Any Wire, Anywhere: Just Virtualize 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 meter device. In this paper, we argue that decoupling the two channels enables new sensing scenarios (e.g. attributing line losses to the loads that cause them) and offers safer and simpler aggregate (e.g. whole house) and dense (e.g. plug load) monitoring. However, decoupling the channels raises a new question: how should they be recombined? A long run of wires? High-rate, real-time, wireless sensor data transfer? Separate current and voltage channels that time-stamped and logged in real-time, but recombined later? Of the various approaches, we propose the voltage channel be virtualized: a single (root) node in the network measures the voltage magnitude, frequency, and phase. The phase is time-stamped relative to a global clock and disseminated wirelessly, along with the voltage and frequency measurements, to all other meter nodes in the network. The other nodes synthesize a suitably scaled replica of the voltage waveform locally, based on the parameters reported by the root, and combine it with locally-measured current readings. By measuring voltage near a home's service entrance and current at a load, for example, wiring losses can be attributed to the load. Or, by measuring current at the service drip loop, whole-house power can be measured easily.

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

Submitted to: BuildSys'10, November 2, 2010, Zürich, Switzerland.

1 Introduction

Today, in-building power measurements are either centralized (e.g. placed near a breaker box) or distributed (e.g. placed near a plug load). The former provides *full* coverage of all the loads while the latter provides *detailed* coverage of only the instrumented loads [7]. In this paper, we argue that these two design points force a false dichotomy, and that a much richer set of power metering options are possible 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, attributing line losses to a particular load 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 is often challenging since feeds are 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 a host of 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 mains 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. This approach allows us to: (i) achieve both full and detailed coverage, (ii) easily attribute, for the first time, wiring losses to each load, and (iii) simplify whole house power meter deployment.

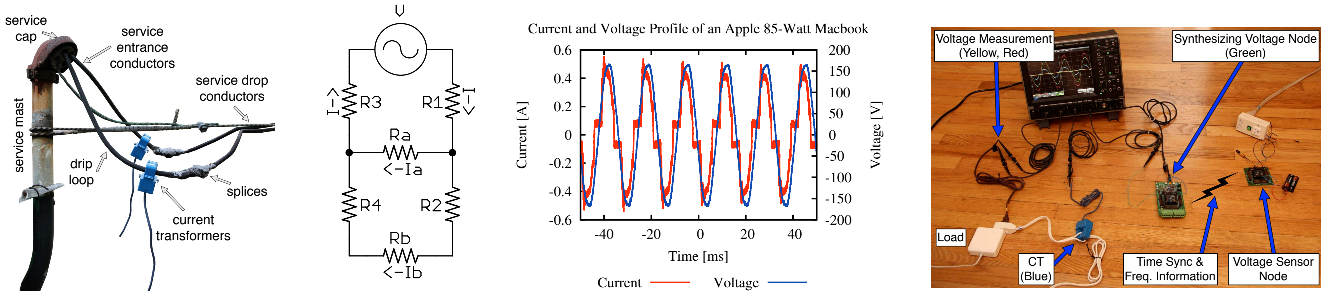

(a) Drip loop sensing. (b) Line losses. (c) MacBook power profile. (d) Distributed power meter design.

Figure 1: (a) Service drip loop details; (b) Circuit model for line loss analysis; (c) MacBook voltage/current profile. (d) System architecture.

2 Motivating Applications

To motivate the need for a distributed power meter, we present two applications that benefit from decoupling current and voltage: whole house metering and line loss attribution.

2.1 Whole House Power Metering

One of the challenges to end-user-deployed, whole house power metering is getting access to the current channel. Most approaches require professional installation due to the safety concerns of attaching current transformers near uninsulated wires inside of a circuit breaker box [10]. We note, however, that the near-universal use of drip loops at the power service entry point provides a convenient attachment point for current transformers, as Figure 1(a) shows. The drip loop also presents 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 is true indoors. Distributed power metering addresses the problem by relaxing sensor placement.

2.2 Load-Level Line Loss Attribution

Attributing in-building losses to the loads that cause them is difficult because these losses depend on the wiring topology, wire gauges, load profiles, and load attachment points. Although the absolute wiring losses are required by code to be less than 7%, the actual figures are rarely known or attributed to specific loads, but 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 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 link.

Patel et al. propose a variant of this approach that streams non-timestamped current readings at a lower rate (1 kHz) and computes an RMS value of the current and voltage samples before multiplication [10]. Unfortunately, this technique works for resistive but not reactive or switching loads, suggesting a different approach is needed.

Another possibility might be to parametrize the current or voltage waveforms, and wirelessly transmit only the 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 [6] 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 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 [9] 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 current implementation, the scope data are post-processed for evaluation, but an integrated design could perform all of the functions on-line and in real-time.

4 Evaluation

We now explore the viability of our distributed power meter design. Specifically, we explore the foundational assumption that the voltage phase is constant throughout a circuit, that the voltage waveform is sufficiently sinusoidal, that the voltage waveforms can be synthesized accurately, that tight time synchronization can be maintained for extended periods of time, and that changes in the voltage period 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 and constant phase throughout a circuit. We first explore how consistent the magnitude, phase, and frequency are across several loads distributed throughout 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, at nominally 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 ensure that even large loads do not change the phase or frequency at the different measurement points. The only expected difference is a voltage depression near a large load due to increased line losses.

Figure 2 shows the measurement results for three different scenarios. Figure 2(a) is the baseline when the total household energy consumption is 0.5 kWh. This load includes lights, a fridge, 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 kWh, the average RMS error between the measured curves is 6.9 V.

To quantify the effect of line losses, we connect a 1.4 kWh 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 remains at the previously measured levels.

In the final scenario, we explore whether the line voltages and waveform across a house are affected when a large load is switched on. We add a 2.4 kWh load to the house by turning on an air conditioner. This additional load does not affect the measured RMS voltage, nor does it introduce a phase change between two power legs, as Figure 2(c) shows.

The results of these three distributed measurements show that adding loads principally affects the RMS voltage, and primarily the voltage on a particular circuit, but not the phase. This suggests that a single-point voltage measurement near the service entry point is sufficient to compute plug load power, even when individual loads are very different.

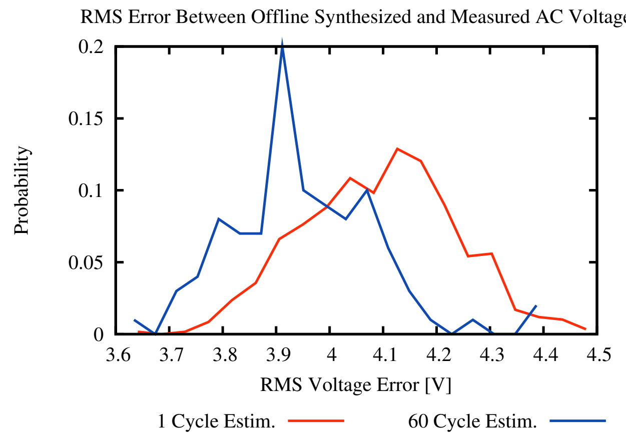

Figure 3: RMS error distribution between AC voltage and a synthesized sine wave. We either used 1 cycle, or 60 cycles to estimate the RMS of the AC line voltage to synthesize a sine wave with same amplitude and phase. While the difference is only minor, using 60 cycles for RMS estimation achieves a smaller RMS error than using a per-cycle calculation.

We next explore how "sinusoidal" the voltage waveforms are. 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 parameterized using these three metrics, then the waveform synthesis will introduce an error, and power (and power factor) estimation will suffer. We compute the RMS value of each voltage cycle and use this value to synthesize one cycle of a sine wave whose amplitude is $\sqrt{2}$ times greater. We then subtract the synthesized sine wave from the captured waveform and compute the RMS value of this error. In a second run, we use 60 voltage cycles to compute the voltage RMS and compare this to 60 cycles of a synthesized sine wave. Figure 3 shows both error distributions. These results show that the voltage waveform is highly sinusoidal and can be parameterized by magnitude, frequency, and phase with less than 4.5 V (3.75%) RMS error.

4.2 Synthesizing Voltage Waveforms

A power meter requires accurate time synchronization and the ability to synthesize a sinusoidal waveform based on the magnitude, frequency, and phase reported by the voltage sensor. The accuracy of this 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) [9] with a 32768 Hz timebase. Prior work has shown that with modern radios, a regression history of four values, and a resynchronization period of 10 s, accuracies of $< \pm 1$ tic ($< \pm 30.5$ μs for a 32768 Hz clock) can be achieved.

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 measures the degree of variation in a clock signal over a time window τ, and is used to describe clock stability over time.

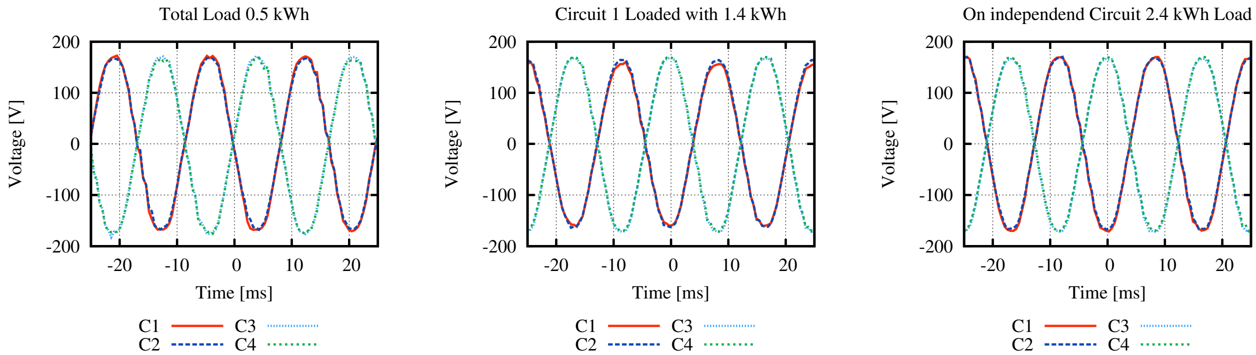

(a) Light 0.5 kWh total load.

(b) Large 1.4 kWh load near C1.

(c) A/C 2.4 kWh on separate circuit

Figure 2: Synchronized voltage measurement at four different outlets across a single phase one family house. 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.

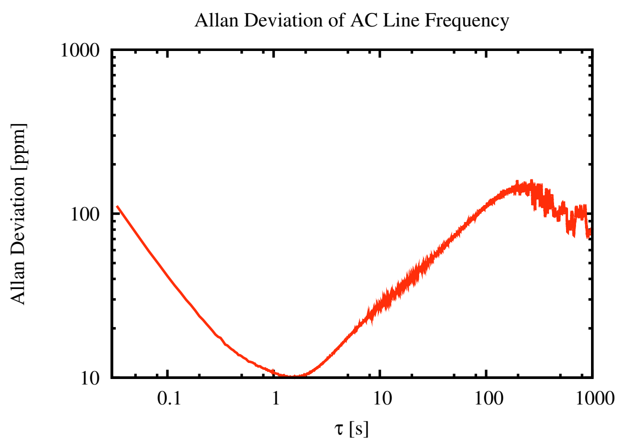

Figure 4: Allan Deviation of the line frequency. The Allan Deviation measures the degree of variation in a signal over a time window τ. We observe that the Allan Deviation rapidly decreases initially, achieving a minimum at about one second, before increasing. Over very large time frames (hours to days) 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.

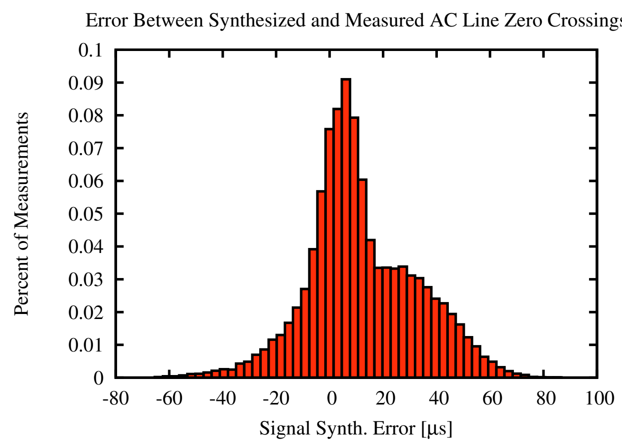

Figure 5: Histogram of measured synthesizing error of the zero crossing time. A voltage sensor measures the zero crossing times of the AC voltage and periodically sends these timestamps to a time synchronized power meter node that synthesizes a corresponding zero crossings waveform. This histogram shows the phase errors resulting from the power meter's synthesis of the zero crossing signal based on a 32768 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. As the measurement was performed in an air-conditioned room with minimal temperature changes, the oscillator frequency is assumed to be stable. The graph shows that the most stable frequency accuracy can be found over a time interval of one second. Thus, for the remainder of our experiments, the voltage sensor distributes magnitude, frequency, and phase parameters of the voltage waveform every second. The reported parameters are based on the average value of the measurements observed over the preceding second, and the phase is measured relative to the global time provided by FTSP, that resynchronizes every 10 s.

An important measure of our system is the accuracy with which a node can recreate the zero crossing times measured by a voltage sensing node. We programed our synthesizer node to output the estimated zero crossing times using a gpio pin on the microcontroller. To precisely toggle the pin, we used the timer capabilities of the TI MSP430 chip used in our system. The voltage sensor node was sending out the frequency and phase information of the zero crossings at a 1 s interval. Figure 5 shows a histogram of the time difference between the actual zero crossing and the zero crossing of the generated signal. We measured this difference over 33'000 zero crossings using an oscilloscope. We can see that the error has a bimodal distribution with a large peak around 5 μs, and a second smaller peak around 20 μs.

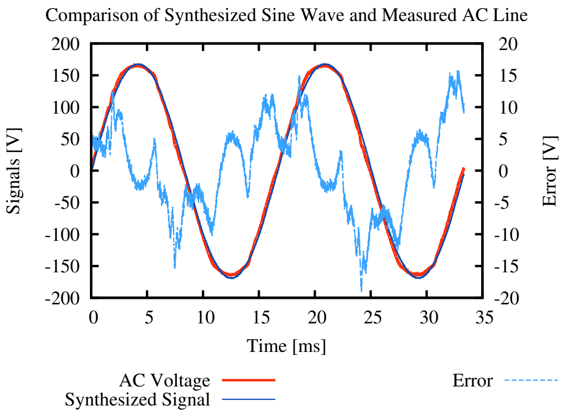

Figure 6: Comparison between scaled synthesized voltage wave of a power meter node, and the actual AC line voltage. The power meter node uses the synchronization information from the voltage sensing node to output a synchronized sine wave using the microcontroller's DAC unit. The scaling of the sine wave was done offline using Octave and the RMS of both signals. The RMS of the error between the two signals is 6.8 V.

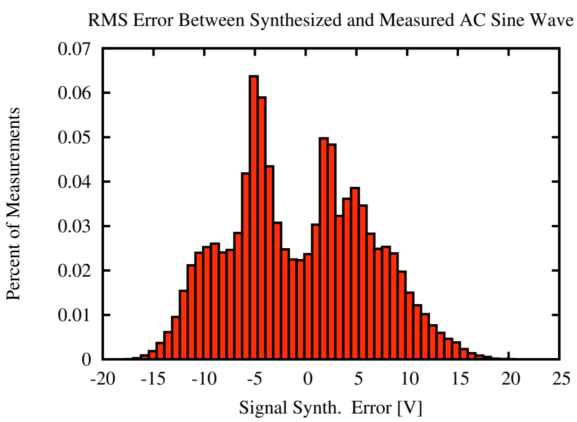

Figure 7: Distribution of voltage synthesis errors. The errors are bimodal with peaks around -5 and 5 V. These errors are due to time sync error and its impact on phase offset. The errors are symmetric, so they largely cancel.

We now use the zero crossing estimation to trigger a DMA transfer of a precomputed sine wave to the DAC of the microcontroller. Using all hardware peripheral reduces the risk of missing interrupts and increases the timing accuracy of the generated signal. We used an oscilloscope to measure the AC line voltage and the generated DAC signal at the same time. Figure 6 shows the two measured waves plus the error between them. Figure 7 shows the PDF of the error distribution. Note that the scaling of the generated sine wave was performed offline using the RMS voltage ratio between the generated and the AC line signal.

The results so far suggest that a remote node can synthesize the voltage sine wave to about 5-10% accuracy given a voltage sensing node sends out frequency, phase, and magnitude every 1 second, and that we run a time synchronization process with a resynchronization rate of 10 seconds. Thus, we now investigate how these errors impact power and power factor estimation.

Trace	AC Pwr. Power	Synth. Power	Synth. Error	RMS Pwr. 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.

4.3 Macroscale Evaluation

To measure the error introduced by using the synthesized, instead of directly measured voltage, we conduct two experiments. We compare the true and RMS power by calculating it directly from traces obtained with an oscilloscope, and by using the measured synthesized signal. We connected the voltage sensor node to an unloaded circuit in a different room, transmitting phase and frequency information every 1 second. We used a current transformer to measure the current drawn by a 150 Watt light bulb and a 85 Watt MacBook. Concurrently, the oscilloscope measured the synthesized waveform from the DAC of a synthesizer node. We collected several 2 second long traces sampled at 5 MS/s and calculated 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. In summary, using the synthesized voltage waveform results in an average true power estimation error of 1.2% for the MacBook, and 0.7% error for the light bulb. As the MacBook is a switching load, RMS power estimation overestimates the power drawn by more than 6.5%, highlighting the need for power factor estimation.

4.4 Limitations

Our hardware relies on 32768 Hz clocks as a time reference. Switching to a higher speed clock will increase synchronization accuracy, and thus reduce the phase error between the generated and measured voltage wave. 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 combo, 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 measuring node to the power meter node. The scaling of the synthesized sine wave was performed offline. The following is a sketch of how this scaling could be performed. Instead of changing the sine wave output on the DAC0, we use a second DAC1 output 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 computation load on the microcontroller as it doesn't have to recompute the sine wave samples stored in memory every time the RMS voltage changes.

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 Kill-A-Watt [1], Plogg [2], and Watts Up [5] devices. Research plug load meters include the ACme [6], Plug [8], and SmartSocket [12]. Nearly all of these meters are plugged into a wall socket and measure the voltage and current of an attached load to compute 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. Although the MTU's current sensors could be replaced with virtual current measurements, we show in this paper that virtualizing the voltage channel makes much more sense. Since the MTU uses the voltage channel for measurement, power, and data, it is not amenable to interfacing with a synthesized, logic-level voltage, rendering virtualization impractical.

Patel et al. designed a whole house, contactless power meter that is architecturally similar to our work [10]. 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 overestimates average power by nearly 7% for a MacBook, but ours 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 μs accuracy [11]. 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. We show that a distributed power meter can be used to attribute line losses to the loads that cause them. More importantly, this design supports dense, accurate, in-building power metering: *any wire, anywhere* can be instrumented quickly and non-invasively. Since a circuit need not be broken, nor must live wires be tapped, this approach is safe and inexpensive.

7 References

[1] Kill-a-watt. `http://www.p3international.com/products/special/P4400/P4400-CE.html`, Aug. 2010.

[2] Plogg Wireless Energy Management. `http://www.plogginternational.com`, Aug. 2010.

[3] PowerCost Monitor. `http://www.bluelineinnovations.com/Products/`, Aug. 2010.

[4] TED: The Energy Detective. `http://www.theenergydetective.com/`, Aug. 2010.

[5] Watts Up Meters. `https://www.wattsupmeters.com/secure/products.php?pn=0`, Aug. 2010.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.