Software or Hardware: The Future of Green Enterprise Computing

Abstract

Over the last few years, interest in "green computing" has motivated research into energy-saving techniques for enterprise systems, from network proxies and virtual machine migration to the return of thin clients. This paper tries to answer a possibly contentious question: would we be better served by the embarrassingly simple approach of replacing every desktop with a laptop? To answer this question, we use power and utilization data collected from more than 100 devices over durations up to 15 months. We find that choosing the right computing systems – laptops – would save more energy than state-of-the-art power management software or thin clients. Furthermore, the marginal savings of applying software techniques on top of laptops is so small that it is probably not worth the trouble.

When selecting computers, there are many other considerations than just energy, such as computational resources, and price. We find that these factors generally do not reduce the attractiveness of a laptop-based enterprise. We discuss current trends in enterprises today, and how our conclusions might affect their directions, sketching a future of how a cost-efficient enterprise might become a hybrid system entwining laptops and tablets with a computing cloud.

1 Introduction

In the enterprise, a system's energy consumption has become a significant portion of its total lifetime cost. For example, after three years, a 120 Watt desktop costs an additional $360: for a $500 desktop, energy constitutes 40 percent of its 3 year total cost of ownership. Motivated by this trend, *green computing* has emerged as a new area of research with a focus on understanding and reducing the energy consumption of everyday computing systems. The search for green computing has inspired a wide spectrum of interesting software systems research, from sleep proxies [12, 18, 34, 36, 43] to thin clients [8, 14] to dynamic virtual machine migration [16, 19].

In the context of this flurry of interest, it may be useful to take a step back and look at how energy conservation has evolved in other domains. The 1940s, 1973, and 1979 in the United States all had governmental and social programs to conserve energy in response to a shortage of oil (from World War II and political events in the Middle East, respectively). The first International Energy Conservation Month was October, 1979 [17]; President Carter created the Department of Energy and set the White House thermostat to 68 degrees.

In response to these energy crises, the U.S. (and other) governments took a two-pronged approach: reducing usage and improving efficiency. Reducing usage involves convincing people to alter their behavior: turning off lights, lowering the thermostat, and allowing right turns at red lights. Improving efficiency involves changing inherent energy costs: fluorescent bulbs, better insulation, and higher gas mileage.

In the forty years after the 1970s energy crises, usage reduction has proven less effective than improving efficiency [27, 40]. Subsiding fluorescent bulbs reduces energy consumption more than asking people to turn off their lights; promoting high mileage vehicles is more successful than convincing drivers to slow down.

This paper asks the question: are computing systems different? Is reducing use – software techniques to power cycle computers and change user behavior – more effective than improving efficiency by purchasing the right equipment? Or, could it be that computing is like any other energy domain: we should focus on optimizing hardware and buying the right hardware for the job, with complex software or behavioral techniques taking second place, if useful at all.

In order to evaluate energy conservation techniques for enterprise computing in a crisp, quantitative way, we take a cynical view and consider a purely economic standpoint. Rather than consider moral or social motivations, we ask a simple question: how much money does conservation save? To answer this question we use power data collected over a year from over a hundred computing devices in an enterprise setting, combined with measurements of the utilization and workload on those devices.

The basic conclusion this paper reaches is that computing is no different than any other energy consumer: improving efficiency is more effective than reducing usage. Rather than deploy software solutions to manage desktop energy states, it is far more cost effective to simply replace desktops with docked laptops. Over three

years, replacing a desktop with a laptop can save on average of $300: over 80% of the energy cost and 32% of the total cost of ownership. A laptop-based approach also provides greater savings than thin client solutions, due to their need for back-end servers.

Furthermore, applying current software techniques on top of these savings does not lead to significant additional savings. Assuming the best reported results from a recently proposed virtual machine migration scheme would save no more than an additional $12 over one year. These comparatively modest savings have to be weighed against the complexity of managing a distributed software architecture.

The basic conclusion from these results is that influencing purchasing is more important and effective than trying to add complexity on top of existing systems. This observation raises interesting questions going forward, such as the implications to hardware architectures and the relationship between end-user systems and cloud computing. We present some conjectures on how a future enterprise computing system might look, finding that current trends towards cloud computing – driven by ease of management, low investment cost, and scalability – may also lead towards green computing.

Energy efficiency in enterprise computing systems has received much less attention than data centers – despite the fact that enterprise systems are at least as large a portion of the U.S. power budget [2]– due to the greater administrative challenges of enterprises . The results in this paper suggest that very simple policy changes, if followed, could change this situation in the near future.

2   Energy Conservation Today

The combination of government mandates and rising energy costs [24] has led to many efforts from both research and industry for energy-efficient computing, especially in the enterprise. Enterprise computing is diverse, both in workloads as well as equipment. Correspondingly, these computing systems today use a huge variety of techniques to save energy by reducing use. This section provides a brief overview of common practice and related research.

2.1   Enterprises Today

Desktops still constitute a significant fraction of enterprise computing. For example, in a 2008 purchase the Department of the Interior made for Tennessee Valley Authority, 66% of the devices were desktops [25]. While some technology-oriented enterprises have a larger fraction of laptops, desktops still dominate in more traditional enterprises.

Modern operating systems can put a machine in sleep mode when idle. In practice, however, users rarely activate these features [38] and IT departments often disable these features to make patching, backup, and maintenance easier [37].

IT departments with a greater eye on saving energy can reduce use with a centralized power management tool. Compiling data from Energy Star case studies for 7 deployments of 11,000 - 499,000 machines, we find sleep scheduling was able to save between $10.75 and $95/computer per year [22]. These deployments used a combination windows built-in sleep function, group policy, and software systems such as PC Powerdown, EZ GPO, Tivoli systems, BigFix, 1E NightWatchman, Computer Associates UAM, and Landesk Management Suite.

2.2   Sleep Proxies

Current IT approaches have two basic costs. The first cost is the administrative complexity of managing the power state of all of the devices. The second is that desktops, when powered down, cannot be accessed remotely.

Sleep proxies, also known as network proxies, attempt to address this second limitation. Sleep proxies are always-on hosts on the same LAN that intercept packets targeted at a sleeping host and answer on its behalf [18].

For more complex tasks, a network proxy can wake up the sleeping PC. Sleep proxies can keep machines in sleep for up to 50% of the time while providing uninterrupted network access for a limited set of protocols [34]. In a real-world enterprise setting, this architecture achieved energy savings of about 20% [36].

Somniloquy [12] augments a single desktop with a low-power embedded computer sleep proxy. Somniloquy device runs stripped-down versions of common applications (e.g., file transfer), resulting in savings of up to 65%. SleepServer [43], in contrast, proxies applications in trimmed-down virtual machines, reducing energy consumption by up to 60%.

2.3   Clients and Servers

Thin clients improve efficiency by consolidating many user environments into a small number of servers. The Sun Ray client [8] and Sun Fire server [7] comprise an example thin client product. Software-based remote desktop clients can also provide a thin-client-like system when run on local, low-power commodity hardware that connects to a remote server. Examples include Windows Remote Desktop [10], RealVNC [6], and NX [4].

Virtual machine migration is a related approach that works by migrating a full user environment running in a virtual machine (VM). When the user is at their computer, the virtual machine executes locally; once the PC

Figure 1: Evolution of idle (top) and sleep (bottom) power for Energy Star qualified desktops and laptops.

becomes idle, the VM migrates to a server and the PC sleeps. LiteGreen [19] uses full VM migration, while partial desktop migration can be used to reduce migration overhead and latency [16].

3   Hardware or Software?

The prior section described a variety of hardware and software solutions for conserving energy. Our hypothesis is that computing is no different than any other energy domain: efficiency is more effective than reducing energy usage. This section tries to answer a simple question – 'Might buying laptops for everyone be the easiest and most cost-effective solution?'

To reduce energy savings to a concrete value, we use the price of commercial electricity in California, 11.5¢/kWh. This price is on the lower end of pricing in California, but is a reasonable approximation of the national average. At this price, it costs approximately $1 to run a 1 Watt device for a year:

1W × 8760h/year × 11.5¢/kWh = $1.007/year.

3.1   Energy Star Data

As a basic sanity check on our hypothesis, we examine data sets of Energy Star certified devices [20]. These data sets represent a huge diversity of devices. To be Energy Star certified, a machine must stay below a yearly energy consumption threshold, given a ratio of off, sleep, and idle states. Many levels of the US government are required to purchase Energy Star hardware [21] as well as enable Energy Star features [24].

Figure 1 shows historical trends for the idle and sleep power of Energy Star-certified desktops and laptops [23]. Each data point in Figure 1 is the averaged sleep or idle

power draw of the Energy Star certified models for that year. Figure 2(a) shows the full distribution of idle power numbers for the most recent data set, released in March, 2011.

Both plots show that, for both laptops and desktops, sleep power is tiny. Sleep power near zero watts implies that aggressively putting machines to sleep may yield meaningful savings. Applying techniques to reduce use, such as LiteGreen [19] or SleepServer [43], can save significant energy. Their success, however, hinges on the fraction of time they are applied.

Both plots also show a large power gap between desktops and laptops that is consistent over time. Laptops are, as one might expect, far more efficient. The median Energy Star desktop draws 45 Watts, and costs $45 a year if run continuously, while the median laptop costs $11 per year. The 4:1 ratio of median idle power means that an energy conservation approach based on reducing use must keep a desktop asleep *at least 75% of the time* to match laptop efficiency.

3.2   Back of the Envelope

We know the energy expected savings from improved laptop efficiency: 75%. We can extract the expected savings for two sample software approaches from their respective papers. Together, these numbers allow us to perform a very rough, back-of-the envelope calculation on their comparative effectiveness.

LiteGreen states –

"Our findings from a small-scale deployment comprising over 3200 user-hours of the system as well as from laboratory experiments and simulation analysis are very promising, with energy savings of 72-74% with LiteGreen compared to 32% with existing Windows and manual power management."

SleepServer reports –

"The measured energy savings across all machines in our deployment for the month of September range from 27% to 86%, with an average savings of 60%."

For completeness, we also consider manual power management and applying two of the best use reduction technique to laptops. Table 1(a) summarizes the savings. *A one-time switch to laptops is at least as effective as the best of software techniques.* Additionally, while combining techniques can optimistically provide an 93% savings over the current cost, the incremental cost reduction over laptops is a mere $8.50 per year.

3.3   Enterprise Measurements

The Energy Star data provides a good sanity check against a wide range of devices. However, the data has
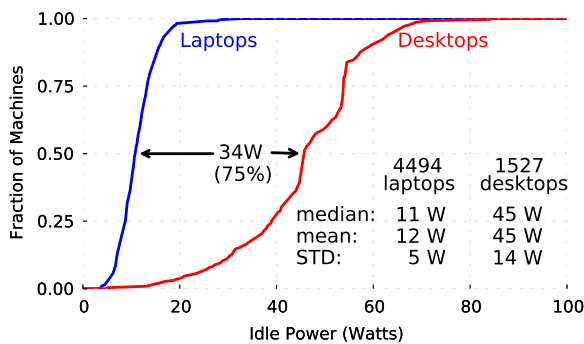
(a) Data for Energy Star-qualified computers as of March 2011

(b) Measurements from computers in our building

Figure 2: Idle power distribution for desktops and laptops. The desktops and laptops in our building have higher idle power than the median Energy Star qualified product. This is because in practice enterprises buy equipment beyond the list of green-certified products.

many flaws that prevent us from reaching a hard conclusion. The data is biased: every device is Energy Star certified. Because it is simply a list of devices, the data has a further bias in that it represents the distribution over all measured devices, not the more meaningful distribution of devices sold. Additionally, Energy Star only deals with the power draw of machines which are either off, asleep, or idle; it does not consider the energy performance of machines under any kind of load. Finally, the data is self-reported by companies that take measurements in a lab, not in a real work environment.

To determine whether the results in Table 1(a) hold in practice, we use a data set collected over the past 15 months from our university's computer science department. The data comes from a wireless sensor network that monitors 250 devices, including servers, desktops, laptops, printers, fax machines, wireless APs, routers, switches, and other electronic devices. For the purposes of this paper, we concentrate only on desktops and laptops. Appendix A provides details on how the sensor network collects and stores this data.

We used the department's IT database of registered computers to determine that there are approximately 650 desktops and 100 laptops. More precise numbers are difficult to obtain due to the continual churn within the department. Our power measurements cover 86 machines (69 desktops and 17 laptops), or roughly 10% of all end-user computing systems.

Figure 2(b) shows the cumulative distribution of power draws of the measured desktops and laptops. Unlike the Energy Star numbers, which were single numbers taken at idle, these come from averaging over all measurements above 2 watts (when the device was on or charging). Each average may include samples from a

range of utilization levels, anywhere from idle to 100%. For desktops, the measured range includes all samples, with only one exception, a desktop that occasionally slept. For laptops, the data excludes data points when the laptop was in sleep mode. Desktop and docked laptops do not include the power cost of an external monitor, if one was present.

There is an almost tenfold difference between the lowest- and highest-power desktops. The tail of the desktop curve extends much further than the Energy Star curve: the department has some high-power desktops which would not pass Energy Star certification. The lowest-power desktop that we measured was a Mac Mini, drawing about 27 watts. The most power-hungry machine was a custom-made desktop with a high-end graphics card drawing over 250 watts. The median power draw in our long-term measurements is 100 watts for desktops and 17 watts for laptops – a ratio of almost 6 to 1.

Using this data, we can repeat the comparison of software techniques to the all-laptop approach, as we did with the Energy Star data. Table 1(b) summarizes our findings. Compared to the Energy Star data, the ratio between the mean desktop and laptop power increases from 4.2 to 5.8, making laptops appear to be an even more attractive choice. In our department, the one-time action of replacing desktops with laptops would save significantly more energy than any of software techniques.

3.4 Why We Could Be Wrong

While our analysis above indicates that a green computing enterprise should simply purchase laptops rather than employ software techniques, it makes simplifying assumptions which might not be true in practice:

4

Approach	Annual $	$ Saved	% Saved	Approach	Annual $	$ Saved	% Saved
Status Quo: desktops	$46.00	–	–	Status Quo: desktops	$100	–	–
Manual	$31.30	$14.70	32%	Manual	$78	$32	32%
SleepServer	$18.4	$27.60	60%	SleepServer	$40	$60	60%
LiteGreen	$12.00	$34.00	74%	LiteGreen	$26	$74	74%
Laptops	$11.50	$34.50	75%	Laptops	$17	$83	83%
Laptops+LiteGreen	$3.00	$43.00	93%	Laptops+LiteGreen	$5	$95	95%

(a) Computations based on Energy Star   (b) Computations based on empirical data

Table 1: Hardware efficiency (using laptops) offers higher savings compared to reducing usage via software techniques. The combination of efficiency and software achieves the highest relative savings, but adds little absolute savings.

- **What about thin clients?** While laptops are a great candidate for reducing energy costs, other options exist. A mixed software and hardware approach, taking advantage of thin client computing, might save more energy.

- **Can laptops handle desktop workloads?** Laptop hardware is usually less powerful than desktops. It could be that replacing desktops with laptops would harm productivity because they are not powerful enough.

- **Do laptops save money?** This initial analysis only considers power. If laptops cost significantly more than desktops, then the dollar savings from their energy savings might be undone by up-front purchase price.

- **Why not combine techniques?** If, at the end of the day, we show that for some types of workloads and equipment, laptops can provide both energy and monetary savings, it does not necessarily mean that software techniques should be thrown away. If software is still valuable, what benefits can we expect if it is applied on top of this new, more efficient hardware

- **Does energy consumption of computing systems matter at the building level?** The energy usage of enterprise computers is important at the national level [2]. It might turn out that computers are not the only component we should focus on compared to other equipment such as networks, LCD screens, and compute clusters.

The following sections examine each of these questions in turn, using careful analysis of power measurements, utilization characteristics, and market data.

4   Thin Clients

Desktops and laptops are part of a distributed computing infrastructure. While most enterprises today rely on each person having their own personal machine, there are alternative approaches. For example, thin client systems use lightweight terminals that have no local computation; they simply display graphics and handle user input. In a thin client setting, back-end servers centralize the computation for many end-user devices. From an energy perspective, the thin client approach appears promising; clients require little power, and consolidating workloads leads to higher utilization on multiplexed servers. Thin client systems represent an extreme version of systems such as LiteGreen: the user VM never migrates off the centralized server.

To evaluate whether thin client solutions would be more energy efficient than simply using laptops, this section presents data from thin client systems from two different universities, one in the United States and one in Germany.

4.1   United States VMWare Deployment

Our first data set comes from an administrative department on our university campus that handles paperwork associated with research grants. The transition to thin clients was prompted by a desire to have a greener infrastructure. Currently, admin workloads are supported by a mixture of Dell thin clients/servers, desktops, and a few laptops. Of those, we collected power measurements for 12 clients, 6 desktops, 2 laptops, and 2 servers.

The desktops are a selection of Dell machines with average power draw of 55 to 100 watts. These measurements reinforce those in Section 3 in that desktops handling identical workloads can have a factor-of-two difference in their energy footprint and therefore energy cost.

Machine	# VMs	Min	Max	Avg
Server 1	21	311W	373W	328W
Server 2	23	332W	410W	348W

Table 2: Power draw statistics for two servers handling 44 virtual desktops via VMWare's ESX Server OS. Each VM corresponds to a user with a 15-watt thin client at their desk. The average total per-user computing cost is 30 watts.

Two members of the department have opted for laptops, presumably for mobility, as their workload involves meetings and presentations in addition to regular office applications. On average, the laptops draw between 12 and 15 watts, immediately resulting in an energy reduction in comparison to the desktops of their co-workers.

The remainder of our power sensors measure the thin client setup. The deployment includes Dell FX-160 Diskless thin clients, together with two Dell PowerEdge servers running VMware ESX Server. The servers have two quad-core, 3GHz CPUs and 32GB of RAM. Each user has their own virtual desktop. The system load balances these VMs, migrating them between the servers if necessary.

The thin clients all average between 15 and 17 watts, with a standard deviation of less than 1 watt. The thin clients, despite having no real computational capabilities, draw the same amount of power as laptops. The two servers handle a total of 44 thin clients. To calculate the power overhead due to servers, we measured their power draw without any VMs running, as well as over several weeks under normal use. Table 2 shows power data from each server as well as the number of virtual desktops running on each machine. The collected data shows that the server overhead for each client is roughly 15 watts, 18 watts at peak.

The result is an average per-user power cost of 30 watts, 50% more than laptops. Furthermore, users cannot take their thin clients to meetings or otherwise use them to work while traveling.

4.2 German Computer Science Lab

The above measurements suggest that thin clients are inferior to laptops. But we reached this conclusion from measurements on one isolated thin client setup.

To validate the results from the American university and see whether another thin client setup is more efficient, we measured the computer science department at a German university. The department has been using a thin client setup for the past ten years. The current deployment features 26 Sun Ray clients [35] of types 1, 1G, and 2FS to support professors, researchers, students, and admin staff. The clients are connected to two Sun Fire X4200 servers and client virtual machines (VMs) migrate between the two servers based on load balancing and redundancy policies. The servers are backed by an HP ProLiant DL360 G7 NFS file server and a storage unit that provides 7.2TB of secondary storage.

We measured the power draw of four clients and one of the two servers in the setup with the help of four Plugwise [5] power meters. We collected data over a period of three weeks. The server power draw is 302 watts on average, with a standard deviation of 4 watts. We monitored the CPU utilization of the server in parallel to its power draw. While load is highly dependent on the time of day, power draw remained almost constant, with peak power draw of 345 watts.

The power draw of the clients is also almost constant; the newer 1G and 2FS clients draw 10 watts and the older type-1 clients draw 18 watts. The file server and attached storage unit draw another 325 watts.

These measurements mean that on average, the backend infrastructure draws 627 watts to support 26 clients, 24 watts per terminal. We assume the best case scenario, when the system has 32 users, its maximum intended load. With 32 users, per-client overhead goes down to 20 watts. Adding to the 10-18 watt client draw, the power budget per user is 30 to 38 watts: the European thin client setup has a higher energy cost than the American one. As with the U.S. thin client system, this number is lower than traditional desktops, and it is still higher than laptops.

4.3 Conclusion

While competitive with laptops in terms of power, they also do not perform any computation. In contrast, there is tremendous engineering pressure to improve the energy efficiency of laptops in order to maximize lifetime on batteries. Given than thin clients do not have such pressure, it seems unlikely that they will become competitive. Furthermore, supporting many clients requires large, powerful servers. Even when shared across a dozen or more clients, server energy consumption is on par with per-client laptop figures.

From a pure power standpoint, we conclude that thin clients are inferior to a laptop-based environment. Of course, there are some practical benefits to thin clients, which may make them attractive. A thin client setup has the advantage of redundancy and centralized management – considerations that may sway a purchasing decision. Nevertheless, thin clients are not a silver bullet: a laptop-based system is more efficient.

	Percentile CPU		
Machine Type	5^{th}	50^{th}	95^{th}
high-end custom-built	0%	1%	57%
Dell Optiplex 745	1%	9%	58%
Dell Precision T3400	0%	4%	29%
Dell Precision T3400	0%	1%	13%
Dell Inspiron 530	1%	1%	8%
HP Pavilion Elite m9250f	0%	0%	25%
Dell Precision T3400	0%	1%	7%

Table 3: CPU utilization for 7 student machines collected over 11 months reveals high under-utilization.

	Percentile CPU		
Machine Type	5^{th}	50^{th}	95^{th}
Dell OptiPlex SX 280	0%	0%	10%
Dell OptiPlex SX 280	0%	0.75%	5.45%
Dell OptiPlex 745	0%	1.55%	9.25%
Dell Dimension 9200	0%	0.75%	3.1%
Dell Precision 690	0%	0.7%	3.9%
Dell Dimension 9200	0%	1.55%	7.7%
Dell OptiPlex 760	0%	0%	5.45%
Dell OptiPlex 760	0%	1.55%	16.9%

Table 4: CPU utilization for administrative staff machines recorded once per second for a month.

5 Utilization

If the argument is that the desktops described in Section 3 can be replaced with laptops, there are two main questions: 'Can a laptop meet user workloads as well as the desktop?' and 'Will the total cost of ownership of the new hardware scenario indeed save money?' We now tackle the first of these questions.

There are a variety of ways to measure computer utilization, from more coarse metrics such as CPU utilization to detailed observations of computing resources correlated with user input. For the purposes of this paper we want to understand utilization well enough to determine what energy conservation techniques are possible.

We present two different views of utilization: CPU load, as a coarse-grained measure of computing needs, and specific workloads, as described by the list of system processes.

5.1 CPU

Since the computing needs of students and research staff are likely to differ from those of administrative staff, we consider these separately. Tables 3 and 4 show the CPU utilization of a number of student and staff desktops. As expected, students require more processing power, but even so, the median CPU usage does not exceed 10%. Looking at the staff data, utilization is even lower, with a 95^{th} percentile between 3% and 16% CPU.

Furthermore, the most power-hungry staff desktop, drawing over 150 watts, has the lowest CPU utilization – 3.1% 95% of the time. Anecdotally, the user of this machine is a recent department hire and with a newly-purchased desktop. The result was a machine that draws almost 2.5 times more power than older machines which handle the same type of work. Generally, people tend to upgrade to more and more powerful machines, yet typical workloads hardly tax the available CPU resources.

5.2 Workloads

As hinted above, CPU utilization does not capture the full story. A more intuitive way of understanding whether utilization matches the type of equipment we buy is to look at typical tasks users perform. The rest of our data analysis concentrates on staff computing because it is more representative of a general, enterprise computing environment. We collected workload data from four Windows machines by logging the list of processes every second.

Table 5 shows the most popular workloads on administrative machines, excluding various Windows services, virus checks, and the occasional game. The popularity of a given process is calculated as the percentage of samples in which it appeared, or in other words, the cumulative time over 1 month that the process was running. Instead of averaging out the results over the four machines, we give a range. In addition, we show range of averaged CPU percentage that each process consumed. For example, we detected that Acrobat Professional was active 1% to 4% percent of the time, and it consumed 5% CPU on average. One percent of active time is roughly 7 hours accumulated over one month.

This workload data once again raises the question of mismatched user needs and technology. There is no reason why an entry level laptop cannot perform the same basic tasks (document editing, web browsing, PDF viewing) as a quad-core, 150-watt desktop.

6 Economic Feasibility

Enterprise procurement policies primarily aim to optimize the total cost of ownership (TCO). While we have evaluated how much laptops can reduce energy costs, it is still widely believed that they are too expensive to pur-

Machine	Processor	Memory	Price	Avg Power	1-yr TCO	5-yr TCO
Dell OptiPlex 760	Intel Core 2 Duo (2.6GHz)	3GB RAM	$509	64 W	$573	$827
Dell OptiPlex 780	Intel Core 2 Quad Q9400 (2.66GHz)	3GB RAM	$688	72 W	$760	$1046
Dell OptiPlex 980	Intel Core i5-660 (dual core, 3.33Ghz)	4GB RAM	$999	96 W	$1094	$1476

Table 6: The technical specification and TCO of the desktops currently used in our department. The third machine is the most recent IT hardware recommendation.

Process	% of time active	CPU
Acrobat Professional	1% to 4%	5% to 6%
Firefox	0.5% to 4%	2% to 10%
Internet Explorer	0.3% to 2%	3% to 12%
MS Excel	1% to 2%	2% to 8%
Thunderbird	0.4% to 1.2%	2% to 4%
MS Word	0.2% to 0.8%	2% to 17%
Outlook	0.4%	4%
Acrobat Reader	0.3%	5% to 15%
Explorer	0.01% to 0.3%	2..5% to 9%

Table 5: 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.
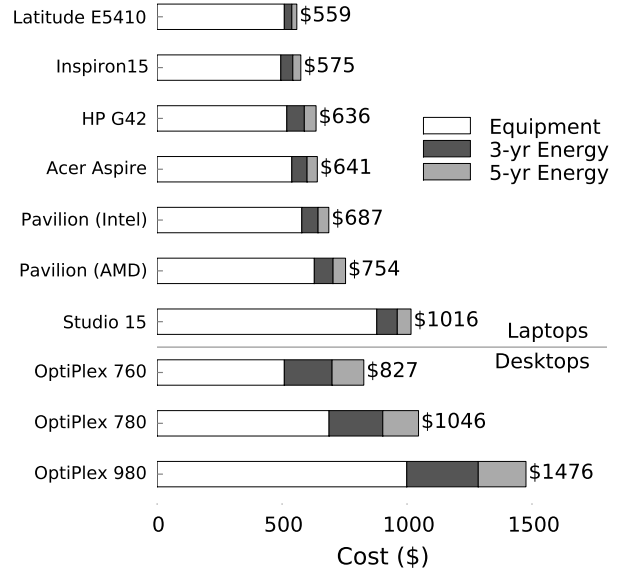

Figure 3: Breakdown of equipment cost for current desktops and their laptop alternatives. All but one laptop have lower purchase price than the desktops. The lower electricity cost due to higher energy efficiency causes the 5-yr TCO for laptops to become significantly lower than for the desktops.

chase. This section examines the interplay of energy-efficiency and cost of equipment.

6.1 Desktop TCO

The total cost of ownership of a machine is the sum of all the costs incurred from purchasing, updating, supporting, and managing the machine, including the cost for power over the lifetime of the machine. For the purposes of our analysis, we concentrate on the sum of purchase price and energy cost and refer to it as TCO.

Like many universities, our university has a purchasing policy that discourages equipment diversity. Dell is one of the preferred suppliers and the policy recommends a few models. The two most common desktop models used in our building are Dell OptiPlex 760 and Dell OptiPlex 780. The energy costs are 38% and 34% of the 5-yr TCO for these two popular models as shown in Table 6. For future purchases, the policy recommends a Dell OptiPlex 980 system. The energy cost for this model is expected to be 32% of its 5-yr TCO of $1476. Thus, we expect energy to contribute at least 30% to the TCO of desktops in the foreseeable future.

6.2 Laptop TCO

Our analysis of user workload in Section 5 found that most desktops are underutilized and can be replaced with laptops. Replacing desktops with laptops will only reduce the per-user cost, if we can find laptops that a priced similarly to desktops – a task that has not been easy in the past.

We identified several laptop models that can handle the workloads seen in our building. We obtained the purchasing price (as of March, 2010) and power draw data of these laptops from both the manufacturer's websites and

---

[1]The power draw of the first two laptops includes the LCD screen at 70% brightness, while the rest had the screen set to the lowest brightness level.

Machine	Processor	Memory	Price	Avg Power[1]	1-yr TCO	5-yr TCO
Dell Latitude E5410	Intel Core i3 380M (2.53GHz)	3GB	$509	10 W [33]	$519	$559
Dell Inspiron 15	Intel Perntium Dual Core (2.3GHz)	4GB	$495	16 W [30]	$511	$575
HP G42 Notebook	AMD Athlon II (2.1Ghz)	3GB	$519	23.5 W	$542	$636
Acer Aspire Notebook	AMD Athlon II P320 (2.1GHz)	3GB	$539	20.4 W	$559	$641
HP Pavilion dv6	AMD Turion II P520 (2.3Ghz)	4GB	$579	21.7 W	$601	$687
HP Pavilion dv6	Intel Core i3-350M (2.26GHz)	4GB	$629	25 W	$654	$753
Dell Studio 15	Intel Core i3-350M (2.26Ghz)	4GB	$879	27.5 W	$906	$1016

Table 7: The technical specification and TCO of the laptops proposed as desktop replacements. Equipment prices of laptops are converging to those of desktops.

a local electronics store. For the two models sold online, we use power measurements from battery performance reviews. We used our equipment to measure power draw of the remaining models.

Table 7 shows these laptops, their power draw, and their 5-year TCO. We notice that laptops prices have gone down, approaching those of entry-level desktops. Although the recommended laptops have similar purchase price as the recommended desktops, the 5-year TCO for laptops is much lower than for the desktops due to the large energy use reduction. Figure 3 shows the relative comparison of desktop and laptop cost.

The, somewhat surprising, conclusion is that choosing laptops makes economic sense in addition to greening the computing infrastructure.

7 Software on Top of Hardware

Looking back at the high-level data in Section 3, Table 1, one may wonder why this paper does not advocate for a mixture of hardware efficiency and software techniques. After all, applying something like SleepServer or Lite-Green on top of laptops will reduce the energy usage even more, than simply using laptops.

What happens when we apply software to an already hardware-efficient computing infrastructure? Even though the percentage of additional cost reduction will be high, 60% to 74% for the example approaches, the absolute dollar savings will be low. For example, applying LiteGreen on top of Energy Star laptop data saves $8.50 a year, per machine and $12 for our empirical enterprise measurements.

This means that if the software requires more than roughly 15 minutes of attention from IT staff, or harms a user's productivity by more than a half hour, the savings are a net *loss*. Section 9 discusses how our future research on software techniques can be re-targeted to provide higher energy-efficiency rather than trying to build

What?	% of Building	Monthly Cost
Network	3%	$1,080
Displays	5%	$1,800
PCs	20%	$7,200
Servers	27%	$9,720
Other[1]	45%	$16,200
Total	**100%**	**$36,000**

Table 8: Power measurements and survey techniques allow us to compute an estimated breakdown of our building's electricity consumption. We find that computers contribute 20%.

an array of power-saving solutions with little impact.

8 Whole-building Energy

An important question to answer is whether reducing the energy consumption of computing systems is a worthwhile problem from the standpoint of an individual enterprise building. Our department pays roughly $36,000 for electricity every month; if only a small portion of that goes toward computing, then maybe other parts of the computing infrastructure should be the priority.

We use our extended set of power measurements and survey techniques similar to those in Kazandjieva et al. [32] to estimate contribution of computing systems relative to the total building electricity consumption. Table 8 summarizes the results. We find that computing systems are responsible for roughly 55% of the bill, with PCs contributing 20%, or over $7,000 a month; hence, computers are a considerable fraction of our building's budget, and worth optimizing.

But what about the other 35% due to computing? It turns out that server machines are an even larger por-

tion of the electricity budget, about 27%. This finding is not surprising since the context of our measurements is a computer science department with two large server rooms located in the basement. However, there are implications for both enterprise environments like ours and for ones without servers. If a building has a server room, those machines are likely to contribute significantly to the energy bill of the whole building, and work in the data center domain will have relevance to the enterprise. Section 9 tackles this further. For server-free buildings, the implication is that the relative share of PC energy will go up.

The remaining 8% of electricity use due to computing comes from LCD screens (5%) and networking equipment (3%). Our measurements suggest that even though a 30-inch LCD can draw as much as 140 watts, its energy footprint is much lower than that of a 140-watt desktop due to effective duty cycling: our year-long measurements show that user's monitors go to sleep after a period of user inactivity and are turned off at the end of the work day. There is little place for improvement, unless the future sees drastic changes to LCD technology.

Last but not least, networking equipment certainly has room for improvement, since switches draw constant power regardless of traffic demands. However, compared to PCs, the network is a relatively small contributor to the overall electricity consumption.

In summary, we find computing is responsible for a large portion of a building's electricity costs. In our department of 600 people, PCs are a fifth of the bill.

9 Looking Forward

The data and analysis in this paper indicates that improving efficiency is a more effective path to greening enterprise computing than reducing use. In this way, enterprise computing is like almost every other energy domain.

This section explores the implications of this observation, considering a computing landscape where desktops disappear and enterprise personal computers become energy-efficient, mobile devices. We separate our discussion of this possible future into four broad topics: green computing in the enterprise, the relationship between the enterprise and the cloud, future research directions these changes might bring, and policy changes they might suggest.

9.1 Green Enterprises

The fundamental issue with our current desktops is that they consume all this energy when no work is being done, because the machine is idle. Laptop-class and low-power devices shine in this area, because they currently consume less active power. Of course, being low-power also means you can also work off of a battery, providing greater flexibility. Batteries, however, place a reasonably hard constraint on the power draw of a mobile device. Unless battery energy density goes up – lithium ion is already very high – we can't put more energy in portable devices. The demand for greater runtime exerts a steady downward pressure on both idle and active power. Correspondingly, it seems unlikely that end-user devices, once they are mobile, will see power draw that creeps upwards anytime soon.

Our implication is that end-user computing devices, in the enterprise domain, will *never* become a significant-enough energy consumer to justify usage software usage reduction techniques such as network proxies. In this domain, we should improve efficiency, not reduce usage. More broadly, we should re-architect our systems to allow for more flexible computation rather than try to build power-saving systems for niche issues.

We saw in Section 5 that these mobile devices can meet the majority of enterprise workloads, but a concern worth addressing is that applications will occasionally benefit from greater performance than what a single mobile device – a laptop or a tablet – can provide. In this vein, of re-architecting the entire desktop computing ecosystem, can we envision a mixed system that supports typical usage efficiently, yet provides access to greater resources when needed? Alternately, perhaps using the laptops as thin clients to access virtual desktops on shared servers is a better approach. In a third scenario, the browser is the primary application. In each of these scenarios, back-end servers for off-loading computation take a greater role; we call this approach "hybrid computing", where a laptop does more than a thin client but still leverages remote resources when available and appropriate.

9.2 Enterprises and the Cloud

As these personal computers begin offloading work to resources in the cloud, demands on data centers will rise, both to on-campus private data centers and massive, off-campus, public ones. The wrinkle is that enterprise workloads correlate strongly with workday hours. This creates a tension between the need to place resources nearby to support interactivity (i.e. placement constrained to close time zones for speed-of-light propagation reasons and with high bandwidth) and the desire to spread load evenly (i.e. reducing peak provisioning to increase average utilization and reduce infrastructure costs). If these data centers must be provisioned more to support nearby enterprise workload during the day, what should be done with these spare resources at night? And,

more relevant to this paper, how might savings from a focus on efficiency compare to those from reducing usage?

One efficiency approach would be to simply buy newer data center hardware. This is already happening, driven by the economics of both energy usage and provisioning at scale. Another efficiency approach would be to sell unused capacity, a la Amazon EC2 Spot instances, where the price varies with load, to encourage flexible computations to be done on off-peak hours. This strategy is preferred whenever the offered price for usage exceeds the cost of electricity. However, the security costs of managing spot capacity count against this approach. Furthermore, while the propagation delays of one or two time zones might not be significant, it seems unlikely that users in Finland will offload highly-interactive applications to Australia. Instead, this spare capacity rotating with the shadow of the sun can sustain jobs that are not as latency sensitive. This approach would improve efficiency without reducing usage.

The alternative, reducing usage, would be to dynamically adjusting capacity to demand, so that unneeded components can be powered down. Server consolidation, combined with VM migration and a data center manager, can be used not only to balance VMs across physical machines, but also to compress and shut down unneeded physical machines. Today, needing to transfer the state of entire VMs presents a large barrier to rapid migration. But, if this approach became practical and trusted, it could help with campus power efficiency, especially in settings (hospitals, universities, companies) that for security reasons may not be willing or able to use a public cloud. Within a virtualized data center, central control over hardware resources is common, and these software approaches begin to make sense.

This places an interesting counterpoint. Because enterprise devices are personal, reducing use is ineffective in comparison to improving efficiency. But the large aggregation of data centers makes reducing use through migration and load balancing attractive. Software techniques will still be important – they will just be important at scale and on back-end systems.

9.3 Future Research

The effect of this shift is that *every* technique which improves overall data center efficiency, regardless of its classification as efficiency or usage reduction, now becomes highly relevant to enterprise computing efficiency. In addition, techniques that help to enable a thinner model of laptop computing – which seemingly have nothing to do with efficiency – can indirectly enable this change to a more efficient desktop computing infrastructure. This shift motivates both research opportunities and challenges in a range of systems areas.

9.3.1 Operating Systems

Hybrid computing represents a chance to make new use of decades of process migration research and multicore process placement. An operating system that aims to support hybrid computing would benefit from knowledge of which applications are better suited for local execution and which processes could be executed faster or more efficiently from the cloud. For example, a video editor working on NFS-mounted data would be likely to benefit from remote execution on a beefier CPU with more cores. A smarter OS might even anticipate the need for migration and make it as seamless as possible; it might even learn from historic user workloads and predict whether an application should start out running locally. As data and code become more distributed, the key role of the operating system is deciding what to cache and where in both the local and wide area. As a counterpoint, the specific OS used in a browser-focused world could become irrelevant – in theory, at least.

In hybrid computing scenarios, state migration might make more sense because its focus is personal workloads and applications, not necessarily server-size data sets or services. The growing relevance of Digital Rights Management (DRM) on mobile platforms such as the iPhone might lead to more trusted applications that can engage cloud resources on your behalf. An OS, sitting on top of the hardware, will be the central control point for managing this security.

One critical aspect of hybrid computing is the wide range of devices and capabilities. While there has been some recent work in multicore operating systems, hybrid systems have the additional complexity of heterogeneity. Correspondingly, techniques like those adopted in Barrelfish may become part of a distributed, hybrid OS [15].

9.3.2 Virtualization

Virtualization is now a standard way to encapsulate the state of a server computer, and it is becoming available for desktops too. Should virtualization software change when run on a battery-powered mobile device? Can it be tuned for mobile processors, which often have less cache? Is it a good idea to migrate a full VM over a wireless channel, and how might that process be optimized? Wireless communication is often less expensive in terms of energy than computation. It could be that mobile devices become wireless thin clients in order to maximize lifetime.

There is ongoing work to optimize the process of VM migration to minimize downtime, possibly even doing the migration in a way that considers network conditions or prioritizes the migration over other traffic. Any technique that helps an application be responsive at a dis-

tance, and more tolerant of packet drops and latency, becomes an enabler to hybrid computing. Optimizing virtualization software to understand and reduce delays induced by hypervisor-kernel transitions would help here [42].

9.3.3 Networks

Anything that reduces the effect of network latency improves hybrid computing. One approach is to hide latency by improving applications that provide interactive access to remote applications, such as NX and VNC [4, 6]. At some point, latency is limited by the speed of light. Even for most interactive applications, with the exception of perhaps some competitive games, the speed of light provides significant flexibility in the geographic and network placement of computation: 1,000 miles is a 10ms round-trip latency.

Current trends in network hardware may add complexity to the situation. The depth of router buffers on a congested network path can have a major effect on observed latency and TCP throughput [13]; this phenomenon has recently been coined "bufferbloat" [9]. Recent work in the IETF to enable bandwidth-intensive, latency-insensitive applications to work well with real-time ones (e.g., the LEDBAT working group) could become critically important [3]. Other work on has network operators provide information about network topology (e.g., the ALTO working group) to assist hybrid applications in placing themselves [1].

9.3.4 Hardware

A future where desktop-class processors are only used almost entirely for multiplexed servers in data centers (again, assume traditional desktops have gone away) implies that the idle power rating of a CPU is no longer relevant. A CPU is unlikely to ever be idle when multiplexed, and power draw at average utilization is more important. As end user workloads are highly bursty, dynamic clock and frequency scaling will be of limited use: instead, sleep modes and their transition times will dominate. In contrast, the fact that cloud systems will always be seeking to operate at capacity means that the opposite is true for them.

9.4 Policy

Hardware and software modifications – combined with a push for power efficiency – change our models of computer usage, and corporate and government policies have to not only adapt but also incentivize these changes. These policies can go further than simple taxes on desktops, subsidies for laptops, or fine-grained billing within an organization.

Mobile work and telecommuting has become commonplace [39], and policy changes have the potential to further encourage the computing equipment TCO savings that these can yield. For example, although work can be taken anywhere on a laptop, without cellular data access, Internet connectivity is not ubiquitous. Hence, corporations might want to subsidize cellular data service for laptops and tablets. This policy does not a represent a big leap from the current practice of subsidizing email access.

Work-equipment policies can become more accommodating for personal use. If work is blending into personal lives, a policy to support the efficient laptop model, to discourage employee desktops left powered on at home, seems reasonable, yet some current policies can strictly forbid this [39]. Companies can set up the office environment to be laptop friendly as well. For example, laptops and desktops should not be exclusive one-time choices. Powerful desktops, large displays, and shared docking spaces could be made available and shared for those days when you really need the computing power.

Another area where policy innovation could spur greater use of mobile computing is in supporting applications for tablets and smartphones. The US government is starting to invest in these apps, with a number of them listed in [11].

10 Conclusion

This paper argues that hardware efficiency is more important than software solutions from both an energy and an economic perspective. Based on this conclusion, we propose replacing desktops with laptop equipment that can meet user workloads at a lower energy cost. We examined energy footprint, utilization, workloads, and ownership costs of real computing, and showed how detailed power measurement data builds the case for a policy of laptop replacement. Two measured examples of thin client computing do not provide savings over laptops, but we find that the centralized paradigm of thin clients has its merits.

However, the specific numeric results of our analysis must be taken with a grain of salt, as they depend on numbers that could change quickly. A new thin client, laptop, or desktop could come out tomorrow and yield a new "winner." We have made our best attempt to measure a representative subset of our computing landscape and consider the issues that might arise from laptop replacement. However, debating the exact power numbers is not our goal. Instead, our goal is to measure our office building and use that concrete data to indicate directions for future systems research, such as operating systems that can migrate individual processes; and tuning virtualization for mobile systems.

11 References

[1] Application-Layer Traffic Optimization (alto). https://datatracker.ietf.org/wg/alto/.

[2] Department of Energy, Annual Energy Review 2008. http://people.virginia.edu/~gdc4k/phys111/fall09/important_documents/aer_2008.pdf.

[3] Low Extra Delay Background Transport (ledbat) Working Group. https://datatracker.ietf.org/wg/ledbat/.

[4] NX Distributed Computing Infrastructure. http://www.nomachine.com/documentation/html/intr-technology.html.

[5] Plugwise Wireless Power Meters. www.plugwise.com.

[6] RealVNC - VNC Remote Control Software. http://www.realvnc.com/.

[7] Sun Fire X4270 M2 Server. http://www.oracle.com/us/products/servers-storage/servers/x86/sun-fire-x4270-m2-server-ds-079882.pdf.

[8] Sun Ray Virtual Display Client. http://www.oracle.com/us/products/servers-storage/desktop-workstations/036088.pdf.

[9] The Bufferbloat Project. http://www.bufferbloat.net/projects/bloat/wiki/Bufferbloat.

[10] Windows Remote Desktop. http://www.microsoft.com/.

[11] Mobile Apps for the Government. http://apps.usa.gov/?v=all, 2011.

[12] Yuvraj Agarwal, Steve Hodges, Ranveer Chandra, James Scott, Paramvir Bahl, and Rajesh Gupta. Somniloquy: augmenting network interfaces to reduce PC energy usage. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, pages 365–380. USENIX Association, 2009.

[13] G. Appenzeller, I. Keslassy, and N. McKeown. Sizing router buffers. *ACM SIGCOMM Computer Communication Review*, 34(4):281–292, 2004.

[14] R.A. Baratto, L.N. Kim, and Jason Nieh. Thinc: A virtual display architecture for thin-client computing. *ACM SIGOPS Operating Systems Review*, 39(5):290, 2005.

[15] A. Baumann, P. Barham, P.E. Dagand, T. Harris, R. Isaacs, S. Peter, T. Roscoe, A. Schüpbach, and A. Singhania. The multikernel: a new OS architecture for scalable multicore systems. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pages 29–44. Citeseer, 2009.

[16] Nilton Bila, Eyal De Lara, Matti Hiltunen, Kaustubh Joshi, and HA Lagar. The Case for Energy-Oriented Partial Desktop Migration. *2nd USENIX Workshop on Hot Topics in Cloud Computing*, 2010.

[17] Jimmy Carter – Proclamation 4681: International Energy Conservation Month. http://www.presidency.ucsb.edu/ws/index.php?pid=31312#axzz1GjB6zYt0, 1979.

[18] K. Christensen. Improving the Energy Efficiency of Ethernet-Connected Devices: A Proposal for Proxying. *Ethernet Alliance*, 2007.

[19] Tathagata Das, Pradeep Padala, Venkata N. Padmanabhan, Ramachandran Ramjee, and Kang G. Shin. Lite-green: saving energy in networked desktops using virtualization. In *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*, USENIX-ATC'10, pages 3–3, Berkeley, CA, USA, 2010. USENIX Association.

[20] Energy star. http://www.energystar.gov/, 2011.

[21] The Energy Policy Act of 2005. The United States Congress, 2005.

[22] Success Stories: Energy Star. http://www.energystar.gov/index.cfm?c=power_mgt.pr_power_mgt_ss.

[23] Energy Star Version 5 Specification. http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/Version5.0_Computer_Spec.pdf, 2009.

[24] Strengthening Federal Environmental, Energy, and Transportation Management. US Executive Order 13423, 2007.

[25] Federal Agencies' EPEAT Purchasing Takes Off. EPEAT Press Release, April 2009.

[26] Philip Levis et al. T2: A Second Generation OS For Embedded Sensor Networks. Technical Report TKN-05-007, Telecommunication Networks Group, Technische Universitat Berlin, 2005.

[27] Gerald T. Gardner and Paul C. Stern. The Short List: The Most Effective Actions U.S. Households Can Take to Curb Climate Change. *Environment: Science and Policy for Sustainable Development*, 50(5):12–25, September 2008.

[28] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection Tree Protocol. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys'09)*, November 2009.

[29] Jonathan W. Hui and David Culler. The dynamic behavior of a data dissemination protocol for network programming at scale. In *Proceedings of the Second ACM Conference on Embedded networked sensor systems (SenSys)*, 2004.

[30] Dell Inspiron 15 Review. http://www.notebookreview.com/default.asp?newsID=5165&p=2, 2011.

[31] Xiaofan Jiang, Stephen Dawson-Haggerty, Prabal Dutta, and David Culler. Design and implementation of a high-fidelity ac metering network. In *The 8th ACM/IEEE International Conference on Information Processing in Sensor Network*, San Francisco, CA, USA, 2009.

[32] Maria Kazandjieva, Omprakash Gnawali, Brandon Heller, Philip Levis, and Christos Kozyrakis. Identifying Energy Waste through Dense Power Sensing and Utilization Monitoring. Technical Report CSTR 2010-03, Stanford University, 2010.

[33] Dell Latitude E5410 Review. http://www.notebookreview.com/default.asp?newsID=5897&p=4, 2011.

[34] Sergiu Nedevschi. Skilled in the Art of Being Idle : Reducing Energy Waste in Networked Systems. *NSDI*, 2009.

[35] Oracle. Sun Ray 2 Virtual Display Client. *http://www.oracle.com/us/products/servers-storage/desktop-workstations/030726.htm*, pages 1–4, 2010.

[36] Joshua Reich, Michel Goraczko, Aman Kansal, and Ji-

tendra Padhye. Sleepless in Seattle No Longer. *USENIX Annual Technical Conference*, 2010.

[37] Power Management Software Helps Slash Energy Costs. http://www.windowsitpro.com/article/news2/power-management-software-helps-slash-energy-costs.aspx, 2008.

[38] After-hours Power Status of Office Equipment and Inventory of Miscellaneous Plug-Load Equipment, . Lawrence Berkeley Laboratory, 2004.

[39] Telework Program Policy Guide. Home & Communities, US Department of Housing and Urban Development, November 2008.

[40] Tim Miles. Green IT. US Department of Commerce, International Trade Administration, May 2010.

[41] Watt's up internet enabled power meters. https://www.wattsupmeters.com/secure/products.php, 2009.

[42] J. Whiteaker, F. Schneider, and R. Teixeira. Explaining packet delays under virtualization. *ACM SIGCOMM Computer Communication Review*, 41(1):38–44, 2011.

[43] Rajesh Gupta Yuvraj Agarwal, Stefan Savage. SleepServer: A Software-Only Approach for Reducing the Energy Consumption of PCs within Enterprise Environments. *USENIX Annual Technical Conference*, 2010.

APPENDIX

A Data Gathering Methodology

The data in this paper is a result of a larger scale effort to characterize the energy consumption of enterprise computing systems. Our sensor deployment collects power and utilization data for the computing infrastructure of an academic building with students, professors, and administrative staff. Currently, the deployment includes 250 single-outlet, wireless and wired meters that sample device power draw once per second. Each meter connects to exactly one computing device, such as a PC, laptop, display, or network switch. For the purposes of this paper, we concentrate on the data collected from computers in the building – about 90 sensing points. The duration of the measurements varies from 7 to over 15 months depending on when the power meters were installed.

In addition to the main building deployment, we have two more monitored environments, both providing data on thin clients and the servers associated with them. The first of these deployments uses the same wireless technology as our department's power meters, for a total of 16 meters. The second is located in Europe and provides more limited data from four Plugwise [5] off-the-shelf power meters, sampling once per second. These meters were purchased to accommodate the European electrical system at 220 V.

Next, we provide more detail on the power and utilization sensors, the and back-end infrastructure.

Wired Meters. Commercially-available Watts Up .NET meters [41] transmit measurements over Ethernet, up to once a second, over the existing building network. Each meter posts data via HTTP to a server process on the back-end server. These meters were a useful first step in gathering power data, though the practical issues of scaling to hundreds of devices, high cost, and proprietary software hindered further deployment. By now many of the meters have either failed or have been upgraded to custom-made wireless sensors. About 50 nodes remain active, in a sparse deployment around the building, covering faraway network closets and student offices.

Wireless Meters. In contrast, the wireless deployment is dense, covering a large fraction of the power outlets on the second, third, and fourth floor of out computer science department building. These custom-made low-power wireless meters transmit data via an ad-hoc multihop network. Each meter is a modified version of the open-source ACme meter [31].

The meter software, built on TinyOS [26], includes sampling, routing [28] and dissemination [29] capabilities. The top-level application reads power draw every second and sends a data packet after buffering ten samples. To our knowledge, this deployment is the largest, longest-term, and highest-density indoor wireless sensor networks. As such, the deployment has provided immense amounts of data, characterizing the energy footprint of a large portion of the building. We will openly share these datasets with the community.

Utilization Monitoring. In addition to power, we also monitor the utilization of a number of desktop computers in the department. We collect the raw CPU utilization via a script that runs either as part of Cron (Linux) or as a Windows service. For a limited number of PCs we also record the lists of active processes. The data is collected once per second. Utilization monitoring is key to understanding what portion of the energy in enterprise computing is wasted due to idle and underutilized machines. It is also key for observing that sacrificing limited amount of performance, by buying laptops, may be acceptable.

Data Storage, Access and Analysis. The wired and wireless meters, and software workload monitors send the data to a central server with two 1.8 GHz cores and 2 GB of RAM. With over 250 sensors reporting as often as once per second, data piles up quickly. Currently, we have 93GB of utilization data and over 230GB of power data. This data stream provides near-real-time feedback to building residents, equipment purchasers, and system administrators through a website. A TV display in the building lobby provides information about the project, along with graphs showing real-time power consumption of categories of devices, such as monitors, servers, and network equipment.