

Cyber-Physical Energy Systems: Focus on Smart Buildings

Jan Kleissl

Dept. of Mechanical and Aerospace Engineering
University of California, San Diego
9500 Gilman Dr.
La Jolla, CA, USA
kleissl@ucsd.edu

Yuvraj Agarwal

Dept. of Computer Science and Engineering
University of California, San Diego
9500 Gilman Dr.
La Jolla, CA, USA
yuvraj@cs.ucsd.edu

ABSTRACT

Operating at the intersection of multiple sensing and control systems designed for occupant comfort, performability and operational efficiency, modern buildings represent a prototypical cyber-physical system with deeply coupled embedded sensing and networked information processing that has increasingly become part of our daily lives. In this paper, we look at modern buildings entirely as a cyber-physical *energy* system and examine the opportunities presented by the joint optimization of energy use by its occupants and information processing equipment. This paper makes two contributions: one, a careful examination of different types of buildings and their energy use; two, opportunities available to improve energy efficient operation through various strategies from lighting to computing. Using a modern 150,000 sq feet office building as a closed system, we detail different strategies to reduce energy use from LEED certification to zero net energy use.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

General Terms

Measurements, Performance

Keywords

Energy Management, Buildings, Smart Grid, Cyber-physical, Embedded, Energy Metering, LEED, ZNEB

1. INTRODUCTION

Modern buildings are system of systems consisting of interacting heat exchange, airflow, water, safety, access/security, and movement control subsystems. These subsystems are increasingly coupled using embedded sensing and control systems where state information from one system is directly

used to make operational decisions in another subsystem. Viewed as an *energy* system, that is, one with energy input/outputs and internal energy flows to achieve building performability requirements, a modern building presents an example of a deeply coupled system of energy usage, comfort and work derived. At a macroscale, buildings use approximately 70% of total electricity usage and emit approximately 40% of greenhouse gases (GHG) annually in the United States[7]. According to the Smart 2020 report, by far the greatest potential of reducing GHGs by application of information and communication technologies reside in building smart grids and smart buildings, savings from each of which will easily exceed the GHGs attributed to entire Information and Communication Infrastructure (ICI) use worldwide (at 1.48 GtCO₂)¹.

It is this opportunity – of applying deeply coupled information processing to the design and operation of buildings – that is the chief driver for our work in exploring buildings as a cyber-physical system (CPS). In particular, we have focused our attention on ICI in buildings, both as energy users as well as energy/operations optimizers, for its potential in achieving energy efficient buildings. Once again, according to Smart 2020, by 2020 the energy use of personal computers and laptops in offices and homes will exceed over 3x the energy use by all data centers combined. Thus the potential of efficiency improvements in operations of buildings when viewed as an integrated space for humans and computing machines is significant. It naturally raises the question: what are the interdependencies and opportunities for optimization across the two modalities of energy use? This paper seeks to answer this question by posing total energy use of a building against its performance criteria as a quantitative challenge towards zero-net energy buildings (ZNEB).

A ZNEB is a building with zero net annual energy consumption. The electricity grid merely acts as a buffer that allows excess energy production to be “stored” and reclaimed during times of insufficient local energy generation. A ZNEB is ultimate goal that requires achieving high energy efficiency together with renewable energy production on-site. ZNEB cannot be achieved solely with physical infrastructure in buildings such as insulation of the building envelope, windows with low solar heat gain coefficients, and energy efficient lighting and devices. ZNEB is as much a challenge in effective renewable local cogeneration as it is a grand challenge for sensing, information processing and control for the

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.

DAC'10, June 13-18, 2010, Anaheim, California, USA
Copyright 2010 ACM 978-1-4503-0002-5/10/06 ...\$10.00.

¹www.smart2020.org




Figure 1: California building energy use survey (CEUS) results for electricity use in Large Office Buildings [11].

coupled CPS that ensures minimum energy loss, storage and timely availability of energy to various subsystems. Managing the energy consumption by Information Technology (IT) equipment, ventilation, and lighting in response to internal and external environmental conditions is essential in reducing the energy use of modern buildings that are increasingly becoming mixed-use [3, 5]. In this paper we will provide a review of various smart building technologies and use an example of a densely metered mixed-use office building to lay out strategies to make such a building a ZNEB.

2. QUANTIFYING THE BASELINE ENERGY USE IN BUILDINGS

A significant portion of modern buildings today can be classified as “mixed-use” since they have a mix of both human occupants and supporting ICI. The thermodynamics of such building systems therefore entail not only modeling of heat and airflow conditions, but also heat generation in the computational processes that convert electric energy to information and heat. This is especially true of large office buildings, which are the most prevalent new commercial construction and have the largest potential for implementation as cyberphysical energy systems. Building energy use is regulated by the American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) standard 90.1[4], which applies to building components such as building envelope, electrical systems, lighting and HVAC schedules, and mechanical and plumbing systems. The Leadership in Energy and Environmental Design (LEED) rating system provides up to 10 points in category EAc1, energy optimization, for reduction in *simulated* energy usage compared to a similar building built to ASHRAE 90.1 standards.

Figure 1 shows the results of the California Commercial End Use Survey (CEUS) for large office buildings in California in 2005. Interior lighting, office equipment, cooling, and ventilation make up about 83% of the total electricity use. These components are also the largest at any time during the day (Figure 2), but cooling accounts for the largest daytime contribution - especially during peak time - and office equipment accounts for the largest nighttime (or baseline) electrical load contribution. **Baseline electrical load** is load that always occurs, even during times of low occupancy such

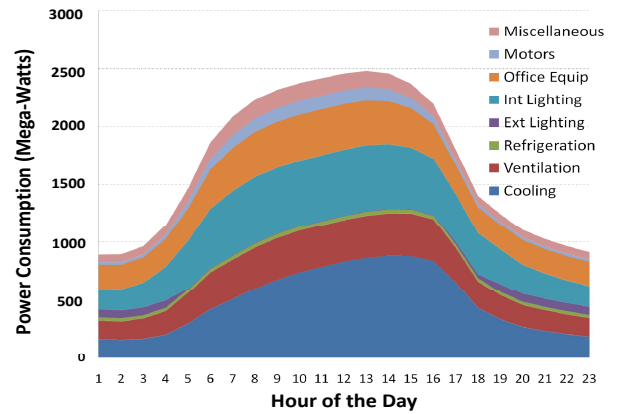


Figure 2: Typically summer day electricity usage of a large office building [11].

as nights and weekends. Baseline load is an important metric in energy systems design because it is met by the most economical, reliable, but continuous power source such as coal or nuclear power plants. The baseline to peakload variations are generally subject of innovations in energy storage, renewable energy and their dynamic control by building automation systems.

As part of a large scale energy monitoring project [3], we have built an Energy Dashboard² that visualizes and compares energy use across sixty buildings on the University of California, San Diego, campus. In Table 1 we compare power the average, the peak and the baseline electrical consumption for six buildings ranging from a very IT-dominated building (San Diego Supercomputer Center, SDSC) to mixed use, office, and residential buildings. For all these buildings in Table 1, baseline load is a significant fraction of peak load and almost equal to their average power consumption. This is expected for SDSC as computing jobs run 24/7, but most of the other buildings are lightly or unoccupied and unused at night. With the exception of a few “mission-critical” components, such as refrigerators and automated research equipment, there should be no reason why energy use should be so large at night. This presents an opportunity for energy conservation without impacting comfort or productivity of the building or its occupants.

3. PROBLEM SCOPE AND COMPONENTS

Each of the main electricity consumers in buildings - interior lighting, office equipment, cooling, and, ventilation - has potential for increases in energy efficiency. Some of these, as our results show later, can be dramatic. However, these measures must be balanced against building performance to ensure design of the so-called performative architecture. For instance, as a caution note that a 2% decrease in the productivity of office building occupants has the same economic impact as all building maintenance and energy expenditures. Thus, while dimming all lights would result in large energy savings, the loss in productivity and comfort in buildings would not justify this measure during work hours. Of course, interior lighting energy usage can be reduced by efficient lighting technology (fluorescent lights, LEDs) and daylighting. Care has to be taken to minimize glare. Fea-

²The Energy Dashboard: <http://energy.ucsd.edu>

	Building	Type	Average Power Use	Peak Power Use	Baseline Power Use
(1)	SDSC	Super Computer Center	1873 KW	1930 KW	1841 KW
(2)	CAL-IT2	Mixed-Use + Nano-Fab	918 KW	978 KW	886 KW
(3)	Natural Science	Mixed-Use	602 KW	746 KW	537 KW
(4)	CSE Building	Mixed-Use	399 KW	521 KW	255 KW
(5)	Bio-Medical	Mixed-Use + Medical Equipment	273 KW	348 KW	230 KW
(6)	Pepper Canyon	Residential Hall	56KW	83 KW	39 KW

Table 1: Average, peak and the baseline power usage of various buildings across the UC San Diego campus for 2009. Baseline values are the average of the lowest 10% and peak values are the average of the highest 10% of the values over the year.

tures such as automated dynamic exterior shading devices and dynamic dimming of interior fixtures can result in dramatic reduction in energy usage.

Another source of optimization is the energy use by office equipment which is increasingly dominated by information and communication infrastructure (ICI), such as servers, networking, and personal computers. Tremendous gains in energy efficiency have been achieved in the design of ICI resulting in a modest increase in its energy use compared to the increase in performance. Most laptops today use less power than an incandescent light bulb. While ICI has to be responsive and resilient during daytime, energy efficient and smart standby mechanisms can be employed to reduce electricity usage at night and on weekends.

Cooling and ventilation of buildings is another technology area where sophisticated cyber-physical modeling and tools have significant potential for impact on building energy efficiency. Cooling needs are determined primarily by heat conduction and solar radiation through the building envelope, cooling of outdoor air to ventilate indoor spaces, and indoor heat production by humans, lighting, and ICI. While sophisticated models exist to simulate these processes in the building design phase, model predictive control (MPC) needs to be implemented to simulate these processes during building operation so as to actuate building cooling. An MPC formulation enables system regulation while accounting for constraints and producing pseudo-optimal policies for buildings. Knowledge of these processes can also inform strategies for automated demand response. Weather data, time-of-use energy prices, and energy load forecasts could be used to better tune control systems increasing energy efficiency and reducing operating costs. Also forecasts of local renewable generation can inform decisions to use local power generation versus purchasing power from the grid.

In the following sections we will quantify energy use due to lighting, ventilation, and office equipment in a modern mixed-use office buildings and explore opportunities for energy conservation and renewable energy production.

4. A MIXED-USE BUILDING

Our experimental testbed literally consists of a large university campus at the scale of a small city. The University of California at San Diego (UCSD) campus is situated in the coastal community of La Jolla in ASHRAE climate zone 3 (warm) in a marine climate influenced by the sea breeze. The campus features a variety of renewable energy generation and storage elements: self-generating over 80% of its use using two 13.5 MW gas turbines, one 3 MW steam turbine, and 1.2 MW of solar photovoltaic (PV). With a daily population of over 45,000 people out of which over 29,000 are students (comprised of 10,000 on campus residents), the

campus presents a rich testbed to explore energy use and performance metrics of buildings. There are over 450 buildings scattered across the UCSD campus, most of which are metered by commercial grade meters to measure the electricity and cooling load of each building (most buildings are cooled through district cooling using a chilled water loop). Since building-level metering makes it hard to separate the loads into subsystems we submeter a mixed-use building, the Computer Science and Engineering building (CSE).

CSE, completed in 2004, incorporates various energy saving features and design concepts, and provides a crystal ball to the future of emerging mixed-use buildings. "Mixed-use" describes the collocation of human occupancy and ever increasing IT infrastructure in a building. CSE is a large office building, spread across five floors encompassing an area of 149,804 basic gross ft² and 132,382 net usable ft². The CSE building nominally holds approximately 600 occupants in regular offices, which increases to about 1000 people at times due to classes and labs. CSE is ventilated by a closed loop system that provides zonal and floor-by-floor control of air flow, temperature, and lighting conditions. It also provides dynamic control of window shading to control the amount of incident sunlight and associated cooling load into the offices.

The total electrical energy consumption for the CSE building in 2009 was 3414 MWh (11,648 mmBTU). At a utility electricity rate of 13 cents/kWh this translates to almost \$450,000 in electricity costs annually. We have deployed an energy measurement infrastructure within CSE that allows a detailed breakdown of electricity usage into four logical subsystems (Fig. 3): *Lighting loads*, *Mechanical load* (includes all climate control equipment such as air handlers, pumps, etc.), *Machine-room loads* (includes >400 servers, networking equipment, UPS) and finally all the *Plug Loads* (include everything else plugged into the wall sockets; a large fraction of the plug load energy use can be attributed to >1000 laptop and desktop PCs).

Peak power consumption occurs during working hours on weekdays with a reduction in power consumption during nights and weekends. There are several important observations from the data presented in Figure reffig:cse-breakdown-week. First, the mechanical load on the building has the largest dynamic variation and it is the only load that conforms to the occupancy cycle since it only increases from 6AM to 8PM during weekdays. Second, the baseline load of the building, during nights and weekends, remains at around 325 kW. Finally, and most importantly, power consumed by IT equipment (machine room load and a major portion of the plug loads) accounts for 50% to 80% of this baseline load of the CSE building. This implies that most IT equipment is not turned off or put into low power sleep modes dur-




Figure 3: Detailed breakdown of the energy consumption of the CSE mixed-use building for 3 weeks in March 2010. The thermal load on the building is provided by a separate campus-wide chilled water and a hot water loop and is not included.

ing off-hours, consistent with previous studies [1, 15]. We believe that the fact that the other subsystems of the CSE buildings are energy efficient actually makes the energy used by IT equipment stand out even more. CSE has an Energy Usage Intensity (EUI) of 80 kBtu ft^{-2} . For the sake of comparison, in a study of 121 Leadership in Energy and Environmental Design (LEED) certified buildings in the U.S., the median measured EUI of office buildings was 62 kBtu ft^{-2} [16]. On the other hand the average building uses about 95 ft^{-2} [8]. However, for the warm-hot ASHRAE climate zones (1-3) LEED buildings were no better than the average building at around 75 kBtu ft^{-2} . Given the large variation in climate and building use, these numbers merely serve as a guide. Despite its relatively recent construction and energy efficiency features, the CSE building uses more energy than comparable office buildings due to energy intensive research and computing operations occurring in the building. This emphasizes the need for further increase of energy efficiency using coupled measures targeting IT equipment and climate control systems.

We will now conduct a virtual experiment to determine how CSE could become a ZNEB. Since ZNEBs produce energy on-site to avoid import from the grid, the objective is to power CSE from local renewable energy sources such as solar, geothermal or wind. Energy efficiency measures often provide the most cost-effective means for addressing electricity use. In addition, a judicious mix of energy efficiency measures not only reduces electricity demand but also helps reduce the size and, therefore, the required capital outlay for solar photovoltaic (PV) systems. Based on the detailed analysis of energy use in CSE, we examine the following measures:

1. The desktop PCs and the machine room servers will utilize an energy saving architecture called Somniloquy[1].
2. Using more efficient LED bulbs and motion sensors to control the lighting, the energy usage for lighting could be reduced.
3. Reduce mechanical subsystems energy demands by changing the settings of the climate control system to adjust air conditioning to actual demand and to replace current fans and pumps with more efficient units. In the

absence of detailed specifications of the building we assume that a 20% reduction in mechanical subsystem energy use could be obtained.

4. Install a solar PV array on the roof of the CSE building as a means to generate energy locally.

4.1 Somniloquy to Reduce IT Energy Use

Maintaining continuous network presence and availability and utilizing low power sleep modes are at odds with each other. Our solution to this issue are two energy saving architectures, Somniloquy [1] and SleepServers[2], that enable PCs to respond to network traffic and even run some applications in low power sleep states. In a pilot project we have deployed Somniloquy and SleepServers to thirty users across CSE. We have been collecting detailed energy consumption traces for these users, and have observed energy savings ranging from 60% to 80% [1, 2]. Based on this small deployment we can extrapolate the possible energy savings with a wider deployment across the entire CSE building. Using network level techniques we have identified over 1000 desktop PCs in the department. Assuming an average of 100 Watts per desktop[1, 12] this translates to the total plug load of PCs as approximately 100 kW. Note, as can be seen from Figure 3, the total power consumed by all the plug loads is higher at 130 kW. The total energy use for all the PCs assuming they remain on is $100 \text{ kW} \times 8760 \text{ hours/year} = 876 \text{ MWh}$ per year. Now assuming a conservative 60% average energy savings per PC using Somniloquy, the total energy consumption across all PCs in CSE reduces to 350 MWh per year.

A large fraction of the total energy consumed by IT equipment in the CSE building is consumed by servers in the machine room. The average power consumption for the machine room across 2009 was 142 kW. The use modalities of servers is, however, quite different than desktop PCs. Servers are usually a shared resources among different users or applications, and can even be used to run processing jobs or serve data during nights and weekends. However, previous research has shown that even servers are often under utilized [9, 6] and techniques such as workload consolidation or virtual machine migration can be used to reduce the total number of physical servers that need to be active at any given time [14]. Techniques such as Somniloquy can be used in addition to consolidation based approaches to put

	Lighting	Plug Loads	Mechanical	Machine Room	Total
Status Quo					
Annual Energy Usage	438 MWh	1139 MWh	571 MWh	1,244 MWh	3,392 MWh
Average Power	50 kW	130 kW	65 kW	142 kW	387 kW
After Energy Efficiency Upgrades					
Annual Energy Usage	96 MWh	613 MWh	456 MWh	622 MWh	1,787 MWh
Average Power	11 kW	70 kW	52 kW	71 kW	204 kW

Table 2: Annual energy use (1st line in each field) and average power draw (2nd line) use of different building components in CSE and reduction after implementation of energy efficiency measures. 30 kW was added to the plug load number calculated in section 4.1 to reflect energy use from other miscellaneous office equipment (non-computer related) that is not reduced in our case study.

un-utilized machines to sleep to save energy. At status quo, the total energy consumed by the servers in the machine room at CSE is measured to be: $142 \text{ kW} \times 8760 \text{ hours/year} = 1,244 \text{ MWh}$ per year. Based on the server baseloads, we estimate that this consumption can be reduced (conservatively) by 50% to 622 MWh per year using a combination of server consolidation and Somniloquy.

4.2 Lighting

Lighting causes a peak load of 75 kW (0.5 W ft^{-2}) during workday afternoons, but lighting also has a large base load contribution. During these times, except for safety lights, lighting could be turned off 30 minutes after motion sensors detected the last movement. This would be practical during off-hours when the occupancy throughout the building is minimal. While all lighting could be turned off during off-hours in principle, for safety reasons lighting in hallways, bathrooms and other areas that provide access throughout the building have to be always on. Based on building plans, it was estimated that in 20% of the building the lights should be kept on for safety reasons and 80% of the 75 kW peak load could be eliminated during off hours (6pm-8am). The overall energy use with motions sensors is estimated as $75 \text{ kW} \times 50 \text{ hours a week} \times 52.1 \text{ weeks per year} = 196 \text{ MWh}$ per year during work hours and $0.2 \times 75 \text{ kW} \times 118 \text{ hours a week} \times 52.1 \text{ weeks a year} = 92 \text{ MWh}$ per year during off-hours resulting in a total of 288 MWh. Currently fluorescent or CFL lighting is used in the CSE building. LED lights could be installed which – for the same light output in the building – reduce the power consumption by 1/3. Consequently energy use with the sensor system and LED lighting of 288 MWh \times (1/3) = 96 MWh.

4.3 Solar PhotoVoltaics

While geothermal systems can be used to generate hot water and heat or cool buildings, such systems do not work in all soils and cannot produce electricity at a reasonable cost. For wind power, on the other hand, the roughness of the urban terrain causes a dramatic reduction in wind speed and enhanced turbulence which are not conducive to wind power except for on tall skyscrapers. Consequently, solar PV is the only viable option to generate renewable electricity for most buildings. Unless building-integrated PV were used, the potential solar power that can be collected is a function of roof area. We estimated roof space using Google Earth as 2700 m^2 . Assuming a solar irradiance to AC conversion efficiency of 15%, typical for high-end commercially available panels and inverters, a 405 kW_p system could be installed on the roof. This is larger than the average power draw of the building (Table 1). However, 405 kW is the power output at 1000 W m^{-2} solar irradiance which only occurs during sunny

days around noon. To get the overall output over the year the optimal solar panel tilt of 33.5° and an optimal azimuth angle of 10° West [13] are applied to obtain an irradiance of 2.1 MWh m^{-2} of roof area (or 111 kBtu ft^{-2} or 4.4 W ft^{-2} per square foot of building area). Thus the total output of the fixed tilt PV system would be $2.1 \text{ MWh m}^{-2} \times 2700 \text{ m}^2 \times 0.15 = 850 \text{ MWh}$. This is only 25% of the current annual energy use and still smaller than the total energy use after energy efficiency upgrades.

An opportunity to increase the harvesting of solar energy is through tracking solar panels. In San Diego a 2-dimensional tracking solar panel could receive 2.6 MWh m^{-2} of solar energy. Consequently a tracking solar panel with a 48% solar to AC conversion efficiency would be able to make CSE a zero net energy building (ZNEB). If energy efficiency measures were implemented then a 25% conversion efficiency would be sufficient to power CSE. For reference the thermodynamic limit for solar conversion efficiency is 31% for single junction solar photovoltaic (the vast majority in use today) and 66% for multi-junction solar PV. The best solar conversion efficiencies demonstrated in the laboratory are at 32% for multi-junction solar PV [10].

5. DISCUSSION

Modern buildings represent a prototypical cyber-physical system with deeply coupled embedded sensing and networked information processing that has increasingly become part of our daily lives. Modeling a building as a cyberphysical energy system will play a critical role in achieving and operating zero net energy buildings (ZNEB). Our analysis has shown that even though most buildings are lightly or not occupied at night, nighttime energy use in buildings is significant and includes a large and increasing fraction of energy use by the computing and networking equipment. We made the case that for 'mixed-use' office buildings, that is, building containing substantial energy use by humans and computing equipment, there exists a significant potential for increases in energy efficiency by joint optimization of the coupled subsystems. To substantiate this, we presented a case study of an intensely metered mixed-use office buildings. We showed that even though a modern large office building included a variety of energy efficient components, due to computing-intensive components it had a larger energy use intensity than the existing building stock. The large energy use was found to be related to server and PC energy use which constitute more than 50% of total electricity consumption. Only if off-hour energy saving features are implemented for the computing load, this building's energy use will be reduced below that of the existing building stock and would make it possible to achieve LEED energy

optimization credits for the building.

Achieving net zero use, however, is a much harder goal. Potential energy savings were estimated to be 80% for lighting, 60% for computing, 50% for server rooms, and 20% for mechanical loads. Once these measures are enacted, then covering the roof with efficient, but commercially available tracking solar photovoltaic panels would have the potential to create a ZNEB based on our measurements and modeling of the incident solar radiation. At current efficiency levels, however, dramatic improvements in solar conversion efficiency would be necessary to harvest enough energy to create ZNEBs based on solar energy harvesting alone.

Costs are an important factor in driving the conversion of buildings to ZNEB. If a utility allows net metering, a NZEB would incur no electricity cost resulting in annual savings of \$450k. Covering the roof with solar PV with a conversion efficiency of 15% would cost about \$1.9M after tax credits and incentives. The cost of energy savings measures are difficult to estimate. Based on an Internet search we estimate \$332k for LED lighting, \$50k for motion sensors and installation (we assume that the reduced maintenance cost of LEDs would be balanced by the higher maintenance cost for motion sensors), \$80k for more efficient HVAC pumps and fans, and \$70k (including maintenance) for Somniloquy in server rooms and desktop PCs. Consequently the simple payback time for these investments would be 5.3 years, which is too long for most businesses.

Power use at night has seen a resurgence of interest as state renewable portfolio standards (RPS) mandate a large fraction of intermittent energy production on the grid (e.g. 20% renewables in California by 2010). A large base-load at nighttime will cause grid operators to rely on coal and nuclear power plants to provide that load at low-cost. However, since cycling or even turning off these plants is economically and technologically prohibitive, they cannot serve to even out daytime fluctuations in solar power output. Consequently, a decrease in the difference between daytime peak power use and nighttime base load will limit economical solar energy penetration to about 10-20% of the total grid energy. Thus – in absence of energy storage – one strategy to make solar energy more attractive is to align the energy use profile with the solar cycle which means minimizing nighttime energy use. The ability to make the average nighttime energy use constant on the grid (e.g. by scheduling computer sleep modes according to real-time pricing signals) will lower greenhouse gas emissions by reducing required load following capacity, thus allowing use of only the most cost- or environmentally efficient power plants at night.

6. ACKNOWLEDGMENTS

We are indebted to Rajesh Gupta for his invaluable feedback and direction to help us write this paper. We want to acknowledge Thomas Weng who contributed significantly to build the UCSD Energy Dashboard. We also wish to thank John Dilllitt, Robert Austin and Bob Caldwell for their help in setting up access to the UCSD energy use data. Parts of this paper are based on an analysis by UC San Diego undergraduate environmental engineering majors Stephanie Chen, Jennifer Krawczyk, Robert V. Reith, and Ya-Ting (Sheila) Tsai for their MAE125A design project. This work is supported by Multiscale Systems Center (MuSyc) under the Focus Center Research Program (FCRP) supported by DARPA/MARCO.

7. REFERENCES

- [1] Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta. Somniloquy: Augmenting Network Interfaces to Reduce PC Energy Usage. In *Proceedings of USENIX Symposium on Networked Systems Design and Implementation (NSDI '09)*. USENIX Association Berkeley, CA, USA, 2009.
- [2] Y. Agarwal, S. Savage, and R. Gupta. SleepServers: A Software-Only Approach for Reducing the Energy Consumption of PCs within Enterprise Environments. In *To Appear at the USENIX Annual Technical Conference (USENIX '10), June 2010*. USENIX Association Berkeley, CA, USA, 2010.
- [3] Y. Agarwal, T. Weng, and R. Gupta. The Energy Dashboard: Improving the Visibility of Energy Consumption at a Campus-Wide Scale. In *ACM BuildSys '09*, 2009.
- [4] R. American Society of Heating and A.-C. Engineers. *ANSI/ASHRAE 90.1-2007: Energy Standard for Buildings Except Low-Rise Residential Buildings*. 2008.
- [5] California Energy Commission. CEC End-Use Survey, CEC-400-2006-005, March 2006. <http://buildingsdatatobook.eren.doe.gov/>.
- [6] J. Chase, D. C. Anderson, P. Thakkar, A. Vahdat, and R. P. Doyle. Managing Eenergy and Server Resources in Hosting Centers. In *Proceedings of SOSP '01*, 2001.
- [7] DOE. Buildings Energy Data Book, Department of Energy, March 2009. <http://buildingsdatatobook.eren.doe.gov/>.
- [8] Energy Efficiency Administration. Commercial Buildings Energy Consumption Survey (CBECS). <http://www.eia.doe.gov/emeu/cbecs/>.
- [9] X. Fan, W.-D. Weber, and L. A. Barroso. Power provisioning for a warehouse-sized computer. In *Proceedings of ISCA '07*, 2007.
- [10] G.W. Crabtree and N.S. Lewis. Solar Energy Conversion. *Physics Today*, 2007.
- [11] Itron Inc. California Commercial End-Use Survey. <http://capabilities.itron.com/ceusweb/>.
- [12] J. Koomey. Estimating Total Power Consumption by Servers in the US and the World. *Final Report. LBNL. February 2007*, 2007.
- [13] M. Lave and J. Kleissl. Maps of optimal tilt and azimuth for PV systems and advantages of tracking over the U.S. *in preparation for submission to Solar Energy*, 2010.
- [14] R. Nathuji and K. Schwan. VirtualPower: Coordinated Power Management in Virtualized Enterprise Systems. *ACM SIGOPS Operating Systems Review*, 41(6):278, 2007.
- [15] K. Roth and K. McKenney. Energy Consumption by Consumer Electronics in US Residences. *Final Report to the Consumer Electronics Association (CEA)*, 2007.
- [16] C. Turner and M. Frankel. Energy Performance of LEED for New Construction Buildings, report for US Green Building Council. *U.S. Green Building Council*, 2008.