A Profit Optimization Framework of Energy Storage Devices in Data Centers: Hierarchical Structure and Hybrid Types

Xue Lin*, Massoud Pedram*, Jian Tang† and Yanzhi Wang†
*Department of Electrical Engineering
University of Southern California, Los Angeles, CA 90089
Email: xuelin@usc.edu, pedram@usc.edu
†Department of Electrical Engineering and Computer Science
Syracuse University, Syracuse, NY 13244
Email: ywang393@syr.edu, jtang02@syr.edu

***Abstract*—This paper investigates the hierarchical deployment and over-provisioning of energy storage devices (ESDs) in data ceners by (i) adopting a realistic power delivery architecture (from Intel) for centralized ESD structure as the starting point; (ii) presenting a novel and realistic power delivery architecture, borrowing the best features of the centralized ESD structure from Intel and distributed single-level ESD structures from Google and Microsoft, and supporting the case that different types of ESDs are employed for each of the data center, rack, and server levels; (iii) providing an optimal design (i.e., determining the ESD type, and ESD provisioning at each level) and control (i.e., scheduling the charging and discharging of various ESDs) framework to maximize the *amortized profit* of the hierarchical ESD structure. The amortized one-time capital cost (capex), operating cost (opex), and cost associated with battery aging and replacement are considered in the profit optimization. Constraints on ESD volume and realistic characteristics of ESDs and power conversion circuitries are accounted for in the framework. (iv) conducting experiments using real data center workload traces from Google based on realistic data center specifications, demonstrating the effectiveness of the proposed design and control framework.**

I. Introduction

With the increase of the current large-scale and data-intensive Internet services, cloud computing has been gradually becoming a major paradigm in IT industry [1]. Recently, data center as a principal infrastructure to support this technology has been widely scrutinized and studied in both academia and industry areas.

The one-time construction expenses (called capital cost or capex) and monthly recurring operating expenses (opex) constitute the major investments of a data center [2]. The capex of a data center is proportional to the total provisioned IT facility power and estimated as $10-20 per Watt, as each Watt of computing power requires associated supporting equipment (cooling, backup, monitoring, etc.) [3]. The opex is charged by utility company based on dynamic energy pricing policy and high power tariff scheme [4]. Although the capex of a modern data center is significant, the opex of powering a data center has been steadily rising and could even surpass the capex of the data center itself [5]. As a result, reducing both capex and opex of a modern data center has become a key enabler to ensure its economic success.

Since the capex is dependent on the largest provisioning power and up to 40% opex is caused by the peak power tariff [6], power capping is generally employed in the modern data centers to reduce the peak power as to decrease the capex and opex. The power capping mainly focused on (i) exploiting server-level performance knobs such as DVFS (dynamic voltage and frequency scaling) in CPUs [7], [8], (ii) workload scheduling, server consolidation, and postponing delay-insensitive workloads under dynamic energy prices [9], and (iii) improving non-peak/idle power efficiency of the servers in data center [2], [10]. However, all the methods stated above incur performance overhead.

Recently, there is a trend in introducing and over-provisioning energy storage devices (ESDs) in data centers for facilitating the power provisioning and capex/opex reductions, which usually will not cause any performance overhead. The ESDs in a data center are commonly made of lead-acid batteries and utilized as centralized [11] uninterruptible power supply (UPS), which provides backup power to bridge the time gap between the power failure and the diesel generator startup. Generally, the time gap ranges from a few seconds to minutes. For over-provisioned UPS/ESDs, if the remaining energy is above some threshold to power the data center for a few minutes, it would be safe and beneficial to further exploit ESDs for power capping (power shaving) [2], [6], [12]. The profit analysis of over-provisioning ESDs for power capping and power shaving has been conducted in [2] and [12], demonstrating that the benefits from capex and opex reductions outweigh the extra costs associated with over-provisioning ESDs.

Some data centers recently built by large CSPs such as Google [13], Facebook [14], and Microsoft [15] employ the distributed single-level ESD structure, where ESDs are incorporated into the rack level or the server level of data center and directly connected to the corresponding DC power buses. The distributed single-level ESD structure demonstrates advantages over the centralized counterparts: it (i) eliminates a potential single point of failure, (ii) enables finer-granularity in energy storage control [2], and (iii) achieves less transmission line power loss and thereby higher power conversion efficiency due to the elimination of AC-DC-AC double conversion [4]. However, the distributed structure may encounter volume/real-estate constraints since the space inside each rack of data center is precious and limited, thereby limiting the ESD size and capability in performing power capping. Hence, a hierarchical ESD structure that places ESDs at data center, rack, and server levels has been proposed to take the advantages of both centralized and distributed ESD structures [4]. Moreover, reference work [4] also pointed out that different types of ESDs (e.g., Li-ion batteries, lead-acid batteries, supercapacitors, flywheels) have the potential to be integrated together in order to achieve high capability in power capping and high energy capacity simultaneously, since by proper deployment and control of multiple types of ESDs it is possible to exploit the benefit of each type of ESD while hiding its weaknesses [18], [19].

However, the prior work stated above has not taken into account a realistic power delivery architecture of the data center with a hierarchical ESD structure. In fact, without a proper design of power delivery architecture, it may fail to avoid AC-DC-AC double conversion or fail to directly connect rack or server-level ESDs to proper DC buses, thereby significantly degrading the overall energy efficiency (to

2159-6190/16 $31.00 © 2016 IEEE
DOI 10.1109/CLOUD.2016.88

be no higher than the centralized ESD structure.) Moreover, some key characteristics in the ESD system have been neglected in the control framework of the prior work, such as rate capacity effect, the most significant cause of power losses in lead-acid batteries [16], power losses in various AC/DC and DC/DC converters, and battery aging as a function of charging/discharging, etc. Finally, the incorporation of potentially different types of ESDs poses a unique research problem of selecting the most appropriate type of ESDs at each level of the data center, based on a comprehensive profit analysis.

To shed some light on the realistic benefits of hierarchical ESD framework, this paper adopts a realistic power delivery architecture from Intel [11] as the starting point, and makes the following contributions:

- We present novel and realistic power delivery architectures with hierarchical ESD structure, borrowing the best features from the centralized ESD structure from Intel [11] and distributed single-level ESD structures from Google [13], Facebook [14], and Microsoft [15], and avoiding AC-DC-AC double conversion to enhance efficiency.
- We provide realistic and highly efficient power delivery architectures with hierarchical ESD structure, supporting the case that different types of ESDs are employed for each of the data center, rack, and server levels.
- We provide a framework for the optimal design (i.e., determining the ESD type, and ESD provisioning at each level) and control (i.e., scheduling the charging and discharging of various ESDs) to maximize the *amortized profit* of the hierarchical ESD structure, or equivalently, minimize the *amortized total cost* including the amortized capex, opex, and cost of ESD aging and replacement. The proposed framework accounts for constraints on ESD volume, power losses due to rate capacity effect and conversion circuitry, and aging mechanisms of ESDs.
- We conduct experiments using real data center workload traces from Google based on realistic data center specifications, demonstrating the effectiveness of the proposed design and control framework of the hierarchical ESDs.

II. Hierarchical ESD Structure

In the centralized or distributed single-level ESD structures, ESDs are mainly utilized as the UPS for providing backup power to bridge the grid to diesel generator transition during power outage. Actually, it would be cost-effective to over-provision the ESDs and exploit them for power capping [2], [12]. Hierarchical ESD deployment is a natural choice for ESD over-provisioning because (i) rack-level or server-level ESDs can achieve high efficiency without redundancies due to direct connection to DC power buses, and (ii) the available space for rack-level or server-level ESDs alone may not be sufficient for performing power capping.

From the starting point of the Intel structure [11], we propose a realistic hierarchical ESD structure, borrowing the best features of centralized ESD structure from Intel and distributed ESD structures from Google [13] and Microsoft [15]. Please note that this hierarchical ESD structure can also accommodate the case that different types of ESDs are employed for each of the data center, rack, and server levels.

The proposed hierarchical ESD structure is shown in Fig. 1. Different from the centralized and distributed (single-level) ESD structures, ESDs are allocated to multiple levels, i.e., data center, rack, and server levels. High power supply reliability can be achieved without redundancies (due to the direct connections of rack and server-level ESDs). In addition, we employ a new type of data center-level UPS connection method that can be operated in either double-conversion mode or high efficiency mode by effectively controlling a set of programmable switches [17]. As reported in [17], the high efficiency mode, which bypasses input power from the grid to the PDU, could improve the power efficiency by up to 10% compared with the double-conversion mode. The time to switch between the double-conversion mode and the high efficiency mode is only one AC cycle (16.7ms in a 60Hz grid), which can be handled by rack/server-level ESDs or the server exception handlers.

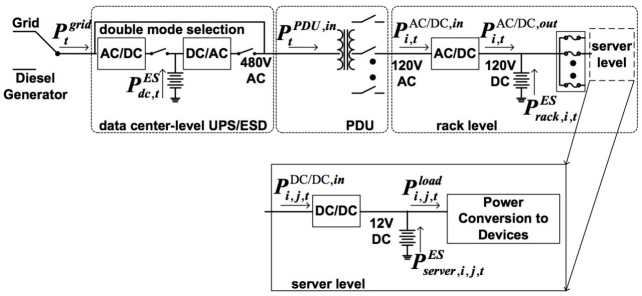

Fig. 1. The power delivery architecture for the hierarchical ESD structure.

The power delivery facilities can be classified into four stages. (i) 480V AC power after the data center-level UPS connection. (ii) The PDU transforms 480V AC into 120V AC and distributes to each rack in the data center. (iii) The 120V AC power is first rectified into 120V DC and then distributed to each server inside the rack. The rack-level ESDs are directly connected to 120V DC buses without power conversion circuitry. (iv) For each server, the 120V DC power is converted to 12V DC to feed the server, and the DC/DC converter is uni-directional due to capital cost considerations. The server-level ESDs are directly connected to 12V DC buses. The power conversion efficiency of conversion circuitries will be similar to the results from [11] (because similar power conversion circuitries are utilized) except for the case of high efficiency mode, which is close to 100% (97% in practice as reported in [11]).

A. Power Regulation in the Data Center

In this section we present the power regulation method in the data center, from the server level up to the data center level. The basic principle is that power conversion circuitries can regulate (control) their output voltage or current, and thereby regulate their output power.

At the server level, the load power demand at each time t, denoted by $P^{load}_{i,j,t}$ (of server j in rack i), is determined by the server workloads. The server-level DC bus is 12V and is determined by the terminal voltage of the directly connected server-level ESD. The server-level DC/DC converter regulates its output current and then its output power can also be determined. The output power of each server-level ESD, denoted by $P^{ES}_{server,i,j,t}$, can be determined as the difference between the output power of server-level DC/DC converter and load power demand. The rack level is similar. The load power demand for a rack is the summation of input power levels of each server-level DC/DC converter. The rack-level DC bus voltage is dictated by the rack-level ESD. Each rack-level AC/DC rectifier regulates its output power through regulating output current, and the output power $P^{ES}_{rack,i,t}$ of rack-level ESD can be determined accordingly.

At the data center level, a set of programmable switches will decide the operating mode: high efficiency mode or double-conversion

mode. In the high efficiency mode, the power drawn from the utility grid is (almost) the same as the input power $P_t^{\text{PDU},in}$ of the PDU. In the double-conversion mode, the input power of DC/AC inverter is determined by its output power, i.e., $P_t^{\text{PDU},in}$, and the conversion efficiency. The AC/DC rectifier regulates its output power through regulating output current, and the output power $P_{dc,t}^{ES}$ of data center-level ESD will be determined accordingly.

III. Energy Storage Devices

In this paper, lead-acid and Li-ion batteries are considered as candidates for ESD over-provisioning, due to their unique characteristics. Some important ESD characteristics are listed as follows. Data are derived from [18], [19].

Rate Capacity Effect and Charging/Discharging Efficiencies: In a data center, ESDs will incur a significant part of power loss besides the power conversion circuitry. Batteries suffer from the rate capacity effect, which is the major cause of battery power loss [16]. Battery's rate capacity effect specifies that the discharging (charging) efficiency of a battery decreases with the increase of the discharging (charging) current.

We denote the change rate of stored energy in battery by $P_{bat,in}$, which can be positive (discharging the storage), negative (charging the storage), or zero. Based on the modified Peukert's formulae [16], [21], the relationship between $P_{bat,in}$ and the output power of battery, P_{bat}, is given by

$$P_{bat} = \begin{cases} V_{bat} \cdot I_{bat,ref} \cdot (\frac{P_{bat,in}}{V_{bat} \cdot I_{bat,ref}})^{\beta_d}, & \frac{P_{bat,in}}{V_{bat} \cdot I_{bat,ref}} > 1 \\ -V_{bat} \cdot I_{bat,ref} \cdot (\frac{|P_{bat,in}|}{V_{bat} \cdot I_{bat,ref}})^{\beta_c}, & \frac{P_{bat,in}}{V_{bat} \cdot I_{bat,ref}} < -1 \\ P_{bat,in}, & \textit{otherwise} \end{cases} \tag{1}$$

where V_{bat} is the battery terminal voltage and is assumed to be (nearly) constant, I_{ref} denotes the reference current level for discharging and charging which is typically the current that can fully deplete the battery in 20 hours, and coefficients $\beta_d < 1$ and $\beta_c > 1$ are Peukert's coefficients.

For the convenience of expression, the above relationship in Eqn. (1) will be referred to as $P_{bat} = f_{bat}(P_{bat,in})$. One important observation is that this function is a concave and monotonically increasing function over the input domain $-\infty < P_{bat,in} < \infty$. The lead-acid battery suffers from more significant rate capacity effect (i.e., lower β_d and higher β_c values, or equivalently, lower discharging/charging efficiencies) compared with Li-ion battery.

Battery Aging: Another significant portion of power loss for batteries [16] is due to the state-of-health (SoH) degradation, i.e. the charge capacity of a battery will slowly degrade as the battery ages. The amount of SoH degradation, denoted by D_{SoH}, is defined as:

$$D_{SoH} = \frac{C_{full}^{nom} - C_{full}}{C_{full}^{nom}} \times 100\% \tag{2}$$

where C_{full}^{nom} is the nominal charge capacity of a new battery and C_{full} is the charge capacity of the battery in its current state. The state-of-charge (SoC) of a battery is defined as:

$$SoC = \frac{C_{bat}}{C_{full}} \times 100\% \tag{3}$$

where C_{bat} is the remaining charge stored in the battery.

The SoH degradation of a battery relates to a set of long-term electrochemical reactions inside the battery. These effects strongly depend on the operating condition of the battery such as number of charge/discharge cycles, SoC swing, average SoC, charging/discharging currents, and operation temperature [16], [18]. In order for effectively solving an optimization problem like ours, we adopt the SoH degradation model proposed in [20], which calculates the SoH degradation based on charge/discharge cycles and shows a good match with real data. The SoH degradation in one charge/discharge cycle, denoted by $D_{SoH,cycle}$, depends on the SoC swing (the difference between highest and lowest SoCs) and average SoC in the cycle, denoted by the function $D_{SoH,cycle}(SoC_{swing}, SoC_{avg})$. More specifically, it can be observed in [20], [21] that $D_{SoH,cycle}$ is a *superlinear and convex function* of SoC_{swing} and SoC_{avg}, in which SoC_{swing} has the dominant effect.

In literature, a battery reaches its end of life when its cumulative SoH degradation reaches a specific threshold (typically 20% to 30%). In general, if the ESD experiences 1 - 2 charge/discharge cycles in a day and the SoC swing (or called depth-of-discharge or DoD) in each cycle is restricted within 40% - 60%, a lead-acid ESD can operate for 1.5 - 2 years while a Li-ion ESD can operate for more than 5 years.

Self-Discharge Rate: The self-discharge rate is a measure of how quickly an ESD will lose its energy when it simply sits on the shelf. Typically lead-acid or Li-ion batteries exhibit negligible self-discharge compared with other "leaky" ESDs such as supercapacitors or flywheels.

Unit Capital Cost: The unit capital cost of a battery (or supercapacitor) ESD, given by \$/kWh, will significantly affect the ESD energy capacity provisioning results. The unit capital cost is \$50-150/kWh for lead-acid battery and \$400-600/kWh for Li-ion battery.

Energy Density: Energy density is calculated as the maximum stored energy divided by the volume of an ESD. The energy density is 50-80kWh/m^3 for a lead-acid battery, and 200-500kWh/m^3 for a Li-ion battery.

Power Density: Similarly, power density of an ESD is defined as the rated output power divided by the volume/weight. The power densities of lead-acid batteries and Li-ion batteries are given by 75-300W/kg and 250-340W/kg, respectively.

IV. Data Center ESD Design and Control Framework

We formulate an optimal data center ESD design (i.e., determining the ESD type, and ESD provisioning at each level) and control (i.e., scheduling the charging and discharging of various ESDs) problem to maximize the *amortized profit* of the hierarchical ESD structure, or equivalently, minimize the *amortized total cost* including the amortized capex, opex, and the cost of ESD aging and replacement. The proposed framework accounts for constraints on ESD volume (at each level), power losses due to rate capacity effect and conversion circuitry, as well as aging mechanisms of ESDs. We adopt a slotted time model in the problem formulation, in which the optimization time horizon (typically one day) is divided into T time intervals of equal length Δ_t.

A. Inputs

Data Center Specification: Let M denote the number of racks in the data center, and N denote the number of servers in a rack. Servers may be homogeneous or heterogenous in terms of computing, memory, and power characteristics.

Workload (Power Demand): There are many prior works on server workload modeling and prediction [22]. In this work, we assume that prior work on server load prediction can be leveraged and combined with power modeling work [23] to derive accurate power consumption time series at the server granularity over an optimization horizon. More specifically, let the power demand of the j-th server in i-th rack be denoted by $P_{i,j,t}^{load}$, where $t \in \{1, 2, ..., T\}$, $i \in \{1, 2, ..., M\}$,

and $j \in \{1, 2, ..., N\}$. We will adopt real Google cluster workloads [24] for evaluations.

Power Conversion Circuitry Efficiency: Let $\eta_{\text{AC/DC},dc}$, η_{PDU}, $\eta_{\text{AC/DC},rack}$, and $\eta_{\text{DC/DC}}$ denote the power conversion efficiency of the data center-level AC/DC rectifier and DC/AC inverter, the efficiency of PDU, the conversion efficiency of rack-level AC/DC rectifier, and the efficiency of server-level DC/DC converter (converting 120V DC to 12V DC), respectively. Please refer to Section II for typical efficiency values.

ESD Constraints: We are given the unit capital cost of each type of ESD by $cost^{LI}$ and $cost^{LA}$ in \$/kWh, where LI and LA denote "Li-ion" and "lead-acid", respectively. Similarly, the energy density values are denoted by d^{LI}_{energy} and d^{LA}_{energy} in kWh/m^3, and the power density values of ESD by d^{LI}_{power} and d^{LA}_{power} in kW/m^3. For example, the unit capital cost, energy density, power density of the lead-acid battery are \$50-150/kWh, 50-80kWh/m^3, and up to 400kW/m^3, respectively [18], [19]. For the space constraints, the available spaces for ESD at the three levels are given by L_{dc}, L_{rack}, and L_{server}, respectively.

Other ESD Characteristics: We are given rate capacity effect formula (Eqn. (1)) for Li-ion and lead-acid batteries, as well as SoH degradation functions $D^{LI}_{SoH,cycle}(SoC_{swing}, SoC_{avg})$ and $D^{LA}_{SoH,cycle}(SoC_{swing}, SoC_{avg})$ for Li-ion and lead-acid batteries, respectively.

Capex and Opex Components: The capex of a data center is proportional to the total provisioned IT facility power and estimated as \$10-20 per Watt (denoted by the parameter $CapEx\_per\_Watt$), as each Watt of computing power requires associated supporting equipment [3]. For the opex of data center, we assume a general and realistic day-ahead energy price function comprised of a *peak price* part (the peak power tariff) $Price\_Peak$ in \$/kW to charge the peak power consumption in a billing period (a day), as well as a time-of-day *energy price* part given by $Price_t$ in \$/kWh for $t \in \{1, 2, ..., T\}$.

B. System Equations

We list the system equations of the hierarchical ESD framework based on Kirchhoff's current law and energy conservation as follows.

$$P^{\text{DC/DC},in}_{i,j,t} = \frac{P^{load}_{i,j,t} - P^{ES}_{server,i,j,t}}{\eta_{\text{DC/DC}}} \tag{4}$$

$$P^{\text{AC/DC},in}_{i,t} = \frac{\sum_{j=1}^{N} P^{\text{DC/DC},in}_{i,j,t} - P^{ES}_{rack,i,t}}{\eta_{\text{AC/DC},rack}} \tag{5}$$

$$P^{\text{PDU},in}_{t} = \frac{\sum_{i=1}^{M} P^{\text{AC/DC},in}_{i,t}}{\eta_{\text{PDU}}} \tag{6}$$

$$P^{grid}_t = \begin{cases} P^{\text{PDU},in}_t, & \text{if } P^{ES}_{dc,t} = 0 \\ \dfrac{P^{\text{PDU},in}_t / \eta_{\text{AC/DC},dc} - P^{ES}_{dc,t}}{\eta_{\text{AC/DC},dc}}, & \text{otherwise} \end{cases} \tag{7}$$

for $i \in \{1, 2, ..., M\}$, $j \in \{1, 2, ..., N\}$, and $t \in \{1, 2, ..., T\}$. Eqns. (4), (5), (6) represent power flow equations at the server level, rack level, and PDU, respectively. On the other hand, Eqn. (7) represents power flow equation for the data center-level ESD. In the high efficiency mode ($P^{ES}_{dc,t} = 0$), we have $P^{grid}_t \approx P^{\text{PDU},in}_t$. In the double-conversion mode ($P^{ES}_{dc,t} \neq 0$), power from the grid must flow through the AC/DC rectifier and DC/AC inverter and then reach the PDU.

We use $E_{dc,t}$, $E_{rack,i,t}$, and $E_{server,i,j,t}$ to denote the energy storage at the end of time slot t in the data center ESD, ESD in i-th rack, and ESD in j-th server of i-th rack, respectively. Their initial values are represented by $E_{dc,0}$, $E_{rack,i,0}$, and $E_{server,i,j,0}$, respectively. For an example, the relationship between $E_{dc,t}$ and $E_{dc,0}$ satisfies:

$$E_{dc,t} = E_{dc,0} - \sum_{t'=1}^{t} P^{ES,in}_{dc,t'} \cdot \Delta_t \tag{8}$$

$E_{rack,i,t}$ and $E_{rack,i,0}$, and $E_{server,i,j,t}$ and $E_{server,i,j,0}$ satisfy similar relationships, which are omitted due to space limitation.

C. Optimization Variables

Given our goal to jointly address provisioning and subsequent control problems of the hierarchical ESD structure, we choose optimization variables that capture both the type and provisioning as well as the operational aspects of ESDs.

Design Problem: We need to decide the type and capacity of ESDs at three different levels. We use $Type_{dc}$, $Type_{rack}$, and $Type_{server}$ to denote the type of data center-level, rack-level, and server-level ESD, respectively, and use E^C_{dc}, E^C_{rack}, and E^C_{server} to denote the corresponding energy capacity values.

Control Problem: For the data center ESD power management, the control/optimization variables are the discharging/charging powers of all ESDs at different levels of the data center. Let $P^{ES}_{server,i,j,t}$ denote the ESD discharging power of the j-th server in the i-th rack, and the corresponding decrease rate of ESD stored energy is $P^{ES,in}_{server,i,j,t}$. The relationship between $P^{ES}_{server,i,j,t}$ and $P^{ES,in}_{server,i,j,t}$ is given by the rate capacity effect Eqn. (1). Similarly we define $P^{ES}_{rack,i,t}$'s, $P^{ES,in}_{rack,i,t}$'s, $P^{ES}_{dc,t}$'s and $P^{ES,in}_{dc,t}$'s.

D. Objective Function

The objective function to maximize is the *amortized profit* over the billing period. Equivalently, we **minimize the amortized total cost** over the billing period $[0, T \cdot \Delta_t]$, given by:

$$\frac{CapEx}{Number\_of\_Days} + OpEx + ESD\_Degradation\_Cost \tag{9}$$

where $CapEx = CapEx\_per\_Watt \cdot \max_t P^{grid}_t$ is the capex for provisioning supporting equipments (cooling, backup, monitoring, etc.), and $Number\_of\_Days$ represents the total operating days of the data center. Hence $\frac{CapEx}{Number\_of\_Days}$ is the amortized capex over the billing period. Moreover, $OpEx$ consists of a peak price part and a time-of-day energy price part, given by:

$$OpEx = Price\_Peak \cdot \max_t P^{grid}_t + \sum_{t=1}^{T} Price_t \cdot P^{grid}_t \cdot \Delta_t \tag{10}$$

Finally, the cost of ESD aging and replacement, $ESD\_Degradation\_Cost$, is calculated as:

$$cost^{Type_{dc}} \cdot E^C_{dc} \cdot D_{SoH,dc} + cost^{Type_{rack}} \cdot E^C_{rack} \sum_{i=1}^{M} D_{SoH,rack,i} + cost^{Type_{server}} \cdot E^C_{server} \cdot \sum_{i=1}^{M} \sum_{j=1}^{N} D_{SoH,server,i,j} \tag{11}$$

where $D_{SoH,dc}$, $D_{SoH,rack,i}$, $D_{SoH,server,i,j}$ are the SoH degradation values of the data center-level ESD, ESD in rack i, and ESD in server j of rack i, respectively. In order to quickly estimate such SoH degradation values for run-time control optimization, we make the assumption that **each ESD only encounters one charging/discharging cycle in each day**, which is typically true since the data center power consumption often exhibits single peak in a day [24]. As an example of SoH degradation estimation under such assumption, we have

$$D_{SoH,server,i,j} \approx D^{Type_{server}}_{SoH,cycle}(SoC_{swing,server,i,j}, SoC_{avg,server,i,j}) \tag{12}$$

where $SoC_{swing,server,i,j}$ and $SoC_{avg,server,i,j}$ represent the SoC swing and average SoC of the ESD in server j of rack i, respectively.

E. System Constraints

We provide a list of constraints in the ESD design and control framework as follows:

Volume Constraint: The ESD volumes at data center, rack, and server-level are constrained by the available spaces at those three levels:

$$\frac{E^C_{dc}}{d^{Type_{dc}}_{energy}} \leq L_{dc}, \frac{E^C_{rack}}{d^{Type_{rack}}_{energy}} \leq L_{rack}, \frac{E^C_{server}}{d^{Type_{server}}_{energy}} \leq L_{server} \quad (13)$$

where $Type_{dc}$, $Type_{rack}$, and $Type_{server}$ can be either "LA" or "LI".

Maximum Power Constraint: The absolute value of the maximum output power of data center, rack, and server-level ESDs are constrained by power density constraints:

$$\left|P^{ES}_{dc,t}\right| \leq \frac{E^C_{dc}}{d^{Type_{dc}}_{energy}} \cdot d^{Type_{dc}}_{power}, \left|P^{ES}_{rack,i,t}\right| \leq \frac{E^C_{rack}}{d^{Type_{rack}}_{energy}} \cdot d^{Type_{rack}}_{power}$$
$$\left|P^{ES}_{server,i,j,t}\right| \leq \frac{E^C_{server}}{d^{Type_{server}}_{energy}} \cdot d^{Type_{server}}_{power} \quad (14)$$

for $i \in \{1, 2, ..., M\}$, $j \in \{1, 2, ..., N\}$, and $t \in \{1, 2, ..., T\}$.

Energy Storage Constraint: The energy storage in a data center, rack, and server-level ESD cannot be lower than a lower bound LB and cannot exceed 100% of corresponding energy capacity:

$$LB \cdot E^C_{dc} \leq E_{dc,t} \leq E^C_{dc}, LB \cdot E^C_{rack} \leq E_{rack,i,t} \leq E^C_{rack}$$
$$LB \cdot E^C_{server} \leq E_{server,i,j,t} \leq E^C_{server} \quad (15)$$

for $i \in \{1, 2, ..., M\}$, $j \in \{1, 2, ..., N\}$, and $t \in \{1, 2, ..., T\}$. The lower bound is included to ensure the availability of ESDs, i.e., for providing backup power when power outage occurs.

Operation Constraint: Because uni-directional conversion circuitry is utilized in racks and servers (due to cost considerations), the following constraints hold:

$$P^{ES}_{server,i,j,t} \leq P^{load}_{i,j,t} \quad (16)$$

$$P^{ES}_{rack,i,t} \leq \sum_{j=1}^{N} P^{\text{DC/DC},in}_{i,j,t} \quad (17)$$

for $i \in \{1, 2, ..., M\}$, $j \in \{1, 2, ..., N\}$, and $t \in \{1, 2, ..., T\}$.

Moreover, the data center-level ESD connection does not support selling electric power back to grid, i.e.,:

$$P^{ES}_{dc,t} \leq \frac{P^{\text{PDU},in}_t}{\eta_{\text{AC/DC},dc}} \quad (18)$$

for $t \in \{1, 2, ..., T\}$.

V. Control Optimization of Hierarchical ESD Framework

In this section, we provide detailed formulation and solution method for the optimal ESD control in the hierarchical ESD framework, whereas the design optimization problem will be described and solved in the next section. For the control optimization problem, the type and energy capacity of ESDs at different levels, i.e., $Type_{dc}$, $Type_{rack}$, $Type_{server}$, E^C_{dc}, E^C_{rack}, and E^C_{server}, are given parameters (from the design optimization problem.) The objective is to derive the optimal values of run-time control variables $P^{ES}_{dc,t}$'s, $P^{ES}_{rack,i,t}$'s, and $P^{ES}_{server,i,j,t}$'s of all ESDs in the coming billing period, in order to minimize the amortized total cost defined in Eqn. (9). We also need to satisfy system equations, i.e., Eqns. (4) - (8), and constraints, i.e., Eqns. (14) - (18).

A. Three-Step Solution of the Control Optimization Problem

1) Server-Level Control Optimization: First, we formulate and provide optimal solution of the server-level control optimization. For the j-th server in the i-th rack, we aim to derive the optimal values of control variables $P^{ES}_{server,i,j,t}$'s, so as to minimize the amortized total cost **seen from** this specific server, given by:

$$CapEx\_per\_Watt \cdot \max_t P^{\text{DC/DC},in}_{i,j,t} / Number\_of\_Days +$$
$$Price\_Peak \cdot \max_t P^{\text{DC/DC},in}_{i,j,t} + \sum_{t=1}^{T} Price_t \cdot P^{\text{DC/DC},in}_{i,j,t} \cdot \Delta_t +$$
$$cost^{Type_{server}} \cdot E^C_{server} \cdot$$
$$D^{Type_{server}}_{SoH,cycle}(SoC_{swing,server,i,j}, SoC_{avg,server,i,j}) \quad (19)$$

where $P^{\text{DC/DC},in}_{i,j,t}$ is a linearly decreasing function of $P^{ES}_{server,i,j,t}$ as defined in Eqn. (4). The four items in the objective function represent amortized capex, peak price component of opex, time-of-day energy price component of opex, and ESD degradation cost, respectively. We need to satisfy system constraints (14), (15), and (16).

The above-mentioned server-level control optimization problem is non-convex because Eqn. (15) is not a convex constraint. We transform the server-level control problem into standard convex optimization problem with convex objective function and inequality constraints, and linear equality constraints. More specifically, we use $P^{ES,in}_{server,i,j,t}$'s instead of $P^{ES}_{server,i,j,t}$'s as optimization variables for solving the problem[1]. The server-level control problem is equivalent before and after changing optimization variables due to the one-to-one correspondence between $P^{ES}_{server,i,j,t}$ and $P^{ES,in}_{server,i,j,t}$. We provide the proof on the convexity of the server-level control problem after changing optimization variables:

Proof: We need to prove that (i) the objective function is convex function, (ii) inequality constraints are convex, and (iii) equality constraints are linear. The objective function (19) is a convex function of $P^{ES,in}_{server,i,j,t}$'s due to the following reasons:

(i) $P^{ES}_{server,i,j,t}$ is a concave function of $P^{ES,in}_{server,i,j,t}$, and thus $P^{\text{DC/DC},in}_{i,j,t}$ is a linearly decreasing function of $P^{ES}_{server,i,j,t}$. Thus $P^{\text{DC/DC},in}_{i,j,t}$ is a convex function of $P^{ES,in}_{server,i,j,t}$.

(ii) The capex and opex components in (19) are convex functions of $P^{ES,in}_{server,i,j,t}$'s because the pointwise maximum of a set of convex functions is convex.

(iii) $SoC_{swing,server,i,j}$ and $SoC_{avg,server,i,j}$ are linear functions of $P^{ES,in}_{server,i,j,t}$'s (Please refer to Section IV). The function $D^{Type_{server}}_{SoH,cycle}(SoC_{swing,server,i,j}, SoC_{avg,server,i,j})$ is convex and increasing with respect to $SoC_{swing,server,i,j}$ and $SoC_{avg,server,i,j}$. Thus $D^{Type_{server}}_{SoH,cycle}(SoC_{swing,server,i,j}, SoC_{avg,server,i,j})$ is convex function of $P^{ES,in}_{server,i,j,t}$'s [25].

For the constraints, Eqn. (15) becomes a linear constraint because $E_{server,i,j,t}$ is a linear function of $P^{ES,in}_{server,i,j,t}$'s (described in Section V-C). Eqn. (14) and (16) are still linear constraints because of the one-to-one correspondence between $P^{ES}_{server,i,j,t}$ and $P^{ES,in}_{server,i,j,t}$. □

Since we have proved the convexity on the server-level control problem after changing optimization variables, the problem can be solved optimally in polynomial time complexity using standard convex optimization solvers such as CVX [26] or fmincon function in MATLAB. The number of optimization variables is equal to the number of time slots (288 if each time slot is 5 minutes) in the control

[1]Please note that the system control still controls storage output powers $P^{ES}_{server,i,j,t}$'s during system operation. The only modification is that we use internal energy changes $P^{ES,in}_{server,i,j,t}$'s as optimization variables when deriving the optimal control solution.

optimization problem for each server, and therefore the solution has reasonable time complexity.

2) Rack-Level and Data Center-Level Control Optimizations: After the optimal control problems of all servers in a rack have been solved, we proceed with the rack-level control optimization. For each rack i, the power flowing into each uni-directional DC/DC converter inside the rack, $P_{i,j,t}^{\text{DC/DC},in}$, is already given (from server-level optimizations), and we aim to derive the optimal values of control variables $P_{rack,i,t}^{ES}$'s, in order to minimize the amortized total cost **seen from** this specific rack, given by:

$$CapEx\_per\_Watt \cdot \max_t P_{i,t}^{\text{AC/DC},in}/Number\_of\_Days+$$
$$Price\_Peak \cdot \max_t P_{i,t}^{\text{AC/DC},in} + \sum_{t=1}^{T} Price_t \cdot P_{i,t}^{\text{AC/DC},in} \cdot \Delta_t +$$
$$cost^{Type_{rack}} \cdot E_{rack}^{C} \cdot$$
$$D_{SoH,cycle}^{Type_{rack}}(SoC_{swing,rack,i}, SoC_{avg,rack,i}) \quad (20)$$

where $P_{i,t}^{\text{AC/DC},in}$ is a linearly decreasing function of $P_{server,i,j,t}^{ES}$ as defined in Eqn. (5). Constraints of rack-level control optimization problem include Eqns. (14), (15), and (17).

Similar to the server-level control optimization, we transform the rack-level control optimization problem into standard convex optimization by using $P_{rack,i,t}^{ES,in}$'s as optimization variables. The proof is similar to that for server-level control optimization and omitted due to space limitation. Thus the rack-level problem can also be solved optimally in polynomial time using standard convex optimization solvers. The number of optimization variables is again equal to the number of time slots for rack-level control optimization.

After the optimal control problems of all racks have been solved, we proceed with the data center-level optimization. At this time the power flowing into the PDU in each time slot, $P_t^{\text{PDU},in}$, is already given (from rack-level optimizations), and we aim to derive the optimal values of control variables $P_{dc,t}^{ES}$'s, in order to minimize the amortized total cost seen from the whole data center, i.e., the amortized total cost calculated in Eqn. (9). Constraints of data center-level control optimization include Eqns. (14), (15), and (18).

Similarly, we transform the data center-level control optimization into standard convex optimization, which can be solved optimally in polynomial time complexity, by utilizing $P_{dc,t}^{ES,in}$'s as optimization variables. The proof of convexity is similar. The number of optimization variables is again equal to the number of time slots.

VI. Design Optimization of Hierarchical ESD Framework

The design optimization problem of hierarchical ESD framework properly determines the optimal type of ESDs at different levels, i.e., $Type_{dc}$, $Type_{rack}$, and $Type_{server}$, along with optimal energy capacities of ESDs at different levels, i.e., E_{dc}^{C}, E_{rack}^{C}, and E_{server}^{C}. A straightforward search-based algorithm would search all the possible values of the six variables and result in a complexity of $O(S^3)$ where S is the search precision level of each of variables E_{dc}^{C}, E_{rack}^{C}, and E_{server}^{C} (on the other hand, searching $Type_{dc}$, $Type_{rack}$, and $Type_{server}$ has constant complexity because they are binary variables.) Although more efficient search algorithms like ternary searches can be employed to enhance the search speed, more efficient solution algorithm of the design optimization is still necessary.

We propose optimal and effective design optimization algorithm to reduce search complexity based on the idea of dynamic programming. More specifically, for each possible value pair of $Type_{server}$ and E_{server}^{C}, we solve the server-level control optimization problem (described in Section VI-B) and minimize the total amortized cost seen from the servers. We record the total amortized cost seen from the servers as a function of $Type_{server}$ and E_{server}^{C} in a table. Next, for each possible value pair of $Type_{rack}$ and E_{rack}^{C}, we derive (i) the best $Type_{server}$ and E_{server}^{C} values and (ii) the rack-level control optimization (described in Section VI-B), in order to minimize the total amortized cost seen from the racks. This procedure is achieved by (i) searching all possible $Type_{server}$ and E_{server}^{C} values in the table and (ii) solving the rack-level control optimization as a convex optimization problem. Again, we record the (best found) total amortized cost seen from the racks as a function of $Type_{rack}$ and E_{rack}^{C}. Finally, for each possible value pair of $Type_{dc}$ and E_{dc}^{C}, we derive (i) the best-suited $Type_{rack}$ and E_{rack}^{C} values and (ii) the data center-level control optimization (described in Section VI-B), in order to minimize the total amortized cost of the data center. This procedure is achieved by (i) searching all possible $Type_{rack}$ and E_{rack}^{C} values in the table and (ii) solving the data center-level control optimization as a convex optimization problem.

It can be observed that the proposed design optimization algorithm has a reduced solution complexity of $O(S^2)$. This is because we successfully iterate over only two levels (server and rack levels, or rack and data center levels) instead of three levels, by effectively storing the partial solution in a table.

VII. Experimental Results

Experimental results are provided for both the control optimization problem and design optimization problem of the hierarchical ESD structure. Control optimization evaluation is shown in Section VIII-A while design optimization evaluation is shown in Section VIII-B. We consider a realistic data center setup similar to that in [4], which comprises 8,192 servers placed in 256 racks with 32 servers/rack, and exhibits 4 MW peak power consumption. The power delivery and hierarchical ESD structure has been discussed in Section III. In the experiments, we set the billing period (time horizon) to be one day and time slot $\Delta_t = 5$ minutes.

Google cluster trace is used as realistic workloads for our evaluations [25]. The Google cluster trace released in 2012 is measured in a 29-day period including 672,075 jobs and more than 48 million tasks. The (normalized values of) CPU, memory, and disk utilizations of the server cluster are measured and recorded in every 5 minutes. We derive load power consumption $P_{i,j,t}^{load}$ of each server based on CPU and memory utilization traces and accurate server power modeling [28].

For the capex component in the objective function, we use parameter $CapEx\_per\_Watt = \$15$ per Watt, and $Number\_of\_Days = 365 \times 20 = 7,300$ (assuming 20 years life-time of the power delivery infrastructure). For the opex component, we adopt a realistic dynamic energy pricing policy similar to the one of LADWP [29], which comprises a time-of-day energy price component and a peak price component (high power tariff). The energy price component is given by: 0.01879 \$/kWh during 00:00 to 09:59 and 20:00 to 23:59, 0.03952 \$/kWh during 10:00 to 12:59 and 17:00 to 19:59, 0.04679 \$/kWh during 13:00 to 16:59. The peak price component is given by 0.575\$/kW to charge the peak power consumption over the whole day. For the battery unit capital cost and SoH degradations, related parameters are imported from reference [18], [19].

A. Control Optimization Evaluation

We provide evaluation results on the control optimization given ESD types and provisioning results in this section, using all 29 days of Google cluster trace. We consider two hierarchical ESD structures, one using lead-acid batteries and the other with Li-ion batteries. For

both (lead-acid based and Li-ion based) structures, we use 0.24kWh as the energy capacity for each server-level ESD, which corresponds to 3L in volume for lead-acid batteries when assuming an energy density of 80kWh/m^3, or 0.75L in volume for Li-ion batteries when assuming an energy density of 320kWh/m^3. Let 32×0.24kWh be the energy capacity for each rack-level ESD for both (lead-acid based and Li-ion based) structures, where 32 is the number of servers in a rack. Finally, let 8192×0.24kWh be the energy capacity for each data center-level ESD for both structures, where 8192 is the number of servers in a data center.

We compare the two proposed systems with three baseline systems. Baseline 1 adopts the lead-acid battery-based centralized ESD structure used in [11] and the ESD/UPS is not utilized for power capping. Baseline 1 serves as the basis for our comparisons because the state-of-the-art ESD/UPS in data centers are made of lead-acid batteries and are not yet utilized for power capping. **The *amortized profit* of our proposed systems are compared with respect to Baseline 1.** Baseline 2 also uses lead-acid battery-based centralized ESD structure and the ESD/UPS is optimally controlled for amortized cost minimization/amortized profit maximization. The energy capacity of the UPS in Baseline 2 equals to the summation of energy capacities of all ESDs in the proposed system. Baseline 3 adopts the same hierarchical ESD structure and uses ESDs for amortized cost minimization/amortized profit maximization. However, the control algorithm does not take into account rate capacity effect and battery aging for ESDs.

Fig. 2 provides the comparison results among the two proposed systems and three baselines. More specifically, Fig. 2 (a) provides the (daily) amortized total cost of the two proposed systems and three baselines, which is an average of the 29-day period. On the other hand, Fig. 2 (b) provides the amortized profit with respect to Baseline 1 which does not utilize centralized ESD/UPS for power capping. It can be clearly observed that the two proposed systems significantly outperform the three baselines in amortized total cost saving, with a maximum reduction in amortized total cost of about 72.6%. Between the two proposed systems, the Li-ion battery-based system outperforms the lead-acid battery-based one due to its higher efficiency (less significant rate capacity effect). Among all baselines, Baseline 1 performs the worst since it uses the least optimized structure (centralized ESD structure) and does not use ESDs for power capping, while Baseline 2 performs the best since it adopts the optimal control algorithm of the centralized ESD to perform amortized cost minimization/amortized profit maximization.

For more detailed comparison results, Fig. 3 (a) provides the amortized capex component of the two proposed systems and three baselines. Fig. 3 (b) and (c) provides the comparison results on the opex and ESD degradation cost of each day among the five systems. It can be observed that (i) the proposed systems significantly outperform baselines in terms of amortized capex reduction (due to the control optimization algorithm), although they may result in higher ESD degradation cost compared with baselines (due to the more often charging/discharging of ESDs); (ii) The Li-ion battery based system outperforms the lead-acid battery based system in capex reduction (due to the higher efficiency of Li-ion batteries) while it only incurs slightly higher ESD degradation (although Li-ion batteries have higher unit capital cost compared with lead-acid batteries, they also have longer life time/lower SoH degradation rate.)

B. Design Optimization

In this section, we provide evaluation results on the design optimization problem, which decides the optimal type and provisioning

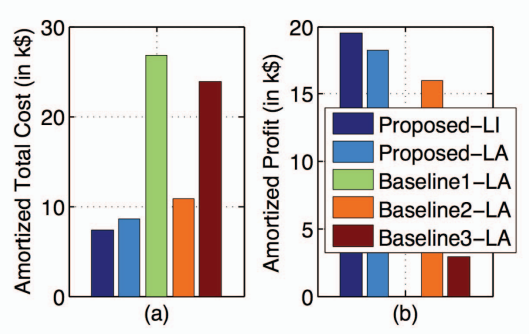

Fig. 2. Comparison results on (a) amortized total cost, (b) amortized profit among the two proposed systems and three baselines.

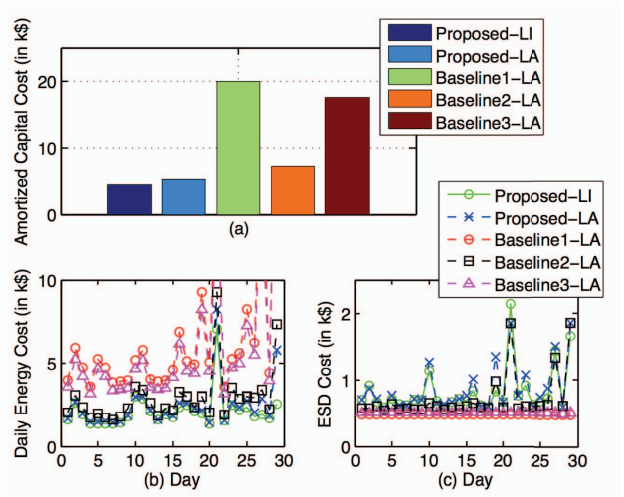

Fig. 3. Comparison results on (a) amortized capital cost component, (b) opex of each day, and (c) ESD degradation cost of each day among the two proposed systems and three baselines.

results of data center, rack, and server-level ESDs. In the experiments, the volume constraint of data center-level ESD is 8192×3L and we change the volume constraints on server-level and rack-level ESDs. This is for realistic consideration because in reality server and rack-level ESDs are subject to more stringent volume constraint compared with data center-level ones [4] (data centers often use outside located UPS's for data center-level UPS/ESD). We use two sets of volume constraints for the server-level and rack-level ESDs. The first set of constraints uses 1L as the server-level ESD volume constraint and 32×1L as the rack-level ESD volume constraint, whereas the second set of constraints uses 3L as the server-level ESD volume constraint and 32×3L as the rack-level ESD volume constraint. We compare the optimally designed hierarchical ESD system with two baselines: Baseline 1 adopts the hierarchical structure and uses the largest possible data center, rack, and server-level ESDs using lead-acid batteries, whereas Baseline 2 adopts the hierarchical structure and uses the largest possible data center, rack, and server-level ESDs using Li-ion batteries.

Fig. 4 (a) and (b) demonstrate the comparison results on amortized total cost and amortized profit, respectively, between the optimally designed system with two baselines. The server-level ESD volume constraints constitute the x-axis of Fig. 4. When the server-level vol-

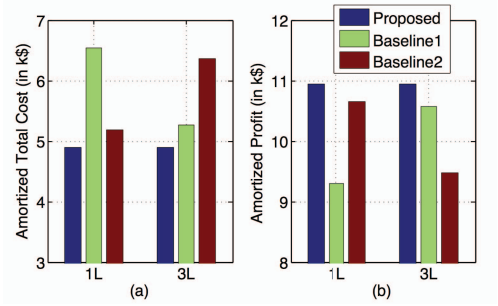

Fig. 4. Comparison results on (a) amortized total cost and (b) amortized profit for the design optimization problem.

ume constraint is 1L (and rack-level 32×1L), the optimally designed hierarchical ESD system uses $10\% \times 8192 \times 3$L, $40\% \times 32 \times 1$L, and 1L as volumes of data center, rack, and server-level ESDs, respectively, all using Li-ion batteries. When the server-level volume constraint is 3L (and rack-level 32×3L), the optimally designed system uses $10\% \times 8192 \times 3$L, $40\% \times 32 \times 1$L, and 1.04L as volumes of data center, rack, and server-level ESDs, respectively, all using Li-ion batteries. One can draw the following conclusions: (i) The optimally designed system consistently outperforms baseline systems in amortized total cost reduction or equivalently amortized profit maximization; (ii) Li-ion batteries are more preferred than lead-acid batteries (or a combination of Li-ion and lead-acid ones) because the benefits of higher efficiency and longer life-time outweigh the disadvantage of higher unit capital cost; (iii) Server-level ESD is more preferred in the optimal design of the hierarchical ESD system than rack-level ESD, whereas data center-level ESD is the least preferred; (iv) The optimal solution under 1L and 3L server-level volume constraints are close to each other, illustrating that large-size Li-ion battery based ESDs are not preferred due to the high associated battery degradation cost; (v) Comparing the two baseline systems, Li-ion battery based baseline performs better when server-level volume constraint is 1L whereas lead-acid battery based baseline performs better when server-level constraint is 3L. This is because of the higher unit capital cost of Li-ion battery compared with lead-acid ones.

VIII. Conclusion

In this paper we presented a design and control framework of hierarchical ESD structure. The proposed framework especially supports the integration of different types of ESDs in order to exploit the advantages of each type while hiding the weaknesses. More specifically, we (i) adopted a realistic power delivery facility from Intel for centralized ESD structure as the starting point; (ii) presented a realistic power delivery architecture, borrowing the best features of both centralized and distributed single-level ESD structures, and supporting the case that different types of ESDs are employed for each of the data center, rack, and server levels; (iii) presented an optimal design (i.e., determining the ESD type, and ESD provisioning at each level) and control (i.e., scheduling the charging and discharging of various ESDs) framework to maximize the amortized profit of the hierarchical ESD structure; (iv) conducted extensive experiments based on real data center workload traces from Google and realistic data center specifications.

References

[1] M. Armbrust, A. Fox, R. Griffith, et al., "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50-58, Apr. 2010.

[2] L. A. Barroso, and U. Holzle, "The case for energy-proportional computing," *IEEE Computer*, vol. 40, no. 12, pp. 33-37, Dec. 2007.

[3] V. Kontorinis, L. E. Zhang, B. Aksanli, et al., "Managing distributed UPS energy for effective power capping in data centers," *Proc. of ISCA*, Jun. 2012.

[4] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in data centers: what, where, and how much?" *Proc. of SIGMETRICS*, pp. 187-198, London, UK, Jun. 2012.

[5] X. Fan, W. D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," *Proc. of ISCA*, Jun. 2007.

[6] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Levaraging stored energy for handling power emergencies in aggressively provisioned datacenters," *Proc. of ASPLOS*, pp. 75-86, Mar. 2012.

[7] R. Cochran, C. Hankendi, A. K. Coskun, and S. Reda, "Pack & cap: adaptive DVFS and thread packing under power caps," *Proc. of MICRO*, 2011.

[8] A. Gandhi, V. Gupta, M. Harchol-Balter, and M. Kozuch, "Optimality analysis of energy-performance trade-off for server farm management," *Proc. of SIGMETRICS*, 2010.

[9] A. Gandhi, M. Harchol-Balter, R. Das, and C. Lefurgy, "Optimal power allocation in server farms," in *Proc. of SIGMETRICS*, 2009.

[10] D. Meisner, C. M. Sadler, L. A. Barroso, W. Weber, and T. F. Wenisch, "Power management of online data-intensive services," *Proc. the 38th Annual International Symposium on Computer Architecture (ISCA)*, 2011.

[11] M. Ton and B. Fortenbury, "High performance buildings: data centers uninterruptible power supplies (UPS)", Dec. 2005. http://energy.lbl.gov/ea/mills/_archive/HT-3-13-08/documents/UPS/Final_UPS_Report.pdf

[12] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and limitations of tapping into stored energy for datacenters," in *Proc. the 38th Annual International Symposium on Computer Architecture (ISCA)*, 2011.

[13] Google Datacenter Video Tour, 2009. http://www.google.com/about/datacenters/efficiency/external/2009-summit.html

[14] Facebook Open Compute Project, 2011. http://www.opencompute.org

[15] Microsoft Reveals its Specialty Servers, Racks, 2011. http://www.datacenterknowledge.com/archives/2011/04/25/microsoft-reveals-its-specialty-servers-racks

[16] D. Linden and T. B. Reddy, *Handbook of Batteries*, McGraw-Hill, 2002.

[17] M. Ton, B. Fortenbery, and W. Tschudi, "DC power for improved data center efficiency," Mar. 2008. http://energy.lbl.gov/ea/mills/HT/documents/data_centers/DCDemoFinalReport.pdf

[18] H. Chen, T. N. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding, "Progress in electrical energy storage system: A critical review," *Progress in Natural Science*, 2009.

[19] M. Pedram, N. Chang, Y. Kim, and Y. Wang, "Hybrid electrical energy storage systems," in *Proc. of International Symposium on Low Power Electronics and Design (ISLPED)*, 2010.

[20] A. Millner, "Modeling Lithium Ion battery degradation in electric vehicles," in *Proc. of IEEE CITRES*, 2010.

[21] Q. Xie, X. Lin, Y. Wang, M. Pedram, D. Shin, and N. Chang, "State of health aware charge management in hybrid electrical energy storage systems," in *Proc. of Design, Automation, and Test in Europ (DATE)*, 2012.

[22] M. F. Arlitt and C. L. Williamson, "Internet web servers: workload characterization and performance implications," *IEEE Trans. Networking*, 1997.

[23] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: a framework for architectural-level power analysis and optimizations," *Proc. of ISCA*, Jun. 2000.

[24] "Google cluster data" 2012, [Online]. Available: https://code.google.com/p/googleclusterdata/.

[25] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[26] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21." http://cvxr.com/cvx, Feb. 2011.

[27] Los Angeles Department of Water & Power, Electric Rates, http://www.ladwp.com/ladwp/cms/ladwp001752.jsp.

[28] M. Pedram and I. Hwang, "Power and performance modeling in a virtualized server system," in *Green Computing Workshop conjunction with International Conference on Parallel Processing*, 2010.

[29] Consolidated Edison Company of New York, Inc., "Service Classification No.1 - Residential and Religious", 2012.