This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.

Datacenter Peak Power Management with Energy Storage Devices

Baris Aksanli
Electrical and Computer Engineering Department
San Diego State University
baksanli@sdsu.edu

Abstract—**Datacenters are warehouse-scale computing buildings that require high amounts of power. This results in prohibitively high electricity costs, increasing operational costs. Recently, researchers proposed using energy storage devices (ESDs) in datacenters to reduce their maximum power demand. ESDs enable datacenters to set smaller power budgets because they can provide additional energy when the demand is expected to be higher than this budget. This paper leverages a detailed survey of previous studies and analyzes the economical feasibility of this methodology from three different perspectives. We first compare different ESDs based on their ability to manage datacenter peak power. Then, we demonstrate that peak shaving benefits might differ due to datacenter usage and ownership. Finally, we analyze the option where datacenters participate in electric utility programs with their ESDs to obtain additional savings.**

Keywords—datacenter, peak power, energy storage device, cost savings

I. Introduction

Datacenters are power-hungry, warehouse-scale buildings with thousands of servers to serve hundreds of thousands of clients simultaneously, which constantly require significantly high power. This need places heavy burdens on 1) datacenter owners, due to utility bills, 2) electric utilities, as they have to timely respond to the demands of these buildings, 3) environment, due to carbon emissions from energy consumed. Researchers proposed several methods to address each issue. These solutions include datacenter energy/power management mechanisms (i.e. virtual machine (VM) management, dynamic voltage and frequency scaling (DVFS)-based management, load balancing, etc.), datacenter participation in the utility ancillary service markets (i.e. regulation, demand-response, etc.), and using renewable energy (i.e. solar and wind) to account for datacenter power consumption, either fully of partially. We focus on another well-known method, peak power management, for datacenters. The idea is to reduce the maximum instantaneous power draw. This way, 1) datacenters pay less electricity bills due to smaller demand charges, 2) utilities will have lower total peak demand and thus activate fewer power plants, 3) with fewer power plants at a time, adverse environment effects are reduced.

Many mechanisms have been proposed for peak shaving, including DVFS [1], [2], VM-based management [3], online job migration [4], [5], [6], and batteries [7], [8], [9], [10] or other energy storage devices (ESD), including thermal ESDs [11], super-capacitors [12], flywheels [13], etc. Among these, ESDs are particularly useful, as they do not introduce the performance overhead associated with meeting the power budget.

Figure 1 illustrates two strategies for peak shaving with ESDs. The horizontal axes represent a 24-hour interval and the vertical axes show the power consumption. In Figure 1-a, the dotted horizontal line denotes the datacenter power budget. The lower curve indicates the power consumption of a nominal size datacenter without peak shaving. A significant amount of provisioned power is wasted during low activity periods, resulting in lower profit. The upper curve adds extra servers and handles oversubscribed power with peak shaving, so that the power utilization is higher. Peak shaving prevents the power consumption from exceeding the peak power budget by the shaded region. The dashed line illustrates how much power the datacenter would consume without peak shaving, which would then incur as additional demand charges. Peak shaving increases the revenues by adding more machines to service more users and prevents utility-facing power consumption from exceeding the provisioned power with no performance cost.

Figure 1-b uses peak shaving to decrease the peak power without increasing the number of servers. The upper horizontal line represents the original peak and the lower one shows the power cap. The difference between the original peak and the power cap corresponds to savings, as the datacenter would pay less demand charges. If the power demand is greater than the power cap, the batteries provide energy. During low demand, the batteries recharge in preparation for the next peak.

This article comprehensively and comparatively investigates the effectiveness of datacenter peak shaving using ESDs through existing literature. We first show that different ESDs can be used for peak shaving but their efficiencies and peak shaving performance (i.e. percentage reduction in observed peak), cost (i.e. cost of using each method and electricity bill) and performance overhead (i.e. workload response times) depend on datacenter properties. Then, we analyze datacenters that are owned and used differently, e.g. a single owner datacenter vs. a colocation datacenter with multiple tenants. Potential savings become significantly different for these setups. We also show that the datacenters can help both themselves with reduced electricity bill and electric utilities with the reduced peak demand when utilities balance electricity supply and demand.

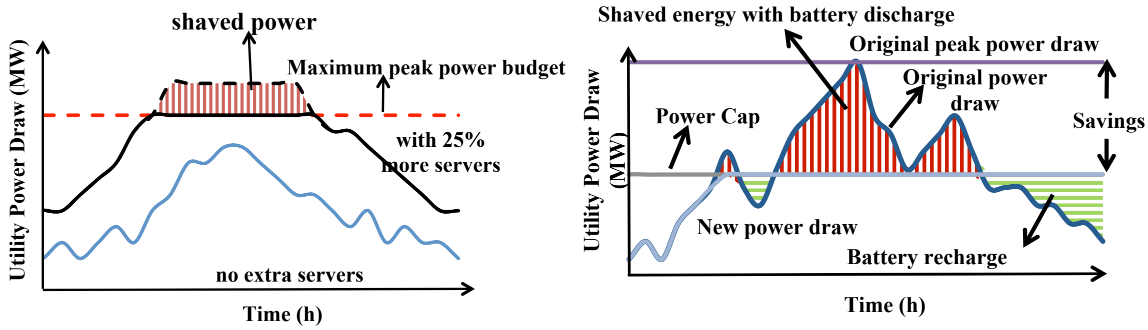

Figure 1. Sample peak shaving a) with and b) without extra servers [10]

II. Peak Power Shaving vs Energy Storage Technology

All ESDs can store and discharge energy whenever needed, with different characteristics. Some have high energy density (e.g. batteries, fuel cells), i.e. they can provide high energy over long time, whereas some have high power density (e.g. super-capacitors), i.e. they can provide high power over short time. Figure 2.i from US Energy Information Administration depicts energy vs. power characteristics of various ESDs. The x-axis shows the discharge time, i.e. how long a device can sustain its power, and the y-axis shows the capacity, i.e. the maximum power an ESD delivers. The graph's upper right portion shows the ESDs suitable for energy management, where providing reliable energy over long durations is important, e.g. compressed air (CAES) and pumped hydro systems. They require significant space and are not generally used for managing the peak power of individual systems. The graph's left portion has ESDs more suitable for power management, where a device needs to provide some amount of power (depending on the system size) over a relatively short period to respond to short-timescale (seconds, minutes, sometimes hours) power fluctuations. Examples include batteries, capacitors, and flywheels. These devices are used as backup power supplies (UPS) that can sustain power to the whole system for a while when needed (e.g. during a power outage).

Studies in datacenter peak power management leverage ESDs to sustain additional power during peak periods. This is because 1) ESDs can provide variable and high amounts of power in short time intervals, 2) they are space-efficient (as compared to pumped hydro) and 3) they already exist in most systems as UPS. Datacenters determine a power budget based on their financial and physical parameters, and whenever the power consumption is expected to be over the budget, ESDs provide the excess amount (see Figure 1). Figure 2.ii demonstrates sample peak demands over time. In each subfigure, the solid line is the demand and the dashed line is the power budget. Examples include a tall (in magnitude) but short (in time) spike (upper left), relatively tall and long spikes (lower left), multiple, short and frequent spikes (upper right) and tall with very long spikes (lower right).

Multiple studies, [13] [12] [8] [10] [7], analyzed the suitability of various ESDs for different peak demand shapes. These studies show that if the peak duration is short, super-capacitors are the best option because of their high power density, fast recharge-time and practically infinite charge/discharge cycles [13]. This applies to the upper graphs in Figure 2.ii. However, if the peak duration gets longer, super-capacitors become infeasible. We need ESDs that can sustain power longer. The first option is batteries, as they already exist in datacenters as UPS. Batteries are good for relatively shorter peak intervals (lower left in Figure 2.ii) as their cost becomes too high for larger capacities. Also, for longer peak intervals, the batteries need to discharge deeper, reducing their effective lifetime [10]. If the peak duration gets longer (lower right in Figure 2.ii), we can use devices with better energy density, such as flywheels and fuel-cells. The device choice here depends on the expected peak magnitude and duration. For smaller and longer peaks, we can use fuel-cells (as their energy density is good but power density is not big); and for taller and relatively long peaks, we can use flywheels (as they have better tradeoff between energy vs. power density). Some studies show that even CAES is a feasible option if the expected peak width is very long [13].

Datacenter power consumption, over time, can exhibit multiple characteristics due to changing nature of workload requests. In this case, we can use combinations of ESDs to satisfy different peak requirements. Combinations include super-capacitors and batteries [12]; thermal ESD and batteries [11]; super-capacitors and flywheels [13]. For example, in the super-capacitor/battery combination, super-capacitors can mostly be used before batteries to account for sudden power changes and limit battery usage (to increase battery lifetime) [12]. Another solution is to place ESDs at different locations in the datacenter power delivery hierarchy. Examples include distributed, per server batteries [8] [10] [14] [9] and batteries with grid-tie connections [10]. It is also important to select the most appropriate option for ESDs with multiple subtypes. Studies [8] [10] show that lithium-iron-phosphate (LFP) batteries are more profitable than lead-acid (LA) batteries, the most common type in datacenters, although the cost of LFP is considerably higher than LA [8]. This is because expected lifetime of LFP is longer than LA under frequent usage required for peak shaving. Renewable energy (solar, wind, etc.) is another option for peak shaving, supporting ESD devices or deployed individually [15].

Another important factor is the ESD capacity. Although previous studies demonstrated that datacenters could leverage overprovisioned ESDs [8] [10], their cost can get prohibitively significant. Kontorinis et al. [8] show that 20Ah LA or 40Ah LFP capacity per server (considering size and volume limitations) is optimal when using batteries for peak shaving. However, the battery only solutions need to cleverly manage the batteries to maximize their lifetime. Aksanli et al. [10] [16] show that the state-of-charge (SoC) and state-of-health (SoH) of each battery should be constantly monitored and batteries should be explicitly selected (limiting the SoC value and choosing the maximum SoH value) during a peak shaving event to maximize overall battery lifetime. Zheng et al. [11] add up to a 232k-gallon capacity, centralized thermal ESD to increase the savings. Liu et al. [12] examine different ratios between super-capacitor and battery capacities and conclude that a 6:4 ratio (super-capacitor:battery) is the best. Wang et al. [13] determine that the optimal ESD type/capacity combination depends on the workload type.

Traditional workload management mechanisms can be used as complementary to provide better peak shaving guarantees, including DVFS [17], online job migration [17], load balancing [9], etc. These methods can reduce the peak power level whenever the ESD capacity falls short. The server peak level can be reduced with a lower frequency setting (DVFS), or multiple jobs can be consolidated on a subset of all servers and the rest could be put in the sleep mode (job migration). However, these methods come with negative performance effects. Therefore, previous studies including these methods [17] [7] [9] set explicit quality of service constraints to limit these adverse effects, e.g. setting a tolerance factor in the 95th percentile job throughput or response time. In contrast, ESD-only solutions do not interfere with job scheduling decisions and thus, do not affect job performance.

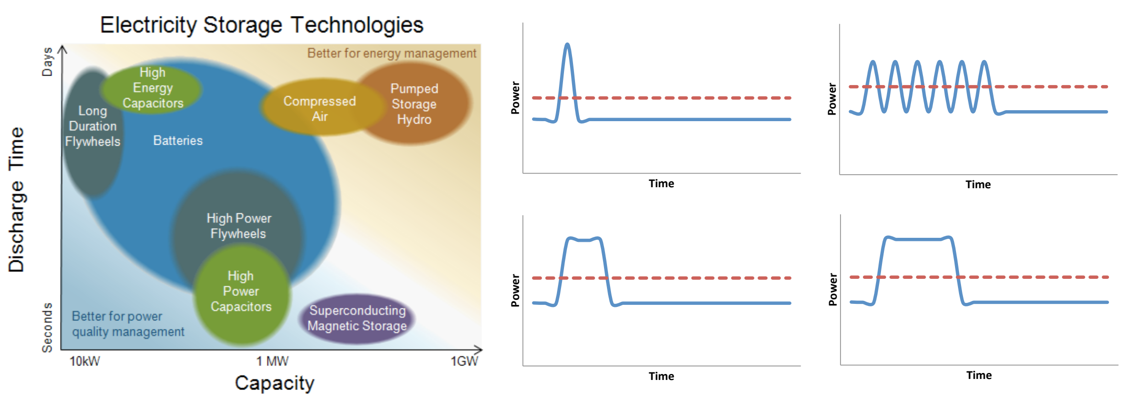

Figure 2. i) ESD characteristics: power vs. energy [18], ii) Various peak demands over time

III. Peak Power Shaving vs Datacenter Ownership

We study this part based on the peak shaving type a datacenter uses (with or without extra servers, Figure1). Co-location providers rent their equipment and space to retail customers. This applies to companies that require a datacenter-like system but do not want to build their own. A well-known example for a co-location renter is content delivery networks (CDNs) [9]. These renters make long-term power contracts with co-location providers and pay based on their provisioned power, instead of their actual consumed power. Thus, decreasing their peak consumption immediately translates to savings (Figure 1-b). The savings is proportional to the percentage peak reduced, however, the additional ESD cost should be subtracted to obtain net savings. The main purpose of peak shaving is to reduce the provisioned power level so that the co-location renters can contract for less power.

There are several companies that own their datacenters, where they still have peak budgets to control their electricity cost. However, they achieve this peak value rarely and underutilize the provisioned power. A solution is to add more servers, which improves the power utilization, with increased peak level. A peak shaving mechanism ensures that the provisioned power level is not violated with additional servers. The provisioned power level does not decrease but both the provisioned power and the datacenter equipment can be used to host more servers and thus total cost of ownership (TCO) per server reduces. Assuming that each server brings a constant amount of revenue, the total profit increases [8].

We leverage the TCO calculator in [8] for this analysis, which is based on APC's commercial calculator [19]. This model computes the TCO/server by itemizing it, calculating each part separately and analyzing how each part changes with more servers within the same power budget. These components include space/server/ESD depreciation, power/cooling infrastructure depreciation, opex costs, etc. When applying peak shaving with overprovisioned ESDs, the relevant TCO component is ESD depreciation (UPS depreciation in [8]), accounting for ESD costs. Some reasons for high ESD depreciation are short average lifetime (frequent replacements) or using an inappropriate ESD for peak shaving (low energy density, short service time, etc.). Further details of the TCO model are covered in detail in [8].

The savings both methods achieve depend on 1) the device lifetime, 2) the device unit price, 3) device usage. The last item is determined by how ESDs are placed in the datacenter (e.g. centralized vs. distributed) and the algorithm controlling their usage. In Table 1, we show a sample expected savings analysis for single vs. multi-user (colocation) datacenters with different ESD configurations. For clarity, we include only battery-based configurations (with more device types, layout information and different lifetime values, this table grows significantly). The colocation savings are calculated in terms of peak power budget reduction. The single-user cost savings is based on TCO/server, thus, even though its savings percentage is smaller, it translates into a larger value since it is multiplied for every server in the system, e.g. every 1% TCO/server reduction is equal to $27000/month savings for a 10MW datacenter, whereas 1% colocation cost savings for the same datacenter is around $3700/month [8].

ESD lifetime significantly affects savings because shorter lifetime means more frequent ESD replacement. This lifetime depends on how frequently and deeply the device is used every time. Similarly, device type and layout properties are important. Some configurations have better savings, but harder to manage, e.g. in a colocation datacenter, it is difficult have a centralized system that requires complete knowledge and control over individual ESDs. Similarly, for a very large single user datacenter, it is hard to manage distributed batteries because communication and network issues start becoming a serious bottleneck [16].

Another important issue is to manage the system downtime with additional ESDs because system downtime jeopardizes the profits and hence ultimately increases infrastructure amortization time. Studies quantify the system downtime to measure the effects of using ESDs for peak shaving [20] [12]. Govindan et al. [20] calculate the ESD cost and system availability relationship, and show that a hybrid (server + rack level) battery placement achieves the minimum cost with highest availability. Liu et al. show that using super-capacitors with batteries reduces the system downtime possibility by 42% [12]. Another strategy is to allocate some ESD capacity for only emergency handling and not use it for peak shaving [17].

ESD Type	**Additional Layout Information**	**Lifetime (years)**	**Colocation Cost Savings (Reducing peak power budget)**	**Single User Cost Savings (Adding more servers) TCO/server**
Battery – Lead acid	Distributed placement	1.5	2.7%	0.9%
		3	6.4%	2.65%
	Centralized with a grid tie	2.5	5.5%	1.87%
Battery – Lithium Iron Phosphate	Distributed placement	4	6.4%	1.86%
		10	15.5%	6.24%
	Centralized with a grid tie	6.5	11%	2.77%

Table 1. Peak shaving savings using different device/layout properties for single- and multi-user datacenters

IV. Peak Power Shaving and Utility Ancillary Services

The massive power demand of datacenters provides a different opportunity for them that are also beneficial to the electricity providers. Datacenters can participate in special programs, coordinated by utilities that help utilities maintain the electricity delivery system. These programs include demand response (DR), spinning and non-spinning reserves and regulation services. DR is a program where the participants get rewarded by reduced consumption. Spinning reserves are similar to DR, but require explicit contracts between participants and utilities. Regulation services help utilities balance the electricity supply and demand. For every program, the utility compensates the participant proportionately by the capacity dedicated to the program. It is evident that to have a significant enough effect on the electricity delivery system, participant's power demand should be sufficiently high. Datacenters, with as maximum demand as 100MW per site [21], can provide promising opportunities for ancillary services.

The key requirement when participating in an ancillary service is to have flexible demand and reduce the peak power whenever required. Although datacenters can employ several power management methods (DVFS, job migration, etc.), these methods may not provide significant flexibility in power demand and may result in severe performance hit to the workloads. In contrast, ESDs, especially when overprovisioned for peak shaving, can provide significant flexibility without negatively affecting the workloads.

Figure 3 illustrates an example, where a datacenter participates in regulation services using ESDs in two ways. In Figure 3.i, the datacenter does not change its original demand (dark blue curve). The original power budget is the light blue, horizontal curve. The red area around the nominal power is the regulation capacity promised to the utility. The depth of this area at any time represents the magnitude of the regulation capacity and determines the additional savings. Since the utility can ask the datacenter to adjust its power consumption to any point in the red area, the datacenter needs to charge (to have higher consumption than original) or discharge (to have lower consumption than original) available ESDs to adjust its consumption. There are some possible values that are greater than the power budget. Thus, the regulation capacity should be chosen carefully so that the savings are not neutralized by increased peak cost. The arrangement in Figure 3.ii eliminates this issue by making sure that the power budget is not violated. The regulation capacity (red area) is always within the power budget. However, it results in smaller regulation capacity, hence lower savings. Additionally, the battery discharge periods are mostly used to keep the power demand under the power budget, thus, we can only use shift the recharge periods to provide power flexibility. If we still want to use discharge intervals to provide additional flexibility, we can either 1) add more ESD capacity or 2) use the ESDs more aggressively (requiring more frequent replacements). Both options increase the capital costs. Although Figure 3 focuses on regulation services, the main idea is similar in other ancillary services.

A 21MW datacenter can save $40000/month from regulation services using its batteries [14]. Here, the batteries are overprovisioned for peak shaving and the savings are in addition to peak shaving savings. Another study estimates savings as much as 15% of the monthly electricity bill [22]. Similarly, a datacenter can enroll in DR programs using its power flexibility. Benefits can be as high as $2500/day assuming a 100 cluster datacenter with hundreds of servers per cluster [23]. The net savings depend on battery installation and replacement costs. Thus, battery usage should be optimized to limit frequent replacements.

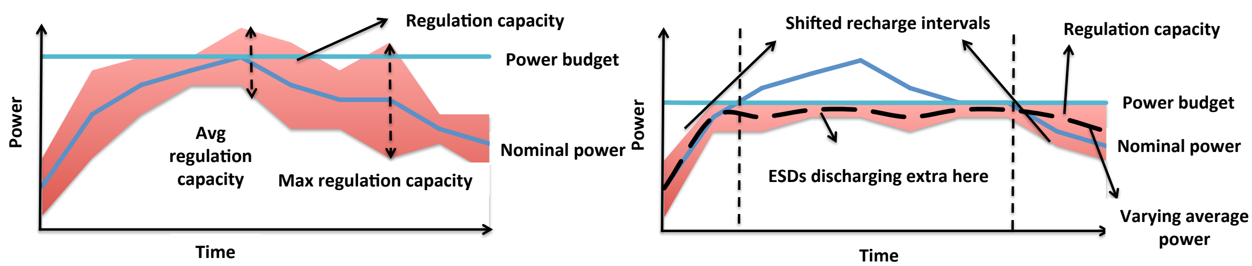

Figure 3. Regulation using ESDs, with i) flexible and ii) fixed power budget

V. DISCUSSION AND FUTURE DIRECTIONS

Previous studies prove that datacenters can leverage ESDs to reduce their peak level and electricity bills. This paper first surveys these studies, different ESD options for datacenters and their cost feasibility. In addition to the reduced peak level, these ESDs also provide significant power flexibility that can be exploited by electric utilities. As the power generation and distribution system is getting more distributed, individual buildings that act smartly with large power flexibilities become very crucial. In Section IV, we show that datacenters with ESDs can achieve this flexibility without harming their workload performance, becoming a very large-scale smart building with controllable peak power. With the addition of renewable energy sources, these large systems can act as a small microgrid, or become more active in the energy trade markets as they get more independent of traditional energy sources, e.g. recent zero-carbon cloud systems. However, datacenters still need the traditional power management mechanisms to provide system availability guarantees due to intermittent nature of renewable energy and finite ESD capacities. Thus, the research in the intersection of workload management and power system control is very crucial for self-sustaining datacenter environments. ESDs have multiple important roles in this scenario: reducing the maximum peak level, helping handling emergency cases, and providing bidirectional energy flexibility whenever needed.

VI. AUTHOR BIO

Baris Aksanli is an assistant professor at Electrical and Computer Engineering department of San Diego State University. He was a postdoctoral researcher in the Computer Science and Engineering Department at University of California, San Diego (UCSD). He received his PhD and MS degrees in Computer Science from UCSD, and two BS degrees in Computer Engineering and Mathematics from Bogazici University, Turkey. His research interests include energy efficient cyber physical systems, human behavior modeling for the Internet of Things, big data for energy efficient large-scale systems. He won the Internet2 IDEA Award with his work in Lawrence Berkeley National Laboratory and Spontaneous Recognition Award from Intel.

REFERENCES

[1] X. Fan, W. Weber, and L.A. Barosso, "Power provisioning for a warehouse-sized computer," I*SCA*, 2007.

[2] D. Meisner, C.M. Sadler, L.A. Barroso, W.D. Weber, and T.F. Wenisch, "Power management of online data-intensive services," *ISCA*, 2011.

[3] R. Nathuji and K. Schwan, "Vpm tokens: virtual machine-aware power budgeting in datacenters," HPDC, 2008.

[4] Niv Buchbinder, Navendu Jain, and Ishai Menache, "Online job-migration for reducing the electricity bill in the cloud," *NETWORKING*, 2011.

[5] Lei Rao, Xue Liu, Le Xie, and Wenyu Liu, "Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment," *INFOCOM*, 2010.

[6] Yanwei Zhang, Yefu Wang, and Xiaorui Wang, "Capping the electricity cost of cloud-scale data centers with impacts on power markets," *HPDC*, 2011.

[7] Sriram Govindan, Di Wang, Anand Sivasubramaniam, and Bhuvan Urgaonkar, "Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters," *ASPLOS*, 2012.

[8] Vasileios Kontorinis, Liuyi Zhang, Baris Aksanli, Jack Sampson, Homayoun Houman, Eddie Pettis, Dean Tullsen, and Tajana Rosing, "Managing Distributed UPS Energy for Effective Power Capping in Data Centers," *ISCA*, 2012.

[9] Darshan Palasamudram, Ramesh Sitaraman, Bhuvan Urgaonkar, and Rahul Urgaonkar, "Using Batteries to Reduce the Power Costs of Internet-scale Distributed Networks," *ACM SoCC*, 2012.

[10] Baris Aksanli, Eddie Pettis, and Tajana Rosing, "Architecting Efficient Peak Power Shaving Using Batteries in Data Centers," *MASCOTS*, 2013.

[11] Wenli Zheng, Kai Ma, and Xiaorui Wang, "Exploiting thermal energy storage to reduce data center capital and operating expenses," *HPC*, 2014.

[12] Longjun Liu, Chao Li, Hongbin Sun, Yang Hu, Juncheng Gu, Tao Li, Jingmin Xin, and Nanning Zheng, "HEB: deploying and managing hybrid energy buffers for improving datacenter efficiency and economy," *ACM SIGARCH Computer Architecture News*, 2015.

[13] Di Wang, Chuangang Ren, Anand Sivasubramaniam, Bhuvan Urgaonkar, and Hosam Fathy, "Energy Storage in Datacenters: What, Where, and How much?," *SIGMETRICS*, 2012.

[14] Baris Aksanli and Tajana Rosing, "Providing Regulation Services and Managing Data Center Peak Power Budgets," *DATE*, 2014.

[15] Daniel Gmach, Jerry Rolia, Cullen Bash, Yuan Chen, Tom Christian, Amip Shah, Ratnesh Sharma, and Zhikui Wang, "Capacity planning and power management to exploit sustainable energy," *CNSM*, 2010.

[16] Baris Aksanli, Tajana Rosing, and Eddie Pettis, "Distributed Battery Control for Peak Power Shaving in Data Centers," *IGCC*, 2013.

[17] Sriram Govindan, Anand Sivasubramaniam, and Bhuvan Urgaonkar, "Benefits and Limitations of Tapping into Stored Energy For Datacenters," *ISCA*, 2011.

[18] US Energy Information Administration. Electricity storage technologies can be used for energy management and power quality. http://www.eia.gov/todayinenergy/detail.php?id=4310

[19] APC. InfraStruxure Total Cost of Ownership, Infrastructure cost report. http://www.apc.com/tools/isx/tco/

[20] Sriram Govindan, Di Wang, Lydia Chen, Anand Sivasubramaniam, and Bhuvan Urgaonkar, "Towards Realizing a Low Cost and Highly Available Datacenter Power Infrastructure," *HotPower*, 2011.

[21] G. Cook and J. Horn, "How Dirty Is Your Data? A Look at the Energy Choices that Power Cloud Computing," 2011.

[22] Yuanyuan Shi, Bolun Xu, Baosen Zhang, and Di Wang, "Leveraging energy storage to optimize data center electricity cost in emerging power markets," *ACM e-Energy*, 2016.

[23] A. Mamun, I. Narayanan, D. Wang, A. Sivasubramaniam, and H.K. Fathy, "Multi-objective optimization of demand response in a datacenter with lithium-ion battery storage," *Journal of Energy Storage*, 2016.