Understanding Desktop Energy Footprint in an Academic Computer Lab

Dong Han
Department of Computer Science
University of Houston
Houston, TX, USA
donny@cs.uh.edu

Omprakash Gnawali
Department of Computer Science
University of Houston
Houston, TX, USA
gnawali@cs.uh.edu

***Abstract*—In this paper, we present results from our measurement and analysis of energy use and user behavior in an academic computer lab. We use wireless power sensors to collect power readings from 22 computers in the lab. We use software monitor to log various user activities in the computer. We collected a total of 59.6 million power readings and 220.3 million user activity logs over one month. We analyze the data collected from this instrumentation to not only understand how much energy is used but to also drill down and reveal a detailed understanding of which machines, processes, and users contribute the most to computing energy footprint of the lab. Our results show the power draw on different machines in the lab are different despite the identical hardware and software settings. Our study attributes this difference to different users presenting different types and lengths of load to the computers, preferring specific physical computers, which leads to some computer being used more than others. The results show that the majority of energy was wasted while the computer was left in idle mode, and individual user behavior affects the energy consumption.**

***Keywords*-Power Measurements; Computer Monitoring; Wireless Sensor Network; Energy Efficient Computing.**

I. Introduction

Energy-efficient computing has emerged as a major area of research and engineering in the recent years. As a result, hardware as well as software has become more energy efficient over the years. This progress is partly made possible by careful study of energy consumption of various components within a computing system. Identification of energy hotspots in hardware and software components helps us focus our effort in the areas that are likely to maximize the impact on computer power draw.

There has been recent interest in understanding power draw of a collection of machines, e.g., in data centers [7], [22], [6], [3] or computers in a building [20], [8], [25]. These studies provide measurement-based models of computer power use in buildings. These models can be used to test new approaches to make computing infrastructures in buildings more energy-efficient. While these datasets are extremely useful to the community, they are limited by the setting in which those measurements are taken. Thus, studying and modeling power profile in more settings can be a valuable asset for the green computing research community.

In this work, we conduct a measurement study of power use in an academic shared computing lab environment. In academic buildings, these shared computer labs contribute for a sizable fraction of total energy use. Two factors differentiate this setting from the settings profiled in prior studies. First, the computer labs in a university setting are largely homogeneous: the labs have one or a small number of desktop models. This homogeneity makes this environment more similar to data centers than a typical population of computers in an academic research building. Second, the computers are shared across a number of users. The computers in a shared lab are not personal computers used by a single user. During the course of a day tens of students might log in and use a given computer. Thus, understanding power use of a single computer requires accounting for different users and their different computing requirements.

In our study, we design an instrumentation for the computer lab. The instrumentation consists of two sets of sensing systems. First, there are wireless power meters that continuously monitor power draw of each desktop and transmit to a server using wireless network. We use power meters [15] and RPL protocol [24] to build our wireless energy sensing system. Second, a small service installed on each desktop PC monitors major user events and logs this information in a database. Using data collected with this infrastructure, we can develop a detailed understanding of power draw and the user activities on the computers that drive power use on the computers.

Our measurement study logged power use and user activity on 22 computers in one of the computer labs for over 30 days. We collected 59.6 million power readings and 220.3 million user activity readings. Analyzing this data, we found a considerable heterogeneity in power use despite identical hardware and software configurations on the computers. This difference in power use across the machines is the result of several factors which we quantify in this paper. Different users present a different type of computing load on the computers. They might use computers for different lengths of time. Finally, although the computers are expected to be identical, errors or misconfigurations cause these computers to become different and hence might result in different energy use. Our results also show that the computers were only used for a small fraction of uptime, which means the majority of energy used in the computer lab was wasted.

In this paper, we make four contributions:

- Design and deploy power and user activity instrumentation in a computer lab.
- Present a large dataset describing computing power data in a shared homogeneous computing environment.
- Understand the relation between user activities and power draw.
- Identify energy waste in computing in an academic lab

setting.

The remainder of this paper is organized as follows: Section II presents an overview of the related works. Section III presents the instruction design. Section IV presents the results on the measurement of energy consumption and user activities. Section V discusses alternative methodologies and additional insights from the study and Section VI concludes this paper.

II. Related Work

In this section we give an overview of research related to wireless power meter, measurement of energy consumption, and machine idle proportion in computer lab.

A. Power sensors

Since power consumption has become a significant concern in the development of all kinds of computing laboratories and data centers, much progress and various measurement methods have been designed to measure the power used by the systems, such as Cornil et al. [2], who uses the Fluke ampere meter to measure the current from the power supply, Serra et al. [23] use an AD7757 IC chip based on the shunt resistor from Analog Devices for measuring electric power. Lifton, et al. [19], introduced the MIT Plug sensor network, which embodied the idea of designing sensor nodes to seamlessly become a part of their environment.

Unlike tradional power meters, where the results can only be displayed on local LEDs or saved as data onto flash drives and read later, wireless power meters can transmit the power readings to a remote database to be processed later. There are several wireless powers meter that have been designed in recent years, such as ACme [12] or commercial meters such as *The UFO Power Center*. Jiang, et al. [13] deployed ACme meters for high-fidelity monitoring of electrical usage in a building. Krioukov, et al. [17] presented a personalized smartphone application designed to control the lighting, heating and cooling in user's vicinity, by using ACme for the sensing and actuation. PowerNet [15] is a platform from Stanford University used for collecting, viewing, and analyzing plug-level power data collected. The goal of PowerNet is aiming at answering questions about total power usage, variation, and efficiency. In our work, we use PowerNet meters to collect power readings.

B. Energy Measurement studies

Many measurement studies have contributed to understanding energy consumption in commercial buildings and households. [14] reported a hybrid sensor network based on PowerNet for monitoring the power and utilization of computing systems. Their 3-month monitoring and measurements revealed the IT-related power waste and savings opportunities. Dawson-Haggerty et al. [4] developed a stratified sampling methodology for surveying energy use and conducted a year-long, 455-meter deployment of wireless plug-load electric meters in a large commercial building. They found that the interior of a commercial building is a dynamic environment and confirmed the value of point-to-point routing in a real sensor network deployment. ViridiScope [16] is an indirect power monitoring system.

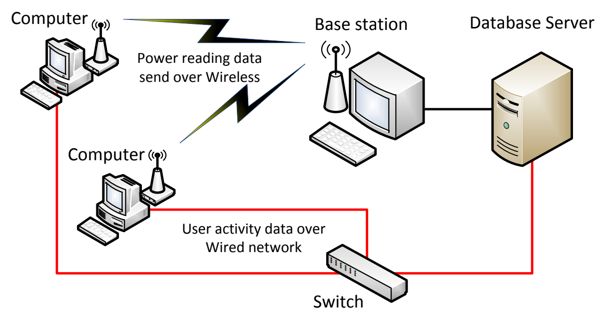

Figure 1. Measurement system overview

It estimates an appliance's power consumption by placing a magnetic sensor near a power cord, based on the fact that the appliance emits measureable magnetic signals when it consumes energy. The eMeter system [21] provides device level energy consumption in a household that is based on a single sensor. This system can provide real time energy usage to a user's smart phone, which makes it possible to get consumers to save energy. Hnat, et al. [10] studied a large-scale residential sensing system for monitoring people's energy consumption in their homes. In this project, the team experienced significant connectivity and access challenges in the home-environment. In this measurement study, we collect and analyze power readings from an academic computer lab environment and put the results in the context of user behavior.

C. Computer idle proportion

Recent studies have shown that the vast majority of workstations and desktop computers remain idle most of the time, which the average CPU idleness at 97.9% in classroom, while the average unused memory is 42.1% [5]. Heap et al. [9] performed 15-minute periodic resource monitor studies on Windows and Unix servers. The study found that Windows servers are idle for approximately 95% of the day respectively, while Unix servers had an average of 85% CPU idleness. Another study showed the average idle time for desktop machines was up to 80% of the day [1].

While some papers focus on measuring the power consumption for different machines, and others tried to summarize the relationship between power consumption and performance, however, to the best of the authors' knowledge, there have been few reports in the literature to date about the relationship between user behavior and energy consumptions, on the same type of computers.

III. Power and User Activity Instrumentation

We instrumented a computer lab to collect two sets of information: power draw of computers and user activity on the computers. The overview of sensing system is shown in Figure 1. The power readings and user activity logs were sent to a database server by wireless and wired network.

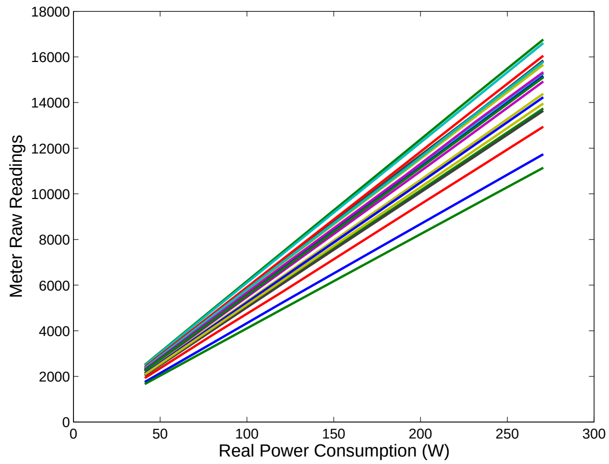

Figure 2. Meter raw reading before calibration

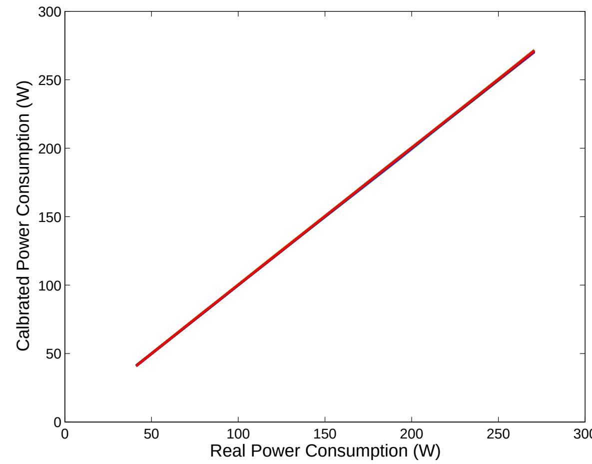

Figure 3. Meter reading in Watt after calibration

Undergraduate and graduate students visit this lab for academic purpose, such as to finish programming homework or to access remote server by using ssh client. There are no scheduled classes in the lab. There are 22 desktops made by Gateway located in 4 rows, equipped with Intel Core2 duo 1.88G CPU and 2048 MB memory, running Windows XP. These desktops have CPU and monitor in a single package. There is no policy to restrict students from using any specific computer, but students are required to login with their personal student ID before using it. And this log in, log out activity was automatically recorded in database.

A. *Wireless Energy Meter*

We use PowerNet [15] nodes to measure power draw of a computer. Computer power cable is plugged to a PowerNet node and the PowerNet node is plugged to an AC outlet. These meters have energy metering ICs and MSP430 micro-controller for sensor control and data processing. The meter can sample power draw at up to 14 KHz. The power meters also have a IEEE 802.15.4 radio chip CC2420 running at 2.4 GHz unlicensed spectrum. We use RPL [24] running on TinyOS [18] for collecting power measurements.

We programmed PowerNet nodes to sample current at 10 Hz. The nodes pack 20 readings into a single packet and send it to the base station. Each reading is 2 bytes. 20 readings and metadata results in a 58 byte application payload. Metadata include a local sequence number, time stamp, and node ID. We increment the local sequence number after sending each packet. We used a local timer value as time stamp. Although this time is not globally synchronized, it is sufficient to study the time gap between the packets.

B. *Process Monitor*

We wrote a C++ application to measure CPU usage and installed as a Windows Service on all computers in the lab. Every second, the process monitor calculates each processs CPU usage by using Windows Management Instrumentation (WMI) API [11]. The process monitor then transmits the list of processes and their CPU utilization to a database server over wired Ethernet. We install the process monitor as a windows service so that it starts automatically during bootup and continuously collects information regarding processes running in the computer even when no one is logged in to the computer. The process monitor itself uses in average 0.45% (with a peak of 0.8%) of CPU resource, which we can safely ignore from our calculation.

C. *User Authentication Monitor*

We use a C# application to monitor user authentication. It records user log-in and log-out activities, user ID, machine name, and then saves this information on a database server.

D. *Meter Calibration*

We calibrated all the power meters before deployment. We performed six point calibration with resistive loads from 40 to 260 watts. This range is within the power draw range of a desktop, typically between 80 and 180 watts. The ground truth in Watts of resistive loads was calculated by using its instantaneous current multiplied by the potential difference across this component, which was measured by a high-precision multi-meter.

We first connect all of the resistive loads through one power meter. We then turn on the first load, wait for a few seconds until the temperature is stabilized, then take 50 raw readings using the meter and calculate the average. Then, we repeat the same process by turning the load one by one. After we have 6 average raw readings at 6 different loads, we use Polynomial curve fitting function in Matlab to calculate the coefficients of raw readings of degree 1 that fits the ground truth. Figure 2 shows raw readings from 24 meters before calibration and Figure 3 shows the result after calibration. After calibration, all of the 24 lines overlap as expected because they are all measuring the current through the same resistive loads. After calibration, the Mean Square Error (MSE) in power readings across all of the meters was less than 0.2.

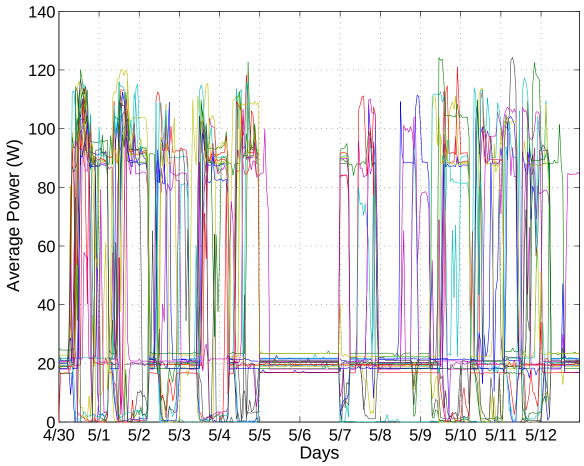

Figure 4. Average Power draw of each machine across time.

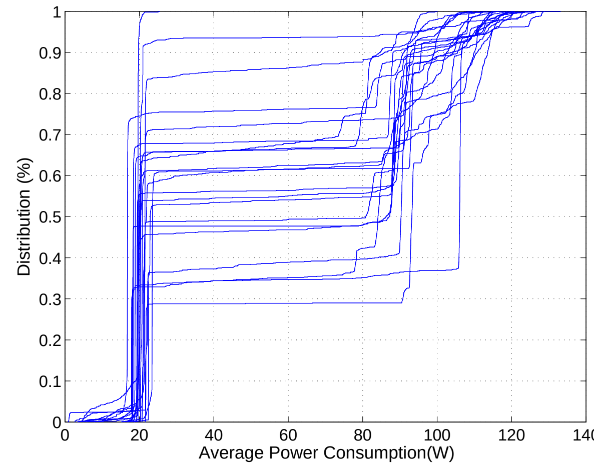

Figure 5. CDF of Power draw of each machine.

E. Data collection rate

Initially, the process monitor reported all the processes and their utilization to the database server. After analyzing utilization data for two weeks, we found that 85% of the readings had a utilization of 0%. We optimized network transmission and storage for these processes using a simple compression, which is tag a list of processes with a single utilization number of 0% rather than (process, utilization) tuple for each process. This optimization reduced the number of rows in the database by around 90%.

IV. Results

Our analysis of data collected in this study reveals up to 13.42 times difference in energy footprints of different computers in the lab. Furthermore, we found up to 575.23 times difference between the energy used by different students. In this section, we elaborate on these findings.

A. Power across time

We first study the temporal trends in power draw of the computers in the lab. We found that the power draw can change by as much as 197.31% over the course of a day. When idle, most computers required around 70 watts. The maximum power requirement was 138.12 watts. The power changes depending on the load on the computer was due to user activities. Figure 4 shows the power draw of each machine for 12 days. Different machines show different temporal patterns. In this figure, day 6 and day 7 are weekends. No student was allowed to use the computers in the lab during that period, which is why the power was stable. Figure 5 plots the distribution of power draw across time for each computer. It shows that all computers spend at least 29% of the time drawing power less than 70 watts.

Aggregating the power draw from all the computers, we find that the total power drawn by the computers changes across time as shown in figure 6. It shows a minimum of 400 watts and a maximum of 1580 watts as total power used by computers in the lab.

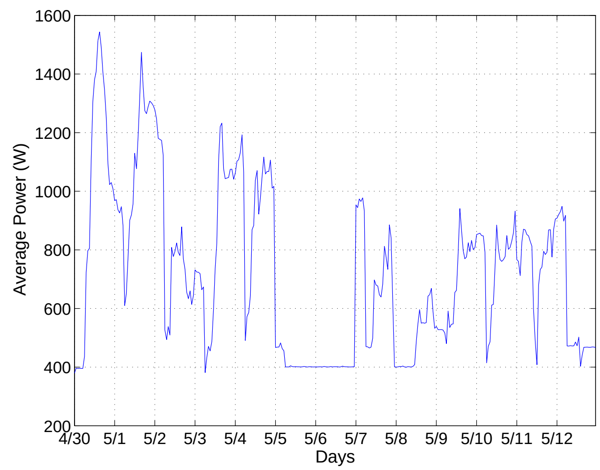

Figure 6. Total Power used by all the computers in the lab.

B. Power across activities

One of the reasons for different power draw across time is the changing user activities. The activity most relevant to understanding power draw is a student logging into the computer, launching applications, and after some period, logging off. In figure 7, we plot the CDF of average power for sessions during which users were logged in and were not logged in. The power is generally higher when the user is logged in compared to when the user is not logged in, approximately 80-140 watts compared to 20-80 watts on most of the machines. The reason it consumed more energy when user was logged in, is not only that more CPU operations were performed, but also hard disk read-write operations, graphics card calculations, and network transmissions which expend more energy. Some of the blue lines stay around 20 watts for large fraction of time. That means those machines were left in power off status. For a

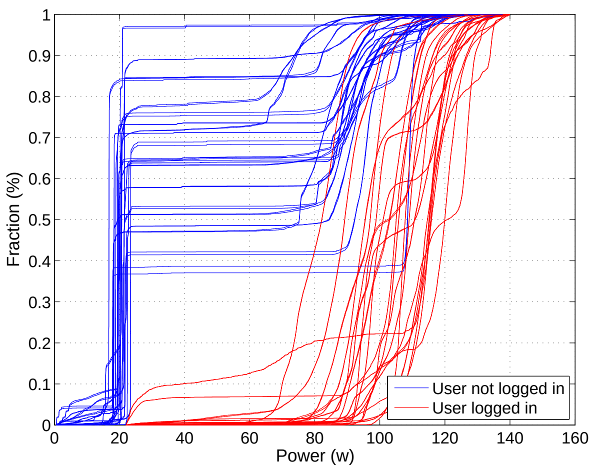

Figure 7. CDF of Power draw of each machine when a user is logged in and not logged in to the machine.

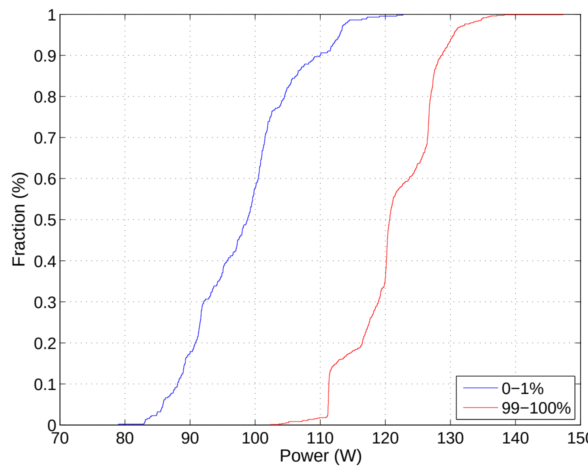

Figure 8. CDF of Power draw when CPU utilization is 0-1% and 99-100%

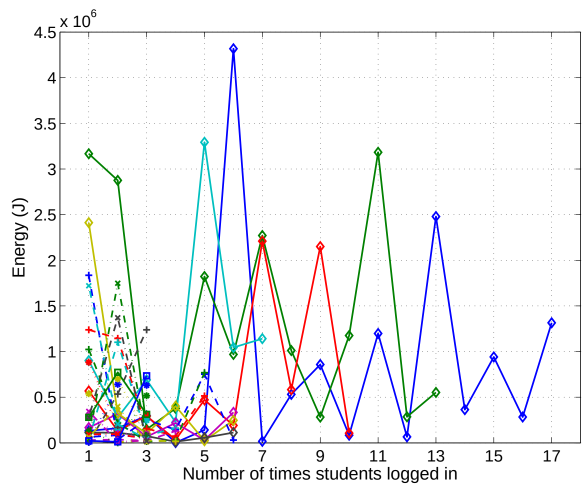

Figure 9. Energy consumed by each student.

small fraction of the dataset, the computers use more power when a user is not logged in compared to when a user is logged in. This is because the Windows was performing system updates or scheduled virus scan. Our test indicates this model of machines consume approximately 15 - 22 watts power even when power is turned off; that was mainly because the internal capacities are still charging.

C. Power across machines

Although all the computers in the lab have identical manufacturer specification and software installation, each computer is slightly different due to manufacturing difference, hardware abuse and errors over time, and unintentional errors and updates on the software. For example, some machines do not have McAfee VirusScan installed, while others do. It is believed that such discrepancy is due to software configuration errors accumulating over time.

To understand the difference between the machines, we study two sets of trace. We first plot the distribution of power draw when no user is logged in and CPU utilization is 0-1% in figure 8. The figure shows even when all the computers are idle, the power draw across the machines is in 85-115 watts range. In the same figure, we also plot the power draw when a user is logged in and the CPU utilization is 99-100%. When the machines are fully utilized, the power draw across the machines are in the 108-136 watts range. Thus, our results show that even though computers of similar manufacturer specification and software installation are subjected to similar CPU loads, the power draw can be significantly different.

D. Power across students

In figure 9,we plot the total computing energy used by 30 students who logged in the most number of times during the two weeks. Each line represents the total energy used by a student. The dots shows the value of energy consumed during that log-in period. All of these numbers were captured during the same period as the figures showed in pervious sections. We can see that around 80% of students used less than 1,000 kilojoules, while only 0.5% of students used more than 3,000 kilojoules, which is mostly caused by students having been logged in for a significantly longer period, compared to other students.

From figure 10 we can tell each student has her unique average power usage, but 70% of these values are in the range of 100-120 watts. In some cases the average power is below 100 watts mainly because after the user logged in, the user left the computer idle or performed some simple tasks like writing emails. The user activity log indicates that no high workload task was performed on that computer during that period.

E. Power across processes

Different processes perform different tasks in a computer and have different energy profile. We now explore power

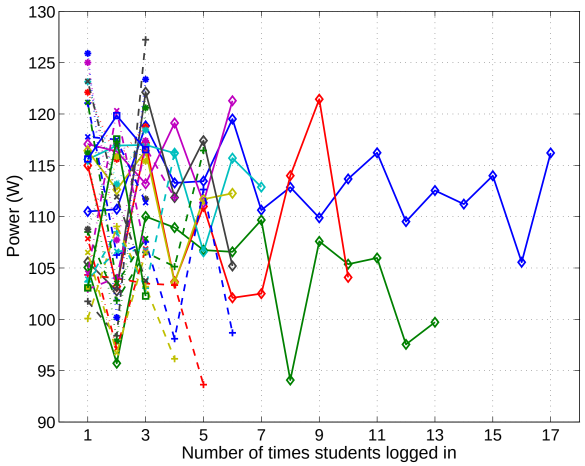

Figure 10. Average power used by each student.

Type	Process Name	Power (W)	Energy(Kj)
System	WPFFontCache	70.44	812.65
	wmiprvse	56.38	18262.63
	searchindexer	47.76	209.02
	winlogon	47.50	30.02
	explorer	38.51	69.48
	svchost	38.48	11920.26
	Total	299.07	31304.06
User	McScript_InUse	59.14	69.78
	netbeans	41.92	33.53
	mcshield	36.91	226.21
	firefox	34.62	1104.40
	POWERPNT	34.18	34.25
	iexplore	32.59	220.93
	Total	239.36	1689.10

Table I
TOP 12 PROCESSES THAT CONSUMED MOST POWER.

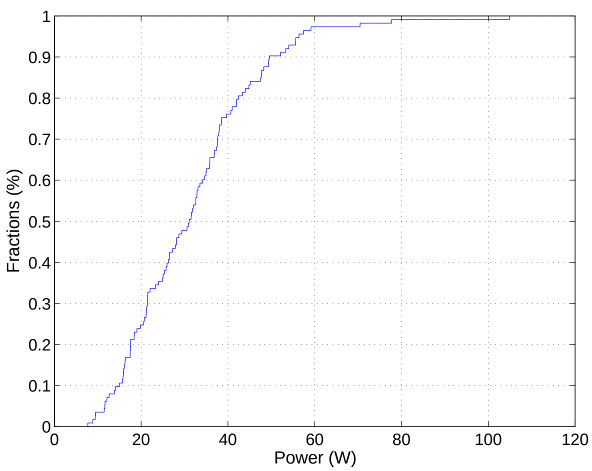

Figure 11. CDF of average Power draw across processes.

consumption by each process. We combine all of the data from 22 machines, but only consider the data which has non-idle processes with more than 5% CPU utilization, and divide energy proportional to CPU usage, then calculate the average power consumption and total energy usage for each process.

In Figure 11, we plot the distribution of average power across the processes. It shows that less than 90% of processes have their average power consumption, evenly distributed between 8 watts and 50 watts.

It is not surprising that there is a large difference between the energy consumption by different processes. Table I shows the details for top 12 processes ordered by average power consumption, divided into two parts: System or User processes. Process *wmiprvse* consumed 18,262 kilojoules, the highest energy consumed by a process. Since *wmiprvse* is a Windows Management Instrumentation component that provides operating system management information and control in an enterprise environment, this process was executed in the background regularly. The Process Monitor periodically calls the API provided by *wmiprvse* service. This may be the reason why the CPU utilization is high for *wmiprvse*. Process *winlogon* is only executed for user authorization and windows activation checks when user tries to login. Table I indicates although the average power consumption of *wmiprvse* is only 18.69% more than *winlogon*, the total energy used by the former is 608.35 times more than the latter since the total running time of *wmiprvse* is much longer than *winlogon*. While System part consumed 31,304 kilojoules, the User part only used 1,689 kilojoules. This not only shows that efficiency of these computer is low, but also that these computers were left in idle mode for much longer than they were used by users. From analysis of the energy usage by different application, we found that students prefer Firefox to IE for web browsing in this lab.

F. Computing energy footprint

Until now, we studied the impact of individual factor on power draw of computers. We now study the combination of all these factors to understand the total computing energy footprint of the computer lab.

Figure 12 shows total energy consumption on all computers during the period of 2 weeks. We divide the total energy into three parts. The red bar shows the energy used when a user is logged in, the blue shows the energy when a user is not logged in, and the green shows the energy when the computer was powered off. The x-axis was sorted in increasing order of each machine's energy use. Even though these machines have the same specification, the total energy consumed is distributed from approximately 900 to 6,000 Kilojoules. There was, as expected, a significant amount of energy waste, while the computer is powered off. The energy used while there is no user is logged in is 200% to 300% as compared to the period which the user is logged in. There are some computers that consumed much more energy compared to others, which indicates that those computers were selected

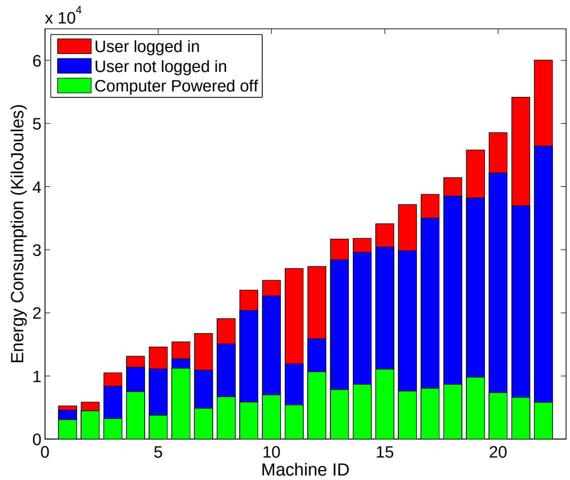

Figure 12. Total Energy Consumed by each machine.

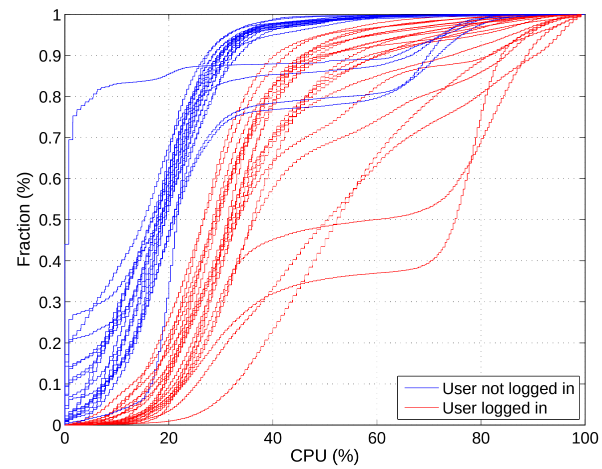

Figure 13. CDF of CPU Utilization on all machines

by students more often. One of the potential reasons is the convenient location for student, such as closer to the lab entrance, or beside the aisle. Another reason is the lab admin manually turning off a whole row of computer using per-row power switch. This happens only when the lab admin is in the lab and notices all the computers in a row not being used for a long time. Students may later use the manual power switch to power the computer.

V. Discussion

Instrumenting all the computers in a lab is a capital and time intensive process. It is natural to ask if we can use CPU utilization as a proxy for power use. To answer this question, we try to interpret the power draw shown in Figure 13 in context of corresponding CPU utilization distribution. Each line in the figure represents individual machine. It is not surprising to see that the fraction of CPU utilization while user is not logged in stays below 30%-40% on most machines, which was caused by some background processes like real-time back up and viruses scans.

On the other hand, when a user is logged in, the CPU utilization is distributed in the range of 40% - 90%. While users are using the machine, they will consume more CPU resource, since every user task requires additional CPU operations.

Figure 14 shows the distribution of total energy consumption during the whole measurement period across students. It shows around 75% students consumed equal to or less than 1,000 kilojoules of energy, approximately 20% of users consumed the energy between 1,000 and 3,000 kilojoules. From these numbers we can see that, around three-quarter of students used this computer lab only for a short of period. Perhaps they use this lab for printing or submitting homework. While 5% of users consumed more than 3,000 kilojoules of energy. They occupied the computers for long period and probably use the computer lab to finish their programming homework.

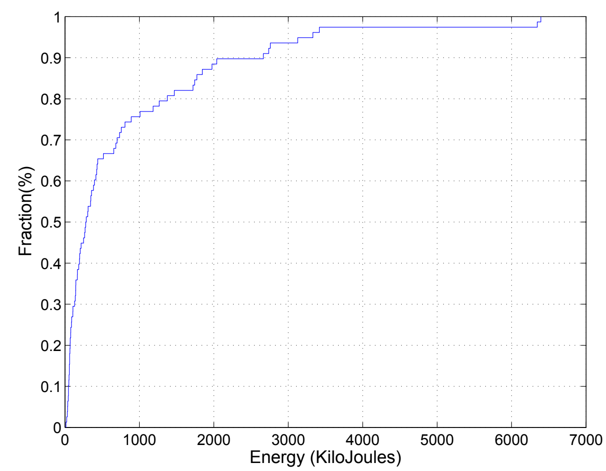

Figure 14. CDF of Total Energy Consumption across students

VI. Conclusions

In this paper, we described the sensing system that measures the energy consumption and user activity in an academic computer lab. We measured the power consumption, CPU utilization, and user activity across a homogeneous set of desktop computers. We studied the factors that drive the heterogeneity in energy use across the desktops and lessons learnt from the real-world measurement on 22 machines over one month. We found that the energy consumption of each computer is highly related to individual user behavior, and the 60% of energy consumed every day was during the computer was on and no one was logged in.

We are currently working with the department to continue the measurement in other labs. As the department migrates to thin clients, we also plan to extend our study to thin clients.

Acknowledgments

Thanks to Maria Kazandjieva and Philip Levis for their help with power instrumentation infrastructure. Thanks to Tom Cumpian and Babu Sundaram for providing access to the computer lab for measurements.

References

[1] Anurag Acharya, Guy Edjlali, and Joel Saltz. The utility of exploiting idle workstations for parallel computation. *SIGMETRICS Perform. Eval. Rev.*, 25(1):225–234, June 1997.

[2] M. Cornil, L. Colman, and G. De Mey. Power consumption measurements on portable pc's. In *Mixed Design of Integrated Circuits and Systems (MIXDES), 2011 Proceedings of the 18th International Conference*, pages 549 –551, june 2011.

[3] P. Cremonesi, A. Sansottera, and S. Gualandi. Optimizing cooling and server power consumption. In *Intelligent Computer Communication and Processing (ICCP), 2011 IEEE International Conference on*, pages 455 –462, aug. 2011.

[4] S. Dawson-Haggerty, S. Lanzisera, J. Taneja, R. Brown, and D. Culler. @ scale: insights from a large, long-lived appliance energy wsn. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 37–48. ACM, 2012.

[5] P. Domingues, P. Marques, and L. Silva. Resource usage of windows computer laboratories. In *Parallel Processing, 2005. ICPP 2005 Workshops. International Conference Workshops on*, pages 469 – 476, june 2005.

[6] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. Power provisioning for a warehouse-sized computer. In *Proceedings of the 34th annual international symposium on Computer architecture*, ISCA '07, pages 13–23, New York, NY, USA, 2007. ACM.

[7] Ravi Giri and Anand Vanchi. Increasing data center efficiency with server power measurements. http://resources.spiceworks.com/banners/ibm/europe/whitepapers/ServerPowerMeasurement.pdf.

[8] G.W. Hart. Residential energy monitoring and computerized surveillance via utility power flows. *Technology and Society Magazine, IEEE*, 8(2):12 –16, june 1989.

[9] DG Heap. Taurus-a taxonomy of actual utilization of real unix and windows servers. *IBM White Paper GM12-0191*, 2003.

[10] T.W. Hnat, V. Srinivasan, J. Lu, T.I. Sookoor, R. Dawson, J. Stankovic, and K. Whitehouse. The hitchhiker's guide to successful residential sensing deployments. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, pages 232–245. ACM, 2011.

[11] Windows Management Instrumentation. http://msdn.microsoft.com/en-us/library/windows/desktop/aa394582(v=vs.85).aspx.

[12] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *Information Processing in Sensor Networks, 2009. IPSN 2009. International Conference on*, pages 253–264. IEEE, 2009.

[13] X. Jiang, M. Van Ly, J. Taneja, P. Dutta, and D. Culler. Experiences with a high-fidelity wireless building energy auditing network. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, pages 113–126. ACM, 2009.

[14] M. Kazandjieva, O. Gnawali, B. Heller, P. Levis, and C. Kozyrakis. Identifying energy waste through dense power sensing and utilization monitoring. Technical report, Stanford University, Tech. Rep. CSTR 2010-03, 2010.

[15] M. Kazandjieva, B. Heller, D. Gal, P. Levis, C. Kozyrakis, and N. McKeown. Powernet: A magnifying glass for computing system energy. In *Proc. Stanford Energy & Feedback Workshop: End-Use Energy Reductions through Monitoring, Feedback, and Behavior Modification*, 2008.

[16] Y. Kim, T. Schmid, Z.M. Charbiwala, and M.B. Srivastava. Viridiscope: design and implementation of a fine grained power monitoring system for homes. In *Proceedings of the 11th international conference on Ubiquitous computing*, pages 245–254. ACM, 2009.

[17] Andrew Krioukov and David Culler. Personal building controls. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, IPSN '12, pages 157–158, New York, NY, USA, 2012. ACM.

[18] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, et al. Tinyos: An operating system for sensor networks. *Ambient intelligence*, 35, 2005.

[19] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J.A. Paradiso. A platform for ubiquitous sensor deployment in occupational and domestic environments. In *Information Processing in Sensor Networks, 2007. IPSN 2007. 6th International Symposium on*, pages 119–127. IEEE, 2007.

[20] Alan Marchiori and Qi Han. Using circuit-level power measurements in household energy management systems. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '09, pages 7–12, New York, NY, USA, 2009. ACM.

[21] Friedemann Mattern, Thorsten Staake, and Markus Weiss. Ict for green: how computers can help us to conserve energy. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, e-Energy '10, pages 1–10, New York, NY, USA, 2010. ACM.

[22] Meikel Poess and Raghunath Othayoth Nambiar. Energy cost, the key challenge of today's data centers: a power consumption analysis of tpc-c results. *Proc. VLDB Endow.*, 1(2):1229–1240, August 2008.

[23] H. Serra, J. Correia, A.J. Gano, A.M. de Campos, and I. Teixeira. Domestic power consumption measurement and automatic home appliance detection. In *Intelligent Signal Processing, 2005 IEEE International Workshop on*, pages 128 – 132, sept. 2005.

[24] Ed. T. Winter and Ed P. Thubert. Rpl: Ipv6 routing protocol for low power and lossy networks. http://tools.ietf.org/html/draft-ietf-roll-rpl-19, 2011.

[25] Chuyuan Wei and Yongzhen Li. Design of energy consumption monitoring and energy-saving management system of intelligent building based on the internet of things. In *Electronics, Communications and Control (ICECC), 2011 International Conference on*, pages 3650 –3652, sept. 2011.