding Aging: Aging-Driven 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 cause 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

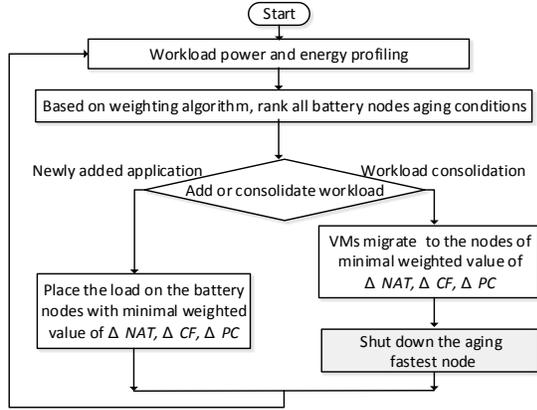


Fig. 8. The flow chart of BAAT aging hiding

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 availability is severely threatened as a battery node that ages faster can more easily cause unexpected server downtime. Therefore, it is wise to carefully balance battery aging variation issue.

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.

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, energy 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 while considering the aging issue. Many datacenter applications can provide coarse granularity power profile, such as long time running services (web searching, memcached, etc.) and/or periodic/repetitive workloads (e.g., web crawling) [33]. In our study, using power profiling contributes to better estimation of its impact on battery aging as well as aging handling.

b) Battery Aging Consideration

As mentioned in the last section, BAAT 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 runtime logs. BAAT 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 [31, 32]. 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/server for the given

workloads/applications. 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.

TABLE III. RELATION BETWEEN POWER DEMANDS AND AGING FACTORS

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

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% of the peak power. Otherwise, we define the power demand as “Small”. Similarly, we classify the energy demand as either “More” or “Less”. They implies 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 CF and PC, whereas the AT is more likely to decrease when the energy request becomes “Less”.

We use “High”, “Low” and “Medium” to 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 are set as 50% in the “High” scenario, 30% in the “Middle” scenario and 20% in the “Low” scenario respectively. Our extensive training and experiments shows that these weighting factors are fairly effective when evaluating battery aging. 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 Figure 8.

In summary, the battery aging-driven workload scheduling can effectively balance the battery aging processes across the datacenters. As the inaccurate power profiling may lead to wrong load placement or consolidation sometimes, we can further leverage aging slowing down technique to remedy it.

C. 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 overall battery lifetime, we need to slow down battery aging at appropriate timestamp.

2) Implementation Details

Our algorithm periodically checks two metrics: deep discharge time (DDT) and high discharge rate (DR). If a battery

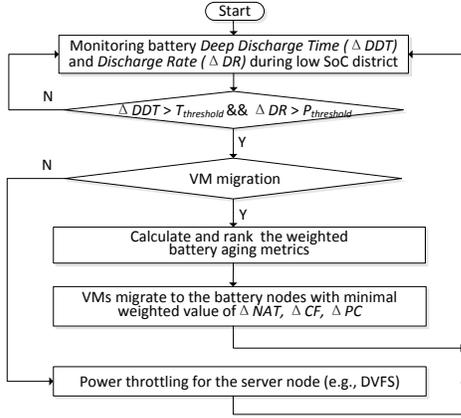


Fig. 9. The flow chart of BAAT aging slowdown. The $P_{threshold}$ is the maximal current that can sustain discharge for 2 minutes ($T_{threshold}$)

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 budget, battery voltage may reduce to the cut-out line and lead to server downtime. Therefore, we set thresholds for DDT and DR to avoid unplanned battery cut out.

When the SoC of battery drops below 40%, BAAT periodically checks DDT and DR to see if they reach a preset thresholds. 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 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 elsewhere in the datacenter (e.g., CPU/Memory/Disk), 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 is depicted in Figure 9.

D. 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 to “shift” some performance from the unused portion of the battery’s lifetime to the used time, which refers to planned aging.

In Figure 10 we show the cycle life data from different manufacturers (Hoppecke Battery, Trojan Battery and UPG Battery). It shows that the battery cycle life decreases by 50% if it is frequently discharged at a DoD above 50%. 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

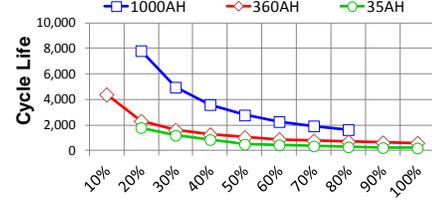


Fig. 10. Battery cycle life under varying depth of discharge (DoD)

capacity and the expected battery life cycles to calculate the needed DoD as show below:

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

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 $(1 - DoD_{goal})$. Also, the planned aging does not exclude hiding aging. The hiding aging technique still places and consolidates workload across datacenter to balance all the battery node aging based on the planed aging rate. The planned aging technique can effectively help us synchronize the battery aging rate with the end-of-life of datacenter infrastructures.

V. EXPERIMENTAL METHODOLOGY

A. System Prototype

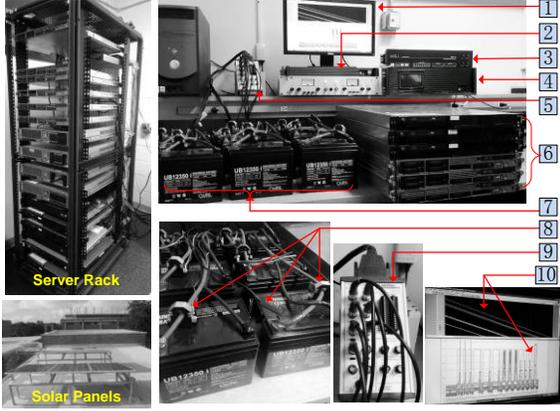
We have built a heavily instrumented energy storage system. It enables us to analyze various battery aging issues and evaluate the proposed techniques. We synergistically integrate our energy storage system with a scaled-down green datacenter prototype also built from scratch. As shown in Figure 11,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

1) **Battery Module.** Our system employs emerging distributed energy storage architecture. Each server is equipped with individual battery unit. We use multiple new sealed lead-acid batteries with the same capacity and voltage rating in our experiment. The same initial condition of the batteries allows us to conduct comparable experiments with different aging management policies.

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 [35] plugged-in server motherboard via PCI-e interface is used to collect the data from sensors to computer hard disk. Based on the two devices, the sensor data can be viewed by LabVIEW [38] in real time.

3) **Control Module.** Our control server is a customized i7 low power server, which runs our BAAT 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 [37]. 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.



Battery array	
(7)	Twelve 12V 35Ah sealed lead-acid batteries
Sensor devices	
(5)(9)	Shielded connector block (NI BNC-2110[34]) for sensor signal transmitting from front sensors to sensor data acquisition card (NI PCI-6221[35])
(8)	Front sensors for batteries for collecting battery voltage, current and temperature [36]
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[37], 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
Computing nodes	
(6)	Three IBM X series 330 and three HP ProLiant servers
Display	
(1)	Displayed information includes data captured by sensors, system log trace, and various aging metrics calculated by the control server in real time
(10)	

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

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. 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. The display module can visualize the data captured by sensors via LabVIEW and the aging impact factors calculated by control server in real time.

B. Workload Deployment

We deploy six datacenter workloads. Three of them come from Hibench [39]: Nutch Indexing; K-Means Clustering, and Word Count. 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 [40]: Software Testing, Web Serving, and Data Analytic. For example, Software Testing is a resource-hungry and time-consuming application that allows us to stress our servers and distributed batteries.

TABLE IV. THE COMPARISON OF FOUR BATTERY MANAGEMENT POLICIES

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

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.

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

VI. EVALUATION RESULTS

In this section, we evaluate the impact of various power management schemes on battery aging. We first compare the BAAT with three baseline power management policies as shown in Table 4. Among those, *e-Buffer* 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 scheme. *BAAT-s* only focuses on slowing down the battery aging processes, while *BAAT-h* mainly emphasizes battery aging hiding.

A. System Runtime Profiling

We first profile the system runtime of our prototype across different solar generation scenarios by analyzing the generated logs and collected system runtime traces. As expected, due to the intermittent solar power budget and different server power demands, the usage frequency of the six battery pack varies significantly, as shown in Figure 12-(a), which leads to different battery aging rates.

To quantify the aging processes, we present the variation of several aging-related metrics of one battery node, as shown in Figures 12(e) to 12(k). We present the values of NAT, CF, and PC under different weather conditions, as shown in Figures 12(b) to 12(d). The total energy budget for Sunny, Cloudy, and Rainy is 8 kWh, 6kWh, and 3kWh, respectively.

It is clear that the battery nodes yield less Ah-throughput in sunny day than the other two weather conditions. This is because the solar energy can afford most server power demands and the batteries are rarely used. 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.

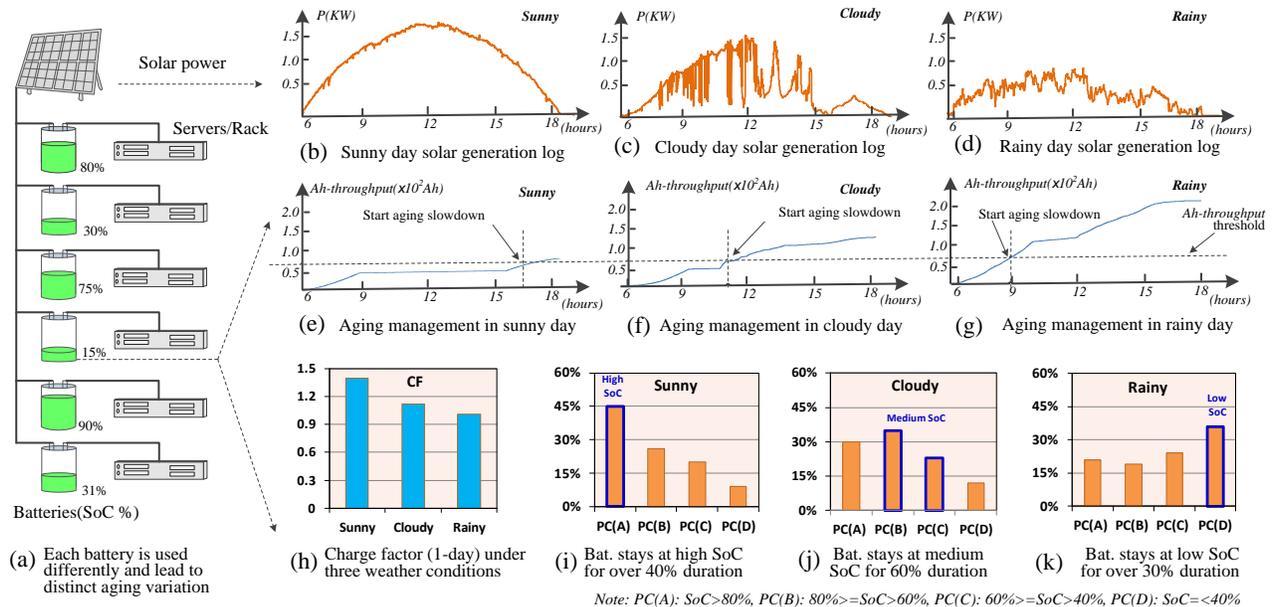


Fig. 12. The variation of different aging metrics collected from our solar energy powered green datacenter prototype under different weather conditions

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 threshold. The slowdown time varies in different weathers, as marked in Figures 12(e) to 12(g). Based on the logs of different aging metrics, our BAAT power management framework can combine load power profiling information and weighted aging ranking to effectively hide and slow down battery aging.

B. Aging Progress Comparison

In this subsection, we present the impact of different power management schemes on battery aging progresses by comparing different aging metrics. Specially, we compare our aging-aware power management algorithm BAAT with three baseline algorithms. We select two typical weather conditions (sunny and cloudy day) and “young” and “old” battery aging stages (new and aged batteries) to evaluate the four management algorithms.

In April, after the setup of our prototype system, we start by using the new batteries to conduct our aging experiments. We deploy and iteratively run the workloads hosted in virtual machines on our computing server nodes. Each power management scheme is run one day and 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 (e.g., NAT, CF, PC, DDT, and DR) of six battery nodes. Note that, as the solar power varies every day, to obtain the same solar generation condition for comparing the four power management experiments, we run our experiments multiple days and record all the logs. By analyzing the solar power state logs (e.g., per day maximal, minimal, and average values, as well as the total energy), we are able to find the most similar solar generation scenarios across the multi-groups of experiment logs and compare the results of four schemes. Figures 13-(a) and 13-(b) show the impact of the four power management algorithms on

the battery aging when batteries are “young”. Note that we select the worst battery node that has the most Ah-throughput in each power management policies.

In October, when batteries became “old”, we conduct the experiments with same methods and the results are showed in Figures 13-(c) and 13-(d). Note that, during April to October, we regularly use the batteries and make them gradually and synchronously aging in our prototype. Our system collects various log data during runtime.

By comparing the variation of the aging metrics of the four power management schemes, we can conclude the following results. (1) Battery ages faster in the harsh usage conditions, for example, the Ah-throughput of *e-Buffer* algorithm in cloudy is increased by 35% than the sunny day on average. (2) The “old” battery should demands more attentions to take care of the aging issue than the “young” battery. Especially for the *e-Buffer* power management scheme, we can see that the PC value of “old” battery is reduced by 17% than the “young” battery on average, which means the battery node stays low SoC and is prone to wear out. (3) By leveraging our BAAT battery aging aware power management algorithm, the battery aging progresses can be effectively hidden and slowed down.

The *e-Buffer* 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 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 6.3). The BAAT-h power 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 aging metrics) and the

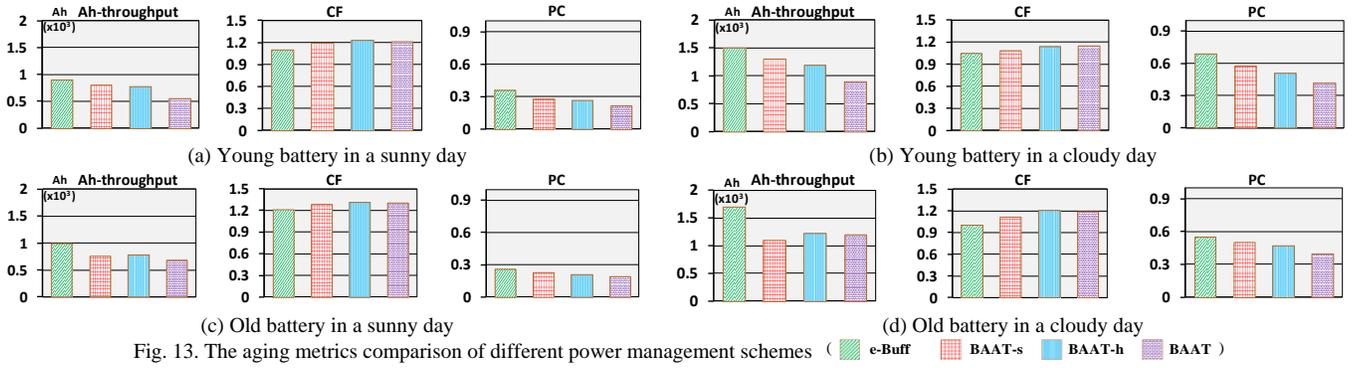


Fig. 13. The aging metrics comparison of different power management schemes (e-Buffer BAAT-s BAAT-h BAAT)

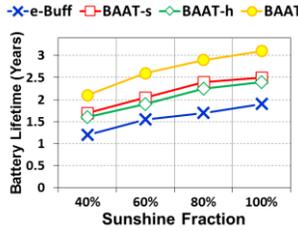


Fig. 14. Battery life under different green energy

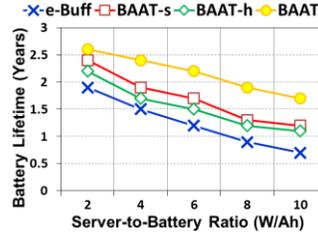


Fig. 15. Battery life under different capacity

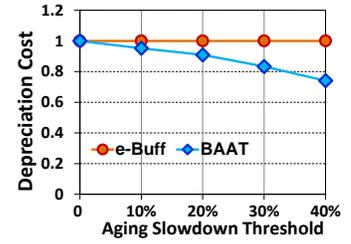


Fig. 16. BAAT reduces annual depreciation cost

migration is unaware the aging state of other battery nodes, which make the migration become random and low efficiency. Our BAAT algorithm is a holistic battery aging-aware management policy which calculates and ranks the weighted aging metrics of all battery nodes. It also refers the work-load power profiling information to manage battery aging. Based on these techniques, our BAAT can dynamically slow down 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 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 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 can reduce battery aging speed in the worst case by 38%.

C. Lifetime Improvement

Since the battery cycling frequency is greatly affected by renewable energy availability and load power demand, we evaluate battery lifetime improvements under different solar energy potentials and server capacities (Figures 14 and 15).

We first consider geographic locations that have different solar energy availabilities (represented by sunshine fraction, the percentage of time when sunshine is recorded [41]). 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 power output is high. In fact, sometimes the entire load can be directly powered by solar panels without using the stored green energy. BAAT

could extend battery life by 69% on average, compared to e-Buffer. The lifetime improvement for BAAT-s and BAAT-h is 37% and 29%, respectively. Our results show that aging slowdown has a larger impact on battery lifetime, compared to aging balancing.

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 35% (Figure 15). This is mainly because heavy server loading is more likely to create power spikes, which cause deep battery discharge and high discharge rate. (2) The optimization effectiveness of BAAT 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 (compared to e-Buffer) grows from 37% to 1.4X. This indicates the benefits of our battery management 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 Figure 15, doubling the installed battery can cut the battery-to-server ratio by half, but may result in less than 30% 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.

D. Cost Benefits of BAAT

Increasing battery lifetime can greatly increase the return on investment (ROI) due to the reduced battery depreciation cost. By varying the threshold in 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-Buffer, BAAT can achieve 26% cost reduction, as

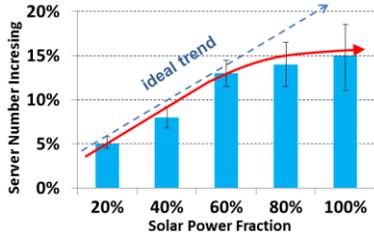


Fig. 17. Green datacenter can economically trade off battery lifetime for server capacity

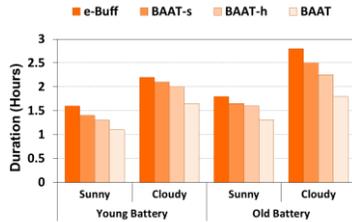


Fig. 18. Low-SoC duration comparison. BAAT effectively reduces deep discharge

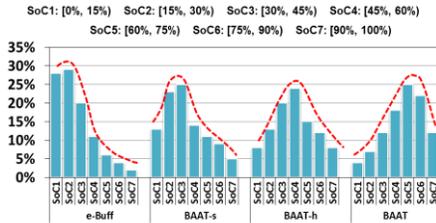


Fig. 19. Distribution of battery SoC under different battery power management schemes

shown in Figure 16. Note that aggressively applying the aging slowdown algorithm is not wise since it may cause unnecessary performance degradation.

An interesting observation is that BAAT 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 life can actually be used to purchase more servers. Figure 17 evaluates 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 15% 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.

E. Datacenter Availability

This section focuses on the availability of server system from the perspective of battery aging management. The key aging factor that directly correlates with server availability is deep discharge time (DDT). Prior work has shown that datacenter must leave 2 minutes of reserve capacity in UPS battery for high availability [42]. A low SoC means less reserved energy, which is dangerous when green servers plan to draw large amount of power from the battery.

We monitor and collect the low-SoC duration of different power management schemes from the experiment logs, as shown in Figure 18. As can be seen, the e-Buffer 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 can dynamically slow down and balance the SoC across all the battery nodes, which can effectively eliminate SPOF. Our results show that BAAT could increase battery availability by 47% based on the statistics of low-SoC duration of the worst-case battery node.

In Figure 19 we further evaluate the distribution of deep discharging over 6 months. It is clear that e-Buffer tends to create low-SoC batteries, whereas BAAT can shift the most likely SoC region towards 90%-100%. Therefore, BAAT increases the resiliency and emergency handling capability.

F. Workload Performance

This section presents the impact of different battery power management schemes on workload performance. The experimental method is the same as Section 6.B. We evaluate the total compute throughput of the four power management schemes in one day, as shown in Figure 20.

Intuitively, the e-Buffer 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-Buffer. The BAAT-s power management policy always leverage 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 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 algorithm can dynamically slow down battery aging and schedule workloads according. Compared to e-Buffer in the worst case (cloudy and old battery), BAAT can improve the performance by 28%.

G. Benefits of Planned Aging

Another distinctive feature of BAAT 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 [43] and 10~15 years [44], respectively. Without aging planning, it is highly likely that datacenter operators have to discard the latest replaced batteries before they wear out.

BAAT 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 Figure 21, the performance improvement is not linear with DoD variation. When the DoD increases from 40% to 60%, the performance improvement is more visible than when the DoD increases from 70% to 90%. This is because compared with a normal DoD range of 20%~40%, the latter will yield too low battery SoC, which leads to reduced battery lifetime.

We consider different durations of the expected battery service life (i.e., from the installation time of battery to the end-of-life of the data center). By aggressively utilize the battery, one can improve datacenter productivity by up to 33% (Figure 22), compared to e-Buffer. Our results show that the performance benefits decrease when the newly replaced battery is too close to the datacenter/server end-of-life. This is because there is an upper bound of battery discharge (i.e., over 90% DoD) and the performance improvement is heavily restricted by the available operating duration. On the other hand, the benefits of aging planning also decrease when the battery installation time is far from the expected datacenter end-of-life. This is mainly because

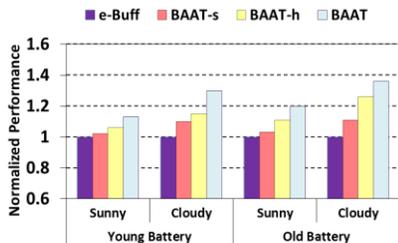


Fig. 20. Performance improvement

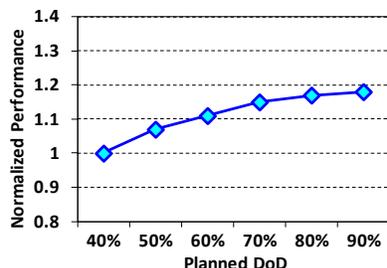


Fig. 21. Performance impact of planned DoD

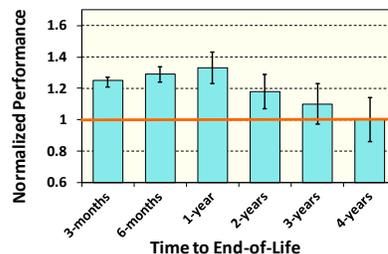


Fig. 22. Performance benefits of planning aging

the unused portion of the battery’s lifetime that can be shifted is too small.

VII. RELATED WORK

To the best of our knowledge, this is the first extensive analysis of battery aging and failure in green datacenters. The relevant prior work are summarized as bellow.

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 evaluation [14, 15]. The main focus is to leverage mathematical formula to model battery behavior. (2) Battery aging mechanism studies [16, 17]. Among those, [16] presents an overview of lead-acid battery aging mechanisms. [17] focuses on the sulfation aging in battery. (3) Battery lifetime prediction model [18, 19]. To our knowledge, none of prior work addresses the battery aging issues at datacenter level.

Battery management in micro-systems: As the primary power source in micro-systems, battery usage and lifetime management have been studied in light of mobile/embedded environment [45-48] and hybrid electric vehicles (HEV) system [49]. Prior work mainly focuses on battery-aware energy management and task scheduling. For example, [48] presents a stochastic model of battery and a framework for estimating the battery lifetime in mobile embedded systems. [49] formulates an optimal control problem to minimize the fuel consumption as well as battery aging in HEV system. Most of their works focus on single and micro system control. In contract, our work explores the aging issues of battery systems in green datacenters.

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, 50] and emerging renewable energy powered datacenters [4, 5, 12, 13, 51, 52]. 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). These studies rarely manage battery aging dynamically. Similarly, while many recent studies [4] also employ battery to buffer renewable power in datacenters, there has been no work that explores battery aging-aware power management in the context of green datacenters.

Our work distinguishes itself from other studies in four aspects: (1) It investigates emerging battery aging issues at the datacenter level. (2) It comprehensively analyzes battery aging mechanisms from the perspective of architecture and system designers (rather than a electrochemical point of view). (3) It proposes a novel battery aging management framework that can

jointly balance, slow down and plan the aging process. (4) It builds a real prototype to evaluate the benefit and efficiency of battery aging-aware power management schemes.

VIII. CONCLUSIONS

In this study we explores battery aging issues on a scaled-down prototype over six months and proposes a novel power 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.

We propose BAAT, a novel battery aging management framework which can jointly hide, slow down and plan battery aging. 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 can leads to 38% battery aging rate reduction even in the worst case and can improve battery lifetime by 69% on average. Meanwhile, BAAT could reduce the performance overhead caused by inefficient battery management, thereby improving workload throughput by 28%. Our design allows datacenter to reduce 26% battery annual depreciation cost.

ACKNOWLEDGEMENT

This work is supported in part by NSF grants 1423090, 1320100,1117261, 0937869, 0916384, 0845721(CAREER), 0834288, 0811611, 0720476, by SRC grants 2008-HJ-1798, 2007-RJ-1651G, by Microsoft Research Trustworthy Computing, Safe and Scalable Multi-core Computing Awards, by three IBM Faculty Awards and by the Program of Introducing Talents of Discipline to University grant B13043. Chao Li is also supported in part by a Facebook Fellowship and a SJTU-MSRA Faculty Award.

REFERENCES

- [1] <http://www.cnet.com/news/google-unlocks-once-secret-server-10209580/>
- [2] <http://www.datacenterknowledge.com/archives/2011/04/25/microsoft-reveals-its-specialty-servers-racks/>
- [3] Facebook. Hacking Conventional Computing Infrastructure. <http://opencompute.org/>, 2011.
- [4] C. Li, Y. Hu, R. Zhou, M. Liu, L. Liu, J. Yuan and T. Li, “Enabling Datacenter Servers to Scale Out Economically and Sustainably,” in *International Symposium on Microarchitecture (MICRO)*, 2013.
- [5] I. Goiri, W. Katsak, K. Le, T. Nguyen and R. Bianchini, “Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy,” in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2013.
- [6] V. Kontorinis, L. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, T. Rosing and D. Tullsen, “Managing Distributed UPS Energy

- for Effective Power Capping in Data Centers,” in *International Symposium on Computer Architecture (ISCA)*, 2012.
- [7] S. Govindan, A. Sivasubramaniam and B. Urgaonkar, “Benefits and Limitations of Tapping into Stored Energy for Datacenters,” in *International Symposium on Computer Architecture (ISCA)*, 2011.
- [8] Energy Storage for Wind and Solar Integration, Navigant Research, Technical Report, 2013.
- [9] Forecasting Energy Demand: Global Market Profile Report, DatacenterDynamics Industry Census, 2011.
- [10] S. Govindan, D. Wang, A. Sivasubramaniam and B. Urgaonkar, “Leveraging Stored Energy for Handling Power Emergencies in Aggressively Provisioned Datacenters,” in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2012.
- [11] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, “Energy Storage in Datacenters: What, Where, and How Much,” in *SIGMETRICS Performance Evaluation Review*, 2012.
- [12] N. Sharma, S. Barker, D. Irwin and P. Shenoy, “Blink: Managing Server Clusters on Intermittent Power,” in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2011.
- [13] C. Li, W. Zhang, C. Cho and T. Li, “SolarCore: Solar Energy Driven Multi-core Architecture Power Management,” in *International Symposium on High-Performance Computer Architecture (HPCA)*, 2011.
- [14] S. Buller, M. Thele, E. Karden and R.W. Doncker, “Impedance-based non-linear dynamic battery modeling for automotive applications,” in *Journal of Power Sources*, 2003, 113(2), pp. 422-430.
- [15] Y. Sun, H. Jou and J. Wu “Aging Estimation Method for Lead-Acid Battery,” in *IEEE Transaction on Energy Conversion*, 2011, 26(1), pp. 264-271.
- [16] P. Ruetschi, “Aging Mechanisms and Service Life of Lead-Acid Batteries,” in *Journal of Power Sources*, 2004, 127(1), pp. 33-44.
- [17] H. Catherino, F. Feres and F. Trinidad, “Sulfation in Lead-Acid Batteries,” in *Journal of Power Sources*, 2004, 129(1), pp. 113-120.
- [18] D. Sauer, and H. Wenzl, “Comparison of Different Approaches for Lifetime Prediction of Electrochemical Systems Using Lead-Acid Batteries as Example,” in *Journal of Power Sources*, 2008, 176(2), pp. 534-546.
- [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,” in *Journal of power sources*, 2000, 87(1), pp. 39-56.
- [20] Y. Kuroda, A. Akai, T. Kato, and Y. Kudo, “High-Efficiency Power Supply System for Server Machines in Data Center,” in *International Conference on High Performance Computing and Simulation (HPCS)*, 2013.
- [21] 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 *Workshop on Power-Aware Computing and Systems (HOTPOWER)*, 2014.
- [22] APC White Paper: Comparing UPS System Design Configurations, 2008.
- [23] R. Lizotte, Data Center VRLA Battery End-of-Life Recycling Procedures, APC White Paper #36.
- [24] J. Garche, “Corrosion of Lead and Lead Alloys: Influence of the Active Mass and of the Polarization Conditions,” in *Journal of Power Sources*, 1995, 53(1), pp. 85-92.
- [25] K. Constanti, A. Hollenkamp, M. Koop and K. McGregor, “Physical Change in Positive-Plate Material - an Underrated Contributor to Premature Capacity Loss,” in *Journal of power sources*, 1995, 55(2), pp. 269-275.
- [26] A. Jossen, J. Garche and U. Sauer, “Operation Conditions of Batteries in PV Applications,” in *Solar Energy*, 2004, 76(6), pp. 759-769.
- [27] V. Svoboda et.al., “Operating Conditions of Batteries in Off-Grid Renewable Energy Systems,” in *Solar Energy*, 2007, 81(11), pp. 1409-1425.
- [28] U. Sauer et.al., “Analysis of the Performance Parameters of Lead/Acid Batteries in Photovoltaic Systems,” in *Journal of Power Sources*, 1997, 64(1), pp. 197-201.
- [29] H. Kim, and K. Shin, “Scheduling of Battery Charge, Discharge, and Rest,” in *Real-Time Systems Symposium (RTSS)*, 2009.
- [30] M. Nispel, “Important Considerations When Reducing The Run-Times of VRLA UPS Batteries,” 2011, <http://www.battcon.com/>.
- [31] K. Yabuta, T. Matsushita, and T. Tsujikawa, “Examination of the Cycle Life of Valve Regulated Lead Acid Batteries,” in *International Telecommunications Energy Conference (INTELEC)*, 2007.
- [32] H. Bindner, T. Cronin, P. Lundsager, J. Manwell, U. Abdulwahid, and I. Gould, “Lifetime Modelling of Lead Acid Batteries,” Technical Report, Risø National Laboratory, 2005.
- [33] D. Wang, C. Ren, and A. Sivasubramaniam, “Virtualizing Power Distribution in Datacenters,” in *International Symposium on Computer Architecture (ISCA)*, 2013.
- [34] <http://sine.ni.com/nips/cds/view/p/lang/en/nid/1865>
- [35] <http://sine.ni.com/nips/cds/view/p/lang/en/nid/14132>
- [36] <http://www.u-rd.com/technical/tech2.html>
- [37] <http://www.wti.com/c-34-high-amp-c19-switched-pdu.aspx>
- [38] <http://www.ni.com/labview/>
- [39] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang, “The HiBench Benchmark Suite: Characterization of the MapReduce-Based Data Analysis,” in Data Engineering Workshops, *International Conference on Data Engineering (ICDE)*, 2010.
- [40] The CloudSuite 2.0, <http://parsa.epfl.ch/cloudsuite>.
- [41] L. Wang and W. Chow, “Solar Radiation Model,” in *Applied Energy*, Vol 69, 2011, pp. 191-224.
- [42] S. Govindan, D. Wang, L. Chen, A. Sivasubramaniam and B. Urgaonkar, “Towards Realizing a Low Cost and Highly Available Datacenter Power Infrastructure,” in *Workshop on Power-Aware Computing and Systems (HOTPOWER)*, 2011.
- [43] S. McCluer, “Battery Technology for Data Centers and Network Rooms: Lead-Acid Battery Options,” APC White Paper #30.
- [44] “Lifecycle Costing for Data Centers: Determining the True Costs of Data Center Cooling,” White Paper, Emerson Network Power, 2003.
- [45] A. Acquaviva, A. Aldini, M. Bernardo, A. Bogliolo, E. Bonta and E. Lattanzi, “Assessing the Impact of Dynamic Power Management on the Functionality and the Performance of Battery-Powered Appliance,” in *International Conference on Dependable Systems and Networks (DSN)*, 2004.
- [46] L. Cloth, M. Jongerden and B. Haverkort, “Computing Battery Lifetime Distributions,” in *International Conference on Dependable Systems and Networks (DSN)*, 2007.
- [47] C. Ma and Y. Yang, “A Battery-Aware Scheme for Routing in Wireless Ad Hoc Networks,” in *IEEE Transaction on Vehicular Technology*, 2011, 60(8), pp. 3919-3932.
- [48] T. Panigrahi, D. Panigrahi, C. Chiasserini, S. Dey, R. Rao, A. Raghunathan and K. Lahiri, “Battery Life Estimation of Mobile Embedded Systems,” in *International Conference of VLSI Design*, 2001.
- [49] L. Serrao, S. Onori, A. Sciarretta, Y. Guezennec, and G. Rizzoni, “Optimal Energy Management of Hybrid Electric Vehicles including Battery Aging,” in *American Control Conference (ACC)*, 2011.
- [50] L. Liu, C. Li, H. Sun, Y. Hu, J. Xin, N. Zheng and T. Li, “Leveraging Heterogeneous Power for Improving Datacenter Efficiency and Resiliency,” in *IEEE Computer Architecture Letters (CAL)*, 2014.
- [51] L. Liu, H. Sun, Y. Hu, J. Xin, N. Zheng and T. Li, “Leveraging Distributed UPS Energy for Managing Solar Energy Powered Data centers,” in *International Green Computing Conference (IGCC)*, 2014.
- [52] L. Liu, C. Li, H. Sun, Y. Hu, J. Gu, T. Li, J. Xin and N. Zheng, “HEB: Deploying and Managing Hybrid Energy Buffers for Improving Datacenter Efficiency and Economy,” in *International Symposium on Computer Architecture (ISCA)*, 2015.