

Meter Any Wire, Anywhere by Virtualizing the Voltage Channel

Thomas Schmid[†], David Culler[‡], Prabal Dutta[†]

[†]Computer Science & Engineering Division
University of Michigan
Ann Arbor, MI 48109
{thschmid, prabal}@eecs.umich.edu

[‡]Computer Science Division
University of California, Berkeley
Berkeley, CA 94720
culler@cs.berkeley.edu

Abstract

AC power meters require both voltage and current to be sampled concurrently to obtain real, reactive, and apparent power. Typically, the two measurements are taken in close physical proximity and fed into a single power metering device. In this paper, we explore the viability of decoupling the voltage and current channels, and placing them in physically disparate locations. Such decoupling could ease the installation of metering infrastructure and enable new sensing scenarios. However, decoupling the voltage and current channels raises a new question: how should they be recombined? Of the various approaches, we propose the voltage channel be virtualized: a voltage sensor measures the voltage magnitude, frequency, and phase, typically near the root of a circuit branch. The extracted phase is time-stamped relative to a global clock and disseminated wirelessly, along with the magnitude and frequency measurements, to power meters throughout the network. The power meters synthesize a suitably scaled replica of the voltage waveform locally, based on the parameters reported by the voltage sensor, and combine it with locally-measured current readings. This paper demonstrates – through empirical characterization of the line voltages, a proof-of-concept power meter implementation, and house-scale evaluation – that the design holds promise and offers substantially lower measurement errors than other distributed power metering approaches for non-resistive loads.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: General

General Terms

Design, Measurement, Performance

Keywords

Energy metering, power factor measurement, time synchronization, wireless sensing

1 Introduction

Today, in-building power measurements are usually centralized (e.g. in breaker box) or distributed (e.g. at a plug load). The former provides *full* coverage over the aggregate load while the latter provides *detailed* coverage of only instrumented loads [8, 11]. In this paper, we argue that these two design points are too restrictive, and that a richer set of emerging power metering options [9, 13] would be greatly aided if only the current and voltage channels could be decoupled and placed individually at physically distinct locations. However, naively placing current and voltage sensors throughout a house is not practical today because these signals must be recombined to obtain power measurements.

For example, allocating line losses to the responsible loads in a house requires measuring voltage at the service entry point but current at the load (measuring both current and voltage at the load fails to account for wiring losses in the house). As another example, measuring whole house current is easy if measured outdoors, using a split-core current transformer attached to the service drip loop, but voltage is difficult to access outdoors. Conversely, voltage is easy to measure indoors, but access to whole house current may be difficult since feeds are often encased in conduits or hidden behind circuit breaker panels. A power meter that allows the current and voltage channels to be decoupled and virtualized enables these and many other measurement scenarios.

To explore these new measurement scenarios, we present a distributed power meter that decouples the current and voltage channels, allowing each to be measured in the most expedient manner for a particular application, and recombines them using a low-rate wireless channel. Our design measures the magnitude, frequency, and phase of the line voltage using a *voltage sensor* typically placed near the the root of the circuit subtree (e.g. near a building's service meter). The voltage sensor disseminates the voltage parameters over a wireless network to one or more *power meters* (e.g. a drip loop meter or plug load meter). The power meters locally synthesize a suitably scaled replica of the voltage waveform, based on the parameters reported by the voltage sensor, and combine the synthesized voltage with current measurements from a locally-connected *current sensor* to compute real, reactive, and apparent power. Among its many benefits, this design enables novel sensing scenarios, simplifies installation, and can proportionally allocate wiring losses to loads.

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'10, November 2, 2010, Zürich, Switzerland.

Copyright © 2010 ACM 978-1-4503-0458-0/10/11/02...\$10.00




Figure 1: (a) Service drip loop where current sensing is easy but voltage sensing is not; (b) Circuit model for line loss analysis; (c) MacBook voltage/current profile illustrating harmonic and displacement distortion in current. (d) System architecture for virtualizing the voltage channel and recombining it with current.

2 Example Application Scenarios

There are many applications of a distributed power meter. They include simplifying installation, allocating line losses to the responsible loads, measuring true power for hard-to-reach loads, and improving safety by eliminating the need for inline measurements. We detail two such scenarios.

2.1 Metering Whole House Power

One challenge with end-user-deployed, whole house power metering is accessing the current channel. Some approaches require professional installation due to the safety concerns of installing a current transformer (CT) inside a breaker panel [13]. In some buildings, however, the drip loop at the overhead power service entry provides a convenient current sense point, as Figure 1(a) shows. The drip loop may be the last point before the conductors enter the service mast conduit, where they remain inaccessible until emerging inside the breaker box. The voltage channel, however, is inaccessible outdoors since the conductors are insulated, while the converse may be true indoors. Distributed power metering addresses the problem by relaxing sensor placement.

2.2 Allocating Line Losses to Loads

Allocating line losses to the responsible loads is difficult because these losses depend on the wiring topology, wire gauges, load profiles, and attachment points. Absolute losses must be less than 7%, but the actual losses are rarely allocated to specific loads. We show how they could be.

Figure 1(b) shows a source supplying a voltage, V , and current, I , to two loads, R_a and R_b , over wires with resistance R_1 , R_2 , R_3 , and R_4 . Measuring I and V at the source, as whole house meters do, fails to disambiguate loads or allocate losses. Measuring current and voltage at the loads, as plug load meters do, fails to account for line losses. However, measuring current at the load (e.g. I_a or I_b) and voltage at the source (e.g. V) allows load disambiguation and per-load line loss attribution (e.g. R_a or R_b), as we next show.

Total circuit power dissipation, including load and line, is $P = IV$. Expanding the current term I into its constituent components gives $P = (I_a + I_b)V = I_aV + I_bV$. This clearly shows the power draw due to each load, including line losses. However, it requires a voltage measurement at the root of the load tree and current measurements at the loads – challenging with today’s meters, but feasible using a distributed power meter with decoupled current and voltage channels.

3 System Design and Implementation

Our distributed power meter design advocates separating the current and voltage sense channels, and installing them independently. The key challenge that arises is how to correlate and recombine the physically disparate sensor data streams. One approach would be to time-stamp the disparate current and voltage samples and stream them to a single point for combining into real, reactive, and apparent power. The main problem with this approach is that it requires streaming data at a high rate over a typically low rate wireless network.

Patel et al. propose a variant of this approach that streams non-timestamped current readings at a 1 kHz rate and computes an RMS value of the current and voltage samples before multiplication [13]. Unfortunately, multiplying RMS rather than instantaneous values works for resistive loads but introduces errors for switching or reactive loads that exhibit harmonic or displacement distortion.

Another possibility might be to parametrize the current or voltage waveforms, and wirelessly transmit only the extracted parameters. Figure 1(c) shows the voltage and current of a typical switching power supply; the current waveform exhibits significant harmonic and displacement distortion, implying considerable processing to parametrize. Note, however, that *the voltage waveform exhibits low total harmonic distortion, and is therefore well-suited to compact parametrization*. We use this technique – virtualizing voltage – to implement a practical distributed power meter.

Figure 1(d) shows our prototype implementation. The load current and voltage are monitored using a CT and differential probes, respectively, and captured using a LeCroy WaveRunner oscilloscope. An ACme meter [7] is plugged in near the root of the load tree, and is modified to export its zero-crossing (ZC) signal. The ZC signal is fed into a battery-powered Epic [6] mounted in a breakout board. Collectively, the ACme+Epic is the *voltage sensor*. A second Epic mounted in a development board is wirelessly time-synchronized [12] with the voltage sensor, establishing a common timebase. The voltage sensor extracts the voltage phase (relative to the timebase) and frequency, and periodically transmits these parameters to the second Epic, which uses them to synthesize a voltage waveform that is captured by the scope. In our implementation, the scope data are post-processed for evaluation, but an integrated design would perform the functions on-line and in real-time.




Figure 2: Synchronized voltage measurement at four different outlets across a single-phase one-family home. Figure (a) shows that two measurement points (C1, C2) are on one power leg while the other two (C3, C4) are on the second power leg. The phase difference of 180° comes from center tapping the 240 V phase-to-phase line in the utility transformer. With only a light total load, the voltages on the two phases are almost identical. Figure (b) shows the voltages when a 1.4 kWh space heater is placed next to measurement point C1. We observe a voltage depression at this point due to line losses. However, if a large load (air conditioner of 2.4 kWh) is turned on, the measured voltages are not impacted, as Figure (c) shows, since the air conditioner is on its own circuit branch.

4 Evaluation

We now explore the viability of our distributed power meter design. Specifically, we explore the foundational assumptions that the voltage waveform is sufficiently sinusoidal, that the voltage waveform can be compactly represented, that the voltage waveform can be synthesized accurately, that tight time synchronization can be maintained for extended periods of time, and that changes in the line frequency are small relative to the underlying time synchronization stability. Finally, we explore if all of these assumptions hold, then whether an accurate, distributed AC power meter can be implemented.

4.1 Characterizing Voltage Waveforms

The proposed approach depends critically on a sinusoidal AC voltage waveform. We first explore how consistent the voltage waveform is across several loads distributed through a single-family home fed by a single power phase that is center tapped at the utility transformer to provide two supply legs, 180° out of phase, nominally at 120 V.

We measure the voltage waveform across four loads using four digital storage oscilloscopes (Rigol model DS1052E) that are jointly triggered via their external trigger input. The voltage is sampled at 5 MS/s, and the scope data are collected using the scopes' USB ports. We conduct the measurements under different load scenarios to explore the effects of loading on magnitude, frequency, phase, and line loss.

Figure 2 shows the measurement results for three different scenarios. Figure 2(a) is the baseline when the total household power draw is 0.5 kW. This load includes lights, a refrigerator, computers, and the scopes. Two of the measurement points (C1, C2) are on one power leg, while the other two (C3, C4) are on the second. The RMS voltage between neutral and each power leg is ~ 122 V, while the RMS voltage between the two power legs is ~ 245 V. To measure the similarity of the sine waves, we calculate the six pairwise RMS errors by subtracting the waveforms from each other. With a total load of 0.5 kW, the average RMS error between the measured curves is 6.9 V, suggesting that load-induced line losses are responsible for the difference.

To quantify the effect of line losses, we connect a 1.4 kW load (a space heater) next to measurement point C1. Figure 2(b) shows the resulting voltage at all four measurement points. The RMS voltage at C1 drops to 116 V, or about 6 V compared to the unloaded case (a 5% loss), while the other RMS voltages remain at the previously measured levels.

In the final scenario, we explore the degree to which the line voltages across a house are affected when a large load is switched on. We add a 2.4 kW load to the house by turning on an air conditioner. This additional load does not affect the measured RMS voltage at the four disparate points, C1 through C4, nor does it introduce a phase change between two power legs, as Figure 2(c) shows.

These data suggest that the air conditioner load is parallel to our instrumented circuits and that there is no appreciable loss in the service drop conductors delivering power from the utility pole to the house. If there had been losses in the service drop, then the voltage of the instrumented circuits would have shown a depression, as can be observed in the scenario with the 1.4 kW space heater. These data confirm that in-building line losses are non-negligible.

We next explore the distortion in the voltage waveforms. This metric is important because the voltage magnitude, phase, and frequency are estimated by the voltage sensor and disseminated to the power meters. If the voltage waveform cannot be accurately parametrized using these three values, then the waveform synthesis will introduce an error, and power (and power factor) estimation will suffer. To evaluate distortion, we compute the RMS value of each voltage cycle and use this to synthesize one cycle of a sine wave whose amplitude is $\sqrt{2}$ times greater. We then subtract the synthesized waveform from the empirical one, and compute the RMS error. In a second run, we use 60 voltage cycles to compute the voltage RMS and compare this to 60 cycles of a synthesized waveform. Figure 3 shows both error distributions. The results show the voltage waveform exhibits low distortion and can be parametrized by magnitude, frequency, and phase with less than 4.5 V (3.75%) RMS error.




Figure 3: RMS error distribution between the AC line voltage and a synthesized waveform. We use either 1 or 60 cycles to estimate the RMS of the line voltage for synthesizing a waveform with the same amplitude and phase. While the difference is minor, using 60 cycles for RMS estimation achieves a smaller RMS error than using per-cycle calculations because the phase noise is reduced by averaging.

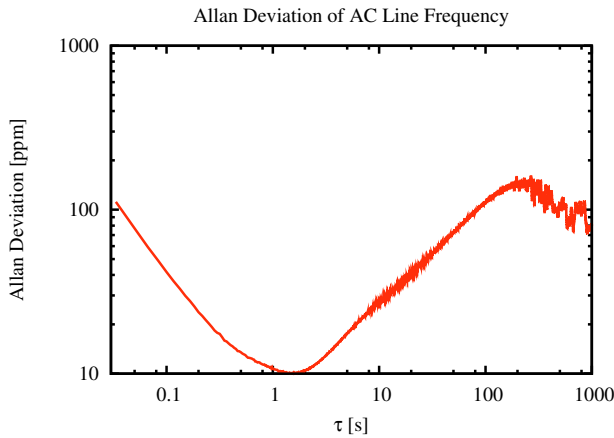


Figure 4: Allan deviation of the line frequency. The Allan deviation measures the degree of variation in a signal over a time window τ . We observe that the Allan deviation rapidly decreases initially, achieving a minimum at about one second, before increasing. Over very large time frames (hours to days) the frequency is controlled by the grid operators to achieve a long-term 60 Hz average, but for our purposes, averaging over one second suffices.

4.2 Synthesizing Voltage Waveforms

A power meter must be accurately time-synchronized and have the ability to synthesize a sinusoidal waveform based on the magnitude, frequency, and phase reported by a voltage sensor. The accuracy of the synthesized waveform depends on the time synchronization accuracy, as well as the measurement accuracy at the meter. We use a modified version of the Flooding Time Synchronization Protocol (FTSP) [12] with a 32,768 Hz timebase. Prior work has shown that by using modern radios, a regression history of four values, and a resynchronization period of 10 s, accuracies of $< \pm 1$ tick ($< \pm 30.5 \mu\text{s}$ for a 32,768 Hz clock) can be achieved. This figure translates to less than $\pm 0.2\%$ phase error.

While this establishes the required time synchronization interval, we still must determine how often the voltage sensor should disseminate the voltage parameters to power meters. To explore this question, we analyze the Allan deviation of line frequency. The Allan deviation is a measure of the frequency stability of a periodic signal over a time window τ .




Figure 5: Histogram of the phase error in the synthesized waveform. A voltage sensor measures the zero crossing times of the AC voltage and periodically sends these timestamps to a time-synchronized power meter that synthesizes a corresponding waveform of zero crossings. This histogram shows the phase errors resulting from the power meter's synthesis of the zero crossing signal based on a 32,768 Hz timebase.

Figure 4 shows the Allan deviation of the line frequency. We produce this figure by measuring the zero crossing rate of the line voltage using an 8 MHz oscillator. Since the measurement is performed in an air-conditioned room with minimal temperature variations, the oscillator frequency is assumed to be stable. The data show the minimum Allan deviation occurs over a 1 s window, balancing phase noise and long-term frequency variation contributions. Thus, for the remainder of our experiments, the voltage sensor distributes magnitude, frequency, and phase parameters of the voltage waveform every second. The parameters we report are the average value of the measurements observed over each second; the phase is measured relative to the global time provided by FTSP, which resynchronizes with a 10 s period.

An important measure of our system is the accuracy with which a node can synthesize the zero crossing times measured by a voltage sensor. To explore this question, we program a node to output the estimated zero crossings by toggling a GPIO line. To precisely toggle the pin, we use the timer capabilities of the TI MSP430 used in our system. The voltage sensor extracts and transmits the frequency and phase information of the zero crossings every second. Figure 5 shows a histogram of the time difference between the actual line zero crossings and the synthesized zero crossings. This difference is measured over 33,000 AC cycles using an oscilloscope. The data show a bimodal error distribution with one large peak near $5 \mu\text{s}$ and a second small peak near $16 \mu\text{s}$, the latter of which may be due to timer quantization.

We next use the synthesized zero crossings to trigger a DMA transfer of a precomputed sine wave to a DAC output. Using an all-hardware approach reduces the risk of missed interrupts and increases the timing accuracy of the synthesized signal. We use an oscilloscope to concurrently capture the AC line voltage and the synthesized DAC signal. Figure 6 shows the two waveforms, and their difference. Figure 7 shows a histogram of the errors. The synthesized waveform is scaled offline using the ratio between the RMS voltage of the synthesized waveform and the AC line signal.




Figure 6: Comparing the scaled, synthesized voltage waveform generated by a power meter and the actual AC line voltage. The power meter uses synchronization data supplied by the voltage sensor to generate a synchronized sinusoidal waveform using a DAC. The synthesized waveform is scaled offline using the RMS of both signals. The RMS of the error between the two signals is 6.8 V. The sharpness and periodicity of the error waveform suggests a third-order harmonic riding on a first-order phase offset error.

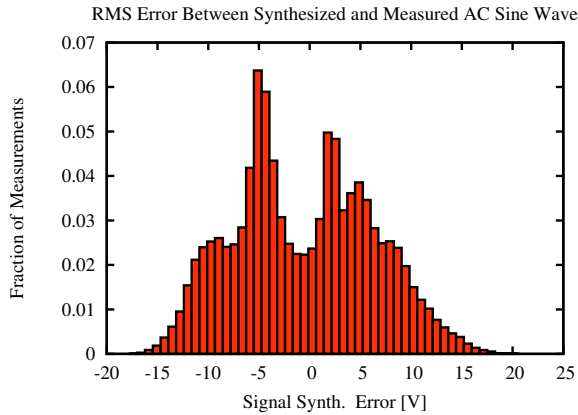


Figure 7: Voltage synthesis error distribution. The errors are bimodal with peaks near -5 and 5 V. The errors are due in part to time synchronization errors and in part to the presence of a third harmonic in the voltage signal. Although the errors are nearly symmetric with approximately zero-mean, they still may not cancel once multiplied with the current waveform.

These results suggest that a remote node can synthesize the voltage waveform to within 5-10% accuracy given a voltage sensor that transmits frequency, phase, and magnitude every second, and that a time synchronization process with a resynchronization interval of 10 seconds exists. We next investigate how these synthesis and timing errors impact power and power factor estimation.

4.3 Macroscale Evaluation

To measure the error introduced by using the synthesized waveform rather than empirical one, we conduct two experiments. We compare the true power ($P = IV$) and RMS power ($P_{rms} = I_{rms}V_{rms}$) by calculating them from traces obtained with an oscilloscope, and by sampling the synthesized signals. We connect the voltage sensor to an unloaded circuit in a different room and transmit phase and frequency every second. We use a current transformer to measure the current drawn of a 150 W light bulb and a 85 W MacBook. Concurrently, the oscilloscope measures the synthesized waveform generated from a DAC output.

Trace	True Power	Synth Power	Synth Error	RMS Power	RMS Error
MacBook 1	88.4 W	87.4 W	1.1%	94.4 W	6.7%
MacBook 2	88.3 W	87.2 W	1.2%	94.1 W	6.5%
MacBook 3	88.6 W	87.4 W	1.3%	94.5 W	6.7%
MacBook 4	88.4 W	87.3 W	1.2%	94.3 W	6.7%
Bulb 1	98.5 W	98.0 W	0.5%	99.1 W	0.6%
Bulb 2	98.5 W	97.7 W	0.8%	99.1 W	0.6%
Bulb 3	98.4 W	97.6 W	0.8%	99.0 W	0.6%
Bulb 4	98.3 W	97.7 W	0.6%	98.9 W	0.6%

Table 1: Comparison of measured and synthesized power calculation.

We collect several 2 s long traces sampled at 5 MS/s and calculate the true AC power from direct measurement, the true power by using the synthesized waveform, and the RMS power for both the light bulb and MacBook. Table 1 shows the results. Using the synthesized voltage waveform results in an average error of 1.2% in estimating the true power draw of the MacBook, and 0.7% average error for the light bulb. As the MacBook is a switching load, RMS power estimation overestimates the power drawn by more than 6.5%, showing that the RMS approach previously advocated [13] introduces substantial errors for reactive or switching loads.

4.4 Limitations

Our hardware relies on 32,768 Hz crystal as a time reference. Switching to a higher speed clock will increase synchronization accuracy, and thus reduce the phase error between the synthesized and measured voltage waveforms. Additionally, a better DAC triggering mechanism would reduce jitter in the start of the DMA transfer. With the current hardware, it is not possible to trigger a DMA transfer on a timer capture event and then switch over to a clock signal for the rest of the transfer. A more advanced DAC and DMA pair, potentially even specialized hardware, could offer such support and reduce the RMS error significantly.

We currently do not transmit RMS voltage measurements from the voltage sensor to the power meter. The scaling of the synthesized sine wave is currently performed offline, but we sketch how the scaling could be performed online. Instead of changing the sine wave output on the DAC0, we could use a second output (DAC1) as reference voltage to the first DAC0. Thus, DAC0 is only concerned with generating a sine wave at the right frequency, while DAC1 scales the amplitude of that signal. This method significantly reduces the processing load on the microcontroller and obviates the need to rescale the sine wave table stored in memory for every change in the RMS voltage.

5 Related Work

A number of commercial and research power meters have been developed for plug load energy monitoring, with many more on the way. Commercial plug load meters include the Kill-A-Watt [1], Plogg [2], and Watts Up [5] devices. Research plug load meters include the ACme [7], Plug [10], and Smart-Socket [15] devices. Nearly all of these meters are plugged into a wall socket and measure the voltage and current of an attached load to compute true and RMS power. Their integrated design does not permit easily decoupling the current and voltage channels, limiting their measurement scenarios.

Several commercial and research power meters have been developed for whole house monitoring as well. Blue Line Innovations' PowerCost Monitor [3] clamps to an existing Watt-Hour meter typically installed on the outside of a building. PowerCost tracks energy consumption using either an optical output port on an electronic meter, or by counting revolutions of a spinning disk on a mechanical meter, making it unable to support the more interesting metering topologies we envision. The Energy Detective (TED) [4] uses split-core current transformers (CTs) that are installed inside a circuit breaker box. The current sensors are connected to a measuring transmitting unit (MTU), which also exposes two wires that connect to phase A, and neutral, to measure the line voltage. These wires are also used to power the MTU and transmit post-processed, low-rate data over the in-building power lines to a gateway or display unit. Since the MTU uses the voltage channel for measurement, power, and data, it is not amenable to interfacing with a synthesized, logic-level voltage waveform, rendering voltage virtualization impractical.

Patel et al. designed a whole house, contactless power meter that is architecturally similar to our work [13]. Their design uses a pair of magnetometers to estimate current flow in the two bus bars of a circuit breaker box using a custom, peel-and-stick sensor. The sensor is externally powered but transmits its readings wirelessly at 1 kHz to a Bluetooth-enabled PC. The PC computes RMS current, which it then multiplies with the RMS voltage obtained from the AC mains. By averaging the current and voltage measurements, their approach does not require the two channels to be tightly synchronized, but the approach is also susceptible to power calculation errors. Computing RMS current and voltage and only then multiplying the values may work for resistive loads but will introduce errors for switching loads (with harmonic power factors) and reactive loads (with displacement power factors). We show that their approach of using RMS current and voltage to compute power overestimates the average power by over 6% for a MacBook but our approach does not.

Finally, synchronized phasor measurement units (synchrophasors or PMUs) are devices that can extract current and voltage phase relative to a GPS clock with 1 μ s accuracy [14]. PMUs are expensive devices used to monitor grid-scale operation to identify voltage sags and phase offsets that could indicate pending grid instability. Our design shares some similarities with synchrophasors, but it does not require GPS clocks, perform DFT or symmetrical component transformations, or support wide-area deployment.

6 Conclusion

We present a distributed power meter that decouples the voltage and current sense channels, and recombines them wirelessly for power calculations. Disaggregating the power meter in this manner allows us to install voltage and current sensors in individually optimal locations, and it also supports a one-to-many relationship between the sensors. This design supports dense, accurate, in-building power metering; any wire, anywhere can be instrumented quickly and non-invasively to provide real and reactive power measurements. Since a circuit need not be broken, nor must live wires be tapped, this approach is safe, inexpensive, and practical.

7 Acknowledgments

Special thanks to Mani Srivastava for his many insightful comments and to the anonymous reviewers for their feedback. This material is based upon work supported by NSF Awards #CNS-0964120 and #CPS-0932209. NSF support was also provided under Award #CNS-1019343 to the CRA for the CIFellows Project. This work was supported in part by an SNSF Fellowship and an MSR Graduate Fellowship.

8 References

- [1] Kill-a-watt. <http://www.p3international.com/products/special/P4400/P4400-CE.html>, 2010.
- [2] Plogg Wireless Energy Management. <http://www.plogginternational.com>, 2010.
- [3] PowerCost Monitor. <http://www.bluelineinnovations.com/Products/>, 2010.
- [4] TED: The Energy Detective. <http://www.theenergydetective.com/>, 2010.
- [5] Watts Up Meters. <https://www.wattsupmeters.com/secure/products.php?pn=0>, 2010.
- [6] P. Dutta, J. Taneja, J. Jeong, X. Jiang, and D. Culler. A building block approach to sensor networks. In *SenSys'08: In Proceedings of the Sixth ACM Conference on Embedded Networked Sensor Systems*, Nov. 2008.
- [7] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *IPSN '09: Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, pages 253–264, Apr. 2009.
- [8] X. Jiang, M. Van Ly, J. Taneja, P. Dutta, and D. Culler. Experiences with a high-fidelity wireless building energy auditing network. In *SenSys'09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, pages 113–126, Nov. 2009.
- [9] 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 '09: Proceedings of the 11th international conference on Ubiquitous computing*, pages 245–254, 2009.
- [10] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso. A platform for ubiquitous sensor deployment in occupational and domestic environments. *IPSN'07: Proceedings of the 6th international conference on Information processing in sensor networks*, Apr. 2007.
- [11] A. Marchiori and Q. Han. Using circuit-level power measurements in household energy management systems. In *BuildSys'09: Proceedings of First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, Nov. 2009.
- [12] M. Maróti, B. Kusy, G. Simon, and Á. Lédeczi. The flooding time synchronization protocol. In *Sensys'04: Proceedings of the 2nd ACM Conference on Embedded Networked Sensor Systems*, pages 39–49, Nov. 2004.
- [13] S. N. Patel, S. Gupta, and M. S. Reynolds. The design and evaluation of an end-user-deployable, whole house, contactless power consumption sensor. In *CHI'10: Proceedings of ACM Conference on Human Factors in Computing Systems*, Apr. 2010.
- [14] A. G. Phadke. Synchronized phasor measurements—a historical overview. In *Transmission and Distribution Conference and Exhibition 2002: Asia Pacific. IEEE/PES*, volume 1, pages 476 – 479, 6-10 2002.
- [15] B. Priyantha, A. Kansal, M. Goraczko, and F. Zhao. Tiny web services: Design and implementation on interoperable and evolvable sensor networks. In *Sensys'08: Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems*, Nov. 2008.