

HEB: Deploying and Managing Hybrid Energy Buffers for Improving Datacenter Efficiency and Economy

Longjun Liu^{1,2,†}, Chao Li^{3,2,†}, Hongbin Sun^{1,*}, Yang Hu², Juncheng Gu², Tao Li², Jingmin Xin¹ and Nanning Zheng¹

¹ School of Electrical and
Information Engineering
Xi'an Jiaotong University, China

{longjun.liu@stu.xjtu.edu.cn},
{hsun, jxin, nnzheng}@mail.xjtu.edu.cn

² Department of Electrical and
Computer Engineering
University of Florida, USA

{huyang.ece, gujuncheng}@ufl.edu,
taoli@ece.ufl.edu

³ Department of Computer
Science and Engineering
Shanghai Jiao Tong University
Shanghai, China
lichao@cs.sjtu.edu.cn

Abstract

Today, an increasing number of applications and services are being hosted by large-scale data centers. The massive and irregular load surges challenge data center power infrastructures. As a result, power mismatching between supply and demand has emerged as a crucial issue in modern data centers which are either under-provisioned or powered by intermittent power sources. Recent proposals have employed energy storage devices such as the uninterruptible power supply (UPS) systems to address this issue. However, current approaches lack the capacity of efficiently handling the irregular and unpredictable power mismatches.

In this paper, we propose Hybrid Energy Buffering (HEB), the first heterogeneous and adaptive strategy that incorporates super-capacitors (SCs) into existing data centers to dynamically deal with power mismatches. Our techniques exploit diverse energy absorbing characteristics and intelligent load assignment policies to provide efficiency- and scenario- aware power mismatch management. More attractively, our management schemes make the costly energy storage devices more affordable and economical for datacenter-scale usage. We evaluate the HEB design with a real system prototype. Compared with a homogenous battery energy buffering system, HEB could improve energy efficiency by 39.7%, extend UPS lifetime by 4.7X, reduce system downtime by 41% and improve renewable energy utilization by 81.2%. Our TCO analysis shows that HEB manifests high ROI and is able to gain more than 1.9X peak shaving benefit during an 8-years period. It allows datacenters to adapt to various power supply anomalies, thereby improving operational efficiency, resiliency and economy.

Categories and Subject Descriptors:

C.0 [Computer System Organization]: General

Keywords: Datacenter, Power Management, Energy Storage

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISCA'15, June 13-17, 2015, Portland, OR USA

1. Introduction

It is predicted that the power consumption of world data centers alone will approach 1,000TWh within a decade (2013-2025), which is more than the total now used for all purposes by Japan and Germany combined [1]. The huge power demands not only imply significant electricity cost but also lead to tremendous carbon emission.

Driven by the enormous amount of power cost and environmental concern, industry and academia alike are focusing more on the new perspectives of improving data center power infrastructures. Currently, there are two primary techniques: (1) aggressively under-provisioned datacenter power infrastructures (a.k.a., power under-provisioned data centers), which has been proved as a meaningful methodology to dramatically reduce infrastructure capital expenses (CAP-EX) and monthly recurring operating expenses (OP-EX) [2-8], and (2) renewable energy integration into data center facilities. To effectively reduce carbon emission, not only academia has started to study the intermittent energy power management schemes [9-18], but also many IT companies (Microsoft, IBM, Apple, Google, HP, etc.) have begun to build renewable energy data centers [19-23].

Although the above two power provisioning schemes can significantly reduce electricity cost and carbon emission, we notice that power mismatches are more prone to occur since (1) power under-provisioned data centers intentionally subscribe lower power supply infrastructures, which may lead to power budget violations due to the irregular and bursty service requests, and (2) the nature of renewable power sources is intermittent and fluctuated, and it may exceed (i.e. valley power) or lower (i.e. peak power) than power demands even if the latter are stable.

Existing proposals to handle the power mismatching issue can be classified into two categories: (1) performance scaling techniques on the power demand side, and (2) energy sources tuning mechanisms on the power supply side. Among those, the performance scaling techniques primarily leverage server power state tuning (e.g., DVFS and ACPI techniques [24-27]) and workload scheduling to accommodate runtime power budget or track the time-vary renewable energy budget [11, 14-16, 28]. These approaches can forcefully cap power mismatches at the cost of performance degradation. Recently, a new tuning knob on the power supply

*Corresponding author: Hongbin Sun. †Authors contribute equally

side, the energy storage devices (e.g., UPS batteries), is repurposed to shave peak power mismatching [6, 8, 29-33]. Compared with performance scaling schemes, this technique can effectively mitigate performance penalty.

When used to address the power mismatching issue, existing UPS batteries manifest several disadvantages: (1) batteries have limited lifetime cycle (approximate 2000 to 3000 cycles [34]). Frequent charging/discharging can lead to a much shorter lifetime [35]; (2) large discharge current may lead to less usable capacity (known as the Peukert's law effect) [36]; and (3) to avoid battery overheating during charging, batteries cannot be re-charged very fast with large charging current. In addition, the low energy efficiency is another major drawback of batteries - the round trip energy loss of batteries can reach to 15%-20% [37]. Therefore, *can we find a new way to gracefully handle the power mismatching on the power supply side while avoiding these limitations of batteries?*

In this paper, we propose a different power provisioning scheme - HEB, which explores the benefits of incorporating hybrid energy buffering technologies into data centers to intelligently and economically handle power mismatching. Specifically, we integrate super-capacitors (a.k.a., ultra-capacitors) with conventional UPS systems to provide an additional layer of safety in the event of unexpected power mismatches. Super-capacitors (SCs) have emerged as a promising alternative to batteries [38]. They have the following advantages: (1) high efficiency and low round-trip energy loss, (2) allowing fast charging and discharging with a high current, and (3) two to three orders of magnitude more life cycles than batteries [37, 38]. However, currently SCs are still too expensive for the large-scale, exclusive deployment in data centers. As a result, the heterogeneous energy buffering systems, which combine batteries and SCs, provide a more feasible and attractive solution.

When transmitting from homogeneous to heterogeneous energy storage technologies, challenges arise as the latter requires more intelligent power management schemes between the two types of energy buffers to achieve efficiency and economy: (1) for a given peak power mismatching scenario, there exists an optimal schedule of discharging that could provide the longest discharging duration. Note that the optimal discharging point often shifts as the available stored energy changes in either batteries or SCs, and (2) for a given valley power charging opportunity, the energy buffers should be rapidly charged so that they can supply enough energy prior to the following peak power mismatching. What is more, from the perspective of energy efficiency, the ideal usage pattern of heterogeneous energy buffers also depends on power mismatching scenarios. For instance, when the power peaks are small and narrow, it is better to exclusively use SCs to provide power shortfall since SCs can be easily charged within a short duration while suffering negligible round-trip energy loss. In sum, we need an emergency-aware and workload-aware power management strategy, which can effectively coordinate the utilization of het-

erogeneous energy buffers and intelligently assign load power demands to hybrid energy buffers.

This paper makes the following contributions:

- We explore super-capacitors (SCs) as a new tuning knob on the data center power supply side to handle the irregular power mismatches. By comparing the amortized cost, life cycle, charging/discharging rate, and energy efficiency between SCs and batteries, we demonstrate the design feasibility of combining SCs as hybrid energy storage buffering in data centers.
- We present HEB, a novel, heterogeneous energy buffering based power provisioning architecture that enables data centers to effectively and economically incorporate batteries and SCs to handle power mismatching. The architecture of HEB is based on distributed and reconfigurable energy storage scheme which is easy to scale out and configure.
- We propose a tailored power management framework, which can intelligently assign different ratio of the server loads to appropriate energy buffers to achieve high energy efficiency and low performance degradation during power mismatching events. The power management framework can auto-tune the load assignment and self-optimize its assignment performance.
- We implement a scale down version of HEB as a prototype research platform and our real-system based experiments show that HEB could improve energy efficiency by 39.7%, extend UPS lifetime by 4.7X, reduce server downtime by 41%, and improve renewable energy utilization by 81.2%. HEB also manifests high CAP-EX, ROI and is able to gain more than 1.9X peak shaving benefit during an 8-years period.

The rest of this paper is organized as follows. Section 2 provides the background and motivation of handling power mismatching. Section 3 characterizes the heterogeneous energy buffers and highlights the key design considerations. Section 4 presents our heterogeneous power provisioning architecture and compares it with conventional design. Section 5 proposes the power management policies for HEB. Section 6 describes our prototype and experimental methodology. Section 7 presents the evaluation results with our prototype system and TCO analysis. Section 8 discusses related work. Finally, Section 9 concludes this paper.

2. Background and Motivation

To mitigate power cost and carbon emission, both under-provisioned and renewable energy powered data centers have attracted growing attentions in recently years. In this Section, we first introduce the two emerging power provisioning schemes. Then, we discuss the power mismatching issues in the two data centers.

2.1 Power Under-Provisioned Data Centers

Conventional data center power infrastructures are commonly over-provisioned based on the nameplate rating power of all the servers, but this incurs significant power

overhead and low power infrastructure utilization [2]. Therefore, many data centers today start to under-provision power infrastructures [2-8]. To detail the benefits and disadvantages of the schemes, we analyze the different power provisioning rates based on a Google cluster workload trace [2, 32], as shown in Figure 1(a). We assume four different power provisioning rates (P1-P4). Among those, P1 is an over-provisioning scheme and can cover all peak demands. P4 is an under-provisioning scheme and only supplies 40% power budget for the data center loads. We define the maximum power provisioning utilization ($MPPU$) as: $MPPU = \sum t / \sum T$ ($\sum t$ is the total time, during which power demands reach the provisioned budget and $\sum T$ is the total load running time).

Figure 1(a) shows that aggressively under-provisioning power infrastructure can yield high $MPPU$ and low infrastructure capital cost (capital cost is proportional to the provisioned IT power facility, estimated as \$10-20 per Watt [3, 6, 8]). Nevertheless, the under-provisioning power infrastructure incurs more power mismatching, which degrade load performance if they are forcedly capped. Therefore, to reduce power related CAP-EX and to improve $MPPU$ while avoiding performance degradation; the power mismatching caused by under-provisioned power infrastructure should be gracefully handled.

2.2 Renewable Energy Powered Data Centers

Provisioning clean renewable energy into data centers can alleviate their carbon emission. However, due to its intrinsic fluctuation, intermittent power mismatching is one of the greatest challenges for integrating renewable energy. Recent proposals leverage load deferment and load scheduling [9-12] to match demand to the supply, which may violate the service level agreement (SLA) and are not suitable for performance oriented data centers. Another approach is to utilize large-scale battery farms to regulate the power mismatches for performance consideration.

As shown in Figure 1(b), during the peak power, the load can draw additional energy from batteries, and during the valley power, the surplus renewable energy can recharge batteries. Since the renewable energy generation is time-varying, it is critical for batteries to make the most of the opportunities of each power valley to store more energy. Therefore, the renewable energy utilization (REU) is a crucial consideration to maximally utilize the green energy for power mismatches handling. The REU can be defined as $REU = (\sum B_{RE} + \sum L_{RE}) / \sum S_{RE}$, where $\sum B_{RE}$ is the renewable energy stored in batteries, $\sum L_{RE}$ is the renewable energy for load and $\sum S_{RE}$ is the total amount of renewable energy generation. However, typically batteries have the upper bound of charging current and cannot timely absorb all the renewable energy during the very deep power valleys, which wastes renewable energy and leads to low REU . Consequently, we need alternative energy storages without the




Figure 1. Modern datacenters in under-provisioned power infrastructure and renewable generation environment

limitation of upper-bound charging current, which can take advantage of the deep valley power mismatching opportunities to maximally absorb intermittent power in renewable energy powered data centers.

3. Hybrid Energy Buffers: Characterization and Key Design Concerns

In this section, we characterize different energy storage devices and discuss the design concerns and opportunities on leveraging hybrid energy buffers in data centers.

3.1 Characterization

We first build up an energy storage characterization test-bed which consists a group of super-capacitor modules (Maxwell 16V, 600F [39]) and LA-batteries (12V, 4AH), as shown in Figure 2. Our experimental test-bed allows us to charge and discharge SCs and batteries alone for characterizing their behavior. In addition, we can jointly utilize SCs and batteries to power server loads.

Energy Efficiency Analysis: One of the primary reasons for using SCs to buffer energy is that they incur negligible round-trip energy loss [37]. Our experimental measurements indicate that SCs can achieve 90%-95% round-trip energy efficiency, as shown in Figure 3. In contrast, lead-acid batteries have less than 80% efficiency even in the best case in our experiments. The energy efficiency calculation is based on detailed charging/discharging logs of our system with different server power demands.

In fact, the efficiency of batteries can be even worse, depending on their usage patterns. There is a so-called recovery effect: batteries cannot release all of their stored energy in a one-time, high-current discharging – part of the stored energy seems to be “lost”; during periods of no or very low discharge, they can recover the energy “lost” to a certain extent [40]. Figure 3 shows our characterization of different discharging scenarios with one, two and four servers, which reflect different power demands and battery discharging currents. The one-time discharging efficiency of




Figure 2. Hybrid energy buffers characterization test-bed




Figure 3. Energy efficiency comparison with one, two and four servers




Figure 4. Cost comparison of different energy storage technologies

the lead-acid battery decreases as we add more servers (i.e. increase the power demand). Given additional discharge cycles and enough recovery time, the battery efficiency can increase significantly (i.e. by 6%~24%). However, this does not mean that one should always cap load power demand and wait for the battery to recover. This is because the energy waste due to server on/off cycles can be significant (i.e. account for nearly half of the recovered energy), as shown in Figure 3. Therefore, in order to improve energy efficiency, it is wise to use SCs to deal with power mismatching.

Cost Analysis: SCs commonly have much longer life cycle. They can sustain hundreds of thousands charging/discharging cycles without degradation. Consequently, the amortized cost of SCs to each charging and discharging cycle (\$/KWh/cycle) is competitive. As shown in Figure 4, the initial cost of most UPS batteries is about 100-300 \$/KWh, while for SCs it is about 10K-30K \$/KWh [34, 37, 38]. However, the amortized cost of SCs is very close to NiCd and Li-ion batteries (about 0.4 \$/KWh per cycle) and is higher than lead-acid batteries. In sum, SCs have high initial capital cost but very competitive amortized cost. Note that the technology improvement of SC is much faster than that of the lead-acid battery, which makes the cost gap between SC and battery smaller in the future [41]. On the other hand, from the perspective of availability, it would be worthwhile to deploy efficient energy storage buffers in data centers to avoid the even more expensive service downtime, which has reached to \$100,000 per hour on average in 2010 and increased by 38% between 2010 and 2012 [42].

Charging and Discharging Comparison: Batteries and SCs manifest completely different charging/discharging features as battery stores energy electrochemically while there is no chemical reaction in SCs. SCs can be charged very fast without the limitation of upper-bound charging current, but neither does battery. We compare different discharging scenarios of batteries and SCs with different numbers of servers, as shown in Figure 5. Our results show that the SC discharging voltage shows linearly declining trend irrespective of power demands. However, batteries exhibit a sharp voltage drop in light of large power demands since the chemical reaction process in batteries is slow and cannot release more power with a short time period. When handling power mismatching, the large peak power demands may cause battery voltage to transiently drop, which poses serious threat to server uptime. Therefore, it is important to




Figure 5. Comparison of SC and battery discharging with different power demands

Figure 6. Discharge duration (m/n means m servers for SC and n servers for battery)

avert using batteries to handle the large peak power mismatches. On the contrary, with the linear discharging properties, SCs are more stable and controllable for those scenarios.

3.2 Implications from Characterization

Based on the characterization above, we can see that leveraging SCs can deliver the high current levels needed for dealing with large power mismatching while being recharged quickly between the events with high energy efficiency. Nevertheless, this does not necessarily mean that one should only employ SC as its current cost is still high. It is crucial to use hybrid energy buffers for exploiting the merits of both while averting their shortcomings.

We further perform experiments to explore how to jointly utilize SCs and batteries to power servers. We first vary the number of servers assigned to the batteries and SCs to measure the maximum server runtime with constant power demands. In the experiments, whenever one energy storage device is depleted, the other will take over the entire load immediately via power switches. As Figure 6 shows, there is an optimal load assignment that can provide the longest discharging time. It is clear that one should not heavily rely on either SCs or batteries. For example, by assigning heavy load on SCs, the server cluster runtime (uptime) can be decreased by 25% on average. Therefore, we should identify an optimal ratio to assign servers that powered by batteries or SCs for maximizing the server runtime.

The challenge of such load power assignment is that there is not a fixed optimal operating point. The optimal server assignment actually depends on the current capacity of the heterogeneous energy buffers and the time-vary shape




Figure 7. The comparison of different energy storage system architectures in datacenters

of power peaks. Therefore, we should dynamically identify the optimal operation point to distribute appropriate loads between SCs and batteries upon a power mismatching event.

4. System Architecture

Many research efforts focus on optimizing power delivery topology to enhance data center efficiency. In this section, we first analyze the pros and cons of leveraging current energy storage system to handle power mismatching. Then we present our heterogeneous power architecture in detail.

4.1 Energy Storage System Architecture Analysis

The centralized (Figure 7(a)) and distributed (Figure 7(b)) topologies are two primary energy storage architectures in data centers currently. In a centralized battery energy storage system, the UPS battery system locates on the critical path between the Automatic Transfer Switch (ATS) and the Power Distribution Units (PDU). When used to deal with the peak power mismatching (similar to [8]), it can only provide load shifting for the entire data center but cannot handle the peak shaving in a fine-grained manner. Moreover, the centralized UPS system commonly works on-line and always performs double converting (AC-DC-AC), which leads to 4-10% power losses [29]. It is also not easy to scale out in large data centers.

At present, IT giants such as Google, Microsoft and Facebook have explored the distributed power topology (Figure 7(b)) in their data centers. For instance, Facebook employs a cabinet of batteries for every 6 racks, or a total of 180 servers [43]. Their design is scalable in rack level and allows data centers to shave peak power by using a fraction of the installed batteries. To avoid power double converting, it needs customized servers that support DC power. Google mounts a battery in every server after Power Supply Unit (PSU) [44]. This design can completely avoid the battery double converting energy loss when shaving the peak power mismatching [29]. However, each server is assigned to a dedicated battery and multiple servers cannot share battery energy with each other to assist peak shaving. Furthermore, as the batteries are deployed in the inner chassis of servers, they are constrained by limited capacity. Note that both of the existing designs for data centers are exclusively based on the homogeneous batteries and inevitably suffer the drawbacks of battery. Figure 7(c) depicts our heterogeneous energy buffer topology, which provides opportunities to employ the pros and evade the cons of batteries and SCs

when handling the power mismatching. The power switch based control enables data centers to dynamically determine the distribution of server power demands between batteries and SCs. The batteries will offer bulk energy to the load since they can deliver large amount of energy slowly over a longer period of time while the SC pool will handle the transient peak power mismatching since they can be charged and discharged very quickly. The detailed architecture and hardware implementation scheme are presented as below.

4.2 HEB Architecture

Figure 8(a) illustrates an overview of our HEB power provisioning architecture. The renewable/utility power charges batteries and SCs when the load power demands are lower than the provisioned power budget. The HEB Controller (*hControl*) is a key decision-making component that monitors and controls other components. The voltage & current of batteries and SCs collected from the sensors are transmitted to the *hControl*. The power switch states (i.e. ON/OFF) as well as all the server power demands information (measured by the Intelligent Power Distribution Unit or IPDU) are transferred to the *hControl* too. With the above state feedbacks, the *hControl* makes operation decisions and sends the control signals to each power switch to distribute energy sources for each server. In our current implementation, the *hControl* is a low power server that hosts our heterogeneous power management algorithms, such as dynamic scheduler and optimizer (detailed in Section 5).

Note that the *hControl* and heterogeneous energy buffers can be deployed either at cluster- or rack- level in data centers. Figure 8(b) illustrates the cluster-level deployment, which only uses one *hControl* and one group of heterogeneous energy buffers. In this case, the *hControl* controls all the servers. The DC/AC converter is needed due to the long distance power delivery from the energy buffers to each server, which inevitably leads to energy conversion loss. Figure 8(c) shows the rack-level deployment, which consists of several *hControls* and multiple groups of heterogeneous energy buffers. This can avoid the DC/AC conversion as the DC power can be directly delivered from the energy buffers to each server. The disadvantage is that each group of energy buffers is independent and cannot share their energy.

With both deployment strategies, the *hControl* can coordinate available heterogeneous energy buffers and assign appropriate energy to servers within its domain based our heterogeneous power management framework.

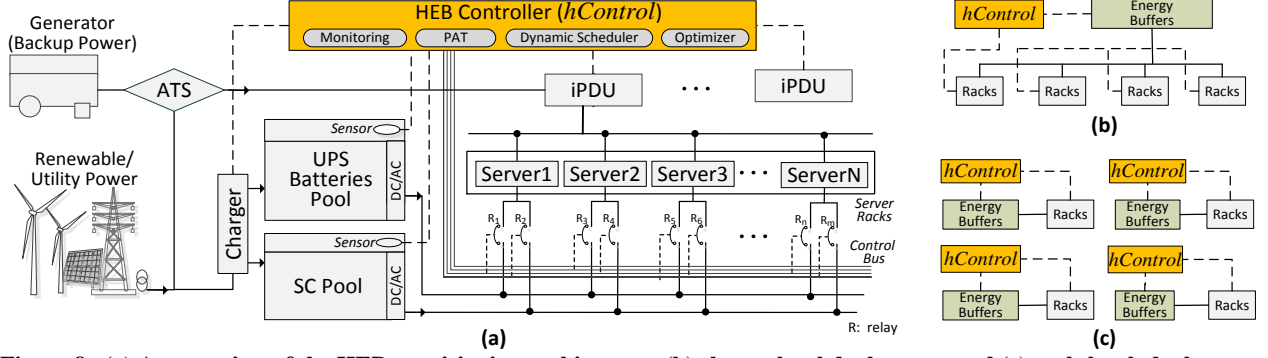


Figure 8. (a) An overview of the HEB provisioning architecture, (b) cluster-level deployment and (c) rack-level deployment

5. Power Management Framework

In this section, we present the heterogeneous power management framework for HEB, which primarily integrates power mismatching prediction method, dynamic load assigning mechanism and heterogeneous energy buffer optimization strategy. The proposed power management framework allows the hybrid energy buffer system to efficiently and economically handle the unregularly and unpredictable power mismatching.

5.1 Problem Formulation

At the beginning of each control interval, HEB controller obtains the current available capacity of batteries (ΔBA) and SCs (ΔSC) based on the feedback information from the sensors. We assume the power mismatching during the control interval is ΔPM . We define R_λ as the ratio of servers powered by SCs, therefore, the number of server powered by SCs is $NumS * R_\lambda$, where $NumS$ is the total number of servers. Likewise, the number of server powered by batteries is $NumS * (1 - R_\lambda)$. The HEB controller assigns the energy buffer based on the above four variables (i.e. ΔBA , ΔSC , ΔPM , and R_λ) to handle the peak power mismatching events, and it can calculate the energy efficiency (EE) and server downtime (SD) at the end of the control interval. We can formulate the peak power mismatching handling in HEB as following:

$$(SD, EE) = f_T(\Delta BA, \Delta SC, \Delta PM, R_\lambda)$$

where EE is the overall efficiency of heterogeneous energy buffers. SD is the aggregated time that servers have to shut down as the stored energy is insufficient to shave the peaks. As the values of ΔBA , ΔSC and ΔPM are not fixed during each power mismatching period, our power management goal is to minimize the SD and maximize the EE by adjusting the ratio R_λ . We solve the problem by constructing a dynamic load assignment table and continuously optimizing the assignment, which can be implemented in our real system prototype.

5.2 Dynamic Load Scheduling

A fundamental problem in the design space of HEB is how to assign load power to the most appropriate heterogeneous buffers. During each time-slot (e.g., 10 minutes by default), HEB dynamically distributes load power demands between batteries and SCs. Figure 9 shows the power management framework.

Prediction: To identify the average peak power characterization (e.g., small peaks or large peaks) of next time-slot, we employ time series prediction (TSP) method [45]. Specially, we leverage the classical triple exponential prediction (Holt-Winters exponential prediction) algorithm [46] to periodically predict the power demands, which can analyze the nature of the history and current data, extract meaningful statistics trend and predict future values. The algorithm maintains two groups of series data: the peak power and valley power. It predicts the peak power demands (P_{peak}) and valley power (P_{valley}) of next time-slot. The difference of P_{peak} and P_{valley} ($\Delta PM = P_{peak} - P_{valley}$) is the net amount power that needed from the energy buffers. Note that we select a time series prediction method that is effective for the data center power consumption patterns, but any sophisticated prediction approaches can be integrated into our power management framework.

Small Peaks Handling: When the average height of predicted power mismatching is mild and the duration is short, the HEB controller treats the batteries and SCs as a two-tier energy storage system. Either batteries or SCs can handle these small peak power mismatches. In order to enhance energy efficiency, the HEB controller preferentially assigns all loads power on SCs ($R_\lambda = 1$). This is because SCs have much better round-trip energy efficiency and they can be swiftly charged and discharged without degradation. Only when all the SCs are used up, the HEB controller will turn on all the battery relays and assign all server loads on batteries ($R_\lambda = 0$) to compensate the energy shortages. In brief, SCs are aggressively used to handle the small peak power mismatching for high energy efficiency while maintaining minimal server down time by employing batteries as supplement during the interval when SCs are used up.

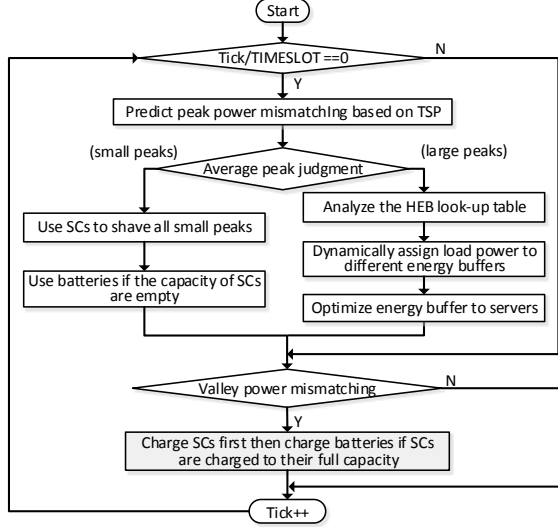


Figure 9. An overview of HEB power management framework (includes prediction, small peak/large peak handling and utilizing valley power mismatching for charging at any time)

Large Peaks Handling: The HEB controller treats batteries and SCs as a unified energy buffer when the predicted average peak power mismatching is significant and the duration is long (large peaks). In other words, the HEB controller schedules all the loads on batteries and SCs simultaneously to jointly shave large peaks. To maximize energy efficiency and minimize server downtime, we should carefully allocate an optimal R_λ ($0 < R_\lambda < 1$).

To this end, the HEB controller maintains a power allocation table (PAT) for its heterogeneous energy buffers. This table specifies initial and coarse-grained load assignments on batteries and SCs. Each entry of the power allocation table contains the available energy levels of the battery and SC pools, power demands and the server ratio that indicates the fractional servers powered by SCs. The initial value of each entry is obtained via profiling in a pilot scheme like Figure 6. The profiling values in the table are not fixed all the time, and they can be optimized and updated (detailed in Section 5.3). Figure 10 shows the pseudo code of the server loads assignment (Lines 1-11). Based on the available energy buffer and predicted average power mismatching value at each time-slot, the HEB controller can find the energy allocation ratio R_λ or similar R_λ in the PAT and dynamically control the on/off power switches to assign different ratio servers powered by SCs or batteries. However, as it cannot profile all scenarios of available energy buffer and power demands, the number of entries in PAT is limited. Therefore, it may be difficult to find an optimized energy allocation ratio R_λ in such initial PAT.

5.3 Optimizing Energy Buffering Allocation

As mentioned above, the PAT table cannot always guarantee the optimal load assignment results because (1)

```

//At the beginning of each time-slot
1. Obtain current SC capacity:  $SC_{initial}$ , Battery capacity:  $BA_{initial}$ , and predicted
   power mismatching  $\Delta PM (\Delta PM = P_{peak} - P_{valley})$ ;
2. For table index = 1 to n // search the look-up table PAT
3.   If ( $SC_{index} == SC_{initial}$  &&  $BA_{index} == BA_{initial}$  &&  $P_{index} == \Delta PM$ )
4.     find_index = index;
5.   End
6. End
7. If (find_index == 0) //does not find a matched entry
8.   find_index = Similar( $SC_{initial}$ ,  $BA_{initial}$ ,  $\Delta PM$ ); //search the most similar value
9. End
10. Server ratio  $R_\lambda = R_\lambda(\text{find\_index})$ ; //Find the ratio in PAT;
11. Allocate different numbers of servers to SC and BA based on  $R_\lambda$ ;
12. Collect running results at the end of the time-slot.
13. If (index == 0) //new entry (new energy buffer capacity & power demand)
14.   Round( $SC_{initial}$ ,  $BA_{initial}$ ,  $P$ ); //format data, P is the actual power demand
15.   Add ( $SC_{initial}$ ,  $BA_{initial}$ ,  $P$ ,  $R_\lambda$ ) to the PAT look-up table;
16. Else //update the existing entry of the PAT table
17.   If ( $SC_{end}/BA_{end} > SC_{initial}/BA_{initial}$ )
18.      $R_\lambda = R_\lambda + \Delta r$ ; //SC receives increased server assignment
19.   Else If ( $SC_{end}/BA_{end} < SC_{initial}/BA_{initial}$ )
20.      $R_\lambda = R_\lambda - \Delta r$ ; //BA receives increased server assignment
21.   End
22.   Update ( $SC_{initial}$ ,  $BA_{initial}$ ,  $P$ ,  $R_\lambda$ ) in the PAT look-up table;
23. End

```

Figure 10. Algorithm for smartly handling large peaks

the limited profiling data are based on a pilot run and can be less accurate, and (2) with the battery and SC aging, their ability of handling power mismatching will decline. Therefore, the table needs to be dynamically updated.

To ensure effectiveness, the HEB controller updates the PAT table during runtime. Figure 10 shows the pseudo code of the two kinds of optimization operations (Lines 12-23): (1) adding new entries into the table, and (2) updating the existing entry. It first collects the running results at the end of the time slot, which includes the real power mismatching value and server load allocation ratio R_λ of current time slot.

When adding a new entry, the results are formatted and become coarse-grained to avoid too many entries in the table. When updating the existing entry, the HEB controller checks the remaining capacity in SCs and batteries. If the actual battery capacity decline rate (Line 17) is faster than expected (e.g., due to internal wear-out, batteries were assigned too much load and have higher discharge rate than SCs), the HEB controller will increase the load ratio by $\Delta r=1\%$ (default value) to increase the usage of SCs in future allocation. If the actual battery discharging rate is slower (Line 19) than expected, HEB will reduce the load ratio to decrease the usage of SCs. This optimization operation is to balance the using of SCs and batteries for archiving minimized server downtime. The optimization algorithm can progressively correct any inaccuracies caused by profiling or energy buffer aging in its following iterations. As a result, the HEB controller can self-optimize its performance by fine-tuning load assignment effectiveness over the lifetime.

6. Evaluation Methodology

Based on the proposed heterogeneous energy buffer architecture, we build a scale-down prototype to evaluate our




Figure 11. Full-system prototype of HEB - a heterogeneous energy storage research platform

design and power management framework. As shown in Figure 11, the platform includes several small and large batteries/SCs connected by relays to power different servers. There are six two-way relays in our prototype which can simultaneously connect to six servers. The servers are mounted on the rack and respectively connected to IPDU. The IPDU can switch ON/OFF server power supply, report the server power draw every second and send it to the controller by SNMP commands over the Ethernet. Any power management algorithm can be integrated in the controller to monitor and control all components in our prototype. Our platform can be deployed in either conventional power under-provisioned data centers or renewable energy powered data centers to handle the power mismatching.

We choose various datacenter workloads from Hibench [47] and CloudSuite [48]. Hibench contains nine typical Hadoop workloads (including micro benchmarks, HDFS benchmarks, web search benchmarks, etc.). CloudSuite benchmarks are based on real-world software stacks and consist eight popular applications in today's data centers. As shown in Table 1, we select eight workloads from five classified categories. Within each experiment, a workload can be executed iteratively.

Based on the total capacity of energy storage buffers, we select six lower-power computing nodes that use Intel Core i7-2720QM 4-core CPU. Our servers support dual-corded power supplies, one is from the energy storage source and one is from the utility power via IPDU. The measured idle power and peak power of each server are 30W and 70W, respectively. The low-power servers are matched with our energy storage prototype system.

Workloads (Abbr.)	Category	Peak
Page Rank Algorithm of Mahout (PR)	Web Search Benchmarks	Large Peaks
Word Count Program on Hadoop (WC)	Micro Benchmarks	
Data Analysis (DA)	CloudSuite Benchmarks	
Web Search (WS)	CloudSuite Benchmarks	
Media Streaming (MS)	CloudSuite Benchmarks	Small Peaks
Dfsioe (DFS)	HDFS Benchmarks	
Hivebench (HB)	Data Analytics	
Terasort (TS)	Micro Benchmarks	

Table 1. The evaluated workloads [47, 48]

Our server system kernel can be configured with the on demand frequency scaling governor. We can set the low frequency as 1.3GHz and the high frequency as 1.8 GHz. To fully evaluate our small and large peak power management policies, we divide the eight workloads into two groups, one group runs on the high frequency and the other group runs on the low frequency. In this way, we can construct two general peak shapes (small peaks and large peaks) to fully evaluate our power management policies. Note that our method is similar to [8], which leverages SPECjbb to construct various peak demand curves for fully evaluating the power management algorithms.

In our experiments, the controller can collect the utility power consumption of all the servers via IPDU. We set a maximum power drawn from utility (utility power budget, e.g., 260W for six servers). Whenever the server demands exceed the budget (peak occurs), the controller would tap into the energy stored in the energy buffers. Oppositely, the remaining energy can charge energy buffers when the server power demands are lower than the budget.

7. Experimental Evaluation

This section evaluates the benefits of provisioning heterogeneous energy buffers for data centers. To be more specific, we compare *HEB* to five kinds of power management schemes as summarized in Table 2. Among those, *BaOnly* is a representative peak power management technique similar to prior work [8], which only uses homogeneous UPS batteries to shave peak power. Note that with *BaOnly*, the serv-

Schemes	Architecture	Method description
<i>BaOnly</i>	Battery only	Only use battery to handle power mismatch
<i>BaFirst</i>	Hybrid (Battery+SC)	Discharge batteries first, then SCs if the capacity of batteries are empty
<i>SCFirst</i>	Hybrid	Discharge SCs first, then batteries
<i>HEB-F</i>	Hybrid	Load-aware assignment based on power demand value of the last time-slot
<i>HEB-S</i>	Hybrid	Load-aware assignment based on statics and limited profiling information
<i>HEB-D</i>	Hybrid	Load-aware assignment based on our dynamic and optimized power management framework

Table 2. The evaluated power management schemes

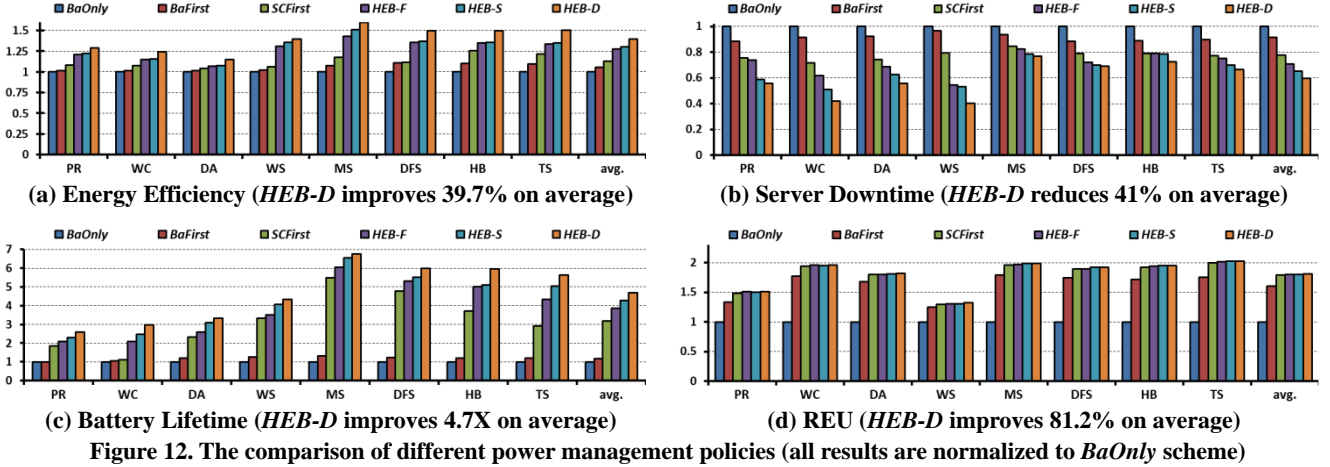


Figure 12. The comparison of different power management policies (all results are normalized to *BaOnly* scheme)

ers are still mainly powered by utility grid when there is no peak power. Although *BaFirst* and *SCFirst* both use hybrid energy buffers, they lack intelligent server allocation policies and only employ a priority-based method to handle power mismatches. The *HEB-F* and *HEB-S* are two naïve implementations of *HEB*. The *HEB-F* assigns the heterogeneous energy buffers to different servers based on the power demand information of last time-slot. The *HEB-S* assigns load power based on a static profiling table that has limited entries. The *HEB-D* is our proposed dynamic and optimized power management framework.

The purpose of comparing *HEB-D* with *HEB-F* and *HEB-S* is to understand the impact of reduced prediction error rate on performance improvement. To fairly compare the performance improvement of battery only and hybrid energy buffers, their total capacity is set to the same by configuring the small and large SCs and batteries in the prototype (the initial ratio of SCs and batteries is 3:7).

Note that this study mainly compares systems with equal storage capacity (so that they have the same worst-case emergency handling capabilities). The reason why we did not compare “equal size” or “equal cost” systems is that they are technology-/vendor- dependent. The capacity of SCs has a direct impact on the performance and lifetime of our systems. For “equal-cost” and “equal-size” designs, it is very hard to tell if the improvement is a result of our optimization scheme or a result of the capacity change due to different SC technologies.

By running six different power management policies in our controller, we evaluate the overall performance of peak shaving in under-provisioned datacenters and renewable energy utilization (REU) in renewable energy datacenters. We further vary the total capacity of hybrid energy buffers and the ratios between SCs and batteries to evaluate their performance impact. At last, we analyze the cost breakdown, return on investment (ROI, CAP-EX benefit) and peak shaving revenue (OP-EX gain) of *HEB*. In the following paragraphs, we present the detailed performance comparisons of *HEB* with other five baseline power management policies under different metrics.

7.1 Energy Efficiency

To improve the efficiency of energy storage systems, one must carefully assign and utilize both SCs and batteries to obtain maximal energy efficiency.

Figure 12(a) shows the overall energy efficiency measurement. Compared to a conventional battery-only power provisioning scheme, the heterogeneous energy buffers yield a visible efficiency improvement. The reason why *BaFirst* is very close to a battery only design is that *BaFirst* always charge/discharge battery first which reduces the chances of SCs utilization. If we always discharge the SC first, we can greatly reduce energy loss such as *SCFirst*, but when the SCs are depleted, batteries would have to handle all the high current drawn which still leads to efficiency degradation. Therefore, employing load-aware assignment to balance the usage of SCs and batteries can achieve better efficiency improvement (e.g., *HEB*). The energy assignment of *HEB-F* is based on the former power demand information which is a naïve prediction scheme and may lead to incorrect energy assignment. The errors in prediction decrease energy efficiency. The *HEB-S* often makes a sub-optimal energy assignment as it only has a coarse-grained profiling table. In contrast, *HEB-D* can achieve better energy efficiency improvement. In addition, *HEB-D* manifests higher efficiency on both small peak workloads (as SCs are preferentially used) and large peak workloads (as loads are dynamically allocated with energy between batteries and SCs) via our proposed policy. The improvement is 52.5% for small peaks and 27.1% for large peaks on average.

7.2 Server Performance

Mitigating performance degradation is one of the key goals of leveraging energy storage devices to handle power mismatching. We employ the server down time as the primary performance metric. In our experiment, server downtime is the aggregated time during which server power demands exceed power budget but the energy buffers do not have enough power to shave the peak. We chose the least recently used servers to shut down when we have to. Note that in this paper we do not use other control knobs such as




Figure 13. The impact of different capacity ratios (m:n means the ratio between SCs and batteries. All the metrics are normalized to the ratio of 3:7)

DVFS for simplifying the performance comparison of different power management schemes. Therefore, our evaluation of server downtime reflects the average availability yielded by a power management scheme. To compare the server downtime of different power management policies, we intentionally lower the utility power budget to trigger server downtime. Due to the Peukert law's effect, it is difficult to adopt *BaOnly* to handle the large peaks. Doing so may lead rapid drop of battery voltage, especially when the batteries have low SoC (State of Charge). The server downtime can be mitigated with integrated SCs in the energy buffer, as shown in Figure 12(b). As can be seen, *HEB* can always maintain the longest discharging duration by dynamically adjusting the server assignment between SCs and batteries. The *HEB-D* can reduce more server downtime (41%), especially for the large peak workloads.

7.3 Battery Lifetime

One of the original intentions of introducing SCs as heterogeneous energy buffers is to protect batteries from large current discharging and prolong their lifetime. We use the Ah-Throughput Battery Lifetime Model [49] to present the anticipated battery lifetime based on detailed battery usage logs. As shown in Figure 12(c), the SC preferential power management policy has more battery life cycle since batteries are used as backup (e.g., *SCFirst* and *HEB*). The *HEB* has better battery lifetime improvement as it only uses SCs to shave small peaks and jointly utilizes SCs and batteries to shave large peaks for protecting batteries from large current discharging. The *HEB-D* can improve the battery lifetime by 4.7X compared to the *BaOnly* scheme. Compared to the lifecycle of SCs, battery lifetime is the bottleneck of heterogeneous energy system lifespan. Longer battery lifetime implies lower replacement and maintenance cost of *HEB*.

7.4 Energy Utilization in Renewable Data Centers

We further present the benefit of heterogeneous energy buffer provisioning in light of renewable powered data centers. As mentioned in Section 2, it is critical to improve the renewable energy utilization (REU) for storing more green and clean energy to handle power mismatches in renewable data centers. Compared with pure battery equipped systems, SCs can absorb renewable energy without upper-bound of charging current, which can achieve more energy utilization. We tap into solar power to our prototype system instead of




Figure 14. The impact of capacity growth on (a) energy efficiency, (b) server downtime, (c) battery lifetime and (d) REU

utility power to evaluate the REU. Note that we have a small solar generation system on the roof of our Lab, which can provide real solar power for our experiments. The evaluation results show that only if introducing SCs to the energy buffer, the REU can be significantly improved (e.g., *BaFirst*, *SCFirst* and *HEB*). As *BaFirst* gives the first priority to batteries, it may lose some chances to absorb renewable energy with large charging current. *SCFirst* and *HEB* always utilize SC first to absorb renewable energy; they have very similar REU as shown in Figure 12(d) and all of them improve the REU about 81% on average compared with the pure batteries provisioning scheme.

7.5 Capacity Planning

We further evaluate the impact of different capacity provisioning for heterogeneous energy buffers. Firstly, by keeping the constant total capacity of the energy buffer, we adjust the capacity ratio between SCs and batteries. In detail, we adjust the Depth-of-Discharge (DoD) of energy buffers to generate different capacity ratios for batteries and SCs. For example, given 8Ah battery, we set the targeted DoD level as 60%. Its useable capacity is 4.8Ah (8Ah*60%). Our controller can disable the utilization of batteries once it hits its DoD threshold. We iteratively run the eight workloads with *HEB-D* power scheme and respectively obtain the average performance of energy efficiency, server downtime, battery lifetime and REU, as shown in Figure 13. The results show that the more ratios of SCs can obtain better performance improvement. Moreover, the impact of the capacity ratio is different across the four metrics. The battery lifetime has the most significant improvement as more SCs can be used to shave peaks. The improvement of energy efficiency and server downtime gradually becomes constant.

Secondly, we keep the capacity ratio between SCs and batteries constant (3:7), and increase the total installed capacity of energy buffers to evaluate its effect. In our experiments, we set a low DoD (40%) of SCs and batteries, and then gradually decrease the DoD (40%, 50%, 60%, 70%,




Figure 15. Cost analysis: (a) Cost breakdown, (b) ROI of hybrid energy storage and (c) The benefit of shaving peak power by ESD (in an 8-years operation period, *HEB* achieves 1.9X revenue from peak shaving benefit compared to *BaOnly*)

80%) to mimic the system capacity growth. Note that the higher DoD has less useable capacity. We run the eight workloads with *HEB-D* power scheme and measure the performance as shown in Figure 14. It's observed that the larger capacity can improve the efficiency and system resiliency. However, the relationship between performance and capacity may be non-linear. The results contribute to the right-sizing of the heterogeneous energy buffers for the real systems as the cost of provisioning energy buffers grows with the increased capacity.

7.6 TCO Analysis

Cost Breakdown: Figure 15(a) shows the cost breakdown of our *HEB* prototype. The energy storage devices (SCs and batteries) are the most expensive components (account for 55% of the overall expenditure). With our existing setup a HEB node powers six servers and its total cost is less than 16% of the server total cost (approximate \$4,850).

Return-On-Investment (ROI): We further evaluate in light of under-provisioned power infrastructure, whether it is worth to invest hybrid energy storage to reduce CAP-EX. Similar to [6], we define the cost of procuring hybrid energy buffers to sustain e hours of peaks as $e * C_{HEB}$ (\$/Watt), and the CAP-EX cost of the power infrastructure to under-provision by C_{cap} (\$/Watt). The ROI for hybrid energy buffer can be calculated as: $(C_{cap} - e * C_{HEB}) / (e * C_{HEB})$, where the C_{HEB} is the total cost of SCs and batteries. We assume the battery cost C_{bat} is 300\$/KWh and SC cost C_{sc} is 10K\$/KWh, as reported in [32, 37, 38]. The hybrid energy cost is: $C_{HEB} = C_{bat} * x + C_{sc} * y$, where x and y are the ratios of batteries and SCs and we set $x=0.3$ and $y=0.7$ based on our prototype. The C_{cap} is reported to grow by \$10-25 for every provisioned Watt.

We vary a wide range of C_{cap} from 2 to 20 (\$/Watt) and calculate the ROI in different peak durations as shown in Figure 15(b). Note that the corresponding cost is amortized during the lifetime (e.g., battery: 4 years, SC: 12 years and infrastructure: 12 years). We observe a positive ROI across most of the operating regions. This suggests that deploying hybrid buffer is worthwhile.

Gain from Peak Shaving: Utilities often charge datacenters expensive peak cost [8]. Energy storage buffer can be used to shave peak power and save the OP-EX cost [6, 8, 32]. We assume a 100KW datacenter deployed with 20KWh

homogenous batteries or hybrid energy buffer (SCs account for 30% and batteries account for 70%). The peak tariff is 12\$/kW. Applying different peak shaving policies to the two types of energy buffers, we compare their revenues due to peak cost reduction within 8 years, as shown in Figure 15(c). The break-even point (in year) for *BaOnly* (battery cost is 300\$/KWh) is 4.2 year, similar to [8]. Taking *BaOnly* as baseline, we calculate the peak shaving gain of other three heterogeneous schemes. Our *HEB* scheme can improve energy efficiency and reduce server downtime by 39.7% and 41% respectively, which are proportional to the harvested peak shaving benefit. The break-even points of *BaFirst*, *SCFirst* and *HEB* are 6.3, 4.9 and 3.7 years respectively. Even through the hybrid energy buffer has expensive initial CAP-EX cost than battery only buffer, with the highly efficient peak shaving policy of *HEB*, we can earn more than 1.9X revenue from peak shaving benefit by accumulating and then averaging the per-year net profit within 8 years. On the contrary, if not appropriately managed, leveraging hybrid energy buffer may be less profitable than utilizing homogenous buffer (e.g., the net profit of *BaFirst* is less than that of *BaOnly*).

In sum, our *HEB* technology is very cost-effective, with the hybrid energy storage devices (ESDs) and the efficient peak shaving algorithm, we can achieve considerable investment return from CAP-EX cost and earn great peak shaving benefit from OP-EX cost.

8. Related Work

We summarize the leading-edge studies on power provisioning and energy storage techniques in datacenters.

Novel Data Center Power Provisioning Schemes: With the increasing of scale and capacity, modern data centers become more power-constrained and carbon-constrained. To address the issues, many novel power provisioning schemes begin to spring up recently [2-8, 13-18, 30]. Wang et al. [30] proposed to virtualized power provisioning scheme in data centers, their vPower can significantly improve system utilization and application performance when working in under-provisioned power infrastructure. Pelley et al. [5] presented a dynamic power provisioning scheme in data centers. Their Power-Routing exploits shuffled topologies to dynamically connect the servers and diverse PDUs while balancing the workload across the PDUs for reducing

the power infrastructure provisioning cost. Meanwhile, there are many renewable power provisioning schemes in data centers to reduce carbon emission [13-18]. We propose heterogeneous power provisioning scheme in data centers, especially, we focus on dispatching heterogeneous energy buffering to dynamically and efficiently handle the power mismatching in the power under-provisioned data centers and renewable energy powered data centers.

Emerging Energy Storage Techniques: Recent efforts start to repurpose UPS batteries [6, 8, 29, 31, 32, 50, 51] to address peak power mismatching issue in data centers for decreasing power cost while maintaining load performance.

Govindan et al. [6] discussed the benefits and limitations of leveraging energy storage device (ESD, e.g., lead-acid batteries) in data centers to reduce data center peak power cost. Nonetheless, the proposed centralized architecture can incur 10-15% energy loss due to double-conversions. Kontorinis et al. [29] proposed distributed energy storage system (per-server UPS) to store energy during low load activity periods and use the energy to shave each server's peak. However, as battery has many constraints, several recent works have tried to explore new tuning knob. Zheng et al. [31] exploited centralized thermal energy storage (TES) to shaving peak power in data centers. As limited by the response time, they also combine the conventional UPS system to handle the frequent and transient peaks. Likewise, SCs have grabbed certain attention in recent characterization work [32, 52]. Wang et al. studied the multiple ESD technology provisioning and placement options in data centers. However, no real implantation and power management policies have been proposed. Distinguished from prior works, we build a real heterogeneous energy buffer prototype and develop new operation optimization algorithm to make SCs more affordable and economic for large-sale deployment in datacenters.

9. Conclusions

In this study, batteries and SCs are first pooled and dynamically dispatched in data centers as heterogeneous energy buffers. We investigate the characterization of different energy buffers (SCs and batteries) with a test-bed. After analyzing the current energy storage architecture, we propose HEB, a novel heterogeneous energy buffering power provisioning architecture that enables data centers to flexibly deploy SCs and batteries. To efficiently utilize the heterogeneous energy buffers, we tailored a power management framework to intelligently and dynamically assign different ratio server loads between SCs and batteries for achieving higher energy efficiency and lower performance degradation when handling power mismatching events. Based on the HEB design, we implement a scale-down prototype from scratch. We evaluate different power management policies with the prototype and the results show that HEB could improve energy efficiency by 39.7%, extend UPS lifetime by 4.7X, reduce system downtime by 41%,

and improve renewable energy utilization 81.2%. HEB manifests high CAP-EX ROI and is able to gain more than 1.9X peak shaving benefit during an 8-years operation period. We believe that deploying HEB to emerging data center power infrastructures could significantly improve their efficiency, resiliency and economy.

ACKNOWLEDGMENT

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, and by three IBM Faculty Awards. Chao Li is also supported in part by Yahoo! KSC Award, Facebook Fellowship and SJTU-MSRA Faculty Award.

References

1. M. Mills, "The Cloud Begins with Coal: Big Data, Big Networks, Big Infrastructure, and Big Power," in *National Mining Association & American Coalition for clean Coal Electricity*, 2013.
2. X. Fan, W. Weber and L. Barroso, "Power Provisioning for a Warehouse-Sized Computer," in *International Symposium on Computer Architecture (ISCA)*, 2007.
3. U. Hoelzle and L. A. Barroso, "The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines," in *Morgan and Claypool Publishers*, 2009.
4. F. Xing, X. Wang and C. Lefurgy, "How Much Power Oversubscription is Safe and Allowed in Data Centers," in *International Conference on Autonomic Computing (ICAC)*, 2011.
5. S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch and J. Underwood, "Power Routing: Dynamic Power Provisioning in the Data center," in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2010.
6. 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.
7. D. Wang, S. Govindan, A. Sivasubramaniam, A. Kansal, J. Liu and B. Khessib, "Underprovisioning Backup Power Infrastructure for Datacenters," in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2014.
8. 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.
9. 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.
10. 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.
11. I. Goiri, R. Beauchea, K. Le, T.D. Nguyen, M. Haque, J. Guitart, J. Torres and R. Bianchini, "GreenSlot: Scheduling Energy Consumption in Green Datacenters," in *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2011.

12. Z. Liu, M. Lin, A. Wierman, S. Low and L. Andrew, "Greening Geographical Load Balancing," in *SIGMETRICS Performance Evaluation Review*, 2011.
13. C. Ren, D. Wang, B. Urgaonkar and A. Sivasubramaniam, "Carbon-Aware Energy Capacity Planning for Datacenters," in *International Symposium on Modelling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2012.
14. I. Goiri, K. Le, T. Nguyen, J. Guitart, J. Torre and R. Bianchini, "GreenHadoop: Leveraging Green Energy in Data Processing Frameworks," in *European Conference on Computer Systems (EuroSys)*, 2012.
15. C. Li, R. Zhou and T. Li, "Enabling Distributed Generation Powered Sustainable High-Performance Data Center," in *International Symposium on High-Performance Computer Architecture (HPCA)*, 2013.
16. 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.
17. M. Haque, K. Le, I. Goiri, R. Bianchini and T. Nguyen, "Providing Green SLAs in High Performance Computing Clouds," in *International Green Computing Conference (IGCC)*, 2013.
18. 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.
19. <http://www.microsoft.com/environment/>
20. <http://www.ibm.com/ibm/environment/climate/>
21. <http://www.apple.com/environment/renewable-energy/>
22. <http://www.google.com/green/bigpicture/>
23. <http://www.hpl.hp.com/news/2012/apr-jun/netzero.html/>
24. D. Meisner, B. Gold and T. Wenisch, "PowerNap: Eliminating Server Idle Power," in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2009.
25. R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang and X. Zhu, "No Power Struggles: Coordinated Multi-level Power Management for the Datacenter," in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2008.
26. X. Wang and M. Chen, "Cluster-level Feedback Power Control for Performance Optimization," in *International Symposium on High-Performance Computer Architecture (HPCA)*, 2008.
27. C. Isci, A. Buyuktosunoglu, C.-Y. Cher, P. Bose and M. Martonosi, "An Analysis of Efficient Multi-core Global Power Management Policies: Maximizing Performance for a Given Power Budget," in *International Symposium on Microarchitecture (MICRO)*, 2006.
28. D. Lo, L. Cheng, R. Govindaraju, L. Barroso and C. Kozyrakis, "Towards Energy Proportionality for Large-scale Latency-critical Workloads," in *International Symposium on Computer Architecture (ISCA)*, 2014.
29. 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.
30. D. Wang, C. Ren and A. Sivasubramaniam, "Virtualizing Power Distribution in Datacenters," in *International Symposium on Computer Architecture (ISCA)*, 2013.
31. W. Zheng, K. Ma and X. Wang, "Exploiting Thermal Energy Storage to Reduce Data Center Capital and Operating Expenses," in *International Symposium on High-Performance Computer Architecture (HPCA)*, 2014.
32. 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.
33. 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.
34. S. Schoenung and W. Hassenzahl, "Long-Term vs. Short-Term Energy Storage Technologies Analysis: A Life-Cycle Cost Study," in *Sandia National Laboratories*, 2003.
35. M. Pedram, N. Chang, Y. Kim and Y. Wang, "Hybrid Electrical Energy Storage Systems," in *International Symposium on Low Power Electronics and Design (ISLPED)*, 2010.
36. D. Baert and A. Vervaeke, "Lead-acid Battery Model for the Derivation of Peukert's Law," in *Electrochimica acta*, 1999.
37. H. Chen, T. Cong, W. Yang, C. Tan, Y. Li and Y. Ding, "Progress in Electrical Energy Storage System: A Critical Review," in *Natural Science*, 2009.
38. D. Rastler, "Electricity Energy Storage Technology Options: A White Paper Primer on Applications, Costs, and Benefits," in *Electric Power Research Institute*, 2010.
39. <http://www.maxwell.com/products/ultracapacitors/products/16v-large-modules>.
40. D. Linden and T. Reddy, "Handbook of Batteries," in *McGraw-Hill (third edition)*, 2002.
41. B. Gregory, "Ultra-Capacitor Sizing and Packaging for Cost Effective," in *White Paper, IOXUS Inc.*, 2013.
42. "Datacenter Downtime: How Much Does It Really Cost?" in *White Paper from Aberdeen Group*, 2012.
43. <http://opencompute.org/>.
44. <http://www.cnet.com/news/google-uncloaks-once-secret-server-10209580/>.
45. Box, G. EP, G. M. Jenkins and G. C. Reinsel, "Time Series Analysis: Forecasting and Control," in *John Wiley & Sons*, 2013.
46. P. Goodwin, "The Holt-winters Approach to Exponential Smoothing: 50 Years Old and Going Strong," in *Foresight*, 2010.
47. 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.
48. The CloudSuite 2.0, <http://parsa.epfl.ch/cloudsuite>.
49. H. Bindner, T. Cronin, P. Lundsager, J. Manwell, U. Abdulwahid and I. Gould, "Lifetime Modelling of Lead Acid Batteries," in *the Technical Report of Risø National Laboratory*, 2005.
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)*, 2015.
51. L. Liu, H. Sun, Y. Hu, J. Xin, N. Zheng and T. Li, "Leveraging Distributed UPS Energy for Managing Solar Energy Powered Datacenters," in *International Green Computing Conference (IGCC)*, 2014.
52. A. Mirhoseini and F. Koushanfar, "HypoEnergy: Hybrid Supercapacitor-battery Power-supply Optimization for Energy Efficiency," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2011.