ReRack: Power Simulation for Data Centers with Renewable Energy Generation

Michael Brown and Jose Renau
Dept. of Computer Engineering, University of California Santa Cruz
http://masc.cse.ucsc.edu

ABSTRACT

Data centers operating cost are dominated by their power consumption. Renewable energy sources can reduce the operating costs when correctly selected. Nevertheless, this is a non trivial task because it should consider different energy sources (wind, solar), storage alternatives (batteries, grid-tie), workload, historical weather patterns, incentives, and service agreement penalties. These are the basic factors, but the model should be extensive to consider other factors like power gating support.

This paper introduces ReRack, an extensible simulation infrastructure that can be used to evaluate the energy cost of a data center using renewable energy sources. ReRack is composed of two main components, a simulation and an optimization component. The simulation component explores the cost effectiveness of renewable energy generation. It requires a model that can simulate both the data center power usage and the location-dependent variability of the power generation source (solar and wind). The ReRack optimization component finds the best ratio of renewable energy sources for a given location and workload.

1. INTRODUCTION

Building a data center is an expensive proposition, and the energy to power it is a significant part of the expense. More so since every watt of power used for computation requires more power to be spent turning a fan to keep a data center's temperature inside the operating range. 2010 was expected to be the year in which lifetime electricity costs associated with owning new servers would exceed the price of the hardware itself [1]. Recently there has been a more concerted effort to increase efficiency of power generation to cut costs. This has made for a resurgence of renewable energy sources as a viable alternative [2]. Taking advantage of this opportunity is not trivial; renewable sources are inherently variable and have different cost characteristics depending on location (weather and local laws). Renewable generation sources can have large variation in output depending on the local weather patterns, so exploring their feasibility in data centers requires some way of estimating the cost of running a data center under different power generation configurations. To be feasible in data centers one would need to significantly decrease the renewable energy cost due to the generation capacity and adapt to energy cost fluctuations found in renewable sources. This paper proposes ReRack to provide such simulation infrastructure.

ReRack can be used to model and optimize capital investment and operational costs associated with power in a data center using renewable energy sources. This paper is an overview of what and how the model works, and the full source code will be made available upon the paper's publication. ReRack can evaluate the cost for a given mix of renewable sources, the weather parameters over a period of time which affects the renewable energy generation, and the workload of real-time and batch jobs during the same time period. For example, to know that power associated costs for a data center powered by 500kW of solar photo-voltaic capacity, 250kW of wind turbine capacity, 400kWh of vanadium redox flow battery storage, and local grid energy, doing the same workload for one year would cost $417,086 in California, $390,902 in Texas, and $306,220 in Michigan. ReRack not only evaluates the cost, but it also optimizes to find the best mix of renewable sources that minimize cost for workload and a location's weather history.

2. RERACK APPROACH

ReRack has two components, the simulator and the optimizer. The simulator estimates the energy cost for a given configuration. The optimizer looks for the best combination of energy sources to minimize the cost of the simulation. With this partitioning, the problem can be handled by fairly well known techniques.

2.1 Simulator

ReRack is a simulation infrastructure that can be used to predict the relative cost of powering a data center by exploring a multi-dimensional range of configurations to find the most inexpensive available. The ReRack simulation currently considers three power generation sources (wind, solar, and grid) as well as power storage, given a fixed server set and characterization and a known daily distribution for demand. The infrastructure is built to be easily extended to allow optimizing over server size and characterization, varying demand distributions, other power generation sources, or nearly any other consideration.

The simulator cost function would be $f(a,b,c,...) = cost$. Where the inputs (a, b, c, etc.) are the list of all variables considered pertinent to the cost of a data center over the time span of a year. This could include, but is not limited to, workload, energy generation cost, data center energy consumption model, and renewable sources. However, instead of a simple equation, the cost function is actually a simulator calculating the actions and effects of all the variables used on a data center's power usage and performance. Usage of grid power is recorded, as is the cost of grid power at the time it was being used. Time steps are set at 1 minute to match the granularity at which power is metered.

To calculate this cost the simulation does the following in a loop for every minute of a full year:

1. Calculate power supply from renewable sources
2. Calculate power demand from active machines
3. If (power is in surplus)
   (a) Store it (or sell it - California & Michigan)
4. Else
   (a) Tap batteries and pay to draw the remaining difference from the grid

5. Calculate how much workload was satisfied
6. Determine how to power up machines to for the next minute
   (a) Decision based on likely power generation, workload demands, the cost of power from the grid, the cost of breaking a service level agreement, and how far the current time is through the billing period

At the end of the simulation the money spent buying power from the grid is added to the amortized costs of purchasing and maintaining the renewable sources. This is returned as the cost for that particular solution.

2.2 Optimizer

The optimizer use is a genetic algorithm, though the particular optimizer is not important and others could be used. The cost function is not an equation, but instead a simulation using an amalgamation of batteries (non-linear, truncation), weather (random data), and goal-oriented decisions about uptime (logic). All of that "bad" behavior in the variables makes for a simulation which can not be expressed as an equation and a solution space which contains lots of local minima and discontinuities; hence the use of an optimizer like the genetic algorithm. Though, like all non-analytical methods, the confidence in the answer as optimal is proportional to the effort applied to the search. Maximally pre-computed inputs greatly assist in keeping search times manageable and by extension, result accuracy high. Additional factors can be considered in the model by just incorporating the relevant new variables into the cost function, and increasing the search time allowed for optimization. In this way ReRack provides a compartmentalized environment where new (such as cooling costs) or more detailed [3,4] component models may be inserted with minimal effort.

The optimizer, in this instance a genetic algorithm, attempts to find the combination of inputs to the cost function that yields the optimal output. Here the cost function is:

$$f\begin{pmatrix}\text{weather, cost incentives, solar panels,}\\ \text{wind turbines, batteries, grid rate}\end{pmatrix} = \text{cost}$$

Some optimizers find provably optimal solutions. A genetic algorithm will not yield a provably optimal solution. However, if applied correctly, in practice it will yield "good" results which are at or near the optimal. Part of applying it correctly is allowing the algorithm enough time to get coverage of the search space. If the search is limited to a single site at a time then three of the six inputs to the cost function become constants (weather, grid rate, and cost incentives; all location specific). Lowering the search space to three dimensions, pre-computing as many of the inputs to the cost function as possible, and optimizing the cost function code makes good coverage of the search space much easier to achieve.

The genetic algorithm code base we used was GAUL [5]. The seed population was 1000 unique configurations spaced in a 3D-grid, each dimension of the grid having a range from zero to twice the total power needs of the data center if met by that resource alone. The algorithm runs, mutating solutions and keeping the most optimal, until the population converges to an optimal fitness. There may or may not be more than one unique solution that produces this optimal fitness so the top 20 solutions are output, though generally only the first is used.

2.3 Sample Utilization

An example situation of what ReRack is intended to model would be the following: A data center builder wants to compare the energy costs including amortization for three different locations (Table 1). Each location has a different wind, solar, energy costs, and grid power cost model. Each site has its own merits. California's central valley gets a lot of sun during the day and has relatively cheap power at night. Very important, California has several financial incentive programs that lower the capital investment necessary. Texas's gulf coast gets about the same amount of sun exposure, but also has vast amounts of wind. Though, Texas has none of the incentive programs found in California. Finally, northern Michigan (not the lakes) gets neither sun, nor wind in amounts useful for renewable energy generation. However, the price of energy generated by fossil fuels is significantly cheaper in Michigan than in either of the other two states, which could out-compete renewable sources all together.

	California	Texas	Michigan
Grid power (cents/kWh)	8.3-14.0	10.2	7.8-2.2
Avg. wind (m/s/day)	< 4.0	6.5	7.5
Avg. sunlight (h/day)	5.5	4.5	4.2

Table 1: Location characteristics.

In addition to the energy costs itself, another key factor is the capacity to sell back the excess electricity generated with the renewable sources, and the energy cost model. The excess generated electricity has very different models depending on the state. While in California the utility company is required to buy back up to 1MW of power, Michigan only requires 20kW, and Texas has no such program whatsoever. Also state dependent is the energy cost model. The energy bill is not only a factor of the kWh consumed, but the consumption pattern. More details are provided in the following section about the input set.

When the fixed data center is simulated in ReRack under each of the three configurations, the tool determines that relative to the other two choices Michigan would be the most cost effective site for the data center, followed by California and last, Texas. Any change in the data center itself, such as workload characteristics, would require a new run through the tool to see if the new traces corresponded better or worse to the availability of each of the power sources. The new fit would not only mean a different set of costs, but a new suggestion on the most efficient way to allocate funds when sizing generation capacity. Winds tend to blow when temperature is in flux, while the sun has its own fairly predictable window of use. ReRack is meant to take information about specific candidate sites, simulate minute by minute power usage while trying to make the best use of the available resources and find the most efficient means of powering the data center.

3. RERACK INPUT SET AND MODELS

The previous section has provided a high level view of the ReRack approach to simulate the energy cost for a given data center location. This section provides more details about the input sets and models used in the ReRack simulation.

ReRack model can be divided into a data center energy consumption model, workload, service level agreement, energy generation for wind and solar which is dependent on the weather patterns, battery model, grid energy cost, and a grid buy-back cost model.

Data Center Energy Consumption: Particulars about the data center hardware will also make a difference in modeling it. Power and performance characteristics of the system used are taken from [6]. Energy consumption for the system was measured and reduced to one simple equation with inputs of percent utilization or activity of the CPU, memory, storage device, and network. A well characterized workload can provide all of the necessary inputs for the equations.

Systems are powered up and down between available DVFS states, always preferring the aggregate state that gives the data center as a whole the most efficiency in power, and by extension, cost per computation.

Workload: For this work we have generated workload traces at 1 minute granularity for CPU utilization. We have two traces, one for compute heavy load which was directly measured from a data center, and the other for interactive load which was modeled after data from Google [7]. Figure 1 has a one day sample of the interactive workload. An important part of the simulator is speculatively determining the number of machines to keep powered up for the next minute of operation. For this, queues are maintained to represent the amount of work, in the form of jobs, waiting to be done by the data center. A set of queues is maintained for two different types of workload. Interactive load is a high priority demand on the data center which must be satisfied immediately. Examples of this would be web browsing traffic or typing in user terminals. Interactive load is kept in a single queue representing the quantity of work remaining to be done, and is assumed to be satisfied in the order of arrival. The other kind of demand is batch compute load. Large, lower priority simulations like weather forecasts which, while not of immediate priority, do have service level agreements which impose fines if not met [8]. This type of load has a separate queue for each minute the work was first registered. In this way, we can know how long a job has been waiting and whether it has or will meet the service level agreement.

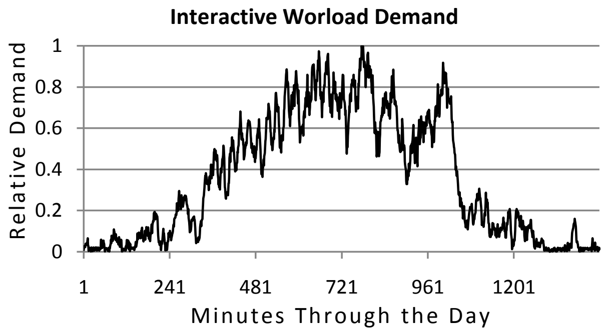

Figure 1: One day sample relative interactive workload used in simulation.

Service Level Agreement (SLA): The option of completing the work, at a cost related to the highest power load spike on the grid so far and to the current power rates and generation, is weighed against the option of putting off some or all of the work, at a cost related to the service agreement [8] and the job's time already spent sitting in the queue. Interactive workloads are penalized for any delay while batch work is allowed a two hour window for completion. Statistical analysis of average daily load and a two-hour weather forecast are used with information about the system's state to speculate the optimal number of systems to keep powered up. A feedback loop corrects for immediate errors in forecasting, while more complicated logic is used to target trade-offs between intricate power billing schedules and service level agreements.

Weather Pattern: The ReRack uses site specific details to determine when and in what quantity generating power sources will be available to the data center. In the ideal situation, a two year survey would be taken of a site which would record power output from a small solar array and a few small wind turbines. This data would then be used directly in the feasibility analysis of the site's power production capabilities. Since this information is not available in most cases, readings of solar intensity and wind speed can be used in conjunction with solar cell efficiency and wind speed-to-power generation curve to determine the power availability of both sources [9]. Solar intensity and wind speed can both be acquired easily from any number of sources (US NOAA NCDC records [10], WeatherUnderground, individual instrument owners). Likewise, solar efficiency and wind speed curves are easily acquired from any manufacturer, typically posted online as part of the spec [11, 12].

Solar Model: The solar power generation is a function of the weather patterns. Given a solar activity with an hourly granularity, the simulator estimates available power using the efficiency of the cell found in the data sheet. This is interpolated during simulation to a minute-by-minute granularity, and scaled to match the size of the solar array being simulated.

A major source of cost of solar energy is the initial or upfront cost for the solar modules. The solar model cost includes amortization over the lifetime of the product. By default, the simulator assumes a 20 year lifespan for solar cells and a 3.5% discount rate (use whatever rate is appropriate).

Wind Model: Wind power model is similar to the solar model. The amortization is the same, but the wind power generation is dependent on wind instead of sunlight. Data about the cost of wind turbines for a large scale deployment comes from [13]. The available wind energy is pre-calculated from hourly wind speed and the turbine's speed-to-power curve, taking into account a minimum cut-in speed as well as a maximum cut-off speed. The a year's worth of hourly data points are, like solar data, interpolated to a minute granularity during simulation.

Randomized cloud and wind gust models were incorporated into early versions of ReRack, but removed because no data could be found to verify their accuracy. With more focused data from site scouting these models could be re-introduced for increased accuracy.

Battery Model: Batteries or energy storage can be added to a data center to reduce the energy consumption peaks, but they have a cost and maximum power draw constraints. Battery backup should lower the cost of operation by storing renewable energy generated in excess of demand for later use, but should not offset the cost beyond the benefit. This is one of the more difficult things to determine since mathematically a battery is a non-linear element. Once the battery backup is fully charged it stops accepting power. Rather than modeling the aggregate behavior of batteries charging across the entire data center, the model uses a simpler view of charge rate as a straight line which truncates at maximum capacity. Maximum charge and discharge rates then become constants, which speeds simulation. The quantity of renewables and batteries to purchase is determined later by the optimizer, as is the outcome of which site will be the most cost effective choice. For the simulation in this paper, data sheets were gathered for characterization of batteries [14], solar panels [12, 15], and wind turbines [11]. ReRack extensibility allows to evaluate future or alternative storage sources.

Grid Energy Cost: The grid energy cost is particular to the local site. Typically, there is a rate schedule of the power company that services the area [16–18], as well as any applicable local regulations or rebates [19]. A power company has as many as thirty or more different schedules. Having estimated the power needs of the data center from the service demand, the peak possible power draw can be determined and a schedule chosen from there. In California, utility companies have adopted what is called "time-of-day" rates which provide financial incentives for customers to conserve energy during the peak draw hours of noon to 6pm, and use more power during the hours of 10pm to 8am (9,14). While currently the standard in California, many states are considering adoption by 2013. Also, if there is uncertainty as to which schedule should be targeted, the simulation of the data center system can be repeated with alternative schedules to see which would be the best fit for the customer.

Grid Buy Back Cost: Grid buy-back is when the customer, in this case the data center facility, generates more power than it can use and the power provider is legally required to buy that power from the customer. Buy-back not available in all states, and the quantity of power which the utility company is required to buy varies wildly between the states for which it is available. Texas utility companies are not required to participate in any buy-back programs. Michigan utility companies are required to buy up to 20kW of supply from an individual customer. California companies must be prepared to buy up to 1MW of excess generation from a customer, placing a credit toward their utility bill equal to the value of the power at the time of day at which it was sold. This can mean a net draw off the power grid at no cost for the facility if power is sold during peak hours and drawn in the night hours as is common place for solar and to some degree, wind.

4. RESULTS

Figure 2 is a two day window of a particular simulation. In California solar power is more heavily subsidized than any other power source [19], even requiring public utility companies to buy excess generation back from the customer. As Figure 2 shows, credits that can be used against the utility company's bill build up through the middle of the day. When the solar generation drops in the evening, the grid credits are cashed back in as the system draws power from the grid.

Another feature of note is that at hours 14 and 17 the service level agreement is broken for a small amount of the work. April 31st is the last day of the month and as the billing period resets on May 1st, the system determines that the penalties paid for breaking the service agreement are less than the cost of increasing the peak kW demand for this month, which would be required to complete all of the work on time. At the start of the new month the peak is readily increased because there are many days remaining in the month which could, at that time of day, benefit from having the peak raised. The higher peak means the freedom to power on more of the system to satisfy sudden spikes in demand. (The "peak" line shown is the record of peak demand for the month of May, not April, hence it is 0 in April). Table 2 contains the best

	California	Texas	Michigan
Solar (kW)	1216	0	0
Wind (kW)	0	545	97
Battery (kWh)	0	0	0
Relative Cost	**1.823**	**1.714**	**1**
Costs After 1 Year			
Renewable Generation ($)	128,799	48,105	8,562
Energy Cost ($)	275,991	393,378	212,917
Service Agr. Penalties ($)	0	0	0
Peak kW Demand Fee ($)	113,545	46,167	62,926
Total Estimate ($)	518,335	487,650	284,405

Table 2: Costs found by optimizer

configurations found by the optimizer and their associated costs and breakdowns. The total workload satisfied over the course of the year was the same for all three configurations as were all elements of the data center itself. What changed was the location and all of the variables that change when location does; power billing schedules, weather, price subsidies on renewable energy sources. These all created different influences on cost in ways that are not necessarily intuitive. ReRack is designed to make exploration of power costs across any factor of the data center simple. As a case study we implemented an additional feature; Table 3 is the same setup implementing near-instantaneous, perfect node-level power gating in the data center.

	California	Texas	Michigan
Solar (kW)	1207	0	0
Wind (kW)	0	322	24
Battery (kWh)	0	0	71
Relative Cost	**1.154**	**1.057**	**0.231**
Costs After 1 Year			
Renewable Generation ($)	127,846	28,422	12,379
Energy Cost ($)	107,172	231,619	0
Service Agr. Penalties ($)	5,691	2,365	2,344
Peak kW Demand Fee ($)	87,619	38,147	51,116
Total Estimate ($)	328,328	300,553	65,839

Table 3: Costs found by optimizer, after implementing node level power gating in the data center

5. RELATED WORKS

[20] did a survey of recent research on power efficient data centers. [21–25] studied power allocation within data centers and [26–30] created simulation, sensor and management tools to enable powering down servers under various conditions. [31–33] looked at how to take advantage of geographic and temporal changes in the cost of energy. [34] looked at quality-of-service aware scheduling policies using time, energy, and a system of task rewards or values. [35] optimized data center resources for performance per watt while we have optimized for performance per cost. [3, 4] created models to characterize renewable supply and workload demands but did not include a mechanism to search for an optimal configuration. [36] created a comprehensive cost model for computation in data center which, like ReRack, optimized for performance per cost, but focused on the choice of server size, utilization, and machine density. This differs from ReRack which focuses on taking advantage of the temporal changes in power availability for a data center largely powered with renewable sources. Also, [9, 37] have made detailed models for predicting renewable source output from weather data.

6. CONCLUSIONS

ReRack was designed to simplify exploration of costs associated with power in data center using renewable generation sources. It uses a simulator to predict the cost of any given configuration by following the data center for a full year and amortizing investment capitol over the lifetime of the resources purchased. Then ReRack uses an optimizer to search the design space for the most cost efficient solution. Several additional case studies were done but excluded for brevity. Likewise, this paper has been an overview of ReRack and the full code will be made available upon publication.

7. ACKNOWLEDGMENTS

We like to thank the reviewers for their feedback on the paper. This work was supported in part by the National Science Foundation under grants 0751222; and ARP 2010 award. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the NSF.

8. REFERENCES

[1] Jon Haas, "Greening it," arstechnia.com/business/news/2009/10/datacenter-energy-costs-outpacing-hardware-price.ars, 2009, Panelist speaker at Emerging Technologies 2009 (EmTech@MIT).

[2] Baryonyx Corp., "Baryonyx Corp. Website," www.baryonyx.com/projects.html, 2011, Company with active projects building 100% renewable energy data centers.

[3] Daniel Gmach, Jerry Rolia, Cullen Bash, Yuan Chen, Tom Christian, and Amip Shah, "Capacity planning and power management to exploit sustainable energy," *6th Int'l Conference on Network and Services Management*, October 2010.

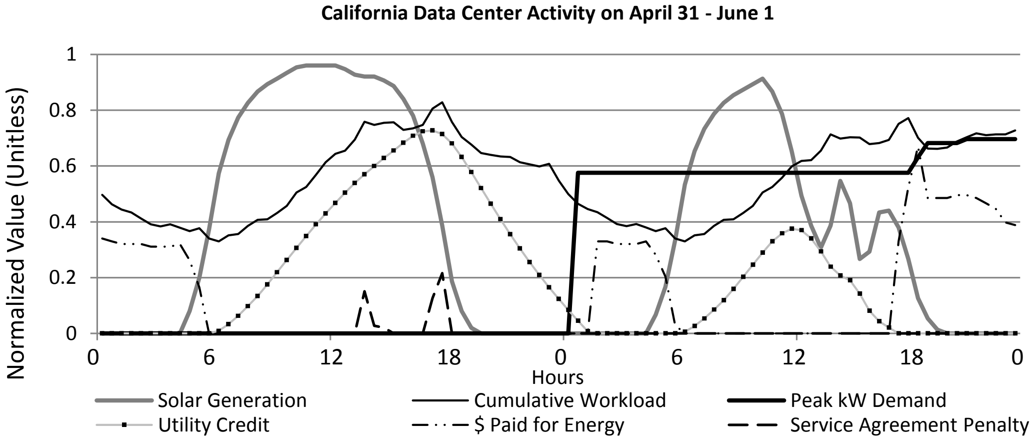

Figure 2: A simulation of a data center in California with 1MW of solar capacity. Here the data center saves money by breaking its service level agreements on the last day of the month, rather than incur a fee for higher peak KW usage on the last day of the billing cycle.

[4] Christopher Stewart and Kai Shen, "Some joules are more precious than others: Managing renewable energy in the datacenter," *2nd Workshop on Power-Aware Computing Systems*, Oct 2009.

[5] Steward Adcock, "Genetic algorithm utility library (gaul)," gaul.sourceforge.net, 2009, Genetic algorithm code base.

[6] Dimitris Economou, Suzanne Rivoire, and Christos Kozyrakis, "Full-system power analysis and modeling for server environments," in *In Workshop on Modeling Benchmarking and Simulation (MoBS)*, June 2006.

[7] Daniel Gmach, Jerry Rolia, Ludmila Cherkasova, and Alfons Kemper, "Workload analysis and demand prediction of enterprise data center applicatoins," in *IEEE Int'l Symposium on Workload Characterization (IISWC)*. June 2007, IEEE.

[8] Amazon Web Services, "Amazon Elastic Compute Cloud," aws.amazon.com/ec2, 2010.

[9] Chris Hatfield and David Umphress, "Solarcast: Using widely accessible environmental variables to forecast solar power," *46th ACM Southeast Conference*, March 2008.

[10] National Oceanic and Aironautics Administration, "NOAA Natioal Climatic Data Center," www.ncdc.noaa.gov/oa/ncdc.html, 2011, Climate Data Archive.

[11] Norwin America Inc., "Norwin 29-STALL 200kW Power Curve and Energy Production," 2010, Product Data Sheet.

[12] Kyocera Solar Inc., "KD210GX-LP Module," 2010, Data Sheet.

[13] Ryan Wiser and Mark Bolinger, "Annual report on u.s. wind power installation, cost, and performance trends," Tech. Rep., United States Department of Energy, 2007.

[14] Christopher Freitas, "Reducing the cost of electricity infrastructure: Redflow's zinc-bromide battery solutions," in *Smart Grid Summit*. June 2010, National Rural Electric Cooperative Association.

[15] SunRGI Inc., "SunRGI Inc. Website," www.sunrgi.com, 2011, Solar panel pricing.

[16] Pacific Gas and Electric, "E-20 Rate Schedule," 2010, Electricity Rate Schedule.

[17] Indiana Michigan Power Company, "Schedule of Tariffs Governing the Sale of Electricity," 2010, Electricity Rate Schedule.

[18] Oncor Electric Delivery Company, "Tariff for Retail Delivery Service," 2011, Electricity Rate Schedule.

[19] Pacific Gas and Electric, "California solar initiative," www.pge.com/myhome/saveenergymoney/solarenergy/csi/howitworks/, 2011, Financial Incentive Program for Renewable Energy Sources.

[20] Dilip Kandlur and Tom Keller, "Green data centers and hot chips," *46th Design Automation Conference (DAC)*, pp. 888–890, June 2009.

[21] Mark Femal and Vincent Freeh, "Boosting data center performance through non-uniform power allocation," *32nd Int'l Symp. on Computer Architecture (ISCA)*, pp. 250–261, June 2005.

[22] Anshul Gandhi, Mor Harchol-Balter, Rajarshi Das, and Charles Lefurgy, "Optimal power allocation in server farms," *Special Interest Group on Performance Evaluation (SIGMETRICS)*, June 2009.

[23] Ramakrishna Kotla, Soraya Ghiasi, Tom Keller, and Freeeman Rawson, "Sheduling processor voltage and frequency in server and cluster systems," *Int'l Parallel and Distributed Processing Symposium (IPDPS)*, p. 234, April 2005.

[24] Wes Felter, Karthick Rajamani, Tom Keller, and Cosmin Rusu, "A performance-conserving approach for reducing peak power consumption in server systems," *19th ACM International Conference on Supercomputing*, pp. 293–302, June 2005.

[25] Justin Meza, Mehul Shah, Parthasarathy Ranganathan, and Mike Fitzner, "Tracking the power in an enterprise descision support," *Int'l Symp on Low Power Electronics Design (ISLPED)*, August 2009.

[26] Rajarshi Das, Jeffery Kephart, Charles Lefurgy, Gerald Tesauro, David Levine, and Hoi Chan, "Autonomic multi-agent management of power and performance in data centers," *10th Int'l Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2008.

[27] Xiaorui Wang, Ming Chen, Charles Lefugy, and Tom Keller, "Ship: Scalable hierarchical power for large-scale data centers," *18th Int'l Conference on Parallel Architecture and Compilation Techniques (PACT)*, pp. 91–100, September 2009.

[28] Prith Banerjee, Chandrakant Patel, Cullen Bash, and Parthasarathy Ranganathan, "Sustainable data centers: Enabled by supply and demand side management," *46th Design Automation Conference (DAC)*, pp. 884–887, June 2009.

[29] Navin Sharma, Sean Barker, David Irwin, and Prashant Shenoy, "Blink: Managing server clusters on intermittent power," *16th Int'l Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2011.

[30] Chao Li, Wangyuan Zhang, Chang-Burm Cho, and Tao Li, "SolarCore: Solar energy driven multi-core architecture power management," *17th Int'l Symposium on High-Performance Computer Architecture (HPCA)*, February 2011.

[31] Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs, "Cutting the electric bill for internet-scale systems," *Special Interest Group on Data Communications (SIGCOMM)*, August 2009.

[32] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven Low, and Lachlan Andrew, "Greening geographical load balancing," *Int'l Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, June 2011.

[33] Kien Le, Ricardo Bianchini, Thu D. Nguyen, Jingru Zhang, Jiandong Meng, and Yogesh Jaluria, "Reducing electricity cost through virtual machine placement in high performance computing clouds," Tech. Rep., Rutgers University, November 2010.

[34] Cosmin Rusu, Rami Melhem, and Daniel Mosse, *The Interplay of Reward and Energy in Real-Time Systems*, Ph.D. thesis, 2006.

[35] Khargharia Bithika, Salim Hariri, and Marzin Yousif, "Autonomic power and performance management for computing systems," *Cluster Computing*, pp. 167–181, 2008.

[36] John Karidis, Jose Moreira, and Jaime Moreno, "True value: Assessing and optimizing the cost of computing at the data center level," *Int'l Conf on Computing Frontiers*, pp. 185–192, May 2009.

[37] Navin Sharma, Jeremy Gummeson, David Irwin, and Prashant Shenoy, "Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems," *7th Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, June 2010.