Contactless Sensing of Appliance State Transitions Through Variations in Electromagnetic Fields

Anthony Rowe† Mario Berges‡ Raj Rajkumar†

† Electrical and Computer Engineering Department
‡ Civil and Environmental Engineering Department
Carnegie Mellon University
Pittsburgh, PA 15213

{agr,mberges,rajkumar}@andrew.cmu.edu

Abstract

Non-Intrusive Load Monitoring (NILM) is a promising technique for disaggregating per-appliance energy consumption in buildings from aggregate voltage/current measurements. One major limitation of the approach is that it typically requires a training phase during which users must manually label device transitions. In this paper, we present an inexpensive contactless electromagnetic field (EMF) event-detector that can detect appliance state changes within close proximity based on magnetic and electric field fluctuations. Each detector wirelessly transmits state changes to a circuit-panel energy meter, which can then be used to label and disambiguate appliance transitions detected from the aggregate signals as well as to track the associated energy consumption. Our EMF sensors are able to detect significant power state changes from a few inches away making it possible to externally monitor in-wall wiring to devices (e.g., overhead lights). We experimentally evaluate our proposed EMF sensor in terms of power consumption, accuracy and detection range on a variety of appliances to demonstrate its effectiveness towards augmenting NILM systems. We show that accurately detecting 100W loads from $10cm$ away is possible while maintaining multiple-year battery life from a coin-cell battery.

Categories and Subject Descriptors

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

General Terms

Measurement, Performance, Experimentation

Keywords

Load Tree Analysis, Non-Intrusive Load Monitoring

BuildSys 2010 November 2, 2010, Zurich, Switzerland.

1 Introduction

Understanding where energy is being used is an important first step towards energy conservation and efficiency. There are currently two main approaches for determining the breakdown of appliance electrical energy usage in a building. The first approach is to connect appliances directly to wired or wireless energy meters. This unfortunately can be expensive in terms of hardware, operation and installation costs. Moreover, many appliances like overhead lights are hard-wired and would require an electrician to install each individual sensor. A less invasive approach being explored by many smart meter manufacturers and researchers is the idea of Non-Intrusive Load Monitoring (NILM). The most common approach to NILM uses a single energy meter installed at the circuit-panel or other strategic points in the distribution system to carefully analyze voltage and/or current transients generated when appliances change their state and identify the appliance that caused it. This allows the NILM system to estimate, solely from measurements made at the distribution panel, what types of appliances are currently running and how much power they consume. The drawback of this approach is that although many signatures are common across device types, there will inevitably be certain appliances that require manual site-specific training. NILM performance can also suffer when many appliances are used concurrently making it more likely to encounter temporally overlapping transients.

In this paper, we present a contactless electromagnetic field (EMF) sensor that can detect appliance power consumption state changes within close proximity, based on magnetic and electric field fluctuations. We recommend that this type of sensing be used to aid in the training of NILM systems. A FET-based electric field sensor in conjunction with a coil-based magnetic field sensor is used to robustly identify changes in appliance power consumption. Each detector wirelessly transmits state change information about a local appliance to a main circuit-panel energy meter. A NILM system then uses these external events, along with information about what appliance the field sensor relates to, in order to label and disambiguate the events detected from the aggregate signals. As compared with other energy monitoring solutions, such as pluggable or inline 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. Our proposed event detection sensor is compact and consumes on average $45\mu W$ making it ideal for long-term battery operation. In permanent installations, these sensors can provide continuous feedback for a NILM system to adjust or re-train appliance signatures as devices change over time or if new devices are added.

2 Related Work

Multiple research projects have investigated improving the observability of energy. The MIT Plug [1] provided users with power and sensor information by means of a smart surge protector. In [2], the author's present experiences using the ACme wireless plug sensor in an office environment. In [3], the author's 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 only detecting appliance state changes rather than trying to directly measure power. This information, if time-synchronized with panel meter data, can be used by a NILM system 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 [4]. One of the primary focuses of this work is on intelligent local processing at the sensor which would benefit these other systems.

Commercial efforts are underway to add sensing devices in homes so as to provide users with energy-usage feedback. Google PowerMeter [5] is a software package that interfaces with smart metering technology to display household energy usage. Companies like Tendril Inc., AlertMe, Trilliant and GreenWave Reality are taking a more proactive approach by offering monitoring devices that home owners can install themselves to monitor energy. Many of the concepts in this work will benefit all of these systems by providing lower-cost collection of richer data with added insight about how users consume energy. Companies like ArchRock and Sentilla have started using wireless sensor networking technology to monitor energy in commercial buildings and data centers. In general, these approaches either require end-device metering or focus on circuit-panel level granularity.

Much work has addressed the problem of disaggregating electrical load using device signatures. In the early stages of the research area, steady-state changes in real and reactive power were used as signatures for appliance state-transitions [6]. Improvements to this work were presented in [7] by introducing a transient pattern matching approach to account for appliances with similar characteristics in the real and reactive power signature space. There have also been some efforts to commercialize the technology by companies like Enetics. In general these approaches require site-specific training and become less accurate in environments with a large number of switching appliances, continuously variable loads or very similar appliances. Of late, there have been a number of new research projects aimed at addressing some of these issues. Some of them also attempt to provide an automated training approach. Many of these proposed solutions can benefit by the EMF detector we present in this paper.

3 System Components

In this section, we discuss the various components of our system required to collect and correlate appliance on/off events with circuit-level power data. First, we designed a custom circuit-level meter with built-in communication that can be used as a gateway to collect event data. Next, we describe the EMF event detector hardware and firmware. Both the main circuit-panel meter and the EMF event detector are built around FireFly wireless sensor nodes that use an ATmega1281 micro-controller and the CC2420 802.15.4 radio. After describing the hardware components of the system, we discuss two possible event detection algorithms that run locally on the event detection sensor. It is important to note that all of our experiments were conducted in 120 VAC 60Hz signals.

3.1 Circuit-Panel Meter

We designed a custom three-phase power meter, shown in Figure 1, specifically to perform NILM operations as well as collect data 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 current and voltage on each phase at 5KHz and then computes true power and energy in software. We use this approach to ensure access to cycle and sub-cycle power data that is required for transient-based NILM. The tight coupling between the power sampling and the radio interface reduces timestamp error between signals sent from the appliance state detectors and the power values 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 clip-on style current transformers. We bit-extend the current readings using two ADC channels for each signal (one of which is amplified) to provide 12-bit current and 10-bit voltage resolution. Overall range and accuracy values depend on the particular configuration of the current transformer used.

3.2 EMF Event Detector Hardware

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

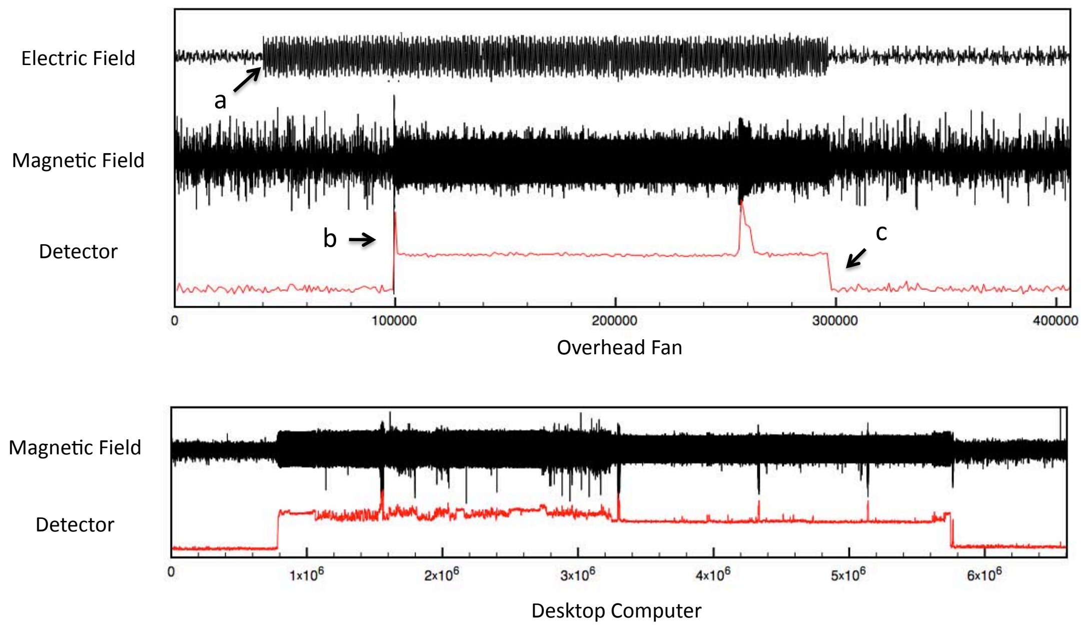

Figure 2. EMF Event Detector Waveforms. Top: Ceiling fan with light switch activated at point (a), manually turned on at point (b), and turned off at point (c). Bottom: Desktop computer is an example of a noisy signal due to switching.

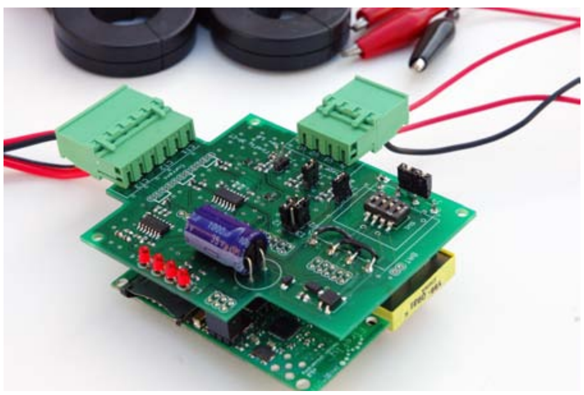

Figure 1. Wireless Three-Phase Circuit-Panel Meter

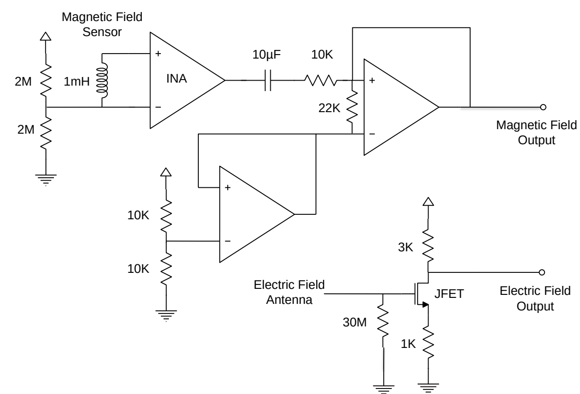

Figure 3. EMF Event Detector Circuit

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 how much power it is consuming (e.g. for example 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 3 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 2 shows two example waveforms received by the circuit when placed near a ceiling fan and a desktop com-

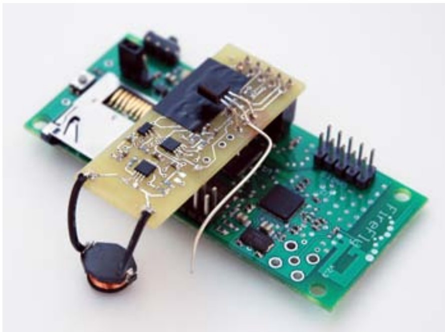

Figure 4. EMF event detector connected to a FireFly sensor node.

puter. Point (a) in the ceiling fan waveform denotes when the wall switch is turned 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 of these cases the edges in the RMS signal are quite pronounced.

Figure 4 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 in Section 4. 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.

3.3 Event Detection Firmware

We consider events to be instants in time when an appliance changes its state (e.g., goes from *off* to *on*). Figure 2, shown earlier, illustrates two such events, *b* and *c*, as captured by the EMF detector in the magnetic field. Although the electric field can provide better context about the nature of the signals, it does not change with the operation of the appliance. The magnetic field is then used for detecting appliance events. Given that the raw signal is periodic, instead of working with it directly, we use the RMS value (H_{RMS}) to detect significant changes. Through experiments described in the section that follows, we determined that a fixed-size window of T samples, where T is a positive integer multiple of the period of these signals, provides the best results in terms of the resulting variance.

We experimented with two different event detectors. The first is a simple threshold detector based on the RMS: $H_{RMS} > t$. The threshold value t was determined during a short calibration period at the moment of installation. After collecting a few seconds of data, t is set to $1/2(max(H_{RMS}) - min(H_{RMS}))$. We will refer to this approach as the threshold detection scheme.

The second is a probabilistic approach based on the Generalize Likelihood Ratio (GLR) test, which makes a Gaussian assumption about the distribution of the signal. We chose this algorithm based on work that has demonstrated its success in detecting HVAC and ligthing events on power signals [8]. Equation 1 is applied to a fixed-size sliding window over H_{RMS}, to detect the sample indices where there is a significant change in mean. This detection window is of size $k-l+1$. The equation requires that σ, the expected standard deviation of the signal, be specified a priori. Just as with the simple thresholding scheme, a few seconds of data were collected after placing the EMF detector at each location to estimate this value.

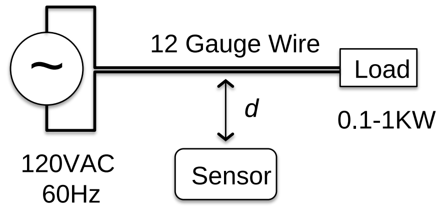

Figure 5. Experimental Setup

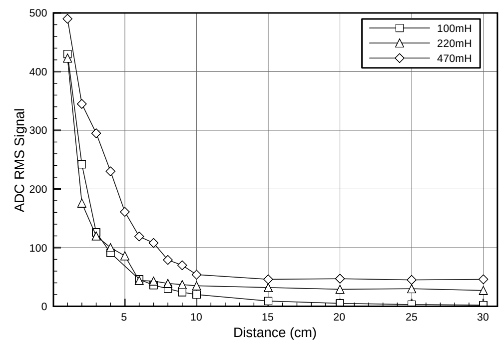

Figure 6. Measured RMS field strength vs distance for various coil sizes

$$e_k = \frac{1}{2\sigma^2} \max_{l \le j \le k} \frac{1}{k-j+1} \left(\sum_{i=j}^{k} (H_{RMS}[i] - \mu_0) \right)^2 \quad (1)$$

The threshold detection scheme has the advantage of being very simple to implement and potentially achieving good results on two-state appliances given that an appropriate calibration period is given. Conversely, the GLR approach is better suited for detecting events of multi-state appliances but requires more processing power. Also, the threshold detection scheme requires a full *on* / *off* transition to occur before it can accurately detect events. In contrast, the GLR approach only requires an expected variance value which can be captured by sampling the signal for a short period of time.

4 Evaluation

In this section, we evaluate the performance of both the magnetic and electric field sensing circuits and discuss the accuracy and energy requirements of the EMF detection node. Figure 5 shows a diagram of our experimental setup intended to test the receiver's sensitivity with respect to various AC loads and at particular distances from a long thin

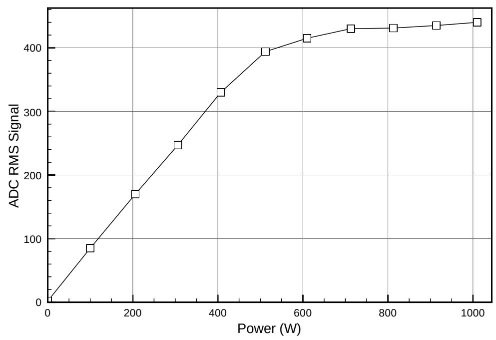

Figure 7. Measured RMS field strength vs load power for 100mH coil at 1cm.

wire. The intent was to simulate the scenario where the sensor is detecting energy either from in-wall wiring or directly from appliance cables. The appliance test load was generated from purely resistive elements.

In order to keep the design both low-cost and compact, we first evaluated the impact of different inductor coil sizes on the magnetic field detectors range. Figure 6 shows how the receiver performs with three different coil sizes given a fixed 1KW load as distance from the wire is varied. Values on the y-axis represent the RMS value of the ADC samples captured by the node's MCU. We see that larger inductors are able to capture smaller signals; however, proximity to the wire plays a more pronounced role in signal intensity. In this case, the $100mH$ inductors is almost one-tenth the size of the $470mH$ one and can compensate for the freespace losses in energy by simply moving slightly closer. We also see that wth a distance of $10cm$, the smaller inductor is still able to detect the load.

In the next experiment, we placed the magnetic sensor with the $100mH$ inductor $5cm$ away from the wire while adjusting the test load. In Figure 7, we see a linear change in signal strength until the sensor begins to saturate at about 500 Watts. Though there is a correlation between current and signal level, as we saw in the previous figure, this is heavily dependent on the distance from the wire (as well as other geometry) and hence would not be accurate for estimating power without calibration.

Next, we evaluate the sensitivity of the electric field detector. In Figure 8, we varied the length of the receiving antenna and plot the corresponding peak voltage that is output by the JFET. We see a relatively linear response, which is expected given a near-field detector operating at a miniscule fraction of the 60Hz wavelength. Figure 9 shows the performance of a 5cm antenna with respect to distance from the wire. Note that since it detects the electric field, the load has no impact on these values.

We deployed the EMF sensor near 8 different types of appliances and collected the raw ADC waveforms for analysis. Since often times the high-frequency switching noise generated by appliances is useful information, we sampled the ADC as quickly as practically possible (in this case 15KHz). We then evaluated the number of required samples in order to ascertain a good estimate of RMS signal strength by varying the window size and computing the average standard devia-

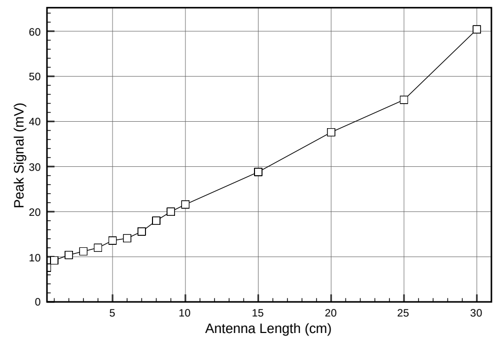

Figure 8. FET antenna length vs distance fixed at 10cm from cable.

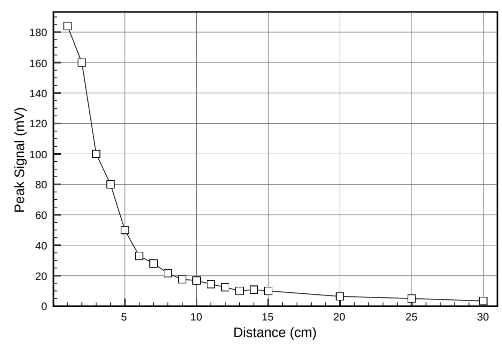

Figure 9. FET peak signal strength vs distance for 5cm antenna.

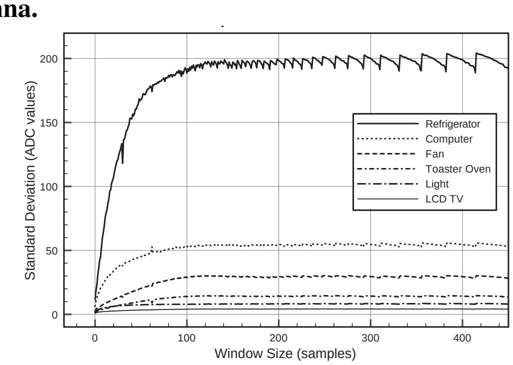

Figure 10. Average standard deviation of the RMS windows during the first 5000 samples of the signal, at various window sizes

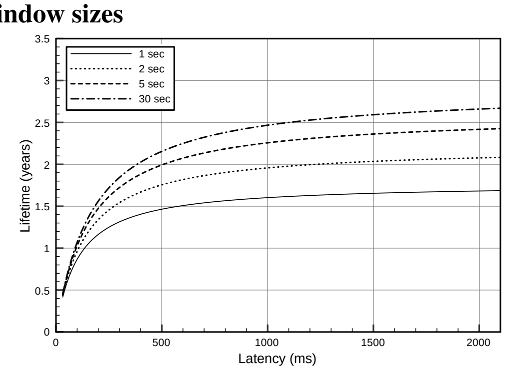

Figure 11. Estimated lifetime of node with respect to latency at various expected event detection periods.

tion of these windows. Figure 10 shows the average standard deviation of the windows used to compute the RMS values for a few representative appliances (other appliances were very similar) over an increasing window size. Initially, the variance is low since the signal does not have time to change, but eventually we see that it starts oscillating around a fixed value, and that these oscillations may be related to the integer multiples of the period T.

From the 8 appliances studied, we captured more than 40 labeled transitions. These appliances included a refrigerator, overhead light, fluorescent light, desktop computer, toaster oven, fan, TV and air conditioner. Of the 40 transitions, all were captured correctly with one false positive due to an unknown noise source, shown in the upper graph of Figure 2, over a total period of 1 hour. In the event of false positives, the NILM system can filter them out if there is no corresponding change in mains power.

By duty-cycling how often the sensor node wakes up to capture a window of data, we can significantly reduce energy at the cost of detection latency. The other major component of power consumption is the radio required to relay information back to the gateway. We assume a model where the radio is only required to wakeup and transmit when an event is detected and it does not have to listen to the channel to route packets hence the expected appliance detection period has a significant impact on the node's lifetime. In Figure 11, we show the impact of latency on node lifetime given various different expected detection periods.

These curves are based on a 650mAh lithium coin-cell battery with the EMF detector being sample by the ATmega1281 and data being transmitted using the CC2420 radio. We also included the leakge current within the battery derived from a 5 year shelf-life. We see that even with expected appliance transitions multiple times per minute and latencies around 1 second, we can expect a multiple-year battery life. Testing conditions for a 1000*ms* latency and 1 packet transmission per second, we measured the actual detector energy consumption to be 0.06mA which is within 30% of the predicted value.

5 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 this 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 suffer from cross-talk with different appliances if the devices or cabling are in close proximity of each other. 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 final limitation to this system is that it requires additional hardware so as to increase the accuracy of a system that could theoretically operate on its own. We believe that until NILM systems get to the point where they no longer require calibration that some form of local event detection will be required.

6 Conclusions

In conclusion, this paper presents a contactless EMF sensing device that can be used to identify and wirelessly relay electrical appliance state transitions. We propose that such a device be used to augment Non-Intrusive Load Monitoring systems allowing them to automatically label, train or tune appliance classifiers at runtime. As compared with other energy monitoring solutions, the combination of NILM along 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 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. We use a secondary electric field sensor to help guide placement of the sensor as well as utilize the electric fields stability to help filter out potential noise. We show that these devices can operate from a coin-cell sized battery for up to 2.5 years depending on desired latency and expected appliance transitions. Further tests need to be conducted to evaluate the effectiveness of the devices in detecting state transitions for more complex cases (e.g., HVAC systems), although the GLR detection algorithm was chosen due to its reliability in these scenarios.

7 References

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] Google PowerMeter: http://www.google.org/powermeter (viewed 2/22/2009).

[6] Hart, G. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, December 1992.

[7] Norford, L.K. and Leeb, S.B. Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms. *Energy and Buildings*, May 1996.

[8] Dong Luo, Leslie Norford, Steven Shaw, and Steven Leeb. Monitoring HVAC equipment electrical loads from a centralized location - methods and field test results. *ASHRAE Transactions*, 108(1):841–857, 2002.