Plugging Into Energy Market Diversity

Asfandyar Qureshi (MIT CSAIL)

ABSTRACT

In North America, electricity prices exhibit both temporal and geographic variation—the later exists due to regional demand differences, transmission inefficiencies and generation diversity. Using historical market data, we characterize the variation and argue that existing distributed systems should be able to exploit it for significant economic gains. We consider pricing in cloud computing systems, and also use simulation to estimate the advantage of dynamically shuffling computation between different energy markets.

1 INTRODUCTION

Electricity is becoming increasingly expensive, and now accounts for a large fraction of the cost of ownership for data centers [1]. It is expected that by 2012, in the US, 3-year energy costs for data centers will be at least twice as much as the server investment [2].

At the same time, deregulation, regional demand variations and energy source diversity have resulted in an uneven and occasionally volatile cost landscape. In the US, electricity prices at two different places can have very different annual averages (figure 1), and prices at a location can vary day to day by a factor of five (figure 4).

Like cost, the *utility* gained by a distributed system's clients may also depend on location. Generally, a client receives less utility if their request is served far away from them. Many existing systems typically maintain multiple replicas, routing clients to nearest replicas, attempting to maximize client utility, while ignoring the geographic variation of cost.

In such replicated systems, it is possible to trade-off between computing in a high cost market versus computing in a lower cost market but with reduced client utility. Shifting clients away from their best replicas, to ones situated in cheaper energy markets, may reduce quality-of-service but yields significant monetary savings.

To some extent, this trade-off is implicit in the placement of large data centers in low-cost energy markets (Google in Oregon and Microsoft in Illinois) rather than in high-demand locations (e.g. New York City). We argue that, due to existing price volatility, this trade-off should be a dynamic choice rather than a static one.

This paper investigates the implications of electricity price volatility and locational variation to Internet scale systems. We argue that there is something to be gained, by building price-sensitive distributed systems, that automatically integrate up-to-date market information, and make cost/performance trade-offs.

We sketch the connection between computing cost and energy cost and establish the significance of locational variation. Using historical electricity market data, we show that the day-to-day, monthly, and yearly variation is substantial. We note that daily prices, at locations near Internet peering points, exhibit exploitable volatility.

We briefly cover how cloud computing providers could increase their margins by being sensitive to geographic variation in energy prices—either with price differentiation or by using cost-optimized routing.

Finally, we use simulation and a 2006-2008 history of US market prices to explore cost/performance trade-offs within Internet-scale replicated systems. We simulate *selective blackouts*, where one or more replicas are deactivated in response to market signals. We quantify possible energy cost savings and discuss practical implications.

To the best of our knowledge, this paper contains the first proposal for distributed systems to use online optimization to algorithmically exploit information from electricity futures and/or spot markets.

2 BACKGROUND

2.1 Concerns about Electricity Cost

Data center energy costs are becoming an increasingly dominant component of overall operating costs. The cost of electricity is poised to overtake the cost of equipment [3]. In the US: in 2000 three-year energy costs were one-tenth the server equipment expenditures; by 2009 the cost of electricity is expected to equal server expenditure; and by 2012, energy is expected to cost at least twice the equipment investment [2]. These expectations take into account recent advances in data center energy efficiency. For a denser non-traditional data center (e.g., Sun's S20 [4]), 2-year energy costs could already exceed the equipment cost, depending on configuration and location.

Additionally, in absolute terms, servers consume a substantial amount of electricity. Servers and their support infrastructure (e.g., cooling) accounted for about 1.2% of US electricity consumption in 2005, about 45 million MWh, or 2.7 billion dollars [5]. By 2010, this is projected to grow to 3% of total US consumption [5].

Consequently, for companies with large computing facilities, even a fractional reduction in electricity costs can translate into a large overall savings. For example, it was estimated that Google owned 450,000 servers worldwide in 2006 [6] and that each server consumed upwards of 200 watts [7]. Each watt used by a computer results in at least two watts drawn from the electric grid [1, 3]. We can, conservatively, estimate that Google servers used around 1.6 million MWh in a year, or 95 million dollars

Location	2004	2005	2006	2007
New York (NYC)	63.1	93.5	70.9	77.1
New England (MA)	53.7	78.6	60.9	67.9
Southwest (Palo Verde)	50.1	67.4	57.6	61.7
Southeast (SERC/FRCC)	48.6	70.8	55.5	59.1
PJM Interconnect (West)	41.7	60.6	50.1	56.9
Northwest (MID-C)	44.5	63.0	50.2	56.6
California (NP-15)	38.4	54.4	43.4	54.6
Texas (ERCOT-North)	42.3	66.5	51.4	52.0
Midwest (Cinergy)		38.4	40.5	46.1

Figure 1: Annual average prices [9], in $/MWh, sorted by 2007 prices.

worth of electricity, at US rates[1]. Therefore, every 1% savings in energy cost could save a large company like Google, a million dollars a year. Google is not alone. Microsoft expects to deploy 800,000 servers by 2011 [6], and the five leading search companies may have already deployed more than 2 million servers [8].

New cooling technologies, architectural redesigns, DC power, multi-core servers, virtualization and energy aware load balancing algorithms, have all been proposed as ways to reduce the energy consumed by a single data center. That work is complementary to ours. However, this paper is concerned with reducing *cost*—our approach can achieve this, even if it causes consumption to rise.

2.2 Electricity Markets

Although market details differ regionally, this section provides a high-level view of deregulated electricity markets, providing a context for the rest of the paper. The discussion is based on markets in North America, but the ideas generalize to other regions with diversified markets.

Electricity is produced by government utilities and independent power producers using a variety of sources. In the United States, this includes nuclear (about 10%), coal (around 30%), natural gas (nearly 40%) and hydroelectric (roughly 8%) [10].

Producers and consumers are connected to an electric *grid* of transmission lines. Electricity cannot be stored easily, so supply and demand must continuously be balanced. In addition to connecting nearby nodes, the grid can be used to import and export electricity from/to distant locations. The United States is divided into ten markets [9], with varying degrees of inter-connectivity. Congestion on the grid, transmission line losses, and market seams issues either limit how electricity can flow, or influence the price at a given location [11].

The existence of rapid price fluctuations reflects the fact that short term demand for electricity is far more elastic than short term supply. Electricity cannot always be efficiently moved from low demand areas to high demand areas, and power plants cannot always ramp up easily. In contrast, we have long used high performance networks and load balancing techniques to relocate computation. We can move our demand closer to a low-cost supply.

[1] $450,000 \times 200W \times 2 \times 24h \times 365 = 1.5678 \times 10^{12} Wh$ @ $6¢/kWh$

Location	Nearby City	Hub	Market
A	San Jose, CA	NP15	California
B	San Diego, CA	SP15	California
C	Portland, OR	MID-C	Northwest
D	Chicago, IL	Illinois	Midwest
E	Ashburn, VA	PJM-West	PJM
F	Houston, TX	ERCOT-H	ERCOT
G	Miami, FL	Florida	Florida

Figure 2: Seven locations, near different Internet exchange points, and in different electricity markets.

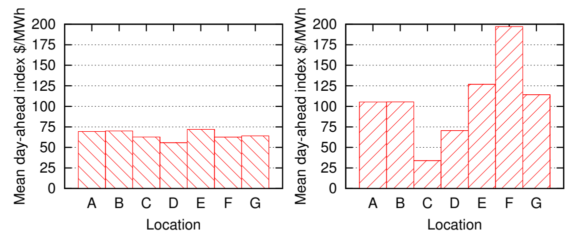

(a) October 2007 **(b) June 2008**

Figure 3: Monthly variation in wholesale market prices can be substantial and geographically dissimilar. For example, comparing (b) with (a): prices at F trebled, while those at C halved.

While short-term and long-term contracts may account for most of what is consumed, electricity can also be bought in *wholesale* markets. In most regions, day-ahead, hour-ahead, and spot markets exist. In this paper we focus on day-ahead markets. Such markets allow consumers to determine the price of electricity the day before it is delivered. Day-ahead prices are forward signals, that can be used to decide how much to consume.

A caveat: companies running data centers may have contracts with electricity providers, do not buy directly from the wholesale market, and so may be buffered from the price volatility we are looking to exploit. Contractual details are hard to come by; this paper ignores contracts.

In reality, there is a great deal more complexity, but our market model is simple: a futures market exists; day-ahead prices are accessible and variable; and different locations see prices that are not perfectly correlated.

2.3 Computation Cost

A service provider accepts requests, performs some computations, and produces responses. The provider incurs some cost in fulfilling this demand.

We model the total computation cost (C) incurred by a service provider at a given location as follows: a large fixed component, the infrastructure cost (I), and a significant variable component (V), which is a monotonically increasing function of demand.

In this formulation, I includes the amortized infrastructure investment, staff salaries, etc. V includes both network and energy costs, but we ignore network costs. Studies have shown that electricity consumption closely follows CPU utilization [12]. Using techniques like multi-core CPUs and virtualization, resources can be allocated on-demand, causing electricity use to step up.

The *marginal* computation cost is the incremental cost

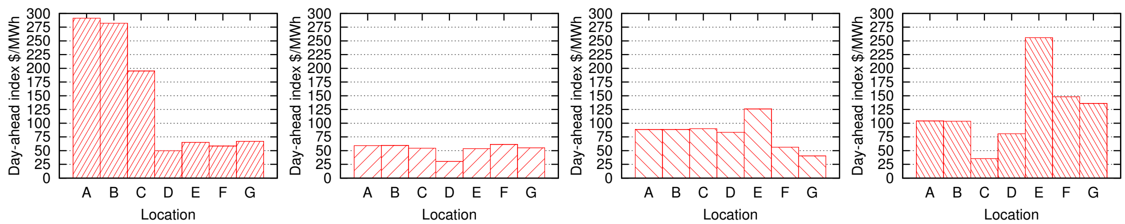

(a) July 05, 2007 **(b) Sept. 14, 2007** **(c) Jan. 22, 2008** **(d) June 10, 2008**

Figure 4: Day-ahead wholesale market prices exhibit significant volatility. For example, prices at E were $54/MWh in (b) but $256/MWh in (d).

of handling one more request, or the derivative of C w.r.t. demand. The *average* cost is the mean cost per request, or C divided by demand.

Differences in electricity prices will always show up in the marginal computation cost for different locations, assuming constant server energy efficiency. If electricity costs are a large enough fraction of overall cost, price volatility will begin to palpably affect average cost.

3 Electricity Price Variation

This paper posits that electricity prices vary dynamically, that prices at different locations are not perfectly correlated, and that differences can be exploited for economic gain. Rather than presenting a theoretical discussion, we take the experimental approach, grounding ourselves in historical market data, from multiple sources [9, 10, 13].

We begin with average annual electricity prices, tabulated in figure 1 for several locations. In 2007, Northeast prices were over 1.5 times Midwest prices, contributing to the impracticality of large data centers in the Northeast.

The remainder of the paper focuses on the smaller set of seven locations from figure 2, all of which are near major Internet exchange points (IXPs) and cover a number of diverse electricity markets.

Apart from annual variation, prices also exhibit seasonal and monthly variation. Figure 3 shows average prices for two different months. In the South, in June '08 the energy needed to handle a million requests would have cost twice as much in Houston (location F) compared to Miami (G). In October '07, the cost difference would have been relatively insignificant. Similarly, on the West coast, in June, electricity in California (A) was thrice as expensive as electricity in Oregon (C), but in October prices were roughly the same. Furthermore, the relative ordering of prices was very different in the two months. Houston (F), for example, moved from the second cheapest market to the most expensive.

Part of the market diversity arises because different regions produce electricity in different ways. For example, in 2006: in Oregon, natural gas accounted for 8% and hydroelectric for 68% of the summer generation capacity; whereas in Texas, natural gas accounted for 71% and coal for 20% of the summer capacity [10]. Consequently, record high natural gas prices in 2008 have had much larger impact on Texas than on Oregon.

Prices in wholesale markets also exhibit significant day-to-day volatility, for a variety of reasons. For example, a localized event such as a heat wave in California could drive up local demand, elevating West-coast prices. Figure 4 shows day-ahead prices for four different days. Price spikes such as those shown in figure 4a and figure 4d occasionally occur. Price volatility has many hard-to-predict causes (e.g., accidents, equipment malfunctions, weather, fuel costs, demand volatility, market manipulation, etc.). Figure 5 shows a more detailed picture for some locations, plotting the evolution of day-ahead market prices from January 2006 through June 2008. Some notable features: seasonal effects, short-term spikes, and only partially correlated behaviour. A detailed discussion is beyond the scope of this paper.

In this paper we restrict ourselves to day-ahead market prices. However, significantly more price volatility exists in hour-ahead and spot markets [11]. Traditional consumers cannot respond quickly enough, but distributed systems can re-route computation at millisecond scale, to modulate their consumption. Beyond our findings in this paper, there may be opportunities within spot and hour-ahead markets, that traditional electricity consumers cannot exploit, but distributed systems can exploit.

4 Pricing in Clouds

With the rise of web-based computing and the computing-as-a-utility model, many companies are renting out their infrastructure to third-party applications. Examples include Amazon's EC2, Google's AppEngine and Sun's Grid. Applications are billed by the resources they consume: computation cycles, network I/O and storage.

How much does it cost a provider to perform one unit of work on behalf of a hosted application? How much does it cost Amazon to handle a single client request on behalf of a hosted web application?

Cost depends on where the request is routed. We have already established that marginal computation costs can differ radically with location and in time. Furthermore, refer back to the cost model from section 2.3. Large cloud providers (Amazon and Google) will already need to absorb their fixed costs. They need to build multiple data centers, and keep machines up and running, to support their own primary services. The cost to them of performing some incremental work on behalf of a hosted applica-

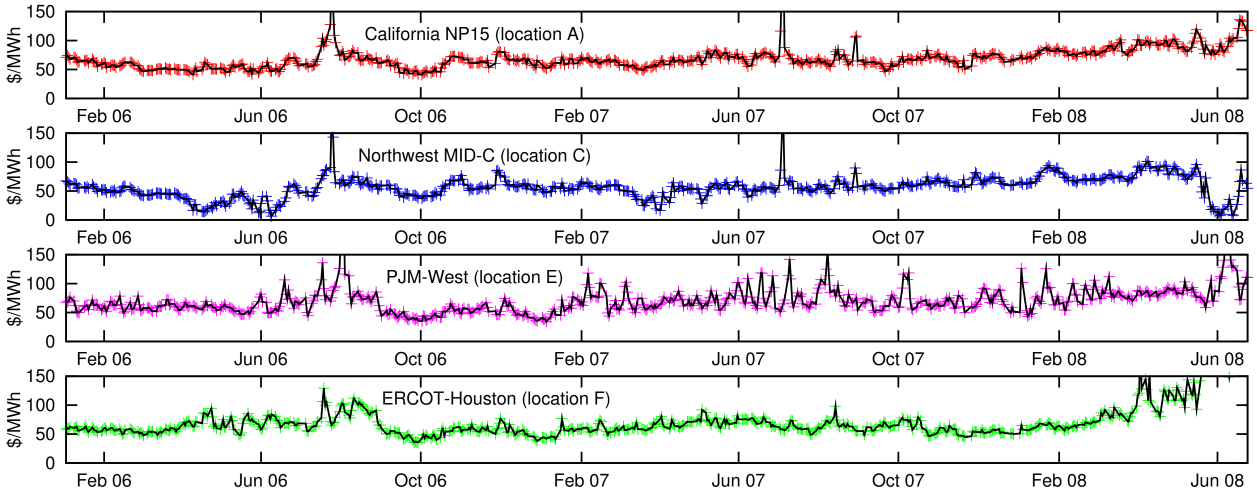

Figure 5: Day ahead index prices at different hubs, from January 2006 through June 2008 [13]. Note the seasonal dips in the hydro dominated Northwest, and the 2008 upward trend in California and Texas, both of which are heavily dependent on natural gas. Price spikes reached $350.

tion will be dominated by the marginal cost, mainly the cost of the additional watts they expend.

By charging a fixed compute-price, while being able to decide where to buy electricity, cloud providers are missing an opportunity. With a price structure that embraces energy cost diversity, and by using a cost-conscious replication strategy, cloud providers can increase their margins or lower their prices.

Buyers care about how much they are charged and what performance their users receive. Providers can build energy cost differences into some pricing plans, allowing buyers to make trade-offs. For example, free applications should always be hosted in the lowest cost locations, capacity permitting. Additionally, some buyers may be willing to pay premiums for regionally optimized performance. The *Dallas Morning News* website, having regionally concentrated demand, values proximity, and can therefore be billed to compensate for elevated prices.

These ideas can be mapped to content distribution networks. For instance, a CDN provider could charge a premium for hosting content in high energy cost markets.

5 Selective Blackouts

Internet-scale systems composed of replicas in different electricity markets can exploit price disparities to substantially reduce their total energy costs, by using Information from energy futures markets, and dynamically shifting consumption away from high-cost regions. Through simulation, we show that an approach based on this idea could yield considerable monetary savings.

5.1 System Model

In the systems we focus on, storage and computing infrastructure can be decomposed into a number of blocks, where each block is a complete replica of the system[2].

[2]Less flexible but acceptable: strict subsets are complete replicas.

The blocks may be:

- For large companies, the blocks are large data centers, owned and operated by the company. Each block can have many thousands of physical machines, and easily consume 4500 kW [1].
- The blocks can be much smaller data centers. In the extreme, blocks may be one or more of Sun's data-center-in-a-container [4], each with fewer than 300 machines and 500 kW of peak consumption.
- For small providers, the different blocks can be leased floor-space in data centers owned by other parties[3]. The main difference between this case and the above cases is control over infrastructure: in the earlier cases if the provider decided to turn off the machines, they can also shut off cooling etc.

An incoming client request to such a system can be served by any of the replicas. In existing systems, replicas tend to be placed near IXPs, such as the locations in figure 2. Conventionally such systems attempt to keep all replica locations active. In order to maximize performance, client requests are routed to their closest replicas.

In the discussion that follows, we assume that the system is over-provisioned: some subset of the replicas has enough capacity to handle the peak load.

We also make a number of simplifying assumptions. We model demand as constant and uniformly distributed. When some blocks are deactivated, we assume client requests will be spread evenly over the remaining blocks and that total energy use is therefore constant. Furthermore, we assume that a deactivated block consumes zero energy, and that the startup/shutdown process also consumes no energy[4]. Finally, we assume that shutting down one replica does not affect prices at any other replicas.

[3] Our work is only relevant when electricity charges are metered.

[4] This ignores the cost of synchronization during replica reactivation.

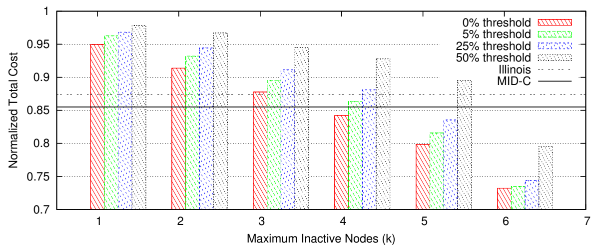

Figure 6: Total electricity costs for seven replica simulations, using 2006-2008 market data. A cost of 1.0 represents running all seven.

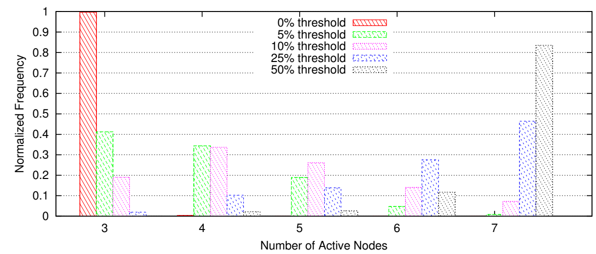

Figure 7: The distribution of the number of active replicas, from simulations using 2006-2008 market data, with $n = 7$ and $k = 4$.

We use the number of active replicas as a first-order approximation for performance. We defer a proper analysis of the performance impact of our proposal.

5.2 Selective Blackouts

With enough excess capacity, one or more replicas can be turned off. This will result in suboptimal system performance and reduce reliability, but can also significantly reduce energy costs.

Deciding which replicas should be active on any given day can be modeled as an optimization problem. Each day, day-ahead market prices can be fed into an automated mechanism that determines which replicas should be deactivated the next day. The set of active replicas changes infrequently, at most once per day, making this compatible with existing routing techniques (e.g., DNS).

Given n replicas, we constrain that no more than k replicas can be deactivated on any given day. Thus the $(n-k)$ lowest cost replicas are always active, regardless of absolute prices. This provides a consistent performance baseline. Replicas remain active as long as their prices are close to the highest price we must pay for baseline performance. We only force deactivation when a significant price disparity exists.

More formally, given day-ahead prices, we derive the set of active replicas A as follows:

$$\begin{aligned} L &= \{(n-k) \text{ lowest cost replicas}\} \\ \phi &= max(\{price_r \text{ for } r \in L\}) \\ A &= \{\text{replica } r \text{ iff } price_r \leq (1+\tau)\cdot\phi\} \end{aligned}$$

τ is a threshold parameter, expressing our sensitivity to price disparity, as a percentage of the baseline price ϕ.

5.3 Simulation Results

We simulated the above selective blackout mechanism using historical prices, wholesale market data from 2006 through 2008 [13], and found that significant cost savings were possible. Demand was modeled as being constant in time and uniformly distributed in space.

North America Seven. We first simulated a seven-node system, one node at each location from figure 2. Figures 6 and 7 summarize the results.

Simulations imply that adding a single redundant node can reduce total electricity costs by 5% (see figure 6, 0% threshold and $k = 1$). These savings are the result of being able to dynamically deactivate nodes during periods of locally elevated prices. Statically picking the best six locations is not enough. Section 3 already illustrated that an always optimal set of six may not exist, and our simulations reinforced this: all nodes were active some of the time. The most active 6-subset accounted for 27% and the next most active 23%, so no subset dominated.

For $k \geq 4$, blackouts can result in total energy costs lower than the average cost of the least expensive market (MID-C line in figure 6). With only one baseline node and six redundant nodes ($\tau = 0$, $k = 6$), energy cost is 85% that of the cheapest node. This is a savings of 27%, compared to running all seven.

The threshold parameter τ can be used to trade off between cost and performance. Figure 7 shows how the number of active replicas depends on τ for $k = 4$. With a threshold of 5%, the median number of active replicas is 4 ($\mu = 3.9$) and the total cost roughly matches the cheapest market (see figure 6). With a threshold of 25%, the median number of active replicas is 6 ($\mu = 6.1$) and the cost is close to the second-cheapest market. At the same time, in contrast with building a large data center in a cheap market, computation resources are now more likely to be near an IXP that provides a fast path to a random client. This can dramatically improve performance.

West Coast Three. With all seven nodes, we can take advantage of regional diversity (e.g., a heat-wave in California does not put pressure on the Illinois hub). Even though, nearby locations in the same market tend to have correlated prices, selective blackouts can still be useful.

To demonstrate this, we simulated a three node west-coast system (NP15, SP15 and MID-C). With blackouts (50% threshold, $k = 2$) the resulting total cost is 6% lower than the cost of continuously running all three, and 8% higher than the cost of computing everything in Oregon. The median number of active replicas is 3 ($\mu = 2.7$). For this to work, Oregon must retain maximum capacity—on some days it is the only active replica. See figure 8.

5.4 Some Lessons

Building extra, occasionally deactivated, replicas will incur some additional infrastructure cost. However, we have shown extra replicas can reduce total energy costs.

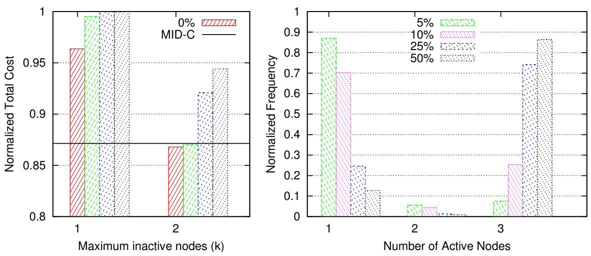

(a) Electricity Costs **(b) Active replica distribution, k=2**

Figure 8: West coast simulation results.

If expected energy savings exceed the up-front infrastructure investment and added maintenance costs, it makes economic sense to use some variant of this blackout mechanism. If data center energy costs in the US double in the next four years [2], or if the replicas are modular data centers [4] with low fixed costs, dynamic blackouts could make a tangible impact on operating costs.

Additionally, apart from yielding savings, blackouts can reduce risk, by dampening the impact of unanticipated price fluctuations. The mechanism described here, for example, would automatically integrate market information and route around multi-day weather problems.

The choice of where to build a data center is typically seen as a *static* optimization problem. If energy costs continue to rise relative to equipment, it may be better modeled as a *dynamic* problem. Despite the economies of scale inherent to large data centers and the possibility of local tax incentives, a company looking to build a monolith should consider building many smaller blocks (e.g., [4]) spread over different energy markets. Redundant capacity is already built into these systems. It may be better to spread these resources, rather than concentrating them.

6 Conclusion

We set out to show that the diversity and day-to-day volatility of today's electricity markets can be exploited in some meaningful way by existing distributed systems. Using data from wholesale electricity markets and simulation, we were able to show that replicated systems can make meaningful cost/performance trade-offs and may be able to achieve substantial energy cost reductions. Many possibilities for future work exist within this area.

In order to understand the trade-offs, a good performance model is necessary. We use the number of active replicas as a coarse performance metric. A better approach would have been to analyze the network latencies between clients and active replicas, assuming a uniform client distribution, using census data, or using server logs. The impact on reliability should also be considered.

Another convenient simplification was to assume constant demand. In reality, demand varies regionally and temporally [14, 15]. Depending on the situation, there may be ways to exploit features within demand signals.

We presented selective blackouts as an illustration of a price-conscious optimization mechanism, rather than as a proposed design. A mature mechanism should synthesize information from both supply (cost) and demand (performance/utility) and derive the best way to use available resources. In addition, hour-ahead and spot-prices are more volatile than day-ahead prices, so more frequent optimization should yield higher savings.

Further, our idea of relating energy costs to computation costs implies that auctions within computing grids can be used to match buyers and sellers, to increase the total economic surplus in the computing utility market.

Finally, contracts complicate the picture, making it unclear who would reap the savings we calculated. Power providers may be willing to index charges to market prices, since this transfers some risk to consumers. If, on the other hand, contracts fix the cost of electricity, a deactivated data center would allow the producer to sell the surplus electricity on the wholesale market. While this would not impact the data center's bottom line, the provider would benefit, and—if resource scarcity has caused the price elevation—the public would benefit.

Acknowledgements

This paper was improved by helpful comments from John Guttag, Hari Balakrishnan, Srikanth Kandula, Vladimir Bychkovsky, and the anonymous reviewers. This work was supported in part by Nokia.

References

[1] J. G. Koomey, K. G. Brill, P. Turner, J. Stanley, and B. Taylor, "A Simple Model for Determining True Total Cost of Ownership for Data Centers," white paper, Uptime Institute, 2007.

[2] K. G. Brill, "The Invisible Crisis in the Data Center: The Economic Meltdown of Moore's Law," white paper, Uptime Institute, 2007.

[3] L. Marchant, "Data Center Growing Pains," in *USENIX Large Installation System Admin. Conference*, 2007. Presentation.

[4] Sun Microsystems, Modular Datacenter `http://www.sun.com/products/sunmd`.

[5] J. G. Koomey, "Estimating Total Power Consumption by Servers in the U.S. and the World," tech. rep., LBNL, 2007.

[6] J. Markoff and S. Hansell, "Hiding in Plain Sight, Google Seeks an Expansion of Power," *the New York Times*, June 2006.

[7] L. A. Barroso, "The price of performance," *ACM Queue*, 2005.

[8] G. Gilder, "The Information Factories," *Wired mag.*, Oct. 2006.

[9] United States Federal Energy Regulatory Commission, Market Oversight. `http://www.ferc.gov`.

[10] United States Department of Energy, Official Statistics. `http://www.eia.doe.gov/fuelelectric.html`.

[11] P. L. Joskow, "Markets for Power in the United States: an Interim Assessment," tech. rep., MIT CEEPR, Aug. 2005.

[12] X. Fan, W.-D. Weber, and L. A. Barroso, "Power Provisioning for a Warehouse-sized Computer," in *ACM International Symposium on Computer Architecture*, 2007.

[13] Platts, "Day-Ahead Market Prices," in *Megawatt Daily*, McGraw-Hill. 2006-2008.

[14] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng, "Understanding User Behavior in Large-Scale Video-on-Demand Systems," in *European Conference in Computer Systems*, 2006.

[15] M. Arlitt and T. Jin, "A workload characterization study of the 1998 world cup web site," *IEEE Network*, vol. 14, no. 3, 2000.