

Energy Dumpster Diving

Maria A. Kazandjieva, Brandon Heller, Philip Levis, Christos Kozyrakis
Computer Systems Lab, Stanford University
Stanford, CA 94305, USA.
{mariakaz,brandonh}@stanford.edu, {pal,christos}@cs.stanford.edu

Abstract

Power data alone cannot identify sources of energy inefficiency. However, correlating power data with utilization statistics can reveal where power is used well and where it is wasted. We describe a sensing infrastructure, PowerNet, that monitors power and utilization in a building environment. The deployment includes both wired and wireless sensors and covers offices, networking closets, and server racks. We present PowerNet’s architecture, then generate initial insights from each monitored environment. Analyzing PowerNet data traces identifies contexts where electricity consumption can be reduced without cost, and others which call for rethinking system designs altogether.

1. INTRODUCTION

Electricity powers our everyday computing, encompassing desktops, servers, networks, and more. When studying a large computing infrastructure, a single bill does not show how much power each device consumes. Therefore, a number of power meters have appeared on the market [2, 6], providing a detailed breakdown of electricity usage by device, category, floor, or user [3, 4, 9].

Still, power data alone is not enough to expose device inefficiencies. Utilization information correlated with power consumption can reveal system designs that are not energy-proportional and usage scenarios that waste energy.

To understand the efficiency of computing systems, we have designed and deployed PowerNet, a monitor-

ing infrastructure that collects correlated power and utilization measurements. The current deployment includes 85 power meters reporting data for the past six months. In addition, 15 desktops and 10 servers are reporting CPU load, and five switches measure network traffic. The collected data allows us to examine case studies from an office environment, a data center server rack, and a small networking closet, all subjected to real use from students and faculty in the Stanford Computer Science building.

The preliminary observations that PowerNet data has enabled are only the beginning of grasping power usage in large-scale computing systems; our end goal is comparing high-level system choices. For example, are PCs more efficient than thin clients providing access to server-hosted virtual machines? Should an office be networked using wiring closet switches or a dense mesh of wireless access points? How should a cluster store data and schedule processing to minimize power usage? By measuring today’s infrastructure, we can inform future designs.

In our energy dumpster diving excursion, utilization data has helped sift through power that was well-used versus power that was wasted away. So far, we have identified that monitor configuration can reduce consumption by as much as 25% and that identical server machines can have different power draw, not explained by differences in load. We have collected desktop data that can help decide when machines should enter sleep mode, and have examined network switch energy-proportionality. The rest of this paper dives into more details on how and why computing systems use power.

2. POWERNET INFRASTRUCTURE

PowerNet is a large-scale distributed sensing infrastructure that provides per-device energy and usage statistics in an office-building environment. Figure 1 shows the main components of PowerNet: power meters, utilization modules, storage system, and public data access interface.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 2009 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

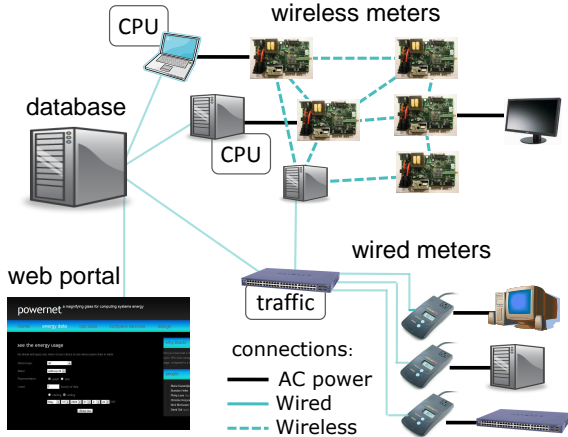


Figure 1: PowerNet architecture for power and utilization monitoring.

2.1 Power Meters

Plug-level power meters sit between the electrical outlet and the system being measured, providing the finest granularity of data outside of modifying each device. PowerNet includes two types of plug-level meters – wireless and wired. These meters can characterize consumption at the device, category, room, user, and floor levels without interfering with devices’ everyday use. One type of meter can validate the other, by connecting both meter types to the same devices. Data validation is important for any large-scale sensing deployment [14], and our tests have confirmed that the meters produce consistent readings.

2.1.1 Wired Power Meters

The Ethernet-enabled Watt’s Up .NET meters, commercially available for \$235, provided a fast way to bootstrap the PowerNet deployment. Each meter has a configurable sampling rate of up to 1Hz and accuracy of +/- 1.5%.

Unfortunately, the practical difficulty of finding available Ethernet ports, combined with the administrative burden of configuring the meters on the Stanford network, slowed the deployment. Furthermore, the power consumption of the infrastructure is about 3 watts per meter, which is excessive - and as many pointed out, against the spirit of the project.

2.1.2 Wireless Power Meters

The limitations of closed, expensive, wired meters make them unattractive candidates for larger-scale deployments; to simplify future deployments, we designed custom wireless sensors.

The communications portion of these meters includes a low-power processor (1mA when active, 1uA

when asleep), radio (2.4 Ghz unlicensed spectrum, 802.15.4-based), and an integrated antenna. The sensing portion includes current and voltage sensors, plus a digital power meter chip that multiplies the sensor values to get an instantaneous power reading, accumulating the readings to determine energy [1]. The meters can sample at 14 KHz, enabling harmonics and power quality analysis, as well as controlled experiments where device utilization varies faster than once per second.

The second, and current, board revision is a result of our experiences from the first PowerNet boards, combined with lessons learned by colleagues at UC Berkeley working on a similar device, the ACme [9]. Revision 2 includes an expansion port with a range of serial interfaces to support additional sensors and storage, such as dust sensors and flash memory. The new boards cost about \$120 apiece.

2.2 Utilization Metering

PowerNet employs software to monitor device use. Volunteers run a Python script that tracks machine CPU utilization. Similar data is monitored on the cluster server machines. The switch data is obtained through available hardware counters and monitored via SNMP. The logs include CPU load and traffic statistics for each port.

2.3 Data Collection

The PowerNet motes run custom drivers and software on top of TinyOS [10], an operating system for low-power wireless devices. The wireless meters form a mesh network and route power readings to PC base stations. The software stack is still being fleshed out and will include collection-tree routing [5] to data sinks, binary dissemination [8] for updating software, and user applications for controlling how and what data is sampled.

All data is synchronized by NTP timestamps and collected at a central server. Normally, the wired meters report to the company’s website, but we point them to a custom server process that logs readings to a MySQL database. The wireless meters route data to a base station, which then sends TCP packets to the same database.

2.4 Data Availability

At the front-end of the PowerNet infrastructure sits a publicly available web server that makes sensor readings accessible. Line-chart visualizations can show an arbitrary time interval of power and utilization data, for one or more sensors. One can pull up data by the meter name, type, or device category. Most importantly, graphs show the correlated

power and utilization data, making it easy to identify energy-proportional devices and idle machines wasting power. In addition, the page displays up to date information on PowerNet’s status and the cumulative energy consumed by all monitored devices. In the coming months, we will be installing a display in the lobby of the Computer Science building to present findings and suggestions for more efficient use of the computing infrastructure. In addition to visual data representations, the full sensor output can be directly downloaded as a text file. This ensures that all data is available to researchers outside our group.

3. INITIAL INSIGHTS

Analyzing the power and utilization data collected through PowerNet has revealed insights into the Gates Hall computing infrastructure. The following subsections present several preliminary case studies.

3.1 Case Study: Desktops

PowerNet data has enhanced our understanding of desktop power consumption. Data gathered over five offices shows that desktops draw between 100 and 300 watts. In addition, the high idle energy, as shown in Figure 2, validates the need for putting desktops to sleep when they are not actively used. In fact, Stanford IT services already encourages students to download software that puts machines to sleep when idle. Devices such as the network proxies presented in [7] go a step further and maintain network reachability even when desktops are in sleep mode.

While such techniques reduce energy consumption, it is not clear exactly when they should be employed. In [11], lack of keyboard and mouse input signifies an idle machine, but this is not sufficient: both user and CPU should be idle. Figure 2 shows power and CPU load for a Dell desktop with 4 Intel cores. There are active and idle periods spread over the entire 24-hour measurement period, suggesting that some machines might not be suited for simple sleep schedules. Large amounts of correlated power and utilization data can be used to create machine usage models. There is already interest in using the PowerNet datasets in conjunction with machine learning algorithms to predict when it is prudent to turn machines off.

Figure 2 also exposes the high correlation between power consumption and energy usage. The linear regression R-squared value is above 0.9 and is similarly high for other measured machines, reaching as high as 0.985. The implication is that after factoring out the large baseline energy, power and CPU data can be used to improve existing power modeling techniques [12].




Figure 2: Desktop case study. CPU utilization and power consumption show a strong correlation. We also observe a very high baseline energy corresponding to 0% CPU usage.


3.2 Case Study: Monitors

While power consumption of desktop machines has been studied extensively [13, 7, 11], computer monitors have not received much attention. Through PowerNet we have discovered that the power draw of large monitors becomes comparable to that of desktops. Monitors are often left on, even when users are away from their desks, consuming anywhere between 40 and 130 watts.


We began by collecting power data from seven 30-inch monitors in a students’ office, expecting to see a simple on-off pattern. Instead, there was significant variation in the power consumption when monitors were turned on and power data alone was not enough to explain it. Our hypothesis was that, opposite to common belief, LCD screens require additional energy to display brighter colors.

Device configuration and utilization data could help understand the variations in power consumption. The hypothesis was verified by correlating power data with different desktop color schemes and monitor brightness settings. Figure 3(a) shows data from a controlled experiment on a 30-inch Dell monitor with 14 brightness settings. Even minor adjustments to the brightness setting lead to a large – 25 watts – decrease in power draw. For 11 users, reducing the brightness did not visibly change user experience. A smaller, but still significant reduction in consumption – 9 watts – was observed when color schemes were switched from white to black.

These findings prompted several users to lower their monitor brightness, as well as change their desktop backgrounds. Figure 3(b) shows typical data from one such user who only modified desktop color schemes. The monitor’s power usage is shown over a working



(a) Various monitor settings



(b) Altered user behavior

Figure 3: Brightness level and color scheme have a significant effect on monitor power consumption. Adjusting the settings can reduce energy by 10%-28% without affecting usability.




Figure 4: Network switch case study. There is no correlation between power and traffic.

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 hours a week, 400Wh are conserved. While this number is not large on its own, multiplied by the hundreds of monitors in every office building, it begins to make a difference.

3.3 Case Study: Network Switches

The previous two studies show correlation between how a device is used and the resulting energy consumption. In other cases, data reveals the opposite.

Figure 4 shows data from three network switches – two 1U switches from NEC and one chassis switch from HP. The HP and one of the NEC switches have 23 active ports, and the other NEC switch has 47 active ports. All ports are 1 Gbps. Each data point shows traffic load and power consumption averaged over a 1-minute interval.

The HP consumes 3 times the power of the smaller

NEC, because it has additional fan loads and overhead for the backplane switch fabric. If all 12 of the HP’s linecards were installed, rather than just 1, the per-port power would drop. The NEC switches show per-port energy proportionality, with the 47-port one consuming twice as much power as the 23-port one.

None of the switches show any correlation between usage and energy cost. While the traffic load goes from zero to 10 and even 50 Mbps, power draw remains the same. In other words, network equipment in practice is not energy-proportional, consuming the same amount of resources regardless of how many packets are getting sent.

In this case, the most efficient usage scenario is one in which every switch always handles traffic close to its maximum capacity. This insight, together with measurements of smaller, per-room switches and wireless APs, can inform a more energy efficient network infrastructure design.

3.4 Case Study: Server Rack

The PowerNet deployment monitors ten identical 1U servers. The ten machines are next to each other in a 40-server rack, stored in the Gates building’s data center. Initial power data showed nearly identical readings for all machines, with one noticeable exception consuming 308 watts, 20% more, compared to 245 watts for the other servers.

The utilization data showed that all machines were at idle. Running a demanding, balanced workload resulted in the same increase in power for all machines. Was the special 308W server misconfigured, did it have a malfunctioning component, or was its position in the rack affecting its power draw? To test the latter theory, we swapped the special server, which was at the top of the rack, with the bottom server.

After the swap, the special server's power consumption dropped back to 245W, while its replacement increased from 250W to 270W, confirming that cooling at the top of the rack was an issue.

It is apparent that utilization alone does not completely explain the power variation within a server rack. Thus, the PowerNet infrastructure will aim to monitor the entire context of a device - including load, temperature, configuration, and others - to gain a complete picture of power usage.

4. FUTURE WORK

The preliminary deployment of PowerNet has already allowed us to collect correlated energy and usage data and draw useful insights into energy efficient computing. Looking forward, we expect the following topics to be the focus of our work.

Large-scale deployment: We are currently extending the PowerNet infrastructure to cover the entire Gates building. The expansion will include a greater variety of monitored devices - end-user devices (notebooks, printers), infrastructure devices (wireless APs, per-room switches), and clustered servers in the building basement. We estimate that a full deployment will require a couple thousand wireless sensors. In addition to hardware power meters, we will expand the set of utilization statistics to include lower-level OS information such as number and type of processes. When necessary, we will also monitor application-level performance (latency or bandwidth) and quality of service.

Due to the scale and experimental nature of the deployment, we believe that custom-built, extensible meters together with open-source software is the right way to proceed. They allow for a low-cost, flexible platform that can be customized in hardware and software to meet evolving research needs. To further explore such capabilities, we plan to teach a class at Stanford in Winter '09-'10, where students will add new capabilities to the PowerNet board, and design applications to visualize and interact with the data.

Specific case studies: We want to explore other case studies using PowerNet. They focus on understanding how different computing setups and infrastructure decisions affect energy and performance efficiency. For example, we intend to compare alternative methods for providing network connectivity within a portion of the building as well as methods for providing compute resources, networked storage, backups, and other administration services.

Analysis and mining: Focused case studies will guide data visualization and analysis. Nonetheless, it is difficult to predict all interesting patterns that will come from the large amounts of PowerNet data

because this level of visibility into energy and utilization has not been available before. Thus, we will also use data mining techniques, such as clustering, to identify interesting trends or anomalies in the data.

Modeling: We will use the PowerNet data to develop models for devices, systems, and user behavior. Predictive models of system performance and energy consumption (e.g., a client-server system) are the key to scheduling for energy efficiency or improving future designs. In collaboration with colleagues specializing in machine learning, we are also considering the use of Markov Decision Processes to model user interaction with both the computing and HVAC systems. Learning from real-world traces to infer future behavior will enable the creation of non-intrusive, automated, energy-saving techniques. We can use the resulting models to accurately determine when a particular computer should go to sleep, how many servers should be on-line at any point, or when a network proxy should be employed.

5. REFERENCES

- [1] ADE7753: Single Phase Multifunction Energy Metering IC with di/dt Input. www.analogdevices.com.
- [2] Arch Rock. www.archrock.com.
- [3] Google Power Meter. www.google.org/powermeter.
- [4] Lucid Design Group Building Dashboard. www.luciddesigngroup.com.
- [5] TEP 123: Collection Tree Protocol. <http://www.tinyos.net/tinyos-2.x/doc/>.
- [6] Watt's up internet enabled power meters. <https://www.wattsupmeters.com/secure/products.php>, 2009.
- [7] Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta. Somniloquy: Augmenting network interfaces to reduce pc energy usage. In *6th USENIX Symposium on Networked Systems Design and Implementation (NSDI'09)*. USENIX, April 2009.
- [8] J. W. Hui and D. 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.
- [9] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *Proceedings of The 8th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN '09)*, San Francisco, CA, USA.
- [10] P. Levis and D. G. et al. T2: A Second Generation OS For Embedded Sensor Networks. Technical Report TKN-05-007, Telecommunication Networks Group, TU Berlin, 2005.
- [11] S. Nedeveschi, S. Ratnasamy, J. Chandrashekar, B. Nordman, and N. Taft. Skilled in the art of being idle: Reducing energy waste in networked systems. In *Networked Systems Design and Implementation*, 2009.
- [12] S. Rivoire, M. Shah, P. Ranganathan, C. Kozyrakis, and J. Meza. Models and Metrics to Enable Energy-Efficiency Optimizations. *IEEE Computer*, pages 39–49, Dec. 2007.
- [13] A. Vahdat, A. Lebeck, and C. S. Ellis. Every joule is precious: the case for revisiting operating system design for energy efficiency. In *EW 9: Proceedings of the 9th workshop on ACM SIGOPS European workshop*, pages 31–36, New York, NY, USA, 2000. ACM.
- [14] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *Proceedings of the 7th symposium on Operating systems design and implementation (OSDI)*, 2006.