Hierarchical Deployment and Control of Energy Storage Devices in Data Centers

Shuo Wang, Yanzhi Wang, Xue Lin, and Massoud Pedram
Department of Electrical Engineering
University of Southern California, Los Angeles, CA 90089
{wang639, yanzhiwa, xuelin, pedram}@usc.edu

***Abstract*—Recent work has presented hierarchical deployment of energy storage devices (ESDs) at the data center, rack, and server levels within a data center, along with a corresponding control framework for peak power shaving and energy cost reduction under (time-of-use) dynamic energy pricing policies. However, the prior work does not use a realistic power delivery architecture of the data center with hierarchical ESD structure, and fails to account for some key characteristics such as rate capacity effect of batteries and power losses in various AC/DC and DC/DC converters in the power delivery architecture. This paper aims to overcome these shortcomings by (i) adopting a realistic power delivery architecture (from Intel) for centralized ESD structure as the starting point; (ii) presenting a novel power delivery architecture for data centers with hierarchical ESD structure, borrowing the best features of the centralized ESD structure from Intel and the distributed single-level ESD structures from Google and Microsoft; (iii) providing a mathematical framework for the optimal design (i.e., ESD provisioning) and control (i.e., scheduling the charging and discharging of various ESDs) of the hierarchical ESD structure to minimize overall energy cost under dynamic energy pricing functions. This framework accounts for constraints on ESD volume (for each level) and the overall (annually amortized) capital cost, and power losses due to the rate capacity effect and conversion circuitry. The ESD design problem is solved by using a search-based algorithm, whereas the ESD control problem is formulated and solved as a hierarchical convex optimization algorithm. Experiments have been conducted using real Google cluster workload based on realistic data center specifications, demonstrating the effectiveness of the proposed optimal design and control framework.**

Keywords*—*data centers; energy storage devices (ESDs); hierarchical ESD structure

I. Introduction

Cloud computing as the long-held dream of computing as a utility [1] has demonstrated its success in the commercialization nowadays. The data center hardware and software are what we will call a *cloud*, and the service being sold is *utility computing* [2]. The cloud computing system collects computation and storage requests from distributed and decentralized clients/users over the Internet, and processes these requests in a centralized manner in the data center. A data center has been effectively managed through the wide utilization of virtual machines (VMs).

In a cloud computing system, cloud service providers (CSPs) benefit from charging the clients for their requested services and from the efficient and cost-effective centralized computation and storage resource management. On the other hand, the clients are attracted by reducing their expenses on building and maintaining new servers on site and by the convenient and high-quality computation and storage services via the high-speed modern network. The clients must have guarantees from the CSPs on the quality-of-service (QoS) as specified in the Service Level Agreements (SLAs), which include requirements and guarantees on computing power, storage space, network bandwidth, availability, security, etc [3], [4]. In addition, it is essential for a market-oriented data center to reduce the operational cost by judicious resource provisioning and task dispatching to ensure commercial success.

The operational cost of a data center due to energy consumption has been steadily increasing as the development of high performance platforms, whereas the hardware cost of a data center has been maintained stable [5]. Recently, it is found that the energy cost of a data center during its lifetime accounting for the dynamic energy pricing and high power tariff could possibly surpass the hardware cost [4], [6], [7]. Therefore, it is imperative to suppress this trend and maintain the operational cost of a data center due to energy consumption within a reasonable range. There have been plenty of studies on reducing the energy cost of a data center or server cluster while satisfying the SLA constraints, through a variety of methods: (i) exploiting server-level performance knobs such as DVFS (dynamic voltage and frequency scaling) in CPUs [8], (ii) scheduling, placement/consolidation, and migration of computation across servers [9], [10], [11], and (iii) reducing energy loss within the overall power infrastructure [12]. Recent work [13], [14] have proposed to reduce the data center electric bill under dynamic electric energy pricing, by performing workload scheduling/postponement to match electricity price.

Recently, there is a trend in introducing energy storage devices (ESDs) in data centers for facilitating the power management, which usually will not cause any performance overhead. The ESDs in a data center are commonly made of lead-acid batteries and utilized as a centralized 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 seconds to minutes. Therefore, the energy stored in the UPS should be abundant for powering the data center for a few minutes. Furthermore, the ESDs can be exploited for peak power shaving of data centers, i.e., storing energy during off-peak periods and providing energy during high-peak periods [15], [16], under dynamic energy pricing policies. This technique can effectively reduce the energy cost and power infrastructure cost of a data center under the dynamic energy pricing and high power tariff mechanisms.

Some data centers recently built by large CSPs such as Google [17], Facebook [18] and Microsoft [19] employ the distributed single-level ESD structure, in which 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 counterpart: (i) it achieves larger degree of reliability (i.e., higher fault-tolerance capability) and enables finer-granularity in energy storage control [20], (ii) it achieves less transmission line power loss and thereby higher power conversion efficiency due to the elimination of AC-DC-AC double conversion [21]. However, the distributed single-level ESD structure may encounter serious volume/real estate constraints since the space inside each data center rack is precious and limited, thereby restricting the ESD size and capability. Therefore, 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 [21]. However, the prior work does not use a realistic power delivery architecture of the data center with hierarchical ESD structure. In fact, without a proper design of power delivery architecture, it may end up failing to avoid AC-DC-AC double conversion or failing to directly connect rack-level or server-level ESDs to proper DC buses, thereby significantly degrading the overall energy efficiency (to be no higher than centralized ESD structure.) Also, some key characteristics in the ESD system have been neglected in the control framework in the prior work, such as rate capacity effect, the most significant cause of power losses in

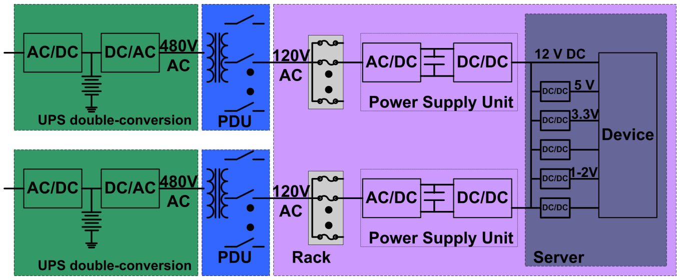

Fig. 1. The power delivery architecture of a data center with the centralized ESD structure.

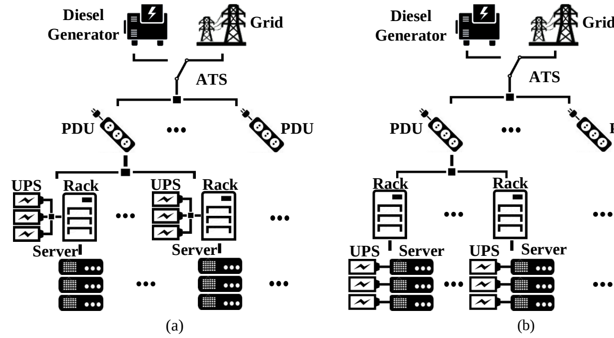

Fig. 2. The power delivery architectures of a data center with the distributed single-level ESDs proposed by (a) Microsoft [19] and (b) Google [17].

lead-acid battery storage [24], power losses in various AC/DC and DC/DC converters, and so on.

To mitigate the gap in literature and shed some light on the realistic benefits of hierarchical ESD framework, this paper (i) adopts a realistic power delivery architecture from Intel [22] as the starting point, (ii) presents a novel and realistic power delivery architecture for data centers with hierarchical ESD structure, borrowing the best features of the centralized ESD structure from Intel [22] and distributed ESD structures from Google [17] and Microsoft [19], and effectively avoiding AC-DC-AC double conversion to enhance efficiency, (iii) provides a mathematical framework for the optimal design (i.e., ESD provisioning at the data center, rack, and server levels) and control (i.e., scheduling the charging and discharging of various ESDs) of the hierarchical ESD structure to minimize overall energy cost under (time-of-use) dynamic energy pricing functions. This framework accounts for constraints on ESD volume (for each level) and the overall (annually amortized) capital cost, and power losses due to the rate capacity effect and conversion circuitry. The ESD design problem is solved by using a search-based algorithm, whereas the ESD control problem is formulated and solved as a hierarchical convex optimization algorithm [30]. Experiments have been conducted using real Google cluster workload [28] based on realistic data center specifications [21], demonstrating the effectiveness of the proposed design and control framework of the hierarchical ESD structure.

II. Background

A. Centralized ESD Structure

Generally, in a data center with the centralized ESD structure, the power delivery facilities can be classified into four stages. Let us use the data center with the centralized ESDs proposed by Intel [22] as an example to illustrate the power delivery architecture. As shown in Fig. 1, the four stages of the power delivery facilities are as follows: (i) The 480V AC power provided by the grid or alternatively by the diesel power generator goes through the AC-DC-AC double-conversion with the centralized ESDs (i.e., the UPS) connected in between. (ii) The output power (480V AC) of the AC-DC-AC double-conversion structure is transformed to 120V AC and distributed to each rack in the data center through the power distribution unit (PDU). (iii) The 120V AC power at each rack is then distributed to each server inside the rack. (iv) At the server level, the 120V AC power is rectified into DC and then converted to 12V DC power through the power supply unit (PSU). The 12V DC power is supplied to the server and may be further converted into lower-voltage DC power by the voltage conversion modules (VCMs) to feed different components/devices in the server. A survey conducted by Intel shows that the power delivery facilities in a data center may result in more than 30% power loss [22]. In other words, the efficiency of the UPS (AC-DC-AC double-conversion) is 85%-90%, the efficiency of the PDU is about 98%, and the efficiency of the PSU is 73%-90% [22]. The PSU results in the lowest conversion efficiency because of the large difference between input and output voltages (120V and 12V, respectively.)

The power delivery facilities in a data center are designed to fulfill the power quality and reliability requirement. When power outage happens, the power supply will be switched from the grid to the alternative diesel power generator by the automatic transfer switch (ATS). However, bringing the diesel generator online takes seconds or even minutes, during which the ESDs (e.g., batteries or flywheels) in the UPS are used to provide backup power. Therefore, the capacity of ESDs in the UPS should be large enough for covering the transition time gap from grid to diesel generator. In addition, the double-conversion structure of the UPS can guarantee seamless transition from the grid to the ESDs and then to the diesel generator. Furthermore, a redundant UPS (as shown in Fig. 1) is adopted to further improve the power reliability of a data center in case one UPS malfunctions or temporarily shuts off for maintenance. The PDU and PSU are also configured with redundancies accordingly.

B. Distributed Single-Level ESD Structure

To mitigate the considerable power loss and the high capital cost of the power delivery facilities with centralized ESDs, the distributed single-level ESD structure has been proposed with simplified power delivery architecture, higher power supply reliability, and higher power delivery efficiency. In contrast to the centralized ESD structure, the UPS's are relocated from the data center level down to the rack/server level, and therefore the ESDs are distributed to each rack/server. Fig. 2 (a) shows the distributed rack-level ESD structure proposed by Microsoft [19], in which ESDs are directly connected (without conversion circuitry) to the rack-level DC bus, and Fig. 2 (b) shows the distributed server-level ESD structure proposed by Google [17], in which ESDs are directly connected to the server-level DC bus[1]. When power outage happens, the distributed ESDs can immediately take over the power supply to each rack/server during the transition time gap from grid to diesel generator.

The distributed ESD structure demonstrates multiple advantages over the centralized counterpart: (i) Since the malfunction or failure of some of the UPS's only affects part of the racks/servers, high power supply reliability can be achieved without UPS redundancy, which reduces the capital cost of the data center. (ii) The ESDs are distributed to each rack/server and directly connected to the corresponding DC buses, and therefore the AC-DC-AC double-conversion structure in the centralized structure is removed, which increases the overall power delivery efficiency of a data center. Google reported that the achieved efficiency improvement corresponds to a yearly cost saving of about $30/server [17]. (iii) The distributed ESD structure has the potential for finer-granularity energy storage control in data center power management [20], [21]. However, the distributed single-level ESD structure may encounter serious volume/real estate constraints since the space inside each data center rack is precious and limited, thereby restricting the ESD size and capability.

[1] On the other hand, Facebook uses dedicated "UPS servers" as distributed ESD structure [18].

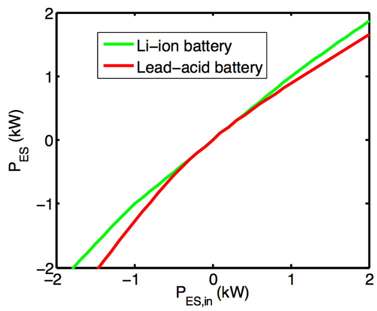

Fig. 3. The rate capacity effect of Li-ion battery and lead-acid battery.

C. Energy Storage Devices

In a data center, the ESDs will incur a significant part of power loss besides the power conversion circuitry. Batteries (lead-acid or Li-ion batteries, especially lead-acid ones) are the widely adopted ESDs due to their good reliability, high energy density, low self-discharge rate, etc. However, batteries suffer from the rate capacity effect, which is the major cause of the battery power loss [23]. The battery rate capacity effect specifies that a high battery discharging current may reduce the amount of energy to be extracted from the battery and a high charging current may also reduce the amount of energy to be stored into the battery. To put simply, the discharging (charging) efficiency of a battery decreases with the increase of the discharging (charging) current. The Peukert's formula captures the rate capacity effect by specifying the battery discharging (charging) efficiency i.e., $\eta_{rate,d}$ ($\eta_{rate,c}$) as a function of the discharging (charging) current i.e., I_d (I_c) [24]:

$$\eta_{rate,d}(I_d) = \frac{1}{(I_d/I_{ref})^{\alpha_d}}, \qquad (1)$$

$$\eta_{rate,c}(I_c) = \frac{1}{(I_c/I_{ref})^{\alpha_c}},$$

where α_d and α_c are Peukert's coefficients ranging from 0.1 to 0.3 for different type of batteries, and I_{ref} denotes the reference current proportional to the nominal capacity C_{nom} (in Ah) of the battery. Typically, I_{ref} is set as $C_{nom}/20$, indicating that it takes 20 hours to fully discharge the battery if using a discharging current of $I_{ref} = C_{nom}/20$. Eqn. (1) is not accurate when the normalized discharging (charging) current i.e., I_d/I_{ref} (I_c/I_{ref}) is less than one, since the calculated efficiency is greater than 100% in this case. To fix this problem, a modified Peukert's formula is proposed in [23] to set the charging/discharging efficiency to 100% when I_d/I_{ref} (I_c/I_{ref}) is less than one.

Let us denote the battery discharging power by P^{ES} and the decrease rate of the battery stored energy by $P^{ES,in}$. Please note that P^{ES} and $P^{ES,in}$ can be positive for battery discharging and negative for battery charging. Then, according to the modified Peukert's formula, the relationship between P^{ES} and $P^{ES,in}$ is given by

$$P^{ES} = \begin{cases} V^{ES} \cdot I_{ref} \cdot (\dfrac{P^{ES,in}}{V^{ES} \cdot I_{ref}})^{\beta_1}, \text{ if } \dfrac{P^{ES,in}}{V^{ES} \cdot I_{ref}} > 1, \\ P^{ES,in}, \text{ if } -1 \le \dfrac{P^{ES,in}}{V^{ES} \cdot I_{ref}} \le 1, \\ -V^{ES} \cdot I_{ref} \cdot (\dfrac{\left|P^{ES,in}\right|}{V^{ES} \cdot I_{ref}})^{\beta_2}, \text{ if } \dfrac{P^{ES,in}}{V^{ES} \cdot I_{ref}} < -1. \end{cases} \qquad (2)$$

In Eqn. (2), V^{ES} is the battery terminal voltage, I_{ref} is the battery reference current, coefficient β_1 is in the range of 0.8-0.9, and coefficient β_2 is in the range of 1.1-1.3. The rate capacity effect of Li-ion battery and lead-acid battery is shown in Fig. 3, from which we can observe that lead-acid battery based ESDs have more significant power loss due to the rate capacity effect than Li-ion battery based ESDs.

Some other battery characteristics need to be taken into account for incorporating such ESDs into the data center, listed as follows. Data are derived from [32].

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

Energy and Power Densities: Energy density is calculated as the maximum stored energy divided by the volume of an ESD. Similarly, power density of an ESD is defined as the rated output power divided by the volume. For example, the energy density is 50 - 80kWh/m^3 for a lead-acid battery, and 200 - 500kWh/m^3 for a Li-ion battery. On the other hand, both types of batteries have relatively high power density.

Cycle Life: The cycle life of an ESD is defined as the number of charge/discharge cycles an ESD can perform before its capacity drops to a specific percentage (60% - 80% typically) of its initial fully-charged capacity. It is the key performance parameter as an indication of the expected working lifetime of an ESD. In general, if the ESD experiences 1 - 2 charge/discharge cycles in a day and the depth-of-discharge (DoD) in each cycle is restricted within 40% - 60% of the full state-of-charge range, 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 quick 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.

Although Li-ion batteries have superior efficiency (less significant rate capacity effect), higher energy/power density, and longer cycle life compared with lead-acid batteries, the latter is more widely adopted for data center usage because of the capital cost and safety considerations.

III. Hierarchical ESD Structure

In the centralized or distributed single-level ESD structure discussed previously, the ESDs are mainly utilized as the UPS for providing backup power to bridge the grid to diesel generator transition during power outage. Actually, the ESDs also have potential applications in peak power shaving and power demand shifting under dynamic energy pricing mechanism [20]. Moreover, in the previously discussed structures, the ESDs are allocated to single level, i.e., data center level in the centralized structure and rack or server level in the distributed structure. However, the power demands at different levels of a data center should be treated distinctly due to the hierarchy of a data center [7], and the available space for rack-level or server-level ESDs alone may not be sufficient for performing peak power shaving. In this work, from the starting point of the centralized ESD structure from Intel [22], we propose a realistic hierarchical ESD structure to not only guarantee the power supply reliability but also achieve finer-granularity power management of a data center, borrowing the best features of centralized ESD structure from Intel [22] and distributed ESD structures from Google [17] and Microsoft [19].

The proposed hierarchical ESD structure is shown in Fig. 4. Different from the centralized and distributed (single-level) ESD structures, ESDs are allocated to multiple levels, i.e., data center level, rack level and server level. In addition, we employ a new type of data center-level UPS connection method presented by some UPS manufacturers that can be operated in either double-conversion mode or high efficiency mode by effectively controlling a set of programmable switches for the data center-level power management

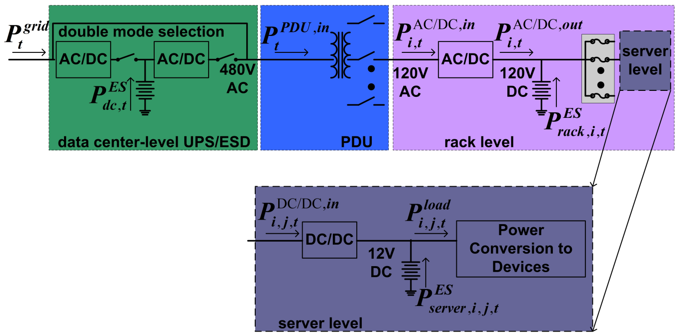

Fig. 4. The power delivery architecture of a data center with the hierarchical ESD structure.

[25][2]. As reported in [25], 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 [25]), which can be handled by rack/server-level UPS or the server exception handlers.

The power delivery facilities of a data center with the hierarchical ESD structure 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[3]. The rack-level ESDs are directly connected to 120V DC buses without power conversion circuitry. We adopt bi-directional AC/DC rectifier (i.e., a rectifier together with a DC/AC inverter) between 120V AC power and 120V DC power to allow the power flowing between different rack-level ESDs [21]. (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. Similar to the results from [22], the power conversion efficiency of the data center-level UPS is 88% in the double-conversion mode and close to 100% (97% in practice) in the high efficiency mode [25]. The efficiency of the PDU is about 98% [22]. The power conversion efficiencies of the rack-level AC/DC rectifier/inverter and the server-lever DC/DC converter are about 95% and 90%, respectively.

The hierarchical ESD structure combines the advantages of both the centralized and the distributed ESD structures while hiding their weaknesses. When power outage happens, the hierarchical ESDs in each level can immediately provide backup power during the grid to diesel generator transition time gap. Also, high power supply reliability can be achieved without redundancies. More importantly, due to the hierarchy, we can realize more flexible and finer-granularity power management for a data center taking into account the battery rate capacity effect, power losses in various AC/DC and DC/DC converters, ESD volume and capital cost constraints, etc.

IV. Data Center ESD Design and Control Framework

To take full advantage of the proposed hierarchical ESD structure, in this section we formulate a data center ESD design (provisioning) and control problem to minimize the data center energy cost, taking into account the dynamic energy pricing mechanism, ESD volume and (annually amortized) capital cost constraints, power losses due to the rate capacity effect and conversion circuitry, etc. More specifically, the ESD design problem provisions ESDs at the data center, rack, and server levels within a data center, whereas the ESD control problem determines how each ESD in the hierarchical architecture is utilized, i.e., charged and discharged, with a given ESD provisioning result.

A slotted time model is adopted in the following problem formulation, and all input values, decision variables, and system constraints are provided with T discrete time intervals of equal length Δ_t, and the time horizon of our optimization is $T \cdot \Delta_t$.

A. Inputs

Data Center Specification: Let M denote the number of racks in the data center of interest, and N denote the number of servers in a rack. A realistic data center may be homogeneous or heterogenous in terms of server characteristics [28].

Workload (Power Demand): There are considerable prior studies on server workload characterization and prediction [26]. For this work, we assume that prior work on load prediction can be leveraged and combined with power modeling work [27] to derive reasonable and accurate time-series of power consumptions at the server granularity over the given optimization horizon (e.g., one day). Specifically, we assume the power demand time-series of the j-th server in the i-th rack is given by $P^{load}_{i,j,t}$, where $t \in \{1, 2, ..., T\}$, $i \in \{1, 2, ..., M\}$, and $j \in \{1, 2, ..., N\}$. Please refer to Fig. 4 for the precise definition of $P^{load}_{i,j,t}$. We will utilize real Google cluster workloads [28] in our 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 AC/DC rectifier and DC/AC inverter for data center-level UPS, the efficiency of the PDU, the power conversion efficiency of the rack-level bi-directional AC/DC rectifier/inverter, and the efficiency of the server-level uni-directional DC/DC converter (converting 120V DC to 12V DC for each server), respectively. Please refer to the analysis in Section III.

Dynamic Energy Pricing: We assume a day-ahead energy price function given by $Price_t$ in \$/kWh for $t \in \{1, 2, ..., T\}$. Our framework is generic enough to accommodate other types of price function, such as the peak price function [29].

ESD Constraints: We are given the unit capital cost of ESD by $cost_{ESD}$ in \$/kWh, the energy density of ESD by d_{energy} in kWh/m^3, the power density of ESD by d_{power} in kW/m^3, and the cycle life time of ESD by $CycleLife$. For instance, the unit capital cost, energy density, power density, and cycle life time of the lead-acid battery are \$50-150/kWh, 50-80kWh/m^3, up to 400kW/m^3, and 1.5 - 2 years, respectively [32]. Due to the space limitation, at each level the available space for a single ESD are given by L_{dc}, L_{rack}, and L_{server}. The limit of the annually amortized capital cost of ESDs in the whole data center is $CostLimit_{dc}$.

B. Optimization Variables

Given our goal to jointly address provisioning and subsequent control problems, we choose optimization variables that capture the operational aspects of ESDs as well as the decisions about their sizing and placement.

For the Design Problem: In order to design the hierarchical ESDs for the data center, we need to decide the capacity of ESDs at different levels. We use E^C_{dc}, E^C_{rack}, and E^C_{server} to denote the energy capacity of each data center-level, rack-level, and server-level ESD, respectively.

For the Optimal Control Problem: In the data center ESD power management, the control variables are the discharging/charging powers of all the ESDs at different levels of the data center. We use $P^{ES}_{server,i,j,t}$ to denote the time-series of discharging power of the ESD of the j-th server in the i-th rack, and the corresponding decrease rate of the 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 defined in the rate capacity effect Eqn. (2). 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.

[2]Fig. 4 uses switches to illustrate the high efficiency mode (bypassing) and double-conversion mode, which captures the basic principle. However, the actual implementation is more sophisticated [25] to ensure instant switching between these two modes.

[3]Please note that it is different from the Intel infrastructure shown in Fig. 1 since the Intel structure first distributes 120V AC power to each server and then rectify to DC power.

C. System Model

In this part, we derive the mathematical system model of the hierarchical ESD framework based on Kirchhoff's current law and energy conservation.

For each j-th server in the i-th rack, the power flowing into the uni-directional DC-DC converter powering this server is given by

$$\frac{P_{i,j,t}^{load} - P_{server,i,j,t}^{ES}}{\eta_{\text{DC/DC}}} \tag{3}$$

Then the power flowing out of the bi-directional AC/DC rectifier/inverter for each i-th rack, denoted by $P_{i,t}^{\text{AC/DC},out}$, is the summation of currents flowing into all uni-directional DC-DC converters in this rack minus the output power of rack-level ESD:

$$P_{i,t}^{\text{AC/DC},out} = \sum_{j=1}^{N} \frac{P_{i,j,t}^{load} - P_{server,i,j,t}^{ES}}{\eta_{\text{DC/DC}}} - P_{rack,i,t}^{ES} \tag{4}$$

Then the power flowing into each bi-directional AC/DC rectifier/inverter for the i-th rack, denoted by $P_{i,t}^{\text{AC/DC},in}$, is given by:

$$P_{i,t}^{\text{AC/DC},in} = \begin{cases} \dfrac{P_{i,t}^{\text{AC/DC},out}}{\eta_{\text{AC/DC},rack}}, & \text{if } P_{i,t}^{\text{AC/DC},out} \geq 0 \\ P_{i,t}^{\text{AC/DC},out} \cdot \eta_{\text{AC/DC},rack}, & \text{if } P_{i,t}^{\text{AC/DC},out} < 0 \end{cases} \tag{5}$$

Since the AC/DC rectifier/inverter is bi-directional, the power values $P_{i,t}^{\text{AC/DC},out}$ and $P_{i,t}^{\text{AC/DC},in}$ could be positive (when power flowing into this rack) or negative (when power flowing out of this rack.)

Given all the $P_{i,t}^{\text{AC/DC},in}$ values, the power flowing into the data center PDU at time slot t, denoted by $P_{t}^{\text{PDU},in}$, is calculated by:

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

Finally, the data center power consumption drawn from the grid, denoted by P_t^{grid}, satisfies:

$$P_t^{grid} = \begin{cases} P_{t}^{\text{PDU},in} - P_{dc,t}^{ES} \cdot \eta_{\text{AC/DC},dc}, & \text{if } P_{dc,t}^{ES} \geq 0 \\ P_{t}^{\text{PDU},in} - \dfrac{P_{dc,t}^{ES}}{\eta_{\text{AC/DC},dc}}, & \text{if } P_{dc,t}^{ES} < 0 \end{cases} \tag{7}$$

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

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

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

D. Objective Function

We minimize the overall energy cost of the data center over the time horizon $[0, T \cdot \Delta_t]$, given by:

$$\sum_{t=1}^{T} Price_t \cdot P_t^{grid} \cdot \Delta_t \tag{9}$$

E. Constraints

The following constraints need to be satisfied in the data center ESD design and control framework:

Volume Constraint: The volume of ESD at the data center-level, rack-level, and server-level should be restricted by the available spaces at those three levels, e.g.,

$$\frac{E_{dc}^C}{d_{energy}} \leq L_{dc}, \frac{E_{rack}^C}{d_{energy}} \leq L_{rack}, \frac{E_{server}^C}{d_{energy}} \leq L_{server} \tag{10}$$

Capital Cost Constraint: The total annually amortized capital cost of data center-level, rack-level, and server-level ESDs should be restricted by the total annually amortized capital cost constraint $CostLimit_{dc}$, i.e.,

$$\frac{cost_{ESD} \cdot (E_{dc}^C + M \cdot E_{rack}^C + MN \cdot E_{server}^C)}{CycleLife} \leq CostLimit_{dc} \tag{11}$$

Maximum Output Power Constraint: The maximum output/input power of data center-level, rack-level, and server-level ESD should be restricted by the power density constraints, i.e.,

$$\left|P_{dc,t}^{ES}\right| \leq \frac{E_{dc}^C}{d_{energy}} \cdot d_{power}, \left|P_{rack,i,t}^{ES}\right| \leq \frac{E_{rack}^C}{d_{energy}} \cdot d_{power}$$
$$\left|P_{server,i,j,t}^{ES}\right| \leq \frac{E_{server}^C}{d_{energy}} \cdot d_{power} \tag{12}$$

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

Energy Storage Constraint: At any time slot t, the energy storage in a data center-level, rack-level, and server-level ESD cannot exceed 100%, and cannot be lower than a lower bound LB portion, of the corresponding energy capacity, i.e.,

$$LB \cdot E_{dc}^C \leq E_{dc,t} \leq E_{dc}^C, LB \cdot E_{rack}^C \leq E_{rack,i,t} \leq E_{rack}^C$$
$$LB \cdot E_{server}^C \leq E_{server,i,j,t} \leq E_{server}^C \tag{13}$$

for $i \in \{1, 2, ..., M\}$, $j \in \{1, 2, ..., N\}$, and $t \in \{1, 2, ..., T\}$. We add the lower bound constraint (usually about 40%) due to two reasons: (i) Certain amount of energy storage needs to be maintained in each ESD at any time to be prepared for providing backup power when power outage occurs. (ii) It is desirable to maintain the depth-of-discharge of ESDs within certain bound due to state-of-health (SoH) degradation and cycle life considerations [33].

Operation Constraint: Since we adopt uni-directional DC-DC converter for servers (please refer to Fig. 4), we have the following constraint

$$P_{server,i,j,t}^{ES} \leq P_{i,j,t}^{load} \tag{14}$$

for $i \in \{1, 2, ..., M\}$, $j \in \{1, 2, ..., N\}$, and $t \in \{1, 2, ..., T\}$. Also the data center-level UPS connection mode adopted in this work (please refer to [25]) does not support selling electric power back to the grid. In other words, we have the following constraint:

$$P_{dc,t}^{ES} \cdot \eta_{\text{AC/DC},dc} \leq P_t^{\text{PDU},in} \tag{15}$$

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

V. Optimization Framework

In this section, we separate the overall hierarchical ESD design and control framework into a control optimization problem and a design optimization problem, and provide effective solutions for both problems.

A. The Control Optimization Problem

In the control optimization problem, the energy capacity of ESDs at different levels, i.e., E_{dc}^C, E_{rack}^C, and E_{server}^C, are given parameters, and we need to determine the optimal values of run-time control variables $P_{dc,t}^{ES}$'s, $P_{rack,i,t}^{ES}$'s, and $P_{server,i,j,t}^{ES}$'s of all ESDs in order to minimize the overall energy cost.

Let us take an example to illustrate the solution complexity of the control optimization problem. Consider a data center in the size of Google cluster [28] or the one considered in [21] with the operation time horizon of one day. The Google cluster is comprised of 7,000 servers, and the data center considered in [21] is comprised of 8,192 servers. The number of time slots is $24 \cdot 12 = 288$ if each time slot is set to be 5 minutes. Hence, the total number of optimization variables in the optimal control problem exceeds 2,000,000, which is far beyond the capability of standard optimization solving methods (for example, the solution complexity of a convex optimization problem is in the cubic order of the number of optimization variables [30].)

In order to effectively solve the optimal control problem by reducing the number of optimization variables, we adopt a hierarchical solution comprised of server-level, rack-level, and data center-level optimizations. For each j-th server in the i-th rack, we aim to find the optimal values of control variables $P^{ES}_{server,i,j,t}$'s, in order to minimize the energy cost **seen from** the single server, which is given by:

$$\sum_{t=1}^{T} Price_t \cdot \frac{P^{load}_{i,j,t} - P^{ES}_{server,i,j,t}}{\eta_{\text{DC/DC}}} \cdot \Delta_t \tag{16}$$

Constraints in this server-level optimal control problem include (12), (13) and (14).

We transform the server-level optimal control problem into standard convex optimization problems such that it could be optimally solved in polynomial time complexity by standard convex optimization tools such as CVX [31] or fmincon function in MATLAB. More specifically, we utilize $P^{ES,in}_{server,i,j,t}$'s instead of $P^{ES}_{server,i,j,t}$'s as optimization variables in such standard convex optimization solvers[4]. One can observe that the objective function (16) is a convex function of $P^{ES,in}_{server,i,j,t}$'s since (i) it is a decreasing function of $P^{ES}_{server,i,j,t}$'s, and (ii) $P^{ES}_{server,i,j,t}$ is a concave function of $P^{ES,in}_{server,i,j,t}$ according to the modified rate capacity effect Eqn. (2). Moreover, constraints (12), (13), and (14) are all linear (and then convex) inequality constraints of $P^{ES}_{server,i,j,t}$'s. Hence, we conclude that the server-level optimal control problem is a convex optimization problem of optimization variables $P^{ES}_{server,i,j,t}$'s since it has convex objective function and convex inequality constraints. The number of optimization variables is equal to the number of time slots (288 if each time slot is 5 minutes and time horizon is one day) in the server-level optimal control problem for each server, and therefore the solution has reasonable time complexity.

After the optimal control problems of all servers have been solved, we proceed with the rack-level optimization. For each i-th rack, the power flowing into each uni-directional DC-DC converter inside the rack in each time slot, given by $\frac{P^{load}_{i,j,t} - P^{ES}_{server,i,j,t}}{\eta_{\text{DC/DC}}}$, is already given, and we aim to find the optimal values of control variables $P^{ES}_{rack,i,t}$'s, in order to minimize the energy cost **seen from** the single rack, which is given by:

$$\sum_{t=1}^{T} Price_t \cdot P^{\text{AC/DC},in}_{i,t} \cdot \Delta_t \tag{17}$$

where $P^{\text{AC/DC},in}_{i,t}$ for each $t \in \{1, 2, ..., T\}$ is calculated by Eqn. (5). Constraints in this rack-level optimal control problem include (12) and (13). Similar to the server-level optimization, we transform the rack-level optimization into standard convex optimization problems by utilizing $P^{ES,in}_{rack,i,t}$'s as optimization variables in standard convex optimization solvers. The number of optimization variables is again equal to the number of time slots in this rack-level optimal control problem.

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 data center PDU in each time slot, $P^{\text{PDU},in}_{t}$, is already given, and we aim to find the optimal values of control variables $P^{ES}_{dc,t}$'s in order to minimize the overall energy cost given in Eqn. (9). Constraints in this optimization level include (12), (13), and (15). Similarly, we transform the data center-level optimization problem into standard convex optimization problems by utilizing $P^{ES,in}_{dc,t}$'s as optimization variables in standard convex optimization solvers.

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

Algorithm 1: The design optimization solution of the hierarchical ESD structure in a data center.

Input: System inputs given in Section IV.A, annually amortized capital cost constraint $CostLimit_{dc}$, volume constraints L_{dc}, L_{rack}, and L_{server}

Output: Optimal energy capacity of ESDs at different levels, i.e., E^C_{dc}, E^C_{rack}, and E^C_{server}

for *each* E^C_{dc} *value (using ternary search here)* **do**
  for *each* E^C_{rack} *value (using ternary search here)* **do**
    Calculate the corresponding E^C_{dc} value satisfying annually amortized capital cost constraint;
    Check whether the volume constraints are satisfied;
    Execute the control optimization algorithm to minimize the overall energy cost if volume constraints are satisfied;

Return the optimal design configuration E^C_{dc}, E^C_{rack}, and E^C_{server} of the hierarchical ESD framework such that overall energy cost is minimized

B. The Design Optimization Problem

The design optimization problem of hierarchical ESD framework properly determines the optimal energy capacity of ESDs at different levels, i.e., E^C_{dc}, E^C_{rack}, and E^C_{server}. We propose an optimal algorithm for the design optimization problem as shown in Algorithm 1. The proposed algorithm searches over all possible E^C_{dc}, E^C_{rack}, and E^C_{server} combinations with the annually amortized capital cost constraint **just satisfied**, i.e., $\frac{cost_{ESD} \cdot (E^C_{dc} + M \cdot E^C_{rack} + MN \cdot E^C_{server})}{CycleLife} = CostLimit_{dc}$ (in this case E^C_{server} is determined as long as E^C_{dc} and E^C_{rack} are given.) For each possible combination of the three design variables, the proposed algorithm (i) evaluates whether the volume constraints are satisfied, and (ii) executes the control optimization algorithm to minimize the overall energy cost if volume constraints are satisfied[5]. The proposed algorithm then chooses among all combinations of E^C_{dc}, E^C_{rack}, and E^C_{server} to select the optimal design configuration of the hierarchical ESD framework. The ternary search method, which is an extension over the well-known binary search algorithm based on the quasi-convex assumption of the underlying problem, is exploited to accelerate the search of the optimal E^C_{dc} and E^C_{rack} values.

VI. Experimental Results

In this section, we provide experimental results on both the control optimization problem and design optimization problem of the hierarchical ESD structure. Our evaluations use a realistic data center setup similar to [21]. The data center of interest exhibits 4 MW peak power consumption, and is comprised of 8,192 servers placed in 256 racks with 32 servers/rack. The power delivery and hierarchical ESD structure have been discussed in Section III. We use lead-acid battery based ESDs and place those ESDs at the data center, rack, and server levels. In the control optimization results in Section VI.A, we demonstrate the effectiveness of the proposed hierarchical optimization method under given energy capacity of ESD in each level. In the design optimization results in Section VI.B, we optimize the energy capacity of ESD in each level subject to volume and annually amortized capital cost constraints. In these experiments, we set the optimization time horizon to be one day and time slot $\Delta_t = 5$ minutes.

Our evaluations use two different day-ahead (time-of-use) dynamic energy pricing policies as shown in Fig. 5. The first one is a

[5]Please note that our optimization framework is general and not necessarily need to only execute the control optimization problem from time slot 1 to T. We can execute the control optimization problem for multiple days and use the average energy cost for the design optimization problem.

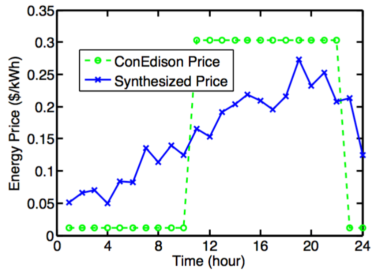

Fig. 5. Two dynamic energy pricing functions: (i) synthesized pricing, (ii) conEdison pricing.

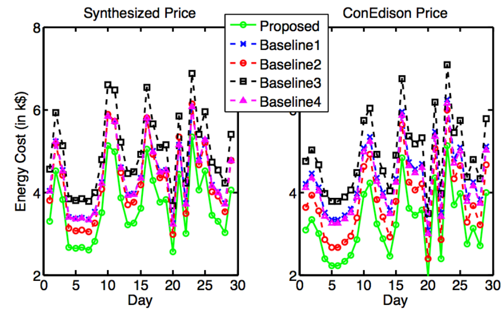

Fig. 6. Energy cost achieved by the proposed hierarchical ESD control versus baselines over all 29 days using Google cluster trace.

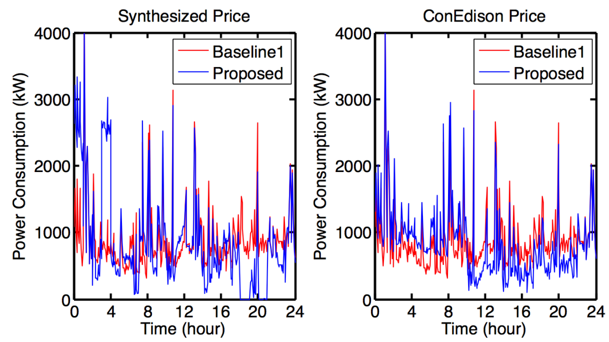

Fig. 7. The data center power drawn from the grid with and without the peak shaving capability of hierarchical ESDs.

synthesized dynamic pricing function, with different electric energy price in each hour. The second one is a real (time-of-use) pricing policy from conEdison [35]. More specifically, the second pricing policy comprises a high-peak time period with energy price \$0.3027/kWh and a low-peak time period with energy price \$0.0116/kWh.

We use Google cluster trace as realistic workloads to evaluate the effectiveness of the hierarchical ESD structure [28]. The Google cluster trace released in 2012 is measured on a 29-day period involving 672,075 jobs and more than 48 million tasks. In the trace, the (normalized values of) CPU, memory, and disk utilizations of the server cluster are measured and updated in every 5 minutes. We derive the power load $P_{i,j,t}^{load}$ of each server based on CPU and memory utilizations based on accurate server power modeling [34].

A. Control Optimization

In this section we demonstrate the effectiveness of the proposed hierarchical optimization method under given energy capacity of ESD in each level. We compare the proposed hierarchical optimization method with four baseline systems: Baseline 1 through 4. Baseline 1 uses the same hierarchical ESD structure but the ESD is not used for peak shaving and energy cost minimization under dynamic pricing (i.e., the ESD is only used for mitigating power outage.) Baseline 2 uses centralized ESD structure in [22] and the ESD/UPS is optimally controlled for energy cost minimization under dynamic pricing functions. 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 uses centralized ESD structure in [22] and the ESD/UPS is not utilized for peak shaving. Baseline 4 uses the same hierarchical ESD structure and uses ESDs for minimizing energy cost. However, the control algorithm of hierarchical ESDs does not account for rate capacity effect of battery ESDs.

We fix the energy capacity of server-level, rack-level, and data center-level ESDs and evaluate the performance of the proposed algorithm versus baselines over all 29 days of Google cluster trace. We use 0.24kWh as the energy capacity for each server-level ESD (corresponding to 3L in volume when assuming a energy density of 80kWh/m^3 for lead-acid batteries.) We use 32×0.24kWh as the energy capacity for each rack-level ESD where 32 is the number of servers in a rack, and use 8192×0.24kWh as the energy capacity of the data center-level ESD where 8192 is the total number of servers in the data center. The comparison results on overall energy cost achieved by the proposed hierarchical control optimization versus baselines are shown in Fig. 6, based on the synthesized dynamic pricing function (Fig. 6 left) and actual dynamic pricing function (Fig. 6 right). We can observe from Fig. 6 that the proposed control algorithm on hierarchical ESD structure consistently outperforms all baselines in terms of energy cost reduction, with a maximum reduction in energy cost of about 45%. Among all baselines, Baseline 3 performs the worst since it uses the least optimized structure (centralized ESD) and does not use ESDs for energy cost minimization. Baseline 2, which uses the least optimized centralized ESD structure but optimally controls the ESD/UPS for energy cost minimization, performs the best among all four baselines. Fig. 7 demonstrates the data center power drawn from the grid, i.e., P_t^{grid}, of the proposed algorithm on hierarchical ESD system versus Baseline 1 (i.e., hierarchical ESD but without peak shaving capability.) We can observe from Fig. 7 (when comparing with price functions in Fig. 5) the functionality of ESDs, i.e., storing excess power when the energy price is relatively low and providing power when the energy price is high.

B. Design Optimization

In this section, we optimize the energy capacity of ESD at each level subject to volume and annually amortized capital cost constraints. More specifically, we fix the overall annually amortized capital cost of the hierarchical ESD structure to be the same as the hierarchical structure (the basic setup) provided in Section VI.A, i.e., using 0.24kWh, 32×0.24kWh, and 8192×0.24kWh for each server-level, rack-level, and data center-level ESDs, respectively. We change the volume constraints on each server-level and rack-level ESDs in the experiments (and do not set a volume constraint on data center-level ESD since in reality data centers often use outside located UPS's.) In the design optimization evaluation, we compare the optimally designed hierarchical ESD system with two baselines: Baseline 1 uses the hierarchical structure (the basic setup) provided in Section VI.A and corresponding optimal control algorithm; Baseline 2 uses the centralized UPS structure in [22] (with the same total ESD capital cost) and corresponding optimal UPS control algorithm. Please note that Baseline 1 may violate the server-level and rack-level ESD volume constraints when such constraints are tight.

Fig. 9 demonstrates the comparison results on overall energy cost between the data center with optimally designed hierarchical ESD structure with two baselines on four sets of volume constraints. The server-level ESD volume constraints are given by 1L, 2L, 3L (corresponding to 0.24kWh), and 5L, respectively, which constitute the x-axis of Fig. 9. The rack-level ESD volume constraints are correspondingly given by 32×1L, 32×2L, 32×3L, 32×5L,

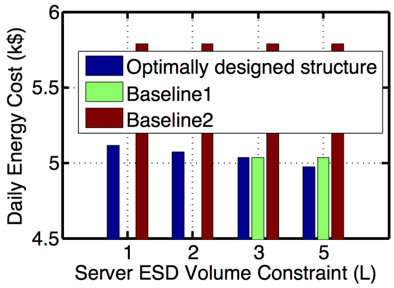

Fig. 8. Comparison results on overall data center energy cost between the optimally designed hierarchical ESD structure with two baselines.

respectively, where 32 is the number of servers in each rack. We can observe that (i) the overall (daily) energy cost of the data center with optimally designed hierarchical ESD structure reduces when the server-level and rack-level volume constraints become loose. This is because it is generally desirable to allocate larger amount of ESDs to the server-level or rack-level than centralized ESDs. (ii) Our basic setup provided in Section VI.A, i.e., using 0.24kWh, 32×0.24kWh, and 8192×0.24kWh for each server-level, rack-level, and data center-level ESDs, respectively, is the optimal design under the 3L server-level ESD volume constraint (and 32×3L rack-level ESD constraint.) (iii) The basic setup fails to meet the server-level and rack-level ESD volume constraints when such constraints are tight.

VII. Conclusion

This paper provides a realistic power delivery architecture along with an optimal design and control framework of the hierarchical ESD structure in order to minimize overall energy cost. First, a realistic and detailed power delivery architecture has been proposed, leveraging the advantages of centralized ESD structure from Intel and distributed single-level ESD structures from Google and Microsoft, and effectively avoiding AC-DC-AC double conversion to enhance efficiency. To take full advantage of the hierarchical ESD structure, we formulated a data center ESD design and control problem accounting for dynamic energy pricing, constraints on ESD volume (for each level) and the overall (annually amortized) capital cost, and power losses due to the rate capacity effect and conversion circuitry, in order to minimize the overall energy cost of the data center. We conduct experiments using Google cluster workload based on realistic data center specifications.

Acknowledgment

This work was supported by a grant from the Software and Hardware Foundations program of the National Science Foundation.

References

[1] D. F. Parkhill, *The Challenge of the Computer Utility*, Addison-Wesley Educational Pub. Co, 1966.

[2] 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.

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

[4] R. Buyya, C. S. Yeo, and S. Venugopal, "Market-oriented cloud computing: vision, hype, and reality for delivering IT services as computing utilities," *Proc. the 10th IEEE International Conference on High Performance Computing and Communications (HPCC)*, pp. 5-13, Dalian, China, Sep. 2008.

[5] X. Fan, W. D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," *Proc. the 34th Annual International Symposium on Computer Architecture (ISCA)*, pp. 13-23, San Diego, California, Jun. 2007.

[6] J. Hamilton, "Internet-scale service infrastructure efficiency," *Proc. the 36th Annual International Symposium on Computer Architecture (ISCA)*, pp. 232-232, Austin, Texas, Jun. 2009.

[7] D. Wang, C. Ren, S. G. A. Sivasubramaniam, et al., "ACE: abstracting, characterizing, and exploiting peaks and valleys in datacenter power consumption," *Proc. the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pp. 333-334, Pittsburgh, PA, Jun. 2013.

[8] A. Gandhi, M. Harchol-Balter, R. Das, and C. Lefurgy, "Optimal power allocation in server farms," in *Proc. the ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2009.

[9] A. Gandhi, V. Gupta, M. Harchol-Balter, and M. Kozuch, "Optimality analysis of energy-performance trade-off for server farm management," *Proc. the ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2010.

[10] Y. Wang, S. Chen, H. Goudarzi, and M. Pedram, "Resource allocation and consolidation in a multi-core server cluster using a Markov decision process model," *Proc. the 14th International Symposium on Quality Electronic Design (ISQED)*, pp. 635-642, Santa Clara, CA, Mar. 2013.

[11] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Greening geographical load balancing," in *Proc. the ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2011.

[12] 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.

[13] J. Li, Z. Li, K. Ren, X. Liu, and H. Su, "Towards optimal electric demand management for Internet data centers," *IEEE Trans. on Smart Grid*, 2011.

[14] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam, "Optimal power cost manageemnt using stored energy in data centers," *Proc. the ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2011.

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

[16] B. Aksanli, E. Pettis, and T. Rosing, "Architecting efficient peak power shaving using batteries in data centers," *Proc. the 21st International Symposium on Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS)*, pp. 242-253, San Francisco, CA, Aug. 2013.

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

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

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

[20] V. Kontorinis, L. E. Zhang, B. Aksanli, et al., "Managing distributed UPS energy for effective power capping in data centers," *Proc. the 39th Annual International Symposium on Computer Architecture (ISCA)*, Jun. 2012.

[21] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in data centers: what, where, and how much?" *Proc. the ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pp. 187-198, London, UK, Jun. 2012.

[22] 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

[23] Y. Wang, X. Lin, and M. Pedram, "Accurate component model based optimal control for energy storage systems in households with photovoltaic modules," *Proc. IEEE Green Technologies Conference (GTC)*, pp. 28-34, Denver, CO, Apr. 2013.

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

[25] 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

[26] M. F. Arlitt and C. L. Williamson, "Internet web servers: workload characterization and performance implications," *IEEE Trans. Networking*, vol. 5, no. 5, pp. 631-645, Oct. 1997.

[27] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: a framework for architectural-level power analysis and optimizations," *Proc. the 27th Annual International Symposium on Computer Architecture (ISCA)*, Jun. 2000.

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

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

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

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

[32] 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*, vol. 19, no. 3, pp. 291 - 312, 2009.

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

[34] 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.

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