Sustainable Computing: Informatics and Systems xxx (2013) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Sustainable Computing: Informatics and Systems

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

Measuring and analyzing the energy use of enterprise computing systems

Maria Kazandjieva [a,*], Brandon Heller [a], Omprakash Gnawali [b], Philip Levis [a], Christos Kozyrakis [a]

[a] Stanford University, United States
[b] University of Houston, United States

ARTICLE INFO

Article history:
Received 3 September 2012
Accepted 29 January 2013

Keywords:
Energy use
Green computing
Power measurements
Enterprise computing

ABSTRACT

Until now, green computing research has largely relied on few, short-term power measurements to characterize the energy use of enterprise computing. This paper brings new and comprehensive power datasets through Powernet, a hybrid sensor network that monitors the power and utilization of the IT systems in a large academic building. Over more than two years, we have collected power data from 250+individual computing devices and have monitored a subset of CPU and network loads. This dense, long-term monitoring allows us to extrapolate the data to a detailed breakdown of electricity use across the building's computing systems.

Our datasets provide an opportunity to examine data analysis and methodology techniques used in green computing research. We show that power variability both between similar devices and over time for a single device can lead to cost or savings estimates that are off by 15–20%. Extending the coverage of measured devices and the duration (to at least one month) significantly reduces errors. Lastly, our experiences with collecting data and the subsequent analysis lead to a better understanding of how one should go about power characterization studies. We provide several methodology guidelines for the green computing community. The data from the Powernet deployment can be found at http://sing.stanford.edu/maria/powernet.

1. Introduction

Common sense tells us that there are opportunities to reduce the energy waste of computing systems. For example, many people leave their computers on overnight, even when they are not needed, or have power-hungry PCs for undemanding tasks such as document-processing and web browsing. Observations like these have motivated recent research into green computing work [8,13,19].

Unfortunately, the data to support and evaluate new green computing solutions remains vastly anecdotal. Until now, power characterization studies have either collected data at the macro scale of a whole building [9], lumping all plug loads into one number, or at the micro scale from a handful of computers and LCD monitors [16]. Data at the macro scale is informative but difficult to act upon – it does not provide visibility into the computing components that can be made more energy efficient. Power data at the micro scale is great at providing a detailed characterization of a single device but fail to show how the individual datapoint relates to the full building energy use.

The green computing research community can benefit from the availability of more extensive power measurements. For example, a single PC power measurement from 2004 [12] has been used in papers as recent as 2010, citing it as a representative value. The aforementioned paper gives the power draw of a 2002 Dell 2350 1.8 GHz computer as 60–85 W. A 2009 paper [7] measured two desktops (102 and 72 W, respectively) and said their measurements were consistent with prior data, citing [12]. Later the same year, a characterization study [9] used 100 W per desktop plus LCD for some of its calculations, citing [7]. In 2010, LiteGreen [13] also referenced [7], stating that the typical PC draws 80–100 W when active. The paper goes on to measure one PC (95W active) and uses it to calculate potential energy savings of their proposed solution.

If we were to continue on citation trails like the one above, we risk using limited and possibly outdated data for new systems' evaluation. In addition, enterprise environments are often heterogeneous and it is beneficial to have power measurements from a larger selection of devices.

This paper helps fill the power data gap by characterizing energy data at the individual– and the building–scale levels. Every month, Stanford's computer science department's pays a $40,000 (330,000 kWh) electricity bill but there is no visibility into exactly where this energy is going and how much of it is spent on wasteful computing systems. This problem has much greater implications than a single department's budget. When one considers

* Corresponding author.
E-mail addresses: mariakaz@cs.stanford.edu (M. Kazandjieva), brandonh@cs.stanford.edu (B. Heller), gnawali@cs.uh.edu (O. Gnawali), pal@cs.stanford.edu (P. Levis), christos@cs.stanford.edu (C. Kozyrakis).

2210-5379/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.suscom.2013.01.009

the problem at scale, computing infrastructures add up to billions of dollars. According to the latest Department of Energy "Annual Energy Review",[1] computing in education and office buildings consume 66 billion kilowatt-hours of electricity [2]. This is 2% of all US electricity consumption and, as with data centers, is increasing.

This paper presents Powernet, a multi-year study on the computing infrastructure in our department. Over two years, we have measured plug loads from over 250 devices as well as utilization rates of a subset of the computers and networking equipment. The data from the deployment is available at http://sing.stanford.edu/maria/powernet. The measurement points were selected carefully instead of being a uniform, random sample. Since power does not necessarily follow a simple distribution over device type, outliers may be extremely important yet difficult to capture with a random sample.

Measuring over a long term allows us to quantify measurement errors over shorter durations. Measuring a large and diverse number of devices allows us to characterize the variation of power draw and utilization across and within device types. We augment our datasets with metadata including network device registrations and explicit equipment inventories. The combination of power data, utilization statistics, and metadata allows us to answer several open questions about green computing:

- What is the contribution of computing systems to an enterprise's overall electricity consumption and waste, and how is this cost distributed across different components of the computing infrastructure?
- Recent green computing research makes power analyses based on isolated research lab measurements: how do different assumptions and methodology techniques hold in a larger enterprise setting?
- We heavily instrumented our infrastructure because we did not know what we would find; now that we have an understanding of the data, how would one design a measurement infrastructure to achieve good accuracy with the least effort?

The answers to these questions form the fundamental contributions of this paper:

- Detailed examination of where energy goes reveals that over 50% of the electricity is spent on computing. PC's account for 17% of the bill despite the fact that their utilization is very low. Networking equipment comes at 3.5% and shows no temporal changes despite variations in traffic load.
- Data analysis shows that estimating savings based on a few isolated desktop measurements is prone to errors due to the wide spread of PC power draws. Assuming that a day of power is representative and using it to calculate yearly values can be off by as much as 20%.
- Our deployment and data studies expose the relative importance of device coverage versus duration of deployment. Once a deployment is past the first month of data collection, one must prioritize the 'what to measure' question over the time scale of the study.

The rest of this paper reviews the current state of green computing data before diving into the analysis of the Powernet datasets. It builds a picture of where energy is going in computing systems and what the utilization looks like. Next, the paper discusses a number of data and methodology insights, and it closes with guidelines for the design of future energy characterization studies.

[1] Released in October of 2011.

2. Background

Up until recently, the green computing community has had to rely on limited energy datasets, requiring researchers to make various explicit and implicit assumptions when analyzing the energy behavior of computing systems.

This section discusses some of the different ways in which related work has procured, used, and analyzed power data in the context of evaluating systems' research. In Section 6 we use the Powernet datasets to study what errors we can expect when using simplifying assumptions or limited data.

In measuring only a small number of devices and generalizing claims, there is an implicit assumption that instances of the same equipment model or size have the same power characteristics. For example, in [16], researchers measure an individual desktop and LCD monitor and determine the ratio of power use between the two. Then, they collect aggregate data from a power strip that has another desktop and LCD of the same size. Applying the precomputed ratio on the aggregate data leads to an estimate of the power draw of the second set of devices. Even though both desktops might have been Dells and both monitors 24-in. ones, we do not know how much variation there is between the seemingly identical devices.

A modeling approach that takes system subcomponents into consideration was used in [20]. Instead of collecting measurement with a meter, the authors use hardware components power models and software counters to calculate the power draw of a PC. This methodology was able to predict the power use of one machine based on a different one with 20% accuracy, indicating that even more sophisticated techniques that take device subcomponents into consideration will show error in estimation when assuming that similar equipment has similar power or usage behavior.

Prior works [8,13,16] have often based their analyses on power data collected from a small number of desktop measurements, 10 or fewer. In some cases, as shown in the introduction, evaluations are based on one or two values only. PCs are usually chosen based on convenience, i.e. desktops in the lab where the research is done, instead of using more deliberate samples.

Related work does not always indicate the duration over which power data was collected. In some cases [9,16] power draw is shown over the course of a day or week, but the rest of the time it is presented as a one-time, instantaneous measurement. Despite the limited duration of measurements, they are used in calculations of long-term benefits of energy-saving techniques. Thus, in Section 6 we explore the accuracy of using a day or month of data to make full-year energy estimates.

Fig. 1

The Energy Star program [3] establishes standards for energy efficient consumer products, including computing systems. As part of the process of obtaining an Energy Star certification, manufacturers of PCs submit power data for each of their devices. The result is an extensive database [6] of computer models and their power draw.

The lack of substantial power datasets within our research community means that it is not uncommon for academic works [12,19] to base their analyses on Energy Star data. These data, however, are not representative of the real-world power characteristics of machines.

The review of recent research reminds us that the more data and the better understanding of methodology we have, the better resulting insights and solutions will be. Powernet addresses both of these needs.

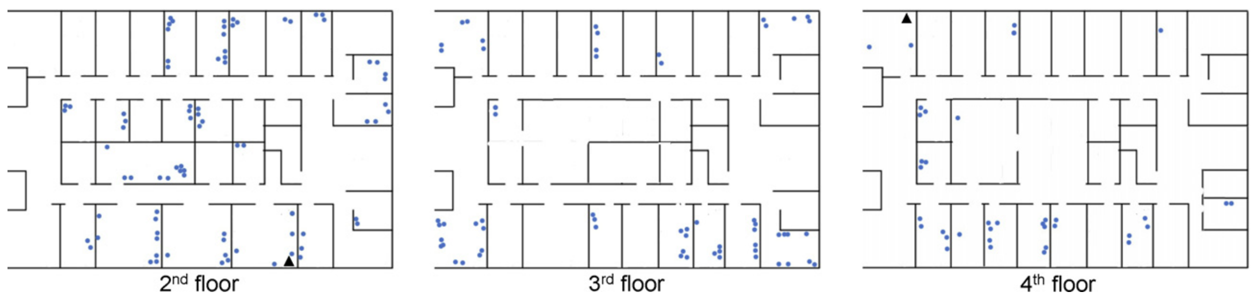

Fig. 1. The wireless power meter deployment spans three floors (shown) and one server room (not shown). Black triangles represent data sinks, while blue dots represent power meters. Most meters are under desks, near the floor. (For interpretation of references to color in this figure legend, the reader is referred to the web version of this article.)

3. Data collection and deployment experiences

Before delving into the analysis of power data and energy study methodology, this section presents the Powernet deployment. Powernet has been active for more than two years, continuously collecting power data from individual computing devices.

3.1. Original deployment

The primary requirement when designing the Powernet measurement infrastructure was that power meters are able to sense individual outlets at high sampling rates. This differs from many residential solutions that track whole-house energy consumption and report data every 10 or more minutes. Commercially-available Watts Up.NET meters were the first power sensors in the deployment, since they met our needs and were easy to obtain [5]. These meters transmit measurements over Ethernet, up to once a second. Meters were placed in wiring closets, the basement server room, and spread-apart offices. While these meters were a useful initial step in gathering power data, deploying and maintaining them proved to be difficult; problems surfaced even before the deployment began.

The first practical issue was the lack of in-field upgradable firmware. When a bug was discovered in the TCP stack, our only option was to pack up four large boxes of power meters and send them back, so that company staff could fix the proprietary code. After several weeks, the meters were back in our possession and the deployment could begin.

It quickly became clear that few offices had an open Ethernet port for each power meter. Many offices required additional small Ethernet switches and extra cables. The volunteer participants were unhappy with the clutter under their desks, due to the size of the meters. Each one weighs 2.5 lbs, with a thick, six-foot-long cord leading to a 7 in. × 4 in. × 2 in. base. Despite the physically clunky deployment experience, we were able to install 80 m.

In the Powernet building, each device must have a MAC address registration to obtain an IP address. Each group within the building has a unique VLAN, and each meter was statically registered to a group. The registrations could not be done all at once, since neighboring offices may correspond to different groups, and we could not know in advance how many meters would be needed for a given office. The network admins were burdened by the power meter registrations, and with this much manual configuration, mistakes happened.

We received an email from a network admin stating that "more than half of all DNS lookups emanating from [the three Engineering buildings] to the campus servers" were coming from the power meters. The solution for the lack of DNS caching was to go back to each meter, plug it into a laptop via USB, and hard-code the IP address of the Powernet server.

In addition to DNS lookups, the meters were also making ARP requests once per second and overwhelming the network security monitoring infrastructure. We received another email from the IT staff, pointing out that "[t]he 70 current meters now account for 20% of total daily recorded flows" by the security system. To work around this problem, the logging server was moved to a special VLAN that was not monitored by the network admins. That resulted in an IP address change, which meant yet another trip to the individual meters to update the hard coded IP address of the server.

Once the deployment was in place, we observed a number of meter software errors. From the 90 power meters, 8 completely stopped working; they did not power up or did not send or display any data. Another set of 5–7 m began reporting incorrect data at some point of the deployment; from the reported numbers we guess it was an integer overflow issue but the closed firmware did not allow us to verify this. The erroneous data was purged from the analyzed data sets. There were also some meters that would stop reporting data over the network until they were rebooted. That again was likely a software problem where the meters were reverting to logging data locally instead of pushing it out via HTTP. Of the original 90, only about 30 are still in operation; a number of residents simply unplugged their meters.

To their credit, the wired meters generally reported accurate data and work well for a dispersed deployment such as the wiring closets. However, three key issues made the wired meters unsuitable for large-scale deployment: the lack of code accessibility and remote firmware upgrade, the overhead of installing the meters within the building network, and user dissatisfaction with clutter and frequent maintenance. These experiences suggest that zero-configuration networks would improve ease of deployment.

3.2. Wireless deployment

To scale the power-monitoring deployment, we designed custom wireless plug-level sensors. The sensing portion of these meters includes current and voltage sensors, plus a digital power meter chip that gives an instantaneous power reading [1]. The communications portion includes a low-power processor, a radio (2.4 GHz, 802.15.4-based), and an integrated antenna. The meter software, built on TinyOS [18], includes sampling, routing [4] and dissemination [15] capabilities. The top-level application reads power draw every second and sends a data packet after buffering ten samples. The meters collect data via an ad-hoc multihop network using the Collection Tree Protocol [4]. The wireless power meters cost about $110 apiece.

The current board, similar to the ACme [16], has been extensively tested and calibrated. We use a WattsUp meter in line with our power meters to calibrate them at different points between 0 and 300-W loads. We find that raw meter values exhibit linear behavior with a r-squared of 0.99 or above for all meters. Of course, our calibration is limited to the accuracy of the WattsUp meters. Complementary to our work, [14] have designed a calibration method that can achieve utility-grade accuracy.

The deployment of the first batch of 85 wireless meters took several afternoons, compared to two weeks for the wired meters. The benefits of the wireless deployment were noticed immediately, and some users even requested that we replace their wired meters with wireless ones. The IT staff was not burdened by meter registrations, and the open nature of the software and hardware made modifications easy. The main meter limitation is transmission distance but as the deployment expanded from a single floor to multiple ones we added two more base stations. The ability of the network to self-organize was key during this step, keeping efforts to a minimum and preventing any disruptions in data collection.

Instrumenting the entire Gates building was not feasible due to the costs and practical challenges associated with monitoring over 2000 devices. Yet, we wanted detailed-enough data to understand where in the building energy is spent and wasted. Several consideration went into deciding what to instrument. We focused our efforts on one of the two building wings, considering it representative of both wings in terms of types of devices and usage cases. Further, we were only interested in computing equipment, therefore we did not include miscellaneous electric loads such as staplers, fridges, coffee makers, or lights and HVAC. This is in contrast to the @Scale deployment [10,14] which adopted a stratified sampling approach in order to avoid a random sample overwhelmed by small, insignificant loads. The main goal of Powernet 's samples was to measure a wide variety of computing equipment to maximize the new information we gain. We did partially follow the stratified approach in allocating meters to go to different device categories such as servers and networking equipment.

To date, we have not observed any hardware failures in the 200+m that have been active at different stages of the deployment. The wireless network has proven a reliable way to collect the data, adapting to varying environment condition as well as load. We were able to seamlessly add additional base stations to improve connectivity without any manual re-configuration of the network. Any lost data was almost exclusively due to downtimes of the back-end server due to upgrades or power failures.

Powernet takes a unique perspective on green computing by measuring not only device power draw but also device usage. We deploy a number of software sensor that collect utilization statistics in the form of CPU, active processes, and network traffic. This utilization data is key for determining energy waste – the cases in which power is drawn but no useful work is done.

Tables 1 and 2 summarizes the different types of devices that our sensing infrastructure measures as well as the total number of datapoints for different sensing categories. The Powernet infrastructure gathers about 1GB of data every day.

Table 1
Powernet covers a variety of devices whose power measurements enable a characterization of the energy consumption of the whole building. Some devices also have CPU utilization or network traffic monitors.

Device type	Count
Desktop	75
Monitor	70
Laptop	28
Network switch	27
Printer	15
Server	36
Thin clients	12
Misc	3
Total	266

Table 2
Summary of collected data, organized by type of measurement.

Sensing type	Num. datapoints
Power data	10 billion
CPU percent	400 million
User processes	2 billion
Network traffic	10 million

4. Energy consumption

This section presents the data collected via Powernet and tackles the problem of extrapolating individual Powernet measurements to the whole building. We divide computing devices into four classes: LCD screens, PCs, networking equipment, and servers. Then, we combine power data with network activity logs, device registration databases, and a survey of building occupants to extend our observation to an energy picture of the full IT system.

We find that computing systems draw between 210 and 259 kW, depending on the time of day, or 47% to 58% of the building's 445 kW load. This aggregate power draw translates to 170,000 kW-h, or 50% of the building's monthly electricity usage.

Table 3 and Fig. 2 summarize our extrapolation methodology and resulting breakdown. Ground truth is provided by aggregate measurements from outside the building, logged every 15 min by campus services. The top curve in Fig. 2 shows one week of this data.

4.1. Personal computers

Personal computers are the second largest contributor to the energy consumption of computing systems, after servers. According to the department's database of registered devices there are about 1250 machines in the building that are actively observed on the network. Of those, roughly 740 are PCs.

In order to extrapolate to the whole building we bin PCs in three classes – laptops, low-end desktops, and high-end desktops. Low-end PCs are those with average power of 80 W or less and include machines such a Mac Minis, Shuttle PCs, Dell Optiplex. Full-size desktops like the Dell Precision are considered high-end machines.

Next, we take the 742 desktop MAC addresses from the network database and cross-referenced them with the university's *whois* service. The whois metadata includes node descriptions, such as PC model and OS, provided upon network registration. Of the 742 nodes, 456 had description that allowed us to classify them as laptops, low- or high-end desktops.

Table 4 shows the number of machines in each PC class; nodes with available description are labeled as 'observed' and breakdown of the other 286 assumes that the observed distribution is representative of the building. While there is no good way of verifying

Table 3
We cross-correlate Powernet measurements with IT databases to extrapolate the energy consumption of all computing systems in the building.

Device type	Measured	Total	Extrapolated via	Total draw (kW)	Uptime (h/day)	Energy/mo. (kWh)	Share (%)
Switches	27	62	Network admin records	15	24	11,000	3.5
Desktops/laptops	83	742	whois, MAC registrations	80	24	61,000	17
LCD displays	70	750	Occupant survey	48	≈8	14,400	4
Servers	32	500	Manual inspection	117	24	86,000	26

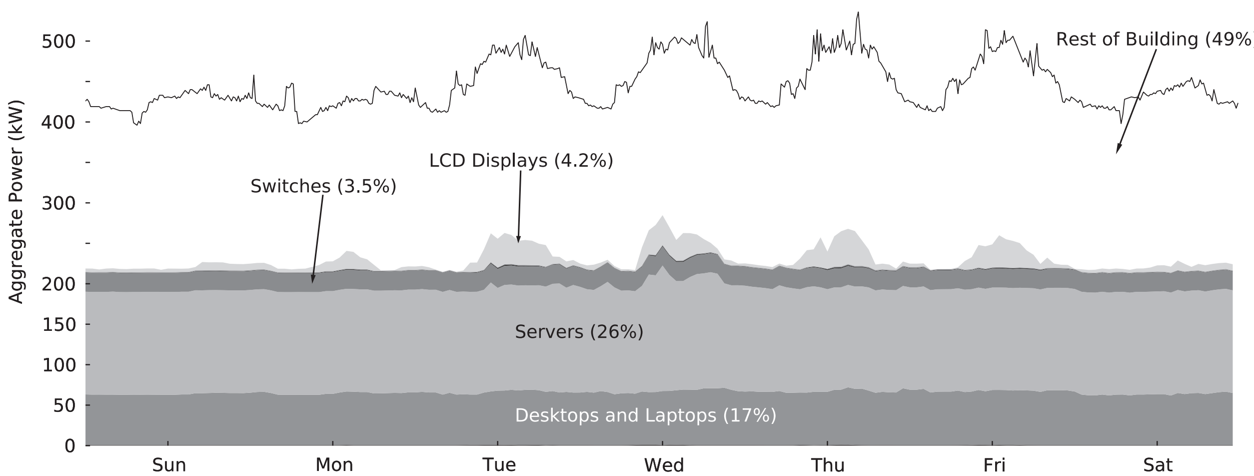

Fig. 2. Aggregate power draw for the entire Powernet building shows diurnal and weekday/weekend patterns. Computing systems account for 51% of the total 445 kW. The given week of data is representative of the building, except for Monday, which was a university holiday (February 15).

this assumption, it is a straightforward way of filling the gaps in inventory information.

Based on Powernet measurements, the median power draw for laptops is 26 W, for low-end machines – 63 W, and for high-end machines – 121 W. This means that the three categories of machines draw 2 kW, 4.4 kW, and 72.1 kW respectively for a total of 78.5 kW a day or 58,500 kW-h a month. The 742 personal computers in the building account for about 17% of total electricity consumption of the Gates building.

4.2. Computer displays

While not often discussed in the green computing literature, displays are just as prevalent as the PCs they are attached to. The trend of using larger size LCDs also means that their energy cost is increasing. For example, it is not uncommon for a 30 in. LCD to draw as much or more power than the average desktop.

To better understand the contribution of computer displays to the overall cost of computing systems, we first study a single LCD. Then we present the full Powernet dataset and extrapolate to the whole building.

Fig. 3(a) shows an hour-long data trace during which we adjusted one 30 in. monitor's brightness and desktop color scheme. Depending on the monitor brightness settings and the colors in the image displayed, the power draw varies by up to 35W (25%). Lowering the brightness by two settings (pressing the '–' button twice) reduced the average power draw from 145 to 117 W, a 19% reduction in consumption. Additionally, LCD power draw is affected by the colors displayed. More energy aligns more liquid crystals in each pixel, permitting more light to shine through and enabling them to display brighter colors. Thus, a 30-in. monitor has maximum power draw, measured at 145 W, when the majority of the screen displays white elements. Switching to a dark background and color scheme or viewing darker web pages reduces the draw to 127 W. Displaying dark colors with the lower brightness setting reduces power draw to 110W.

These findings prompted users participating in the Powernet deployment to lower their monitor brightness, as well as change their desktop backgrounds. Fig. 3(b) shows typical data from one such user who only modified desktop color schemes. (The monitor brightness was already reduced.) The monitor's power usage is shown over a working week day once in April and then again in May. We observe over 10% reduction in energy usage. For a device that is on about 40 h a week, 400 Wh are conserved.

In addition to the controlled measurements described above, Powernet actively collects power data from about 70 LCDs of various sizes. These individual sensing points allow us to quantify the average power draw of different size LCDs. Taking one step further to a whole-building extrapolation requires the distribution of display sizes in the building. To obtain an estimate of this distribution, we use an online survey asking occupants for the number, size, and manufacturer of the computer screens they use. Table 5 presents data from the 169 responses reporting 225 monitors. These responses account for 28% of the building's occupants. The table also shows the power consumption of different displays.

The cumulative power draw of the LCDs reported by users is 15 kW. Scaling that to the whole building yields a power draw of 52 kW or 12% of the building's power demand during daytime. Active duty cycling of screens reduces the energy footprint. Powernet data over time shows that displays are powered on 50–60 h a week. Therefore, over one month, LCD screens consume about 14,000 kWh, or 4.2% of the monthly electricity budget.

4.3. Server machines

Powernet monitors 32 of the 500 servers in Gates Hall. Similar to desktops, servers exhibit varied power profiles. For example, a

Table 4
Personal computers are binned into three categories, and university databases and active network node counts allow us to extrapolate to the whole building.

	Laptops	Low-end PCs	High-end PCs	Total
Observed	47	43	366	456
Estimated	29	27	230	286
Total	76	70	596	742

Table 5
A survey shows that majority of building occupants use mid-sized LCD displays. The number of large (30 in.) monitors is increasing as equipment is upgraded.

Size	Count	Avg. power (W)
< 20 in.	42	30
20 in.–22 in.	40	45
23 in.–25 in.	84	63
26 in.–27 in.	15	80
29 in.–32 in.	44	120

Please cite this article in press as: M. Kazandjieva, et al., Measuring and analyzing the energy use of enterprise computing systems, Sustain. Comput.: Inform. Syst. (2013), http://dx.doi.org/10.1016/j.suscom.2013.01.009

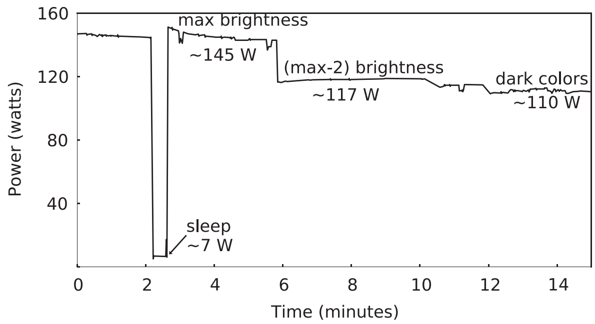

(a) Power draw of a 30" Dell display under different settings.

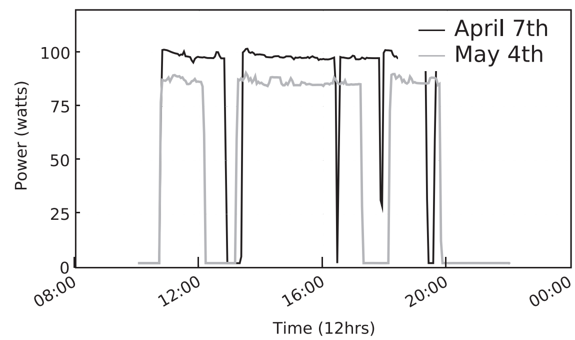

(b) Energy consumption can be reduced by 10%-28% without affecting usability. In this case, a user changed to a dark color scheme.

Fig. 3. Brightness level and color scheme have a significant effect on monitor power consumption. A one-time change in LCD screen configurations can have a large impact on the power draw. (a) Power draw of a 30 in. Dell display under different settings. (b) Energy consumption can be reduced by 10–28% without affecting usability. In this case, a user changed to a dark color scheme.

standard 1U rackmount can have a power draw anywhere between 95 and 275 W. Unlike desktops, the server population is much more homogeneous, e.g. 40 identical 1U machines in a single rack. Therefore, we spread out our measurements to get maximum coverage, with meters measuring identical devices for verification purposes. The average power draw we calculated from the Powernet measurements is 233 W. With about 500 servers, the aggregate draw is 117 kW – 26% of the total building energy consumption per month. In the future, we hope extend the set of server measurement in order to have a more precise extrapolation.

4.4. Networking equipment

Current networking hardware has constant power draw per linecard of Ethernet ports, with some power variation due to CPU load and fan activity [17]. For similar switch models, the number of linecards correlates strongly with power draw. The top two lines in Fig. 4 illustrate this point – even though the HP 5412zl switch has only 24 additional active ports, its power draw is a third higher that of the 6-slot HP5406zl. The power story is different when comparing switches that differ significantly in their make or year of production. The NEC switch is Fig. 4 has half the number of active ports compared to the 12-slot HP, yet only a 5th of its power draw.

At the building level, the network backbone is provided by 2 core switches located in the basement and 26 edge switches spread across the five floors. There are also a number of medium- and small-sized switches that have been deployed on as-needed basis. We account for all major switches and estimate the number of smaller ones with the help IT staff. Table 6 summarizes the types of networking equipment together with their power draw. The power draw of wireless access points is folded into the switch data since they are Ethernet-powered.

This observation together with the relatively small number and homogeneity of devices, lends well to a whole-building calculation. We use Powernet 's measurements and inventory from Table 6 to calculate the daily power draw of all networking equipment, 15.4 kW. This translates to 11,500 kWh per month or 3.5% of the building's total consumption.

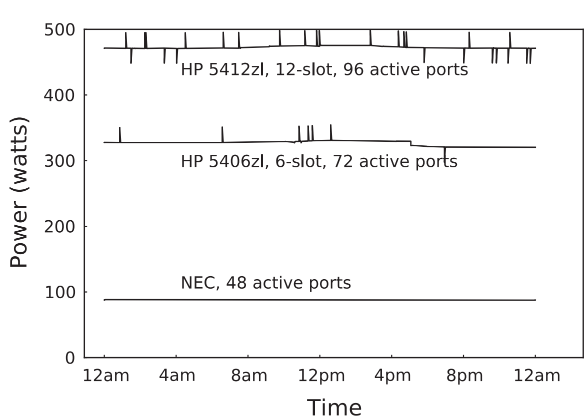

Fig. 4. Switch power consumption is constant, barring transient ups or downs. Differences between switches do exist, even if serving the same amount of traffic.

4.5. Summary

This section presented empirical power data and a new methodology for characterizing energy consumptions, using both plug-level empirical measurements and device metadata, to create a detailed picture of IT energy. We find that 50% of the building's energy goes to computing equipment: 26% goes to servers, 17% to PCs, 4% to displays, and 3.5% to networking.

Our data confirms prior observations and intuition that PCs and servers are major contributors to the energy bill of enterprise buildings [14]. The data also highlights that smaller parts of IT, such as networks and LCD monitors, account for almost 8% of the overall building's electricity use. We find that displays are responsible for 50% of the building's diurnal power draw variation and are the only computing component that exhibits such patterns. This confirms

Table 6
Summary of switch types, quantities, and and estimated individual power consumptions. This inventory includes all major network switches and excludes small per-room switches and hubs.

Type	# Count	Power draw (W)
HP 5406zl (6-slot)	20	325
HP 5412zl (12-slot)	8	500
HP 2724	2	100
Cisco Cat 6509	2	400
Cisco Cat 4000	2	600
Cisco Cat 3750G	2	160
Linksys	2	50
NEC (various)	5	100
Cisco (various)	5	100
Quanta (4-slot)	5	50
Misc (estimated)	100	10
Total major switches	53	

Please cite this article in press as: M. Kazandjieva, et al., Measuring and analyzing the energy use of enterprise computing systems, Sustain. Comput.: Inform. Syst. (2013), http://dx.doi.org/10.1016/j.suscom.2013.01.009

Table 7
CPU utilization of both student and administrative staff machines reveals that processing resources are only lightly taxed. Data was collected once a second for 11 months (students) and 1 month (staff).

Machine type	Percentile CPU		
	5th	50th	95th
Student PCs			
Dell Precision T3400	0%	1%	7%
Dell Inspiron 530	1%	1%	8%
Dell Precision T3400	0%	1%	13%
HP Pavilion Elite m9250f	0%	0%	25%
Dell Precision T3400	0%	4%	29%
High-end custom-built	0%	1%	57%
Dell Optiplex 745	1%	9%	58%
Staff PCs			
Dell Dimension 9200	0%	0.75%	3%
Dell Precision 690	0%	0.7%	4%
Dell OptiPlex 760	0%	0%	5.45%
Dell OptiPlex SX 280	0%	0.75%	5.5%
Dell Dimension 9200	0%	1.5%	8%
Dell OptiPlex 745	0%	1.5%	9%
Dell OptiPlex SX 280	0%	0%	10%
Dell OptiPlex 760	0%	1.55%	17%

that there is room for improvement not only in the IT infrastructure but also in the rest of the building.

5. Utilization

While a breakdown of the electric bill is a useful first step toward finding opportunities for savings, it is difficult to identify specific failures in energy efficiency. Energy data alone is not enough, it is only meaningful if paired with a characterization of systems' utilization. This section examines the workloads of computers and network switches to determine what part of the energy is spent well and how much is wasted.

5.1. Computers

Related work [7,20] suggests that desktop machines are rarely turned off when not in use, and Powernet power measurements over a >1 year-long period support this claim. So far, green computing research has focused on solving the problem of idle PCs. Our utilization data sheds light on an equally wasteful problem – power-hungry machine that even when active, barely tax their resources.

Powernet collects data from both student and staff PCs and since the computing needs of the two groups are likely to differ, we consider them separately. Table 7 shows the CPU utilization of a number of desktops. Computer science students use more of their available processing resources, but even so, in many cases CPU usage is under 30% for 95% of the time. The demand on administrative staff machines is even lower. Since most of the measured computers were left powered on at all times, the 50th-percentile data is not surprising: machines are often idling. What is surprising is that even when PCs are in use, the level of usage is low. If desktops were power-proportional that would not be an issue, but the current high baseline power draw means that the energy cost for a PC that is running at 5–6% of its capabilities is disproportionately high. In one extreme case, measurements showed that the most power-hungry staff desktop (quad-core Dell Dimensions 9200), drawing over 150 W, has the lowest CPU utilization – 3.1% for 95% of the time.

Another way of investigating whether utilization matches the type of equipment we buy is to look at typical tasks users perform. We focus on staff computing because it is more representative of an enterprise computing environment. Table 8 shows the most

Table 8
The most popular workloads on administrative computing systems are general office and web application. These workloads imply that a laptop can be a used instead of a desktop.

Process	% of time active
Acrobat Professional	1–4%
Firefox	0.5–4%
Internet Explorer	0.3–2%
MS Excel	1–2%
Thunderbird	0.4–1.2%
MS Word	0.2–0.8%
Outlook	0.4%
Acrobat Reader	0.3%
Explorer	0.01–0.3%

common workloads on administrative machines, excluding Windows services and virus checks. The percentage of active time is calculated as the cumulative time over one month that the process was running; the range of time captures the minimum and maximum numbers over four computers. The workload data raises the question of mismatched user needs and technology. There is no reason why an entry level laptop or a Mac Mini cannot perform the same basic tasks (document editing, web browsing, PDF viewing) as a quad-core, 150-W desktop.

Characterizing the utilization of computers has revealed that there is a lot more waste than idle machines alone. The baseline power draw of desktops, combined with low use of system resources, means that there are energy-saving opportunities even when PCs are actively used. Powernet 's PC utilization data suggests that future green computing research should tackle all PCs, not just idle ones.

5.2. Network equipment

Section 4 found that the networking infrastructure consumes 3.5% of the building's electricity monthly electricity. This translates to a cost of $15,000 a year just for networking. We also noted that switches consume a constant amount of power due to their hardware design. If the network is operating near capacity, then the 3.5% is energy spent well. Otherwise, if we find that the network operates at, say, 10% capacity even at peak, it means energy is wasted.

This prompts the questions of how much traffic is flowing through the 60 or so switches in the building, and whether smaller or fewer switches could more efficiently meet bandwidth demands.

We begin by examining the traffic coming into one of the four switches on the second floor of our building. This is an HP Procurve switch with 96 1-gigabit active ports, consuming 500 W and serving 50+people. Fig. 5 shows the switch bandwidth over one week, measured once per second. The demand never exceeded 200 Mbps – an amount that could have been handled by a less power-hungry edge switch and additional small switches (2–5 W each) in individual offices to meet port demand. To verify that this is not aberrant behavior, Fig. 6 shows the cumulative distribution of traffic for 7 building switches. Note that the x-axis has a log scale. The number of ports for different switches varies from 24 to 120 and the CDF data was collected over 40–420-day periods.

Similarly to PCs, switches are highly underutilized. For the equipment we measure, total network demand is lower than 1000 Mbps 100% of the time. Network over-provisioning is not a new concept or observation; it provides benefits, including higher throughput, lower loss, and lower jitter. But when the average utilization is under one hundredth of one percent, several questions are worth considering. Is the amount of over-provisioning unnecessarily large? How can we take better advantage of the large amount of bandwidth that today's networks are ready to support? Going forward, there are two ways to address the issue: consolidate

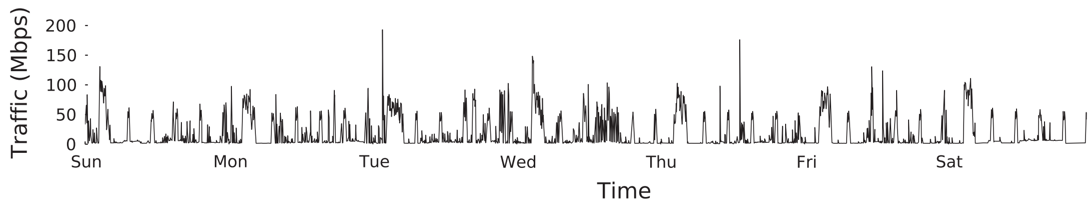

Fig. 5. Typical traffic patterns for one edge switches in the building. Network utilization remain low. Power consumption for this switch remain constant, at approximately 500 W.

equipment and make better purchasing decisions in the future, or make use of the extra available bandwidth.

The story that network traffic tells is no different than that of PC utilization – systems are heavily over-provisioned, often with no regard of expected workloads, leading to large energy wastes. Powernet 's contribution is in bringing such utilization data to light and placing it in the context of green computing.

6. Methodology insights

The previous two sections revealed details about the energy consumption in an office computing infrastructure. This section examines how the data collected via Powernet can help improve both the assumptions about power we make in research and the methodologies we apply to collect more data.

6.1. How does the frequency of sampling affect what the data reveal?

Prior work has rarely discussed in detail the sample interval at which power is measured. If the end goal is to calculate the amount of energy used up over a day or week, the frequency is not of great concern. However, there are inherent tradeoffs associated with the choice interval. Less frequent sampling will result in less stress on the measurement infrastructure and a more manageable dataset. The longer the interval between consecutive samples, the larger the risk of missing interesting events in-between.

An interesting example appeared during the Powernet deployment, illustrating the point above. Fig. 7 shows a one-hour timeline of power draw for a Dell Studio XPS desktop. Each datapoint in the graph is the average power over the last 1 min, for a total of 60 measurements. Minute granularly is not atypical for many commercially available plug-level meters. The flatness of the line indicates that the computer was largely idle for the duration of the data collection period. Furthermore, the value of approximately 100 W is a reasonable power draw for a full-sized desktop – nothing is out of the ordinary.

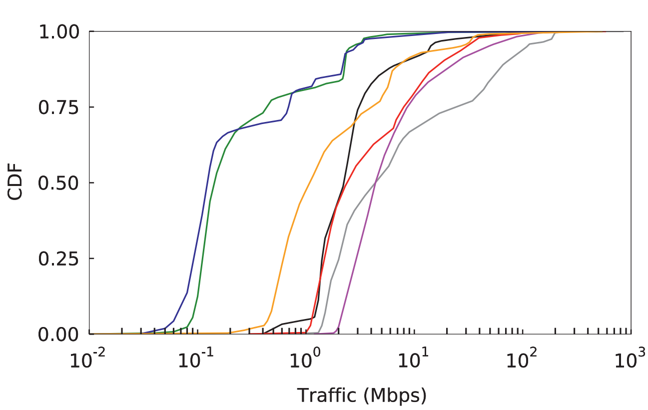

Fig. 6. CDF of traffic for seven switches over 6 months shows that switches are operating well under capacity.

Increasing the granularity of measurements paints a completely different picture, as shown in Fig. 8. This is the same one-hour period but power samples were recorded once a second, for a total of 3600 measurements. The frequent jump in power draw is immediately obvious and taking the FFT of the data confirms that the 30-W spike is regular with a period of 1 min.

Upon further investigation, we were able to correlate the power measurements with CPU spikes caused by the wireless card on the desktop. A quick search online revealed the existence of a bug in the Linux drivers for this specific piece of hardware. Turning off the wireless card solved the periodic CPU/power spike.

Concurrent work [11] has found that 10-s intervals are a reasonable choice, capturing power dynamics without overloading the infrastructure. Our experiences showed that depending on the task, different resolutions are desirable. For many practical uses – visualizing data, computing long-term estimates – even 5-min averages are useful. Higher-resolution data, on the other hand, is needed for correlating utilization metrics with power draw.

While different deployments will have varying resources and goals, experiences with Powernet teach us that having high-resolution data can be valuable not only in energy characterization but also in monitoring for unexpected behavior.

6.2. Devices of the same model or with the same specifications have low variation in power draw.

Prior work has implicitly assumed that instances of the same equipment model or specification have the same power

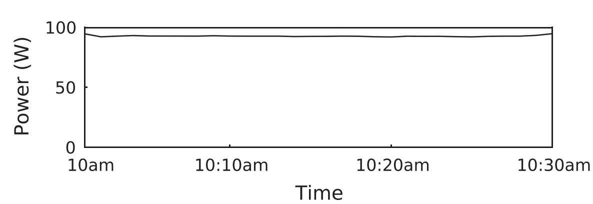

Fig. 7. Five-min averages of power data do not show anything out of the ordinary – the PC is idle at about 95 W.

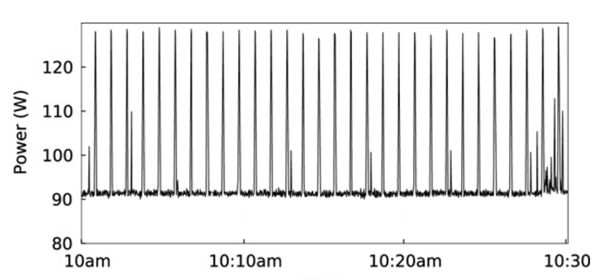

Fig. 8. Power data collected once a second reveals a misbehaving PC. Earlier, 5-min averages hid the anomaly. In certain use cases it is beneficial to have high-resolution data.

Please cite this article in press as: M. Kazandjieva, et al., Measuring and analyzing the energy use of enterprise computing systems, Sustain. Comput.: Inform. Syst. (2013), http://dx.doi.org/10.1016/j.suscom.2013.01.009

Table 9
Average power draw of two different devices with the same model (standard deviation shown in parentheses). Two devices of the same model can differ by as much as 43%. Networking equipment is more uniform than PCs.

Device	#1	#2	% diff
Optiplex 760	60W (9)	34W (24)	43%
Optiplex SX280	68W (12)	56W (8)	18%
Optiplex GX620	71W (8)	63W (13)	11%
Precision T3400	117W (17)	110W (10)	6%
HP 5400zl switch	467W (8)	463W (4)	0.01%

characteristics, simply because of the lack of better data. Under this assumption, measurements taken from one or two devices have been used to reason about other, unmonitored pieces of equipment. Unfortunately, such methodology can yield inaccurate results.

Powernet data reveals that some types of computing systems can exhibit large variations even when comparing two instances of the same device model. Table 9 shows five example devices – 4 Dell desktop models and 1 network switch. The two Dell Optiplex 760 desktops have over 40% difference in their average power draw. In contrast, the two HP switches have almost identical power draw, as well as very low standard deviation over time. In some cases, even though the PCs appear to be the same on the surface, they might have been upgraded with custom components, causing a difference in power draw. Furthermore, while two devices from the same type might have similar motherboards, power supplies, and processors, they can differ in in the user workloads they support, leading to different power profiles.

There is no single solution to error in estimating power, when only partial measurements are available. In the case of desktops, it is not surprising that there is great variability, but putting concrete numbers to it can help anticipate inaccuracies. Additionally, one could augment power measurements with other data, such as PC utilization, to get more accurate understanding of how equipment is used. In the cases when variability is low (e.g. switches), data points from only one or two devices can be treated as much more reliable.

6.3. *Does sampling a few devices in a class provides an accurate average measurement?*

Fig. 9, the Powernet curve, shows the wide distribution of desktop power. It is worth considering what errors can be expect if one were to sample only part of the PC population.

We use ground truth data from 69 desktops to show how the expected error of average power draw changes if based on random samples from the population. The average power draw

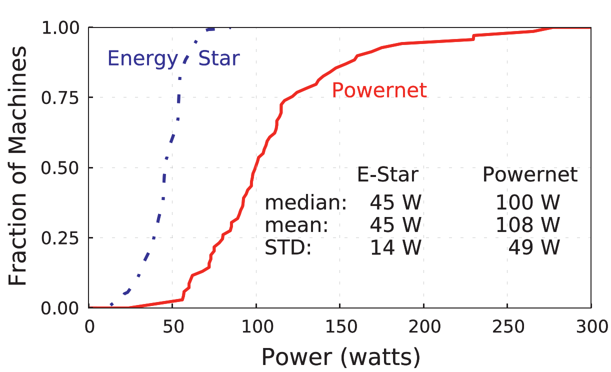

Fig. 9. Energy Star data is not representative of real-world PC power.

of the 69 desktops is 109 W. We generate 1,000,000 random samples of size 5, 10, and 20, drawing from the lists of 69 machines. Fig. 11 shows the resulting histograms of estimated average power. Samples of only 5 desktops can have more than 16% error in estimating the mean power draw. Increasing the sample size from 5 to 20 machines cuts the error by more than a half.

The lesson is that when it comes to PCs, a small sample is not desirable if trying to extrapolate to a large, heterogeneous set. Recent work [10] correctly brings attention to the importance of complete device inventories in order to understand how varied an environment is and targeting measurement points accordingly.

6.4. *Do short-term measurements accurately reflect long-term power draw?*

The Powernet datasets show that while the base load of computing systems is consistently high, month-to-month variations do exist. These changes result in slightly different energy use throughout the month and year. For example, over one year the monthly power draw average of one desktop varied from 183W (min) to 293 (max), with 216-W average over the whole year. Another PC was consistently between 247 and 257 W. One question to tackle is 'How does the duration of power measurements affect a yearly cost estimate for some set of devices?'

Our analysis uses data from 16 desktops; each PC was monitored for one year, from May 2010 to April 2011. The cumulative average power draw of the sixteen PCs is of 1524 W ($1600 for the whole year, at $0.11 per kWh.) Examining the monthly average power draw of each machine reveals that no single month is representative of the whole year. If we were to take one day or week or month data in the hopes of estimating the yearly electricity cost, we should expect to be off. But by how much?

Using a single day of data from the year-long trace allows us to generate over 350 different estimates for the yearly cost. Similarly, for one week or month of data, we can compute multiple cost estimates. We can also repeat the process using sliding windows with size of two to 11 months. Fig. 12 summarizes the results. The x-axis shows what duration of data was used for the estimate. The y-axis shows the average and maximum error of the estimate as percent of the real energy cost. In a worst case scenario, measurements taken for less than a week can have error of 15% or more. At the scale of a building IT systems, such error in predicting costs can be thousands of dollars. On the positive side, the analysis shows that collecting data at the month timescale, as opposed to longer, could yield data with an acceptable error. These results are in line with concurrent work at Lawrence Berkeley National Labs [11] who found that two months of data yields an acceptable tradeoff between deployment effort and accuracy.

6.5. *Is Energy Star data representative?*

For a long time Energy Star was the only large openly available computer power dataset and while this is changing, it is worth to discuss why Energy Star data should be used cautiously. It is composed entirely of devices that have passed minimum energy efficiency requirements. It does not reflect the distribution of devices sold and data is self-reported. Furthermore, Energy Star measurements and certification do not consider PCs under load – they only deal with idle, sleep, and off states.

Fig. 9 illustrates the divide between Energy Star data and the real-world measurements collected by Powernet. The differences are striking – close to 100% of the 4000+Energy Star desktops fall below the 100-W cutoff. In our measurements, that is the median PC

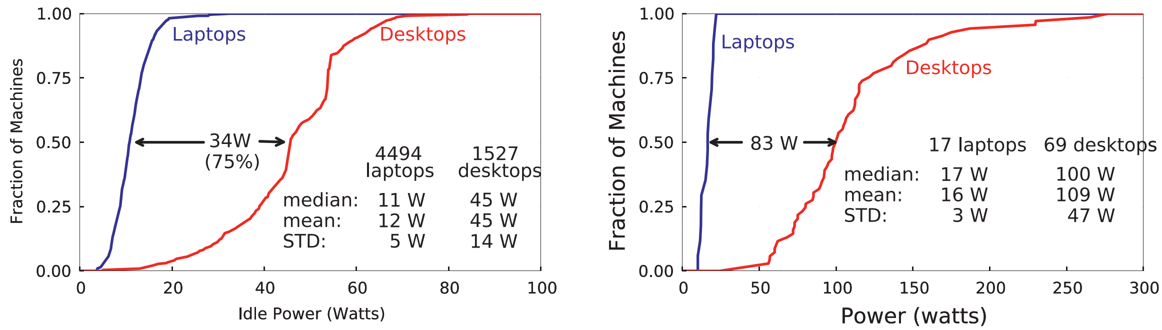

(a) Data for Energy Star-qualified computers as of March 2011

(b) Measurements from computers in our building

Fig. 10. We find a difference not only in the desktop measurements of Energy Star and Powernet, but also in the laptop ones. This is significant because the Energy Star dataset underestimates the true savings a desktop/laptop switch could have: (a) data for Energy Star-qualified computers as of March 2011 and (b) measurements from computers in our building.

power draw. The Energy Star dataset has the benefit of a lot more data points and real-world distributions might shift from building to building. Using Energy Star in lieu of real measurements is likely to underestimate energy costs in most contexts but can be extremely useful in advocating lower-power machines.

We observed one more way in which the Energy Star datasets differ from the smaller, but real-world, Powernet ones. Fig. 10(a) and (b) show Energy Star and Powernet data respectively. In addition to the desktop measurements presented earlier, laptop power data is also plotted. The 50th-percentile difference between desktops and laptops is much greater for the Powernet data – 83 W compared to 34. In an enterprise scenario where different hardware options are being weighed in, the Energy Star data set is likely to de-emphasize the tremendous savings to be had from switching away from desktops to using laptops instead.

Fig. 11.

Fig. 12

7. Measuring computing power accurately

The prior section pointed out several data and methodology considerations which, if not though out carefully, can lead to inaccurate analysis of energy use. The results from Powernet, however, represent only one point in time. As computing continues to evolve, green computing research will need to periodically re-measure energy consumption and waste. This raises the follow-up question: 'Given limited time, money, and effort, how should one measure computing system energy consumption in order to minimize error?' This section presents methodology guidelines to aide future green computing research.

7.1. Step 1: characterization

Not all device classes are equal: some require much more effort to measure accurately than others. Table 9, for example, showed a 43% variation in the power draw of Optiplex 760 PCs but a 0.01% variation in the power draw of HP 5400zl switches. An approximate ordering of the different devices in terms of variability places desktops as the most diverse, followed by servers, laptops, LCD monitors, and lastly, switches. Rather than distribute measurement points uniformly, one should measure the high variation device classes more densely. But device classes change quickly: Dell, for example, no longer sells Optiplex 760 PCs. Being able to determine which device classes have significant variation requires up-to-date, current measurements.

To understand where to measure, one first needs to know which device classes are high variation and which are not. This can be done quickly, as a series of point measurements made over a day. For example, suppose that an enterprise has a large number of a new Dell PC. One can randomly select 10 of these PCs and measure each of them booting. This will provide a large dynamic range of power measurements within the class as well as across the class. If the 10 show significant differences, then they might need to be measured densely. One can use the observed power draw distributions and statistically compute what deployment of sensors will lead to the lowest observed error.

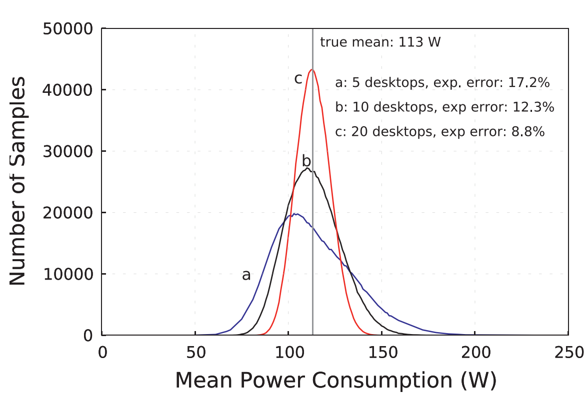

Fig. 11. Desktop diversity requires the measurement of a large sample of the population. In this experiment, if only 5 desktops are used to estimate the power of all 69, then the expected error is over 16%.

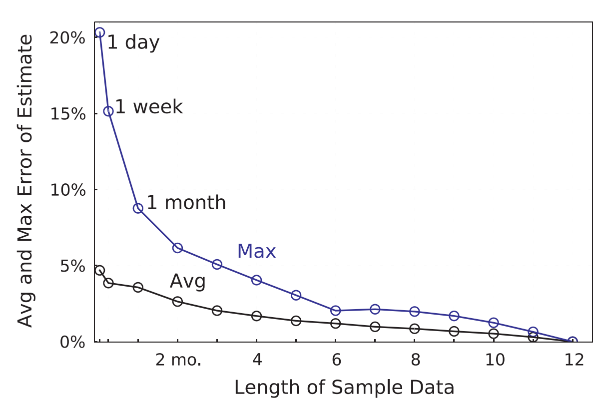

Fig. 12. As the number of months of data increases, the standard deviation of error in estimates decreases. Even if only one month of data is used over 16 desktop, the year approximation will be within 4% of the true value of $1600.

Please cite this article in press as: M. Kazandjieva, et al., Measuring and analyzing the energy use of enterprise computing systems, Sustain. Comput.: Inform. Syst. (2013), http://dx.doi.org/10.1016/j.suscom.2013.01.009

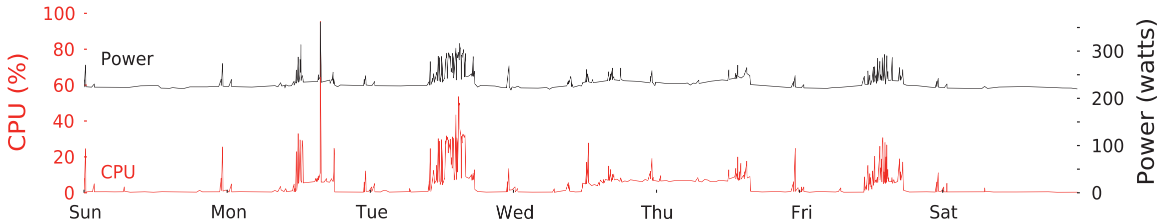

Fig. 13. A week-long trace of power consumption and CPU utilization shows how well the two track each other, with $r^2 = 0.996$.

These point measurements should use simple digital readouts (e.g., Watt's Up or Kill-A-Watt meters) which a person reads and writes down. Depending on a wireless mesh or wired network ports is probably more trouble than it's worth (lack of connectivity, VLANs, etc.).

7.2. Step 2: measurement

Once a short-term characterization study has provided guidance on where to deploy sensors, they need to be deployed for a sufficient duration. We were able to use custom sensors and our own software to collect data over a wireless mesh, but this technology is not commonly available. Our experiences with Watt's Up meters – coordination with IT infrastructure, reconfiguration, failure etc. – was that they are a poor choice for a very large, long-term deployment, but are acceptable at smaller scale.

The results in Fig. 12 showed that energy consumption, especially for personal computing, changes significantly over time. One should measure for at least a week, and preferably for a month. After a month, expected error, even for high-variation devices, drops to 4%.

In addition, while using power data from one machine to estimate another's can be problematic, CPU load can sometimes be used as a proxy for power on a single, calibrated machine. Fig. 13 shows one week of power and CPU data for a desktop. Visually, it is immediately noticeable that CPU tracks power very closely, with $r^2 = 0.996$. Therefore, in the context of desktops, one can collect a limited set of power meter measurements followed by the use of software sensors that report a feature (CPU) that is closely correlated with power.

Overall, given the choice between breadth (number of devices measured) and depth (length of measurement), greater breadth generally leads to more accurate results. At the extremes, it is better to gain a single point measurement of every device than measure one device for a year.

7.3. Step 3: extrapolation

The final step is to take the set of biased measurements and extrapolate to whole system power. Our experiences with Powernet have highlighted the need for data beyond power and utilization measurements. If extrapolation is to be successful, one also needs metadata in the form of equipment inventories and descriptions. Surprisingly, such metadata is not nearly as complete and readily available as we had hoped. Rather, we had to resort to indirect sources such as cross-correlating networked device registrations with active IPs on the network.

In the future, green computing researchers should encourage IT personnel to keep updated and detailed records of what equipment is added to a building.

8. Conclusion

Characterizing the energy use of enterprise computing systems is the first step toward identifying opportunities for improvement. Extensive, empirical data allow researchers to better quantify the problems they are tackling and the potential impact of their proposed solutions. Powernet has provided such data and has shed light on some of the assumptions that we make when faced with the lack of solid measurements.

Despite our best attempts to cover as many computing systems for as long as possible, the Powernet data remain but a single study. While the exact breakdown of energy use and waste might shift from building to building, the overarching methodology and data analysis lessons remain. Going forward, green computing research has not only a reference dataset to use but also a blueprint for how to characterize enterprise building power given limited time and resources.

Acknowledgments

We would like to thank the anonymous reviewers of this and earlier version of the paper as well as our IGCC'11 shepherd, Randy Katz, for their thoughtful comments. We also thank Wanja Hoffer, Mayank Jain, and Jung Il Choi for many insightful discussions and suggestions. This work is supported by the Department of Energy ARPA-E under award number DE-AR0000018. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

[1] ADE7753.;1; Single Phase Multifunction Energy Metering IC with di/dt Input. http://www.analogdevices.com
[2] Department of Energy,;1; Annual Energy Review, October 2011. http://www.eia.doe.gov/aer/
[3] Energy Star;1; http://www.energystar.gov
[4] TEP 123.;1; Collection Tree Protocol. http://www.tinyos.net/tinyos-2.x/doc/
[5] Watt's;1; Up Internet Enabled Power Meters, 2009 https://www.wattsupmeters.com/secure/products.php
[6] Energy;1; Star Computer Power Data, 2010 http://www.energystar.gov/index.cfm?c=revisions.computer_spec
[7] Yuvraj Agarwal, Steve Hodges, Ranveer Chandra, James Scott, Paramvir Bahl, Rajesh Gupta, Somniloquy: augmenting network interfaces to reduce PC energy usage., in: Proc. Networked Systems Design and Implementation, 2009.
[8] Yuvraj Agarwal, Stefan Savage, Rajesh Gupta, SleepServer: a software-only approach for reducing the energy consumption of PCs within enterprise environments., in: USENIX Annual Technical Conference, 2010.
[9] Yuvraj Agarwal, Thomas Weng, Rajesh Gupta, The energy dashboard: improving the visibility of energy consumption at a campus-wide scale., in: ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, 2009.
[10] Richard Brown, Steven Lanzisera, Hoi Ying Iris Cheung, Judy Lai, Xiaofan Jiang, Stephen Dawson-Haggerty, Jay Taneja, Jorge Ortiz, David Culler, Using wireless power meteres to measure energy use of miscellaneous and electric devices in buildings, Energy Efficiency in Domestic Appliances and Lighting (2011).
[11] H.Y. Iris Cheung, Steven Lanzisera, Judy Lai, Richard Brown, Stephen Dawson-Haggerty, Jay Taneja, David Culler, Detailed energy data collection for miscellaneous and electronic loads in a commercial office building, in ACEEE Summer Study on Energy Efficiency in Buildings, 2012.

Please cite this article in press as: M. Kazandjieva, et al., Measuring and analyzing the energy use of enterprise computing systems, Sustain. Comput.: Inform. Syst. (2013), http://dx.doi.org/10.1016/j.suscom.2013.01.009

[12] Kenneth Christensen, Chamara Gunaratne, Bruce Nordman, Alan George, The next frontier for communications networks: power management, Computer Communications (2004).
[13] Tathagata Das, Pradeep Padaala, Venkata Padmanabhan, Ramachandran Ramjee, Kang Shin, LiteGreen: saving energy in networked desktops using virtualization., in: USENIX Annual Technical Conference, 2009.
[14] Haggerty Dawson-Haggerty, Stephen Steven Lanzisera, Jay Taneja, Richard Brown, David Culler, @scale: insights from a large, long-lived appliance energy WSN, in: 11th ACM/IEEE Conference on Information Processing in Sensor Networks, SPOTS Track (IPSN/SPOTS '12), 2012.
[15] Jonathan W. Hui, David Culler, The dynamic behavior of a data dissemination protocol for network programming at scale, in: Proc. ACM Conference on Embedded Networked Sensor Systems (SenSys), 2004.
[16] Xiaofan Jiang, Stephen Dawson-Haggerty, Prabal Dutta, David Culler, Design and implementation of a high-fidelity AC metering network, in: Proc. Conference on Information Processing in Sensor Networks (IPSN), 2009.
[17] Maria Kazandjieva, Brandon Heller, Philip Levis, Christos Kozyrakis, Energy dumpster diving., in: Workshop on Power Aware Computing and Systems, 2009.
[18] Philip Levis, David Gay, et al., T2: A Second Generation OS For Embedded Sensor Networks, Technical Report TKN-05-007, Telecommunication Networks Group, TU Berlin, 2005.
[19] Sergiu Nedevschi, Sylvia Ratnasamy, Jaideep Chandrashekar, Bruce Nordman, Nina Taft, Skilled in the art of being idle: reducing energy waste in networked systems, in: Proc. Networked Systems Design and Implementation, 2009.
[20] Joshua Reich, Michel Goraczko, Aman Kansal, Padhye Jitendra, Sleepless in Seattle no longer, in: USENIX Annual Technical Conference, 2010.

Maria Kazandjieva is a Ph.D. candidate at Stanford University. Her research includes measurement and characterization of energy consumption, and the design of novel, energy-efficient compute systems. In her spare time, Maria enjoys learning about nudibranchs, black holes, and succulents.

Brandon Heller is a Ph.D. student at Stanford with interests in energy-efficient systems and software-defined networking. His recent projects include the ElasticTree data center network optimizer, editing the OpenFlow specification, and building NDB, a debugger for networks.

Omprakash Gnawali is an Assistant Professor at the University of Houston, USA. He was a Postdoctoral Scholar at Stanford University, got his Ph.D. from the University of Southern California, and received his Masters and Bachelors degrees from the Massachusetts Institute of Technology. His research lies at the intersection of low power wireless networks and embedded sensing systems.

Phil Levis is an Associate Professor in the Computer Science and Electrical Engineering Departments of Stanford University. He researches the design and implementation of computing platforms that integrate computers and the physical world, such as sensor networks, 3D virtual worlds, and simulation.

Christos Kozyrakis is an Associate Professor of Electrical Engineering & Computer Science at Stanford University. His current research focuses on resource management in large-scale multi-core systems, energy efficient data centers, and architectural support for security.