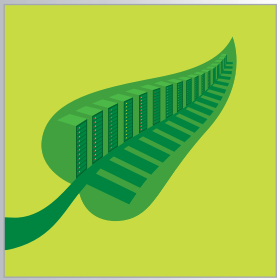

Design and Operational Analysis of a Green Data Center

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 large-scale data centers' energy consumption. Because data center design and operations are a source of competitive advantage, insight into modern data centers is 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.

Prateek Sharma, Patrick Pegus II, David Irwin, and Prashant Shenoy
University of Massachusetts Amherst

John Goodhue and James Culbert
Massachusetts Green High-Performance Computing Center

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 percent 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 effectiveness (PUE), which is defined as

$$\text{PUE} = \frac{\text{Total power consumption}}{\text{IT power consumption}}.$$

Given these large energy demands, improving data center efficiency is an important problem in both academic research[1–3] 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

 1089-7801/17/$33.00

PUE values for their data centers,[4,5] detailed energy measurements and other operational insights aren't widely available, probably because they're considered a proprietary competitive advantage.

In this article, we describe the system design and analyze the power, water, and carbon usage of the Massachusetts Green High-Performance Computing Center (MGHPCC), a 90,000-square-foot, 15-megawatt (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.

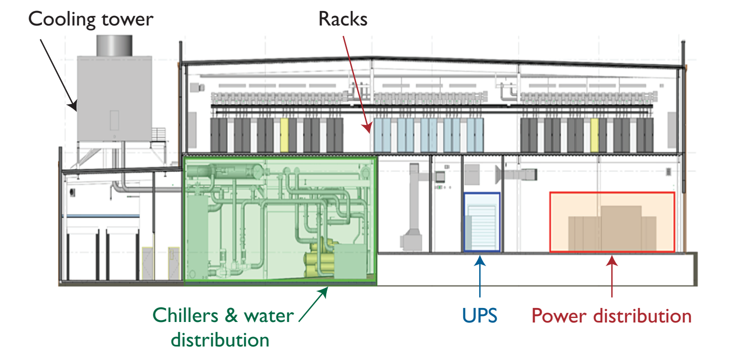

Figure 1. Layout of the Massachusetts Green High-Performance Computing Center (MGHPCC) data center. UPS stands for uninterruptible power supply.

Data Center Design

MGHPCC uses its 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.

Because of 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 *multitenant* 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 percent of available space and 25 percent 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 percent.

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 (see 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 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.

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 13.8 kilovolts (kV), where it's distributed from the main switchboard, transformed to 230 volts (V), and then delivered through power panels in the computer room to the bus plugs that feed the power distribution units (PDUs) in each rack.

Because 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 percent of the available compute load. Because most research computing applications can tolerate occasional outages, and historically the local power infrastructure has been quite reliable, this represented a good tradeoff between energy efficiency, green design, and availability. Thus, only 20 percent 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.

Cooling Infrastructure

Traditionally, data centers have used chillers for cooling. Because chillers consume a significant amount of energy, they're a key contributor to high PUE. Data centers therefore are adopting alternative technologies to lower their PUEs, and continue to explore innovative options.[6] As we mentioned, MGHPCC leverages free cooling (also known as *renewable cooling*) to reduce the amount of energy used by its chillers.

As Figure 2a shows, 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 in-row coolers (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's 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 percent of the year.

Each tenant's rack is configured to use hot aisle containment to prevent hot and cold air from mixing together (see Figure 2b). The cold water in the chilled water loop is circulated through 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 IRCs actively adjust fan speeds and chilled water flow to closely match the computing heat load on nearby racks, thereby enhancing efficiency.

Data Center Monitoring

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

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

Power Usage Analysis

Next, we present an analysis of MGHPCC'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/subsystem level. Direct access to such data is important for developing models and energy optimizations for multitenant data centers, but is often restricted to facility managers only.

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

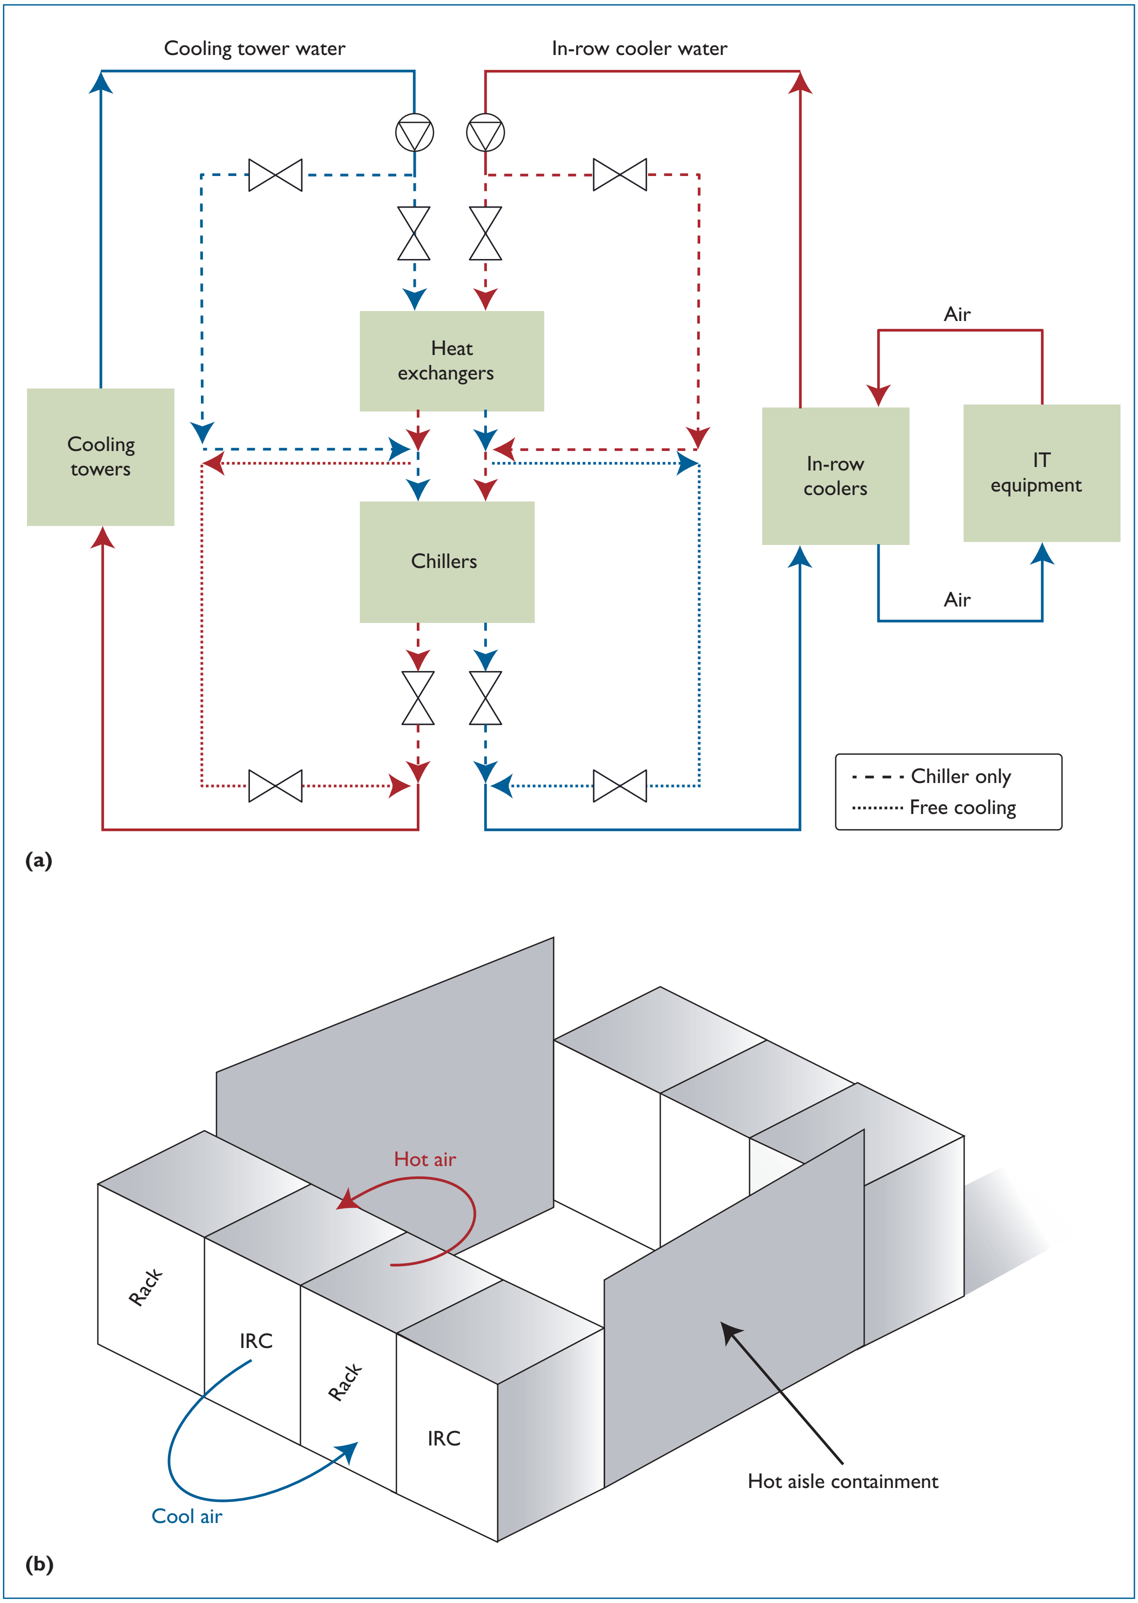

Figure 2. System design of MGHPCC's cooling infrastructure. (a) MGHPCC employs free cooling and chillers for its chilled water. Server racks are air-cooled. (b) Hot aisle containment. In-row coolers (IRCs) extract heat from the racks.

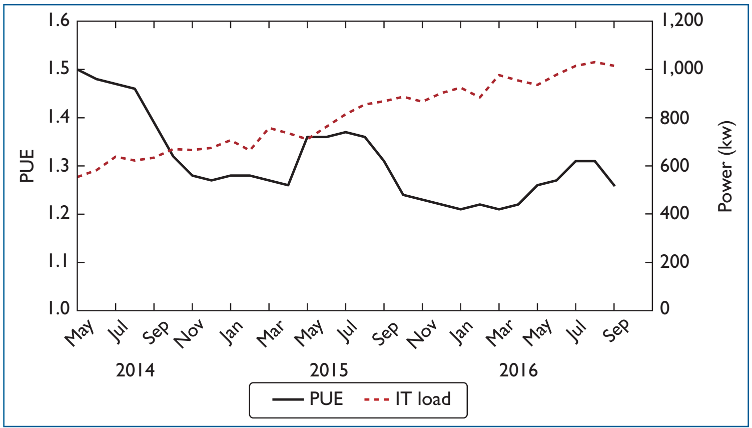

Figure 3. IT load (based on the power consumed by all the IT equipment, given in kilowatts) and power usage effectiveness (PUE). The IT load has been steadily increasing, while PUE sees a general downward trend and sees an increase during summer months.

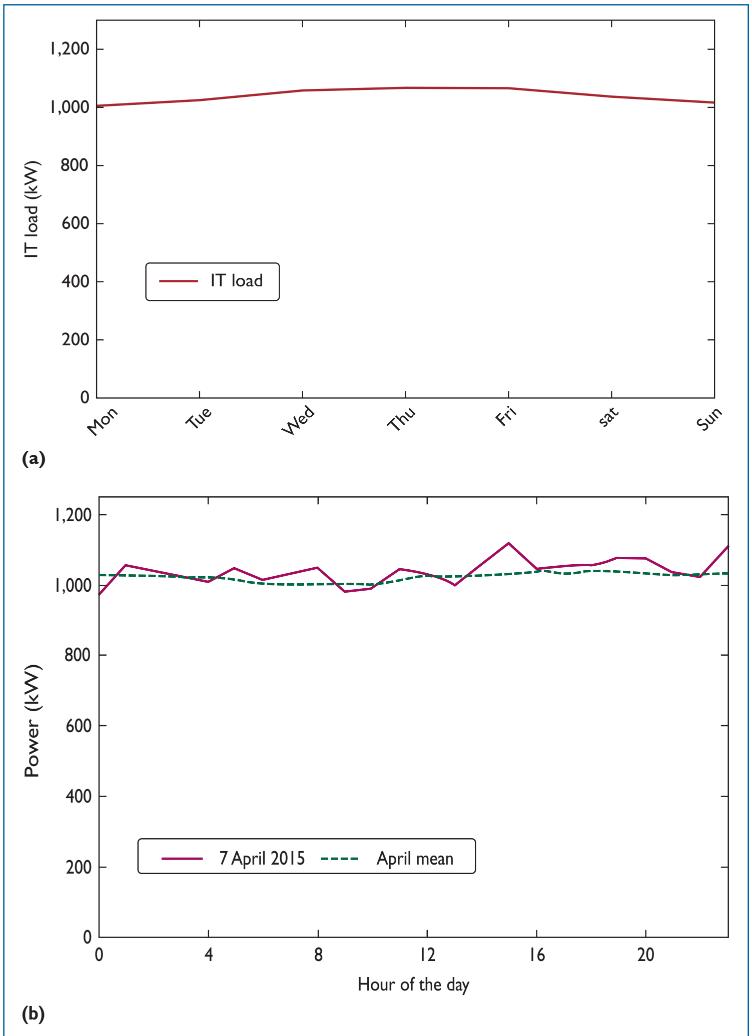

Figure 4. IT load over various time scales. (a) The IT load doesn't show prominent day of week effects. (b) IT load doesn't show prominent time of a day effects.

IT load by aggregating the power usage of each pod (group of racks). The power consumed by all the IT equipment is indicative of the data center's load (see Figure 3). The mean monthly IT load has steadily increased over the last two years, as tenants have commissioned more racks. The MGHPCC IT load is slightly above 1 MW, which is 10 percent 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 (see Figure 4). This is because MGHPCC's research workloads are composed primarily of batch jobs, which can be organized into a steady, high-utilization workload. During MGHPCC's initial operational phase, the job queue was mostly non-empty, which meant that jobs were serviced during both day and night.

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 (Total power)/(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 months when chillers must be used. Second, the yearly decrease in PUE is a result of the increasing IT load. The data center's cooling infrastructure is sized for a much greater load than the present server room occupancy and isn't 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 the result of a change in the building management system software, which eliminated unnecessary air flow 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's common in enterprise data centers in the industry.[7]

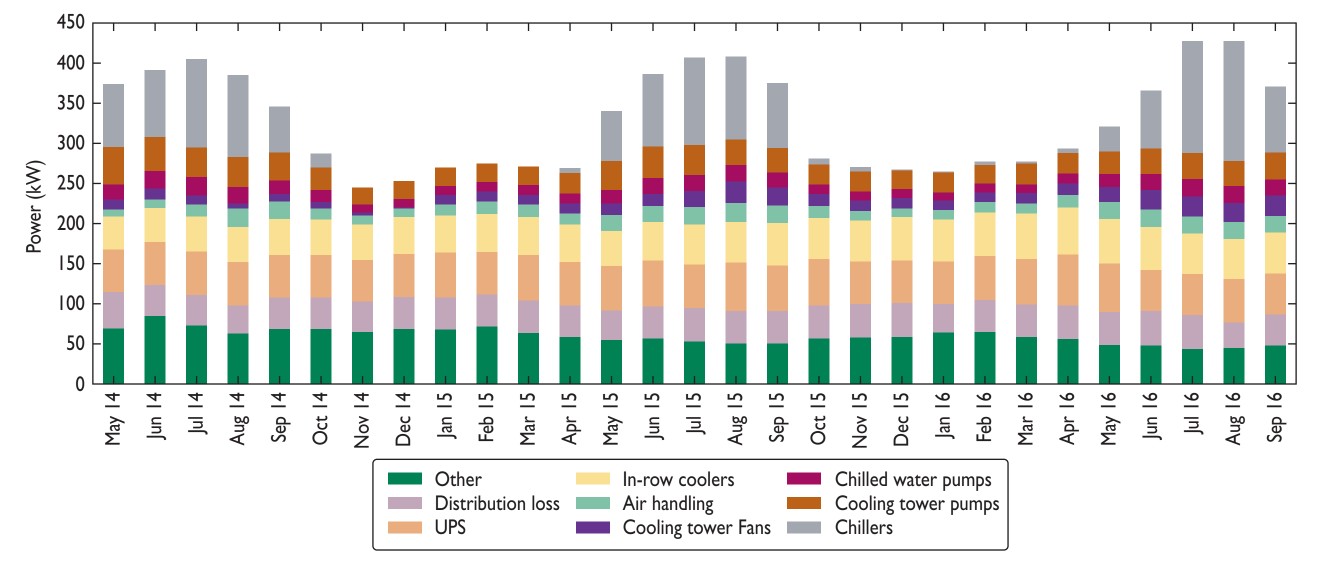

Figure 5. Breakdown of energy consumption. Notice the absence of chiller usage in the colder months.

However, the value isn't as low as the PUEs near 1.1 reported by the newest and most efficient data centers built by large Internet companies such as Facebook and Google.[4,8]

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 percent of the non-IT power in the summer months. Because MGHPCC 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 20 percent of the facility with UPS-backed power. Power distribution losses will increase at a slower rate with an increasing compute load. UPS power consumption is constant. Both will therefore consume a decreasing fraction of overall energy as the compute load increases.

Last, about 15 percent 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's some seasonal variation, annual power consumption for this category is fixed, and its fraction of the total facility load will decrease as the compute load increases. A more detailed power analysis is found elsewhere.[9]

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 (see Figure 2a). While there has been significant emphasis on measuring and optimizing the power usage through 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 water usage effectiveness (WUE) of a data center is defined as liters of water used per kilowatt-hour of energy expended by IT equipment. This can be expressed as

$$\text{WUE} = \frac{\text{Water usage (liters)}}{\text{IT energy usage (kWh)}}.$$

Figure 6a depicts the monthly water usage of MGHPCC. The present water usage varies between 1,000 kL and 2,500 kL per month, depending on the season of the year. Water usage

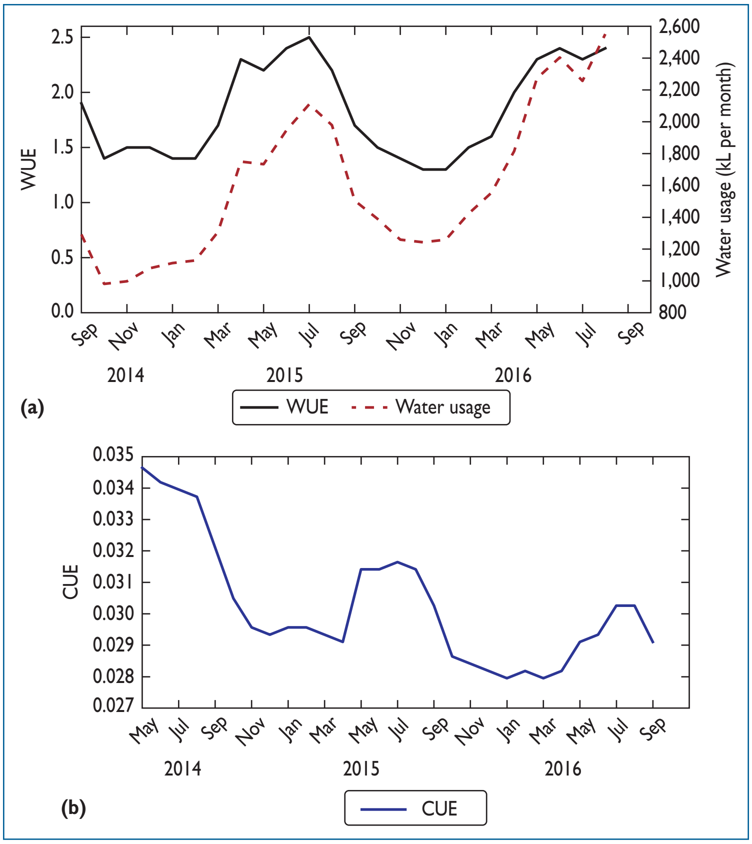

Figure 6. Water and carbon usage. (a) Water usage effectiveness (WUE) is low and shows seasonal trends. (b) Carbon usage effectiveness (CUE).

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's lower during the winter months, when the extra heat loads aren't present and the rate of evaporation from the cooling tower decreases because of 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 amount of water per month. As the 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 per kilowatt hour (L/kWh), improvement will slow as the blowdown rate transitions to a mode where it's 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's little, if any, WUE data available on data centers with cooling tower systems. Facebook has released 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 wasn't usable at MGHPCC, where chilled water for water-cooled computing systems was a design requirement.

Carbon Footprint Analysis

Our final analysis focuses on the carbon footprint of the MGHPCC, because ultimately it's designed to be a green facility. While there are many methodologies to compute the operational carbon footprint of a building, the new carbon usage effectiveness (CUE) metric has been defined explicitly to compute the carbon effectiveness of data centers.[10] CUE is defined as

$$\text{CUE} = \frac{\text{CO}_2 \text{ emissions from the total data center energy}}{\text{IT equipment energy}}$$

$$= \frac{\text{kg CO}_2}{\text{kWh}} \cdot \frac{\text{Total data center energy}}{\text{IT equipment energy}}$$

$$= \frac{\text{kg CO}_2}{\text{kWh}} \cdot \text{PUE}.$$

CUE depends significantly on the carbon emissions caused by 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 onsite or contracted renewable energy, that portion must also be considered in the overall electricity mix.

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. Table 1 shows the mix of generation sources used by HG&E in 2014. Out of the various fuels used, 94.3 percent of the electricity is generated from carbon-free sources: hydroelectric, nuclear, and solar.

The high fraction (94.3 percent) of carbon-free electricity yields a low ratio (of 0.0231) of kilograms of CO_2 emitted per kWh of energy

Table 1. Holyoke Gas & Electric power generation is dominated by hydroelectric, nuclear, and other carbon-free sources.

Fuel type	Energy (MWh)*	Energy (%)	CO_2 (kg)	CO_2 (%)
Oil	1,724	0.4	1,476,897	16.3
Hydroelectric	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

* MWh stands for megawatt hours.

generated. This ratio is nearly an order of magnitude lower than the most carbon-efficient region in the US (0.2 CO_2 per kWh).

MGHPCC's CUE (see 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 US will have 25× higher CUE at the same PUE level (and an even higher CUE at higher typical values of 1.7 PUE).

We presented the design and empirical analysis of the efficiency of a green academic data center — the MGHPCC. Our 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 might 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 doesn't 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. N. El-Sayed et al., "Temperature Management in Data Centers: Why Some (Might) Like It Hot," *Proc. ACM Sigmetrics*, 2012, pp. 163–174.
2. I. Goiri, T.D. Nguyen, and R. Bianchini, "CoolAir: Temperature- and Variation-Aware Management for Free-Cooled Datacenters," *Proc. Architectural Support for Programming Languages and Operating Systems*, 2015, pp. 253–265.
3. L. Wang, S.U. Khan, and J. Dayal, "Thermal Aware Workload Placement with Task-Temperature Profiles in a Data Center," *The J. Supercomputing*, 2012; doi:10.1007/s11227-011-0635-z.
4. Google, *Efficiency: How We Do It*, tech. report, Aug. 2016; www.google.com/about/datacenters/efficiency/internal.
5. J. Hamilton, "Data Center Power Consumption," *Perspectives*, June 2015; http://perspectives.mvdirona.com/2015/06/data-center-power-water-consumption.
6. J. Hamilton, "Data Center Cooling Done Differently," *Perspectives*, Oct. 2014; http://perspectives.mvdirona.com/2014/08/data-center-cooling-done-differently.
7. Y. Sverdlik, "Survey: Industry Average Data Center PUE Stays Nearly Flat over Four Years," *Data Center Knowledge*, 2 June 2014; www.datacenterknowledge.com/archives/2014/06/02/survey-industry-average-data-center-pue-stays-nearly-flat-four-years.
8. J. Parr, *Designing a Very Efficient Data Center*, tech. report, Facebook, 14 Apr. 2011; www.facebook.com/notes/facebook-engineering/designing-a-very-efficient-data-center/10150148003778920.
9. P. Pegus II et al., "Analyzing the Efficiency of a Green University Data Center," *Proc. Int'l Conf. Performance Eng.*, 2016, pp. 63–73.

10. C. Belady et al., *Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric*, white paper no. 32, 2010.

Prateek Sharma is a PhD student in the College of Information and Computer Sciences at the University of Massachusetts Amherst. His research interests focus on cloud computing. Sharma has an MS in computer science from IIT-Bombay. Contact him at prateeks@cs.umass.edu.

Patrick Pegus II is a master's 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 ppegusii@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 include experimental computing systems with a particular emphasis on sustainability. Irwin has a PhD in computer science from Duke University. Contact him at irwin@ecs.umass.edu.

Prashant Shenoy is a professor of computer science at the University of Massachusetts Amherst. His research interests include cloud computing and green computing. Shenoy has a PhD in computer science from the University of Texas at Austin. He's an IEEE Fellow and ACM distinguished member. Contact him at shenoy@cs.umass.edu.

John Goodhue is the executive director of Massachusetts Green High-Performance Computing Center (MGHPCC). Goodhue has a BS 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. Culbert has an MS in electrical engineering from the Massachusetts Institute of Technology. Contact him at culbertj@mghpcc.org.