The Electricity Journal 29 (2016) 61–69

Contents lists available at ScienceDirect

The Electricity Journal

journal homepage: www.elsevier.com/locate/electr

Power usage effectiveness in data centers: Overloaded and underachieving

Nathaniel Horner*, Inês Azevedo

Department of Engineering and Public Policy, Carnegie Mellon University, Baker Hall 129, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:
Available online xxx

Keywords:
Data center
Energy metric
PUE
Information technology

ABSTRACT

The power usage effectiveness (PUE) metric has become an industry standard for reporting energy performance of data centers. However, it is an incomplete metric, failing to address hardware efficiency, energy productivity, and environmental performance. The industry should focus on adopting and systematically reporting more comprehensive metrics, which would allow more insight into data center energy performance.

1. Introduction

The advent of the digital computer brought about the "Information Age," an era in which information has come into its own as a valuable commodity. The accuracy, relevance, and timeliness of an organization's information are—as they have always been—keys to its success. However, the higher speeds, greater traffic, and increased access on the "information superhighway" have made firms hungrier for ever-increasing volumes of data.

Data services lie at the heart of operations for many companies and constitute a core product for others. Google indexes the web to provide search services to users while simultaneously collecting information about search activity to deliver advertising. Social networking entities such as Facebook collect personal data in exchange for hosting virtual communities, aiding interaction among groups. Online retailers like Amazon.com house online inventories and bring buyers and sellers together. Data service providers such as Dropbox and Apple allow users to store their documents, files, and digital content "in the cloud," a distributed storage network. Even traditional retailers like Walmart and Whole Foods use data-intensive processes to manage their inventories in real time.

The hardware and software used to store, transmit, and utilize data to provide e-services are collectively known as information technology (IT) or information and communications technology (ICT). ICT includes computers and software, mobile devices, and communication networks and their components. As digital content has proliferated, so too have the storage mechanisms grown, moving from the lone server to the server closet, the server room, and now the server farm. These storage repositories are collectively known as data centers, which not only provide static storage but also dynamically provide a wide variety of services including hosting web pages and email, streaming multimedia, and running complex business applications like banking management software. Individuals, private firms, universities, and government entities all use data centers of varying scale and complexity to manage the digital information they need to operate.

Expansion of data and services has meant exponential growth in needed storage capacity. IBM reports that global daily data production is 2.5 quintillion bytes and that 90% of the world's data has been produced in the previous two years (IBM, 2013). In addition to the magnitude of data produced, increasing complexity of software has increased its size. As the prevalence of data centers has grown, so have public concerns about their aggregate energy consumption (Markoff and Hansell, 2006). A series of bottom-up estimates has generally found that data centers use on the order of 1–2% of U.S. and global electricity consumption (Koomey, 2008, 2011). An industry report examining the greenhouse gas (GHG) emissions attributable to ICT estimates that, while data center carbon emissions are a negligible part of the global total, they have grown at 8.6% annually since 2001 and will continue to outpace both the global footprint growth rate and the that of other ICT subcategories (networks and end user devices) through 2020 (BCG, 2012).

In response to these concerns, and also to get a better handle on operational costs, the industry has worked to establish metrics to assess data center performance. The most prevalent of these is

* Corresponding author.
E-mail address: nch@cmu.edu (N. Horner).

http://dx.doi.org/10.1016/j.tej.2016.04.011
1040-6190/

power usage effectiveness (PUE), which is the ratio of facility-wide power consumption to power used by the IT equipment. This article discusses the main problems associated with this metric. We first provide a brief background on data center types, as different metrics may be more or less relevant to different sizes and applications of data centers (Section 2). Section 3 defines PUE and highlights its main drawbacks as an overall measure of data center energy performance. Alternative metrics are discussed in Section 4, while Section 5 briefly summarizes different ways in which data center energy performance—once assessed with an appropriate metric—can be improved. Finally, Section 6 summarizes the discussion.

2. Background: data center taxonomy

What is a data center? Lawrence Berkeley National Laboratory (LBNL) describes a data center as a special-purpose facility with the following characteristics and functions (LBNL, 2013):

- "Houses various equipment, such as computers, servers (e.g., web servers, application servers, database servers), switches routers, data storage devices, load balancers, wire cages or closets, vaults, racks, and related equipment.
- Store[s], manage[s], processe[s], and exchange[s] digital data and information;
- Provide[s] application services or management for various data processing, such as web hosting internet, intranet, telecommunication and information technology."

While the LBNL definition is focused on facility contents and function, ICT consulting firm Gartner uses a definition more focused on the organizational role the data center plays, defining a data center as "the department of an enterprise that houses and maintains the back-end [IT] systems and data stores—its mainframes, servers, and databases. In the days of large, centralized IT operations, this department and all the systems resided in one physical space." (Gartner, 2013)

These definitions make no statements about size, and include few restrictions on function. Data centers can range from small server closets to huge server farms and can host a variety of IT services, such as corporate email and filesystems, data archives, and cloud services. The main criterion for a data center seems to be that it houses "back-end" ICT equipment—equipment accessed indirectly by users via a network.

The variety of size and function can make clear classification difficult, though most classification systems rely on some combination of size, criticality of service, and service type. We adopt the following nomenclature for context, loosely based on the taxonomy used by IDC (Patterson, 2010):

1. **Server closets**, or "ad-hoc" data centers, support small businesses or individual projects at larger companies. They may get some support from a corporate-level IT department but may also be configured and operated by non-experts.
2. **Server rooms** are small data centers that support small businesses or special groups or projects of larger entities. They may be administered by central IT staff or "owned" by each project or division.
3. **Localized data centers** provide business-critical applications and have some power and cooling redundancy, though downtime is not catastrophic. Restoration of service on the order of hours is acceptable.
4. **Mid-tier data centers** are large-to-medium-size data centers used to host enterprise-wide applications in support of operations or human resources (e-mail accounts, filesystems, internal data). The data is critical, but incidental to the primary business line. Downtime lasting longer than a few minutes has significant impact on the business. These facilities are operated by the company's central IT department.
5. **Enterprise data centers** are large facilities used, usually by non-ITC companies, in support of core business operations (e.g., banks, health care companies, etc.). These data centers are often in special-purpose facilities and operated under a separate business unit or division. Downtime is catastrophic, and these facilities have highly redundant infrastructure.
6. **Hyperscale data centers, server farms, or warehouse-scale computers (WSCs)** are the very large data centers, usually constructed in their own physical plants, built by ICT companies with a primary business line focused on data (e.g., Google, Apple, Facebook, Amazon, et al.) and, increasingly, cloud-based services. Barroso and Hölzle (2009) coined the term warehouse-scale computer to emphasize the distinguishing large economies of scale, extreme parallelism, hardware and software homogeneity, and aggressive focus on efficiency of these data centers.

Generally, size, infrastructure redundancy, quality of service, and criticality increase as one moves down the list, though these distinctions are necessarily qualitative and somewhat fuzzy in nature. Note that these data center types can be deployed in very different domains, ranging from corporate entities to university- and research-based enterprises.

The Uptime Institute, a corporate data center consultancy, has published a data center classification focused on infrastructure redundancy, ranging from "basic" to "fault-tolerant" (Turner et al., 2008; C7 Data Centers, 2012). Tier I data centers have no redundant systems, whereas Tier IV facilities have duplicate active power and cooling distribution paths, with redundant components on each, so that the center can withstand any single equipment failure. The tiers roughly follow the functional/size-based classification above, with Tier I & II data centers being appropriate for businesses that have no or low obligated quality-of-service requirements and Tier III and IV data centers being appropriate for businesses that need to deliver round-the-clock services with serious consequences for downtime.

When analyzing the energy performance of a data center, it can be difficult to draw a clear boundary around the system, particularly if it is part of a larger multiuse building, as is often the case. Generally the data center encompasses the computing load and associated equipment in such cases, not the entire building. However, if cooling, lighting, and HVAC systems are shared between the data center and spaces dedicated to other uses (e.g., offices), which is especially common in the smallest data centers, then it can be hard to accurately assess energy performance. Inconsistency in what is included in the measurements can make comparisons among different data centers difficult.

3. Measuring energy efficiency: PUE

Estimates of aggregate energy usage provide context for determining if data center energy consumption is a cause for concern. However, these sorts of studies cannot reveal *how* that energy is being used—that is, how efficiently do data centers use energy to deliver the services they provide? To assess energy performance, we must first define a suitable metric.

There have been many different metrics proposed by various industry and research organizations. One review paper cited no less than 30 metrics proposed by various organizations to measure different aspects of data center energy efficiency (Jamalzadeh and Behravan, 2011). However, the industry has converged on power usage effectiveness, or PUE, as the standard of choice (Avelar et al., 2012).

3.1. PUE definition

PUE is the ratio of total power used by the data center facility to the power used by the IT equipment: *PUE = total facility power/IT equipment power.*

Thus, lower PUE values are better, with 1.0 being ideal. A PUE of 1.0 would indicate that 100% of the power delivered to the facility is used by the computing equipment. Power used for lighting, cooling, and other overhead increases PUE. Some have argued that 1.0 is not necessarily a minimum value, as use of recovered waste heat could enable PUE ratings of less than 1.0 (Hamilton, 2009)—although this is perhaps stretching the definition of PUE to encompass energy performance issues beyond its intended application.

Companies like Google and Facebook have aggressively reduced their PUE in recent years through a focus on efficiency and custom hardware design. Google reports a trailing 12-month (TTM) fleet average PUE of 1.12, with individual site ranges from 1.09 to 1.31 (Google, 2015). Facebook does not report a fleet average, but provides dashboards showing real-time PUE measurements for two of its largest data centers, which report TTM averages of 1.08 and 1.09 (Facebook, 2016a,b).

Businesses where data centers are a more ancillary part of operations likely have higher PUEs, and small-to-medium data center operators are less likely to focus intensely on energy efficiency of their installations (Delforge and Whitney, 2014; BCG 2012). According to an Uptime Institute survey, among firms with fewer than 1000 servers, only 50% of operators were concerned with PUE; among firms with over 5000 servers, that number was 90% (Patrizio, 2013). The smaller firms do not have the resources or expertise to achieve the ultra-low PUE benchmarks set by the big firms, and the lower financial rewards put data center greening off the radar of executive leadership. Reported industry-wide average PUE measures, based on survey data, vary between 1.8 (Miller, 2011) and 2.9 (Digital Realty Trust, 2013). The latter survey reported that only 20% of data centers had PUEs of less than 2.0.

Reporting of PUE has become the *de facto* standard—or perhaps more than simply *de facto*, as the EPA has used PUE as the basis for its ENERGY STAR for Data Centers rating program (Miller, 2009)—and has led to a sort of arms race among large data center operators. We now discuss significant problems with over-reliance on this metric.

3.2. PUE as a measure of facility and equipment efficiency

Facility efficiency is exactly what PUE is designed to measure. PUE is improved by reducing overhead (e.g., cooling, lighting) in comparison to the compute load (IT equipment). Calculating PUE requires two measures of power use: one in the numerator (total facility power) and one in the denominator (IT equipment power). While the calculation is a simple one, in practice there is variability resulting from differences in measurement points.

The numerator is intended to be measured at the utility meter (Avelar et al., 2012). However, smaller data centers inside multiuse facilities may not be independently metered, while very large data centers may be metered at higher voltages. In the latter case, measuring at the meter will impact PUE by including losses associated with step-down transformers that occur outside the scope of measurement for other data centers.

The denominator can be affected by what is included in the "IT equipment" calculation. Though they are viewed as infrastructure, components like cooling fans and power supply units may be either counted as IT equipment or not, depending on whether they are housed internally or externally to servers. To get a "correct" value for IT equipment energy, measurements would need to be taken at the component level: CPU and other integrated circuits, memory, disks, etc. (Hamilton, 2009). The fact that such measurements are impractical means that equipment efficiency can muddy the PUE calculations. A low-overhead facility running older, less efficient servers could conceivably achieve a low PUE while still using more energy than it needs. In practice, this may not be much of a concern for larger data centers that refresh their server stock frequently. However, data centers occupying the lower tiers of the spectrum defined in Section 2 may need to look beyond PUE at the hardware they use.

The variation in how the IT equipment "box" is bounded means that PUE measures may not be directly comparable and provides opportunities for organizations to game the rating. Furthermore, actions that improve energy efficiency can perversely increase PUE—reducing IT load through virtualization, for instance (Cole, 2011).

3.3. PUE as a measure of business operational efficiency

It is worth noting that energy efficiency is not the most important measure of effectiveness for data centers, which necessarily prioritize access time, availability, or other such "quality-of-service" metrics. In the early rush to build out storage capacity, energy efficiency was not a primary concern for data center operators. A focus on performance, server uptime, and hardware costs by IT engineers who purchased the equipment left operating costs as an afterthought. Even at companies like Google, an aggressive leader in data center energy efficiency, "it was clear the only way to make [search] work as [a] free product was to run on relatively cheap hardware" according to Urs Hoelzle, the company's vice president of operations (quoted in Shankland, 2009).

However, the growth of data centers means that their operating budgets are an increasing part of overall corporate spending. A heightened focus on efficiency has led to declining PUEs. Unfortunately, for companies focused on improving operational efficiency, PUE says nothing about how well energy is being translated into the services or products the organization delivers. A company operating a server farm with a very low PUE but without an optimal allocation of the computing load to the hardware may be operationally inefficient. Such inefficiencies can result from excessive redundancy in the system or underutilization of the hardware (Klaus, 2012; Armbrust et al., 2009; Kaplan et al., 2008).

3.4. PUE as a measure of 'greenness'

Another reason to care about data center energy use is in the context of "green" operations. Measures that increase energy efficiency to reduce operating costs also tend to reduce greenhouse gas emissions (GHGs), criteria air pollutants, and impacts on water. McKinsey & Co. (Kaplan et al., 2008) estimated that in 2007 data centers were responsible for 170 Mt CO_2 worldwide and projected emissions to quadruple by 2020. GHG emissions attributable to data centers come from four sources: (1) electricity consumption, (2) onsite combustion, (3) refrigerant use, and (4) embodied emissions (Greenhouse Gas Protocol 2012a). We discuss PUE in the context of (1) and (2); (3) and (4) are important factors in overall data center environmental evaluation but are outside the scope of this article.

Greenpeace's 2012 report, "How clean is your cloud?" noted that GHG emissions associated with electricity supply are important, and claimed that several of the largest data center operators relied heavily on dirty electricity (Cook, 2012a). In addressing data center professionals at the Uptime Symposium, the report's author noted that "PUE is a very useful metric and diagnostic tool, but it is not a good indicator of how green you are. . . . It does not speak to the resources you use in the outside world" (quoted in Miller, 2012).

We now illustrate that PUE does not necessarily correlate with carbon efficiency using a small sample of data center measurements. Most companies that measure their data center performance hold results as proprietary information. Here, we use data from 32 data centers in a series of LBNL benchmarking studies (e.g., LBNL, 2003; Greenberg et al., 2006), supplemented with analysis of two federal data centers owned by the Environmental Protection Agency (Patterson, 2010) and the National Renewable Energy Laboratory (Sheppy et al., 2011) as well as self-reported and estimated data for four WSC-type data centers operated by Apple (two) and Facebook (two).

The LBNL case studies were focused on cooling, but they also reported power consumption broken down by use: computing, HVAC, lighting, and UPS losses. We estimate PUE for each facility by dividing total energy use by energy used for computing. As these case studies were carried out by three different organizations, there is some variation in the data reported and the terminology used, particularly with respect to facility size as distinct from building size. Many, though not all, of the data centers occupy multi-use buildings. We made a best effort to calculate PUE based on only the data center portion of the facility (i.e., the "white space," or the space inside the data center cooling envelope used by the IT equipment).

For Facebook, self-reported PUE values were used for the Forest City, N.C., and Prineville, Ore., data centers. Apple does not publish PUE values, though one data center industry insider claims that Apple's Maiden, N.C., facility comes in at about 1.1 (Roberts, 2013). Interestingly, this article implies that measures like the solar arrays and fuel cell plant at the facility give it "an advantage" in the "constant PUE chase." This false conflation of PUE values with carbon efficiency by a data center professional is troublesome but not all that uncommon. Greenpeace has estimated the PUE for the facility at 1.35 (Cook, 2012b), though there is considerable controversy over its assumptions. We use the 1.1 number for Apple's Maiden facility under the assumption that the Greenpeace estimate is too pessimistic and that Apple is generally competitive with Facebook and Google. We use the 1.35 value for Apple's older Newark, Calif., data center. However, given the rapidly changing landscape and lack of detail in the data provided, it would be best to treat these as abstract warehouse-scale data centers rather than accurate representations of specific Apple and Facebook facilities.

Fig. 1 shows PUE plotted against physical size for these data centers. Unsurprisingly, the large WSC data centers have lower PUE ratings. However, PUE is less correlated with size for other data

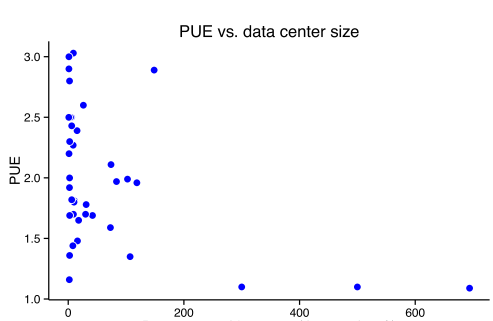

Fig. 1. Data center PUE vs. physical size (white space floor area) for 38 heterogeneous data centers. Data are from LBNL benchmarking studies; one EPA legacy data center; one state-of-the-art data center built at NREL in 2010; and 2012 data from four WSC-type data centers.

center tiers. It is important to note that most of the LBNL case studies are now at least 10 years old, so, although industry reports indicate that average PUE values may still be near 2.0 and that the rate of improvement has stagnated (Digital Realty Trust, 2013; Beltran, 2014), the worst performers may have subsequently adopted best practices and improved their PUE ratings.

Fig. 2 shows that PUE and size are not particularly well correlated with data center energy or carbon performance. The horizontal axis is energy intensity (energy used per square foot); the vertical axis is carbon intensity. Bubble area corresponds to data center size, while bubble shade reflects the PUE of the data center, with lower PUEs (light) being better, and higher PUEs (dark) being worse. Carbon emissions factors, based on the NERC region in which each facility is located, are from Greenhouse Gas Protocol (Greenhouse Gas Protocol, 2012b). The data set includes a wide variation in data center size, type, and use. Facilities are located in 12 states, with over-representation in California and Hawaii.

We note first that energy intensity is highly correlated with carbon intensity. This relationship makes sense, especially given the nature of the calculation in which GHG emissions are calculated directly from energy use. We expect greener data centers to lie along a flatter slope. However, PUE does not generally identify clusters of data centers by either measure of performance. While low-intensity data centers do tend to have poor PUE measures, three facilities with excellent PUE values also show relatively low carbon and energy intensity. The point is made particularly clear by comparing data centers along the 0.5 MWh/sq ft line; data centers with very different PUE ratings have the same energy intensity, while the smallest of these has much worse carbon performance despite having the better PUE value. Fig. 2 demonstrates empirically what Masanet, Shehabi, and Koomey (2013) discuss notionally: namely, that PUE is a poor indicator of overall data center environmental performance. (See, in particular, their Fig. 2 and related discussion.)

4. The quest for a better metric

While these shortcomings of PUE are understood by many in the industry, the metric continues to be emphasized—over 80% of surveyed IT executives track and report it to corporate management (Beltran, 2014)—and many data center operators wrongly report it as a proxy for environmental performance. Industry press has conflated low PUE values with greenness (Roberts, 2013), and even the GHG Protocol draft reporting standard for data centers lists only PUE in its section on calculating operational emissions (Greenhouse Gas Protocol, 2012a). Some governments, notably Amsterdam, a data center hub, have enacted legislation setting maximum PUE standards in a push to be green (Dijkhuis, 2008; Henderson, 2011), while in the U.S., PUE has been adopted as the basis for the ENERGY STAR data center labeling program (Miller, 2009).

The issue is not that these and others are using PUE; PUE is a reasonable and useful metric for assessing the infrastructure efficiency of a data center. The issue is that PUE is far from a comprehensive metric, and quoting PUE values *and nothing else* stops short of providing a truly meaningful sense of data center performance. As a result, the industry has been working on creating additional metrics to round out the picture. In this section, we highlight some of the more promising ideas.

The ultimate energy efficiency metric would measure the amount of *useful computational work* performed per unit of energy used: η_E = *useful computational work/energy consumed*; or, if we are interested in carbon intensity: η_C = *useful computational work/carbonemitted*; or other metrics that represent environmental, health, or other externality concerns.

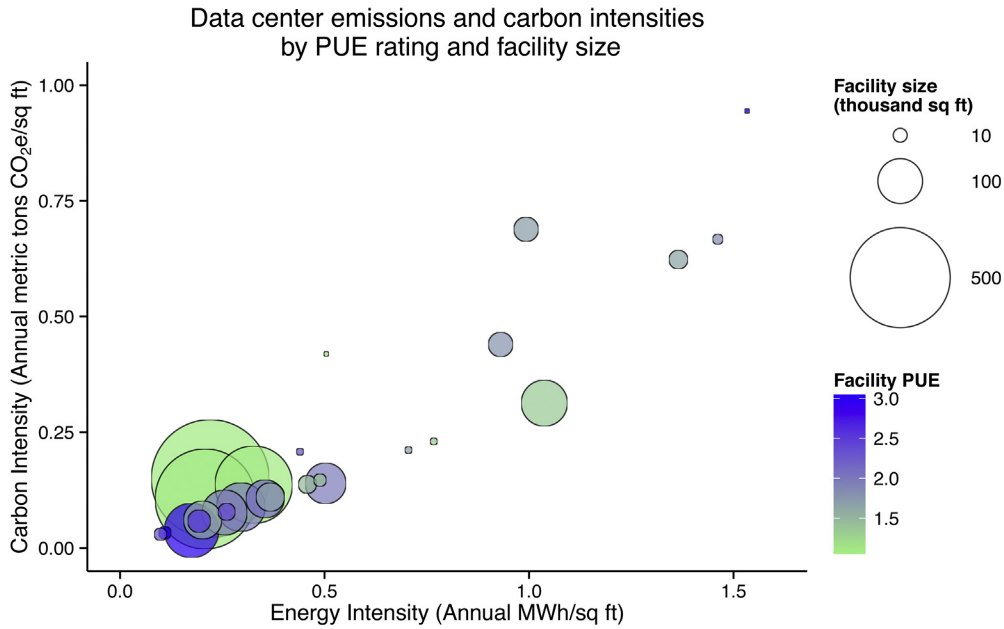

Fig. 2. Annual energy intensity (in MWh/sq ft) vs. GHG emissions intensity (in metric tons CO_2e/sq feet) for 26 data centers in the United States. Bubble area corresponds to data center size square feet; color scale identifies measured PUE. Included data centers are from LBNL benchmarking studies; one EPA legacy data center; one state-of-the-art data center built at NREL in 2010; and 2012 data from four WSC-type data centers. GHG for all data centers except for the WSC facilities are calculated from GHG Protocol grid emissions factors for the year closest to the energy usage report; GHG for the WSC facilities are self-reported by the data center owner. (12 facilities outside the plot bounds, up and to the right, not shown for readability.)

Indeed, the Green Grid has proposed a Data Center Energy Productivity (DCeP) metric along these lines (Anderson et al., 2008). The challenge, of course, is in defining "useful work" in a meaningful, measurable way.

Ebay's Digital Service Efficiency dashboard (which posted current data in 2012 and 2013) (ebay, 2013) hints at the possibilities for productivity metrics, reporting URLs per kWh and revenue per server, per user, and per MWh. The dashboard also reports a PUE of 1.5. Other industry-specific metrics, such as energy per bank transaction, per e-mail, per sale, or per user account, could also be envisioned. The drawback of such special purpose metrics is that they do not allow comparison of data centers that serve different types of loads.

Because of the difficulty in defining a widely applicable metric based on useful work, other intermediate metrics have been proposed to address aspects of energy performance not included in PUE—in essence, treating data center productivity as a composite of other metrics. A white paper from the Uptime Institute defines four different factors in data center efficiency (Stanley et al., 2007):

1. Physical site infrastructure overhead involves facility siting, design, construction, and operation.
2. IT hardware energy efficiency measures the efficiency of individual hardware units (drives, chips, power supplies, etc.).
3. IT hardware asset utilization deals with making sure capacity is fully utilized in operation
4. IT strategy is the overall approach to data center needs during business planning.

PUE measures only the last of these.

Barroso and Hölzle (2009) factor DCeP into three individual components: $1/PUE \times 1/SPUE \times computation/energy$ *to electronic components*.

The first term accounts for facility efficiency (Item #1 in the Uptime Institute list above) as PUE. The second term addresses part of IT hardware efficiency (#2) by creating a new metric, Server PUE (SPUE), which accounts for the overhead of infrastructure at the server level—e.g., power supplies and cooling fans.

Note that different data center architectures blur the line between facility infrastructure and server infrastructure by aggregating cooling fans and power supplies outside of each server box or, going the other way, distributing power backup batteries at the server level. Since it is not always clear on which side of the "IT boundary" components reside, Total-power Usage Effectiveness (TUE), as the product of PUE and ITUE (or SPUE), has been proposed (Patterson et al., 2013).

TUE, comprising the first two terms the equation above, addresses overhead in the energy path from the meter to the IT equipment. The third term addresses the efficiency of the IT equipment itself: processors, memory, and other electronic components. Just as hardware can be benchmarked for performance, several protocols, such as JouleSort (Rivoire et al., 2007) have been developed to benchmark energy use by running predetermined workloads on the hardware while measuring energy consumption.

With such benchmarks, we can get some sense of the "absolute" or "theoretical" energy efficiency of the data center as constructed and provisioned. However, energy efficiency of a data center is heavily affected by how the facility is used operationally. Referring to the Uptime Institute list above, Items #3 and #4 refer to what might be called operational and strategic overhead. Both have to do principally with server utilization. Because data center capacities are often sized for peak load to ensure high quality-of-service, most of the time servers are severely underutilized. Since operators are concerned about handling load spikes, substantial safety margins even beyond the observed peak load are common (Katz et al., 2011). A measurement of 5,000 Google data centers revealed utilization of 60% or less 95% of the time and of 30% or less half of the time (Barroso and Hölzle, 2009). Importantly, Google represents the upper end of utilization; most data centers will see much lower utilization rates. The issue with underutilization is that servers are not energy-proportional—that is, energy usage does not scale linearly with computing load. An x86 server at idle consumes almost 50% of its peak power usage (Barroso and Hölzle, 2009, Fig. 5 .8). Therefore, idle power consumption can be reduced in two ways: reducing the number of idle servers, and working towards energy-proportional hardware. We further discuss methods for increasing utilization in the next section.

While PUE, ITUE, and TUE are important components of data center energy efficiency, we believe that a true data center productivity or efficiency metric should incorporate a measure of utilization.

All of these metrics by definition focus on energy efficiency, which is correlated with both greenness and cost savings. However, if the ultimate goal is to measure environmental footprint, energy efficiency metrics are imperfect. As a supplement to PUE, some companies do focus on emissions, and there are “sustainablility metrics” designed to evaluate carbon performance, such as Carbon Usage Effectiveness (CUE), which is essentially PUE multiplied by a carbon emission factor (Belady, 2010): *CUE = CO_2 emissions from total facility energy consumption/IT equipment energy*.

eBay reports 0.56 t of carbon per MWh on its dashboard. Other companies, such as Facebook and Apple, report carbon usage to some extent in annual environmental footprint statements. Data centers with the same PUE or DCeP could have vastly different CUEs. For instance, consider Apple’s data center in Maiden, N.C., which is co-located with the country’s largest non-utility-owned solar farm and largest non-utility-owned fuel cell plant (Apple, 2014). The very same data center would have a much different carbon performance if it were powered from grid electricity. Another proposed metric is the Green Energy Coefficient, or GEC, which is the ratio of green energy to total energy consumed by the data center (Global Taskforce, 2012).

Primary power reliability also plays a role in a data center’s operational carbon emissions, since interruptions in the main supply are generally met with diesel generator backups. In Northern Virginia, another data center hub, the aggregate capacity of diesel generators is nearly equivalent to that of a nuclear power plant (Glanz, 2012).

Furthermore, evaluating the environmental footprint of a data center requires a lifecycle-cost analysis (LCA) approach to include the embodied emissions of the facility and IT equipment (Raghavan and Ma, 2011), although these carbon emissions are likely small compared to operational emissions (Masanet et al., 2013).

Because the first step in solving a problem is defining it, creating useful, precise metrics is the important first step in improving data center energy performance. Precise metrics help highlight the aspects of the system where investment is most likely to improve performance. We now turn to a brief discussion of some of these target areas.

5. Improving performance

There have been many suggested strategies for improving the energy and carbon performance of data centers, but they fit into the broad categories of *equipment utilization*,*building efficiency*, *equipment efficiency*, and *power sourcing*.

5.1. Equipment utilization

Server utilization in data centers is generally very low. Because uptime, reliability, and fulfillment of service level agreements are the priorities of data center operators, data centers are generally built with extreme redundancy: “McKinsey & Company analyzed energy use by data centers and found that, on average, they were using only 6% to 12% of the electricity powering their servers to perform computations. The rest was essentially used to keep servers idling and ready in case of a surge in activity that could slow or crash their operations” (Glanz, 2012).

One way to reduce server idling is to put servers to sleep when they are not being used. The central concern with such a scheme is that latency involved with waking servers up to meet spiking loads will decrease quality-of-service. However, drawing an interesting parallel between computing loads in data centers and power loads on the electricity grid, Katz et al. (2011) believe that a program of waking and sleeping servers based on a predictable “base load” while maintaining a “spinning reserve” to provide headroom for stochastic, bursty components of the load will work for many types of services. Other researchers, such as Gandhi et al. (2013), are working on formal methods to optimize server provisioning for different load parameters. These sorts of methods could provide data center operators more confidence to sleep or shut down idling servers without risking service quality.

In addition to the designed overhead, there is also an issue of “comatose” servers—those that are no longer needed but are still running because no one can positively determine that the data is old or wants to risk pulling the plug: “[a]necdotal evidence indicates that 10–30% of servers in many data centers are using electricity but no longer delivering computing services. These servers have not yet been decommissioned and are probably not counted in installed base statistics. In many facilities nobody even knows these servers exist….” (Koomey, 2011, 7). Figures for comatose servers may be even higher. A sample at a LexisNexis data center revealed that three-fourths of installed servers used, on average, 10% of their capacity (Glanz, 2012). Both intentional overcapacity and failure to consolidate and decommission old equipment result in very low utilization rates.

In addition to retiring comatose servers, virtualization, colocation, and moving to the cloud have been shown to increase utilization rates. A recent industry survey revealed the following findings regarding migration to the cloud (Stansberry and Kudritzki, 2012):

- Large companies are more likely to pursue cloud computing than small companies, which is interesting because the cloud should be appealing to smaller entities that do not want the overhead associated with running a data center.
- Adoption of cloud computing is heavily skewed towards technology service providers. Traditional large vertical enterprises are much less likely to use the cloud.
- Companies reluctant to move to the cloud cite security, compliance, and reliability concerns.

5.2. Building efficiency

Putting data centers in facilities specifically designed for them improves their efficiency. Standard practices include physical separation of the “hot” and “cold” aisles between server racks, raised floors, and carefully designed cabling conduits. An important contributor to efficiency is initial site selection: data centers in cooler climates or unique locations (e.g., underground or near water sources) allow for “free cooling.” Some advanced data centers reduce the power loss by distributing DC current at the facility level, rather than converting AC to DC in each individual power supply unit.

5.3. Equipment efficiency

Most studies seem to identify other sources of efficiency gains as more important than addressing the efficiency of the IT hardware itself. However, the EPA EnergySTAR program has issued product specifications for enterprise servers, uninterruptable power supplies (UPSs), data center storage, and network equipment (http://www.energystar.gov/products/specs/). Furthermore, if energy-proportional equipment could be achieved, utilization rates would no longer matter, since power usage would scale linearly with load.

5.4. Power sourcing

The source of a data center's electricity has a large effect on its GHG and criteria air pollutant emissions, much like other industrial users. While Apple may not be focusing on driving down PUE to the extent that Google and Facebook are, the company emphasizes its clean energy sources for its data center electricity. If data center operators value a low carbon footprint, they must address the emissions of the power supply either by siting in areas where the grid has a low emissions factor, or by obtaining a separate source for cleaner power.

Which of these four areas provide the most benefit? Masanet et al. (2013) suggest that efficiency potential on the IT side (measured by ITUE and utilization) is larger than that on the infrastructure side (PUE). They also argue that in an environment where renewable energy is limited from the grid, data centers may just be displacing other consumers of clean energy and should therefore focus on improving efficiency rather than on power sourcing.

Several reports (e.g., NRDC, 2012; Stansberry and Kudritzki, 2012) have mentioned that utilization rate is more important than the efficiency of the equipment itself and is where the biggest immediate gains can be made. The NRDC report compared on-premise data centers to the cloud environment with respect to carbon emissions. The model uses PUE, server utilization rate, server refresh period, virtualization ratio, and grid emissions factor estimated for several typical data center deployment scenarios. The study found that the most impact could be gained by targeting server utilization and electricity source emissions factor, and only then by improving infrastructure efficiency. Note that neither of these proposed focus areas are measured by PUE. Perversely, a decision to reduce energy consumption by shutting down idle servers in an existing data center would likely increase (worsen) the facility's PUE (Cook, 2012a). Virtualization and moving to the cloud significantly increase efficiency and reduce GHG emissions. Regardless of whether utilization, equipment, or the facility infrastructure is being targeted for improvement, configuring an energy-optimal data center requires three "pillars": tracking real time environmental and performance variables in data centers and maintaining an accurate record of equipment inventories; following established procedures or best practices for operational management; and experimenting, using simulation software, to understand how configuration changes interact with physical principles to affect performance metrics (Koomey, 2016).

6. Conclusions and discussion

Power usage effectiveness (PUE) has become the industry standard for reporting data center energy performance. While it is a useful measure of facility overhead, it is an incomplete metric. From an energy efficiency standpoint, it does not account for the efficiency of the computing hardware. From a business standpoint, PUE does not measure energy productivity. Finally, from an environmental standpoint, it does not measure energy performance with respect to carbon emissions.

Based on these drawbacks, the industry's past focus on PUE is misplaced. While data center operators should of course continue to adopt best practices related to facility power and cooling, they would be better served by pursuing measures to increase utilization rates, reduce redundancy, and source clean power rather than continuing the race to publish marginally lower PUE numbers. Indeed, the technology leaders in this industry are doing so.

The lack of specific data about how servers are being configured and used as well as specific, benchmarked energy usage for certain types of equipment such as storage and network devices is a significant research issue. This lack of data is a recurring theme in the "future work" sections of literature and was confirmed in a conversation with EPA's products manager for the Energy Star labeling for data center equipment. To deal with this want of data, two strategies are observed in the literature. Most energy use assessments with the goal of estimating an "absolute" value for total energy consumption use a bottom-up approach where variables like server utilization, PUE, and server stocks are estimated using best-available data and expert input (Koomey, 2008). These figures are fed into relatively simple calculations to render an estimate for overall sector energy consumption. Few studies, however, report uncertainty ranges on their results, though several parameterize different scenarios in their models.

Alternatively, some researchers (NRDC, 2012) look at the relative performance of different data center configurations. This approach eliminates the reliance on estimates for server stocks, though estimates for equipment performance, utilization rates, and other parameters are still necessary. These studies do not attempt to calculate absolute energy consumption, but rather to show the magnitude of relative gains that can be made.

The first approach estimates energy use for the aggregate fleet but does not target energy efficiency, while the second evaluates potential efficiency gains for generic systems but does not provide insight on what is deployed and how it is used sector-wide.

We believe this data gap could be addressed by establishing a framework under which data center operators report general data center characteristics and performance metrics. Reported data could be sufficiently anonymized to alleviate security concerns while providing the industry and researchers alike a more complete picture of current performance. Though results from voluntary reporting would likely be skewed towards high-efficiency data centers, such results would still provide a broader window into the data center fleet than what is currently publicly available. Indeed, several governments have programs in place for data center reporting, either under voluntary certification programs or carbon reduction mandates.

The U.S. EPA's ENERGY STAR voluntary data center certification program lists 73 certified facilities, though minimal information about each facility is reported (ENERGY STAR, 2016). The U.S. also has several programs to collect data on federally owned data centers through the DOE Sustainability Performance Office and the Federal Data Center Consolidation Initiative, though these data are not publicly available. The National Australian Built Environment Rating System, which allows building operators to rate and certify their facilities, has a data center category, but only nine data centers are currently reported in the database (NABERS, 2016).

Under the UK's Carbon Reduction Commitment (CRC), large consumers of electric power, including data centers, are required to baseline and report electricity usage annually. The Environment Agency converts electricity usage to CO_2 emissions, which are reported for each firm (Riley, 2014). Data centers are not easily distinguished from other types of facilities, nor are any other performance metrics or characteristics published. The European Commission's Joint Research Centre Institute for Energy and Transport has developed the European Code of Conduct for Energy Efficiency in Data Centres (JRC, 2014), a voluntary program. Participation in the CoC includes a reporting form with useful characteristic and performance data; however, the database of participants is not publicly available.

Given that lack of data seems to be a recurring complaint in the policy and research arenas, it would behoove these communities to come together with industry to establish a reporting framework. While measurement of data center performance has important nuances, and legitimate security and competitiveness concerns exist, the technical and institutional barriers to creating such a data set do not seem insurmountable.

Acknowledgments

The authors thank Dale Sartor, Arman Shehabi, and Jonathan Koomey for helpful discussions that aided in the framing of this article and Rod Mahdavi of LBNL for providing benchmarking data. This work was funded in part by the Center for Climate and Energy Decision Making (SES-0949710 and SES-1463492), through a cooperative agreement between the National Science Foundation and Carnegie Mellon University, and by the Department of Engineering and Public Policy. The authors gratefully acknowledge the NSF and Doug Sicker of EPP for their financial support.

References

Armbrust, M., Fox, O., Griffith, R., Joseph, A.D., Katz, Y., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., Zaharia, M., 2009. Above the clouds: A Berkeley View of Cloud Computing. Online: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.361.8444.

Anderson, Dave, Cader, Tahir, Darby, Tommy, Gruendler, Nick, Hariharan, Rema, Holler, Anne, Lindberg, Chris, Long, Christine, 2008. A Framework for Data Center Energy Productivity. White Paper WP-13. The Green Grid. . http://lpis.csd.auth.gr/smartihu/storage/GreenGridProxies.pdf.

V, Avelar, Azevedo, Dan, French, Alan, 2012. PUE: A Comprehensive Examination of the Metric. White Paper WP-49. The Green Grid. . http://www.thegreengrid.org/?/media/WhitePapers/WP49-PUE%20A%20Comprehensive%20Examination%20of%20the%20Metric_v6.pdf?lang=en.

Barroso, Luiz André, Hölzle, Urs, 2009. The datacenter as a computer: an introduction to the design of warehouse-scale machines. Synth. Lect. Comput. Archit. 4 (1), 1–108. doi:http://dx.doi.org/10.2200/S00193ED1V01Y200905CAC006.

Belady, Christian, 2010. Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric. White Paper. The Green Grid. . http://www.thegreengrid.org/?/media/WhitePapers/Carbon%20Usage%20Effectiveness%20White%20Paper_v3.pdf?lang=en.

Beltran, J., 2014. Data Center Industry Survey Uptime Institute eJournal. Online: https://journal.uptimeinstitute.com/2014-data-center-industry-survey/.

BCG, 2012. GeSI SMARTer: The Role of ICT in Driving a Sustainable Future (Boston Consulting Group). Online: http://gesi.org/portfolio/report/72.

C7 Data Centers, 2012. "C7 Tier Classifications." C7 Data Centers. . (Co., McKinsey & 2008. Revolutionizing Data Center Efficiency) http://www.c7dc.com/articles/tier-classifications/.

Cole, D., 2011. Data center energy efficiency–looking beyond PUE (No Limits Software). Online: http://www.nolimitssoftware.com/docs/DataCenterEnergyEfficiency_LookingBeyond.pdf.

Cook, Gary, 2012a. How Clean Is Your Cloud? Greenpeace International. . http://www.greenpeace.org/international/Global/international/publications/climate/2012/iCoal/HowCleanisYourCloud.pdf.

Cook, Gary, 2012b. A Clean Energy Road Map for Apple. JN 417 UPDATE. Greenpeace International. http://www.greenpeace.org/international/Global/international/publications/climate/2012/iCoal/Apple_Clean_Energy_Road_Map.pdf.

Delforge, P., Whitney, J., 2014. Data Center Efficiency Assessment (Natural Resources Defense Council). .

Digital Realty Trust 2013. North America Campos Survey Results https://c.na6.content.force.com/sfc/dist/version/download?oid=00D300000005uRq&ids=06880000000k573&d=/a/80000000CpC7/k_RJOcsv31zvPC4hgEz9NMQjNd0m4KjS_CzGO5_ni48%3D&viewId=05H80000000EDGV&operationContext=DELIVERY.

Dijkhuis, Pieter, 2008. Government Regulation for Datacenters in Amsterdam. LeaseWeb. . http://blog.leaseweb.com/2008/03/30/government-regulation-for-datacenters-in-amsterdam/.

ebay, 2013. Digital Service Efficiency ebay.com. Online: http://tech.ebay.com/dashboard.

ENERGY STAR, 2016. ENERGY STAR Certified Data Centers. Online: https://www.energystar.gov/index.cfm?fuseaction=labeled_buildings.showDataCenters.

Facebook, 2016a. Forest City Data Center—PUE/WUE. Facebook. . https://www.facebook.com/ForestCityDataCenter/app/288655784601722/.

Facebook, 2016b. Prineville Data Center—PUE/WUE Facebook. . https://www.facebook.com/PrinevilleDataCenter/app/399244020173259/.

Gandhi, Anshul, Doroudi, Sherwin, Harchol-Balter, Mor, Scheller-Wolf, Alan, 2013. Exact analysis of the M/M/k/setup Class of Markov Chains via Recursive Renewal Reward. Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems, Pittsburgh, PA, pp. 153–166. http://dl.acm.org/citation.cfm?id=2465760.

Gartner, 2013. Gartner IT Glossary—Data Center. . http://www.gartner.com/it-glossary/data-center/.

Glanz, J., 2012. Data Centers Waste Vast Amounts of Energy, Belying Industry ImageThe New York Times. . (Sept. 23) http://www.nytimes.com/2012/09/23/technology/data-centers-waste-vast-amounts-of-energy-belying-industry-image.html.

Global Taskforce, 2012. Harmonizing Global Metrics for Data Center Energy Efficiency. White Paper. . http://iet.jrc.ec.europa.eu/energyefficiency/sites/energyefficiency/files/files/documents/ICT_CoC/harmonizing_global_metrics_for_data_center_energy_efficiency_2012-10-02.pdf.

Google, 2015. Efficiency: how we do it. Google—Data Centers, . https://www.google.com/about/datacenters/efficiency/internal/.

Greenberg, Steve, Mills, Evan, Tschudi, Bill, Rumsey, Peter, Myatt, Bruce, 2006. Best practices for data centers: lessons learned from Benchmarking 22 data centers. Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA, ACEEE, August 3: 76–87. http://energy.lbl.gov/EA/mills/EMills/pubs/pdf/aceee-datacenters.pdf.

Greenhouse Gas Protocol, 2012a. Guide for assessing GHG emissions of data centers. GHG Protocol Product Life Cycle Accounting and Reporting Standard ICT Sector Guidance, . http://www.ghgprotocol.org/files/ghgp/Chapter_8_GHGP-ICT%20Data%20Center%20guide%20v2-3%2010MAR2012.pdf.

Greenhouse Gas Protocol, 2012b. GHG Emissions from Purchased Electricity. . http://www.ghgprotocol.org/files/ghgp/tools/GHG-emissions-from-purchased-electricity(Version-4_4_Aug-2012).xlsx.

Hamilton, James, 2009. PUE and total power usage efficiency (tPUE). blog. Perspectives, . http://perspectives.mvdirona.com/2009/06/15/PUEAndTotalPowerUsageEfficiencyTPUE.aspx.

Henderson, Nicole, 2011. EvoSwitch Promotes Amsterdam as European Data Center Gateway. Web Host Industry Review. . http://www.thewhir.com/web-hosting-news/evoswitch-promotes-amsterdam-as-european-data-center-gateway.

IBM, 2013. What Is Big Data? IBM. http://www-01.ibm.com/software/data/bigdata/what-is-big-data.html.

JRC, 2014. Data centres energy efficiency. European Commission Joint Research Centre Institute for Energy and Transport, . http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/data-centres-energy-efficiency.

Jamalzadeh, Morteza, Behravan, Navid, 2011. An exhaustive framework for better data centers energy efficiency and greenness by using metrics. Indian J. Comput. Sci. Eng. 2 (6), 813–822. http://www.doaj.org/doaj?func=fulltext&aId=920933.

Kaplan, J.M., Forrest, W., Kindler, N., 2008. Revolutionizing Data Center Energy Efficiency (McKinsey & Company). Online: http://www.sallan.org/pdf-docs/McKinsey_Data_Center_Efficiency.pdf.

Katz, Randy H., David, E. Culler, Sanders, Seth, Alspaugh, Sare, Chen, Yanpei, Dawson-Haggerty, Stephen, Dutta, Prabal, et al., 2011. An information-centric energy infrastructure: the Berkeley view. Sustain. Comput.: Inf. Syst. 1 (1), 7–22. doi:http://dx.doi.org/10.1016/j.suscom.2010.10.001.

Klaus, Jeffrey S. 2012. Data center energy: past, present and future (Part One). Article. Data Center J. http://www.datacenterjournal.com/it/data-center-energy-past-present-and-future-part-one/.

Koomey, J.G., 2008. Worldwide electricity used in data centers. Environ. Res. Lett. 3 (3), 034008. http://iopscience.iop.org/1748-9326/3/3/034008.

Koomey, J.G., 2011. Growth in Data Center Electricity Use 2005–2010Analytics Press, Oakland, CA. . (August 1: 2010) http://www.twosides.us/Content/rsPDF_218.pdf.

Koomey, J.G., 2016. Three Pillars of Modern Data Center Operations Data Center Knowledge. Online: http://www.datacenterknowledge.com/archives/2016/02/02/three-pillars-modern-data-center-operations/.

LBNL, 2003. Benchmarking: Data Centers—Case Study Reports. High-Performance Buildings for High-Tech Industries. . http://hightech.lbl.gov/dc-benchmarking-results.html.

LBNL, 2013. Definition of Data Center. Data Center Energy Management. . http://hightech.lbl.gov/dctraining/definitions.html.

Markoff, John, Hansell, Saul, 2006. Hiding in Plain Sight, Google Seeks More PowerThe New York Times. . (June 14) http://www.nytimes.com/2006/06/14/technology/14search.html.

Masanet, E., Shehabi, Arman, Koomey, Jonathan G., 2013. Characteristics of low-carbon data centres. Nat. Clim. Change 3, 627–630. doi:http://dx.doi.org/10.1038/NCLIMATE1786.

Miller, Rich, 2009. EPA to Use PUE in Data Center Energy Star. Data Center Knowledge. . http://www.datacenterknowledge.com/archives/2009/04/22/epa-to-use-pue-in-data-center-energy-star/.

Miller, Rich, 2011. Uptime Institute: The Average PUE Is 1.8. Data Center Knowledge. . http://www.datacenterknowledge.com/archives/2011/05/10/uptime-institute-the-average-pue-is-1-8/.

Miller, Rich, 2012. Apple: iDataCenter Power Will Be 100% Green. Data Center Knowledge. . http://www.datacenterknowledge.com/archives/2012/05/17/apple-idatacenter-power-will-be-100-green/.

NABERS, 2016. Rating Register. NABERS. http://www.nabers.gov.au/public/WebPages/ContentStandard.aspx?module=30&template=3&id=310&side=AssessorTertiary.htm.

NRDC, 2012. The Carbon Emissions of Server Computing for Small- to Medium-Sized Organizations: A Performance Study of On-Premise Vs. The Cloud. Natural Resources Defense Council. . http://www.wspenvironmental.com/media/docs/ourlocations/usa/NRDC-WSP_Cloud_Computing.pdf.

Patrizio, Andy, 2013. Data Center Managers Worn Out by PUE Chase. Slashdot. http://slashdot.org/topic/datacenter/data-center-managers-worn-out-by-pue-chase/.
Patterson, Michael K., Poole, Stephen W., Hsu, Chung-Hsing, Maxwell, Don, Tschudi, William, Coles, Henry, Martinez, David J., Bates, Natalie, 2013. TUE, a new energy-Efficiency metric applied at ORNL's jaguar. Supercomputing. Springer, pp. 372–382. http://link.springer.com/chapter/10.1007/978-3-642-38750-0_28.
Patterson, Michael K., 2010. The Green Grid EPA Data Center Assessment. . http://www.thegreengrid.org/~/media/TechForumPresentations2010/EPA_Data_Center_Assessment_Report.pdf?lang=en.
Raghavan, Barath, Ma, Justin, 2011. The energy and emergy of the internet. Proceedings of the 10th ACM Workshop on Hot Topics in Networks, 9 . http://dl.acm.org/citation.cfm?id=2070571.
Riley, Beth, 2014. CRC Annual Report Publication. UK Environment Agency. http://www.environment-agency.gov.uk/business/topics/pollution/146938.aspx.
Rivoire, Suzanne, Shah, Mehul A., Ranganathan, Parthasarathy, Kozyrakis, Christos, 2007. JouleSort: a balanced energy-efficiency benchmark. Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data, ACM, pp. 365–376. http://dl.acm.org/citation.cfm?id=1247522.
Roberts, Tom, 2013. Are You Suffering from PUE Envy? Data Center Knowledge. http://www.datacenterknowledge.com/archives/2013/05/29/are-you-suffering-from-pue-envy/.
Shankland, Stephen, 2009. Google Uncloaks Once-Secret Server. CNET. http://news.cnet.com/8301-1001_3-10209580-92.html.
Sheppy, M., Lobato, C., Van Geet, O., Pless, S., Donovan, K., Powers, C., 2011. Reducing Data Center Loads for a Large-Scale, Low-Energy Office Building: NREL's Research Support Facility (Book). National Renewable Energy Laboratory (NREL), Golden, CO. http://www.osti.gov/scitech/biblio/1031393.
Stansberry, M., Kudritzki, J., 2012. Data Center Industry Survey. .
Stanley, J.R., Brill, K.G., Koomey, J.G., 2007. Four Metrics Define Data Center 'Greenness'. White Paper TUI3009F. Uptime Institute, Santa Fe, NM. http://www.dcxdc.ru/files%5C4ede4eff-13b0-49d9-b4da-b0406bfc190e.pdf.
U.S. Environmental Protection Agency, 2016. Product Specifications & Partner Commitments Search ENERGY STAR. Online: https://www.energystar.gov/products/spec/.

Nathaniel Horner is a Ph.D. candidate in Engineering and Public Policy at Carnegie Mellon University, where his research has ranged across a variety of topics including innovation policy, energy use in the information technology sector, electricity transmission, and shale gas tax policy. He was a 2014 summer associate at the RAND Corporation. Prior to joining CMU, he was a member of the professional staff at the Johns Hopkins University Applied Physics Laboratory, filling systems analysis, software engineering, and project management roles. In addition, Horner was an adjunct lecturer in the Johns Hopkins systems engineering graduate degree program. He holds master's degrees in Systems Engineering from the University of Virginia and Computer Science from Johns Hopkins.

Inês Azevedo is an Associate Professor in the Department of Engineering and Public Policy at Carnegie Mellon, where she serves as co-director for the Climate and Energy Decision Making (CEDM) Center. Azevedo's research interests lie at the intersection of behavioral and decision-making, environmental, technical and economic issues, such as how to address the challenge of climate change and to move towards a more sustainable energy system. She addresses complex problems in which traditional engineering plays an important role but cannot provide a complete answer. She has participated in several reports from National Research Council and from the U.S. National Academy of Sciences.