

Understanding the Role of Buildings in a Smart Microgrid

Yuvraj Agarwal

Computer Science and Engineering
UC San Diego
La Jolla, CA 92093-0404
yuvraj@cs.ucsd.edu

Thomas Weng

Computer Science and Engineering
UC San Diego
La Jolla, CA 92093-0404
tweng@cs.ucsd.edu

Rajesh K. Gupta

Computer Science and Engineering
UC San Diego
La Jolla, CA 92093-0404
gupta@cs.ucsd.edu

Abstract—A ‘smart microgrid’ refers to a distribution network for electrical energy, starting from electricity generation to its transmission and storage with the ability to respond to dynamic changes in energy supply through co-generation and demand adjustments. At the scale of a small town, a microgrid is connected to the wide-area electrical grid that may be used for ‘baseline’ energy supply; or in the extreme case only as a storage system in a completely self-sufficient microgrid. Distributed generation, storage and intelligence are key components of a smart microgrid. In this paper, we examine the significant role that buildings play in energy use and its management in a smart microgrid. In particular, we discuss the relationship that IT equipment has on energy usage by buildings, and show that control of various building subsystems (such as IT and HVAC) can lead to significant energy savings. Using the UCSD as a prototypical smart microgrid, we discuss how buildings can be enhanced and interfaced with the smart microgrid, and demonstrate the benefits that this relationship can bring as well as the challenges in implementing this vision.

I. INTRODUCTION

The smart grid represents one of the most ambitious efforts that many countries are undertaking today, and buildings are one of the most critical components. Envisioned to have a higher degree of fault tolerance, reliability, robustness and efficiency, the smart grid will radically revolutionize the electrical networks that are vital to any nation’s infrastructure. The proliferation of technologies such as embedded sensing and networking at various points in the smart grid make this possible, and information along with electricity is transferred at every stage of electrical production, from generation to transmission and finally to the building. Recently, microgrids have emerged as prototypical testbeds for research and understanding on how smart grids behave. While microgrids are much smaller scale than wider-scale electricity distribution grids, they do provide excellent *controllable* testbeds since they have all the typical components that exist in a larger scale smart grid. For example, the UC San Diego microgrid has local co-generation capabilities to generate power using natural gas turbines as well as solar panels, an extensive power delivery infrastructure spanning across 1200 acres and energy storage solutions, which in the end powers over 450 buildings.

Buildings provide a clear and convenient target for analyzing and reducing energy use in a smart grid. Buildings are also among the dominant users of electrical energy, accounting for over 70% of total electricity use in the US in 2006 [5], [10]. This energy use is evenly split between commercial and residential buildings. The same is also true for the UCSD Smart Microgrid: most of the electricity consumption in the campus is due to buildings (thus making them the load centers). Therefore, studying the various factors that impact the energy consumption of buildings is key and can lead to mechanisms to reduce their energy use and improve the overall efficiency of a microgrid [1], [3].

The composition of energy use within buildings varies greatly based on several factors. The use modality of the buildings is a clear differentiator since it can represent a varying fraction of use factors: human occupants, laboratories, Information Technology (IT) equipment and climate control (HVAC) equipment. The age of the particular building is also important since older equipment usually is more inefficient, and newer buildings are designed with energy efficiency as a primary design goal but at the same time often permit greater utilization of spaces for human activity, thus leading to greater power consumption per unit area. For example, certification by the Leadership in Energy and Environmental Design (LEED) is a desired goal for all new buildings on campus. Traditionally, the Heating Ventilation and Air-Conditioning (HVAC) subsystem required to maintain hospitable environmental conditions is considered to be a dominant consumer of energy use by buildings. However, the energy use of IT equipment is increasingly challenging HVAC in a building, as the number of such equipment is beginning to match or exceed the human occupants in a building.

Thus by examining the ratio of energy use by IT equipment versus total energy use by the building, we can devise building archetypes that correlate with the functional purpose of the buildings. On one extreme is a pure ‘data center’ style building with almost 100% of the electrical going towards powering up compute servers, related IT equipment and HVAC to cool those machines. On the other extreme is a pure ‘occupant-only’ style building with most of the electrical load going towards




Fig. 1: Comparing various buildings in the UCSD micro-grid depending on the fraction of IT loads. The data presented is for a week in August of 2009.

maintaining occupant comfort, such as lighting and HVAC and plug-loads. Most of the buildings, especially modern constructions, fall into the middle category which we call ‘mixed-use’ buildings. These mixed-use buildings comprise of traditional electrical loads to support building occupants such as lighting and HVAC systems but more importantly have a significant amount of IT infrastructure in the form of PCs – desktops, laptops, computer servers and networking equipment – used extensively by the building occupants. Figure 1 compares the energy usage intensity (EUI) of seven different types of buildings across the UCSD micro-grid in terms of Watts/sq-ft. Figure 1 clearly illustrates that pure ‘occupant only’ style buildings such as the campus gymnasium and the residential halls have a lower EUI than the ‘mixed-use’ buildings (CSE) that have higher IT footprints. Buildings with heavy IT and laboratories have even higher EUI (CalIT2, Medical Research Building). Finally, the EUI of the SDSC, a building dominated by the computing equipment of racks of servers, is the largest and remains fairly constant.

As mentioned earlier, within a microgrid environment, optimizing the energy use of buildings and giving buildings finer-grained control of its subsystems is key to achieving overall energy efficiency. There are several challenges however: developing low cost energy metering and control solutions; managing IT energy consumption; making building energy use proportional to their actual use; deeply instrumenting buildings with sensors to detect occupant comfort and presence; data collection and networking to collect this sensor information and the system integration challenge to tie it all together.

This paper puts into perspective our earlier work where we examined in detail the breakdown of energy usage in buildings [3]. The major observations included the unexpected dominant role that IT power consumption plays in overall power consumption by the buildings, leading to particular relevancy of our work on SleepServer even in the context of buildings [19]. We also describe how incorporating additional sensors can optimize the other major building subsystem - HVAC [1]. We first review background on smart grids and

microgrids, and then discuss the challenges of improving building efficiency. Finally we show how such a building can interact with the microgrid, and the benefits that can result.

II. BACKGROUND

Traditional electricity networks are based on centralized generation of energy (e.g. large coal or nuclear plants), which is transmitted across vast distances at high voltages (110kV to 765kV). The transmission grid connects to distribution lines that connect to end users (e.g. homes and buildings) at lower voltages (120V to 240V). Traditional electricity networks however do not contain storage units, which means that the energy generated by the power plants must be balanced with energy consumed by the end-users. Coupled with the fact that current grids do not have real-time monitoring of electricity flow and electrical usage, maintaining this balance can be challenging, and unexpected events can lead to black-outs such as what happened to the US Northeast in 2003.

‘Smart grid’ refers to the modernization of the electrical grid [27]. Chief among the innovations required is real time monitoring of electricity as it is produced and consumed (using so-called smart meters), which will allow grid operators to have real time awareness of the grid. Bi-directional communication between consumers and grid operators is also required so that end-users can effectively respond to energy shortages (or peak-demand conditions) by lowering energy consumption; this is called demand response and is an area of significant focus by the grid operators as a part of their plans in responding to extreme natural or man-made events. Distributed co-generation and storage however has the potential to greatly reduce energy distribution losses and greenhouse gases.

Microgrids are a contained network of energy generation sources and energy storage that are connected to the buildings that consume the electricity. The generation can be from a renewable source, thus reducing the need to produce energy from fossil fuels. Energy Storage allows demand shifting – energy can be stored during lower demand times and used during peak demand periods. In addition, transmission losses are eliminated because energy generation neighbors the energy

consumers. A smart microgrid has real time monitoring of electricity flow and usage, and can control loads based on real time feedback from operators. By themselves, microgrids are ideal research testbeds on how to design and implement smart grid technologies. Connected to the smart grid, these microgrids allow for better accounting of energy as it flows between the two grids, and for increased resilience to anomalies.

One such microgrid is the UCSD campus. The UCSD campus contains over 450 buildings and during the course of a day has over 45,000 people [3]. Total campus power consumption can hit up to 40 MW, and is one of the largest uses of electricity in the city of San Diego. To meet this high demand, UCSD has installed power plants, including a 30 MW co-generation system containing two gas turbines and a steam turbine. Solar panels have been installed on the roofs of three buildings and can produce 3 MW combined. The UCSD central plant also generates chilled and hot water for heating/cooling across the campus. A 2.8 MW fuel cell powered by methane is used to act as an energy buffer during high demand periods. The campus plant is able to hit almost 80% of total campus energy demands, and the rest is imported from the city electrical grid.

The buildings on the campus are serviced by the campus-wide underground distribution grid. In addition, many buildings contain HVAC systems that are cooled by the chilled water loop supplied by the central plant. Buildings operations (such as HVAC control) are supported by the campus BACnet network for many of these buildings. Real time monitoring of the campus grid is an on-going effort. Over 60 buildings have industrial meters installed that can monitor energy consumption and a publicly accessible website called UCSD Energy Dashboard [3] has been released that visualizes the real time electrical consumption across these buildings. Figure 2 shows a campus map with the central plant and solar panels marked, along with the overall electrical flow. The computer science and engineering building is also listed, along with its power consumption characteristics.

III. OPEN CHALLENGES IN SMART MICRO-GRIDS

There are several challenges that need to be addressed in order to realize the vision of a smart microgrid and in particular the role that buildings must play. We outline some of the more pertinent ones and survey some recent research that attempts to address these problems.

A. Energy Metering and Control at Multiple Scales

One of the foremost challenges in a smart microgrid is that of accurate energy metering and accounting at multiple scales – from entire buildings (macro-scale) down to individual plug outlets (micro-scale). Accurate energy metering at a macro-scale of an entire building within a smart microgrid is crucial for several reasons. First it allows a system operator to identify the dominant energy consumers within a smart microgrid for capacity planning purposes as well as to provide system health status by identifying system faults as well as anomalies. Second, by analyzing long term trends, system

operators can identify where their energy saving efforts should be directed to get the most impact. While energy metering at the level of entire buildings provides overall energy usage, it does not provide a breakdown on a subsystem level within a building such as the HVAC or the lighting subsystem. Detailed breakdown of building energy consumption can provide further insights into the specific subsystems, such as HVAC or IT equipment[3]. The challenge with getting detailed breakdown within buildings using additional sub-metering is the cost of installation and deployment. Furthermore, a detailed subsystem level breakdown within a building is only possible if the original electrical layout of the building was done in a manner where the individual subsystems were on independent circuits.

While macro scale energy metering provides building and campus operators an overall picture, it is of limited use to individual occupants. The reasons for this are two fold. First, fine grained energy attribution is hard – i.e. which occupants of the building and to what fraction should the energy usage of a particular building be attributed to. Previous research has shown the pitfalls of simple approaches such as dividing the energy use equally between occupants with a constant base power, or basing it on occupancy measurements within the building [13]. The second challenge is that of granularity of control with macro scale energy metering. From a smart microgrid standpoint, being able to control individual plugged-in devices (by turning them on and off) is critical.

To address both these challenges, finer grained energy measurement and management is needed at an individual plug level. Using individual plug level meters, building occupants can measure and manage their own energy footprint. While several research efforts [14], [15] and commercial energy meters [29] are available, most of them either do not provide mechanisms to power electrical loads on/off or are extremely expensive. One challenge with metering at an individual plug-level is the large number of energy meters required to get significant coverage within a building. As a result, these metering solutions need to be low cost, accurate, easily deployable and conform to safety specifications such as UL. Recently, researchers have also proposed non-intrusive mechanisms to measure and attribute energy use to individual devices within residential environments. These approaches propose a single point of energy measurement augmented with external sensors, network activity measurements or disturbances in the electrical network to disambiguate energy use by specific appliances or occupants[17], [18], [25]. While these techniques of opportunistically exploiting sideband information can be useful to differentiate energy use, their scalability within large enterprise buildings still needs to be investigated.

B. Occupancy Sensing in Buildings

Of particular importance in designing smart buildings is accurate occupancy detection, as the main purpose of mixed-use buildings is to support human occupants. Building operations are typically static to support an “always-occupied” common case, which is wasteful. Therefore, knowing both occupancy patterns and real time occupancy is critical in




Fig. 2: Campus map detailing energy flow through the UCSD microgrid, including breakdown of the CSE building.

optimizing building operations, such as lighting and HVAC scheduling. Auxiliary services are also made possible, such as giving building operators real time awareness of who is in the building during emergency events. Challenges exist in measuring accurate occupancy however, as many of the current technologies have significant weaknesses.

While Passive infrared (PIR) sensors are used often for occupancy detection, they do not actually measure occupancy but rather movement [22]. PIR sensors work by measuring the infrared light from nearby objects, and determining occupancy when there is movement from the sensed objects. Some drawbacks of PIR sensors include the fact that they require line-of-sight to work, exhibit both false positives (detecting a person when no one is actually there) and false negatives (not detecting a person when someone is actually there), and they cannot detect occupancy when a person is still in its detection area [12]. Most systems that utilize PIR sensing will therefore have a timeout period for inactivity to determine when an area is no longer occupied [11]. Ultrasonic sensors, also commonly used, rely on sound and the Doppler effect to detect occupancy [12]. Like the PIR sensors, they detect movement and not actual occupancy, but do not require line-of-sight are able to detect movement around walls. State-of-the-art industrial occupancy systems will use both in a single package; these are called hybrid dual-technology sensors. CO2 sensors have been proposed to detect occupancy, but these can take upwards of 30 minutes to detect changes [28].

In the research community, other methods of occupancy detection have been explored. Camera systems have been explored; combined with sophisticated detection algorithms, it is possible to determine the number of occupants and not just binary occupancy [16]. Camera systems can be accurate, but issues relating to privacy and cost are difficult to solve. Using information technology activity as a proxy for occupancy has also been researched, such as using computer network data to determine if someone is in their office [17]. Recently, we have shown that by combining sensors occupancy detection can be improved significantly. For example, by combining a

PIR sensor with magnetic reed switch we can detect occupancy accurately over 90% of the time in an office environment [1].

C. Data Collection and Management

Wired protocols have traditionally been used within buildings for the majority of the communications. Due to installation costs, many occupancy sensors are therefore actually just tied in to localized controls, such as lighting while other sensors are connected to the building management system. However, wired communications have a very high installation cost, one that becomes even more infeasible after a building has already been commissioned and built. Therefore, for cost reasons in deploying in existing buildings, wireless sensor networks (WSN) are more appropriate.

WSN has been a well researched field over the past decade, and standardized protocols have been used in both industry and research. The IEEE standard 802.15.4 specifies the physical and MAC layer for the most widely used WSN protocols. Like Wi-Fi, 802.15.4 radios transmit in the 2.4 Ghz range, and the MAC layer uses carrier sense multiple access (CSMA/CA) for wireless channel access. A rich ecosystem exists in academia centered around TinyOS [24], an operating system designed for WSN, and many platforms (motes) exist that supports it. In industry, Zigbee, which is built on top of the 802.15.4 protocols, specifies high level communication standards and has seen rapid adoption.

Several challenges exist, however, in deploying WSN widely within buildings. Buildings already contain significant wireless traffic (due to Wi-Fi) [20], which is at odds with the requirement of a smart building (occupancy sensing, electricity metering) that require a substantial number of wireless devices be deployed in a given area. Building protocols that can deal with this wireless contention and interference while sustaining high data rates is therefore essential. In addition, cost is an important factor in deployment, and thus advances in fabrication will enable more sensors to be economically deployed. Managing all of the devices is a very difficult problem, one that has seen some solutions, but new methods, tools, and algorithms will be required to effectively handle

a large number of sensors [4]. Security and authentication is also a critical issue that needs to be addressed. Existing methods revolve around AES encryption and trust center authorization, however due to the sensitive nature of the sensing devices, and the fact that actuation or control of a building depends on these data sources, techniques specific to the smart building environment must be enacted [30]. Finally, the data collected by the WSN must be combined with existing inputs in a building in a consistent representation for analysis and combined actuation. For example, most modern buildings have a building management system or SCADA system that operates to control some of the building operations. These protocols are usually industrial standards such as BACnet or LONWorks. By applying a standard web-based API on top of all of the different protocols and data sources these diverse set of sensors can be composed together [9].

D. Combined Optimization and Operation of Buildings

In the previous sections, we outlined the challenges associated with energy metering and control at different scales, fine grained occupancy detection and data collection, management and visualization of the data. The final, and arguably the most important step, is actuation and control of the various elements within a smart microgrid to optimize for overall energy efficiency while taking into account occupant comfort within buildings. These actuations can be local in scope such as turning off electrical devices within an office during periods of unoccupancy or can be global in scope where actuations are governed by policies and events at the entire microgrid scale such as demand-response conditions.

The first step in actuation is analyzing the sensor data from all the sources within the microgrid. This data includes detailed energy use data within the subsystems of buildings, individual plug loads, and detailed occupancy information using either dedicated sensors or inferred indirectly using proxies (e.g. network traffic). In our previous work [3] we showed that in a modern mixed-use building (CSE department) the dominant loads fall into four categories: electrical and thermal loads related to the HVAC subsystem (10% - 35%); plug-loads which are mostly IT equipment such as computers, monitors, and printers (25% - 40%); lighting load (9% - 15%); and the machine room loads consumed by servers and HVAC equipment for cooling (30% - 36%). While the specific fraction of these loads may change for different mixed-use buildings, our data underscores the importance of managing IT and HVAC related loads within buildings.

The HVAC related energy consumption can be reduced by using the detailed occupancy information such that the HVAC system is duty-cycled and only enabled for particular thermal zones where occupancy is detected. Previous work has investigated the feasibility of occupancy sensors, within a residential context, to detect, model and predict occupancy patterns for HVAC control [21]. Our preliminary work within mixed-use buildings has also shown the feasibility of upto 20% energy savings using our custom built, wireless and inexpensive occupancy sensors[1]. We are currently working

on interfacing directly with the building Energy Management System (EMS) to enable us to actuate specific thermal HVAC zones based on information collected by our occupancy sensors (Section III-B). Lighting loads can also be similarly managed based on occupancy information (many current buildings already implement local lighting controls via PIR sensing). However, from our data we have observed that that lighting loads are actually the least power consuming component of our mixed-use building.

IT loads represent a significant and rapidly growing fraction of energy use [3], [19], [26] within mixed-use buildings. While in theory IT loads, especially PCs, can be duty-cycled by using low-power sleep modes during periods of low utilization, it has been shown that most users chose not to do so for reasons ranging from maintaining remote and administrative access to resources or data or in cases even running active applications while they are away [2], [8]. The proposed solutions, to reduce IT energy use, target this issue of maintaining “seamless network presence” using either lightweight proxies [2] or full-fledged virtual machine images [8] that can act on behalf of PCs while they sleep to save energy. We have deployed our proxying solution, SleepServers, to fifty users within the CSE department and have measured on average 70% reduction in energy use per computer [2]. Techniques similar to SleepServers can also be applied to manage the energy consumption of server loads within the machine rooms of mixed-use buildings. In addition, there has been much work in optimizing server power, especially within the context of data centers such as virtualization and workload scheduling and consolidation that can be leveraged [6], [7], [23].

While independently managing the energy consumption of individual subsystems within buildings is vital, integrating and managing the energy consumption at the scale of an entire microgrid provides even greater opportunities for improving efficiency. For example, by knowing the variation in renewable energy production by photo-voltaics (PV) or by price signals from imported energy from the grid, a system operator can manage energy consumption more efficiently. Periods of cheaper electricity or abundant PV generation can be used for energy storage or by the HVAC system to pre-cool buildings only to be turned off later during periods of low PV generation. In addition, scheduling algorithms can take into consideration energy parameters, time shifting computation as needed.

Another area of microgrid wide cooperation is in enacting demand response mechanisms for handling peak demand at the microgrid level. Buildings plays an important role by reducing their energy consumption automatically based on events relayed by a smart microgrid operator. If a peak demand condition occurs, the building energy consumption can be reduced dynamically by turning off non-essential devices through the smart energy meters, and put unused IT equipment to sleep using architectures such as SleepServers. This is key to truly integrating the building with the smart microgrid, as each building will be able to adjust its electrical consumption in fine-grained detail, ensuring stability of the entire microgrid.

IV. CONCLUSIONS

Smart microgrids are networked small-scale electrical distribution systems that combine local co-generation sources, energy storage solutions, and buildings as energy sinks. Buildings play a critical role in these microgrids, providing opportunities for improving overall energy efficiency and reliability of the microgrid while also improving occupant comfort inside the building. As the dominant consumers of electricity, understanding building energy usage and the associated factors that affect it is therefore key. The majority of electrical loads within buildings can be grouped into IT equipment, lighting subsystems, individual plug-loads and HVAC related equipment. All these subsystems must be optimized to improve the energy efficiency of individual buildings. To do so, some challenges must be addressed which revolve around deploying a network of wireless sensor nodes widely within buildings, collecting the data from those sensors and finally designing control algorithms that analyze this sensor data and in turn actuate the appropriate building subsystems optimally. While independent optimization and control at a subsystem level within a building can lead to substantial savings, combined optimization of different buildings across a smart micro-grid presents even greater opportunities such as responding to demand response scenarios, reducing peak demand of individual buildings and leveraging renewable energy sources.

ACKNOWLEDGMENT

We wish to acknowledge the Synergy Lab members: Bharathan Balaji, Seemanta Dutta, Michael Wei and Sathyanayanan Kuppaswamy for their contributions to the building occupancy project. We would also like to thank John Dilliot, Robert Austin, Walter Richardson and Bob Caldwell for helping us interface with the facilities' systems. This work is supported in part by the Multiscale Systems Center (MuSys) under the Focus Center Research Program (FCRP) supported by DARPA/MARCO and also NSF CCF/SHF grant 1018632.

REFERENCES

- [1] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, "Occupancy-Driven Energy Management for Smart Building Automation," in *Proceedings of the Second ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys '10)*, 2010.
- [2] Y. Agarwal, S. Savage, and R. Gupta, "SleepServer: A Software-Only Approach for Reducing the Energy Consumption of PCs within Enterprise Environments," in *Proceedings of USENIX Annual Technical Symposium (USENIX ATC '10)*, 2010.
- [3] Y. Agarwal, T. Weng, and R. Gupta, "The Energy Dashboard: Improving the Visibility of Energy Consumption at a Campus-Wide Scale," in *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys '09)*, 2009.
- [4] G. Barrenetxea, F. Ingelrest, G. Schaefer, and M. Vetterli, "The hitchhiker's guide to successful wireless sensor network deployments," in *Proceedings of SenSys*. ACM, 2008.
- [5] California Energy Commission, "CEC End-Use Survey, CEC-400-2006-005, March 2006." <http://buildingsdatabook.eren.doe.gov/>.
- [6] J. Chase, D. C. Anderson, P. Thakkar, A. Vahdat, and R. P. Doyle, "Managing Energy and Server Resources in Hosting Centers," in *Proceedings of SOSP '01*, 2001.
- [7] C. Clark, K. Fraser, S. Hand, J. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of Virtual Machines," in *Proceedings of the 2nd Conference on Symposium on Networked Systems Design and Implementation (NSDI '05)*. USENIX Association, 2005, pp. 273–286.
- [8] T. Das, P. Padala, V. Padmanabhan, R. Ramjee, and K. Shin, "Litegreen: Saving Energy in Networked Desktops using Virtualization," in *Proceedings of the 2010 USENIX Annual Technical Conference (USENIX ATC '10)*. USENIX Association, 2010, p. 3.
- [9] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, and D. Culler, "sMAP: a simple measurement and actuation profile for physical information," in *Proceedings of SenSys*. ACM, 2010.
- [10] DOE, "Buildings Energy Data Book, Department of Energy, March 2009." <http://buildingsdatabook.eren.doe.gov/>.
- [11] V. Garg, "Smart occupancy sensors to reduce energy consumption," *Energy and Buildings*, vol. 32, no. 1, pp. 81–87, 2000. [Online]. Available: [http://dx.doi.org/10.1016/S0378-7788\(99\)00040-7](http://dx.doi.org/10.1016/S0378-7788(99)00040-7)
- [12] X. Guo, D. Tiller, G. Henze, and C. Waters, "The performance of occupancy-based lighting control systems: A review," *Lighting Research and Technology*, 2010.
- [13] S. Hay and A. Rice, "The case for Apportionment," in *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys '09)*. ACM, 2009.
- [14] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler, "Design and Implementation of a High-Fidelity AC Metering Network," in *Proceedings of IPSN/SPOTS '09*, 2009.
- [15] X. Jiang, M. V. 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*. ACM, 2009, pp. 113–126.
- [16] A. Kamthe, L. Jiang, M. Dubys, and A. Cerpa, "SCOPES: Smart Cameras Object Position Estimation System," in *Proceedings of the 6th European Conference on Wireless Sensor Networks*. ACM, 2009.
- [17] Y. Kim, R. Balani, H. Zhao, and M. B. Srivastava, "Granger Causality Analysis on IP Traffic and Circuit-Level Energy Monitoring," in *Second ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings (BuildSys)*, 2010.
- [18] 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*. ACM, 2009, pp. 245–254.
- [19] J. Kleissl and Y. Agarwal, "Cyber-physical Energy systems: Focus on Smart Buildings," in *Design Automation Conference (DAC), 2010 47th ACM/IEEE*. IEEE, 2010, pp. 749–754.
- [20] C.-J. M. Liang, N. B. Priyantha, J. Liu, and A. Terzis, "Surviving Wi-Fi Interference in Low Power ZigBee Networks," in *Proceedings of the 8th ACM conference on Embedded Network Sensor Systems (SenSys)*, 2010.
- [21] J. Lu, T. Sookoor, V. Srinivasan, G. Ge, B. Holben, J. Stankovic, E. Field, and K. Whitehouse, "Efficient Application Integration in IP-based Sensor Networks," in *Proceedings of the 8th ACM conference on Embedded Network Sensor Systems (SenSys)*, 2010.
- [22] Michigan State Energy Observer, Dec. 2007, "Occupancy Sensors for Lighting Control," http://www.michigan.gov/documents/dleg/EO_12-07_218809_7.pdf.
- [23] R. Nathuji and K. Schwan, "VirtualPower: Coordinated Power Management in Virtualized Enterprise Systems," in *Proceedings of Twenty-First ACM SIGOPS Symposium on Operating Systems Principles (SOSP '07)*. ACM, 2007, pp. 265–278.
- [24] P. Levis and S. Madden and J. Polastre et al, "TinyOS: An Operating System for Wireless Sensor Networks," <http://www.enocean.com/en/energy-harvesting-wireless/>.
- [25] S. Patel, T. Robertson, J. Kientz, M. Reynolds, and G. Abowd, "At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line," in *Proceedings of the 9th international conference on Ubiquitous computing*, 2007.
- [26] Smart2020, "SMART 2020: Enabling the Low Carbon Economy in the Information Age," http://www.smart2020.org/_assets/files/02_Smart2020Report.pdf.
- [27] US Department of Energy, 2009, "The Smart Grid: An Introduction," http://tonto.eia.doe.gov/ask/electricity_faqs.asp.
- [28] S. Wang and X. Jin, "Co 2-based occupancy detection for on-line outdoor air flow control," *Indoor and Built Environment*, vol. 7, no. 3, pp. 165–181, 1998. [Online]. Available: <http://ibe.sagepub.com/cgi/content/abstract/7/3/165>
- [29] WattsUP, "WattsUP Energy Meters," <http://wattsupmeters.com>.
- [30] Y. Xiao, V. Rayi, B. Sun, X. Du, F. Hu, and M. Galloway, "A survey of key management schemes in wireless sensor networks," *Journal of Computer Communications*, 2007.