A Distributed Energy Monitoring and Analytics Platform and its Use Cases

Neil Klingensmith
Computer Science
Department
1210 West Dayton Street
Madison, WI, USA
naklingensmi@wisc.edu

Dale Willis
Computer Science
Department
1210 West Dayton Street
Madison, WI, USA
dale@cs.wisc.edu

Suman Banerjee
Computer Science
Department
1210 West Dayton Street
Madison, WI, USA
suman@cs.wisc.edu

ABSTRACT

We present Emonix, a distributed, low-cost system for monitoring and analyzing energy consumption patterns in buildings. Emonix is designed with our custom energy sensing hardware and integrated communication units to be efficiently mounted in breaker panels of buildings. In contrast to plug-based monitoring systems, this approach is less intrusive to users because it does not intrude on their physical space, yet it still provides fine-grained real-time energy data in both space and time. The Emonix hardware platform is open, modular, and extensible. It provides an accessible data and configuration API for users, and we believe it is useful to the broad community. To demonstrate the usefulness of this platform, we have deployed this infrastructure on two campus dormitories covering 60 rooms and 120 residents. We have been operating this infrastructure as an energy monitoring service for the residents for more than four months to help them understand their consumption patterns at different timescales. Our results indicate significant temporal variations in energy consumption patterns at different time scales, and that a small fraction of occupants can consume a disproportionately large amount of energy in such buildings.

Categories and Subject Descriptors

C.3 [**SPECIAL-PURPOSE AND APPLICATION-BASED SYSTEMS**]: Real-time and embedded systems

Keywords

Smart Energy, Energy Sensing, Sensor Platforms

1. INTRODUCTION

Energy consumption of buildings has grown significantly and is a dominant contributor to global energy utilization. In the United States for example, energy consumption of buildings has risen by 48% since 1980, and it now represents 74% of the nation's electricity consumption [1].

The first step to control continued growth in energy consumption of buildings is to increase awareness among users and occupants. Understanding energy consumption patterns within buildings has, therefore, been a domain of continued research activity over the years [12, 15, 16]. In this paper, we report on the design and deployment of Emonix[1], a low-cost system for collecting and analyzing fine-grained electrical energy consumption patterns in buildings.

Why a new platform? When we started this project, we searched for a platform that would allow us to cheaply and easily deploy many energy sensors in a commercial building. We needed a platform that would easily permit software modifications to the metering devices. However, we were not able to find a commercial system with an open API that met our cost constraints. Table 1 lists estimated costs for deploying several commercial energy monitoring systems for a panel with 30 breakers. For each system, we included the cost of sensors, current transformers, power supplies, and other ancillary equipment.

Emonix design requirements: With Emonix, we aimed to capture a holistic view of the energy consumption patterns of a large building. One challenge of managing a large building is that different areas are often occupied or controlled by small groups which may not cooperate or even communicate with one another. As such, it may be difficult for one individual or small group to understand the ways that resources such as energy, ventilation, and space are shared among building occupants.

As a first step, we believe it is necessary to understand the electric power consumption of the building. However, we also acknowledge that there are other factors that contribute to a building's energy footprint. Our approach is therefore to design a system that has user-modifiable software which would allow us to easily extend the system's capabilities in the future. For this design, we identified the following requirements:

- **Open and accessible API:** Our energy sensor hardware should have an open and accessible API through which various monitoring tasks can be configured. For instance, it should be possible to instruct the sensors to collect energy samples at different granularities. It

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.
BuildSys'13, November 14 – 15 2013, Roma, Italy
Copyright 2013 ACM 978-1-4503-2431-1/13/11 ...$15.00.

---

[1] Emonix stands for Energy MONitoring and analytIX.

Product	Cost
eMonitor [6]	$899
Modlet [8]	$1065
TED [7]	$3326
Watts Up .Net [11]	$7078
eGauge [2]	$1072
Veris E30 [10]	$6114
Emonix	**$330**

Table 1: Estimated cost of commercial energy sensors for a panel with 30 breakers. The cost for Emonix reflects only price and assembly. However, our devices are purchased at low volume, so our components cost are higher than they would be for large volumes. These figures include estimated costs for all components of the system, including sensors, current transformers, power supplies, etc. For Emonix, these figures only reflect hardware costs, while commercial products necessarily include costs of doing business such as technical support, engineering time, etc.

should also easily permit the addition of new features such as addition of network interface types or new sensor configurations.

- **Expandable sensor ports:** Each breaker panel often has a large number of branch circuits. Each sensor board should have a number of sensor ports to simultaneously collect energy measurements from multiple branch circuits. In addition, it should be possible to attach additional low-cost daughter boards to the main sensing board to flexibly increase the number of branch circuits being monitored. This approach will keep the costs of the overall system low.

- **Fine-grained reporting:** The energy sensors themselves should provide access to as much information as possible about the power distribution system. Ideally, this would include not just access to information about power consumption of a load, but also power quality metrics such as power factor and harmonics as well.

- **Flexible communication alternatives:** There are different alternatives to communication paths out of the sensing hardware. One could use some wireless technology, e.g., WiFi and ZigBee, or some wired counterparts, e.g., power line communication or even RS-232. We have found that none of these individual communication options works well in every scenario. Hence, a board which can be equipped with different communication alternatives is likely useful in diverse scenarios.

- **Accessible data:** The collection backend provides easy access to the measurements taken by the sensors. This may be as simple as exposing an API that allows us to query the database that stores data collected by the system.

There are many commercial products in the market that provide one or two of these features. Table 2 lists the features of a few of the meters we evaluated. Most of the energy monitoring systems we evaluated were targeted either toward plug-level monitoring (for single appliances) or utility entrance monitoring (for an entire building). In principle, it might have been possible for us to modify one of these off-the-shelf devices to monitor the power consumed by a single breaker instead. However, there were two hindrances to such an approach. First, none of the commercial meters provided adequate granular access to the energy data being collected. Second, each of these commercial systems would be very expensive to deploy at scale, primarily because most of the systems on the market are not designed to monitor many breaker panels or individual circuit breakers. For these reasons a more cost-effective monitoring solution is possible.

Hence, we chose to architect a new system that is targeted specifically toward monitoring the power consumed by many branch circuits in a breaker panel. We were able to achieve dramatic cost reductions in the electric current sensing equipment because the overhead of the enclosure, microcontroller, power supply, etc. are amortized over many current sensors. Furthermore, we did not need to outfit every sensing device with an expensive wireless communication interface as is done by many outlet monitors because all of our sensors are located in the same room next to the breaker panel. As a result, an installation of Emonix could cost as little as $10 per breaker, while commercial meters such as Ted or Kill a Watt cost more than $50 per breaker.

In addition, with Emonix, we aim to provide an *open* hardware and a flexible software platform[2] for energy research that is accessible to researchers, regardless of their hardware background. Most existing commercial platforms currently on the market do not expose any sort of API that would make it possible for experimenters to write and test custom software. By contrast, Emonix is fully implemented in the C programming language. Experimenters can modify any component of the system, including the energy sensors themselves, through our open, clean, and well-documented interface. We believe this will allow others to build upon our initial successes.

Use cases through dormitory deployments: To demonstrate the usefulness of Emonix, we have conducted multiple energy monitoring pilots. For these pilots, we have worked with UW-Madison's Housing department. In particular, we have deployed Emonix in two different dormitories on our campus — Cole Hall and Chadbourne Hall. Among them, Cole Hall had a particular interest in our system because its residents are taking an active approach in reducing their energy footprint through sustainability clubs. In each dormitory, two students share a single room. Our deployments have covered 120 residents for a period of over four months, and provide some insights into energy consumption patterns in campus dormitories.

Key contributions: The following are the primary contributions of this work:

- We present our end-to-end electrical energy monitoring infrastructure, including our custom energy sensors for cost-efficient monitoring of branch circuits, with multi-modal communication capabilities, and various software components, e.g., databases and web-based dashboards.

- Our energy sensing hardware will be open, extensible,

---

[2] Source code and design files available at
`http://research.cs.wisc.edu/wings/projects/emonix`

Product	Accessible API	Expandable Sensor Ports	Fine-Grained Reporting	Flexible Communication	Accessible Data
eMonitor		✓			✓
Modlet		✓	✓	✓	✓
TED		✓	✓		✓
Watts Up .Net			✓		
eGauge			✓		✓
Veris E30		✓	✓		✓
Emonix	✓	✓	✓	✓	✓

Table 2: Features of Commercial Power Meters. Emonix provides all features necessary in our deployment.

and flexible, and it implements an open API that allows users to add new modules to the system.

- To demonstrate usefulness of this platform, we have deployed them across multiple dormitories in our campus for a period of more than four months covering 60 rooms and 120 residents, and report on some interesting observations from the measurement study.

2. THE EMONIX ENERGY MEASUREMENT AND MONITORING INFRASTRUCTURE

Emonix is an embedded electric energy monitoring platform that can gather and analyze fine-grained energy consumption patterns at low cost and through minimal intrusion to the end user. Due to the unique constraints of our deployment environment, we chose to design custom energy sensing hardware and communication backend to deliver near real-time power consumption measurements to a database server. These measurements are made available to residents using digital signage in the building as well as a web-based monitoring dashboard with email updates[3].

2.1 Emonix Architecture

The monitoring system is a hybrid of multiple monitoring, communication, and storage technologies. The energy monitoring platform can be broken up into three primary functions:

1) Measure the power consumed by branch circuits,
2) Process the information as necessary for different applications
3) Relay the results to our database backend.

Custom energy measuring hardware:

The real-time sensing is performed by a Freescale ColdFire 51QE microcontroller. Our sensor boards use current transformers to sense the amount of electric current passing through each branch circuit in a breaker panel (see Figure 2). The microcontroller has more than twenty ADCs, allowing us to monitor many inputs on one board. Each analog input is pre-conditioned by a low-pass antialiasing filter that attenuates frequencies greater than 120 Hz. This ensures that high-frequency harmonics present on the voltage or current waveforms will not alias with lower frequencies. Since high frequencies typically carry a small fraction of the power consumed by an appliance, this does not reduce measurement accuracy by a large amount. In fact, calibration experiments done in our labratory indicate that our sensors have less than 1% error in their power computations.

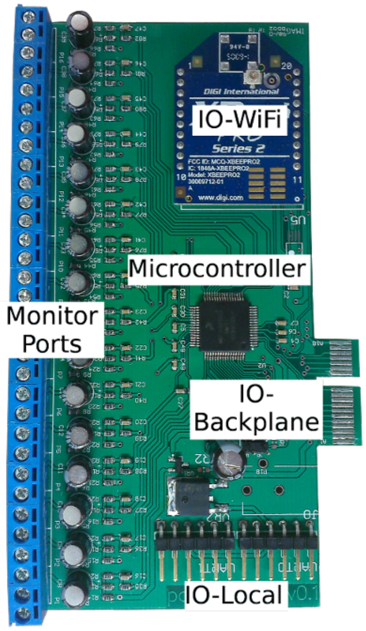

Figure 1: Custom energy sensor board. Monitors up to 16 current transformers and has three distinct forms of communication.

By default, each individual ADC is sampled at a frequency of 1.92 kHz (a harmonic of 60 Hz). The ADCs are sampled at high frequency to ensure that high-order harmonics that may be present in the current waveform do not alias with the 60 Hz fundamental. Harmonics above 1.92 kHz are attenuated by more than 80 dB by the analog antialias filter. After each 64 samples have been collected from each ADC, we apply a simple low-pass digital filter to the sampled waveform and downsample. We then compute the fast Fourier transform (FFT) of the downsampled waveform and transmit the magnitude and phase of the 60 Hz component. Using the API configuration interface, the above settings can be easily modified and even extended.

The communication interfaces on the energy sensor board were carefully designed to allow flexibility for various deployment scenarios. The sensor board is designed to provide an efficient monitoring solution for deployment environments

---

[3] Available to residents through a secure and private website.

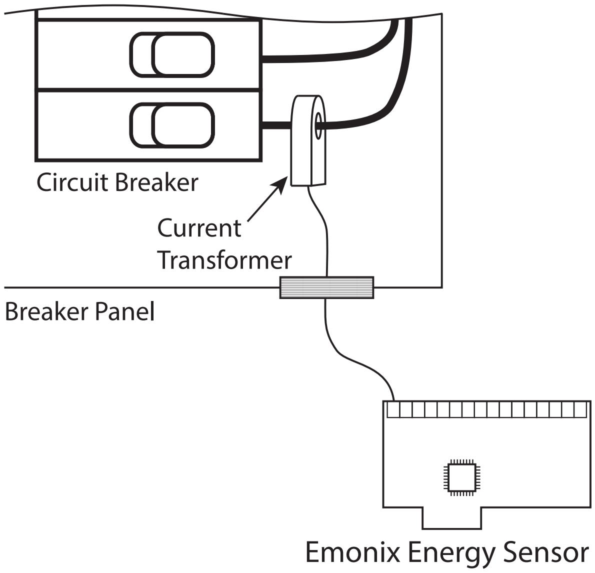

Figure 2: Connection of energy sensor board to monitoring device in breaker panel. Each branch circuit is monitored by a single current transformer, which is an analog device that can sense the current flowing through a conductor. The energy sensor board digitizes the analog signal generated by the current transformer and transmits the data to a gateway.

with varying building wiring configurations, floorplans, and networking topologies. Figure 1 labels the three communication interfaces that the sensor board can use to relay the data it collects to a database server.

`IO-WiFi`: Communicate directly with an available WiFi Access Point to send the data to a server or gateway

`IO-Local`: Communicate directly with a gateway attached with an RS-232 serial cable

`IO-Backplane`: Attach to neighboring sensor boards on a *backplane* board to forward data to one gateway computer

In our residence hall deployments we take advantage of `IO-Backplane` in order to communicate with the gateway. It should be noted that while `IO-Backplane` looks very similar to a PCIe connection, it cannot be plugged into a computer's motherboard. Instead we have designed a "backplane" board, which presents PCIe slots (as seen in Figure 3). This allows up to five[4] sensor boards to share the same enclosure, power supply, and communication methods. These interfaces can also be used to load custom code onto the sensor board.

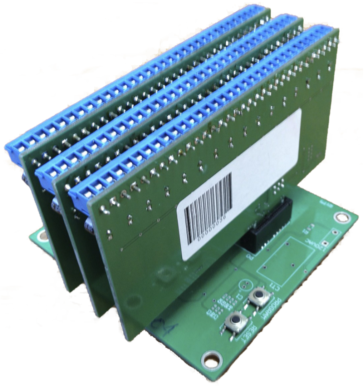

Figure 3: Multiple energy sensor boards connected to a backplane.

API Call	**Implementation**
`pthread` Library	Threads, Mutexes
`sleep`	Full
`time` / `stime`	Get/set UNIX epoch
`malloc` / `free`	Full
BSD Sockets	Partial

Table 3: Features supported by Emonix's API.

API Accessibility: Emonix also provides a complete open-source software set, including drivers for real-time data collection, and a simple RTOS that provides a POSIX-compliant interface to the programmer. Developers can compile custom software using `gcc` [19] and load their code onto a sensor board over the serial or WiFi interfaces. Emonix's API provides a subset of the POSIX calls that UNIX and Linux programmers are familiar with, including pthreads, memory allocation, and networking. Table 3 lists some of the features implemented by Emonix's API.

Using the Emonix API, we have written a few custom applications to run in place of the energy sensing software. In just an afternoon, we were able to completely replace the energy sensing software with an application that monitors and reports environmental data including temperature, humidity, and barometric pressure. The new application was multithreaded and took advantage of the Emonix network stack for exchanging messages with our database server. This exercise was a convincing demonstration of the power of our API.

Transport energy data: Once the data has been initially collected by the sensor board it must be packetized and transmitted to a gateway where it is aggregated, stored locally, and forwarded to a server for later use. In small buildings (single-family homes) with single breaker panels, a single gateway would be sufficient. In larger buildings such as the ones we currently have deployed in, a single gateway is unable to handle the entire distributed set of sensor boards and their continuous stream of energy measurements. Hence, we deploy multiple gateways which are responsible for collecting data from disparate sensor boards and forwarding them appropriately to the server backends. We use low-cost Single Board Computers (SBCs) that consume very little energy as gateways in these deployments.

Process, store, display data: After the data has been aggregated by the gateway, it is transmitted to a database server. Once the data is on the database it can be stored for future use, or processed into a form more amenable for

---

[4] Five is our standard design, however a backplane board could be built to increase this size to practically any number of sensor boards

website viewing. The processing takes fine-grained data and computes commonly viewed data points (daily, weekly, and monthly energy consumption totals) for webserver access. The computation occurs once and is stored rather than each viewing needing to recompute the same data. This processing keeps database queries quick and minimizes latency for website viewing.

Flexibility and communication: Emonix is built to maintain flexibility in the face of unique electrical service layout, as is typical in commercial buildings of varying age. For instance, in one residence hall we are required to aggregate up to 90 breakers in a utility room. In another location, flush mounted breaker panels in common space would require our system to aggregate only 24 breakers and must reside above the ceiling tiles. Due to these varying constraints, we designed Emonix in a scalable way, which is useful to both an individual household as well as large dormitories.

The ability to provide as little as one and as many as 80[5] monitor ports allows Emonix to maintain a level of scalability unique among current energy monitoring systems. The cost-effective nature of our platform is due to several design iterations, which have allowed us to determine the most effective configurations.

Generally we rely on a building's Internet connection to transmit data to the database server. However, many buildings do not always have reliable network connectivity in the utility rooms where breaker panels often reside (for both wired and wireless communication systems). In some cases, these electrical rooms have cinderblock construction that severely limits wireless communication in and out of these rooms. Due to these constraints we typically utilize a combination of communication techniques including: power-line communication (PLC), WiFi, Ethernet, and RS-232.

3. IMPLEMENTATION CHALLENGES

3.1 ADC Sampling

Real-time ADC sampling presented several unique challenges, most of which were timing-related. To get accurate results, the ADC sampling must be exactly right.

Phase Error. Since the ADCs that monitor the current and voltage must be sampled sequentially, there is a phase difference between the samples taken from each ADC channel. Phase error can create a problem when computing power, since the computation depends on the phase angle between current and voltage. We corrected for this problem by adding a different constant phase offset to the phase computed by the FFT before calculating the power. The delay between samples of our ADCs resulted in a phase offset of approximately 0.122π. To adjust for the phase error, we subtract $n \times 0.122\pi$ from the phase of each waveform, where n is the ADC index.

Real-Time Sampling. Since each of sixteen ADCs must be sampled at a frequency of 1.92 kHz, the processor must handle a total of 30,720 ADC conversions per second, each of which is processed by an interrupt service routine. Running at 60 MHz, this leaves only 1,953 CPU cycles per ADC interrupt for code execution, meaning that the CPU can execute roughly 500-1,000 instructions per ADC exception. To put this in perspective, a 64-point FFT takes roughly 1 ms, or 60,000 CPU cycles on our processor.

Delay. In a related problem, we found that other interrupt service routines could delay the execution of the ADC sampling ISR, resulting in sample jitter. This would cause the energy in the 60 Hz component of the signal to spread to neighboring frequencies, resulting in inaccurate current and voltage measurements. To avoid this problem, we had to make the ADC sampling ISR the highest priority event in the system. Unfortunately, this requires us to take processing time away from communication and data processing tasks, which have their own real-time deadlines.

3.2 Integer Math

We chose to do all data processing with integer (fixed-point) math because floating-point support is typically expensive, power-hungry, and slow. All of our data is stored in 16-bit integers, meaning that we need just over 2kiB to store 64 samples from each channel simultaneously[6]. By contrast, IEEE floating point representation would require 4 kiB or 8 kiB depending on the representation standard.

However, integer math introduces more severe errors in the calculations than floating point math would. The errors are most pronounced in trigonometric functions such as sine and arctangent, which are used by the FFT. These errors are the primary contributors to the error in power consumption, and result in an overall error of less than 1%.

3.3 Safety

The national electric code in the United States requires power electronics to be certified by Underwriter Laboratories (UL) of another certified lab. UL standards specifically target equipment that comes in direct contact with conductors carrying line voltage (120V in the US), and are designed to minimize the risk of fire and electrocution. Most electricians will not install equipment that is not certified by UL for liability reasons. Our equipment is compliant with ANSI/UL 61010-1 and UL 796.

4. USE CASES IN RESIDENCE HALLS

We now discuss the data gathered from our deployment in the two dormitories over a period of several months (September to December 2012). There are various temporal trends, variations across residents, and numerous general energy consumption characteristics that we observed.

Thus far we have discussed the advantages of our platform monitoring at the branch circuit level, however there are a few caveats to note. First of all, while we are able to monitor each branch circuit individually, this does not guarantee that 100% of the energy consumed on the circuit belongs to a particular individual. For instance, if two residents in the same dorm room share a refrigerator, the current implementation of Emonix is unable to disaggregate and divide the energy consumed by the device between each roommate in a fair manner. This would also apply to roommates who may share computers, printers, etc.

4.1 Applications

---

[5] 16 monitor ports on each sensor board times five boards per backplane.

[6] $2\ bytes \times 64\ samples \times 17\ ADCs = 2176\ bytes$

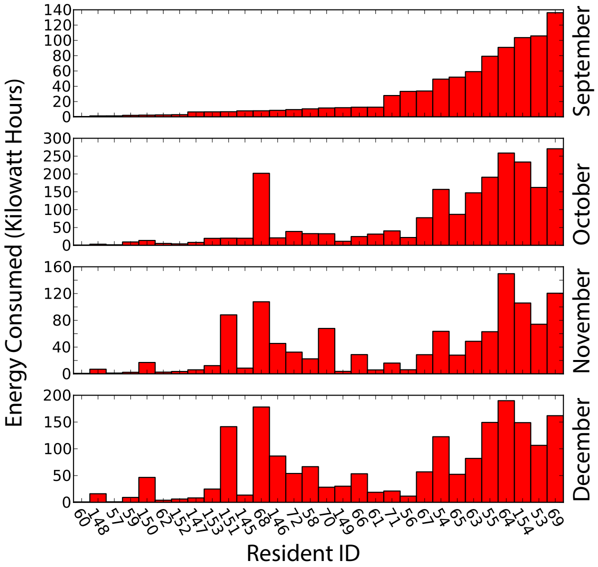

Figure 4: Total energy consumption per resident for the duration of our deployment for each month.

The total data gathered is quite large and diverse, and it is not possible to report all aspects of it in entirety. Instead we will focus on some representative results. The dataset presented here corresponds to a set of 30 residents who have used our system the longest. We use randomly generated resident IDs to anonymize the different residents in this set.

Before we began the study, we took surveys of the building residents to try and understand the number and kind of appliances that were commonly used in a dorm room setting. The most common appliances we encountered were microwaves, refrigerators, laptops, and lamps. The housing department prohibits appliances with heating elements such as toasters and hot plates because they pose a fire hazard, so no residents reported these kinds of appliances. A lack of this class of power-hungry appliances would seem to make university residence halls a fundamentally different type of deployment environment by comparison to other residential buildings.

Temporal behavior across months: Figure 4 shows data we have collected across all four months during our study. The residents are sorted by their energy consumption amount in the month of September, and this ordering is maintained for the remaining months. As is generally true across typical user populations, there is a great diversity among residents in the amount of consumption. In particular, a disproportionately large amount of energy is consumed by a small percentage of residents. This is observed in each month of the study. This implies that energy conservation efforts may benefit most when focused on the small group of residents who manifest high usage patterns.

Previous studies of energy consumption patterns in large populations have found the distribution of energy consumption among users to be lognormal [20]. Our user population is not large enough to make that claim with any degree of statistical confidence, but based on the data we have gathered, it would not be surprising to find that the data were in fact lognormally distributed.

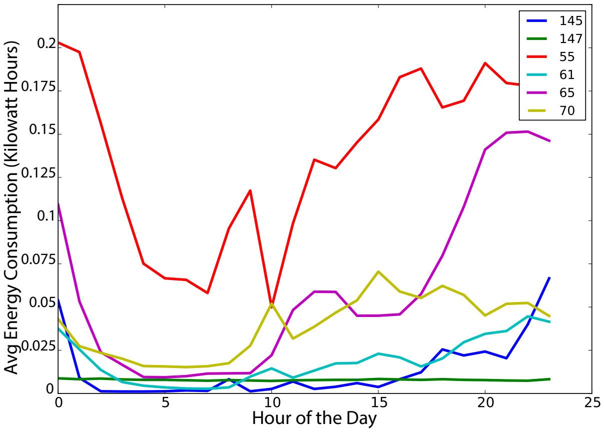

Figure 5: Average energy consumption as a function of hour of the day. Different lines correspond to different randomly chosen residents.

We can also see that the energy usage behavior of many users stay consistent from month to month. However, this trend is not universal, e.g., residents 68 and 151. Finally, as expected, total energy consumption grows from September to December as the seasons change.

Behavior across a day: Most residents follow very consistent daily patterns as is evident in their average energy consumption for each hour of the day in Figure 5. There is a consistent trend which indicates that most users do not maintain the same level of energy consumption for all hours of the day.

Time of Use vs. Flat-Rate Billing: The data we collected from dormitories using Emonix was used by another group on our campus to evaluate electric utility billing policies. Our local electric utility offers its customers two billing rate options: fixed rate and time of use. Under fixed-rate billing, utility customers pay the same rate for the energy they consume at all times of the day, plus a daily customer charge. Under time-of-use billing, customers pay a reduced rate (about 50%) for electricity consumed during off-peak hours (9PM – 10 AM daily) and an increased rate (about 160%) for electricity consumed during peak hours.

Figure 6 shows the hypothetical electricity charges for one week of energy consumption for several of the biggest energy users in this study. We found that everyone we studied would benefit in some degree from using the time-of-use billing scheme. We attribute this to the propensity of college students to stay up late at night, thereby shifting their peak usage to later hours. In fact, many of the students who participated in this study are likely to be absent from their dorm rooms until late at night. Figure 5 clearly indicates that resident 65 routinely uses most of their energy between 8PM and midnight. This type of behavior would result in considerable savings under the time-of-use billing scheme.

5. RELATED WORK

Energy monitoring represents a diverse body of work since it can be approached from many different angles (ie. sensor systems, cyber-physical building controls, feedback tech-

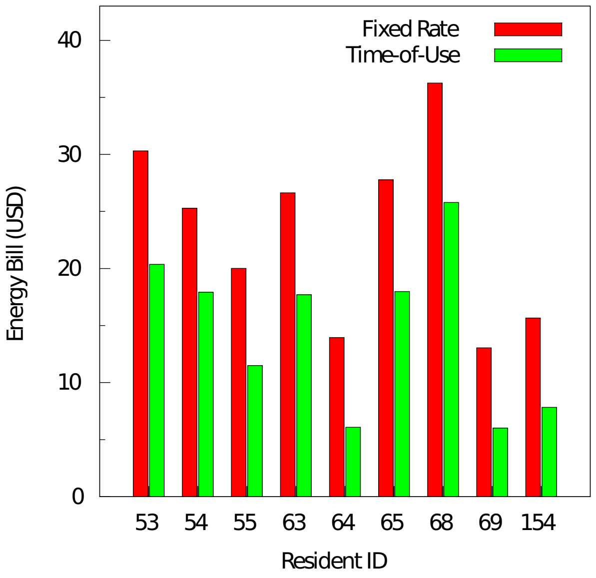

Figure 6: Hypothetical energy bill for several residents. These figures are based on our local electric utility's billing rates [5]. The red bar shows each resident's hypothetical bill under a fixed energy pricing scheme, which is the scheme used for most residential utility customers. The green bar shows the energy bill for time-of-use billing, in which the utility charges higher rates during peak hours and lower rates during off-peak hours.

niques, system architectures). Commercial entities also contribute to energy monitoring with a plethora of their own sensors and architectural paradigms.

5.1 Energy Monitoring Hardware

We do not claim that Emonix is the only effort toward energy monitoring, there are many efforts in several different areas which should be noted. One particular custom monitoring system is very similar to Emonix, however it does not appear to focus on expandable hardware [24]. Other research tends to focus on energy monitoring at the building level without being able to break it down further [16]. In a few cases, an infrastructure has been built around using commercially available hardware, such as [13]. Several commercial entities provide a fully implemented monitoring solution [2, 9, 4, 7], however each deployment is isolated from every other and access to the proprietary API or building data is often difficult. Some research also combines energy monitoring methods like [22], which contains both plug and building level monitoring. The focus of that work was to show how to reduce the energy consumption of the metering hardware using different methodologies. The three-layer architecture used by Emonix is very common for sensor networks, and is implemented in many energy monitoring research such as [16, 13, 22].

In general, there are several reasonable ways to monitor the electricity consumption of a building. Perhaps the simplest approach would be to monitor the total energy drawn by the whole building at the *utility service entrance.* This approach requires only a single monitor and is used in several commercial products we are aware of [9, 4, 7]. However, this approach does not allow residents to see their unique contribution to the overall energy consumption of the building. This makes it difficult for researchers to relate the overall consumption to each occupant. Using building level energy consumption would only allow for building or campus scale predictions and interpretations of said data [21, 12]. Also, we do not want to include electric demand of appliances that residents do not have direct control over such as HVAC systems or hot water heaters because the demand, implementation, and efficiency of such appliances varies from building to building. In other words, we want our measurements to reflect the behavior of individual residents as much as possible.

The second possibility is to monitor electric demand on a per-appliance level. Using this technique, each appliance would have to be individually metered. For instance, one could deploy commercial plug-level energy monitors such as KillAWatt[3] or Watts Up?[11]. These would require one sensor board for each wall outlet and require that residents plug in their appliances to such monitoring units only.

Plug-level monitoring was also not useful for our purpose for many reasons. First, the cost of these appliances in the market is not particularly low. Second, it would be extremely intrusive — deployment of such a system would require explicit cooperation and attention of each resident. Placing and mounting these energy monitors could be disruptive to room layout as well as present aesthetic issues. If the sensors are perceived to be inconvenient, some residents might find it easier to bypass them when installing their appliances, thus skewing the results. For instance, Hnat et al. details the issues associated with deploying large numbers of somewhat intrusive sensors[17]. They discover that many issues are created if the user has the ability to easily bypass the sensor, which we assume would be the case in a dormitory setting. Third, in many commercial systems there is no communication infrastructure associated with plug-level monitors, or if one exists it is difficult to interface with the device, therefore automating the process of energy monitoring would be increasingly difficult. Finally, plug-level monitoring devices have the inherent flaw of only being able to monitor outlets. This is an issue due to the fact that in many commercial buildings the lights and HVAC systems are not connected to electrical outlets, therefore it is impossible for pass-through style plug-level monitors to capture their energy consumption.

While we do not focus on energy monitoring at the plug-level, there are several noteworthy research efforts using such systems. A pass-through plug level monitoring system, ACme [18], is designed around wireless technology with a fully implemented IPv6/6LoWPAN networking stack. A very similar system is proposed in [15], however the authors focus on other sensors to aid in the energy accountability for each person as they use appliances in common space.

5.2 Energy Profiling

Monitoring at the branch circuit level can have many advantages over other alternatives. In [12], the authors compare different buildings using the "Watt/sq-ft" metric. While Watt/sq-ft is a reasonable comparison method for building-level energy monitoring, in some cases it may be unfair. This unfairness could be derived from the heterogeneous nature of buildings from both a physical(age, updates, and maintenance) and utilization standpoint. For instance, can a server farm in a computer science building be directly compared to lab space in a chemistry building? Perhaps several servers could shut down, and their loads distributed to other active servers. However a lab's ventilation system cannot shutdown to save energy while the lab is in use.

The ability to add other parametric analysis techniques to building comparisons could be a powerful tool. Taking advantage of systems like Emonix, or even others such as [13] or [22] to compare people on a more direct level could lead to interesting ways to reduce energy consumption. Finally, using more specialized analysis techniques could result is better intervention mechanisms, which may help influence behavior patterns, such as those discovered in [23, 14].

6. ACKNOWLEDGMENTS

All authors are supported in part by the US National Science Foundation through awards CNS–1040648, CNS–0916955, CNS–0855201, CNS–0747177, CNS–1064944, CNS–1059306, and CNS–1230751. This research was also supported in part by the Department of Energy (DOE) Office of Energy Efficiency and Renewable Energy (EERE) Building Technologies Program (BTP) under the BTP Innovators Program administered by the Oak Ridge Institute for Science and Education (ORISE) for the DOE. ORISE is managed by Oak Ridge Associated Universities (ORAU) under DOE contract number DEAC05–06OR23100. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of NSF, DOE, ORAU, or ORISE.

7. REFERENCES

[1] US Department of Energy, 2012. http://buildingsdatabook.eren.doe.gov.

[2] eGauge, 2013. http://www.egauge.net/products.php.

[3] Kill A Watt, 2013. http://www.p3international.com/products/special/p4400/.

[4] Lucid Design Group, 2013. http://www.luciddesigngroup.com.

[5] Madison Gas and Electric, 2013. http://www.mge.com/home/rates/tou.

[6] Powerhouse Dynamics, 2013. http://www.powerhousedynamics.com/buy-emonitor-dealers/.

[7] TED 5000 Home Electricity Monitor Store, 2013. http://www.theenergydetective.com/residential/store.html.

[8] The Modlet, 2013. http://themodlet.com/buy_office.html.

[9] Uihlein Electric, 2013. http://www.uihleinelectric.com/.

[10] Veris Inudsties, 2013. http://www.veris.com/Item/E30A042.aspx.

[11] Watts Up?, 2013. https://www.wattsupmeters.com/secure/products.php?pn=0.

[12] Y. Agarwal, T. Weng, and R. K. 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, pages 55–60, Berkeley, CA, USA, 2009. ACM.

[13] R. Brewer and P. Johnson. Wattdepot: An open source software ecosystem for enterprise-scale energy data collection, storage, analysis, and visualization. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 91–95, 2010.

[14] R. Brewer, G. Lee, and P. Johnson. The kukui cup: A dorm energy competition focused on sustainable behavior change and energy literacy. In *System Sciences (HICSS), 2011 44th Hawaii International Conference on*, pages 1–10, 2011.

[15] Y. Cheng, K. Chen, B. Zhang, C.-J. M. Liang, X. Jiang, and F. Zhao. Accurate real-time occupant energy-footprinting in commercial buildings. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '12, pages 115–122, Toronto, ON, Canada, 2012. ACM.

[16] M. Dominguez, J. J. Fuertes, S. Alonso, M. A. Prada, A. Moran, and P. Barrientos. Power monitoring system for university buildings: Architecture and advanced analysis tools. *Energy and Buildings*, 59(0):152 – 160, 2013.

[17] 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*, SenSys '11, pages 232–245, Seattle, WA, USA, 2011. ACM.

[18] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, IPSN '09, pages 253–264, Washington, DC, USA, 2009. IEEE Computer Society.

[19] N. Klingensmith. Building gcc as a cross-compiler, 2011. http://pages.cs.wisc.edu/~klingens/files/crossgcc.pdf.

[20] J. Z. Kolter and M. J. Johnson. Redd: A public data set for energy disaggregation research. *SustKDD Workshop on Data Mining Appliations in Sustainbility*, 2011.

[21] J. Z. Kolter and J. F. Jr. A large-scale study on predicting and contextualizing building energy use. *Proceedings of the Conference on Artificial Intelligence, Special Track on Computational Sustainability and AI*, 2011.

[22] D. Lachut, S. Piel, L. Choudhury, Y. Xiong, S. Rollins, K. Moran, and N. Banerjee. Minimizing intrusiveness in home energy measurement. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '12, pages 56–63, Toronto, ON, Canada, 2012. ACM.

[23] G. Peschiera and J. E. Taylor. The impact of peer network position on electricity consumption in building occupant networks utilizing energy feedback systems. *Energy and Buildings*, 2012.

[24] S. Taherian, M. Pias, G. Coulouris, and J. Crowcroft. Profiling energy use in households and office spaces. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, e-Energy '10, pages 21–30, Passau, Germany, 2010. ACM.