Optimizing the Use of Request Distribution and Stored Energy for Cost Reduction in Multi-Site Internet Services

Maja Etinski
Barcelona Supercomputing Center
maja.etinski@gmail.com

Margaret Martonosi
Princeton University
mrm@princeton.edu

Kien Le, Ricardo Bianchini, Thu D. Nguyen
Rutgers University
{lekien,ricardob,tdnguyen}@cs.rutgers.edu

***Abstract*—Electricity bills of large-scale Internet services are a major component in overall operating costs [3]. As a result, it is increasingly important to identify cost-reduction approaches, either through reducing energy consumption or by strategically avoiding periods or data centers with the highest electricity costs. The first group of approaches based on reduction of energy consumption has been explored for many years. However, approaches that exploit various types of electricity price variability have been recently proposed. There is a need for a holistic optimization framework to explore various strategies to mitigate exposure to price variation. While prior work has considered independently aspects of spatial price variability (through an intelligent request distribution for multi-site Internet services) and temporal price variability (via energy storage), ours is the first to do both. In this paper, we propose battery-oriented techniques in combination with geographically-distributed request routing, in order to fully optimize the electricity costs of multi-site Internet services exposed to both types of electricity price variations. Furthermore, we present an optimization based framework designed to analyze the cost-saving leverage afforded by the combination of these techniques. By exploring a range of design decisions, such as battery capacity and service level agreement (SLA) requirements, and operating conditions such as common power pricing policies and price variability, we identify which strategies are most effective and under what conditions. For example, across a range of SLAs, our results show that batteries offer roughly 7% additional savings in multi-site environments already implementing intelligent request distribution.**

I. Introduction

Electricity bills of large-scale Internet services are a major component in overall operating costs [3]. For large data centers, monthly electricity costs can reach as high as 1M USD [3]. Accordingly, there has been considerable research addressing data center electricity usage as well as resulting expenses. Early work in this field aimed to decrease overall energy consumption in order to reduce the electricity bills [5], [7], [8], [14], [15]. Recently, new proposals have appeared that reduce the electricity costs without decreasing energy consumption. These approaches exploit *spatial* and *temporal* electricity price variability.

Spatial price variability is simply the fact that electricity prices vary with geography. The variance can occur due to differences in ease of generation and delivery, power demand, government regulation, or other factors. For an Internet service with several geographically-distributed data centers, each may have quite different electricity prices at any moment. Such large-scale services are supported by multiple data centers to ensure high availability and low response time. Their front-end devices assign user requests to geographically-distributed data centers. Typically, these services route requests to the geographically-nearest available data center, but other options are available. Request distribution policies that exploit difference in price, energy type, or other characteristics among geographically distributed systems have already been proposed [13], [16]. They dynamically move demand (user requests) to data centers with lower price to reduce energy costs.

Temporal price variability occurs at a single site when electricity providers charge electricity rates that vary with time of day or other patterns. While some very large ISPs are able to lock in entirely-fixed power rates, the vast majority of data centers must operate with pricing schemes that have some variability either with time-of-day (peak vs. non-peak), with power demand, or other factors. Furthermore, with different time zones causing peak electricity demand to be staggered at different locations, ISPs with data centers around the world experience richly varied spatial and temporal trends in electricity pricing. Techniques have also been proposed to exploit temporal price variability within a single data center by leveraging energy storage. Traditionally, Uninterruptible Power Supply (UPS) units have been used in data centers to store energy needed in the case of a power outage. Nevertheless, these UPS batteries can be used to store cheaper energy or reduce peak power draw, rather than simply at the onset of a power outage. Different battery recharging/discharging strategies have been proposed depending on the electricity billing model [10], [18].

These two types of electricity price variability have been investigated separately. We argue that there is a need for a holistic approach that combines all available leverages. Accordingly, we propose a new request distribution and power management policy for large-scale Internet services. Furthermore, our simulation framework is used to evaluate the potential of this holistic approach across all electricity billing schemes commonly found in practice.

In this paper, we pose these questions: In addition to

978-3-901882-46-3

the previously-evaluated opportunities afforded individually by cross-data-center request distribution or by within-data-center energy storage, what cost savings might ISPs achieve by composing these techniques together? We investigate whether and how much UPS batteries can help in electricity bill reduction in a geographically-distributed environment where cost-aware request distribution already improves the service operational costs. The use of request off-loading driven by lower electricity price is constrained by the agreed service level agreement (SLA) due to network latency. However, it is not clear how SLA affects the use of stored energy. Can infrastructures deploying energy storage achieve better SLAs under the same cost as when only request off-loading to centers with currently lower prices is used? We also explore how the saving opportunities change under different electricity pricing models.

Our work makes the following contributions. First, we propose a new policy that minimizes the multi-site Internet service electricity costs by implementing both request routing and energy storage in a manner driven by the electricity billing scheme. Second, an optimization framework from prior work was extended to evaluate the energy storage benefits in a distributed environment and especially its interaction with cost reduction through request off-loading. Third, as a result of previous two, we draw the following conclusions. With tighter SLAs, the savings potential from geographic demand redirection rapidly decreases, while strategic use of energy storage (e.g., in UPS batteries) can more stably provide an additional cost reduction of about 7% over various SLAs. In markets with high price variability these savings are 10% or more of the monthly bill. These results were obtained for the electricity prices derived from spot markets. We also show how the battery capacity and different billing models affect the utility of energy storage approaches for multi-site services. While batteries can provide further cost reduction even for an *On/Off peak* billing contract, we find that their potential benefit is significantly reduced in multi-site environments that are billed based on peak power (rather than total energy), because an intelligent request distribution can already avoid excessive power peaks.

The remainder of this paper is structured as follows. In the next section we offer background material on electricity pricing and on UPS batteries. Section III presents our optimization framework for electricity cost reduction, together with a proposed policy for using it. Section IV first gives the methodology used in this paper, and then our results and evaluations. Section V discusses additional related issues, and Section VI discusses related work. Finally, we present our conclusions in Section VII.

II. Background

Before we discuss the details of our optimization framework and its results, this section presents background information on two specific sub-topics related to this work.

A. Electricity Billing Schemes

Various electricity pricing schemes are found in practice. Nowadays, data centers are exposed to rich electricity price variability worldwide. The well-known on- and off-peak pricing charges two different per-KWh prices depending on the time of the day. There are also scenarios where in addition to per-KWh pricing, there are also additional surcharges based on the peak power draw seen at any point in a given billing period. For some data centers, the peak component can account for as much as 40% of the bill [10].

In other pricing schemes, electricity charges are even more variable, drawing entirely from pricing on a spot market [6]. Though electricity spot markets function at the wholesale level, spot prices reflect on eventual retail rates. In some states with advanced demand-side management, retail rates for large customers such as data centers may be based on a direct pass through of the hourly wholesale spot prices. Even if they are not directly passed through, spot prices are related to the eventual retail rates [6].

Our work assumes day-ahead energy markets, in which energy producers offer energy consumers a 24-hour-advance "menu" of per-hour energy prices. Since electricity is not directly storable, prices are much more likely to be driven by spot demand than any other good. Furthermore, price spikes may be motivated by disruption in transmission, generation outages, extreme weather, or a conjunction of these circumstances [9]. Figure 1 shows, however, an April, 2008 example of how electricity prices show quite repetitive behavior [1][1]. In this price trace, we remarked that there are normally two price peaks per day, one around 11 am and another at 9 pm.

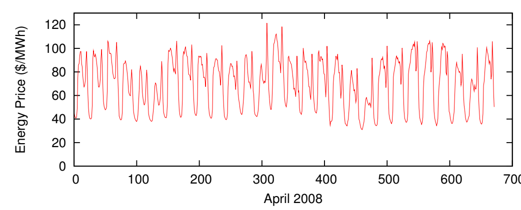

Figure 1. PJM - day ahead market: Hourly energy price for April 2008 shows repetitive behavior.

B. UPS Units

The UPS units contain batteries that are used to bridge the time between a power failure and availability of diesel generators' power. After a failure, generators typically take 10-15 seconds to start and produce enough power for full load [3]. However, UPS units are designed to sustain a minimum of 8-10 minutes of peak load [10]. In practice, the batteries can power the data center much longer. In

[1]Prices are obtained from PJM that operates the world's largest competitive wholesale electricity market and one of of North America's largest power grids.

part, this is because data centers are normally underutilized, consuming less than peak power draw. Furthermore, data centers may also be overprovisioned with UPS batteries in order to provide redundancy in case of UPS unit failures, or to allow data center upgrades over time.

Using UPS batteries for cost-savings, in addition to their original purpose, may raise some concerns which we address here. First, since the UPS units should protect data center availability, our optimization framework can be set up to allow a safety margin of stored battery energy sufficient for safe transition to generator power in the case of a failure. A second issue is that battery lifetime may be affected by extra charges and discharges, and by the depth of discharge. Since batteries are used both for cost savings and for power outages, the number of charging cycles can increase. Furthermore, deeper battery discharges lead to shorter battery lifetimes. Thus, our framework allows one to set parameterized limits on the number of discharges per day. Our default setup uses limits of 2 discharges per day allowing for a battery lifetime of 4 years [10].

Another aspect of battery usage is energy loss. A certain amount of energy is lost while discharging: about 10-15% for lead-acid batteries [18]. Furthermore, the battery can be leaky so that the level of stored energy decreases over time. For simplicity, our work is similar to prior work in assuming there are no energy losses, but these could be added to the optimization framework as well.

III. Optimization Problem

A. Problem Overview

Figure 2 depicts an example of an Internet service with a front-end and three data centers geographically distributed. The front-end gets electricity prices for the next 12 intervals (one scheduling period) from the electricity market and predicts the number of requests for each interval. It solves an optimization problem and decides about the fractions of load to be directed to each data center for each of the intervals. The front-end also determines battery charge/discharge schedules. After 12 intervals, the whole process is repeated for the next scheduling period.

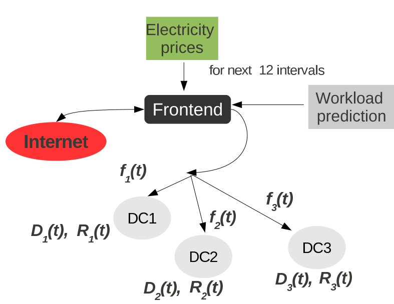

Figure 2. Request distribution: Over an interval t, data center *i* receives $f_i(t)$ fraction of all requests to serve. Over the same interval, the data center discharges/recharges its batteries for $D_i(t)/R_i(t)$ amount of energy.

The goal of our request distribution policy is to minimize the electricity cost of a multi-site Internet service under the constraint that the SLA is satisfied. Besides the electricity price variability across different data centers, it exploits temporal price variability within a data center storing energy in UPS batteries.

B. Optimization Problem Formulation

Given the problem and optimization goals described in the previous subsection, Table I gives optimization problem variables and their meaning. Namely, we aim to **minimize** the total $ElectricityCost$ given by the equation from Figure 3. This function is the total electricity bill paid by the Internet service. It represents the sum of per-data-center, per-interval costs that consist of the cost to execute requests assigned to the data center and the cost to power the data center. While i denotes the data center, t indicates time intervals. The first part of the costs corresponds to the dynamic energy consumption due to request processing $((f_i(t)LT(t)E_i)$, whilst the other represents the energy consumed by powered-on equipment whether idle or busy $(BaseE_i(offered_i(t)))$. We refer to the second type of energy consumption as base energy. Since the request policy estimates the load that will be directed to each data center during an interval, this is used to power only servers sufficient to serve the predicted load during the interval plus a safety margin of 20%. Furthermore, the cost to recharge batteries is also added to $ElectricityCost$ and the savings due to battery discharges are taken into account $((R_i(t) - D_i(t))Battery_i))$. The energy drawn from the grid over an interval at a given data center is multiplied by the price corresponding to the interval and the data center $(EnergyPrice_i(t))$.

The objective function is minimized subject to the constraints 1-9 given in Figure 3. Constraints 1 and 2 ensure that each data center gets a non-negative number of requests and that all user requests are distributed (and served). Each data center has its load capacity measured in the number of requests it can process per second. This number must not be violated (constraint 3). Constraint 4 specifies that the SLA must be satisfied. That is, the percentage of requests that complete within L time must be at least P. Constraints 5 and 6 regulate how much and how often the batteries can be charged/discharged since the function $sign$ gives value 0 when the argument is zero and 1 for positive arguments. Over an interval, batteries of a data center can not be simultaneously recharged and discharged (constraint 7). Note that batteries of a data center can be recharged while they are discharged in another data center. The amount of stored energy is determined by the battery recharges/discharges (constraint 8). Finally, the data center's load computed in constraint 9 is used to determine request completion times and base energies.

Since electricity prices from a day-ahead market are given on hourly basis, we determine request distribution fractions

Symbol	Description
$f_i(t)$	The portion of requests to be forwarded to center i in interval t (value between 0 and 1)
$R_i(t)$	The portion of battery capacity that will be recharged in center i over interval t (value between 0 and 1)
$D_i(t)$	The portion of battery capacity that will be discharged in center i over interval t (value between 0 and 1)
$Battery_i$	Total battery capacity found in center i (in Wh)
Y^i_{min}	Safety minimum given as battery capacity fraction below which batteries in center i are not discharged (value between o and 1)
$Y_i(t)$	Battery charge level in center i at the end of interval t
N_{cycles}	The allowed number of charging/discharging cycles per scheduling period
LC_i	Load capacity of center i (reqs/sec)
$LT(t)$	Expected total service load for interval t (number of reqs)
$LR(t)$	Expected peak service load rate for interval t (reqs/sec)
$offered_i(t)$	Expected peak load assigned to center i for interval t (reqs/sec)
$SLA(L,P)$	SLA given as a percentage P of request that must have latency not higher than L
$CDF_i(L,l)$	Expected percentage of requests sent to center i that complete within L time given l load
E_i	Energy to execute a request in center i due to dynamic power consumption (in Wh)
$BaseE_i(l)$	Base energy consumed over interval t in center i under load l due to static power consumption (in Wh)
$EnergyPrice_i(t)$	Energy price over interval t for center i ($/Wh)

Table I
VARIABLES AND PARAMETERS USED WITHIN THE POLICY ($f_i(t)$, $R_i(t)$ AND $D_i(t)$ ARE DECISION VARIABLES).

Objective function:

$$ElectricityCost = \sum_t \sum_i (f_i(t)LT(t)E_i + BaseE_i(offered_i(t)) + (R_i(t) - D_i(t))Battery_i) * EnergyPrice_i(t)$$

Problem constraints:
1. $\forall t \forall i\ f_i(t) \geq 0$ *i.e. each fraction cannot be negative*
2. $\forall t\ \sum_i f_i(t) = 1$ *i.e. the fractions need to add up to 1*
3. $\forall t \forall i\ offered_i(t) \leq LC_i$ *i.e. a data center should not be overloaded*
4. $\sum_t \sum_i (f_i(t)LT(t)CDF_i(L, offered_i(t)))/\sum_t LT(t) \geq P$ *i.e. the SLA must be satisfied*
5. $\forall t \forall i\ 0 \leq R_i(t) \leq 1 - Y_i(t)\ ,\quad 0 \leq D_i(t) \leq Y_i(t) - Y^i_{min}$ *i.e. battery cannot be recharged more than its capacity is, neither it can be discharged over the security limit*
6. $\forall i\ \sum_t sign(R_i(t)) \leq N_{cycles}\ ,\quad \sum_t sign(D_i(t)) \leq N_{cycles}$ *i.e. batteries of a data center can be recharged/discharged a limited number of times*
7. $\forall t \forall i\ sign(D_i(t)) * sign(R_i(t)) = 0$ *i.e. over an interval batteries of a data center can be either recharged or discharged*
8. $\forall t \forall i\ Y_i(t+1) = Y_i(t) - D_i(t) + R_i(t)$ *i.e. battery level depends on the recharge/discharge amount and its previous state*
9. $\forall t \forall i\ offered_i(t) = f_i(t)LR(t)$ *i.e. the peak load expected in a data center is the assigned fraction of the total peak service load*

Figure 3. The optimization problem formulation.

on an hourly basis as well. Hence, an interval duration corresponds to an hour. Examining the energy price traces, we conclude there are two price peaks per day on average. With this in mind, we solve the optimization problem for the next 12 intervals (hours), though other lookahead windows are possible. One charging/discharging cycle is permitted per data center and problem solution, allowing for a maximum of two charging/discharging cycles per day. In this way, the battery lifetime is not endangered (Section II-B).

Our implementation of the policy uses Simulated Annealing [11] to solve the problem as in previous work [13]. This presents a light overhead for the front-end, occurring only twice per day.

IV. EVALUATION

A. Methodology

We simulate an Internet service that includes a single front-end that distributes user requests among three data centers located on the East Coast, West Coast and in Europe. The front-end is assumed to be on the East Coast. The simulator takes as input electricity price and request traces. It first solves the optimization problem described in previous section. The optimization problem is solved for next scheduling period based on predicted load. Predicted load for each of the intervals of next scheduling period is needed to solve the optimization problem. Workload prediction done in this work is explained below. Once the problem is solved, the simulator gets request fractions and battery charging/discharging schedule. These values are used to simulate the policy and resulting costs using the actual load. The default simulation settings are given in Table II.

Electricity Price Trace. We have obtained hourly energy prices from PJM day-ahead market [1] that correspond to prices on the US East coast in April 2008 (when the request trace was obtained, see below). For the other two data centers, this electricity price trace has been shifted for time zone differences. In this way, each data center has the same electricity price trace. However, in a given moment, each of the data centers may have a different electricity price.

Workload Trace. The request trace represent a 1-month-long real trace from a commercial search engine, *Ask.com*. The trace corresponds to a fraction of requests the service received during April 2008. The load capacity of each data center was assumed to be 250 requests/sec. This load capacity was set to match the intensity of our request trace. While the energy prices are known in advance from the day-ahead markets, request load prediction is used while solving the optimization problem. We use Auto-Regressive Integrated Moving Average (ARIMA) [4] modeling to predict loads. The modeling combines both seasonal and non-seasonal components additively. The non-seasonal component involves 3 auto-regressive parameters and 3 moving average parameters (corresponding to the past three hours) while the seasonal component involves 1 auto-regressive an 1 moving average parameter that corresponds to the same hour of the previous week. This method accurately predicts the load [13].

For simplicity, we assume that all requests are of the same type and can be served by all data centers. Also, we assume that a request takes 400 ms to process on average. The default SLA used in simulations requires that 90% of the requests complete in 500 ms or less, but we also perform sensitivity experiments for different SLAs. The SLA is to be satisfied at the end of scheduling period (12 hours), in all simulations.

Network Latencies. Real experiments with servers located in the 3 previously mentioned regions have been performed to generate a realistic distribution of data center response times [13]. The requests were made to last 400 ms on average at a remote server, according to a Poisson distribution, and they were issued from a client on the US East Coast. The server's response time was measured at 4 utilization levels (20%, 40%, 60% and 80%) to instantiate CDF_i tables. The time between consecutive requests issued by the client followed a Poisson distribution with the appropriate average for the utilization level. The results showed that higher utilization has only a small impact on the response time. Average measured response times were: 412 ms (East Coast), 485 ms (West Coast) and 521 ms (Europe). With respect to our default SLA (500 ms, 90%), only the East Coast server can complete more than 90% of requests within 500 ms. The other can only reply within 500 ms 76% (West Coast) and 16% (Europe) of the time.

Energy Usage Model. In the default case, we assume energy-proportional data centers [2]. An energy-proportional server does not consume base energy i.e. $BaseE_i = 0$. Thus far, servers are not yet perfectly energy-proportional, but base energy is expected to decrease significantly in the next few years because of strong industrial and academic initiatives [2], [14].

Battery Capacity. Battery capacity is considered in terms of the number of requests that can be processed out of the energy stored in the batteries. We evaluate battery capacities of 200K, 400K, 600K, 800K and 900K requests. With the load capacity of 250 requests per second, the data center load capacity is 900 000 requests per hour. Accordingly, the energy to serve 200,000 requests corresponds to the energy consumed over 13.3 minutes at the peak load. The highest evaluated battery capacity (900K) means that the battery is sufficient to run the entire data center at peak load for an hour. Since these capacities might be higher than ones available in data centers, we discuss the amortization of the excess battery capacity in Section V. The given capacities are assumed to be entirely available for electricity cost reduction, i.e. the safety amount of energy necessary for power outages has been subtracted. This is also covered in Section V. Note that all storage capacities are given per data center and that each of the three data centers is assumed to have the same storage capacity in one simulation run.

B. Results

In this section we evaluate UPS battery potential for electricity cost reduction in a distributed environment in which an intelligent request distribution is already done to reduce the electricity costs. The evaluation is given for various scenarios including different pricing schemes, battery capacities and SLAs.

1) Battery Capacity and Price Variability : It has been shown that in a single data center environment, the battery's ability to exploit electricity price variability depends significantly on battery capacity [18]. However, it is not clear how battery capacity affects the costs in a distributed environment where requests can also be directed to a center with low electricity price. Success also depends on the price difference between when the energy is stored and when it is consumed. Price variability is not the same in all spot energy markets, nor in all seasons. In our trace, the highest price difference over the entire trace is a factor of 3.9X. The average daily minimum price of 39.66 $/MWh and the average maximum 99.54 $/MWh give the daily variation

Setting	*SLA*	*Base Power*	LC_i	*Billing Scheme*	*Interval Duration*	*Scheduling Period*
Default Value	(500ms,90%)	0	250 req/s	day-ahead market	1 hour	12 intervals

Table II
DEFAULT SIMULATION PARAMETERS.

factor of 2.5X. Higher price variation observed was reported in the related work [16], [18], reporting variations as much as 10X from one hour to the next [16]. Here, we explore a range of variabilities by scaling only the peak prices by 2 and 5 to obtain approximate daily price variation of 5X and 12.5X respectively.

Figure 4 gives the electricity cost with our request distribution policy for different battery capacities and price variation traces. The electricity cost is normalized with respect to the case with no energy storage and corresponding price variation.

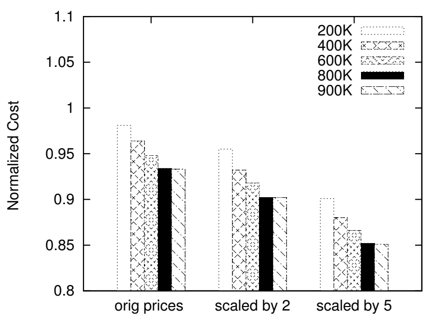

Figure 4. Cost dependence on battery capacity (from 200K requests to 900K) and price variability (original prices, peak prices scaled by 2 and 5).

For the original price trace, using a small battery capacity (200K) that sustains the peak data center load for less than 15 min gives 2% of savings. The results show that for each 15-min-increase in battery capacity the electricity cost is further reduced by slightly more than 1.5%. Assuming batteries that can run the data center for an hour leads to the cost reduction of 6.7%. We tested the policy for a single data center set-up with appropriate load capacity and found this cost reduction to be 6.2%. Hence, batteries can achieve savings even in multi-site environments where data centers with lower electricity prices can be exploited through intelligent request distribution. Moreover, these savings are slightly higher in the multi-site environment since the total number of requests served from stored energy is not limited to 2 hours per day as in the single data center scenario but to 6 hours (though with lower load capacities).

Rather high price variation when peak prices are scaled by 5 leads to the savings of 15%. Cost reduction of 10% is possible if the electricity prices fluctuate similarly to [18] (the case when peak prices are scaled by 2). These results show that exploiting batteries makes the most sense when price variation is quite high. Also, we evaluated an unlimited case in which all three data centers have infinite load capacities and there is no SLA constraint. In this scenario, energy storage can reduce further the costs obtained with request off-loading by 7.6% with the batteries of 900K. Hence, the savings from stored energy are more limited by battery constraints and price variability.

2) Other Electricity Billing Models: So far we assumed electricity prices from day-ahead markets. In this section we look at the other two most common pricing models. With on-peak/off-peak pricing, there are two energy prices, one for *on-peak* hours (weekdays from 8am to 8pm) and another for *off-peak* hours (weekdays from 8pm to 8am and weekends). Figure 5(a) shows the saving potential under this pricing model. Here we assume, the off-peak price of half of the on-peak price.

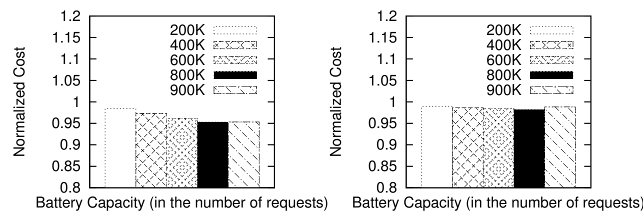

(a) On/Off peak hour pricing (b) Peak power draw based

Figure 5. Normalized cost for different electricity billing.

With this billing scheme, our optimization framework can reduce the cost through energy storage by 4.7%. Taking into account the total amount paid by data centers every month on electricity costs, this still represents substantial savings. Again, if the price difference between *on-peak* and *off-peak* hours were even higher, there would be greater savings. Furthermore, in this scenario, the energy price increases once per day meaning that batteries are recharged/discharged only once per day. This is important taking into account the battery lifetime concerns.

Finally, we look at one more electricity billing scheme that we refer to as *peak-based*. Namely, utilities might bill data centers separately for the peak power drawn over the month. This component of the bill is added to the cost for energy consumed. The peak over the month is the maximum power observed over any shorter interval (e.g. 30 minutes). We had to modify the solver to optimize the electricity cost under this billing scheme. The electricity cost function defined in Section III-B was increased by the sum of peak power draw components for each data center, $\sum_i P_i Price_i$, where P_i represents the peak (maximal) value of power draw averaged over a given duration in the $i-th$ data center and $Price_i$ is a price to be paid for each Watt of the peak draw ($/W).

The problem constraints from Section III-B stayed the same. Again, the policy optimizes the cost locally, for each of 12-hour durations of a scheduling period. Requests are distributed such that spikes in demand directed to any of the data centers are avoided. In the case of battery usage, the optimization policy recharges them when low load is assigned to the corresponding data center and discharges under high load intensity.

In our work, we look at average power over 1-hour durations and track the peak hourly average. The electricity price and the price to be paid in USD per MW of the peak power draw are set in the following way. We assume that the peak draw component of the total bill accounts for 40% of the total electricity bill if the data center half of the time draws the peak power and the rest of the time half of the peak power (based on the analysis from [10]). Note that peak power is tracked for each data center separately and the total bill includes the peak power costs of all data centers.

Achieved savings are much lower than for other two billing models. They range from 1% to 2% as shown in Figure 5(b), though the peak draw component of the cost accounts for 34% of the bill in the non-battery case. Since we limit battery discharges to one hour over 12-hour durations, only the highest peak can be shaved. As there might be more similar peaks, the second highest will determine the cost. It is also important not to generate a new power draw peak due to recharging batteries of large capacity.

Figure 6 depicts request distribution and battery management for the peak-based billing scheme. In Figure 6(a) we can see how the number of requests for each data center forms a smooth line when no batteries are used, meaning that the intelligent request distribution avoids peaks in load within a data center. In Figure 6(b) where batteries of the smallest capacity are assumed, there are peaks in the load that are shaved through battery discharges. However, we were not able to substantially reduce the cost that was achieved through request off-loading. We can conclude that for this billing model, the request distribution already minimizes the cost, leaving little space for battery management to improve it.

3) Varying SLA Strictness: When SLAs are stricter, they restrict the gains possible from geographic request distribution, because sending requests to a far away data center increases latency. On the other hand, shifting power draw via UPS battery storage within a data center does not affect the response time. Thus, batteries may offer better leverage in scenarios with strict SLAs. To test this, we independently vary both SLA parameters, latency and percentage, while applying our policy. Figure 7(a) shows electricity costs for the cases without batteries and with the battery capacity sufficient to run data center's load for an hour. The costs are given for different latencies, normalized to the cost with

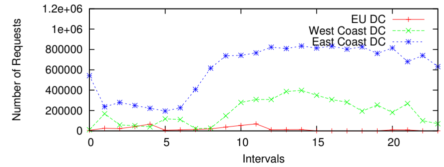

(a) No battery case: SLA (500 ms, 90%)

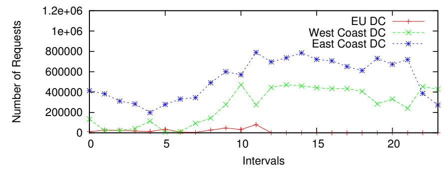

(b) Battery case - 200K: SLA (500 ms, 90%)

Figure 6. Number of requests assigned to each data center for peak draw based billing.

the default SLA (500ms, 90%) without batteries.

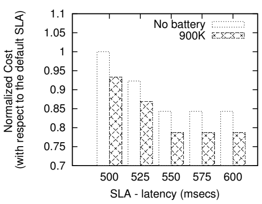

(a) Fixed $P = 90\%$ and variable latency

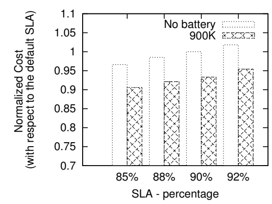

(b) Fixed $L = 500ms$ and variable percentage

Figure 7. Cost dependence on SLA.

We also varied the percentage condition of SLA. Electricity cost for different values of the SLA percentage parameter (assuming target latency of 500 ms) are given in Figure 7(b). Here, the rightmost columns represent the costs for the most strict SLA that requires 92% of requests to be processed within 500 ms.

Looking only at the non-battery results, we can see how much a stricter SLA affects the efficiency of off-loading. When our default latency of 500 ms is increased to 550 ms, the cost drops 15.7%. These results show that savings based only on request distribution highly depend on the SLA. However, for all SLAs the energy storage gives additional savings over the ones achieved by request off-loading since it does not necessarily require demand routing to possibly far away data centers. For all observed SLAs, batteries reduce the electricity bill obtained with request offloading by 6% to 7% meaning that absolute savings through batteries are higher for stricter SLAs.

Energy storage can be used to impose a tighter SLA while paying the same costs. For instance, for the 90% SLA the latency requirement of 525 ms can be set to a more strict one, 500 ms, for the same cost by introducing the battery management (see Figure 7(a)).

4) Energy Proportionality: Here we analyze effects of servers' energy-proportionality. While success of the load offloading approach highly depends on the amount of energy consumed by servers that is load independent [13], [16], the energy storage approach for pricing models different from the peak based should not be very sensitive to the servers' energy-proportionality. Recall that all results we discussed so far assumed no base energy and that we assume that in each data center there are as many active servers as sufficient to serve the expected load plus a margin of 20%.

Figures 8 shows the electricity cost as a function of the amount of base power. Base power of 75 W and 150 W is considered. Since we assumed a dynamic power range of 150 W, the case of base power of 150 W roughly represents today's servers. The cost are normalized to those of the non-battery policy and the corresponding base power consumption. Results are given assuming battery capacities sufficient to sustain 15 min, 30 min and 1 hour of a completely loaded data center. Note that for each of the base energy scenarios the amount of energy that can be stored in batteries differs.

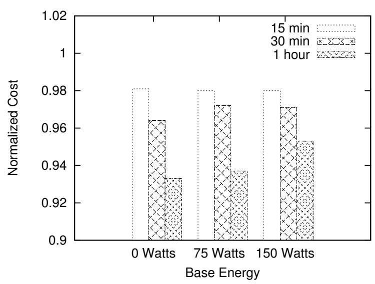

Figure 8. Normalized cost for different base energy.

The results show little difference across different base energy scenarios. An energy-proportional system (that enables complete power draw distribution through request distribution) benefits slightly more than a system with significant idle/static power consumption. With batteries sufficient to run data centers for an hour, the energy proportional system (base power = 0 W) achieves 2% higher savings than the system with high base power (150 W). This difference comes from impossibility to redirect base energy consumption to the data center currently run on batteries (a fraction of the battery benefit is achieved through request off-loading to the currently battery run data centers). This suggests that when servers become more energy-proportional, distributed systems will benefit slightly more from the energy storage than they would today.

V. Discussion

In our evaluation we have used battery capacities higher than the ones normally found in data centers, but our analysis in this section shows that the cost of this battery capacity can be recouped by the electricity price savings it allows. As we explained (see Section II-B), data centers are normally provisioned to run at least for 10 minutes of peak load on UPS batteries though much less time is needed in the case of power outage.

In the case of higher price variability, such as when peaks from the PJM price trace were scaled by 2 or 5, even small battery capacities achieved substantial savings of at least 5%. However, for lower price variability, 40 minutes of battery capacity is necessary to have 5% of gains. Data centers that are not provisioned with this much capacity could invest into additional batteries. We estimate the time needed to recover from such an investment and to start to save through intelligent battery management. The recover time is computed based on the following equation:

$$P_{peak} * AdditionalCapacity * BatteryPrice =$$
$$Savings * P_{avg} * ElectricityPrice * \boldsymbol{RecoverTime}$$

where P_{peak} is the peak data center's power consumption while P_{avg} represents the average data center's power consumption. $AdditionalCapacity$ is the bought battery capacity given in terms of the time over which it can power the data center at peak load. $Savings$ represents the cost reduction achievable with the total capacity (after purchase).

Table III gives the recover time in years for pairs of the previous battery capacity and the capacity after the upgrade. The results are given for lead-acid batteries, the most common battery type today in data centers. In the calculation it was assumed that 5 minutes of batteries are kept for a potential power outage. The battery price used is 100 $/KWh and the electricity price is assumed to be 10 c/KWh [10]. Also, we estimate the average power consumption to be 50% of the peak power consumption. These results are given for the case of original price traces with lower price variability.

		Total capacity after battery purchase			
		400K	**600K**	**800K**	**900K**
Existing	**10 min**	2.29	2.56	2.78	3.12
capacity	**20 min**	1.23	1.83	2.21	2.55
before	**30 min**	0.17	1.09	1.63	1.98
purchase	**40 min**	0	0.36	1.05	1.42

Table III
Amortization time (in years) to recoup additional battery costs and to start to save; estimations are given for the conservative case of low electricity price variation.

For instance, if a data center with battery capacity of 20 minutes is upgraded to battery capacity sufficient for 400K requests plus 5 minutes for power outages, the time it takes to recover from the investment and to start to save is 1.23 years. Note that for higher battery capacity longer time to recover is needed but afterwards savings are higher. Since the average lifetime of this type of battery is 3-5 years [10], all of these recover times mean that after a certain point the additional investment would lead to savings.

VI. Related Work

There has been considerable research on energy conservation in data centers [7], [8], [15]. These works dynamically determine how many servers are required to serve the current load and turn the others off to eliminate idle system power consumption. The problem of idle power consumption can be also tackled through DVFS and low-power modes [14], [17].

Recently, approaches that exploit geographical and temporal electricity price variability have been proposed. Multi-site Internet services can leverage geographical price variability through demand redirection to data centers with lower electricity price [12], [16]. Another policy that routes demand of a multi-site Internet service to reduce the electricity costs taking into account brown energy caps has been designed [13].

Another way to intelligently schedule power draw taking into account the pricing model is by using UPS batteries [10], [18]. Urgaonkar *et al* proposed an optimization based algorithm to minimize the electricity bill by storing energy in UPS batteries when the electricity price is lower and using it when the price is high [18]. Price variability comes from an energy spot market. Govindan *et al* proposed to use batteries for peak shaving with a peak based pricing scheme. Over slots with high load energy from batteries is used to reduce the peak billing component [10]. In both works, the authors look only at a single data center. Hence, it was not clear how much energy storage can contribute in a large-scale distributed system where request distribution can be used to shave both price and load peaks that increase the bill.

VII. Conclusions

In this paper, we presented a framework to investigate the effects of various design decisions and operating conditions on the electricity costs of multi-site Internet services. Furthermore, we proposed an optimization-based request distribution and battery management policy that exploits both geographical and temporal electricity price variability to reduce the electricity bill. An extensive set of simulation experiments based on real traces was performed to evaluate the potential of this holistic approach for different battery capacities, electricity price variabilities, SLAs, pricing schemes and servers' energy proportionality.

Based on our results, we conclude that in order to further reduce the electricity costs in a distributed environment through energy storage either higher capacity batteries or higher temporal price variability are needed. While the energy storage approach is especially effective for billing schemes derived from spot markets, in multi-site environments it is less effective for peak power draw based billing. Also, we found that the absolute cost reduction of this leverage is higher for stricter SLAs when the possibility for requested off-loading to currently cheaper data centers is limited. Furthermore, batteries can be used to impose a stricter SLA for the same electricity costs as when only spatial price variability is exploited.

Acknowledgment

This research was performed while Maja Etinski was a visiting student at Princeton University. We would like to thank the HIPEAC Network of Excellence for supporting this visit through the collaboration grant program.

References

[1] Pjm - day-ahead lmp data. *http://www.pjm.com/markets-and-operations/energy/day-ahead/lmpda.aspx*.

[2] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. *Computer*, 40(12), 2007.

[3] L. A. Barroso and U. Hölzle. The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis Lectures on Computer Architecture*, 4, 2009.

[4] G. Box and G. Jenkins. Time series analysis, forecasting and control. *Holden-Day Incorporated*, 1990.

[5] E. V. Carrera, E. Pinheiro, and R. Bianchini. Conserving disk energy in network servers. In *Proceedings of the 17th annual International Conference on Supercomputing*, ICS '03, 2003.

[6] J. Chandley. How RTOs establish spot market prices. *http://www.pjm.com/~media/documents/reports/spot-market-prices-j-chandley.ashx*, 2007.

[7] J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, and R. P. Doyle. Managing energy and server resources in hosting centers. In *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles*, SOSP '01, 2001.

[8] E. N. Elnozahy, M. Kistler, and R. Rajamony. Energy-efficient server clusters. In *In Proceedings of the 2nd Workshop on Power-Aware Computing Systems*, 2002.

[9] H. Geman and A. Roncoroni. Understanding the fine structure of electricity prices. *The Journal of Business*, 79, May 2006.

[10] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. Benefits and limitations of tapping into stored energy for datacenters. In *Proceedings of the 38th annual International Symposium on Computer Architecture*, ISCA '11, 2011.

[11] S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220, May 1983.

[12] K. Le, R. Bianchini, M. Martonosi, and T. D. Nguyen. Cost- and energy-aware load distribution across data centers. *Workshop on Power Aware Computing and Systems*, 2009.

[13] K. Le, R. Bianchini, T. D. Nguyen, O. Bilgir, and M. Martonosi. Capping the brown energy consumption of internet services at low cost. *International Conference on Green Computing*, 2010.

[14] D. Meisner, B. T. Gold, and T. F. Wenisch. Powernap: eliminating server idle power. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '09, 2009.

[15] E. Pinheiro, R. Bianchini, E. V. Carrera, and T. Heath. Load balancing and unbalancing for power and performance in cluster-based systems. In *Workshop on Compilers and Operating Systems for Low Power*, 2001.

[16] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the electric bill for internet-scale systems. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, 2009.

[17] N. Tolia, Z. Wang, M. Marwah, C. Bash, P. Ranganathan, and X. Zhu. Delivering energy proportionality with non energy-proportional systems: optimizing the ensemble. In *Proceedings of the 2008 Conference on Power Aware Computing and Systems*, HotPower'08, 2008.

[18] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam. Optimal power cost management using stored energy in data centers. In *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, 2011.