A Magnetic Field-based Appliance Metering System

Niranjini Rajagopal † Suman Giri ‡ Mario Berges ‡ Anthony Rowe †

† Electrical and Computer Engineering Department
‡ Civil and Environmental Engineering Department
Carnegie Mellon University
Pittsburgh, PA 15213
{niranjir,sgiri,mberges,agr}@andrew.cmu.edu

ABSTRACT

Understanding where electricity is being used in buildings is an important tool for Cyber-Physical Systems (CPS) used in building energy conservation and efficiency. Current approaches for appliance-level energy metering typically require the installation of plug-through power meters, which is often difficult and costly for devices with inaccessible wires or outlets, or appliances that draw large amounts of current. In this paper, we present an energy measurement system that estimates the energy consumption of individual appliances using a wireless sensor network consisting of contactless electromagnetic field (EMF) sensors deployed near each appliance, and a whole-house power meter. We present the design of a battery-operated EMF sensor, which can detect appliance state transitions within close proximity based on magnetic and electric field fluctuations. Each detector wirelessly transmits state change events to a circuit-panel energy meter, in a time-synchronized fashion, so that the overall power measurements can be used to estimate appliance-level energy usage. Our EMF sensors are able to detect significant power state changes from a few inches away, thus making it possible to externally monitor in-wall wiring to devices. We experimentally evaluate our proposed EMF sensor, three-phase power meter and communication protocol in a residential building collecting data for over a week. The system is able to detect appliance state transitions with an accuracy of 95.8% and estimate the overall energy with an accuracy of 98.1%.

ICCPS'13 April 8-11, 2013, Philadelphia, PA, USA.

Categories and Subject Descriptors

J.2 [**Physical Sciences and Engineering**]: Miscellaneous

General Terms

Measurement, Performance, Experimentation

Keywords

CPS, Energy Metering, Sensor Networks

1. INTRODUCTION

Given that the building sector is responsible for the largest portion of the electricity use in the world (e.g., 75% of the electrical demand in the United States in 2010 [1]), knowing when and where electricity is being used in buildings is an important first step towards improving our energy conservation and efficiency efforts. Monitoring and control strategies used in future smart building CPS require continuous, timely and accurate sensor feedback. In industrial environments, for instance, efficiency experts often use preliminary profiling of appliance usage to suggest equipment updates and/or load scheduling strategies. Similarly, energy audits are performed in many commercial and even residential buildings to discover efficiency or curtailment opportunities. While these strategies provide consumers with a better understanding of their consumption patterns at one point in time, more recently there has been a growing interest in the development and deployment of appliance-level energy monitoring systems with the goal of providing continuous feedback to end users about their consumption patterns with hopes of realizing energy savings that result from a change in behavior triggered by these feedback mechanisms. Some researchers have estimated this feedback-induced savings to be up to 20% [2, 3]. This appliance-level continuous feedback also provides building managers accurate information that can be used to target the most effective update and retrofit strategies.

Current approaches for metering typically require the appliances to be wired inline with energy meters. Unfortunately, this can be expensive in terms of both hardware (requiring full meters at each appliance) and installation cost. For example, many appliances such as overhead lights are hard-wired and would require an electrician to install each individual sensor. Furthermore, many of these plug-through meters have a maximum rated current that prohibits their use with some high-power loads such as those present in heating, ventilation and air conditioning systems, which happen to be some of the top energy consumers in commercial and residential buildings.

To overcome these limitations, in this paper we present a contactless electromagnetic field (EMF) sensor which detects appliance power consumption state changes within close proximity, based on measurements of magnetic and electric field fluctuations. We show how this sensor can be combined with whole-house power measurements obtained at the main electrical panel of the home, to accurately estimate the consumption of individual appliance over time in a low-cost, simple to install way.

Most modern magnetic field measurement devices either have a low sensitivity or are prone to noise, making them only suitable for monitoring large loads such as large electric motors. Our EMF detector, on the other hand, utilizes a highly sensitive auto-gain circuit along with a local event detection algorithm to significantly improve reliability. We also utilize both magnetic and electric field strength measurements independently to determine when significant changes have occurred in the appliance. Data from these sensors are then relayed back to the main meter using a low-latency wireless sensor networking protocol, where changes in the total power consumption of the house are used to determine the power usage of the appliance that caused it. As compared with other energy monitoring solutions, such as plug-through meters, this solution is low-cost, easy-to-deploy and can be installed without disrupting the current operation of appliances. The sensors are able to detect current changes associated with the appliance from a few inches away making it possible to externally monitor in-wall wiring to devices like overhead lights or heavy machinery that might operate on multiple phases of the AC distribution system of the building.

The main contributions of this paper are three-fold. First, we design a low-power and robust battery-operated wireless EMF sensor to detect significant fluctuations that correspond to a change in the operating status of an appliance. We refer to these changes as events, and sometimes call the sensor an EMF event detector. Our proposed event detection sensor is compact and consumes on average $45\mu W$, making it ideal for long-term battery-powered operation. Next, we design and evaluate a high-speed wireless data collection protocol that allows power events to be correlated with whole-house power measurements obtained through a custom power meter board capable of monitoring 3-phase AC lines at the main electric panel of the building. Finally, we evaluate the overall system's ability to estimate appliance-level energy consumption from measurements gathered during a 1-week deployment of the system in a single-family home.

1.1 Organization of the Paper

The rest of the paper is organized as follows. In Section 2, we provide an overview of relevant previous work in electricity monitoring. In Section 3, we describe the various hardware components of the system. Section 4 then discusses the event detector running locally on each EMF sensor as well as the approach used to label appliance power consumption using the three-phase meter. Section 5 describes the sensor networking protocol used to collect the data and Section 6 evaluates the overall performance of the system. Finally, Section 7 provides concluding remarks and suggests future work.

2. RELATED WORK

Multiple research projects have investigated improving the visibility of the electricity demand of buildings. The MIT Plug [4] provided users with power and sensor information by means of a smart surge protector. In [5], the author's present experiences using the ACme wireless plug sensor in an office environment. In [6], the authors present ViridiScope which uses indirect sensing of appliances to estimate per-person energy consumption. This work suggests using magnetic field sensors to estimate the power consumption of a device. This is similar in concept to our EMF event detector except that we perform local processing on a significantly more amplified signal to detect state changes from distances up to a few inches away from wires. We found that the geometry between the cable and the pickup as well as the power factor of the device being tested make it extremely difficult to estimate power consumption without device and installation-specific calibration. Instead, our EMF event detector focuses on detecting appliance state changes rather than trying to directly measure power. This information, if time-synchronized with panel meter data, can be used to quantify the power consumption of the appliance. Other researchers have attempted to address the automated annotation problem by using multi-modal sensor fusion schemes [7]. One of the primary focuses of this work is on intelligent local processing to improve sensing accuracy which would benefit these other systems.

Commercial devices like Energenie [8], Efergy [9] and Kill-A-Watt [10] have already become household names among those interested in monitoring their appliances for power use. These devices are basically portable plug-

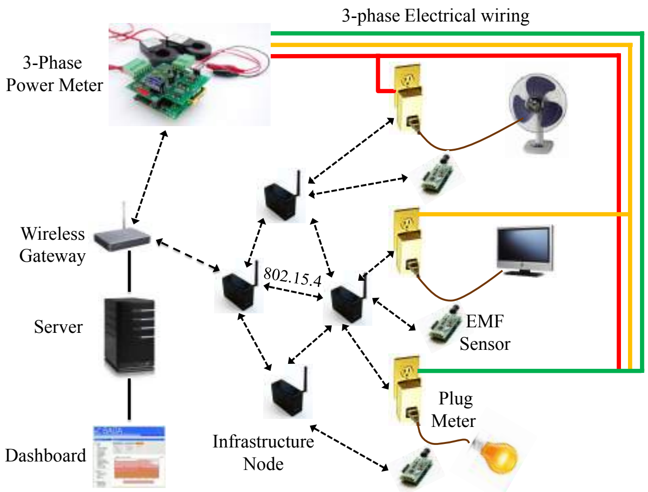

Figure 1: Network Architecture

level meters with a built-in display that provides information on real-time power consumed by the appliance plugged through it. The lack of a central platform for gathering information from multiple such meters means that users sometimes have to go around the house to collect this information. To remedy this, several wireless energy monitoring tools such as ACme [5], Tendril [11], GreenBox [12], EnergyHub [13] and Plogg [14] have been proposed and/or commercialized.

3. SYSTEM HARDWARE

In this section, we discuss the various components of our system that are required to collect and correlate appliance *on/off* events with circuit-level power data. Figure 1 shows an overview of the system architecture which consists of the three-phase meter used for overall mains power metering, plug-meter devices used for ground-truth data collection and the EMF detectors. These components are connected to a networked backend and displayed on the dashboard described in [15].

3.1 Circuit-Panel Meter

We designed a custom three-phase power meter, shown in Figure 2, which employs the cutting edge ADE7878 energy metering chip from Analog Devices, specifically to collect high resolution data which can be correlated with events from our EMF detectors. Off-the-shelf energy meters often make it difficult to capture high-speed raw waveforms. In contrast, our meter samples both the current and the voltage on each phase at 1KHz, and uses an on-chip DSP to compute true, apparent and reactive power, as well as several other energy metrics. The tight coupling between the power sampling and the radio interface reduces timestamp error between the signals sent from the appliance state detectors and the power values

Figure 2: Wireless Three-Phase Circuit-Panel Meter

Figure 3: Wireless Plug-Level Meter

at the circuit panel. A tighter time synchronization also improves the ability of the system to disambiguate temporally close appliance transitions. The main board is powered from either 120 or 240 VAC and can sense voltages as large as 600VAC. Current sensing uses split-core current transformers and both voltage and current values are read at 24-bit resolution. The overall range and accuracy values depend on the particular configuration of the current transformer used, but this configuration typically meets the 0.2% accuracy requirements for most utility billing standards.

3.2 Plug Meter

We use the FireFly plug meter [16] for ground-truth validation and for devices that can benefit from remote actuation. Each plug meter, shown in Figure 3 contains the ability to monitor and control two electrical outlets using wireless communication. The meter uses an efficient switching power supply that draws less than 0.1 watts ensuring that it does not unnecessarily increase the building power consumption. The meter measures true power, apparent power, power factor, frequency, RMS current and RMS voltage with a sampling rate of 1KHz. Two solid-state relays are used to independently control each outlet.

3.3 EMF Sensor

The core principle behind the EMF event detector is the ability to sense an appliance state change, by monitoring changes in nearby electromagnetic fields. From the laws of physics, we know that an alternating current

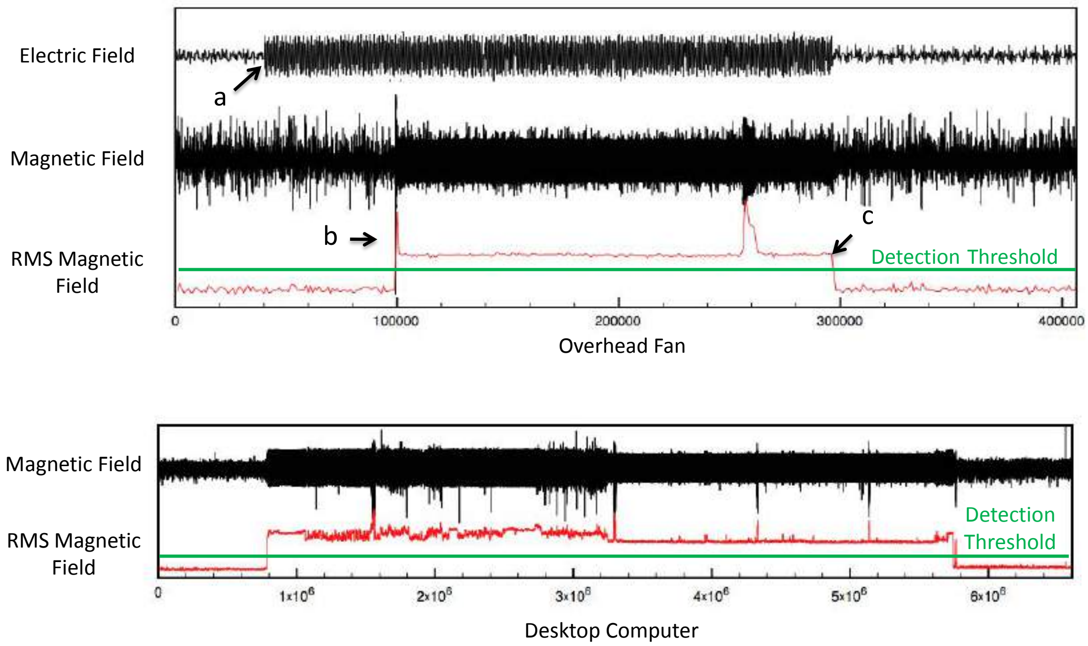

Figure 4: EMF Event Detector Waveforms. Top: Ceiling fan with light switch activated at point (a), manually turned on at point (b), and turned off at point (c). Bottom: Desktop computer is an example of a noisy signal due to switching.

flowing through a conductor will generate a corresponding magnetic field (H). Typically AC wires run as parallel pairs and hence most of the magnetic fields cancel out. However, imbalances in wires and stray currents flowing on ground lines as well as through appliances produce a significant magnetic field. The amplitude of this field is generally small (millivolts), but if sufficiently amplified, one can reconstruct the original source to a reasonable degree of approximation.

Each time an appliance changes its power consumption (transitioning between *on* and *off*), there is a corresponding change in the nearby magnetic field. In contrast, differences in voltages are responsible for creating electric fields. This means that an appliance that is not drawing current may still generate a strong electric field (E). The distinction between the electric and magnetic field is useful for two reasons. First, the electric field can be used to detect if a device is "live" or not. For example, overhead lights often switch the hot AC lines which can easily be detected by inspecting the electric field. Second, if a device is powered, but not active, the electric field strength can be used as a guide to find placement areas where there will be a strong magnetic field once current begins to flow. Since the electric field is not dependent on current flowing, abnormal fluctuations in the electric field tend to indicate potential noisy situations. For example, if people are nearby or touching the sensor, both the electric and magnetic field will be disturbed.

Figure 5: EMF event detector stacked on Fire-Fly3 sensor node.

Figure 6 shows a circuit that detects both magnetic and electric fields. The magnetic field is detected using an instrumentation amplifier (INA) and an inductor. We use a INA with a fixed 1000x gain that then feeds a high-pass capacitively coupled filter that removes DC bias to center the signal given a single ended voltage supply. The amplitude of the analog output generally corresponds to the strength of the magnetic field. The lower portion of the circuit uses a JFET and a small wire acting as a Hertzian antenna to detect potential differences across an electric field. The JFET opens or closes based on the change in force exerted by the electric field. The large-valued resistor between the gate and ground acts as a runoff to remove excess charge buildup from constant nearby fields.

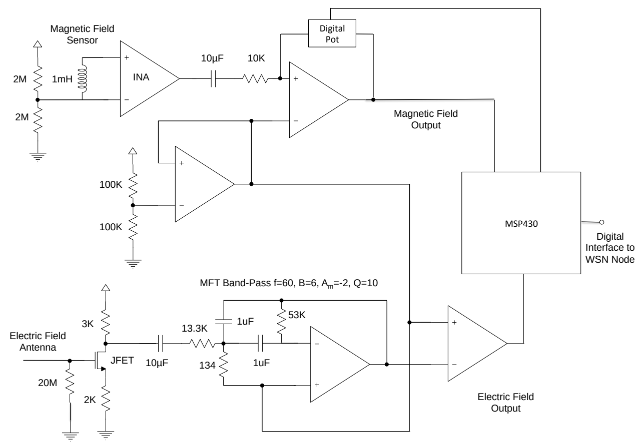

Figure 6: EMF Event Detector Circuit

Figure 4 shows two example waveforms received by the circuit when placed near a ceiling fan and a desktop computer. Point (a) in the ceiling fan waveform denotes the wall switch turning *on* which generates a corresponding electric field. At point (b), the ceiling fan is manually switched *on* (by pulling the hanging cord) causing current to flow and hence generating a magnetic field. The bottom line in the upper graph shows the root mean square (RMS) value of the magnetic field signal averaged over a window of 16ms (1/60Hz). The bottom graph shows the magnetic field and the same sliding RMS value for a desktop computer. In both these cases the edges in the RMS signal are quite pronounced.

Figure 5 shows a picture of the EMF detector hardware connected to a FireFly wireless sensor node. The FireFly node is responsible for periodically sampling the magnetic field in order to report appliance activation events. Since the signal from the EMF detector has a steady-state value associated with the current of the appliance, the FireFly node can duty-cycle its sampling to save energy. We explore this and provide additional evaluation results related to the sensor's sensitivity and detection range in [17]. We measured that the EMF detection front-end consumes approximately $45\mu W$ but this value can vary depending on the strength of the measured magnetic field.

4. SYSTEM SOFTWARE

In this section we first describe and evaluate the event detection algorithm running on each EMF detector. We then describe the approach used to label and record per-appliance energy usage by communicating with the three-phase meter.

4.1 Event Detection Firmware

The EMF sensor locally performs two main tasks: (1) it adjusts a hardware gain setting on the magnetic field sensing front-end to maintain a fixed peak-to-peak value for the sensed signal and (2) it is responsible for detecting significant changes in field strength and reporting those to the wireless sensor node.

Our first challenge was choosing an adequate ADC sampling rate such that the MSP430 could accurately reconstruct the magnetic field signal. Sampling too slowly would lead to poor performance, while sampling too quickly would allow lesser time for the CPU to sleep, resulting in greater power consumption. Figure 9 shows the detection accuracy along with the power consumption of the EMF detector as compared to the sampling frequency used by the MSP430 to sample the ADC. In order to determine a sufficient sampling rate, we validated the detection accuracy as compared to sampling rate using the experimental setup shown in Figure 7. The EMF sensor was placed at a distance of 5 *cm* (maximum range can be significantly larger depending on the load current) from the wire while two different appliances were transitioned (switched *on/off*) 40 times each. The accuracy metric is computed as the number of correct transitions divided by the total transitions detected in order to penalize false positives. For this experiment we used a 60 Watt fan and a 60 Watt incandescent light bulb since one is composed of a largely inductive load and the other almost entirely resistive. Figure 8 shows the raw magnetic field for a few cycles as viewed on an oscilloscope. One can clearly see that the fan has a unique signature that would make it more difficult to detect with a low sampling rate. Based on this graph, we chose to operate the EMF detector at about 1KHz since this maintains reasonable accuracy and is positioned just before the power dramatically increases.

In order to increase robustness on the magnetic field sensing front-end, the MSP430 performs a continuous auto-gain operation to keep the peak-to-peak range of the signal at approximately $\frac{V_{dd}}{2}$ so as to avoid clipping while still capture events of interest. As described above, the MSP430 samples the signal every 1ms over a window of 25 samples. To filter out noise, the peak values found in four such windows are averaged and the gain is adjusted every 100ms using a digital potentiometer. A simple three-stage step function controller is used to continuously track the input signal.

During the data collection, each EMF sensor transmitted its auto-gain value every 640ms. Significant changes in this value indicate that a nearby appliance has changed its power consumption. After inspecting the data offline, it became apparent that a simple $\frac{V_{dd}}{2}$ threshold was sufficient to detect binary appliance state changes. As future work, we intend to investigate detecting more subtle changes based on correlating the gain values with the mains power values to improve labelling of multi-

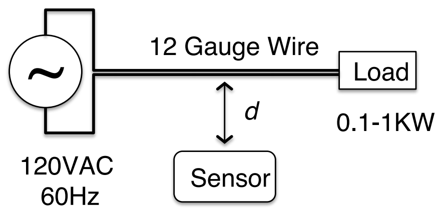

Figure 7: Experimental Setup

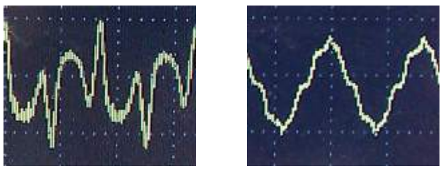

Figure 8: Raw magnetic field waveform for fan (left) and incandescent light bulb (right).

state or continuously variable appliances.

4.2 Appliance Metering

Every time an EMF sensor detects an event, it sends a time-stamped message to the gateway. Our algorithm then determines the change in power across a time window before and after the event (the size of this window is evaluated in Section 6). We assume that the power consumption for this appliance remains constant when the appliance is *on*. This constant value is found by averaging the absolute value of the power change during the *off*-*on* and *on*-*off* transitions, and the energy is found by integrating the power for the *on* period. Figure 10 illustrates this power computation process for a couple of hours while metering a refrigerator. The true power from the plug meter is also shown for validation.

A benefit of estimating power based on events, apart from its simplicity and effectiveness, is that it can handle finite state transitions of an appliance to estimate varying power consumption over a cycle, which has traditionally been a hard problem to tackle. The main limitation is that this approach fails to correctly estimate energy usage for devices whose power consumption varies slowly over time (and between events) or devices that have significant phantom loads.

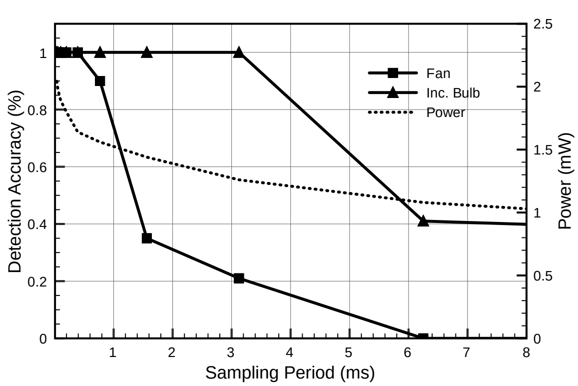

Figure 9: Accuracy and power vs sampling rate given a 50ms window.

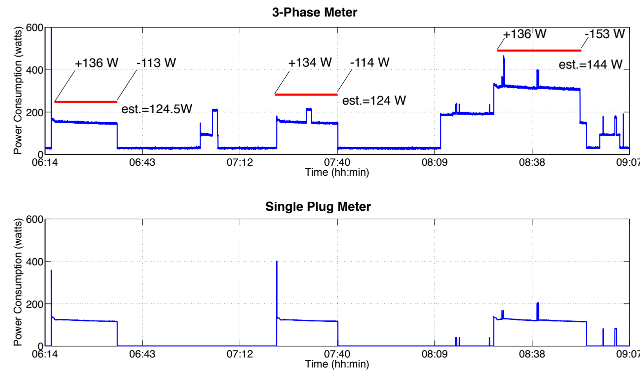

Figure 10: Top: Appliance energy estimator running on a refrigerator. Arrows show EMF event inputs onto the mains power waveform. The bars across the top represent the constant power. Botton: Ground truth from plug meter

5. HIGH-SPEED DATA COLLECTION

In order to facilitate ground-truth data collection, we required a networking protocol that was able to receive data from 64 nodes (our deployment homes had approximately 50 appliances in them), across multiple hops at sampling rates greater than 1Hz. A subset of these nodes were powered by AAA batteries and were hence highly energy constrained. We designed a protocol called mPCF that is loosely based upon the Point-Coordination-Function (PCF) mode of 802.11. The major differences are that we added support for efficient downstream multi-hop time synchronization using the Glossy [18] time synchronization primitive and included a simple multi-hop upstream packet-scheduling scheme. Unlike existing (often complex) TDMA protocols, our goal was to provide a simple implicitly scheduled system that did not require topology collection or scheduling at a master node and that could still provide high-throughput communication between root and leaf nodes.

In 802.11 PCF, a master node transmits a beacon to synchronize a set of client nodes. Each client node is assigned a time-slot in which they can reply in a collision-free manner each cycle. In our implementation, each node uses the lower byte of its MAC address as a reply slot assignment. This imposes the requirement that all nodes have unique subnet addresses (lower byte), which set the implicit schedule. All downstream communication (for example to toggle a plug-meter outlet on or off) is sent at the end of the master's beacon packet and each client node can only reply upstream to the mas-

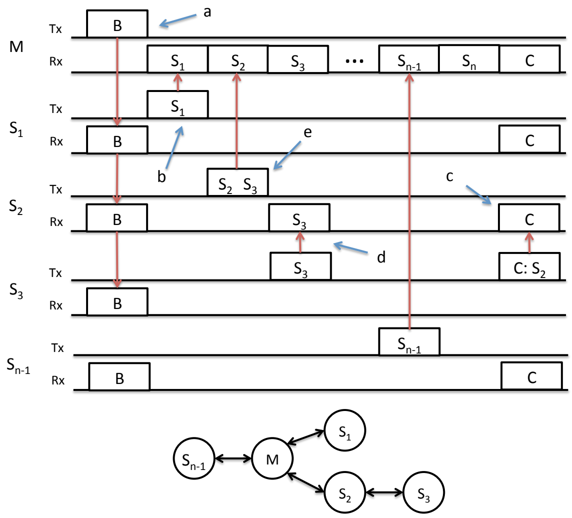

Figure 11: Operation of the mPCF protocol

ter (no peer-to-peer support). The master beacons are flooded across the network using the active construction principle found in the Glossy [18] time synchronization protocol. Each node rebroadcasts the downstream message after decrementing a TTL counter quickly enough that nodes experience constructive interference at the receivers rather than a collision. In Figure 11, point a shows the master's beacon message at the start of each cycle. Point b then shows a client S_1 sending data back to the master in the first slot. The client can then sleep until the start of the next cycle (the master node is assumed to be connected to a gateway and powered). In our implementation, the beacon node includes the number of slots and the size of each slot in the current cycle, to allow the master node to dynamically adjust the throughput on a cycle-by-cycle basis. In a single-hop network, this approach proves to be extremely effective in terms of both energy and throughput. In one of our deployments, we had 64 nodes reporting data every 640 ms with an average radio duty-cycle of 0.03% which is comparable to the idle mode of existing listening/probing [19, 20] protocols. Based on a 50ppm clock crystals, nodes can typically afford to miss 2 or 3 beacons before they become significantly out of sync to cause problems. Once a node does lose synchronization, it will keep its radio active until it receives a new beacon packet.

The 802.15.4 radio has a theoretical throughput of 250Kb/s. In our deployment, we have 64 10ms TDMA slots. The maximum data that can be transmitted per slot, excluding the packet overhead is 100 bytes, thus resulting in a goodput of 100Kb/s. As described in Section 6, we deployed 28 plug meters (51 bytes/packet) and 7 syntonistors (3 bytes/packet), resulting in an effective data rate of 22.64Kb/s.

The two most prominent drawbacks of this approach are the lack of multi-hop support and the requirement that all nodes have unique lower-byte MAC addresses. In order to add multi-hop support, we extended PCF to include a contention-based communication slot at the end of each TDMA cycle. If a node is unable to reliably communicate directly with the master node (downstream communication uses Glossy), it can request a neighbor to forward its messages during a contention slot. Each client node listens on a contention slot at the end of each TDMA cycle. The dwarfed client node can then communicate to its neighbor in the contention slot asking to become its child. Figure 11 point c shows this transaction between node S_3 and node S_2. If the assigned parent has the resources available, it will then begin to listen on the child's transmit slot as shown by point d. The parent node will then aggregate its data along with the child's data (shown by point e). The child node can then overhear its parent's transmission to acknowledge that its data is being forwarded (it listens for this packet to synchronize anyway). At this point, the child has been paired with a parent and will stop sending requests on the contention slot. Its important to note that downstream data from the master can utilize the Glossy primitive because it is a single packet being flooded across the entire network. This would not work for upstream communication.

This extension has two main limitations. First, a parent must share its transmit packet with all children below it which severely limits multi-hop scalability. Nodes could keep packet queues and not transmit new data on each cycle, but this coordination would have to be maintained by the application. Second, now nodes that are eligible for forwarding must listen on an additional contention slot, which consumes more energy. For deployments where relatively small data fragments are being transmitted (single packet per cycle) and the hop depth is reasonable (typical homes and offices often require just a few hops) this approach is simple, effective and enables high-speed data collection from tens or even hundreds of nodes.

6. SYSTEM EVALUATION

In order to evaluate the overall energy metering performance, we deployed our system in a single-family residential building, and collected data for 1 week. We installed our three-phase energy meter on the two 150 Amp mains lines that feed the house. We installed plug-level meters (28 total) on all accessible plug-load appliances. Finally, we installed 7 EMF sensors on the following appliances: LCD TV, Washing Machine, Toaster Oven, Window AC, Laser Printer, Refrigerator and Iron. Each appliance with an EMF sensor also had a plug-meter which could be used to measure ground-truth readings.

6.1 Event Detection Performance

First, we evaluate how accurately each EMF sensor was able to detect appliance transitions. Over the period of 1 week, we compared the EMF sensor's threshold-based event detector with a hand-tuned threshold selected for each plug-meter. When looking at the plug-meter data, we selected thresholds that represented each appliance in either an *on* or an *off* state. Figure 12 shows the confusion matrix for each appliance as well as the overall average across all appliances. The confusion matrix was generated by comparing the amount of time that the EMF sensor categorized an appliance in a state that agreed or disagreed with the ground-truth. This is a more telling metric then doing an event-by-event comparison since one poor event transition could potentially set an appliance in the wrong state for an extended period of time.

From the confusion matrix, we can observe that the system in general performed well in categorizing the appliance-state in agreement with the ground-truth. We also observe that the Iron, TV and Laser Printer show above average misclassification. In the case of the Iron, the EMF-node was accidentally moved away from the appliance around the fifth day in the week of testing, hence it failed to detect 6 of the 29 events. The EMF-detector assumed that every appliance takes only two states, but the TV takes three states - *on*, *off* and standby. When the TV was in the standby mode, the signal strength was not sufficient to be classified as an *on*-state by the EMF-detector, resulting in 3 standby events being classified as *off* events. Notice that all appliances except the Printer perform well while the appliance is *off*. The EMF-sensor corresponding to the Printer was physically close to a monitor and was detecting the signal due to both appliances. Hence the EMF-sensor detected an *on* state even though the Printer was *off*.

6.2 Appliance Metering Performance

We then evaluated how well the system worked at estimating the energy consumption of each appliance. We analyzed how the system performed in two steps. First, we computed the ground-truth energy consumption of each device as recorded by the plug-meters. Next, we ran our energy estimation algorithm based on main measurements and using hand-selected thresholds derived from the plug-meters. The idea was that this would show appliance-metering performance assuming perfect event detection. Finally, we let the system compute the energy for each appliance using the actual EMF sensor data. Figure 13 shows a stacked bar-graph plot of each scheme. In this example, we choose a before-event and post-event window size of 2 seconds for determining the change in mains power for the appliance. Figure 14 shows the error for the EMF sensor as we vary this window size. In certain cases, having too small a window leads to poor performance because the appliance has a slow *on* or *off* transient. For example, an appliance like a refrigerator draws an abnormally high amount of current for a second or two while starting up. If the window is too small, the system miscalculates the energy consumption by assuming the high transient-power for the entire *on* cycle of the appliance. On the other hand, even if we select a large window size, the system runs the risk of miscalculating the energy due to events that occur nearby in time. As seen in Figure 14, a window size of 2 seconds performed the best.

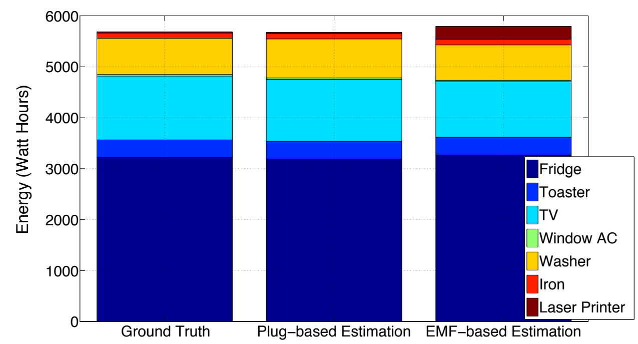

Figure 13: Energy estimation performance

Consistent with the confusion matrix, we observe that the EMF-based energy estimate of the TV and the Iron are lower than the ground truth, due to the failure to detect 20% and 15% of the events respectively. The best-case per appliance energy disaggregation was seen for the refrigerator, with an error of 1.3%. The worst-case disaggregation was seen for the laser printer, with an error of more than 1000%. As mentioned earlier, the corresponding EMF-sensor was also detecting a monitor placed nearby. Though the power-consumption of the monitor is much lower than the printer, the monitor was *on* for a much longer duration, resulting in a several-fold increase in the energy-estimation. The error of the overall energy consumption is 1.9%. It is to be noted that the refrigerator consumed 56.7% of the total energy, whereas the laser printer consumed only 0.4% of the total energy. As seen from Figure 13, the error is lower for appliances that consume more energy and higher for appliances that tend to contribute less to the house as a whole.

6.3 Energy and Latency Trade-off

Each time an appliance event occurs, the EMF sensor node wakes up the FireFly so that it can report the event back to the gateway. This means that the arrival rate of events has a significant impact on the lifetime of

Refrigerator (195)

EMF Sensor \ Ground Truth	On	Off
On	99.8	0.1
Off	0.2	99.9

Toaster Oven (5)

EMF Sensor \ Ground Truth	On	Off
On	99.7	0.08
Off	0.03	99.2

TV (20)

EMF Sensor \ Ground Truth	On	Off
On	92.6	0.08
Off	7.4	99.2

Window AC(2)

EMF Sensor \ Ground Truth	On	Off
On	99.1	0.00
Off	0.09	100

Washer (47)

EMF Sensor \ Ground Truth	On	Off
On	99.5	0.00
Off	0.05	100

Iron (29)

EMF Sensor \ Ground Truth	On	Off
On	83.1	0.50
Off	16.9	99.5

Printer (24)

EMF Sensor \ Ground Truth	On	Off
On	96.5	2.7
Off	3.5	97.3

Total (322)

EMF Sensor \ Ground Truth	On	Off
On	95.8	0.7
Off	4.2	99.3

Figure 12: EMF sensor event detection confusion matrices. Total event count in parentheses after appliance name.

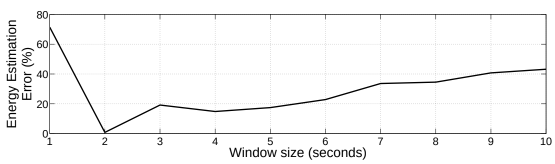

Figure 14: Error vs Mains sampling window

the node. The duty-cycle of the communication protocol also impacts lifetime at the cost of latency. For example, with a slow TDMA duty-cycle the energy requirement might be less, but the time it takes to report an event could much longer. Figure 15 shows the lifetime of our node under different expected event arrival rates while varying the maximum expected communication latency (related to the TDMA cycle size). These curves are based on AA battery while taking into account energy used by the EMF front-end, the FireFly node, radio communication and battery shelf-life leakage. We see that a system sized for 500ms latencies and an event inter-arrival time of 1 second, operates for 1.5 years.

Next, we evaluate how jitter in event arrival (due to potentially delayed packets) impacts the energy estimation. Figure 16 shows that as long as the events are reported within about 2 seconds of occurrence, we are able to maintain nearly the best possible performance for the system.

6.4 Limitations

There are three main limitations to this approach. First, a local event detector still has the challenge associated with determining which internal state transitions are significant. In our system, we were focused on signaling large state changes, but often appliances could have a sequence of small internal states or continuously variable consumption. In these cases, a different type of detection algorithm may need to be investigated; perhaps one that analyzes the signals in the frequency domain. The second limitation is that these devices can

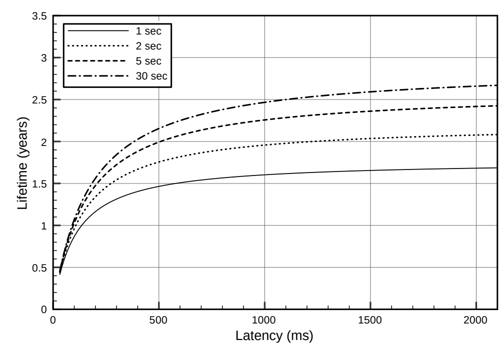

Figure 15: Estimated lifetime of node with respect to latency at various expected event detection periods.

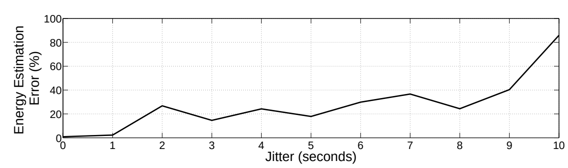

Figure 16: Error vs Event arrival jitter

suffer from cross-talk with different appliances if the devices or cabling are in close proximity of each other. We see this in our evaluation of certain devices present in our experiments. For example, the laser printer suffers from false positives. Part of optimizing the design is to build a device where the range is large enough to detect hard-to-reach wires, but small enough to minimize overhearing other signals. The third limitation is that the system cannot attribute static base-load values to appliances. In certain cases, the energy consumed while the appliance is supposedly *off* may be larger than its active energy. Despite these limitations, we believe that metering appliance-level energy consumption in a simple low-cost manner can still provide significant insight to end-users.

7. CONCLUSIONS

In conclusion, this paper presents the design of a new EMF sensor which can be networked with a mains energy meter to estimate per-appliance energy consumption. As compared with other energy monitoring solutions, the combination of a mains meter with EMF sensing is low-cost, easy-to-deploy and can be installed without disrupting the current operation of appliances. Since the sensors are proximity-based, they can detect current changes associated with the appliance from a few inches away making it possible to externally monitor in-wall wiring to devices like overhead lights. We use a secondary electric field sensor to help guide placement of the sensor as well as utilize the electric field's stability to help filter out potential noise. We show that these devices can operate from AA sized batteries for multiple years and when run in a realistic residential environment were able to detect events with an accuracy of 95.8% and estimate overall energy with an accuracy of 98.1%, as compared to ground-truth plug-meters.

8. ACKNOWLEDGMENTS

We would like to thank Patrick Lazik, Ge Yang and Emre Kara for their help. This work was partly funded by Samsung Electronics, and the National Science Foundation (NSF), grant #09-30868. The opinions expressed herein are those of the authors and not of the NSF.

9. REFERENCES

[1] Committee on America's Energy Future; National Academy of Sciences; National Academy of Engineering; National Research Council. *America's Energy Future:Technology and Transformation*. The National Academies Press, Washington, D.C., 2009.

[2] Sarah Darby. The effectiveness of feedback on energy consumption: A review of DEFRA of the literature on metering, billing and direct displays. Technical report, Environmental Change Institute, University of Oxford, Oxford, UK, 2006.

[3] World Business Council for Sustainable Development. Energy efficiency in buildings: Business realities and opportunities. Technical report, Switzerland, 2008.

[4] Lifton, J., Feldmeier, M., Ono, Y., Lewis, C., Paradiso, J. A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments. *International Conference on Information Processing in Sensor Networks (IPSN)*, April 2007.

[5] Jiang X., Ly M V., Taneja J., Dutta P., and Culler D. Experiences with a High-Fidelity Wireless Building Energy Auditing Network. *SenSys*, November 2009.

[6] Kim Y., Schmid T., Charbiwala Z. M., and Srivastava M. B. Viridiscope: Design and implementation of a fine grained power monitoring system for homes. *UbiComp*, September 2009.

[7] Schoofs A., Guerrieri A., Delaney D. T., O'Har G. and Ruzzelli A. G . Annot: Automatd electricity data annotation using wireless sensor networks. *IEEE Sensor Mesh and Ad Hoc Communications and Networks (SECON)*, 2010.

[8] Sandal Public Limited Company. Energenie. https://energenie4u.co.uk/.

[9] Efergy Technologies Ltd. Efergy. http://www.efergy.eu/, August 2012.

[10] P3 International Inc. Kill a watt. http://www.p3international.com/products.html.

[11] Tendril Inc. http://www.tendrilinc.com/.

[12] Greenbox Technology Inc. http://getgreenbox.com/.

[13] EnergyHub Inc. EnergyHub. http://www.energyhub.com/.

[14] Energy Optimizers Limited. Plogg. http://www.plogg.co.uk/.

[15] Buevich, M., Rowe, A., Rajkumar, R. Tracking and visualization of building energ. *1st International Workshop on Cyber-Physical Systems, Networks, and Applications (CPSNA 11)*, 2011.

[16] Rowe A., Berges M., Bhatia G., Goldman E., Rajkumar R., Soibelman L., Garrett J., Moura J. . Demonstrating Sensor Andrew: Large-Scale Campus-Wide Sensing and Actuation. *Demo Abstract, IPSN*, 2009.

[17] Anthony Rowe, Mario Berges, and Raj Rajkumar. Contactless sensing of appliance state transitions through variations in electromagnetic fields. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, BuildSys '10, pages 19–24, New York, NY, USA, 2010. ACM.

[18] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient network flooding and time synchronization with glossy. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, april 2011.

[19] Polastre J., Hill J. and Culler D. Versatile low power media access for wireless sensor networks. *SenSys*, November 2005.

[20] Prabal Dutta, Stephen Dawson-Haggerty, Yin Chen, Chieh-Jan Mike Liang, and Andreas Terzis. Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 1–14, New York, NY, USA, 2010. ACM.