

Analyzing the Efficiency of a Green University Data Center

Patrick Pegus II*, Benoy Varghese†, Tian Guo*, David Irwin*, Prashant Shenoy*
Anirban Mahanti†, James Culbert‡, John Goodhue‡, Chris Hill§

*University of Massachusetts Amherst, †NICTA, ‡MGHPCC, §Massachusetts Institute of Technology

ppegusii@cs.umass.edu, benoy.varghese@nicta.com.au, tian@cs.umass.edu, irwin@ecs.umass.edu, shenoy@cs.umass.edu
anirban.mahanti@nicta.com.au, culbertj@mghpcc.org, jtgoodhue@mghpcc.org, cnh@mit.edu

ABSTRACT

Data centers are an indispensable part of today’s IT infrastructure. To keep pace with modern computing needs, data centers continue to grow in scale and consume increasing amounts of power. While prior work on data centers has led to significant improvements in their energy-efficiency, detailed measurements from these facilities’ operations are not widely available, as data center design is often considered part of a company’s competitive advantage. However, such detailed measurements are critical to the research community in motivating and evaluating new energy-efficiency optimizations. In this paper, we present a detailed analysis of a state-of-the-art 15MW green multi-tenant data center that incorporates many of the technological advances used in commercial data centers. We analyze the data center’s computing load and its impact on power, water, and carbon usage using standard effectiveness metrics, including PUE, WUE, and CUE. Our results reveal the benefits of optimizations, such as free cooling, and provide insights into how the various effectiveness metrics change with the seasons and increasing capacity usage. More broadly, our PUE, WUE, and CUE analysis validate the green design of this LEED Platinum data center.

1. INTRODUCTION

Data centers form the backbone of our increasingly IT-driven economy, and are commonly used by enterprises to run their IT infrastructure. In recent years, the number and scale of data centers has grown rapidly. While small data centers may host a few thousand servers, the largest ones now host hundreds of thousands of servers. The energy consumed by these servers and their associated IT and network infrastructure is significant—globally, recent estimates attribute 2% of U.S. electricity consumption to data centers [12]. The largest data centers now consume over 100MW and incur monthly energy bills in the millions of dollars [8].

Thus, improving data center energy-efficiency has emerged as both an important academic research topic, as well as a pressing industry need. Over the past fifteen years, there has been much work on improving the energy-efficiency of the servers housed in data centers, e.g., [6, 3, 19, 4]. More recently, researchers have fo-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICPE’16, March 12–18, 2016, Delft, Netherlands.

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4080-9/16/03...\$15.00

DOI: <http://dx.doi.org/10.1145/2851553.2851557>

cused on optimizing the efficiency of data center cooling systems, e.g., by using free cooling from the outside air with air-side economizers [9, 13], since cooling servers consumes a significant fraction of data center energy. Collectively, these advances have led to a steady decrease in the Power Usage Effectiveness (PUE) metric commonly used to quantify data center energy-effectiveness.¹ Recent studies show that older enterprise data centers have PUEs of 1.7 or higher [21], while newer data centers that incorporate energy optimizations for servers and their cooling infrastructure have PUEs near 1.1 [10].

While industry groups and companies, including Google, Microsoft, Facebook, Amazon, and Apple, have published average PUE values across their data centers, detailed energy measurements are not widely available. For example, the published PUE values are typically averages across many data centers over a long period, e.g., the past year, and do not break PUE down spatially, temporally, or across subsystems. Facebook has taken strides to increase access to such data through its OpenCompute project, which includes both hardware and facility designs for its data centers, as well as a public dashboard showing their real-time PUE and Water Usage Effectiveness (WUE) [14] (although without a detailed breakdown). Detailed access to such operational data can enable important insights into data center operations that motivate new research directions. Unfortunately, detailed data on internal data center operations is typically kept confidential, since most companies view data center design as a competitive advantage. Thus, only the few researchers at each company with access to the data are able to identify the real problems that affect the energy-efficiency of data center operations. Our goal is, in part, to democratize research in data center energy-efficiency to enable a much broader set of researchers to make contributions in this area. To do so, this paper presents and analyzes detailed energy measurements from a state-of-the-art 15MW multi-tenant (“colo”) university data center.

While our measurements, analysis, and insights should prove useful to systems researchers, our study is particularly interesting since our data center—the Massachusetts Green High Performance Computing Center (MGHPCC)—is specifically designed to be a “green” facility, and thus incorporates many of same technological advances employed by recent state-of-the-art commercial data centers. The facility uses renewable cooling and renewable hydro-electric power, and is one of only 13 data centers in the country (and the only university data center) to receive a LEED Platinum rating [1]. The data center is jointly owned and operated by a consortium of universities in Massachusetts, including UMass, MIT,

¹PUE is a widely used metric that quantifies the effectiveness with which a data center cools and delivers power to servers. It is not a measure of energy-efficiency, since the performance of the servers is not included in the PUE computation.

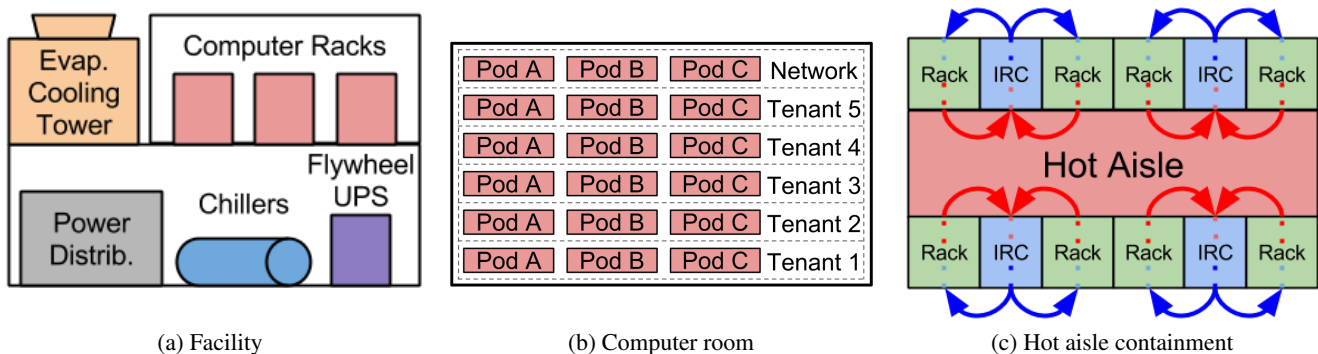


Figure 1: Physical layout of the MGHPCC.

Harvard, Boston University, and Northeastern. The data center is primarily used for research-oriented computing with batch workloads; each university in the consortium is a tenant, which allocates its reserved space within the data center as its own colo facility to house compute clusters owned by various research groups on each campus.

In this paper, we analyze detailed facility-level data, e.g., of energy and water use, from the second year of the MGHPCC’s operation. In doing so, we aim to address the following questions.

- Since the data center’s workload is primarily batch-oriented, how much do the time-of-day and seasonal effects influence the workload’s intensity? What are the implications for many previously proposed energy-efficiency optimizations, which often focus on exploiting these effects?
- What is the facility’s overall PUE and what is the contribution/overhead from each subsystem?
- How does the PUE vary over time and how is it affected by changing weather and the seasons?
- How does the PUE vary spatially across tenants and clusters and what causes such variations?
- What is the WUE of the facility?
- Given its use of renewable energy sources, what is the Carbon Usage Effectiveness (CUE) of the facility?

Unlike prior work on improving data center energy efficiency, we focus on detailed measurements and analysis of data from the facility itself, rather than from individual servers or clusters. While direct access to such data has largely been restricted to facility managers, it is important in developing energy optimizations for multi-tenant data centers. While data centers controlled by a single entity have access to the underlying servers and network infrastructure and are able to implement many previously proposed cluster-level energy optimizations, multi-tenant data centers do not. Thus, these data centers must apply optimizations at the facility-level, similar to how a multi-tenant commercial building generally cannot control the energy usage of its tenants. As a result, facility-level energy optimizations for data centers have more in common with building energy-efficiency optimizations than the server-centric optimizations that have largely been the focus of prior work. In answering the questions above, this paper makes the following contributions.

Detailed Design Overview and Data Collection. In Section 2, we describe details of the design and operation of a medium-sized data center facility, as well as the instrumentation available for facility data collection. The design applies many of the advanced techniques used in recent state-of-the-art commercial data centers. We gather and analyze data from over 900 sensors in the facility to both better understand its operation, and provide a baseline for other re-

searchers studying the efficiency of data centers at the facility level. **Temporal and Spatial PUE Analysis.** In Section 3, we then analyze both the temporal and spatial PUE of the data center over the past year. To estimate per-tenant spatial PUE, we develop a model for partitioning the energy usage of the data center’s centralized cooling system across different pods of racks based on their load. **CUE and WUE Analysis.** Based on our data, we also analyze the WUE in Section 4 and the CUE in Section 5. Our results indicate that, even at low capacity utilization, the data center’s WUE is better than published numbers for an average data center, while its CUE is near the published CUEs of the best commercial data centers.

2. BACKGROUND AND GREEN DESIGN

In this section, we present general background on the MGHPCC, including the design of its power and cooling infrastructure.

Overview. Our study focuses on the MGHPCC, a green data center built by a consortium of universities in Massachusetts for research computing. The data center is located in Holyoke, Massachusetts. The location was chosen based on the availability of abundant and cheap renewable hydroelectric power in Holyoke, the proximity to fibre-optic network backbones, and inexpensive real-estate. Massachusetts has a relatively cool climate with mean summer and winter temperatures of 23°C and -3°C, respectively, although summers months can be hot and humid; this cool climate enables the facility to employ renewable cooling, as explained later.

The data center became operational in November 2012 at the location of a former industrial mill site. The facility has 90,000 square feet of computing space and is provisioned for a 15MW peak load. The present utilization of the data center is less than 10% of its peak load, with an approximate compute load of 1MW and a non-compute load of 0.3 MW. As we discuss later, the compute load is steadily ramping up over time as new colo clusters are installed, and is expected to reach full capacity in a few years.

The data center is jointly operated by the university consortium and is structured as a multi-tenant facility with space pre-allocated to each member university. Each university uses its space to collocate compute clusters owned by various research groups and units on their campus. The multi-tenant colo data center houses a growing number of clusters for research computing with workloads that are primarily *batch-oriented* with different clusters running scientific batch jobs of various flavors.

Green Design The MGHPCC was designed as a green facility, and was the first university data center to achieve a LEED Platinum rating. As noted earlier, the facility is largely powered using renewable hydroelectric power, and employs a number of modern techniques to increase its energy-efficiency and minimize its power usage, as discussed below. For example, the data center employs




Figure 2: The MGHPCC’s cooling infrastructure leverages free evaporative cooling and backup chillers.

hot aisle containment and uses renewable free cooling from outside air whenever possible to reduce its cooling-related energy usage. Finally, the facility employed a number of sustainable practices during the construction process and continues to do so during its operations (described in more detail here [1]).

Physical Layout Figure 1 depicts the physical layout of the data center. As shown in Figure 1(a), the data center comprises of two levels. The lower level contains the main power infrastructure, including a utility sub-station and flywheel-based uninterruptible power supply (UPS) system, and cooling infrastructure including chillers and water pumps. The upper level mainly contains racks for hosting the computing infrastructure. Evaporative cooling towers are housed on the roof of the first floor, adjacent to the computer floor. The data center also includes backup diesel generators that provide power for a fraction of the facility in case of a utility outage. The UPS system mends the gap between an outage and the activation of the diesel generators to prevent servers from losing power. Figure 1(b) shows the layout of the computer floor. There are five main aisles of racks, one for each tenant. A sixth aisle is the “networking aisle” and houses networking equipment to connect each tenant’s computing infrastructure to the incoming fibre optic lines. Each tenant’s aisle has three groups of racks, each referred to as a pod. Racks in each pod are designed for hot aisle containment as shown in Figure 1(c).

Power Infrastructure. The power infrastructure for the data center resembles a small-scale distribution network in the electric grid. The infrastructure comprises of 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 before entering the switchboards at the lowest levels, and is finally delivered to the busplugs that feed the power distribution units (PDUs) in each server 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. First, the facility uses high voltage, and low current, to deliver power, which reduces losses due to power conversion and heat generation. Higher distribution voltages also make it possible to eliminate an entire tier of transformers from the distribution network, further reducing transformer losses. Second, energy losses due to the UPS system are another source of higher PUEs.

Since the data center houses research computing infrastructure, not all of which is “mission critical”, only a fraction (roughly 20%) of each tenant’s racks are backed up by the centralized UPS system. The tenants are then able to choose how to partition their compute infrastructure between UPS and non-UPS racks. The remaining racks are not connected to the centralized UPS system, which naturally avoids UPS losses for the 80% of the racks in the data cen-

ter. Third, in many data centers, the UPS system for UPS-backed racks normally operates in a double conversion mode, which incurs losses in both directions when converting from AC to DC and from DC to AC. Double conversion is often useful in conditioning the incoming power to provide a consistent high-quality AC power signal, i.e., a “tight” 60Hz sine wave, for mission-critical applications. At the MGHPCC, most workloads are not mission-critical and the hydroelectric power offered by the local utility is already high-quality. Thus, UPS systems in the MGHPCC are configured to operate in direct mode, where power is fed directly to racks, rather than through the UPS, and there is a near instantaneous transfer (within tens of milliseconds) to UPS systems when a power failure is detected.

Finally, 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. The data center is provisioned for 18 seconds of UPS power in case of an outage, and standby diesel generators take over within this time period.

Cooling Infrastructure. Traditionally data centers have used chillers to cool the servers in the facility. However, chillers consume a significant amount of energy and their use is a key contributor to high data center PUEs. Thus, modern data centers have begun using alternative technologies to cool their servers and lower their PUEs. The MGHPCC leverages “free cooling” (also known as “renewable cooling”) to cool servers. Specifically the data center uses evaporative cooling technology that essentially uses the outside air to cool servers. Figure 2 depicts the two cooling water loops used in the data center.

The water in the outer loop is cooled using evaporative cooling towers whenever the temperature of the outside air permits it. The inner water loop circulates water through the computer room racks. The water loop is used to extract heat from the hot air ejected by the server, which cools the air. The hot water in the inner loop is then sent to the heat exchanger, where the heat from this water is exchanged with the cold water in the outer water loop. Doing so, transfers the heat from the servers to the outer loop, cooling the water in the inner loop, which is sent back to the computer racks. The hot water in the outer loop is sent to the evaporative cooling tower where it is cooled again using the outside air, through an evaporative process, and circulated back to the heat exchangers.

The cooler climate in Massachusetts permits the use of this free cooling approach for over 70% of the year. Evaporative cooling becomes less feasible or infeasible during the warmer, and more humid, summer months. During these months, the data center falls back on using chillers to cool water in the inner loop. A hybrid mode is also possible where water is partly cooled using evaporative cooling and then cooled further using chillers (when the weather permits part, but not full, free cooling).

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

Finally, the data center maintains the computer floor temperature at 80°F, which is a higher temperature than traditional data centers; doing so, reduces the amount of cooling required, which in

Dataset	Description	Resolution
Power	All IT and non-IT power usage	20 second
Mechanical	Cooling equipment usage	1 minute
Water	Water usage	Monthly
Weather	Temperature, humidity	5 minute

Table 1: Description of our datasets.

turn improves cooling efficiency without impacting the reliability of modern server hardware.

Compute Infrastructure. As noted earlier, the data center is a multi-tenant facility, with each university tenant treating its allocated racks as a co-location facility to house research computing clusters from on-campus groups. Thus, the data center does not own, or exercise direct control, over the type of servers deployed at the facility. Consequently, all energy optimizations “stop” at the rack and the data center cannot mandate the use of any specific server model. This is in contrast to companies, such as Facebook or Google, that are capable of deploying optimizations anywhere in the data center, including energy-optimized servers that may be DC powered, use power supplies optimized for their load levels, or employ local on-board batteries rather than a centralized UPS system.

The compute load has been increasing steadily as new clusters are deployed by each tenant. One consequence of the presently low capacity utilization is that the current PUE of the facility is higher than it would be at full utilization, largely because the cooling infrastructure is sized for a much higher load and is less efficient at lower loads (since it is not energy proportional).

Monitoring Infrastructure. The data center is highly instrumented to monitor all aspects of its operation and has several thousand points of instrumentation that provide real-time data on power, cooling, and water usage within the facility. Note that this facility-level data is separate from the type of data monitored at the level of individual servers and clusters; facility-level data is monitored by the data center staff, while the server and cluster-level data is accessible only to tenants (who own the compute infrastructure), and not to the facility staff.

Table 1 depicts the various datasets that we have gathered from the facility and form the basis of our study in this paper. At the facility level, the power distribution infrastructure is monitored by over 900 networked electric meters that monitor power 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 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 chillers.

The facility’s mechanical systems, which are primarily associated with the cooling infrastructure, are also monitored by a conventional building management system. The available data includes water pump flow levels at various points in the water loops, as well as data from in-row chillers, such as fan speed, water inlet and outlet temperature, and water flow 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. The outside weather data is monitored using a weather station (we also use data from Weather Underground), and the facility’s overall water usage is recorded by a water meter.

Datasets: We use the datasets in Table 1 gathered over a 12 month period from May 2014 to April 2015, which roughly corresponds to the second year of the data center’s operation. As shown in the table, we use four different datasets in our analysis. The power data is gathered from the 900 electric meters deployed within




Figure 3: Average monthly temperature at the data center.

the power distribution system. This data is gathered at 20 second resolution and includes the average power usage data of individual racks, UPS-backed rack usage, aggregate in-row chiller power usage, and the power usage of the cooling infrastructure such as chillers and water pumps.² The mechanical data comprises primarily data from in-row chillers, which includes fan speed, water flow speed, as well water inlet and outlet temperature. Our water usage data includes the monthly water usage of the facility. Finally, weather data consists of outside temperature and humidity at the facility over the year (see Figure 3). Inside temperature data monitored at hot and cold side of individual racks is also available, but not directly used in our analysis. Data collection is ongoing and will last several years to enable a long-term efficiency study of the MGHPC.

3. PUE ANALYSIS

In this section, we analyze in detail the power usage of the MGH-PCC. We first analyze the IT load, e.g., of the server and networking equipment, over different time scales to quantify its impact on the facility’s PUE. We also analyze the various factors contributing to the observed PUE, and consider the impact of seasons on PUE to quantify the benefits of free cooling. In addition to analyzing temporal PUE, we perform a spatial analysis to compute per-tenant PUEs and analyze how the PUE varies across tenants and why.

3.1 IT Load Analysis

Figure 4 depicts the IT power usage of the data center at the time scale of months, a week and a day. The IT power usage is derived by combining the electricity meter data for only those meters that supply power directly to the computing and networking racks. Figure 4(a) depicts the mean monthly IT load from May 2014 to April 2015. As the figure shows, the IT load steadily rose over the one year period, largely due to new colo compute clusters being commissioned at a steady rate by the various tenants. The figure also shows that the mean IT load is less than 1MW, which is less than 10% of the peak provisioned power capacity. Thus, we expect the increasing trend in IT load to continue for the foreseeable future.

As a research computing facility, the workload of the data center is primarily batch-oriented. Hence, analyzing the IT load over the time scales of a week and a day are instructive in determining whether the workload exhibits the time-of-day and week-of-day effects that are common in commercial data centers that host interactive Internet workloads. Figure 4(b) depicts the mean daily IT

²Some older power data is only available at an 8 hour resolution, while all recent data is archived at 20s granularity.




Figure 4: Variation in IT Power load at time scales of months, days and hours. The load has been increasing steadily over the course of the year, but the batch-oriented load does not show any significant day-of-the-week or time-of-day effects.

load for different days of the week for the month of April 2015, while Figure 4(c) shows the load for different hours of the day. The figures show that, unlike interactive workloads, the batch workload of the data center does not show pronounced time-of-day or day-of-the-week effects. As Figure 4(b) shows, there is only a very modest rise in the load in the middle of the week, but no significant weekday-weekend effects. Similarly, 4(c) shows that long-running batch jobs or batch schedulers with a queue of jobs cause the compute load to remain high across both the day and the night. While there is a very small drop ($\sim 2.5\%$) in load in the early hours of the day (4am to 7am), possibly due to the completion of overnight batch jobs, it does not yield any perceptible time-of-day effects.

Result: The explicit goal of many server- and cluster-level energy optimizations is to exploit such time-of-day and day-of-week effects, e.g., by powering down servers when the workload drops to make them more energy-proportional [4, 22]. Our data indicates that these types of energy optimizations are not as applicable to the MGHPPCC, as it does not experience significant time-of-day and day-of-week effects.

3.2 Temporal PUE Analysis

We next analyze the PUE of the data center over the course of the year. The PUE metric is computed as ratio of the facility’s total power usage to the power usage of the IT equipment, e.g., servers and switches. The total power usage is monitored directly by a networked meter that measures the power entering the facility from the grid, while the IT power usage is computed as outlined in the previous section. Figure 5 depicts the monthly PUE of the data center over the course of the year. The figure shows that the PUE varies between 1.285 and 1.509 over the year. We note that a mean PUE value of 1.377 is significantly lower than the average PUE value of 1.7 (or higher) that is common in enterprise data centers in the industry [21]. However, the value is not as low as the PUEs near 1.1 reported by the newest (and most effective) data centers built by large Internet companies, such as Facebook and Google [17, 10].

We analyze the key contributors to our PUE in more detail below. We note that the data center’s cooling infrastructure is sized for a much greater load than the present server room occupancy produces and is not energy proportional. Hence, we believe the facility will achieve further reductions in PUE as its IT load increases to full utilization. Thus, the PUE values we report here are conservative in that the facility’s design is more effective than its current PUE values indicate.

Second, the figure shows a reduction in the PUE from May 2014 to April 2015. This reduction can be attributed to two possible factors. First, since the IT power load has risen during this period, the cooling infrastructure, which is not energy proportional, becomes





Figure 5: Temporal variations in the monthly PUE.

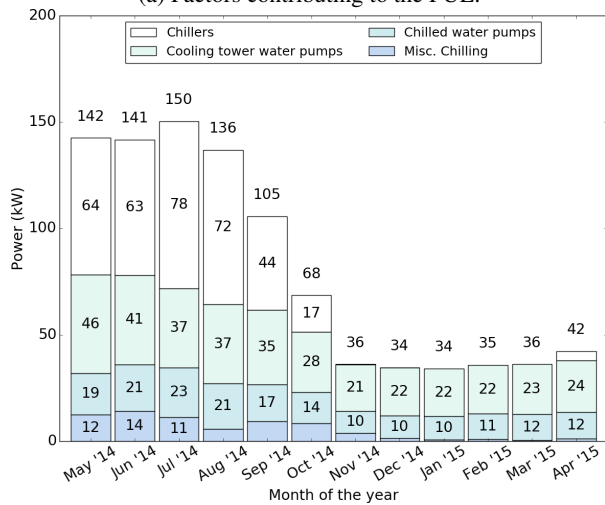
relatively more efficient with increased IT load, yielding a lower PUE. Second, free cooling is feasible only during cooler months of the year and the data center needs chillers to cool the facility in warmer months. This use of chillers will cause a higher PUE in warmer months (May to September), and the PUE is lower for the remaining months when free cooling is used.

Our analysis shows that the second factor dominates, since turning off chillers yields a greater reduction in PUE than the increase in IT power load during this period (since the total IT power load is relatively low, its impact on the PUE is much smaller). Figure 6(a) confirms that chillers are used during the months of May to September and that they consume a significant amount of energy (thereby contributing to a higher PUE in those months). The lack of chiller energy use in other months stems from the use of free cooling during those months, yielding a lower PUE. This result shows the PUE of the data center is 1.413 when chillers are in use and 1.301 when free cooling is used; in other words, PUE decreases 0.112 when not using chillers even at this low capacity utilization.

Next, we analyze the non-IT load of the data center, which in turn reveals the various factors contributing to the PUE. Conventional wisdom has held that there are two main sources of overhead that contribute to the non-IT power load: cooling infrastructure and power distribution losses, including UPS losses. Figure 6 depicts the power usage of various non-IT loads in the data center. As the figure shows, roughly half of the non-compute load can be attributed to cooling and other mechanical systems; a quarter can be attributed to power losses; and another quarter can be attributed to other factors, including measurement error. The measurement error is only a few tens of kilowatts (of the multi-megawatt power usage), but it is almost a quarter of the current non-IT load, since the over-



(a) Factors contributing to the PUE.



(b) Cooling factors contributing to the PUE


Figure 6: Various factors contributing to the PUE, including cooling energy use and power losses.

all capacity utilization is low. This error will become negligible once the data center becomes fully utilized.


The cooling load, which is roughly half or more of the total non-IT power load, comprises of (i) chillers, (ii) cooling tower (“outer loop”) water pumps, (iii) chilled water (“inner loop”) water pumps, (iv) air handlers, (v) in-row chillers, and (vi) miscellaneous cooling equipment. As shown in Figure 6, chillers are the largest component whenever they are operating in the warmer months. Thus, eliminating the use of chillers by using free cooling yields a significant reduction in both cooling power usage and PUE.³

The power losses consist of two key components: power distribution losses that occur when the incoming power flows through various components of the data center’s distribution network, and UPS losses that occur in all UPS systems. The data center has opti-

³As shown in Figure 6(b), cooling tower pumps used more power in the months of May and June than in warmer summer months such as July; this is due to use of hybrid cooling to partially cool water first and then use chillers to cool the rest. As seen, this mode uses more power than using chillers alone in warmer months; facility managers are currently optimizing pump controls to enhance the efficiency of hybrid cooling.



(a) Daily



(b) Hourly

Figure 7: Daily and hourly variations in the data center PUE.

mized UPS losses by not using UPS systems in double conversion mode, which results in losses from AC to DC and DC to AC conversion; instead power is directly fed to the computing racks with a fast fail-over to UPS upon detecting a power loss or fluctuation. Further, since only 20% of the racks have a UPS backup, this naturally cuts down on the total UPS losses in the facility.

Finally, Figure 7 depicts the daily and hourly variations in the PUE for the different days of the week in April and a particular day in April 2015, respectively. As Figure 7(a) shows, the PUE is mostly flat over the course of a week, as the corresponding IT load seen during different days of the week, shown in Figure 4(b), is nearly flat without any week of the day effects. The hourly PUE in Figure 7(b) shows small fluctuations caused by corresponding variations in the instantaneous hourly IT load; the hourly PUE is mostly flat when averaged over the month, as expected, since the hourly IT load is similarly flat without any significant time-of-day effects.

Result: The MGHPCC’s PUE is 1.3773, even at a low 10% capacity utilization, a figure generally considered to be quite competitive, although not as low as some of the most effective data centers in industry. We note that this relatively low PUE is achieved in a multi-tenant colo facility where the data center has no direct ability to optimize server hardware (unlike, say, Google or Facebook data centers where end-to-end optimizations that include the server tier are feasible). Free cooling yields a significant reduction in PUE over periods when chillers have to be deployed, demonstrating the

benefits of optimizing the cooling infrastructure on lowering the PUE. Finally, an interesting artifact of our batch workloads is that they do not exhibit typical time-of-day or weekend effects, pointing to the possibility of higher server utilizations throughout the day.

3.3 Per-tenant PUE Analysis

The previous section shows that the overall PUE of the data center is between 1.285 and 1.509. As noted earlier, the data center is a multi-tenant facility, with each tenant operating an entire aisle of racks independently of the other tenants. Thus, analyzing and comparing the PUE of each tenant to the facility’s overall PUE is useful. Further, different tenants are in different stages of their roll-out of computing equipment at the data center, and per-tenant PUE analysis can provide insights into how the PUE might vary when the racks are at different capacity utilizations.

There is no well-known method to compute the per-tenant PUE. While the IT load is directly metered on a per-rack (and per-tenant) basis, the non-IT load is not. The cooling infrastructure including the evaporative cooling towers and chillers are facility-wide equipment and not deployed on a per-tenant basis. Thus, to compute PUE on a per-tenant basis, we must determine how to apportion to non-IT (and particularly cooling) loads across individual tenants.

Fortunately, in our case, each tenant operates an entirely separate aisle of racks, such that the racks do not mix computing equipment from multiple tenants. Thus, determining the compute and cooling load of all racks in a given aisle is equivalent to determining the computing and cooling load due to that tenant. In other words, spatial analysis of PUE across racks and aisles in our case also yields the PUE of the various tenants.

3.3.1 Per-tenant PUE and IRC Power Models

To determine the PUE of an aisle of racks, or more precisely a pod of racks, we make the following assumptions.

First, we assume that the hot air containment used by the racks to isolate the hot air from cold air is perfect [18]. That is, the hot air from the racks is fully contained and does not impact the temperature, or associated cooling, of the racks in other aisles.

Second, while the cooling infrastructure, such as chillers and cooling towers, are facility-wide equipment, the in-row chillers are deployed to locally cool adjacent racks and represent per-pod (and per-tenant) cooling equipment. Further, in-row chillers directly remove the heat generated by racks in each pod, and hence, the power consumed by in-row chillers is an indirect measure of the cooling needs of that pod. Hence, we can use power consumed by the in-row chillers of a pod to apportion the remaining non-IT load across pods. We note that such a method is an approximation since neither the in-row chillers nor the facility-wide cooling equipment are energy-proportional, i.e., a linear increase in heat generated *does not* result in a proportionate linear increase in power usage of the IRCs or the cooling equipment.

Given these assumptions, the PUE of a pod is given below.

$$PUE_{pod} = \frac{P_{total}^{pod}}{P_{IT}^{pod}} = \frac{P_{IT}^{pod} + P_{non-IT}^{pod}}{P_{IT}^{pod}} \quad (1)$$

In the equation, P_{total}^{pod} denotes the total power usage of a pod, while P_{IT}^{pod} and P_{non-IT}^{pod} denote the IT and non-IT power used by the pod. The IT power consumed by each rack is directly measured, while the non-IT power used by the pod must be estimated. Based on the assumption above, the non-IT power usage of a pod is assumed to be proportional to the power consumed by the pod’s in-row chillers, which itself depends on local cooling demands.

Thus, we estimate the non-IT power usage of the pod below.

$$P_{non-IT}^{pod} = \frac{P_{IRC}^{pod}}{P_{IRC}^{total}} (P_{total} - P_{IT}) \quad (2)$$

Here, P_{IRC}^{pod} and P_{IRC}^{total} denote the power consumed by the IRCs of a pod and the total power consumed by IRC across all aisles and pods, and P_{total} and P_{IT} denotes the total facility power and the total IT power across all tenants (the difference between the two is the total non-IT power usage). If each in-row chiller were individually metered, all of the quantities in the above equation would be known. However, the data center meters IRC power consumption in the aggregate (for groups of IRCs) and thus the power used by individual IRCs is not directly monitored. However, our mechanical dataset monitors the fan speeds of the in-row chillers and it is well known (from IRC manuals) that power consumption of an IRC is a cubic function of its fan speed. Thus, we use a model to estimate IRC power usage from its monitored fan speed, as shown below.

$$P_{IRC} = \alpha \cdot x^3 + \beta \quad (3)$$

Here, α and β are constants that depend on a specific model of an IRC and x denotes the fan speed. Since the aggregate power consumed by a group of IRCs is metered and known, the following relationship holds for power consumed by IRCs in each metered group at time instant t .

$$P_{IRC}^1 + P_{IRC}^2 + \dots + P_{IRC}^n = P_{IRC}^{total} + \epsilon \quad (4)$$

Here, P_{IRC}^i denotes the power consumed by the i^{th} IRC within a metered group and P_{IRC}^{total} denotes the total power consumed by all IRCs within that group. ϵ is a term that captures the measurement error. By substituting Equation 3 for each individual IRC into Equation 4, we obtain a set of equations, one for each measurement interval t , for the unknown constants α and β .

We can then use regression on this set of equations to derive the α and β that minimize the error term ϵ . By deriving α and β , the regression then yields an IRC power model where the power consumed by the IRC is a function of the fan speed x with known constants α and β . That is, $P_{IRC} = \alpha \cdot x^3 + \beta$. Since fan speeds are directly measured and available to us, the power usage of a pod can be estimated using this approach, and this value can be substituted in Equation 2 to estimate the non-IT power usage of a pod P_{non-IT}^{pod} . Since the IT power usage of each pod is directly measured and known, we can compute the PUE of each pod.

We ran the regression on the measured values of IRC fan speeds and the total IRC power consumption to derive the IRC power model as discussed above. Figure 8 depicts the model we learned for the IRC power consumption as a function of the IRC fan speed. To validate our model, we compute the power consumed by individual IRCs using measured fan speeds, and then compare the sum of the computed individual IRC power values to the total IRC power as measured by the electric meter. Figure 9 depicts the estimated IRC power consumption from the model and the actual values from the metered data. As can be seen, there is a close match between the model estimates and the actual values.

3.3.2 Spatial PUE Analysis.

Given our models above to estimate the per-IRC power consumption and the per-pod PUE, we next analyze the *spatial* distribution of computing load and the resulting PUE on a pod-by-pod and tenant-by-tenant basis. We note at the outset that the spatial distribution of servers across racks can impact the PUE, similar to how prior work has shown the spatial distribution of compute load influences cooling costs [16]. To illustrate, consider two different deployments of ten servers, one where all ten servers are housed

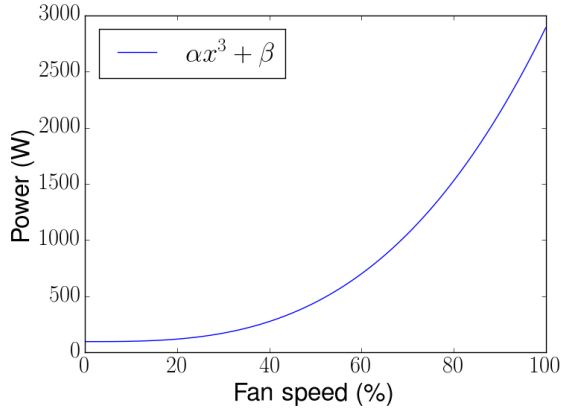


Figure 8: IRC power model learned via regression: the model is cubic in fan speed with parameters $\alpha = 0.00279$ and $\beta = 97.7$.




Figure 9: We validate our IRC power model by comparing the fit of total modeled IRC power to the aggregate meter values.

in a single rack and another where each individual server is deployed on a separate rack, i.e., a “depth-first” versus “breadth-first” deployment. Although the IT power load of these ten servers is independent of how they are placed on the racks, the cooling load depends on the spatial distribution of servers. In the former case, the cooling load is concentrated in one rack and a single IRC can handle the cooling of the servers while the remaining IRCs in the pod can remain idle. In the latter case, the cooling load is spread across multiple racks, and multiple IRCs will need to absorb this spatially by spreading out the cooling load. Thus, the two deployments will result in different PUE values even though both have the same IT load. This toy example illustrates that two pods with identical IT loads may have different PUEs if they have different spatial distributions of servers across racks.

Figure 10(a) depicts the spatial distribution of the IT power usage of server racks in each pod across different pods and tenants. Recall that each tenant has an aisle dedicated to them, and each aisle is partitioned into three pods of racks; let T_i denote tenant i , and the suffix A , B or C denote the three pods allocated to that tenant. The sixth row, denoted by N , houses network equipment to connect the tenants to various network/fibre backbones. As shown, different tenants are at different stages of deployment of their research clusters—the IT load of a pod varies from 2kW to 214kW. Many pods—those with usage of less than 5kW—remain empty. A few pods are moderately loaded and have IT loads of 50-100kW. Only one pod (pod A for tenant 2) is nearing capacity and has a

current load of nearly 220kW. The pods housing networking gear also remain lightly utilized.

Using these IT power loads and our models, we compute the PUE of each pod, which is depicted as a heat-map in Figure 10(b). The data shown is for April 3-11, 2015, where the overall PUE of the data center was 1.28. To compute this figure, we used a more refined PUE model than the one discussed in Section 3.3.1 where the UPS power losses are only attributed to pods with UPS backup power, rather than being uniformly spread across both UPS and non-UPS racks. The figure reveals the following insights.

As expected, the pods that have low utilization also have high PUEs. However, these PUE values are not meaningful since they are associated with an IT load that is close to zero. In general, we observe that the PUE of a pod is inversely proportional to the IT load: as the IT load increases, its associated PUE value falls. Thus, lightly and moderately loaded pods have PUEs that are higher than the facility-wide PUE, while more heavily loaded pods have lower PUEs than the overall PUE. This trend is not surprising, since the cooling equipment and in-row chillers are not energy proportional and operate at optimal efficiency levels at near-peak loads, and are much less efficient at lower loads. Hence, the PUE is much lower (and better) for more heavily loaded pods.

Interestingly, pod B for tenant 4 and pod C for tenant 2 have roughly similar IT load, i.e., 70kW, but have different PUE values of 1.28 and 1.21, respectively. This is a real-world depiction of the toy example above, and demonstrates that the spatial distribution of servers in a pod does matter and can impact a pod’s PUE, i.e., two pods with identical compute loads but different spatial distribution of servers can yield different PUEs.

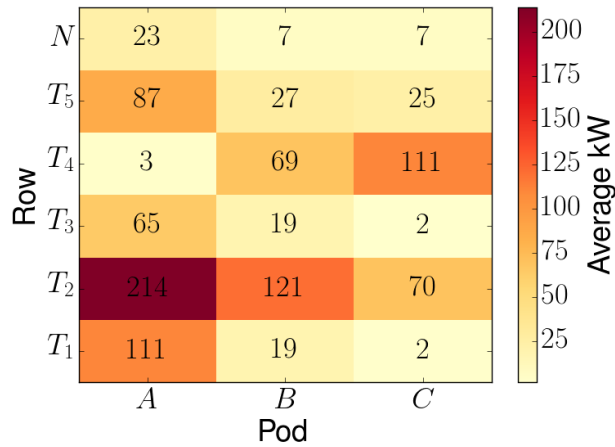
4. WUE ANALYSIS

In addition to consuming significant amounts of power, data centers also typically consume significant amounts of water, mainly as part of their cooling infrastructure. 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 below.

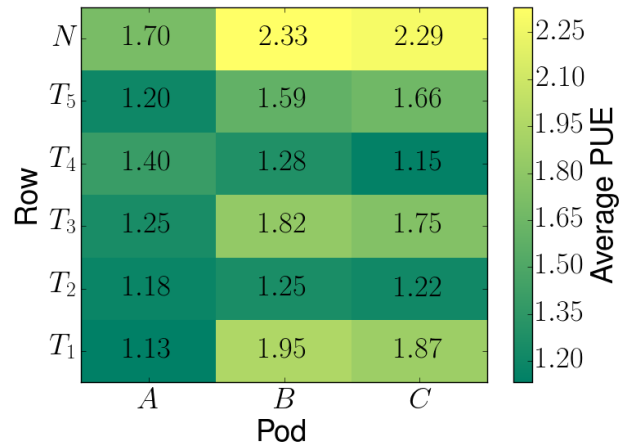
$$\text{WUE} = \frac{\text{Water Usage (liters)}}{\text{IT Energy Usage (kWh)}} \quad (5)$$

Intuitively, WUE is defined as liters of water used per kilowatt-hour (kWh) of energy used by the IT equipment. Unlike its better known PUE counterpart, there is little data published on WUE of data centers. Recently Facebook released data indicating that the WUE of their Prineville data center was 0.28 L/kWh and their Forest City data center was 0.34 L/kWh [7]. In contrast, an average 15MW data center (similar in size of our MGHPCC data center) may consume as much as 360,000 gallons of water each day [15]. Assuming a moderate PUE of 1.5, a fully utilized 15MW data center has an IT load of 10MW, which translates into a WUE of 5.67 L/kWh. In contrast, GreenGrid published a report indicating the average data center has a WUE of 1 L/kWh but did not provide details [11].

Figure 11(a) depicts the monthly water usage of the data center over a 12 month period. The present water usage varies between 1000 kL and 2000 KL per month depending on the season of the year. As the figure shows, the water usage is higher in the warmer months and lower in cooler months. This is not surprising since warmer months lead to more evaporative water loss and the use of chillers in these months consumes more water. In contrast, the data



(a) IT power consumption



(b) pod PUE

Figure 10: Heat-maps showing the mean IT power consumption and mean pod PUE of different pods and tenants across our multi-tenant data center. Data shown is for April 3-11, 2015.

center relies on free cooling in cooler months, resulting in lower water usage in those months. Next, Figure 11(b) depicts the WUE of our data center over a 12 month period. As shown, the WUE values vary between 1.5 L/kWh and 3.0 L/kWh over the course of the year. The WUE rises in the warmer months and falls in the cooler months. We attribute this trend to two factors: rising IT load over the course of the year results in more effective use of water and a fall in WUE, and seasonal effects where there is less loss from evaporative cooling in cooler months.


While a WUE between 1.5 and 3.0 L/kWh is already lower than the average WUE of 5.67 hinted in [15], it is higher than the Green-Grid value of 1. There is little real-world data available to provide a meaningful comparison. We note, however, that the MGHPC is presently operating at only 10% capacity, and we expect a significant fall in WUE as the capacity ramps up, in line with the trends observed in the initial months of 2015. Thus, our hypothesis is that, in the long run, the WUE of the MGHPC will be significantly lower than the “typical” data center, and in line with its green design goal. Although not shown here, the data center uses a number of other measures to optimize its water footprint, including the use of water filtration techniques to maximize the circulation of water in the two water loops as well as use of recycled water for many auxiliary purposes, e.g., for landscaping.

5. CUE ANALYSIS


Our final analysis focuses on the carbon impact of the MGHPC, 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 [5]. The CUE of a data center is defined as below

$$\begin{aligned}
 \text{CUE} &= \frac{\text{CO}_2 \text{ emissions from the total data center energy}}{\text{IT equipment energy}} \quad (6) \\
 &= \frac{\text{kg CO}_2}{\text{kWh}} \cdot \frac{\text{Total data center energy}}{\text{IT equipment energy}} = \frac{\text{kg CO}_2}{\text{kWh}} \cdot \text{PUE}
 \end{aligned}$$

As shown, the CUE depends significantly on the carbon emissions due to the electricity consumed by the data center. The carbon emissions of the electricity consumption, in turn, depend on the generation source mix of the electric utility that supplies power to the data center. In the event that the data center uses on-site or



(a) Monthly water usage



(b) Monthly WUE

Figure 11: Monthly water usage and WUE of our data center.

Fuel Type	Energy (MWh)	Energy (%)	CO ₂ (kg)	CO ₂ (%)
Oil	1724	0.4	1476897	16.3
Hydro	261691	66.7	0	0
Nuclear	61310	15.6	0	0
Solar	6105	1.6	0	0
Contracted (carbon free)	40800	10.4	0	0
Contracted (other)	20592	5.3	7584064	83.7
Total	392222	100	9060054	100

Table 2: Holyoke Gas & Electric power generation sources.

Utility	HG&E	U.S. grid		
		mean	min	max
$\frac{\text{kg CO}_2}{\text{kWh}}$	0.0231	0.559	0.203	0.860

Table 3: Power producer carbon intensity [2].

contracted renewable energy, that portion must also be considered in the overall electricity mix as well.

The MGHPC does not use any on-site renewables and depends entirely on the local utility company for its power needs. The local utility, Holyoke Gas and Electric (HG&E), generates a large fraction of its electricity using hydroelectric power from a sophisticated canal system. This hydro-electric power is not only inexpensive, but also a clean source of renewable energy. The mix of generation sources used by HG&E to generate power for its customers is shown in Table 2. As the table shows, HG&E generates or purchases 94.3% of the electricity from carbon-free sources, two-thirds of which derives from the local canal system. In addition, HG&E operates one of the largest solar deployments in New England, totaling 5.3MW of installed capacity. In addition, 15% of its power derives from nuclear power plants, which are also carbon free. HG&E also purchases electricity in the wholesale electricity market through a variety of contracts. Roughly 16% of its needs are met from these contracted sources, of which 10.4% comes from carbon-free sources. Thus, their high fraction (94.3%) of carbon-free electricity in their generation mix yields a low ratio (of 0.0231) in the amount of kilograms of CO₂ emitted per kWh of energy generated. This ratio, which is nearly an order of magnitude lower than the most carbon efficient region in the U.S., is shown in Table 3.

Consequently, 94.3% of electricity consumed by the MGHPC is carbon-free. Figure 12 depicts the monthly CUE of the data center. The CUE varies from 0.0297 to 0.0349 with an annual average of 0.0318. 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 1.8 PUE). Recently, Apple claimed that 100% of its data centers are powered using renewables and Google has followed a similar strategy of using contracted wind energy for its data centers. Our data center compares favorably to these state-of-the-art data centers in terms of CUE, but has achieved its low CUE by careful choice of location and utility rather than building or contracting renewable energy.

6. CONCLUDING REMARKS

In this paper, we present an empirical analysis of the efficiency of a green academic data center. The data center we study, the MGHPC, is a multi-tenant facility that is designed to house colo research clusters running batch-oriented workloads. Our temporal PUE analysis reveals that the data center has PUE values that range from 1.285 to 1.509, with higher PUEs in warmer summer months. We show that free cooling, which avoids the use of chillers can reduce PUE by as much as 0.224 in cool seasons. Our spatial multi-




Figure 12: Monthly CUE of the data center.

tenant analysis reveals the non-proportional nature of the cooling equipment, which causes its efficiency to increase as each pod of racks becomes fully utilized, yielding lower PUEs. Our water usage analysis shows that the WUE of the data center is between 1.5 and 3 L/kWh. Finally, we show that data center has a CUE of 0.03, which is 25× lower than a typical data center. The low CUE is mainly due to the large portion of renewable energy within the electricity mix supplied by the local utility. Overall, our results validate the green design of the data center and point to further efficiencies in the future at higher utilizations.

This work is a first effort in a long-term efficiency study of the MGHPC. In the future, we will examine energy efficiency at the server level. By calculating metrics, such as Energy-Agility proposed in a companion work [20], we will account for server efficiency that is beyond the scope of PUE. Additionally, we will develop data driven models that will learn the relationship between various control parameters and their impact on electricity and cooling demands. We will use these models to determine inefficient settings and tune these parameters to enhance day-to-day operational efficiency.

Acknowledgements. We would like to thank our shepherd Klaus-Dieter Lange and the anonymous reviewers for their astute comments that improved this paper. This work is supported in part by the UMass Clean Energy Extension, NEAGEP/IMSD Fellowship, and NSF grants CNS-1405826, CNS-1422245, and ACI-1339839.

7. REFERENCES

- [1] Massachusetts Green High Performance Computing Center: Green Design. <http://www.mghpcc.org/about/green-design/>.
- [2] U. E. P. Agency. eGRID 9th edition Version 1.0 Year 2010 GHG Annual Output Emission Rates. http://www.epa.gov/cleanenergy/documents/egridzips/eGRID_9th_edition_V1-0_year_2010_GHG_Rates.pdf.
- [3] D. Andersen, J. Franklin, M. Kaminsky, A. Phanishayee, L. Tan, and V. Vasudevan. FAWN: A Fast Array of Wimpy Nodes. In *SOSP*, October 2009.
- [4] L. Barroso and U. Hözl. The Case for Energy-Proportional Computing. *Computer*, 40(12), December 2007.
- [5] C. Belady, D. Azevedo, M. Patterson, J. Pouchet, and R. Tiple. Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric. *White Paper*, 32, 2010.
- [6] J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, and R. P. Doyle. Managing Energy and Server Resources in Hosting Centers. In *SOSP*, October 2001.

- [7] Facebook. OpenCompute Project: Energy-Efficiency. <http://www.opencompute.org/about/energy-efficiency/>.
- [8] K. Fehrenbacher. The Era of the 100MW Data Center. In *Gigaom*, January 31 2012.
- [9] I. Goiri, T. D. Nguyen, and R. Bianchini. CoolAir: Temperature- and Variation-Aware Management for Free-Cooled Datacenters. In *ASPLOS*, March 2015.
- [10] Google. Efficiency: How we do it. <http://www.google.com/about/datacenters/efficiency/internal/>.
- [11] GreenGrid. Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric. <http://www.thegreengrid.org/en/Global/Content/white-papers/WUE>.
- [12] J. Koomey. Growth in Data Center Electricity Use 2005 to 2010. In *Analytics Press*, Oakland, California, August 2011.
- [13] I. Manousakis, I. Goiri, S. Sankar, T. D. Nguyen, and R. Bianchini. CoolProvision: Underprovisioning Datacenter Cooling. In *SoCC*, August 2015.
- [14] L. McTiernan. Open Sourcing PUE/WUE Dashboards. <https://code.facebook.com/posts/272417392924843/open-sourcing-pue-wue-dashboards/>, March 2014.
- [15] R. Miller. Data Centers Move to Cut Water Waste. In *Data Center Knowledge*, April 9th 2009.
- [16] J. Moore, J. Chase, P. Ranganathan, and R. Sharma. Making Scheduling “Cool”: Temperature-Aware Workload Placement in Data Centers. In *USENIX*, June 2005.
- [17] J. Parr. Designing a Very Efficient Data Center. <https://www.facebook.com/notes/facebook-engineering/designing-a-very-efficient-data-center/10150148003778920>, April 2011.
- [18] A. O. Rabassa III. Economic performance of modularized hot-aisle contained datacenter PODs utilizing horizontal airflow cooling. <http://hdl.handle.net/1721.1/90248>.
- [19] S. Rivoire, M. Shah, and P. Ranganathan. JouleSort: A Balanced Energy-Efficient Benchmark. In *SIGMOD*, June 2007.
- [20] S. Subramanya, Z. Mustafa, D. Irwin, and P. Shenoy. Beyond Energy-Efficiency: Evaluating Green Datacenter Applications for Energy-Agility. In *ICPE*, March 2016.
- [21] Y. Sverdlik. Survey: Industry Average Data Center PUE Stays Nearly Flat Over Four Years. In *Data Center Knowledge*, June 2nd 2014.
- [22] N. Tolia, Z. Wang, M. Marwah, C. Bash, P. Ranganathan, and X. Zhu. Delivering Energy Proportionality with Non-Energy-Proportional Systems: Optimizing the Ensemble. In *HotPower*, December 2008.