

Managing Distributed UPS Energy for Effective Power Capping in Data Centers

Vasileios Kontorinis[†], Liuyi Eric Zhang[†], Baris Aksanli[†], Jack Sampson[†]
Houman Homayoun[†], Eddie Pettis[‡], Dean M. Tullsen[†], Tajana Simunic Rosing[†]

[†] Department of Computer Science and Engineering, UC San Diego

[‡] Google Inc.

Abstract

Power over-subscription can reduce costs for modern data centers. However, designing the power infrastructure for a lower operating power point than the aggregated peak power of all servers requires dynamic techniques to avoid high peak power costs and, even worse, tripping circuit breakers. This work presents an architecture for distributed per-server UPSs that stores energy during low activity periods and uses this energy during power spikes. This work leverages the distributed nature of the UPS batteries and develops policies that prolong the duration of their usage. The specific approach shaves 19.4% of the peak power for modern servers, at no cost in performance, allowing the installation of 24% more servers within the same power budget. More servers amortize infrastructure costs better and, hence, reduce total cost of ownership per server by 6.3%.

1 Introduction

The costs of building and running a data center, and the capacity to which we can populate it, are driven in large part by the peak power available to that data center. This work demonstrates techniques to significantly reduce the observed peak power demand for data centers with distributed UPS batteries, enabling significant increases in data center capacity and reductions in cost.

Modern data center investments consist of one-time infrastructure costs that are amortized over the lifetime of the data center (capital expenses, or capex) and monthly recurring operating expenses (opex) [22]. Capex costs are proportional to the provisioned IT power per facility, estimated at \$10-20 per Watt [9, 35, 43], as each Watt of computing power requires associated support equipment (cooling, backup, monitoring, etc.). Utilities typically charge a power premium that is tied to the peak power. This can become a significant portion of the monthly bill, up to 40% [18]. This paper examines the use of distributed batteries in the data center to reduce both capex and opex costs.

Power infrastructure is commonly over-provisioned in data centers to accommodate peaks and to allow for future expansion. However, to improve common case utilization, we can intentionally over-subscribe (under-provision) the

power infrastructure [13, 22, 24, 27, 33]. Over-subscribing provisions power infrastructure to support a lower demand than the largest potential peak and employs techniques to prevent power budget violations. In the worst case, such violations could trip circuit-breakers and disable whole sections of the data center, causing costly down time. To avoid this, data centers can employ power capping approaches such as CPU capping, virtual CPU management, and dynamic voltage and frequency scaling (DVFS) [27, 30, 36]. CPU capping limits the time an application is scheduled on the CPU. Virtual CPU management limits virtual machine power by changing the number of virtual CPUs. DVFS attacks the peak power problem by reducing chip voltage and frequency. However, all of these techniques result in performance degradation. This is a problem for any workload that has performance constraints or service-level agreements because power management policies apply these performance-reducing mechanisms at the exact time that performance is critical – at peak load.

Govindan, et al. [18] introduce a new approach that has no performance overhead in the common case. They leverage the energy stored in a centralized data center UPS to provide energy during peak demand, effectively hiding the extra power from the power grid. This technique is shown to work well with brief (1-2 hours), high-magnitude power spikes that can be completely "shaved" with the energy stored in batteries; however, it is less effective for long (8-10 hour) spikes. For longer spikes, they suggest a hybrid battery-DVFS approach.

However, many large data centers do not employ centralized batteries. Distributed, per-server batteries represent a more economical solution for battery backup. They scale naturally with the data center size and eliminate a potential single point of failure. Google employs this topology in their state-of-the-art data centers [16].

When leveraging a distributed UPS architecture to shave peak power, challenges arise due to the lack of heavy over-provisioning and the distributed nature of the batteries. The absence of over-provisioned UPSs means we need to justify the use of larger batteries based purely on cost savings from power capping. We need policies to determine how many batteries to enable, which batteries to enable, and when. However, there are also opportunities compared to

prior solutions. In a centralized UPS architecture, all power typically comes from either the battery or the utility. Thus, when batteries are enabled, they supply all datacenter power and drain quickly – if we only supply the over-threshold power, the batteries can sustain longer peaks. This is easily done in the distributed architecture by simply enabling enough batteries to hide the desired peak.

In this work, we discuss the applicability of battery-enabled power capping to distributed UPS topologies. We present details on the sizing and the technology alternatives of per-server batteries and consider several approaches that orchestrate battery charging and discharging while addressing reliability and availability concerns. This research goes beyond prior work by modeling realistic data center workload patterns over a multi-day period and by arguing that battery over-provisioning is financially beneficial. Enabling the placement of additional servers under a given power budget permits reductions of the data center total cost of ownership per server on the order of 6%. This is equivalent to more than \$15M for a datacenter with 28,000 servers.

This paper is organized as follows. Section 2 presents common UPS topologies and the associated trade-offs. Section 3 describes our total cost of ownership analysis. In Section 4 we contrast alternative battery technologies for frequent battery charge/discharge in the data center context and elaborate on their properties. In Section 5, we present our policies. In Section 6, we discuss our experimental methodology and give experimental results in Section 7. Section 8 reviews related work in power capping techniques, and Section 9 concludes.

2 Background

Primary power delivery in data centers is through a utility line. Data centers are also equipped with a diesel generator unit which acts as a secondary source of power during a utility failure. To facilitate switching power between the utility and the diesel generator, an automatic transfer switch (ATS) selects the source of power, which takes 10-20 seconds [18]. During this short and critical interval, the UPS units supply the power to the data center. In the centralized topology shown in figure 1(a), the power from a single UPS is fed to several Power Distribution Units (PDUs) to route the power to racks and servers. To eliminate the transfer time of the power line to the UPS, data centers commonly use double conversion UPSs. With double conversion UPSs, power is transformed from AC-to-DC to be stored in batteries and then from DC-to-AC to be used by the racks and servers. Although this organization has zero transfer time to the UPS (the UPS is always in the power path), the availability of the whole data center is dependent on the UPS. Additionally, double conversion introduces 4-10% power losses during normal operation [16].

The centralized UPS topology in figure 1(a) does not scale well for large data centers. This topology either requires double conversion, so that the power network dis-

tributes AC power, to be converted again to DC, or it distributes DC over the large network, resulting in higher cable losses. The inefficiency of AC-DC-AC conversions becomes more costly at scale. The UPS is also a single point of failure and must be overprovisioned.

Figure 1(b) shows the distributed design adopted by Facebook. A cabinet of batteries for every 6 racks, or a total of 180 servers, replaces the centralized UPS [12]. This design avoids double conversion by customizing the server power supply unit to support both AC power (from the grid) and DC power (from the battery cabinet). DC power is distributed from the UPS to the servers, but in this case that is a much shorter distance. Google goes even further, attaching a battery on every server after the Power Supply Unit (PSU) [16], as depicted in figure 1(c). This design also avoids the AC-DC-AC double conversion, saving energy under normal operation, and brings the AC distribution even closer to the IT load, before it is converted.

Availability in data centers is a function of how often failures happen, the size of the failure domain, and the recovery time after each failure. UPS placement topology impacts the availability of the data center, particularly the associated failure domain. The more distributed the UPS solution, the smaller the failure domain. Thus, the centralized design requires full redundancy, while the Google approach provides none (loss of a single node is insignificant), further reducing cost.

3 Total Cost of Ownership analysis

Modern data centers are typically power limited [43]. This means that the overall capacity (number of servers) is limited by the initial provisioning of the power supporting equipment, such as utility substations, diesel generators, PDUs, and cooling. If we reduce the peak computing power, we can add additional servers while remaining within the same power budget, effectively amortizing the initial investment costs over a larger number of servers. Moreover, extra work done per data center should result in fewer data centers, greatly reducing capex costs.

Distributed UPSs are currently designed to support the whole computing load long enough to ensure safe transition from the main grid to the diesel generator. This time window (less than one minute) translates to batteries with insufficient stored energy for meaningful peak power shaving. Therefore, to enable peak power capping using UPS stored energy in the distributed context, we need to over-provision per server battery capacity. This section discusses the TCO model we use to examine the battery over-provisioning that makes financial sense and maximizes total profits.

The profitability of an investment is defined as the generated revenue minus the associated total cost of ownership (TCO). The data center revenue equals the number of servers times the income per server. We assume constant income per server. Therefore, maximizing the profitability per server is equivalent to minimizing the TCO per server.

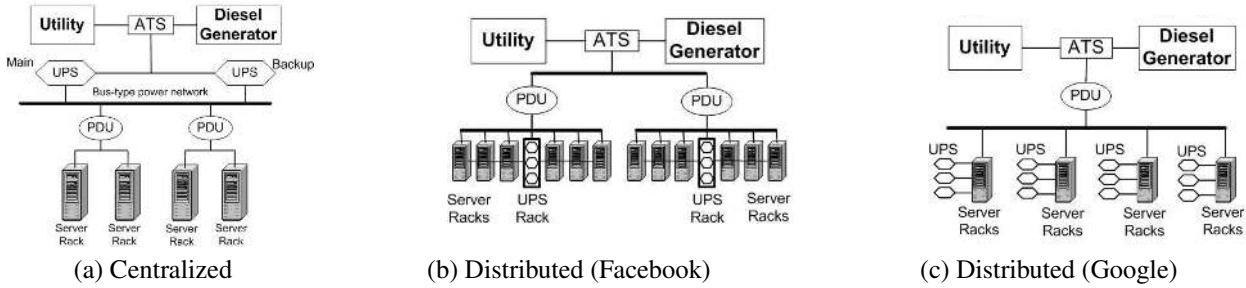


Figure 1: Power hierarchy topologies with (a) centralized UPS and (b,c) distributed UPS solutions.

$$\begin{aligned}
 TCO/server &= (dataCenterDepreciation + dataCenterOpex + serverDepreciation + serverOpex)/N_{servers} \\
 &= ((FacilitySpaceDepr + UPSDepr + PowerInfrastructureDepr + CoolingDepr + RestDepr) + dataCenterOpex \\
 &\quad + serverDepr + (ServerRepairOpex + (ServerEnergyOpex + ServerPowerOpex) * PUE))/N_{servers} \quad (1)
 \end{aligned}$$

We now explain how placing more servers within the same power budget reduces TCO per server. Our TCO analysis is inspired by the data center cost chapter in Barroso and Holze [22]. For simplicity, we assume sufficient initial capital, hence there are no monthly loan payments, and full capacity for the data center (limited by the provisioned power) from the first day. The TCO/server is given by equation 1.

In this equation, data center depreciation is the monthly depreciated cost of building a data center (we assume 10 year straight-line depreciation [22]) The assets required for a data center are land, UPS and power infrastructure (diesel generators, PDUs, back-room switchgear, electrical substation), cooling infrastructure (CRAC, economizers), as well as several other components such as engineering, installation labor, racks, and system monitors that we include in *RestDepreciation*. The data center opex is the monthly cost for running the data center (infrastructure service, lighting). We collect the depreciation and opex cost information for a data center with 10MW provisioned computing power (critical power) from APC’s commercial TCO calculator [5].

Servers typically have shorter lifetimes and are depreciated over 4 years. Server opex consists of server repairs and the electricity bill. Utility charges have a power component and an energy component. The power component is based on the peak sustained power for a 15 minute window over the period of a month [8] while the energy is based on the total data center energy used (different charging models provide similar results). To account for the electricity consumed by infrastructure, excluding servers, we scale the total server peak power and energy by the power usage effectiveness (PUE), assumed at 1.15 [16]. We assume a customized commodity server similar to Sun Fire X4270, with 8 cores (Intel Xeon 5570) at 2.40 GHz, 8 GB of memory, and costing \$1500. The inputs to our TCO model are summarized in table 1.

The table and the pie chart in Figure 2 show the breakdown of TCO/month/server. The major TCO component is server depreciation (40.6%). Infrastructure related compo-

Data center Critical Power	10 MW
Server	Idle Power: 175W, Peak Power: 350W (measured)
Number of servers	28000 (critical power / server peak)
Average Server Utilization	50% [22]
Utility Prices	Energy: 4.7 c/KWh, Power: 12 \$/KW [8, 18]
Server cost	\$1500
PUE	1.15 [16]
Amortization Time	Infrastructure: 10 years, Servers: 4 years [22]

Table 1: TCO model assumptions

nents (facility space, power, cooling, and data center opex) account for more than 35%. In the same table, we also present how the ratio of each TCO component per server changes when we are able to add additional servers within the same power budget. Server depreciation, server opex, and UPS TCO scale with the number of servers and are constant. The energy component of the utility bill also scales with the number of servers, but the power component stays the same and is amortized over more servers. Infrastructure costs are also amortized over a larger number of servers. The UPS cost (estimated as the total cost of the server-attached batteries) represents a very small portion of the TCO; it is marginally visible in the pie chart. Our proposal over-provisions batteries and increases the cost of the distributed UPS. In return, we amortize the cost of several large components over a larger set of servers. The full TCO model described here can be found in [25].

4 Characterizing distributed UPS batteries

Current UPS designs rely on lead-acid batteries because of their ability to provide large currents for high power applications at low cost. In this section, we discuss alternative battery technologies for distributed UPSs, model battery behavior when employed for peak power capping, and elaborate on the selection of parameters (capacity, cost, depth of discharge) to minimize TCO/server.

The spider graph in Figure 3 compares the major competing battery technologies for high power applications, typical for servers, at the range of 12V and 15-30A: lead-acid (LA), Lithium Cobalt Oxide (LCO), and Lithium Iron

TCO component	TCO/month (TCO/month/server)	TCO/server trend with extra servers
Facility Space depreciation	96,875\$ (3.46\$)	Decreasing
UPS depreciation	3,733\$ (0.13\$)	Constant
Power Infrastructure depreciation	169,250\$ (6.04\$)	Decreasing
Cooling infrastructure depreciation	70,000\$ (2.50\$)	Decreasing
Rest depreciation (racks, monitoring, engineering, installation)	255,594\$ (9.13\$)	Decreasing
Data center opex (maintenance, lighting)	213,514\$ (7.63\$)	Decreasing
Server depreciation	875,000\$ (31.25\$)	Constant
Server opex (Service/repairs)	43,750\$ (1.56\$)	Constant
PUE overhead	55,467\$ (1.98\$)	Constant
Utility monthly energy cost	252,179\$ (9.01\$)	Constant
Utility monthly power cost	117,600\$ (4.20\$)	Decreasing
Total	2,152,961\$ (76.89\$)	Decreasing




Figure 2: Total Cost of Ownership (TCO) [5]. TCO/server decreases as we increase servers under same power budget




Figure 3: Comparison of battery technologies [1, 44]. Note that we show the inverse of \$/Wh/cycle and recharge time so that higher is better.

Phosphate (LFP). Other technologies like NiCd, NiMH, or other lithium derivatives are excluded because they are dominated by one of the discussed technologies across all metrics. LA never performs best along any dimension except at low temperatures. While LA is cheapest per Wh, LFP offers an order of magnitude more recharge cycles, hence provides better \$/Wh/cycle than LA. LCO is the most expensive technology and provides comparable recharge cycles to LA. The advantage of LCO technology is its high volumetric density (Wh/l) and gravimetric density (Wh/Kg). Lithium batteries have longer service life than LA and also recharge faster. LFP has higher margins for over-charging and is safer than LA (may release toxic gases when over-charged) and LCO (may catch fire).

Properly selecting the technology and battery size depends on its use. UPS batteries in modern data centers are discharged only during a power outage. According to [34], the number of utility outages that affect data centers ranges from 1.5 to 4.4 per year. Therefore, cost, service life, and size are the most important parameters. The selection criteria become quite different when we re-purpose the batteries to be aggressively charged and discharged. Recharging cycles become crucial because continuous battery use may drastically shorten battery lifetime, resulting in frequent replacement costs that negatively affect TCO/server. Hence \$/Wh/cycle is a better metric than \$/Wh alone. Since LCO does poorly on both cost and cycles, it is not considered further.

We now focus on the per server distributed UPS design and explore the degree of overprovisioning that is most fi-

Input	Value		Reference
	LA	LFP	
Service time	4yrs	10yrs	[45, 11]
Battery Cost per Ah	2\$/Ah	5\$/Ah	[11, 1]
Depth of Discharge	40%	60%	Estimated (see figure 6)
Peukert's exponent	1.15	1.05	[21]
Existing Server Bat. Capacity	3.2Ah		[16]
Recharge Cycles	f(DoD) – Table 3		[45, 42]
Battery Voltage	12V		[16]
Max Bat. Discharge Current	23A		Estimated (ServerPeak * PSUEff / Voltage)
PSUEff	0.8		[7]
Discharges per day	1		Based on data from [15]
Battery losses	5%		[40, 46]

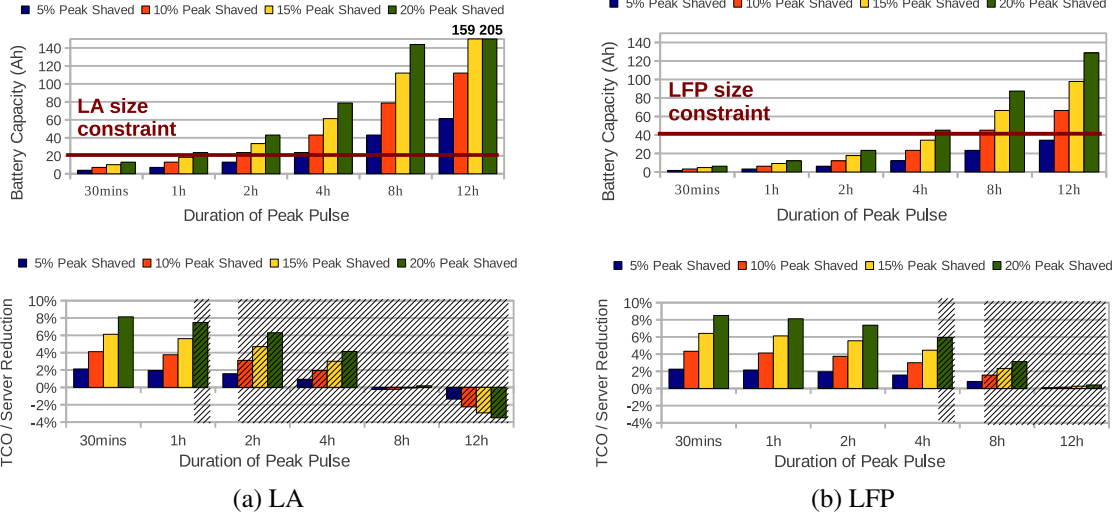
Table 2: Input values for battery cost estimation.

nancially beneficial. We derive comprehensive models for the effective battery capacity, battery discharge time, average battery lifetime, and resulting battery depreciation cost. Due to space constraints, we refer the reader to [26] for those equations.

Battery cost is estimated based on its 20h-rated capacity in Amp-hours (Ah) and the cost per Ah. We derive the required battery capacity based on the amount of power we want to shave and the corresponding energy stored in a battery for a given daily power profile. We derive the cost per Ah from [1, 11]. Tables 2 and 3 show all the inputs for the battery sizing estimation. Those inputs include Peukert's exponent [21], which is used to derive the battery discharge time as a function of the discharge current for a given battery capacity. Peukert's exponent is lower for LFP, meaning it retains more useful capacity at high discharge currents.

Depth of discharge is a critical input. High depth of discharge (e.g., allowing the battery to discharge 90% of its original capacity each time) erodes the lifetime of the battery and possibly makes the system vulnerable to power failures. Low depth of discharge means we need larger, more expensive batteries to shave the same peak power. See [26] for full details.

To derive the required battery capacity, we first set a peak power reduction goal and estimate the total energy that needs to be shaved at the data center level over the period of a day. We assume all batteries get charged and discharged once per day because, according to [15], all the traffic pro-



(a) LA (b) LFP
Figure 4: Battery capacities for different pulse widths and portion of peak power shaved. We also show the monthly TCO per server savings, assuming current battery costs, for the specified capacities of Lead-acid (LA) and Lithium Iron Phosphate (LFP) batteries. When the battery cannot fit within a 2U server, the associated savings are hatch shaded.

DoD (%)	10	20	30	40	50	60	70	80	90	100
Rcycles LA	5000	2800	1860	1300	1000	830	650	500	410	330
Rcycles LFP	100000	40000	10000	6000	4000	3000	2000	1700	1200	1000

Table 3: Recharge cycles as a function of depth of discharge (DoD). Deep battery discharge results in a fewer recharge cycles[42, 45].

files of large distributed applications demonstrate a single peak.

We use our battery models to contrast LA with LFP technologies as we vary the peak time in the power profile, study the effect of decreasing battery cost per Ah, and identify the depth of discharge that minimizes TCO/server. Figure 4 shows the provisioned battery capacity for a given peak power time and a targeted reduction in peak power as well as the respective TCO/server reduction. More energy needs to be stored in the battery to achieve the same reduction in peak power as the duration of peak power demand increases. Hence, the cost of the distributed UPS increases. In the LA case, over-provisioning is no longer helpful when the peak power lasts for 8 hours. This means that the additional distributed UPS cost is similar to the reduction of TCO/server due to amortization of the infrastructure costs on more servers. In contrast, LFP batteries remain beneficial for 8 hour peaks. Size constraints only allow shaving 5% of the 2-hour peak demand in the LA case, while we can shave 5% of an 8-hour pulse with LFP. In the TCO/server diagrams in figure 4, we denote the battery capacities that do not fit in a 2U server by hatch shading the respective columns. For the same spike duration, it always makes sense to shave more peak with a bigger battery, within size limitations. To further quantify these profits, we find using the analysis of section 3 that 6.3% monthly TCO/server re-

duction translates to \$4.8 per month per server, or more than \$15M over the 10-year lifetime of a data center with 28,000 servers.

Figure 5 presents the monthly TCO/server savings as the battery costs change. The projection for LA batteries is that costs do not change, while LFP prices are expected to be reduced due to the push for cheaper hybrid and electric cars [10]. For these graphs we assume that LFP cost reduces yearly at 8% [3]. At 4h peak per day, we achieve 5.5% TCO/server reduction for lead-acid, ignoring space considerations, while this value drops to 1% for a battery that fits within a 2U server design (2h peak). Using LFP batteries today we can achieve 7.25% TCO/server reduction and these savings will increase to 8.75% in the next 6 years.

Figure 6 shows the relation between depth of discharge and the TCO/server gains for both LA and LFP technology. There is a clear peak for the values 40% and 60% DoD, respectively. For low DoD values, the battery costs dominate the savings, because we need larger batteries to provide the same capping. For large DoD values, the lifetime of the battery decreases and more frequent replacements increase the UPS cost. The peak reduction of TCO/server occurs when the number of recharge cycles is equal to the battery service life in days. Note that due to the battery overprovisioning, less than 5% charge can sustain the server for 1 min and ensure data center continuity. Therefore, battery lifetime considerations affect TCO/server well before data center continuity becomes a concern.

To summarize our discussion on battery technologies and battery properties, we conclude: (1) Battery-based peak power shaving using existing batteries is only effective for brief spikes. To tolerate long spikes, larger batteries are necessary. However, the benefits from increased peak power




Figure 5: For the 4 hour pulse we show the projection of savings (ignoring space constraints) as the battery cost changes in the future [3].




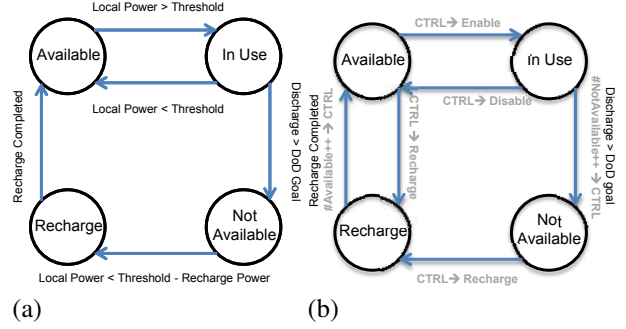
Figure 6: The relation between targeted depth of discharge and the reduction in TCO.

shaving outweigh the extra battery costs even when high demand lasts 8 hours. (2) LFP is a better, more profitable choice than LA for frequent discharge/recharge cycles on distributed UPS designs. This is due to the increased number of cycles and longer service lifetime, better discharge characteristics, higher energy density, and the reduction in battery costs expected in the near future. (3) It makes sense to increase the capacity of the battery to the extent that it fits under the space constraints. This translates to increased power reduction and more savings. (4) For each battery technology, there is a depth of discharge value that maximizes savings (40% for LA and 60% for LFP). This is the point where battery lifetime is no longer limited by the battery service time and needs to be replaced earlier due to frequent charging and discharging.

5 Policies

The analysis in the previous section assumes a simplified model of the power profile and perfect coverage of that peak by the batteries. As we move to a more complex model of real data center workloads and the associated power profiles, we investigate a number of policies for peak power shaving which react to the observed load on the data center power infrastructure.

We evaluate three policies for peak power shaving using distributed, per-server batteries. We examine policies that operate at different levels of the power hierarchy. The first policy budgets power at the server level. The second operates at the PDU level to coordinate the batteries of all the servers powered by the same PDU. Finally, power budgeting at the cluster level coordinates all the machines in a clus-



```

State=[NumAvail,NumInUse,NumNotAvailable,NumRecharging]
/* NumAvail: Batteries with charge currently idle (Available state)*/
/* NumInUse: Batteries with charge currently discharging (Inuse)*/
/* NumNotAvail: Batteries without sufficient charge (NotAvailable)*/
/* NumRecharging: Batteries currently recharging (Recharge)*/

```

```

1: /* Get difference between current and targeted power */
2: delta = load - threshold
3: /* Get difference in batteries */
4: ΔBats = abs(delta)/serverAveragePower
5: if (delta > 0) then
6:   /* Over peak goal */
7:   EnBats = min(NumAvail, ΔBats)
8:   NumInUse += EnBats
9:   NumAvail -= EnBats
10:  Enable EnBats batteries
11: end if
12: if (delta < 0) and (ΔBats > 25) then
13:   /* Under peak goal */
14:   DisBats = min(NumInUse, ΔBats)
15:   NumAvail += DisBats
16:   NumInUse -= DisBats
17:   Disable DisBats batteries
18:   RSlackBats = ΔBats-DisBats
19:   if RSlackBats > 0 then
20:     NumRecharging += RSlackBats
21:     Recharge RSlackBats batteries
22:   end if
23: end if

```

(c)

Figure 7: Per battery state machines for the local policy ServCtrl(a) and the coordinated policies PduCtrl and ClustCtrl (b). The controller logic for (b) is described in (c).

ter. The communication protocol to remotely enable/disable batteries or start recharge can be easily supported with existing network interfaces, such as SNMP or IPMI. The actual control algorithm can be implemented entirely in software. The policies manage battery discharge and also recharge. Recharging the batteries requires appreciable power and is thus best performed when the overall data center load is low. We consider the following policies:

1. Server with Local Controller (ServCtrl) When a server's power exceeds a preset threshold, this policy switches that server from grid power to battery power. The value of the power threshold defines how aggressively we cap power. When the overall power consumption is less

than the threshold and there is sufficient margin to recharge the battery without exceeding the budget, battery recharge is enabled. Each server has its own local LFP battery and a controller that periodically monitors server power. Figure 7(a) shows the state machine for this controller. If measured server power is higher than the local power threshold (peak power cap / number of servers), then the controller switches the server to battery power. Recharge activates when battery depth of discharge reaches the set goal (60% for LFP) and there is sufficient margin between the current server power and the target power cap to accommodate recharge power.

2. PDU with Centralized Controller (PduCtrl). This policy implements a controller per PDU. Each controller coordinates the operation of the batteries associated with the servers under a common PDU in order to match the energy to be shaved with the number of discharging batteries. It periodically estimates the power difference between current PDU power and the targeted PDU peak. As soon as this delta becomes positive, the controller estimates the approximate number of batteries that should start discharging ($abs(delta)/serverAveragePower$). Similarly, when the delta is negative and there are discharging batteries, the local controller will signal a number of batteries proportional to the magnitude of the estimated difference to stop discharging. We introduce an additional condition that the number of batteries we want to stop needs to be more than 25, which provides some hysteresis. The value 25 is a function of how fast the workload changes and how fast our controller responds (controller period is 3 mins).

Figure 7(b) and 7(c) show the state machine for the local battery controller and pseudo-code for the algorithm running on the PDU level controller. Arcs labeled in light color correspond to events sent to or from the centralized controller, whereas the other arcs, such as determining when the DoD goal has been met, remain local decisions. The controller attempts to distribute the enabling and disabling of batteries evenly by employing a static sequence that is interleaved across racks. When no batteries are currently enabled, the controller gradually signals discharged batteries to recharge. The controller also forces batteries that have not yet discharged to the DoD goal, but have not recently been recharged, to begin recharging in anticipation of the next day’s peak. Staggering recharge limits the possibility of power violations during low demand periods due to recharge power drawn from the utility.

3. Cluster with Centralized Controller (ClustCtrl). This policy applies the same logic as *PduCtrl*, but at the cluster level. Data center power delivery involves circuit breakers at several levels of the hierarchy. The previous policy, *PduCtrl*, maintains a power budget at the PDU level allowing additional servers at the PDU level. This policy targets a power budget at the cluster level, enabling over-subscription at the cluster level. We again employ a sequence to enable and disable batteries to evenly distribute

Workload	Service Time Mean	Interarrival Time Mean	Reference
Search	50ms	42ms	[29]
Social Networking	1sec	445ms	[4]
MapReduce	2 mins	3.3 mins	[6]

Table 4: Workloads

the power load across the levels of the power hierarchy.

6 Methodology

Section 3 derives upper bound power savings based on a simplified model of the workload and oracle knowledge of that workload. Here we present the tools used to model a variable, data-driven workload and realistic reactive capping policies that do not rely on oracle knowledge.

We developed a discrete-event simulator that captures the behavior of 1000 server nodes at the back-end of large distributed web applications. Each server is modeled as a queue with 8 consumers (cores) per server to service the incoming requests. Thus, we simulate a large network of M/M/8 queues. Our simulator monitors all levels of the data center power delivery hierarchy, namely the servers, racks, PDU, cluster, and data center. We measured the idle power of a Sun Fire X4270 server with a Hioki powermeter as 175W and the peak while fully utilized as 350W. We model server power as varying linearly between these two values, based on utilization.

For our results, we assume a distributed UPS with 12V LFP batteries attached to each server, provisioned at 40Ah, the maximum capacity that fits within the server size constraints. These batteries can sustain peak load for 92 minutes and take 2 hours to recharge once fully drained, measure 285.27 in³ or 16% of the 2U rack space and should fit in front of the server in the cold aisle. To properly capture Peukert’s effect during discharge, we recalculate remaining charge time every time the power draw on an individual server changes.

Table 4 presents the parameters of the workloads we use in our simulator. We assume a mix of web search, social networking, and MapReduce background computation. To capture the dynamic behavior of our workloads throughout the day, we use the Google Transparency Report [15] and scale interarrival time accordingly. We collect the traffic data for a whole year (10/1/2010-9/30/2011) for two google products in the United States. Google unencrypted search represents search traffic, and Orkut represents social networking traffic (similar to Facebook). MapReduce is a company internal product and, as such, does not appear in the Transparency report. Instead, we reproduce the weekly waveform that appears in figure 3 of [6] and repeat it over the period of a year.

We model a data center which serves all three types of workloads, with relative total demand placed on the servers in the ratios shown in Table 5. The relative loads of search vs Facebook/Orkut is chosen to match worldwide demand as reported by www.alexacom [2]. Note that we use Orkut




Figure 8: On the left we see the variation of data center energy throughout the year. During weekends and the summer traffic is lower. The average energy corresponds to 240.1W per server or utilization of 37.1%. On the right we zoom on the three days with highest energy requirements (11/17/2010-11/19/2010). We show average server power for each service and for MIX, the aggregate load. The average power for these days is 250.5W and the corresponding utilization 43%.

Workload	Relative Normalized Traffic
Search	29.2%
Social Networking	55.8%
MapReduce	15%

Table 5: Relative traffic numbers as obtained from [2]. MapReduce jobs are 15% of the load.

data to define the shape of the social networking demand curve, but use Facebook data to gauge the magnitude of the load. The maximum daily peak of the aggregate load is set to 80% of the data center peak computational capability. This number leaves sufficient computing margin to ensure stability of the system, and is consistent with published data center design goals, as shown in figure 1 in [29]. Note that because of this restriction the peak observed value of the average server power, 315W, is less than than the peak achievable power of 350W.

Figure 8(a) shows the day-to-day variation of the daily data center energy. The yearly daily average corresponds to 240.1W per server and varies moderately throughout the year. Weekends and summer months demonstrate lower traffic. We test our policies on the three consecutive days with the highest demand in energy. Graph 8(b) zooms in on these days (11/17/2010-11/19/2010) and presents the daily power profile for each workload separately, as well as the combined mix. The peaks of Search and Facebook are adjacent resulting in a waveform with broader peak. MapReduce traffic increases the variance of the graph.

We evaluate the workload mix in figure 8(b) under two different web service allocations: 1) restricting each service to its own dedicated set of servers (*split case*), 2) co-locating all web services, with highest priority for search, lower for social networking and lowest for MapReduce jobs (*mixed case*).

Additionally, we emulate the scheduling of jobs across

individual servers. Specifically, we consider a simple round-robin scheduling policy, similar to the operation of a naive web traffic load balancer, and a load-aware policy with knowledge of server CPU utilization. This scheduler is responsible for allocating the work among servers and is independent from the per-server scheduler that maps jobs to specific cores. In our simulated data center, the load-aware policy is extremely effective at distributing the load evenly, probably unrealistically so. Thus, the round-robin scheduler represents a more uneven distribution of work. A deployed load-aware scheduler probably falls somewhere between the two.

7 Results

The capacity of the battery, as well as the targeted DoD level, place an upper limit on the power capping that is possible for a given traffic pattern. In practice, though, the max achievable power capping also depends on the effectiveness of the policies that control the batteries. Setting the power capping threshold aggressively creates lower margins for wasted energy in our solution. There are two sources of battery energy waste: spatial and temporal. Spatial waste enables more batteries than necessary to shave a portion of overall power, while temporal waste enables batteries when capping is not required.

In this section, we gradually lower the peak power threshold until a policy begins to violate it. We show results for the lowest threshold (per server) that succeeds (horizontal line in figures 9, 10). Thus, we can compare policies based on that threshold. Some policies are not effective enough to cap power over a reasonable range. For those we give examples to illustrate why they fail. On an average day, it is to be expected that conservative estimates of peak power will result in a decent margin between battery capacity and the shaved peak load (some days the batteries may not be used at all). However, because we are modeling the worst days of the year, it is reasonable to expect that the available battery capacity is fully utilized. This methodology is reflective of what would happen on those days.

The *ServCtrl* policy (Figure 9(a)) assumes distributed, per-server batteries and does not require any centralized coordination. It relies completely on local server information. It is easy to implement, but due to the lack of coordination this scheme does not make efficient use of battery stored energy. Specifically, *ServCtrl* introduces temporal energy waste when transient effects create imbalances in the load across servers, resulting in battery discharge even if the total data center power does not exceed the threshold, leaving fewer batteries available to hide the real peak. We can even have batteries recharging during the peak. In the round-robin case, we cannot effectively shave peak for any meaningful power threshold,

When very effective load-balancing is in place, we see fewer instances of unnecessary discharge, but we observe a new problem. Once traffic increases to the degree that

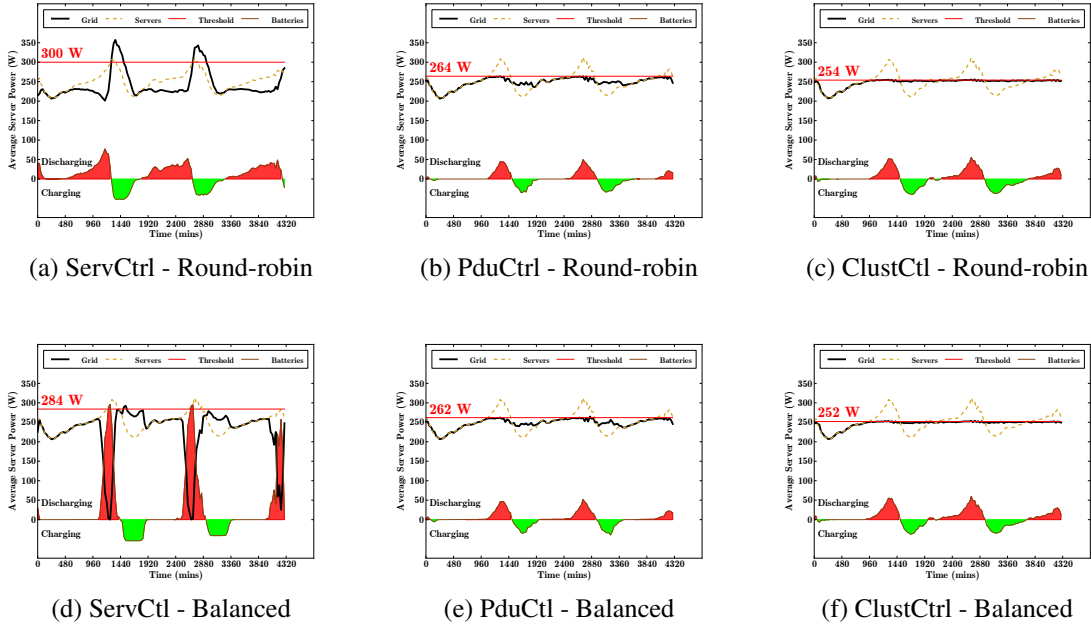


Figure 9: These plots show the average server, grid and battery power during battery discharge and charge. Grid power is equivalent to server minus battery power. Power capping at higher power hierarchy levels is more effective.

the power of each server crosses the threshold, all batteries begin discharging. As a result, a power dip follows. This effect is clearly visible in figure 9(d). Because the batteries reduce overall datacenter power well below the threshold, this overuses the total battery capacity. This is a similar effect experienced with power capping on a centralized UPS that can only produce power from the grid or the UPS, but not both. With this scheme, the batteries cannot sustain the peak, and grid power eventually exceeds the threshold.

There is a trade-off between recharge time and recharge current. Large values for recharge current (power) reduce recharge time but make it harder to find the necessary margin to initiate a recharge without violating the power budget. On the other hand, low recharge current provides ample margin for batteries to recharge, but risk having the battery still charging when the next peak arrives. For the *ServCtrl* policy we use a small recharge current of 3.7A (~0.1C) that corresponds to a charge time of 10 hours. In the coordinated policies, *PduCtrl* and *ClustCtrl*, the controller initiates the recharge of each server battery. It is much easier to find sufficient power slack to recharge a battery without violating the PDU or the cluster power budget respectively. For these policies, we use a high recharge current of 18.5A (~0.5C) that corresponds to a charge time of 2 hours.

Figure 9b shows that the *PduCtrl* policy performs much better than *ServCtrl*, maintaining a power threshold of 264W for round-robin and 262W for the load-aware scheduler. This is the result of coordination among batteries to achieve a local cap at the PDU level. Just enough batteries in each PDU region are activated to reduce power below the threshold, thus preserving other batteries to ride out the full

peak. Battery recharge is similarly coordinated so that no more than the available spare power is used for recharge. Global imbalances in the loads seen by each PDU result in slight noise in the total power; however, because each PDU is enforcing its threshold, that noise only results in grid power varying a little below the threshold.

That result holds when all three services run on all PDUs, because each PDU sees a similar power profile. For the *PduCtrl* we also study the scenario where each service is allowed to run on a subset of the PDUs. In this case, batteries are statically partitioned. As a result, search batteries are not available to help with the Facebook peak, and vice versa. Globally, we have batteries charging and discharging at the same time, which is clearly suboptimal. The lowest power budget that we can enforce in the split case is 291W (figure 10). This analysis motivates resource sharing among applications, despite the associated complexity for fairness and quality of service.

Figures 9c, 9f show the *ClustCtrl* policy applied on the mixed scenario for the round-robin balancing and the load-aware balancing. The lowest power cap for this policy is 254W and 252W for the two cases. Note that both of these results are very close to the ideal scenario which would reduce power to 250W (average power for the worst day). This increased efficiency is a direct result of being able to take a more global view of power. Imbalances between the PDUs no longer result in undershooting the power threshold, allowing us to preserve batteries that much longer.

There are many considerations that might determine the right level to apply our battery control policies. Our results show that the policy becomes most effective as we move up

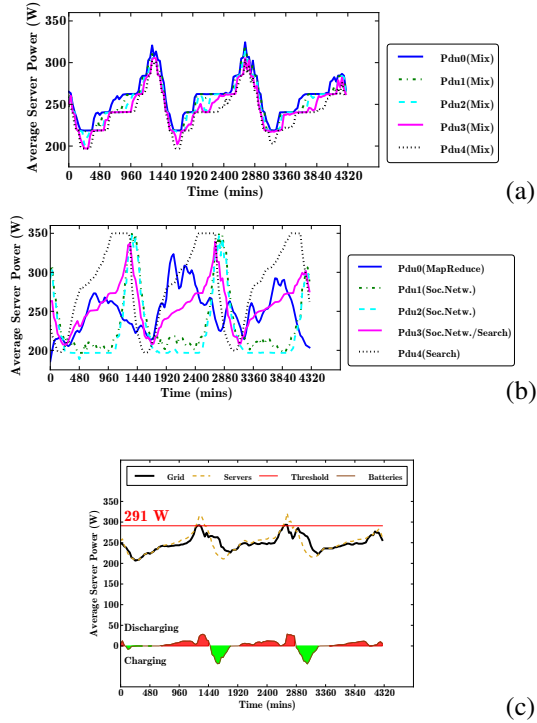


Figure 10: Here we quantify the effect of segmenting webservices into predefined PDUs. In (a) we show the server average power per PDU (without batteries) for the mixed case. In (b) we show the split case. When webservices run on split servers there are fewer available batteries to deal with a power peak. This is why in (c), when we use the batteries we can only guarantee peak power of 291W.

the power hierarchy. Most importantly, the policy should be applied no lower than the level at which the component workloads of the datacenter are mixed together. These results indicate that with properly sized batteries and an effective control policy, we can do much more than shave the extreme peaks of the load – in fact, we almost completely flatten the power profile very close to the average power. Capping peak power from 315W to 254W corresponds to a reduction of 19.4%. This reduction will allow 24% more servers within the same provisioned power and reduce TCO/server by 6.3% (see section 3), resulting in more than \$15M savings in costs for a data center with 28,000 servers over its lifetime.

Guard band and DVFS – when projections fail Prior work on power capping either applied performance degrading techniques, like DVFS, at peak load, or fall back to it as a failsafe when the batteries fail [18]. However, applying techniques such as DVFS at peak load is often an unacceptable option. Many datacenter applications track performance by watching the tail of the performance distribution – e.g., 95th percentile response time. Applying DVFS at peak load, even for a short time, can have a disastrous ef-

fect on those metrics. DVFS not only extends latencies, but also reduces throughput and induces higher queuing delays. Reducing performance at peak load increases the response time of those jobs already at the tail of the distribution.

Our technique does not apply DVFS, even when the peak power exceeds our conservative estimates, nor do we give up and allow the grid power to increase. In all the previous algorithms we disable the batteries once we hit the DoD goal (preserving battery lifetime – see Section 4). However, another benefit of the high DoD limit is additional stored energy in our batteries that can be used in case of emergency. With LFP per-server batteries there is approximately 35% guard band before we are in danger of not having sufficient reserves to survive a grid failure. This guard band can be used on days where the power profile exceeds worst-case peak power estimates. Our projections for the optimal DoD level were based on daily discharge; however, going below 40% to say, 35% or 30%, a couple times a year, or even once a month, will have no significant impact on the lifetime of the battery. Thus, we never need to apply performance-reducing mechanisms at peak load unless our estimated peak power is off by enormous margins. That does not mean that DVFS cannot still be an effective power management technique. But in our system, we would apply it during *low* utilization, where there is slack in the response times. By reducing power at low demand, we create more margin for recharging and accelerating the recharge cycle. This technique is not employed in the results shown in this paper, but would allow us to shave even more power with minimal performance impact.

Failure analysis In large data centers it is common to cluster maintenance operations to reduce costs. This means that a non-negligible portion of batteries may be unusable for peak power shaving purposes before these batteries get replaced. Figure 11 shows how the lowest achievable peak changes when we assume that a portion of batteries has failed. We compare the best policy *ClustCtrl* with and without the use of the additional energy provided by discharging our batteries beyond the DoD goal. The peak threshold gradually increases with a larger portion of dead batteries. However, the increase is relatively small. Even when half of the batteries are dead we can still shave 16% of peak power. For these experiments, we find that we do not need to modify the algorithm of the controller to handle the unusable components. The controller signals a faulty component to start discharging, but no decrease in power takes place. As a result, the controller signals additional batteries in the next round and eventually corrects for the failure without any direct feedback. We also observe that the additional energy from (rare) deeper discharge of the batteries allows us to shave more power with fewer batteries. However, if the datacenter is allowed to enter deeper discharge frequently while dead batteries stack up for an extended period, then it can have an impact on the battery lifetime.

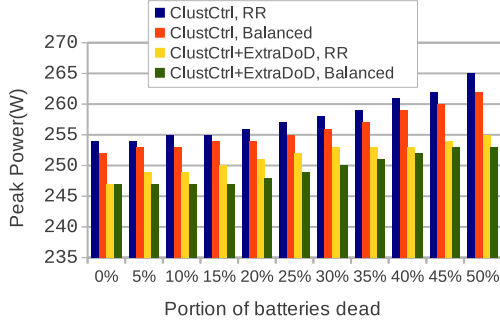


Figure 11: As the number of unusual batteries increase, the lowest possible peak power increases. Allowing to exceed our DoD goal occasionally, permits even higher peak power reduction. Load imbalances discharge batteries at different rates and make power capping harder.

Energy proportionality The server used for this study is a representative modern platform, with idle power close to 50% of peak, based on our measurements. In the future, servers are expected to become increasingly energy proportional. We model the impact of a server that is completely energy proportional.

Energy proportional servers essentially increase the height of the peak, relative to the average power, since power is significantly lower during off-peak periods. Consequently, we can further reduce the power threshold. Our simulations indicate the ability to reduce the peak observed power from 280W to 175W, a reduction of 37.5%. That results in an increase in server capacity of 60%.

8 Related Work

Peak Power Provisioning and Capping: Reducing power consumption in server clusters is a well-studied problem in the literature [39, 31, 28, 17]. The overall idea is to combine CPU throttling, dynamic voltage/frequency scaling (DVFS), and switching entire servers on/off depending on the workload. Raghavendra, et al. [39] note that more efficient power management solutions are possible by managing power at the rack level than at individual blades. They devise proactive and reactive policies based on DVFS to cap power budgets at the rack level. Nathuji and Schwan [31] introduce the notion of power tokens to deal with heterogeneity across hardware platforms. Govindan, et al. [17] combine applications with heterogeneous characteristics in terms of burstiness. As a result, the power budget is exceeded statistically infrequently. DVFS is used as a failsafe mechanism to prevent against lasting peak power violations.

Femal et al. [14] were among the first to use formal control theory to maximize throughput while capping power. Raghavendra, et al. [36] extend the control theory idea to present a hierarchy of coordinated controllers that cap power across different levels of the power hierarchy and minimize performance impact. They argue for nesting controllers that operate at different time granularities to ensure

stability and emphasize the information flow between the controllers.

Using batteries in data centers: Battery power management has been studied in the embedded/mobile system domain with various works proposing techniques to adjust the drain rate of batteries in order to elongate the system operation time [32, 38, 37, 41]. Prior research has also investigated analytical models for battery capacity and voltage in portable devices [32, 41, 23]. Govindan, et al [18] introduce the idea of reducing data center peak power by leveraging the stored energy in a centralized UPS. During peak load, power from the UPS batteries augments the main grid, effectively hiding the peak from the utility service. During low load, the batteries recharge, consuming additional power.

In a follow-up work [20], they extend their prior work to also use distributed UPSs for peak power capping. That work focuses on power capping at the rack, using small lead-acid batteries to shave peak power. This approach allows them to prevent rare, brief power emergencies without performance degradation and relies on DVFS and load migration for peaks longer than several minutes. In our work, we examine solutions at multiple levels of the power hierarchy, show the financial advantages of more aggressive batteries with a more detailed model that incorporates battery lifetime considerations, and employ solutions that sacrifice no performance – the desired solution in a performance-sensitive data center under peak load.

In a separate work [19], the same authors also argue for a distributed UPS solution from a cost and reliability perspective. They find that a hybrid distributed UPS placement, at PDU and server level, yields the most promising topology. They do not consider battery energy for peak power capping in that work, but this finding provides additional motivation for our work on the use of distributed batteries for power capping.

9 Conclusions

State-of-the-art data centers such as Google’s and Facebook’s have adopted a distributed UPS topology in response to the high cost associated with a centralized UPS design. In this work we explore the potential of using battery-stored energy in a distributed UPS topology to shave peak power. We describe how to provision the capacity of the battery and elaborate on how recharge cycles, the depth of discharge, and the workload power profile affect the potential for peak power shaving. We leverage the distributed nature of the batteries and design a controller to use them only when needed and thus prolong the duration of their usage, without violating the targeted power budget. Significant peak power reductions of up to 19.4%, are possible with our technique. These reductions allow us to provision more servers under the same power budget and reduce the TCO per server by 6.3%, significantly increasing the computation that can be done per facility and saving millions of dollars per datacenter.

Acknowledgments

We would like to thank the anonymous reviewers for providing useful feedback. This work was funded in part by Google, MuSys, NSF grants EEC-0812072, OCI-0962997, CNS-0821155, CCF-1018356, SRC Grant 2086.001, Microsoft, Ericsson, and Cisco.

References

- [1] AA Portable Power Corp. Portable Power Product design, assemble and quality control. <http://www.batteryspace.com/lifepo4cellspacks.aspx>.
- [2] Alexa. Web information, traffic metrics, search analytics, demographics for websites. <http://www.alexa.com>.
- [3] D. Anderson. An evaluation of current and future costs for lithium-ion batteries for use in electrified vehicle powertrains. Master's thesis, Duke University, 2009.
- [4] Apache. <http://incubator.apache.org/olilo/>.
- [5] APC. InfraStruxure Total Cost of Ownership. <http://www.apc.com/tools/isx/tco/>, 2008.
- [6] Y. Chen, A. Ganapathi, R. Griffith, and R. Katz. The case for evaluating MapReduce performance using workload suites. In *Technical Report No. UCB/EECS-2011-21*, 2011.
- [7] Climate Savers Computing. Power supply efficiency specifications. <http://www.climatesaverscomputing.org/resources/certification>, 2011.
- [8] Duke Energy. Utility bill. <http://www.duke-energy.com/pdfs/scschedulesopt.pdf>, 2009.
- [9] Dupont Fabros Technology, Inc. Sec filing (s-11) 333-145294, August 9, 2007.
- [10] Economist. In search of the perfect battery. http://www.mitenergyclub.org/assets/2009/9/25/Economist_Batteries2008.pdf, Mar. 2008.
- [11] Electric motor sport. EV construction, thundersky batteries. http://www.electricmotorsport.com/store/ems_ev_parts_batteries.php, 2001.
- [12] Facebook. Hacking conventional computing infrastructure. <http://opencompute.org/>, 2011.
- [13] X. Fan, W.-D. Weber, and L. A. Barroso. Power provisioning for a warehouse-sized computer. In *ISCA*, 2007.
- [14] M. E. Femal and V. W. Freeh. Boosting data center performance through non-uniform power allocation. In *ICAC*, 2005.
- [15] Google. <http://www.google.com/transparencyreport/traffic/>.
- [16] Google Summit. <http://www.google.com/corporate/datacenter/events/dc-summit-2009.html>, 2009.
- [17] S. Govindan, J. Choi, B. Urgaonkar, A. Sivasubramaniam, and A. Baldini. Statistical profiling-based techniques for effective power provisioning in data centers. In *EuroSys*, 2009.
- [18] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. Benefits and limitations of tapping into stored energy for datacenters. In *ISCA*, 2011.
- [19] S. Govindan, D. Wang, L. Chen, A. Sivasubramaniam, and A. Sivasubramaniam. Towards realizing a low cost and highly available datacenter power infrastructure. In *HotPower*, 2011.
- [20] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar. Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters. In *ASPLOS*, 2012.
- [21] F. Harvey. Listing of ev batteries. <http://www.39pw.us/car/batteryTable.html>, 2009.
- [22] U. Hoelzle and L. A. Barroso. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan and Claypool Publishers, 2009.
- [23] M. Jongerden and B. Haverkort. Battery modeling. Technical report, TR-CTIT-08-01, CTIT, 2008.
- [24] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya. Virtual machine power metering and provisioning. In *SOCC*, 2010.
- [25] V. Kontorinis. Battery-aware Data Center TCO models. <http://cseweb.ucsd.edu/~tullsen/DCmodeling.html>, 2012.
- [26] V. Kontorinis, B. Aksanli, T. Rosing, and D. Tullsen. Battery Provisioning and Associated Costs for Data Center Power Capping. Technical report, UCSD, 2012.
- [27] H. Lim, A. Kansal, and J. Liu. Power budgeting for virtualized data centers. In *USENIX*, 2011.
- [28] D. Meisner, B. Gold, and W. Thomas. Powernap: Eliminating server idle power. In *ASPLOS*, 2009.
- [29] D. Meisner, C. M. Sadler, L. A. Barroso, W.-D. Weber, and T. F. Wenisch. Power management of online data-intensive services. In *ISCA*, 2011.
- [30] R. Nathuji and K. Schwan. Virtualpower: coordinated power management in virtualized enterprise systems. In *SOSP*, 2007.
- [31] R. Nathuji and K. Schwan. VPM tokens: virtual machine-aware power budgeting in datacenters. In *HPDC*, 2008.
- [32] D. Panigrahi, S. Dey, R. R. Rao, K. Lahiri, C.-F. Chiasserini, and A. Raghunathan. Battery life estimation of mobile embedded systems. In *VLSI Design*, 2001.
- [33] S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch, and J. Underwood. Power routing: dynamic power provisioning in the data center. In *ASPLOS*, 2010.
- [34] Ponemon Inst. *National Survey on Data Center Outages*, 2010.
- [35] S. press release. Savvis sells asserts related to two datacenters for \$200 million. <http://www.savvis.com/en-US/Company/News/Press/Pages/SAVVIS+Sells+Assets+Related+to+Two+Data+Centers+for+200+Million.aspx>, June 29 2007.
- [36] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu. No "power" struggles: coordinated multi-level power management for the data center. In *ASPLOS*, 2008.
- [37] D. N. Rakhmatov. Battery voltage prediction for portable systems. In *ISCAS*, 2005.
- [38] D. N. Rakhmatov and S. B. K. Vrudhula. Energy management for battery-powered embedded systems. *TECS*, 2003.
- [39] P. Ranganathan, P. Leech, D. Irwin, and J. Chase. Ensemble-level power management for dense blade servers. In *ISCA*, June 2006.
- [40] T. B. Reddy and D. Linden. *Linden's Handbook of Batteries (4th edition)*. McGraw-Hill, 2011.
- [41] P. Rong and M. Pedram. An analytical model for predicting the remaining battery capacity of lithium-ion batteries. In *DATE*, 2003.
- [42] M. Swierczynski, R. Teodorescu, and P. Rodriguez. Lifetime investigations of a lithium iron phosphate (LFP) battery system connected to a wind turbine for forecast improvement and output power gradient reduction. In *BatCon*, 2008.
- [43] W. P. Turner and K. G. Brill. Cost Model: Dollars per kW plus Dollars per Square Floor of Computer Floor, Uptime Inst. White paper, 2009.
- [44] B. University. Online university education about batteries. <http://batteryuniversity.com/>, 2003.
- [45] Windsun. Lead-acid batteries: Lifetime vs Depth of discharge. http://www.windsun.com/Batteries/Battery_FAQ.htm, 2009.
- [46] L.-L. Zhang, G. Liang, A. Ignatov, M. C. Croft, X.-Q. Xiong, I.-M. Hung, Y.-H. Huang, X.-L. Hu, W.-X. Zhang, and Y.-L. Peng. Effect of Vanadium Incorporation on Electrochemical Performance of LiFePO₄ for Lithium-Ion Batteries. In *Journal of Physical Chemistry*, June 2011.