Demo Abstract: Disaggregating End Loads with Energy-Harvesting Sensors and Cloud Analytics

Bradford Campbell, Samuel DeBruin, Meghan Clark, and Prabal Dutta
Electrical Engineering and Computer Science Department
University of Michigan
Ann Arbor, MI 48109
{bradjc,sdebruin,mclarkk,prabal}@umich.edu

ABSTRACT

Obtaining a detailed breakdown of household energy consumption would allow occupants to better understand their energy usage patterns and identify opportunities for energy savings. Current solutions are too course-grained, too difficult to deploy, not networked, or offer poor coverage of hard to meter items, such as ceiling lights. To address these problems, we demonstrate a wirelessly networked, energy-harvesting power metering system that draws zero standby power and is power proportional to the load it is metering. The system is comprised of three different meters: one for plugged-in loads, one for panel-level circuits, and one for hard-to-sense loads, such as ceiling lights. Each meter harvests energy proportionally to the load it is measuring and powers a sensor node intermittently. Together, these sensors create multiple data streams which are aggregated by a receiver. When combined with a calibrated meter that measures total household power, our system can iteratively determine the contributions of each load to the total power usage, allowing users to gain a broad yet detailed view of their energy consumption and costs.

Categories and Subject Descriptors

B.m [**HARDWARE**]: Miscellaneous—*Miscellaneous*

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Energy-harvesting, Power metering, Data aggregation

1. INTRODUCTION

Fine-grained, whole-home energy metering is challenging but important: U.S. residential homes consume 38% of all electricity and 22% of the primary energy in the United States [1]. Decomposing the monthly power bill into a detailed, localized report would enable consumers to better understand their usage and reduce waste [4].

BuildSys'13, Nov 13–November 14 2013, Roma, Italy
ACM 978-1-4503-2431-1/13/11.
http://dx.doi.org/10.1145/2528282.2534160

However, current home energy meters target very specific goals and applications, making it difficult for consumers to gain a detailed breakdown of their entire energy usage. Current consumer power meters fall into three general categories: a household meter (as used by the electric company for billing), circuit breaker panel meters that measure the load of an entire circuit, and plug-load meters that can meter a single device. These devices span the spectrum of deployment simplicity and provided level of metering granularity. Commercially available systems typically target one of the three categories, and therefore only support one point on the ease-of-use versus detail tradeoff. Also, the sensors used to meter electricity are actively and constantly powered, meaning they require expensive AC-DC power supplies and draw their own power even when the load they are metering is off.

In contrast, we propose a system that is easy to deploy, offers fine-grained and automated wireless data collection, and uses inexpensive energy-harvesting sensors which draw zero power when the measured loads are turned off. By deploying these simple sensors at the circuit level, the plug-load level, and at additional points that indirectly measure electricity, we hypothesize that we can disambiguate the relative energy cost of the loads in a home by combining and processing the resulting data streams. We also claim that this system can measure loads that are otherwise difficult to instrument, like overhead lighting, which do not have exposed plugs.

2. POWER PROPORTIONAL SENSORS

To enable power proportional energy meters, we use a new type of sensor that does not directly measure the power of an AC load, as a shunt resistor would, but rather uses a side channel to indirectly measure the load using energy-harvesting principles. The fundamental idea is that the sensor harvests energy at a rate proportional to the load being measured, and once the sensor has accrued sufficient energy, it wakes up and increments a local counter. On certain wakeups the node also transmits a packet containing the counter value. The receiver of these packets can then determine the energy usage of the load based on the interval between packets and the counter values embedded in the packets. For instance, Figure 1 shows how the activation frequency changes with load power for a plug-load meter. This monotonically increasing function enables the receiver to determine consumption based on the wakeup rate.

For the plug load meter, we use a current transformer with one wire of the load's AC supply line wrapped around the coil [2]. When the load is active, the primary's magnetic field induces an AC current that is harvested to eventually power the node core, increment the counter, and send a packet. All of the electronics are housed in an integrated enclosure into which the load can be plugged, as shown in Figure 2(a). When the load is off, no magnetic field exists, and the meter draws zero power.

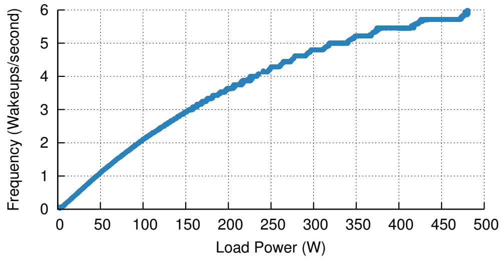

Figure 1: Frequency of wakeups over a range of load wattages for a plug-load meter. Based on the rate of packets from a plug-load sensor the packet receiver can determine the power of the load.

To measure electricity at the circuit panel level, we use a split-core current transformer that clips around the phase wire, as shown in Figure 2(b). This sensor harvests energy in an identical manner to the plug-load version. The required electronics are small and can be easily installed inside of the panel, making installation simple.

This sensing principle applies to other sources of harvestable energy as well, including indoor lights. To measure lighting, we use a solar panel based energy-harvesting power supply [5] as shown in Figure 2(c). By attaching these sensors near ceiling lights, we can detect when lights are on and their relative brightness. As long as the power draw of the light increases with brightness, we can indirectly measure power through activation rate.

3. NETWORKING

This sensing scheme employs wireless transmissions. However, due to the unpredictable nature of sensor activations, the system currently requires an always-on receiver. Typically, a border router in a wireless sensor network is mains-powered, so we leverage this network element to timestamp and forward the packets.

It is desirable for a single border router to be able to handle packets from not only the sensors mentioned but also from future wakeup-based sensors, so we design our network to require only a single, stateless border router. This allows new sensors to be added without reprogramming the border router. We accomplish this by running an IPv6/6LoWPAN stack on each sensor and on the border router. Each energy-harvesting sensor sends a valid 6LoWPAN packet containing its counter value. The border router simply expands the 6LoWPAN packet and routes it to the Internet. Using IPv6 allows us to transparently use any always-on mesh nodes present in the network as they are also able to interpret and forward the packets.

The border router hardware we use is a custom TI CC2520-based 802.15.4 radio shield for the Raspberry Pi mini-computer, shown in Figure 2(d). This platform is inexpensive, reliable, widely available, and runs Linux, which provides many networking and IPv6 primitives. We ported TinyOS to run on the RPi and wrote a CC2520 kernel driver to leverage the existing 6LoWPAN stack.

4. DISAMBIGUATING LOADS

Aggregate load measurements, which represent the sum of all the active loads within a building or within a building subsystem, are relatively accessible and provide high coverage compared to individual loads. Aggregate measurements can take the form of utility meter readings or readings from monitors placed at key locations along the "load tree" structure that frequently characterizes power flow within buildings. However, the resolution of aggregate measurements is too coarse-grained to provide deep insights into the nature of power draw within buildings.

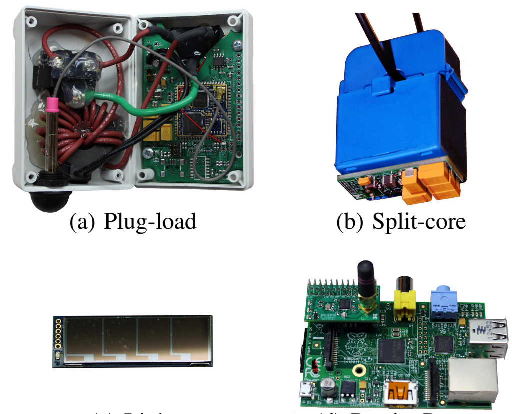

(a) Plug-load (b) Split-core

(c) Light (d) Border Router

Figure 2: Energy-harvesting sensors and 6LoWPAN border router.

Augmenting aggregate measurements with non-invasive sensors allows us to decompose an aggregate load into the individual loads that compose it. By analyzing how sensor wakeup rates change as the aggregate changes, we dynamically discover a system model that enables us to determine the actual load profiles.

We begin by modeling the power draw of each load as a function of the sensor wakeup rate. Following the work of Kim et al., we call this the *calibration function* [3] and assume that the dominant characteristics of the calibration functions of each sensor can be approximated by a strictly monotonically increasing polynomial of degree three or less. We then wish to discover the coefficients for each sensor that minimize the error between the sum of our predicted loads and the actual sum of the loads given by the aggregate load measurement. The aggregate measurements and wakeup rates that are reported to our disaggregation server over time are constraints on the coefficient values. Discovering the coefficients is a linear optimization problem, which can be solved iteratively until the prediction error is sufficiently small.

Abstracting individual loads as wakeup frequencies allows our system to be adaptable and allows us to support interoperability. Any combination of heterogeneous or third-party energy-harvesting sensors can be used, so long as the sensors report wakeup counts and satisfy the reasonable constraint that their calibration function can be approximated by a strictly monotonically increasing polynomial of degree three or less. This system can easily incorporate new sensors on the fly, whenever a new sensor is detected. The user receives near real-time load predictions and accurate historical data while installation and setup costs are kept low.

5. REFERENCES

[1] Dept. of Energy buildings energy data book (2012). `http://buildingsdatabook.eren.doe.gov/`.

[2] S. DeBruin, B. Campbell, and P. Dutta. Monjolo: an energy-harvesting energy meter architecture. In *Sensys'13*, 2013.

[3] Y. Kim, T. Schmid, Z. M. Charbiwala, and M. B. Srivastava. Viridiscope: design and implementation of a fine grained power monitoring system for homes. In *UbiComp'11*, 2009.

[4] A. H. Mcmakin, E. L. Malone, and R. E. Lundgren. Motivating residents to conserve energy without financial incentives abstract:, 2002.

[5] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta. Grafting energy-harvesting leaves onto the sensornet tree. In *IPSN '12*, 2012.