

Design and Operational Analysis of a Green Data Center

Prateek Sharma, Patrick Pegus II, David Irwin, Prashant Shenoy, John Goodhue, James Culbert

Abstract—Data centers are at the heart of the IT-driven economy. Power consumption of a single data center can range from tens to a hundred megawatts, and operational costs can run into millions of dollars a month. Data center operators incorporate careful design and optimizations to reduce energy consumption of large scale data centers. Since data center design and operations are a source of competitive advantage, insight into modern data centers has been scarce. This article describes the design and analysis of a state-of-the-art green university data center, and provides several insights into its operational and efficiency characteristics.

1 INTRODUCTION

The vast computing needs of today’s world are met by data centers, which house thousands of computing, storage, and networking devices. Data centers are crucial components of IT infrastructure, and used for running server-side workloads, data storage, and large-scale data and scientific processing.

Computing equipment is energy-hungry, and today’s large data centers can consume tens or even a hundred megawatts of electricity—which is roughly equal to the power consumption of 10,000 US households. Not all the energy consumed goes into powering the IT equipment. As much as 50% of the energy is consumed for cooling IT equipment, power transmission, and other overhead. The efficiency metric used by data center operators is Power Usage Efficiency (PUE), which is defined as:

$$\text{PUE} = \frac{\text{Total power consumption}}{\text{IT power consumption}} \quad (1)$$

Given these large energy demands, improving data center efficiency is an important problem in both academic research [3], [4], [10] and industry, as it can yield tremendous cost savings. Modern data centers use a plethora of design and optimization techniques to achieve PUE values as low as 1.1. For example, modern data centers often rely on “free cooling”—using outside cold air to keep the data center cool. While industry groups and companies have published average PUE values for their data centers [1], [6], detailed energy measurements and other operational insights are not widely available, likely because they are considered to be proprietary competitive advantage.

In this article, we will describe system design and analysis of power, water, and carbon usage of the Massachusetts Green High Performance Computing Center (MGHPCC), a 90,000 square foot, 15 MW data center that uses recent advances in cooling and power distribution to improve energy efficiency. As part of its research mission, the MGHPCC gathers and analyzes fine-grained resource consumption information, and allows insight into the operation of modern, highly efficient data centers.

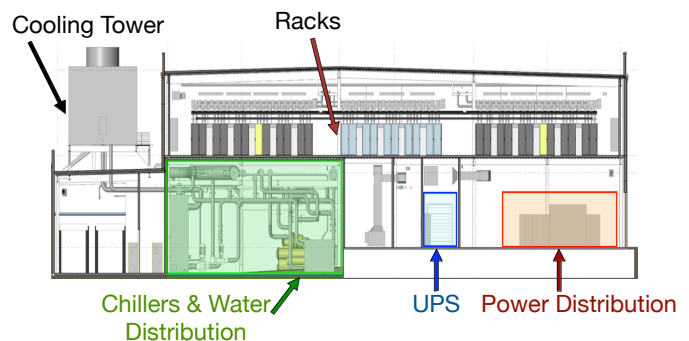


Fig. 1. Layout of the MGHPCC data center.

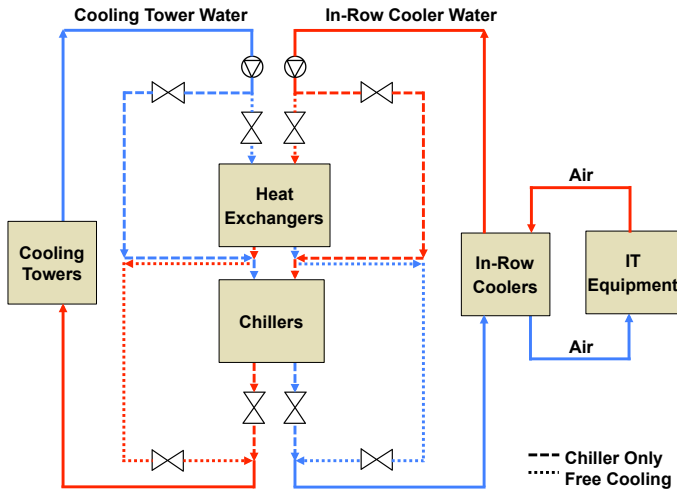
2 DATA CENTER DESIGN

MGHPCC uses favorable geographical location and green design to achieve extremely energy efficient operation. MGHPCC is located in Holyoke in western Massachusetts at the location of a former industrial mill site. Due to its location, MGHPCC enjoys access to cheap and abundant electricity, inexpensive real estate, and proximity to fiber-optic network backbones. Western Massachusetts has a cool climate, with mean summer and winter temperatures of 23°C and −3°C respectively. The cool climate enables the facility to employ free cooling, as explained later. Locating data centers in cold areas is increasingly popular to reduce the cooling energy consumption—Facebook for example recently revealed a 120 MW data center close to the arctic circle in Sweden.


MGHPCC is jointly operated by a university consortium and is a *multi-tenant* facility with floor-space allocated to each university, which is used to co-locate compute clusters. Use of MGHPCC has grown steadily since it opened. As of September 2016, 67% of available space and 25% of available power and cooling are in use. To accommodate more tenants and clusters, the data center has the ability to increase space, power, and cooling by 60%.

2.1 Physical Layout

The MGHPCC facility is a two-story building, with the lower floor containing the power and cooling infrastructure, and the upper floor containing the racks for hosting the computing infrastructure (Figure 1). Evaporative cooling towers are housed on the roof of the upper floor, adjacent to the computer floor. The data center also includes backup diesel generators that provide power for parts of the facility in case of a utility outage. A flywheel-based uninterruptible power supply (UPS) provides power until the generators can come



(a) MGHPCC employs free cooling and chillers for its chilled water. Server racks are air-cooled.



(b) Hot aisle containment. In row coolers (IRCs) extract heat from the racks.

Fig. 2. System design of MGHPCC's cooling infrastructure.

online after an outage.

Each tenant's share of the facility is divided into pods, and each pod contains a set of racks with the necessary power, cooling, and communication support. Each pod contains 20 or 24 racks.

2.2 Power Infrastructure

The power infrastructure for the data center resembles a small-scale distribution network in the electric grid. The infrastructure comprises substations, feeders, transformers, and switchboards that feed power to the computing and cooling infrastructure. Electricity enters the facility at 13.8kV where it is distributed from the main switchboard, transformed to 230V, and is then delivered through power panels in the computer room to the bus plugs that feed the power distribution units (PDUs) in each rack.

Since power conversion losses can be a key source of higher PUEs in data centers, the data center uses a number of techniques to reduce such losses. MGHPCC uses high voltage and low current to deliver power, which reduces transformer losses. Higher distribution voltages also make it possible to eliminate an entire tier of transformers from the distribution network, further reducing transformer losses.

The facility's UPS system stores energy kinetically in spinning flywheels, which is more environmentally-friendly than storing energy chemically in batteries that often contain harmful chemicals, such as lead in lead-acid batteries. This approach affects energy

efficiency, as power is required to keep the flywheels spinning in standby mode. To reduce this load, the data center backs up only 20% of the available compute load. Since most research computing applications can tolerate occasional outages, and the local power infrastructure has been historically very reliable, this represented good tradeoff between energy efficiency, green design, and availability. Thus, only 20% of each tenant's racks are UPS-backed. The tenants are able to choose how to partition their cluster between UPS and non-UPS racks.

2.3 Cooling Infrastructure

Traditionally, data centers have used chillers for cooling. Since chillers consume a significant amount of energy, they are a key contributor to high PUE. Data centers are therefore adopting alternative technologies to lower their PUEs, and continue to explore innovative options [5]. MGHPCC leverages "free cooling" (also known as "renewable cooling") to reduce the amount of energy used by its chillers.

As shown in Figure 2a, there are two main water cooling loops. The water in the cooling tower loop is cooled using evaporative cooling. The chilled water loop circulates chilled water through IRCs that cool the computer room air as they remove it from the hot aisle shown in Figure 2a. If the outside air temperature and humidity are low enough, then it is used by the heat-exchangers to cool the water in the chilled water loop.

To maximize the amount of time in free-cooling mode, the facility maintains the computer room temperature at 26.7°C, which is compatible with modern servers but higher than the temperature settings for traditional data centers. Combined with the cooler climate of Massachusetts, this permits the use of free cooling for more than 70% of the year.

Each tenant's rack is configured to use hot aisle containment to prevent hot and cold air from mixing together (Figure 2b). The cold water in the chilled water loop is circulated through in-row coolers (IRCs), which are deployed adjacent to racks to cool the hot air extracted from the servers. The use of in-row coolers allows for a close coupling of the cooling with the computing heat load—the controls of the in-row coolers actively adjust fan speeds and chilled water flow to closely match the computing heat load on nearby racks, thereby enhancing efficiency.

2.4 Data Center Monitoring

The data center has several thousand instrumentation points to monitor power, cooling, and water usage within the facility. Facility data is available in real time to both facility and computer system operators, though server-level data is often restricted to just the computer system operators.

We use several types of monitoring devices to track resource usage. At the facility level, the power distribution infrastructure is monitored by over 900 networked electric meters that monitor energy usage at different levels of the distribution network at 20 second granularity. These meters monitor the average power usage of individual racks, as well as the aggregate usage at higher levels of the power distribution hierarchy. There are also separate meters to monitor the power usage of the cooling infrastructure, including its associated pumps, chillers, and in-row coolers.

The facility's mechanical systems, which are primarily associated with the cooling infrastructure, are also monitored. The available data includes water pump flow levels at various points in the water loops, as well as data from in-row coolers, such as their fan speed, water inlet and outlet temperature, and water flow

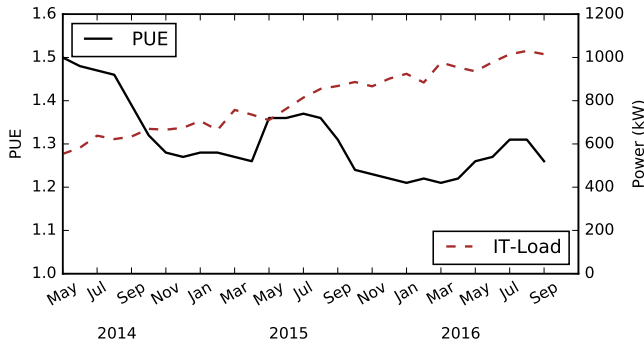


Fig. 3. IT load and PUE. IT load has been steadily increasing, while PUE sees a general downward trend and sees an increase during summer months.

data. This data is generally recorded and available at a one-minute granularity. The temperature and humidity of the computer room floor is extensively monitored using sensors that are deployed on the hot and cold sides of each rack.

3 POWER USAGE ANALYSIS

In this section, we present an analysis of MGHPCCC's power usage. The extensive monitoring infrastructure described earlier helps us collect and analyze facility-level power data as well as fine-grained power usage at a component/sub-system level. Direct access to such data is important for developing models and energy optimizations for multi-tenant data centers, but is often restricted to facility managers only.

3.1 IT Load Analysis

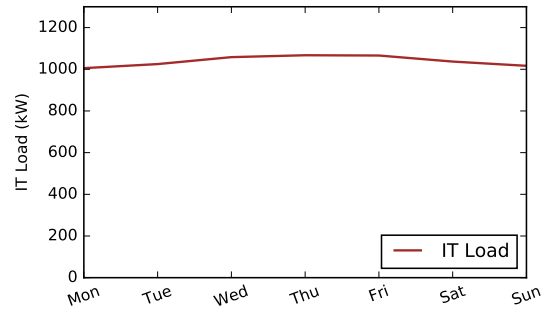
The primary role of a data center is to run computing, networking, and storage devices (collectively called IT equipment). We measure the IT load by aggregating the power utilization of each pod (group of racks). The power consumed by all the IT equipment is indicative of the load on the data center, and is shown in Figure 3. The mean monthly IT load has steadily increased over the last two years, as tenants have commissioned more racks. The MGHPCCC IT load is slightly above 1 MW, which is 10% of the provisioned IT power.

A finer-grained temporal analysis of the IT load shows that there are no prominent day-of-week or time-of-day effects (Figure 4). This is because MGHPCCC's research workloads are composed primarily of batch jobs, which can be organized into a steady, high-utilization workload. During MGHPCCC's initial operational phase, the job-queue was mostly non-empty, which meant that jobs were serviced during both day and night.

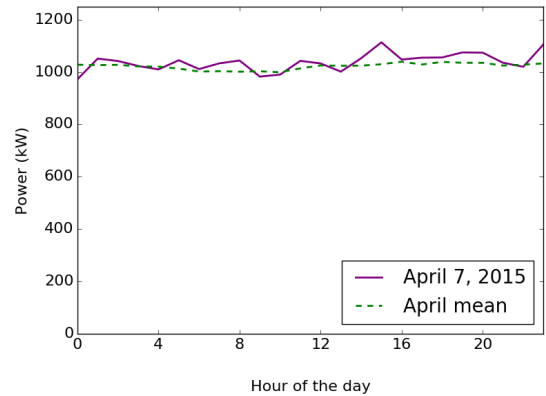
3.2 PUE Analysis

As mentioned earlier, a significant fraction of the total power in a data center is consumed for non-IT tasks such as cooling the IT equipment and power distribution losses. The metric for power efficiency of data centers is PUE, which is defined as $\frac{\text{Total power}}{\text{IT power}}$. To compute the PUE, we measure the total power directly from the networked meter that measures power entering the facility from the grid.

Figure 3 depicts the monthly PUE of the data center over the course of two years. The average PUE for the year ending in September 2015 was 1.37, and the average for the year ending in September 2016 was 1.29. The minimum PUE observed is 1.21. We see three trends. First, the PUE increases during summer



(a) IT load doesn't show prominent day of week effects.



(b) IT load doesn't show prominent time of a day effects.

Fig. 4. IT load over various time scales.

months when chillers must be used. Second, the yearly decrease in PUE is a result of increasing IT load. The data center's cooling infrastructure is sized for a much greater load than the present server room occupancy and is not energy proportional. With an increasing IT load, the cooling infrastructure is better utilized, resulting in better (lower) PUE values. Third, the decline in PUE from May to September 2014 is a result of a change in the building management system software which eliminated unnecessary air mover operation, and increased the amount of time spent in free cooling mode.

An average PUE of 1.29 is significantly lower than the average PUE of 1.7 that is common in enterprise data centers in the industry [9]. However, the value is not as low as the PUEs near 1.1 reported by the newest and most efficient data centers built by large internet companies like Facebook and Google [1], [7].

MGHPCC's fine-grained instrumentation allows us to identify the power consumption of each component. We show a component level breakdown in Figure 5. We see that the cooling equipment: the chillers, pumps, and in-row coolers, all consume almost 60% of the non-IT power in the summer months. Since MGHPCCC uses free cooling when possible, the cooling component reduces in the winter months. For example, the use of chillers is almost completely absent from October through March, resulting in lower non-IT power usage and low PUE values (close to 1.21).

Two other major factors contributing to non-IT power are power distribution losses that occur when the incoming power flows through the transformers and wiring in the facility, and the power consumed by the flywheels and control logic in the UPS units. As noted earlier, the facility reduces the amount of energy needed to run the UPS flywheels and control logic by covering only

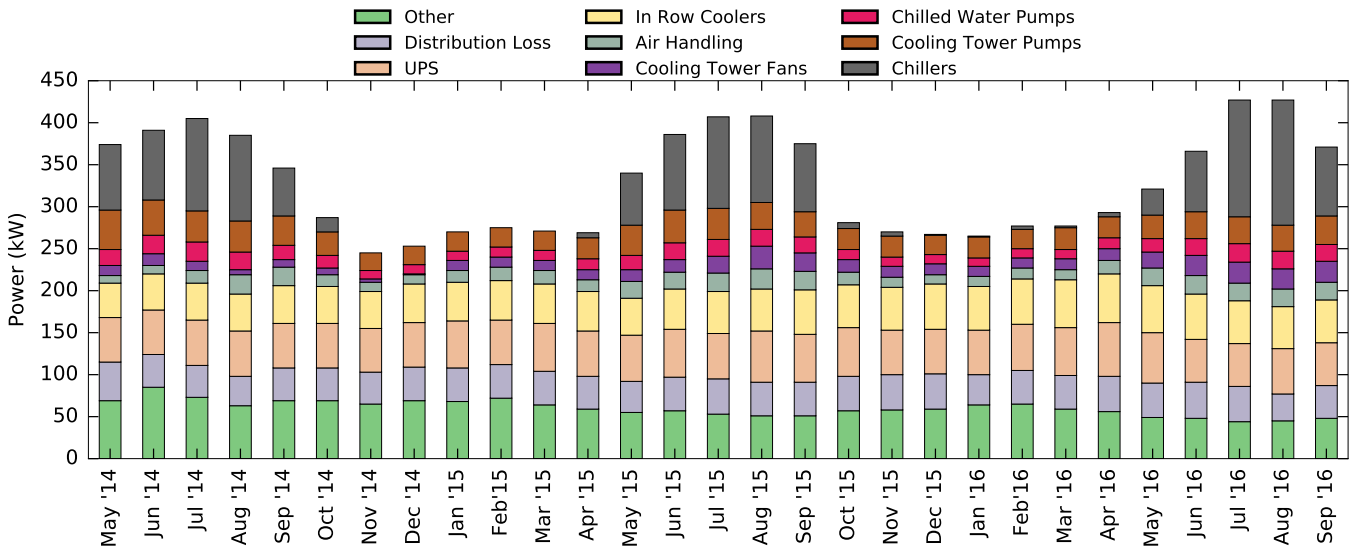


Fig. 5. Breakdown of energy consumption. Notice the absence of chiller usage in the colder months.

20% of the facility with UPS-backed power. Power distribution losses will increase at a slower rate with increasing compute load. UPS power consumption is constant. Both will therefore consume a decreasing fraction of overall energy as compute load increases.

Lastly, about 15% of the non-IT power is consumed by ancillary items such as lighting, utility outlets, generator fuel heaters, loading dock overhead door motors, elevators, and compressed air pumps for the sprinkler system. While there is some seasonal variation, annual power consumption for this category is fixed, and its fraction of the total facility load will decrease as compute load increases. A more detailed power analysis may be found in [8].

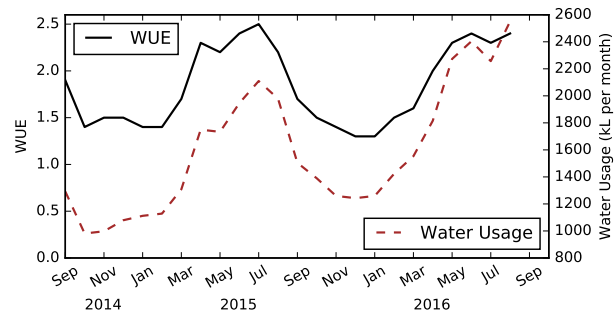
4 WATER USAGE ANALYSIS

In addition to consuming significant amounts of power, data centers also typically consume significant amounts of water, mainly as part of their cooling infrastructure (Figure 2a). While there has been significant emphasis on measuring and optimizing the power usage using metrics such as PUE, there has been less attention on measuring the efficiency of water usage. Recently a new metric to capture the effectiveness of water usage has been proposed. The WUE of a data center is defined as liters of water used per kilowatt-hour of energy used by IT equipment and can be expressed as:

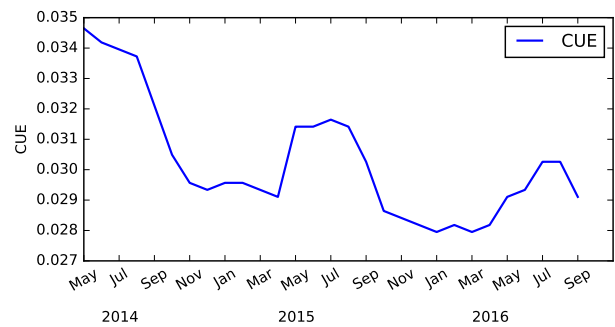
$$WUE = \frac{\text{Water Usage (liters)}}{\text{IT Energy Usage (kWh)}} \tag{2}$$

Figure 6a depicts the monthly water usage of MGHPC. The present water usage varies between 1000 kL and 2500 kL per month depending on the season of the year. Water usage is higher during the warmer months when the chilled loop must support heat loads from the chiller that dehumidifies air throughout the building, and the heat exchanger that cools the air in the office areas. It is lower during the winter months when the extra heat loads are not present and the rate of evaporation from the cooling tower decreases due heat transfer from the warmer water droplets in the cooling tower to the cooler outside air.

Beyond evaporation, water usage includes windage, blowdown, and water filtration backwash, which currently consume a fixed



(a) Water usage is low and shows seasonal trends.



(b) Carbon Usage Efficiency

Fig. 6. Water and Carbon usage.

amount of water per month. As compute load increases, the fixed consumers become a smaller percentage of total water use, driving the rate of consumption per kilowatt-hour down. When WUE reaches approximately 1.4 Liters/kWh, improvement will slow as the blowdown rate transitions to a mode where it is proportional to the evaporation rate instead of its current fixed rate.

MGHPCC’s WUE varies between 1.3 L/kWh and 2.5 L/kWh over the course of the year, and shows similar seasonal trends as the water usage. There is little, if any, WUE data available on data centers with cooling tower systems. Facebook has released

Fuel Type	Energy (MWh)	Energy (%)	CO ₂ (kg)	CO ₂ (%)
Oil	1,724	0.4	1,476,897	16.3
Hydro	261,691	66.7	0	0
Nuclear	61,310	15.6	0	0
Solar	6,105	1.6	0	0
Contracted (carbon free)	40,800	10.4	0	0
Contracted (other)	20,592	5.3	7,584,064	83.7
Total	392,222	100	9,060,054	100

TABLE 1

Holyoke Gas & Electric power generation is dominated by hydro, nuclear, and other carbon-free sources.

data indicating a WUE of 0.28 L/kWh for its Prineville and 0.34 L/kWh Forest City data centers. The direct evaporative cooling and humidification (ECH) misting system used in these data centers delivers impressive results, but was not usable at the MGHPCC, where chilled water for water cooled computing systems was a design requirement.

5 CARBON FOOTPRINT ANALYSIS

Our final analysis focuses on the carbon footprint of the MGHPCC, since ultimately it is designed to be a green facility. While there are many methodologies to compute the operational carbon footprint of a building, the new CUE metric has been defined explicitly to compute the carbon effectiveness of data centers [2]. CUE is defined as:

$$\begin{aligned} \text{CUE} &= \frac{\text{CO}_2 \text{ emissions from the total data center energy}}{\text{IT equipment energy}} \quad (3) \\ &= \frac{\text{kg CO}_2}{\text{kWh}} \cdot \frac{\text{Total data center energy}}{\text{IT equipment energy}} = \frac{\text{kg CO}_2}{\text{kWh}} \cdot \text{PUE} \end{aligned}$$

CUE depends significantly on the carbon emissions due to electricity consumption. The carbon emissions of the electricity consumption, in turn, depend on the electric utility's generation sources. In the event that the data center uses on-site or contracted renewable energy, that portion must also be considered in the overall electricity mix.

The MGHPCC relies on a local utility company, Holyoke Gas and Electric (HG&E). HG&E generates a large fraction of electricity using hydroelectric power, which is an inexpensive and clean source of renewable energy. The mix of generation sources used by HG&E in 2014 is shown in Table 1. 94.3% of the electricity is generated from carbon-free sources : hydro-electric, nuclear, and solar.

The high fraction (94.3%) of carbon-free electricity yields a low ratio (of 0.0231) of kilograms of CO₂ emitted per kWh of energy generated. This ratio is nearly an order of magnitude lower than the most carbon efficient region in the U.S (0.2 kg CO₂ per kWh).

MGHPCC's CUE (shown in Figure 6b) varies from 0.028 to 0.03. By way of comparison, an "average" data center that draws power from the "average" utility mix in the U.S. will have 25× higher CUE at the same PUE level (and an even higher CUE at higher typical values of 1.7 PUE).

6 CONCLUDING REMARKS

We presented the design and empirical analysis of the efficiency of a green academic data center—the MGHPCC. The power usage effectiveness analysis of MGHPCC reveals that it has a minimum PUE of 1.21 and an average of 1.29. MGHPCC's PUE is significantly lower than the current industry average of 1.7, but

not as low as the PUE of 1.1 of the newest and most efficient large scale data centers. As noted earlier, this may be due to the use of direct evaporative cooling and other techniques that deliver better performance than free cooling. Our data analysis shows significant seasonal variation in energy efficiency, but does not show significant diurnal variation.

The fine-grained monitoring infrastructure deployed at MGHPCC opens up the possibilities of further analysis, for example at a server and rack level. This data can be used by researchers to develop data center models to tune system parameters and minimize resource consumption. Aggregate monthly data can be found at <http://traces.cs.umass.edu/index.php/Smart/Smart>.

REFERENCES

- [1] Google data center efficiency. <https://www.google.com/about/datacenters/efficiency/internal/>, August 2016.
- [2] C. Belady, D. Azevedo, M. Patterson, J. Pouchet, and R. Topley. Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric. *White Paper*, 32, 2010.
- [3] N. El-Sayed, I. A. Stefanovici, G. Amvrosiadis, A. A. Hwang, and B. Schroeder. Temperature management in data centers: Why some (might) like it hot. *SIGMETRICS*, 2012.
- [4] I. Goiri, T. D. Nguyen, and R. Bianchini. CoolAir: Temperature- and Variation-Aware Management for Free-Cooled Datacenters. In *ASPLOS*, March 2015.
- [5] J. Hamilton. Data center cooling done differently. <http://perspectives.mvdirona.com/2014/08/data-center-cooling-done-differently/>, October 2014.
- [6] J. Hamilton. Data center power consumption. <http://perspectives.mvdirona.com/2015/06/data-center-power-water-consumption/>, June 2015.
- [7] J. Parr. Designing a Very Efficient Data Center. <https://www.facebook.com/notes/facebook-engineering/designing-a-very-efficient-data-center/10150148003778920>, April 2011.
- [8] P. Pegus II, B. Varghese, T. Guo, D. Irwin, P. Shenoy, A. Mahanti, J. Culbert, J. Goodhue, and C. Hill. Analyzing the efficiency of a green university data center. In *International Conference on Performance Engineering*. ACM, 2016.
- [9] Y. Sverdlik. Survey: Industry Average Data Center PUE Stays Nearly Flat Over Four Years. In *Data Center Knowledge*, June 2nd 2014.
- [10] L. Wang, S. U. Khan, and J. Dayal. Thermal aware workload placement with task-temperature profiles in a data center. *The Journal of Supercomputing*, 2012.

Prateek Sharma is a PhD student in the College of Information and Computer Sciences at the University of Massachusetts Amherst. His current research focuses on cloud computing. Contact him at prateeks@cs.umass.edu.

Patrick Pegus II is an MS student in the College of Information and Computer Sciences at the University of Massachusetts Amherst. His current research focuses on green data centers. Contact him at pegusii@cs.umass.edu.

David Irwin is an Assistant Professor in the Department of Electrical and Computer Engineering at the University of Massachusetts Amherst. His research interests are broadly in experimental computing systems with a particular emphasis on sustainability. Contact him at irwin@ecs.umass.edu.

Prashant Shenoy is a Professor of Computer Science at the University of Massachusetts. His current research focuses on cloud computing and green computing. He is a distinguished member of the ACM and a Fellow of the IEEE. Contact him at shenoy@cs.umass.edu.

John Goodhue is the Executive Director of MGHPCC. He holds a B.S. in Computer Engineering from the Massachusetts Institute of Technology. Contact him at jtgoodhue@mghpcc.org.

James Culbert is the Director of IT Services at MGHPCC. He holds an M.S. in Electrical Engineering from the Massachusetts Institute of Technology. Contact him at culbertj@mghpcc.org.