

Monjolo: An Energy-Harvesting Energy Meter Architecture

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

ABSTRACT

Conventional AC power meters perform at least two distinct functions: power conversion, to supply the meter itself, and energy metering, to measure the load consumption. This paper presents Monjolo, a new energy-metering architecture that combines these two functions to yield a new design point in the metering space. The key insight underlying this work is that the output of a current transformer – nominally used to measure a load current – can be harvested and used to *intermittently* power a wireless sensor node. The hypothesis is that the node’s activation frequency increases monotonically with the primary load’s draw, making it possible to estimate load power from the interval between activations, assuming the node consumes a fixed energy quanta during each activation. This paper explores this thesis by designing, implementing, and evaluating the Monjolo metering architecture. The results demonstrate that it is possible to build a meter that draws zero-power under zero-load conditions, offers high accuracy for near-unity power factor loads, works with non-unity power factor loads in combination with a whole-house meter, wirelessly reports readings to a data aggregator, is resilient to communication failures, and is parsimonious with the radio channel, even under heavy loads. Monjolo eliminates the high-voltage AC-DC power supply and AC metering circuitry present in earlier designs, enabling a smaller, simpler, safer, and lower-cost design point that supports novel deployment scenarios like non-intrusive circuit-level metering.

Categories and Subject Descriptors

B.0 [Hardware]: General; B.4 [Hardware]: Input/Output & Data Communications; H.4 [Information Systems Applications]: General

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Energy harvesting, Power metering, Wireless sensor networks, Data aggregation, Intermittent power

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SenSys '13, November 11–15, 2013, Roma, Italy

Copyright 2013 ACM 978-1-4503-2027-6/13/11 ...\$15.00.

1. INTRODUCTION

A recent National Science and Technology Council study claims that “for the foreseeable future, the greatest national energy saving potential lies with improvements to existing buildings.” [18] Another NSTC study projects that, “If current trends continue, by 2025 buildings worldwide will consume more energy than the transportation and industry sectors combined.” [19] The latter study suggests that, “the building community should view submetering as an essential component of future building operational improvements for energy efficiency and conservation improvements,” but it also notes that there has been “less emphasis placed on submeters,” and that “deployment of large-scale submetering of plug loads is not currently widespread” due to the “high installation costs” and “complexity and challenges of the submetering solutions.” Existing submetering systems are expensive, due in part to the complexity of their power supply and measurement circuitry, and in part to their installation costs, particularly at the panel- and circuit-levels. They are also physically intrusive, due to their need to tap into mains power. Finally, they often exhibit high standby power, which runs counter to the conservation efforts they seek to promote.

We claim that a new energy meter design and metering architecture is needed to support submetering at building scales, and we present Monjolo, a power-proportional, energy-harvesting energy meter architecture as its keystone. The Monjolo principle is simple: a sensor device harvests energy from the magnetic field of a wire supplying power to an AC load. When sufficient energy has accrued, Monjolo powers up a radio and sends a packet, consuming a fixed energy quanta during each such activation. Under steady-state conditions, the activation rate is proportional to the load power, so packet transmission timing conveys the load power. This indirect method of energy metering – observing the dynamic operation of an energy-harvesting power supply – bears a conceptual similarity to the iCount [11] energy meter. But unlike iCount, which fixes the input voltage and measures the regulator switching frequency to estimate load power, Monjolo fixes the energy consumed during each activation cycle and measures the activation frequency.

The Monjolo design offers power-proportionality. When the AC load is off, no current flows in the supply wire, so Monjolo does not wake up. Therefore, Monjolo draws zero standby power. As the load power increases, Monjolo wakes up more frequently, so its power overhead increases, but only in proportion to the AC load’s power draw. Coupled with carefully optimized boot up, this power-proportional operation enables an energy meter that can operate from just a few microwatts and still deliver measurements with minute-level temporal granularity. Conventional energy meters offer neither of these virtues; their active measurement circuitry draws power continuously, regardless of whether the load is active or not, and they cannot operate from microwatts.

The Monjolo design eliminates the high-voltage AC-DC power supply and AC metering circuitry that is present in most conventional AC power meters. The power supply circuitry is difficult to miniaturize due to the large inductors and capacitors needed to store sufficient energy at the low (50 or 60 Hz) mains frequency. The metering circuitry in conventional meters must sample and process a current (and often a voltage) channel at kilohertz frequencies, which increases power draw and complexity. In addition, these circuits expose high voltage and add cost. Monjolo eliminates both of these explicit subsystems and performs their operations implicitly by monitoring the operation of an energy-harvesting power supply. This approach sacrifices some measurement fidelity for improvements in size, power, cost, isolation, and flexibility.

Monjolo redefines energy metering around a foundational hypothesis – *the energy harvester is the energy meter* – that claims the rate at which an energy harvesting power supply can harvest increases monotonically with load power. The harvesting rate is measured in Monjolo activations and timing information about the activations is relayed via packet transmissions from the Monjolo sensors to a data aggregator. The aggregator keeps track of the interval between activations, and uses this information to estimate the load power. Activations are infrequent (a few to tens of Hz for typical household loads) and packet transmissions are quick (one or two milliseconds), so several active loads can coexist without congesting the radio channel within a single collision domain. To support higher deployment densities, greater load power draws, and increased data delivery reliability, we employ two additional techniques. A non-volatile activation counter keeps track of every wakeup and the counter value is transmitted in every packet. A countdown timer circuit limits the transmission rate by squelching packet transmissions during the countdown window. Although unsynchronized transmissions of short packets is incompatible with traditional low-power link layers, Monjolo has access to the mains AC phase, so future versions could synchronize to this signal. Meanwhile, we employ Monjolo using a star topology to demonstrate viability of the fundamental approach.

To evaluate our hypothesis and validate this method of energy metering, we implement Monjolo in three forms – a breakout board, a plug-load meter, and circuit-level meter – and characterize it across a variety of real loads. Our results show that for resistive loads, Monjolo reports power readings with an average of 3.7% error. With reactive loads we find that Monjolo is able to clearly identify the different phases of operation for a given load and that it effectively measures apparent power (as opposed to real power). We present a method for recovering real power for a reactive load using a panel-level energy meter in conjunction with Monjolo. We also compare Monjolo to commercially available energy meters and find that Monjolo’s power overhead is at least one order of magnitude lower, while maintaining similar sampling rates and often better networking properties.

2. RELATED WORK

Growing attention is being focused on both energy harvesting power supplies and miscellaneous electrical load metering. Residential utility customers, in response to both environmental and financial concerns, are expressing a growing interest in intelligently reducing their energy consumption. To meet this growing demand for greater visibility into energy consumption, both industry and academia have designed a variety of energy metering solutions to the plug-load and circuit-panel metering problems. Further, utility companies encourage and incentivize efficiency behaviors to reduce or delay the need to build or operate high-cost, low-efficiency peaker plants.

Cooper Power Systems offers a power line energy harvester [1]. This device attaches to high-voltage overhead power lines and provides power for a separate wireless node. Gupta et al. also explore whether energy can be harvested from the electromagnetic energy radiating from AC power lines [13]. Both use similar energy harvesting technique as Monjolo, but neither offers energy metering.

The information provided by energy usage feedback has been shown to reduce residential energy consumption by as much as 55% [22]. This was observed in a university dormitory environment by providing feedback and incentives in response to conservation. In a real-world application, incentives will come in the form of savings on utility bills but feedback, especially feedback at the plug-load level, is more difficult. In particular, this type of study, in which two particular dormitories were engineered to provide “high resolution” feedback, is infeasible at large scale.

Several academic and industry energy metering platforms attempt to provide this feedback through a typical sense-and-report system. ACme [14] and Microsoft Research’s Smart-Socket [23] use similar architectures to provide high accuracy reporting using Analog Devices energy metering ICs: the ADE7753 [8] and ADE7757 [9], respectively. Plug uses a current transformer fed into an ADC to meter apparent power [16]. Many commercial meters follow this same sense-and-report approach, such as the Kill-A-Watt Wireless [3], SmartHome’s iMeter Solo [6], and the Watts up?.Net [7]. In all of these cases the meter is powered from an AC-DC converter, is always on, and periodically samples both the current and voltage channels (regardless of the current sense method). In contrast, Monjolo is only active when a load draws power.

Several meter designs attempt to reduce the size, component count, complexity, or energy required to meter by using a variety of techniques. One such implementation moves voltage sensing to a central location and only measures current at each plug load [24]. This allows for smaller and lower power sensors. Another approach expands on this idea by metering total power draw at a central location and having each plug-load meter only for ON/OFF state sensing [15, 26]. This data can then be used to determine time-in-mode, and estimate the power consumption of the main loads in the house. In both cases, however, the system is still always on, even when the load is not. An energy harvesting approach from EnOcean [2] solves the active power issue with energy harvesting, but it requires two current transformers and an AC load of at least 0.7 A, higher than Monjolo. This makes it poorly suited for a small plug-load measuring device and for measuring small loads.

Another option is to multitask the load sensing mechanism to also power the sensor node. One such device, a “Stick On” sensor, uses a piezoelectromagnetic (PEM) current sensor to sense the current supplied to the load [27]. Because these PEM transducers can generate several milliwatts, they can be used not only to sense but also to supply power [21]. This “Stick On” sensor system has an optional power supply utilizing this property, allowing the meter to be powered from the load while remaining fully non-contact. While the “Stick On” sensor’s power supply is similar to Monjolo, the implementation uses two transducers, operates only at the circuit breaker level, suffers from adjacent channel crosstalk, and measures current directly rather than indirectly.

Studies have analyzed the possibilities associated with using plug load meters. Energy Lens is a system that uses ACme nodes but explores areas of data visualization and real time aggregation [20]. The Smart Energy Meter is a custom meter that was built not as the focus of the project but as a means to test a system’s ability to intelligently actuate AC loads [25]. Monjolo does not aim to improve or replace this work, but rather provides a hardware sensing layer that could provide the data for these higher level systems.

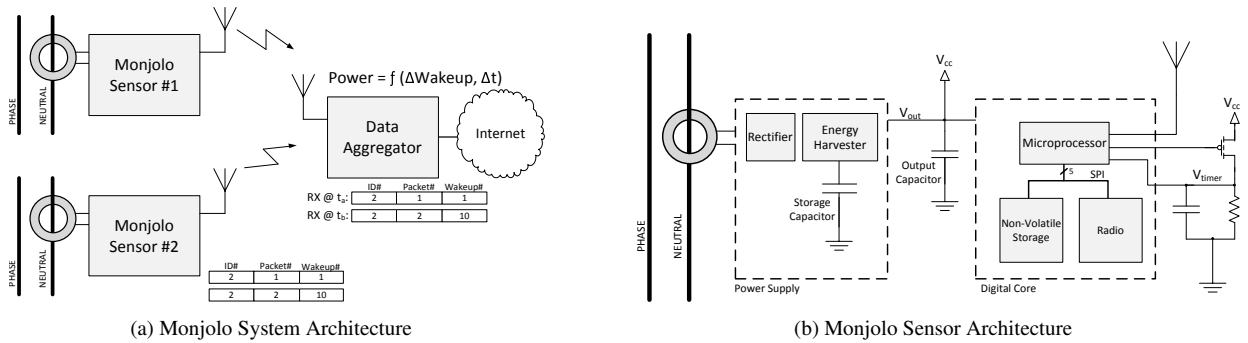


Figure 1: System overview. (a) Two Monjolo nodes transmitting to the data aggregator. The data aggregator uses consecutive values transmitted from a sensor to estimate power and forwards that data to the Internet. (b) A detailed view of the Monjolo sensor architecture. The power supply harvests energy using a current transformer, rectifies the resulting AC current, charges a capacitor, and notifies the digital core when there is sufficient energy stored to enable activation. The digital core then boots up, increments a non-volatile counter, and transmits a packet with the counter value and other identifying information.

All of these systems share in the goal of increased awareness of plug-load energy consumption, but they all share a common drawback as well: they are active devices that are always on and draw power regardless of the whether the load is drawing power. iCount, in contrast, performs “energy metering for free” by counting cycles of a pulse-frequency modulated switching regulator [11]. The key observation underlying iCount is that switching frequency is proportional to load current, provided that the input voltage is held constant. The Monjolo design is similar, but we note that the activation frequency of an energy-harvesting power supply is proportional, but not necessarily linear, to the input power, provided that the load is held constant.

3. SYSTEM OVERVIEW

Monjolo is an energy metering architecture designed to measure AC loads with little or no overhead. It is comprised of an energy harvesting sensor and a data aggregator that collects data from the sensor to estimate the power draw of the load. The energy harvesting sensor has two useful forms: a plug-load meter for miscellaneous electrical loads, and a circuit-level meter for installation in circuit breaker panels or drip loops, that can be clipped around a supply line. Figure 1 shows the system and sensor architecture.

3.1 Sensor

The sensor system itself is composed of two parts: a power supply which harvests and regulates energy, and a digital core that maintains counters and transmits packets. The power supply utilizes a current transformer around one path of the power line running from the AC main to the load itself. The current in this line induces a magnetic field in the current transformer, which in turn induces an AC signal on the secondary. This current transformer output is rectified and harvested to power the digital core. The rate that energy can be harvested is proportional to the power draw of the load being measured. We exploit this relationship to estimate the power draw of the attached load.

The digital core is disconnected from power by the energy harvester while the harvester is in operation. When the energy harvester has collected enough energy to transmit a packet, it powers the core which reads and increments a non-volatile counter and transmits a packet. Because the harvesting capability of the power supply is proportional to the load power, the frequency of packet transmissions is also proportional to the load power. This relationship represents the core operating principle of the energy meter.

Running the energy harvester adds little or no overhead to the main load. In addition to not having an AC-DC power converter, a factor in making this possible is that the Monjolo meters do not themselves sample electrical waveforms or compute complex power. Rather, the data aggregator trivially “computes” power, which reduces the complexity and power requirements of each sensor.

3.2 Data Aggregator

The data aggregator listens for packets from one or more Monjolo sensors. Upon packet reception, it computes the interval between the packets and uses that value to estimate the power of the load attached to the sensor. The data aggregator is then responsible for displaying that information to the user. Data aggregators are also responsible for archiving the measurements for historical analysis or relaying that information to other entities for further analysis, processing, or visualization.

4. ENERGY METER DESIGN

This section presents the design of the Monjolo energy meter. We begin with a simplified version of Monjolo that makes its basic operating principles clear. We then progressively relax simplifying assumptions and introduce real-world concerns to yield a design that works in practice.

4.1 Basic Operating Principle

The core operating principle of the Monjolo meter is quite simple: energy is harvested into a reservoir at a rate proportional to the power draw of a load to be measured. When the reservoir – in this case a capacitor – has accrued a certain amount of energy, the energy harvester activates a microcontroller and radio to quickly transmit a packet using this harvested energy. The microcontroller then ensures that a fixed quanta of energy is discharged from the reservoir over a fixed amount of time before allowing the reservoir to begin recharging. When the reservoir reaches the threshold again, the wakeup and transmission cycle repeats.

We claim that the interval between packets can be used to determine the load’s power draw, even if the packet transmission rate does not scale linearly with load power. As long as the packet transmission rate increases monotonically with load power, and the packets are quickly and successfully received by the data aggregator, we can estimate load power from the inter-packet interval. We explore the precise nature of the relationship between packet transmission rate and load power draw in Section 5.

4.2 Dealing with Data Loss

Using the design in Section 4.1, every wireless packet that is transmitted must be received to accurately reconstruct the power of the load. If packets are lost, the load power estimation error will increase and the load will be underestimated. To mitigate this source of error, Monjolo transmits a monotonically increasing counter value with each packet transmission and the aggregator caches the last packet timestamp and counter value received from each Monjolo sensor. The aggregator can then divide the difference in the packet sequence numbers by the wall time elapsed to account for lost packets. This scheme does not correct for the loss in temporal resolution from dropped packets, as any delay between packets prevents the load power estimate from being updated. Also, we can only estimate the average power of the load during the period in which the packets are dropped by assuming a constant load.

4.3 Moderating Packet Delivery Rates

One problem with the Monjolo design is that its packet transmission rate is a function of load power: as the load power increases, so does the packet transmission rate. Further, the Monjolo system in Section 4.2 has no long-term sense of time so it cannot moderate its own packet delivery rate (its power is disconnected after each activation). If multiple Monjolo sensors are colocated, this could in principle saturate the wireless channel, resulting in data loss from packet collisions, or inequitable data delivery ratios.

We address this problem by decoupling the node activations from the packet transmissions. Node activations continue to increment a counter, but the packet transmission rate is capped. One challenge is how to cap the transmission rate, since this requires keeping track of time in the absence of steady source of power. Our solution notes each packet transmission internally by charging a small capacitor. Over time this capacitor is slowly discharged through a large resistor. If Monjolo’s digital core wakes up before the capacitor has discharged to a low threshold, the core increments the counter but does not transmit a packet. This RC timekeeping circuit provides a simple mechanism to keep time and squelch packet transmissions.

Adding the RC timekeeping circuit decouples packet transmission frequency from primary power. This means that the wakeup counter from any two packets, along with the time difference between them, can be used to accurately determine the average power over that interval. In addition to the activation counter, which can increment many times between packet transmissions, Monjolo also keeps a sequence number that increments with every packet, which enables the aggregator to detect dropped packets. One important note about suppressing packet transmissions is that Monjolo itself must continue to consume the same quanta of energy during each activation cycle regardless of whether it transmits a packet or not.

4.4 Handling Non-linear Recharge Rates

Ideally the recharge rate of the reservoir capacitor – and by extension the activation rate of Monjolo – would be linear with load power. However, even if it was not, if the relationship could be described with a simple transfer function, it would allow for flexibility in reporting load powers for which no exact calibration values exist. In practice, non-linearities in the system, like current transformer core saturation, means that this relationship cannot be described by a simple analytical function. We address this by fitting a model to the empirically-derived relationship between activation rate and load power, and then use this model to estimate the corresponding load power from the instantaneous activation rate (or interval). We further elaborate on the transfer function and on how we find and use this model in Section 5.2.




Figure 2: Sensor prototype used to evaluate the Monjolo principle, including an LTC3588-based energy harvester, Epic mote, FRAM non-volatile storage, RC time keeping circuit, and configuration switches to select between storage (input) capacitances of 200 μF , 300 μF , and 500 μF , output capacitances of 47 μF , 200 μF , and 247 μF , half/full-wave rectification, and regulator output voltage.

5. IMPLEMENTATION

To validate the Monjolo principle and evaluate its accuracy, we implement and test a prototype system. Figure 2 shows the Monjolo sensor prototype and its various configuration options.

5.1 Hardware Platform

Due to the minimalist nature of the Monjolo software, most of the functionality comes from decisions made at the hardware level. Further, most of the configuration options, both for testing and optimizing for different loads, are implemented in hardware.

5.1.1 Energy-harvesting Power Supply

Monjolo’s energy-harvesting power supply serves as the system’s only sensor. The power subsystem harvests energy from the primary AC load and accrues enough energy over time to wake up and transmit a packet. The aggregator uses the frequency of these wakeups to infer primary power. The following constraints are necessary for proper operation: (i) the system must be able to harvest energy over the full spectrum of possible load powers, (ii) the activation rate must be a monotonically increasing function of the load’s power draw, and (iii) the power supply must never supply so much power as to enter an always-on state. In an always-on state, the harvester can supply more power to the system than the system can dissipate. This saturates the harvester, eliminates the wakeups, and leaves the system without its one sensor.

To achieve these goals with low power and space overhead, we employ the CR Magnetics CR2550, a 1:3000 turns-ratio current transformer [10]. The CR2550 magnetically couples to the load’s supply wire, and supplies the Linear Technology LTC3588 energy harvester [17]. We find that six turns of the primary coil, for a turns ratio of 1:500, ensures that the minimum input voltage requirement of the LTC3588 is met at a primary load power of 17 W and the system still alternates between charge/discharge states at 480 W. Hence, Monjolo can meter loads from 17 W to at least 480 W. This range could be extended further by increasing the primary turns and increasing the energy quanta consumed in each cycle. Figure 3 shows the delivered power from the coil across a 430 Ω resistor, which draws an average current that is close to that of the LTC3588 load. The curve is non-linear at load power levels below 200 W and this non-linearity manifests itself in Monjolo’s activation rate.




Figure 3: Transfer function of the current transformer used in the energy-harvesting power supply when connected to a 430 Ω resistor. The power the coil can deliver varies non-linearly as the load power changes. This affects the activation interval of the sensor, causing the activation rate to be non-linear with load power.

The LTC3588 energy harvesting integrated circuit has two piezo inputs and uses an internal full wave rectifier to charge a capacitor over time. By disconnecting one of these inputs and shunting it to ground, the user can force the LTC3588 to only perform half-wave rectification. We use this mode and employ half-wave rectification to minimize the Monjolo’s internal power draw. Using full-wave rectification would only increase the activation rate, and thereby increase the average power draw, while just slightly improving the granularity of the power measurements.

After the AC waveform is rectified, the LTC3588 employs a boost converter to increase the low amplitude signal to charge the storage capacitor (C_{store}). When this capacitor reaches the threshold voltage of 5.1 V (for a 3.3 V output), the LTC3588 enables a buck converter to supply this output on the V_{out} rail. When the voltage on C_{store} drops to 3.8 V, the LTC3588 disables V_{out} and continues to charge C_{store} back to 5.1 V.

Two capacitors play a critical role in system performance: C_{store} and C_{out} , the latter a capacitor on the regulator’s V_{out} line. The values of these capacitors affect the system operation and measurement granularity. Increasing the value of C_{store} increases the time required to recharge and reduces granularity, but it also increases the runtime duration of each activation cycle. Increasing the value of C_{out} reduces voltage ripple on V_{out} and increases runtime, but also increases ramp-up time for V_{out} .

We evaluate several options for the capacitance C_{store} with the goal of finding the minimum capacitance with which the system can complete all of its required tasks. Because the power delivered by the current transformer for a given load power is fixed, capacitance can affect only the activation rate but not overall system power. The activation duty cycle, and by extension the average power draw, is fixed because the period and runtime scale proportionally as the value of C_{store} changes. Therefore, the minimum possible capacitance needed to run the digital core is optimal. A higher capacitance reduces granularity without reducing power draw. For our system’s operating range of 17 W to 480 W, we empirically determine that a value of 500 μ F for C_{store} provides high granularity and sufficient time to complete all operations, including updating the counter, sampling the timing capacitor, and transmitting a packet. We use a tantalum capacitor for C_{store} because tantalums have low leakage and high charge storage density.

We evaluate three possible choices for the output capacitance: 47 μ F, 200 μ F, and 247 μ F. We observe that although runtime does increase with increased input capacitance, the effective runtime actually decreases due to the increase in ramp-up time. Since all three values sufficiently reduce voltage ripple, we find the optimal value to be 47 μ F for our particular design point.

5.1.2 Analog Packet Suppression Timer

To suppress rapid packet transmissions, Monjolo employs an RC timer. The digital core charges a capacitor (C_{timer}) after each packet transmission. This capacitor discharges over time, allowing a sensor to keep an approximate sense of elapsed time since the prior packet transmission. At every activation, Monjolo samples V_{timer} and suppresses packet transmissions if the capacitor has not discharged to a low threshold level.

Our implementation of Monjolo charges a 2.2 μ F electrolytic capacitor which discharges both through its own leakage and through a 4.7 M Ω resistor. The capacitor is sized to limit the inrush current while offering for a multi-second discharge delay. The value of the resistor creates a time constant of about 10 seconds. We find that for the RF conditions and deployment density in our lab, a reasonable inter-packet interval is about 5 s, but designing for a τ of 10 s allows each Monjolo sensor to enact a higher delay in software, if needed, such as when there is a high density of sensors.

5.1.3 Digital Processing Core

Monjolo’s digital core consists of three main components: a microcontroller (TI MSP430F1611), an IEEE 802.15.4 compatible radio (TI CC2420), and FRAM (Ramtron LM25L04B). The microcontroller and radio are elements of the Epic Core [12]. The FRAM handles the non-volatile storage of the activation counter and packet sequence number. FRAM is preferred to flash or EEPROM for its high-speed reads and low-energy writes. To reduce power, the flash chip on the Epic Core is removed.

The software on the node is optimized for fast and simple operation. Each activation consists of an FRAM read and write, an ADC sample, and a potential radio transmission. Because the LTC3588 disconnects the V_{out} rail when the V_{store} voltage falls below a low threshold, the microcontroller cold boots during each activation cycle. This both simplifies the software and corrects transient errors, as the subsequent boot resets the microcontroller state.

5.2 Sensor Operation

Figure 4(a) shows one activation cycle of Monjolo operation. As the LTC3588 harvests from the current transformer, its boost converter increases V_{store} until it reaches threshold at 5.1 V. The LTC3588 then enables an internal buck converter to regulate the storage capacitor voltage to 3.3 V and supply it on the V_{out} rail. This causes the microcontroller to boot up.

During boot, the microcontroller enables the radio to allow for consistent energy consumption across Monjolo activations. After reading and writing the FRAM to increment the wakeup counter, the microcontroller samples the V_{timer} line to determine if sufficient time has passed since the prior transmission. Figure 4(a) shows this situation occurring at 5.6 s while a pulse on SFD (a signal from the radio) indicates a packet transmission.

The microcontroller then restarts the RC timer by turning on a MOSFET to recharge the capacitor. If the node sent a packet, it writes the FRAM with the sequence number for the next transmission. Finally, with its tasks complete, the microcontroller leaves the radio in receive mode to discharge the storage capacitor to 3.8 V, at which point the LTC3588 disconnects the V_{out} rail, allowing the harvester to recharge the capacitor.




Figure 4: Monjolo energy harvesting operation. (a) A trace of the LTC3588 chip harvesting from an 80 W load and charging up the main reservoir capacitor, C_{store} . When the capacitor reaches 5.1 V, the LTC3588 enables the 3.3 V buck regulator and the microcontroller boots up. The microcontroller samples the V_{timer} line to determine whether to send a packet or not. When the node wakes up at 5.6 s there is only approximately 0.6 V on the timing capacitor, causing the node to send a packet before recharging C_{timer} . (b) The charging cycle of two Monjolo sensors with a half-wave rectifier. Monjolo #1 has a slightly higher load attached and is able to charge up slightly faster than Monjolo #2. However, after the discharge phase both Monjolo sensors start charging at the same point in time, causing their wakeup frequency to equalize with only a slight phase shift. This lack of continuous charging leads to a stair-step relationship between activation rate and load power, which can result in a loss of precision, particularly at high load powers.

5.2.1 Concurrent Charging and Discharging

The choice of the LTC3588 has interesting ramifications for overall Monjolo operation. First, the LTC3588 does not disconnect the input power source while the output is enabled. When the current transformer input voltage exceeds the minimum LTC3588 input threshold, the LTC3588 will supply current to C_{store} as usual. Unfortunately, this violates the constraint established in Section 4.1 that a fixed energy quanta must be consumed over a fixed period of time during each activation cycle. Because the LTC3588 is able to charge C_{store} faster at higher load wattages, the duration and energy depletion of each wakeup will increase with larger primary loads. The change in the wakeup duration introduces error into the power estimation algorithm, which is based on activation rate. If a ΔP increase to a primary load P_{load} causes the same increase in both the Monjolo activation rate and duration, then the overall packet transmission rate of both P_{load} and $P_{load} + \Delta P$ will be identical and indistinguishable.

To compensate for this situation, we add two pull-down MOSFETs, one on each output leg of the energy-harvesting current transformer, with the MOSFET gates connected to the regulator's V_{out} line. This causes the LTC3588 to disconnect its own inputs when its buck converter is enabled and the digital core is active, thus ensuring a fixed discharge time and energy during each activation.

5.2.2 Quantization Issues from Harvesting AC

A by-product of using an AC input signal as the energy harvesting source is that there is a period during each 60 Hz cycle where the current transformer output does not meet the minimum LTC3588 input voltage, so Monjolo cannot charge. This seemingly benign observation has deep implications for Monjolo's ability to map different load powers to distinct activation rates. To better explain this phenomenon we use a slightly stylized illustration of the charging cycles of two Monjolo sensors as shown in Figure 4(b). In this simulation, Monjolo #1 is connected to a slightly higher power load than Monjolo #2. Therefore, its input AC waveform is also slightly higher in magnitude. We also assume for the simplicity of the illustration that each Monjolo node transmits a packet during every activation cycle.

Both Monjolo nodes require three cycles of the 60 Hz wave to charge, with the higher power circuit connected to Monjolo #1 reaching 5.1 V slightly sooner. It follows that the data aggregator will receive Monjolo #1's packet equally sooner. Both capacitors discharge at the same rate (a rate which is a fixed property of the system) and finish discharging during a negative "off" half-cycle. At this point both Monjolos *must* wait until the next "on" cycle to start recharging. This delay means that although Monjolo #1's packet will consistently arrive moments before Monjolo #2's, *both will wakeup and transmit at the same rate*. The result is that all load powers that cause Monjolo to finish discharging in the same "off" period quantize to the same activation rate.

This quantization phenomenon has a more pronounced effect at higher load powers, where the LTC3588 requires fewer cycles of the 60 Hz wave to recharge V_{store} . In this operating regime, the amount that the primary load must increase in order to increase the input power such that one fewer 60 Hz cycle is required to meet the activation threshold is much larger than at much lower loads. Stated differently, the amount each charge cycle must additionally contribute to V_{store} to drop from 100 charge cycles to 99 charge cycles is much lower than the amount of extra energy required to drop from 10 cycles to 9. This manifests in the stair-step effect visible at higher load powers in Figure 5. For instance, as the load increases from 445 W to 465 W, the resulting Monjolo activation rate remains constant.

Figure 5 does not exhibit crisp steps, however, because of the consistent time duration and energy use during the Monjolo discharge phase. Without the enhancements described in Section 5.2.1, once Monjolo begins discharging, it continues to discharge until it reaches a point in time at which the rectifier is in an "off" state. That is, Monjolo uses any "on" state to prolong its wakeup duration and does not discharge C_{store} until the subsequent "off" state. This forces every wakeup interval to delay until it can take the next step up. Figure 6 shows that without the MOSFETs, the quantizing effect is significantly worse. By satisfying the consistent discharge requirement, the quantizing effect can be mitigated for many load power ranges, illustrating precisely why the consistent discharge requirement is necessary for this design.




Figure 5: Frequency of Monjolo activations over a range of loads. As the primary load’s power increases from zero to 475 W, the activation rate increases from zero to 8.5 Hz. Although not linear, this monotonically increasing relationship between activation rate and load power suggests that Monjolo may be a viable approach to metering a range of AC loads.

Using half-wave rectification exacerbates the quantizing effect because it increases the duration of the “off” cycle. In practice, however, we notice that using half-wave rectification only displays pronounced steps at a load power greater than 270 W, and the effect on percent error is relatively small – at most $\pm 5\%$ of the true value.

5.2.3 Wireless MAC Layer

Monjolo operates with a simple MAC protocol running in a star topology. Each node simply transmits when it wakes up to an always-on data aggregator. We choose not to use a more advanced MAC layer because Monjolo is unable to support any existing MAC protocol that relies on scheduling wakeups between nodes. This limitation exists for two primary reasons. First, Monjolo has no backup energy storage to supply an on-demand wakeup. Therefore, it cannot activate at an arbitrary time as required by scheduled MAC protocols. Second, Monjolo has no means of keeping a sufficiently accurate timer active to generate these on-demand wakeups. Since the microcontroller is disconnected from power between activations, all timing information is lost.

This limited MAC layer is acceptable for the Monjolo application. The new approach to sensing advocated by the Monjolo principle – essentially a modern interpretation of the well-established pulse counting method used by utility meters – trades accuracy and reliability for simplicity of the node. Further, because Monjolo takes steps to avoid saturating the wireless channel, and the impact of packet loss is only a transient reduction in temporal precision and not a reduction in overall accuracy, a MAC protocol that ensures reliable or collision-free data transmission is unnecessary.

5.3 Data Aggregator

To collect packets from the Monjolo sensors, we use a Raspberry Pi (RPI) [5] embedded computer with a custom TI CC2520 radio shield. The CC2520 is an IEEE 802.15.4 compatible radio and allows the RPI to receive packets from Monjolo nodes. Using the RPI as a base station provides a platform capable of timestamping, processing, and forwarding the data streams from Monjolo sensors.

To estimate the power of the load under measurement, the RPI listens for all Monjolo packets. Upon reception, it computes the elapsed time since the last packet and the difference in the counter value, and feeds that into the model that relates load power to activation rate. The power estimate is timestamped and transmitted to a server for storage, aggregation, and further processing.




Figure 6: Quantized Monjolo activation rate versus load power. If the LTC3588 energy harvester is permitted to continue to supply energy to C_{store} during the discharge phase of operation, then the quantizing effect inherent in the system is substantially more pronounced. Particularly at high load powers, the steps span a nearly 50 W range, which dilutes precision and introduces $\pm 5\%$ error.

5.3.1 Load Power Model

We use a simple model based on data from Figure 5 to convert activation intervals to a load power estimate. Rather than fitting a polynomial or logarithmic function to the curve, we use a lookup table that encodes Figure 5’s data. This table can be locally stored on a Monjolo node. The sensor can then deliver the table to a data aggregator or the aggregator can download it from a server using the sensor’s identifier.

6. EVALUATION

This section evaluates how well Monjolo performs as an energy meter. We also evaluate the effect of varying loads, changing conditions, and noisy RF environments on Monjolo’s accuracy. Next, we show that running Monjolo adds very little overhead to the load under measurement. Further, we investigate the cost of producing Monjolo sensors at scale. Then, we examine the range of load powers that Monjolo supports. Finally, we show two application-specific realizations of the Monjolo sensor.




Figure 7: Calibration testbed. A control board manages nine circuits of light bulbs and power resistors with varying wattages. The values are chosen to allow for 1 W steps in load power. Light switches provide a hard shutoff, but during experimentation, wattages are selected via relays controlled by an Epic Core on the control board shown in the bottom left. The load is independently metered by a PLM-1LP for ground truth.

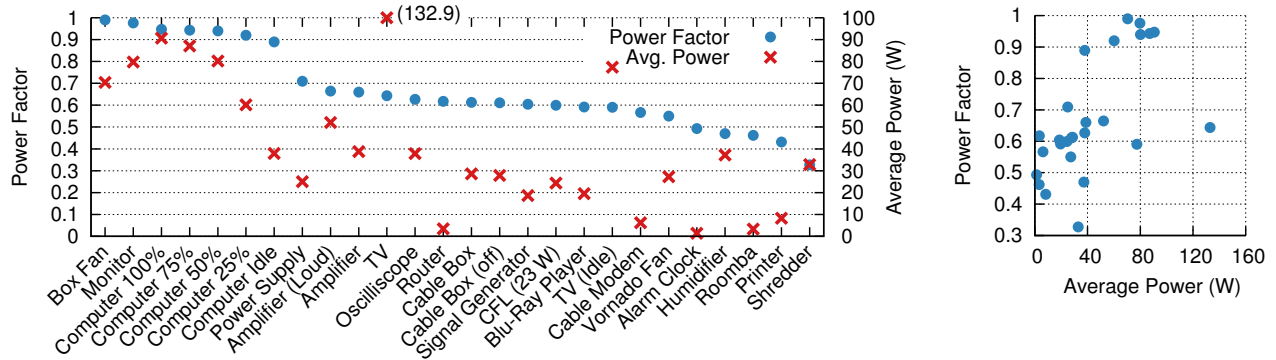


Figure 8: Power factors and average power for a variety of household and laboratory loads. Higher power loads (>75 W) generally have near unity power factors due to mandated power factor correction. We find that 0.6 is a common power factor for a range of common devices.

6.1 Evaluation Methodology

To test Monjolo accurately over a range of load powers, we use a measurement setup based on a PLM-1LP [4] power line meter and a custom programmable AC load. The programmable load, shown in Figure 7, uses nine resistive AC circuits, rated from one to 240 W, to allow the load to nominally range from zero to 480 W in 1 W increments. The load is controlled by its own Epic Core allowing us to program long-running experiments and have the load wirelessly transmit its approximate load over time. Due to the exclusively resistive nature of the light bulbs and resistors, this programmable load has a constant power factor of 1.0 and provide us with a stable testing platform. One note about the load: the 240 W bulb in practice draws 243 W, causing a slight gap in the possible loads of this setup. This gap is reflected in some of the figures.

The test setup does not allow for a programmable power factor due to the complexity in implementing such a system and our desire to accurately reflect the power factor conditions typically found in a house or office in our testing. Figure 8 shows the average wattage and power factor for several home loads. To test with varying power factors, we select devices that represent different points in the typical device power factor space. This allows us to evaluate Monjolo’s accuracy with respect to changes in power factor.

We meter all experiments in series with the PLM-1LP meter, allowing us to collect ground truth voltage, current, wattage, and power factor measurements in real time for all loads.

6.2 Monjolo as an Energy Meter

To determine how well the Monjolo system functions as an energy meter, we run it with different loads while simultaneously collecting ground truth data.

6.2.1 Resistive Loads

Fully resistive loads, or loads with a unity power factor, are characterized by phase-synchronized sinusoidal voltage and current waveforms, as shown in Figure 10(a). These loads draw exclusively real power, meaning the magnitude of apparent power is equal to real power. Our programmable load uses resistive circuits and is one such unity power factor load.

Using this software-programmable resistive load, we run an application that randomly selects power values for the AC load and switches after random intervals. This highlights Monjolo’s response to different load powers and a range of power transitions. Figure 9 shows the results of a 30 minute run of this application. Monjolo tracks ground truth very closely when the load is constant with an error of only 1%.




Figure 9: Accuracy of Monjolo’s power measurements over a range of different loads. When the load is constant, Monjolo’s error is very low at only 1%. When the load changes, Monjolo’s relatively slow response time causes sharp instantaneous error while the data aggregator waits to receive a packet informing it that the load has changed. When the load is very small, as it is at $t = 500$ s, Monjolo is unable to transmit and the resulting error is quite large, although a less naïve power estimation algorithm would be able to correct for this. Ignoring the few sections where Monjolo is drastically incorrect, Monjolo shows an average error of 3.7% across a randomly varying resistive load.

During times that the load is in transition, Monjolo cannot respond quickly enough to communicate these changes. Hence, the load estimation error spikes at these transition points. These spikes typically last for a second or so, with error duration largely dependent on the new wattage. At points where the load transitions to a wattage below the harvesting threshold, Monjolo cannot transmit a packet and the error spikes. Because our simple algorithm does not use the sudden lack of packet receptions to make assumptions about what the load “must” be, transient errors can spike above 100%. This is quickly rectified when the next packet is received.

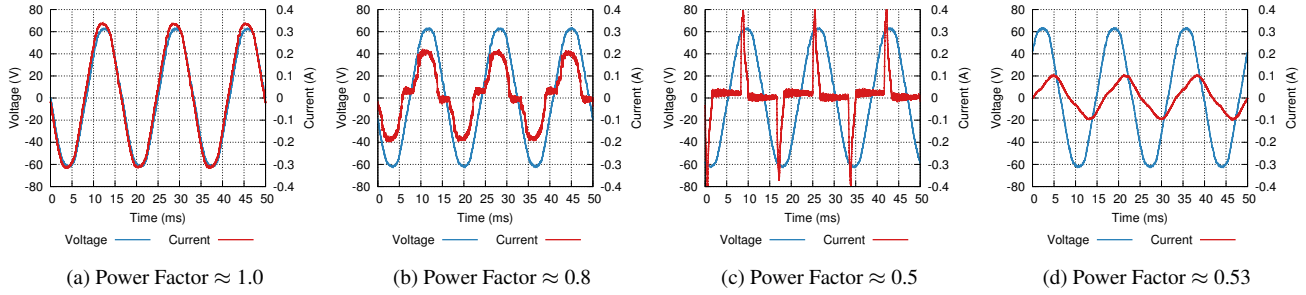


Figure 10: Current and voltage waveforms of four different loads with varying power factors. (a) A resistive load with its phase-synchronized current and voltage waveforms. (b) and (c) Non-unity power factor loads due to harmonic components. (d) A reactive load due to a large displacement between the voltage and current waveforms (and some smaller harmonic content).

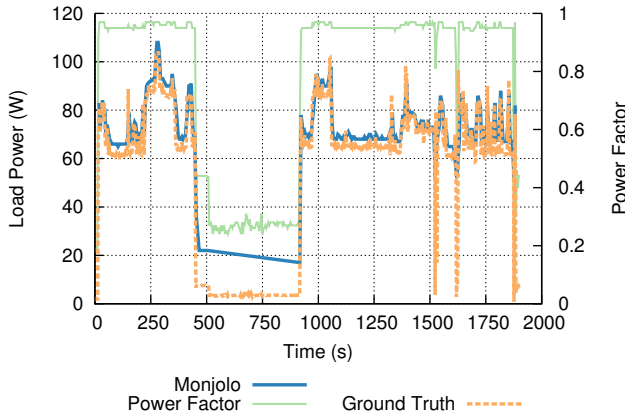


Figure 11: Power measurements of a running desktop computer. Monjolo is able to track the power draw with a slight error in magnitude when the power factor is close to 1.0. When the power factor drops to 0.3, the load power also drops, preventing Monjolo from being able to charge. While the error is larger with this non-unity power factor load than a purely resistive one, Monjolo is off by a relatively constant factor. This still allows Monjolo to clearly identify the load’s mode of operation over the course of the experiment. With some notion of the power factor, we expect Monjolo could more accurately estimate power in the future.

6.2.2 Reactive Loads

Typical loads are not fully resistive; they have a reactive power component and a non-unity power factor. This reactive component may be dominated by harmonic components (Figures 10(b) and 10(c)), displacement components (Figure 10(d)), or both. This non-unity power factor comes from energy being stored in the load during a portion of the 60 Hz cycle and resupplied to the source during another portion, or due to a non-linear (switching) regulator. This causes current to pass through the wire even if it does not pass through the load.

Any time current is flowing in the wire that Monjolo is sensing, it is able to harvest energy. This causes Monjolo to measure apparent power, the magnitude of the vector sum of real and reactive power, as opposed to real power itself [16]. Utility companies charge residential customers based on real power, making it the desirable quantity to measure.¹

¹Industrial customers are sometimes charged a fee associated with their reactive power draws.

Energy companies encourage their commercial and industrial customers to maintain a near-unity power factor because reactive power causes added transmission losses without performing useful work. For the same reason that energy companies encourage unity power factors in their customers, Monjolo measures apparent power: current is nonetheless flowing through the supply wires.

This hardly diminishes the value that Monjolo provides. The report from Monjolo is accurate if the desired quantity is apparent power. Even if not, the table lookup model that Monjolo uses is based on a unity power factor. Apparent power is always greater than or equal to real power, depending on the magnitude of the reactive power, so Monjolo does not underestimate real power. Finally, power factor (PF) is defined as real power divided by apparent power, so Monjolo will therefore be incorrect by a value of $1/\text{PF}$. Mode transitions are still apparent so Monjolo’s time-in-mode reporting remains accurate. Utilizing a single watch-point meter that measures the total power supplied to multiple AC loads being metered by multiple Monjolo nodes may allow the user to interpolate the value of the scaling factor.

To illustrate these properties, we test Monjolo while running a typical desktop computer. Figure 11 shows the actual and estimated power traces while the computer transitions through several operating modes. Between 250 s and 350 s, the CPU is fully active. At time 450 s, the computer is put into suspend mode until time 900 s. After it wakes up, it again runs at 100% and then after an idle period performs some normal activity.

During the periods in which power factor is nearly (but not exactly) unity, (i) Monjolo is off by an average of roughly 6% and (ii) Monjolo exclusively overestimates the load. During the period in which power factor is near 0.3, the error increases dramatically but the transitions into different operating modes are clearly distinguishable, and Monjolo continues to exclusively overestimate.

To further examine the effects of reactive loads, we run Monjolo against ground truth for four other loads, as shown in Figure 12. These loads exhibit a variety of different real-world power profiles. Monjolo is able to track the actual power draw of the various loads but continues to overestimate.

To illustrate the ability to accurately reconstruct the power trace given a single watch-point, we conduct the computer experiment again. This time, however, we run a 75 W resistive load on a circuit in parallel with the computer. Both the computer circuit metered by Monjolo and this additional 75 W load are measured by the PLM. Because the PLM measures the power factor of both circuits combined, we calculate power factor from their weighted average. Figure 13 shows that the total ground truth power tracks directly with Monjolo’s time-in-mode reporting and that the magnitude of the laptop draw can be calculated from the graph.

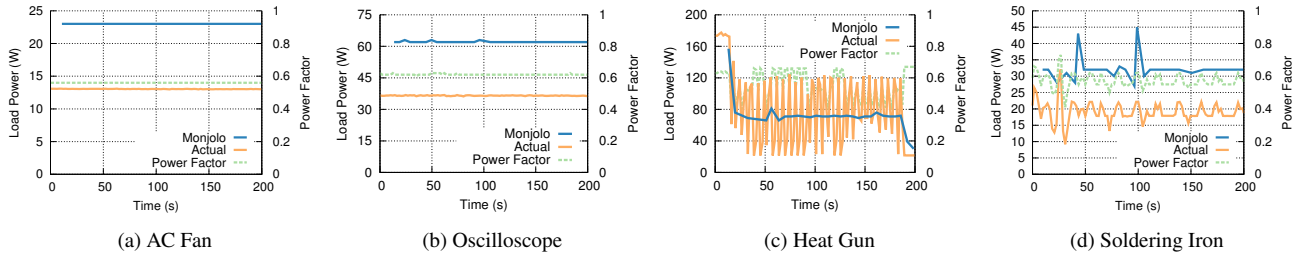


Figure 12: Monjolo performance on loads with non-unity power factors. Each figure shows the ground truth power and power factor measurement of the load, and Monjolo’s estimate of the power. If the power factor remains constant, Monjolo estimates the power with a constant offset, as shown in 12(a) and 12(b). The heat gun in 12(c) is an interesting case when the high frequency oscillations cause Monjolo to estimate an average value. The fluctuating power factor in 12(d) causes Monjolo to sharply overestimate, but it quickly recovers. While non-unity power factors cause Monjolo to lose accuracy, we believe the magnitude offset can be compensated for with other mechanisms.

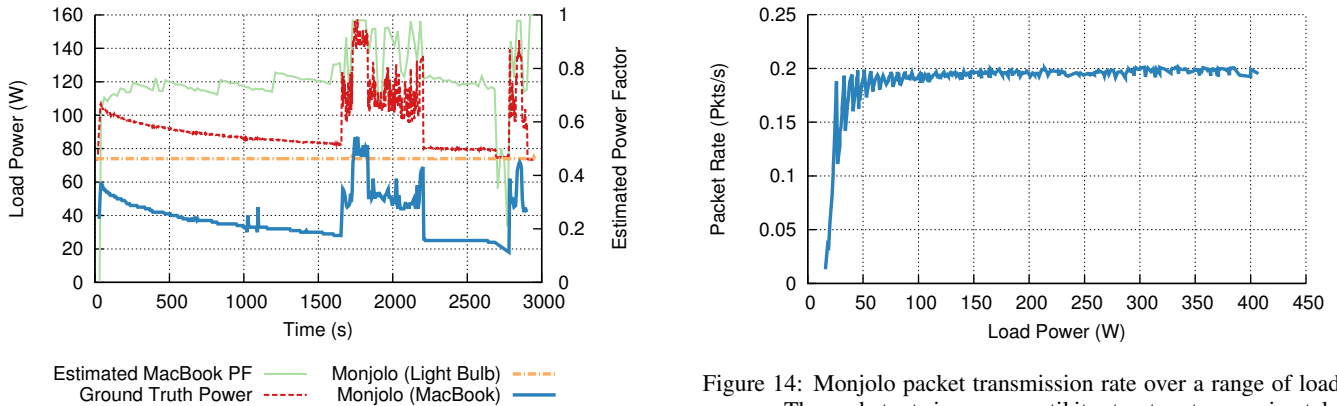


Figure 13: Estimating real power of a laptop power supply from Monjolo estimates and total system power. A MacBook power supply and a 75 W light bulb are plugged into two Monjolo sensors. As described in Section 6.2.2, Monjolo correctly tracks the MacBook’s power curve, but with an offset magnitude. The Monjolo sensor with the light bulb is accurate, however. The central monitor knows the correct total power and the power as measured by the two Monjolo sensors. From this data it can derive a more accurate estimate of the MacBook’s power draw. We claim that Monjolo sensors, along with a circuit level meter, can be used to compensate for Monjolo errors with non-unity power factor loads.

Figure 14: Monjolo packet transmission rate over a range of load powers. The packet rate increases until it saturates at approximately one packet every five seconds. The oscillations occur due to the increasing wakeup frequency interacting with the static 5 s rate limiter. This limit on the transmission rate ensures that a single Monjolo sensor will not saturate the wireless channel. While we use a static rate limit, other incarnations could employ feedback from the data aggregator to dynamically adjust the maximum rate.

6.2.3 Rate-Limiting Packet Transmissions

Timely and successful packet transmissions are essential to Monjolo’s operation. To enable more than one Monjolo to co-exist in a single collision domain, each node must not saturate the shared wireless channel. To verify that our RC timer successfully limits packet transmission rate, we measure the rate over a range of primary load powers. The results are shown in Figure 14. As expected, lower wattages cause Monjolo to transmit infrequently, but as the load power increases, the packet rate also increases until it is capped at one packet every five seconds.

6.3 External Effects on Accuracy

In prior sections, we evaluate Monjolo’s maximum attainable accuracy due to the physical properties of the system and the interactive nature of real, reactive, and apparent power. In this section, we relax the ideal environmental conditions and evaluate how Monjolo performs under more realistic and varied situations.

6.3.1 Temperature Changes

Monjolo depends on a consistent relationship between activation rate and load power to accurately estimate one from the other. Capacitors tend to slightly change their leakage and capacitance values with temperature. If significant, these physical changes would affect Monjolo’s activation rate, requiring a temperature-dependent calibration procedure to ensure Monjolo remains accurate across a range of field conditions.

To test the temperature dependence of Monjolo, we build a temperature control housing containing the Monjolo power supply, a temperature sensor, and a temperature control gun. Figure 15 shows this setup with the top of the housing removed. Using this setup, we accurately control the temperature of the Monjolo power supply under test conditions and determine its temperature dependence.

Figure 16 shows how system temperature affects Monjolo. Over a temperature range spanning 21°C (room temperature) to almost 60°C, there is a slightly downward trend in the power estimate, so Monjolo slightly underestimates power at high temperatures. This estimate, however, accounts for a maximum error of 2%. Given the low error, and the power, time, and cost overhead required to add and sample a temperature sensor, we decide that the temperature effects are not significant enough to warrant adding temperature compensation circuitry or software to Monjolo.




Figure 15: Temperature sensitivity test setup. An instrumented version of Monjolo's power supply is housed in a temperature-controlled enclosure. A fan at the rear ensures constant airflow and an analog temperature sensor between the power supply and the fan measures the temperature near C_{store} .




Figure 16: Effect of changing temperature on Monjolo accuracy. We fix the primary load power at 231 W and run Monjolo while slowly increasing the temperature of the system, including the storage capacitor. This simulates the heating that the system might experience in practice. We observe 2% error over a range of 40°C.

6.3.2 Packet Loss

Monjolo has the potential to operate in a noisy RF environment. Because Monjolo's fundamental ability to sense, and its metering resolution, are predicated upon the data aggregator successfully receiving packets in a timely manner, we evaluate the effect that packet loss has on Monjolo's measurement accuracy.

Figure 17 shows Monjolo's response to packet loss under conditions of 0%, 10%, 30%, and 50% loss rates, as well as the ground truth that the PLM provides. The graph shows that while the most accurate Monjolo report comes from the dataset with 0% packet loss, even 50% packet loss allows the system to track the load very well. The most significant errors occur when the load pulses rapidly and the dropped packet(s) occur during that pulse. During periods of constant load, packet loss has no effect on accuracy as it does not affect average reported power. When packets are dropped as the load changes, temporal granularity is lost but overall average accuracy is maintained.

6.3.3 Wireless Channel Contention

We take an analytical approach to examining the effect of the MAC protocol described in Section 5.2.3 on channel contention and packet loss. Figure 14 shows that the packet transmission rate of a node saturates at one transmission per five seconds for metered loads greater than 50 W. With N nodes metering P power, the worst case (highest collision probability) occurs if P is greater than this cutoff and N nodes each transmit once every five seconds.

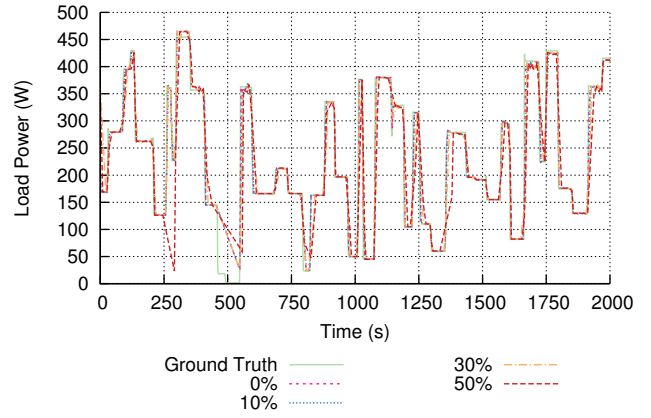


Figure 17: Effect of packet loss on Monjolo's accuracy. To test how Monjolo performs in a lossy and noisy RF environment, we simulate the effect of varying degrees of packet loss. By manually removing packets to simulate 10%, 30% and 50% packet loss, and then running the remaining data through our model, we see that Monjolo is still able to track the actual power draw of the load. Even at 50% packet loss, most transitions are measured with only a few high frequency power changes going undetected. While such high packet loss rates are not expected, this demonstrates that Monjolo is robust to communication failures.

Given $N/5$ transmission per second, what is the probability that two packets collide? Given that a Monjolo takes roughly 1 ms to transmit a packet, this results in an average channel duty cycle of $N/5000$. For a network of 100 nodes, there is a 2% probability of a random collision. This can also be further reduced by lowering the capacitor compare value in software (at the expense of lower temporal resolution).

The situation is worse when nodes become synchronized. Figure 4(b) shows that after a discharge, the charging start time is not random. Multiple nodes will be slightly synchronized by starting their charge at the onset of a positive 60 Hz cycle. Even if this effect increases the collision probability by a factor of four, the probability still remains acceptably low.

6.4 Energy Cost of Measuring

While no system can monitor power draw with zero overhead, Monjolo adds a much smaller power overhead than today's commercial power meters. In particular, when a load is completely off Monjolo draws no power due to its load current based energy harvesting technique. In contrast, commercial meters draw between 180 mW and 1.51 W with no load, as shown in Table 1.

To estimate the power overhead of Monjolo with a load attached, we model Monjolo's overhead at a given load power as:

$$P_{overhead} = f_{wake} \cdot P_{wake} \cdot t_{wake} + P_{loss}$$

P_{wake} is constant at approximately 66 mW (20 mA at 3.3 V). t_{wake} , the time the output voltage regulator is active and the microcontroller is powered on for, is also constant (as required by the system) at 58 ms. This is itself a function of P_{wake} and the capacitance of C_{store} . P_{loss} is the efficiency power losses of the system, including the overhead of the coil, rectifier, and regulators.

The power overhead of Monjolo scales with primary load power as the activation rate increases. This makes Monjolo's operation power-proportional – certainly a desirable trait.

Power Meter	Overhead (mW)	
	0 W Load	60 W Load
Belkin Conserve Insight*	440	560
Ensupra PM0001*	280	360
Kill A Watt PS-10*	180	690
Kill A Watt Wireless	670	720
SmartHome iMeter Solo	810	820
UPM EM100*	240	280
Watts Up? .net	1510	1590
Monjolo	0	≈ 4

Table 1: Power draw of several commercial power meters and Monjolo with no load and with a 60 W load. Meters marked with an asterisk have no ability to transmit their readings, whether wired or wirelessly. Due to the active circuitry and displays common in commercial plug-load meters, these devices consume nonzero power with no load. Monjolo, however, remains completely off when no load is attached. At 60 W the power draw of Monjolo is less than 10 mW (granularity of our measurement system) but draws approximately 4 mW based on the theoretical model.

6.5 Costs Analysis

One important factor in considering the scalability of this system is unit cost. If the cost is prohibitively high, then high-density deployments will be unlikely.

The unit cost of Monjolo is outlined in Table 2, broken down by major subsystems. Monjolo was developed using the Epic Core. However, the MSP430F1611 found in Epic is unnecessarily powerful for the simple Monjolo application. The MSP430G2152 is sufficient for Monjolo, and costs \$8.58 less.

As a means of comparison, Table 3 shows the unit cost of some other power metering options. Even with a 25% markup Monjolo would be price competitive with existing options, and is substantially more affordable than similarly-featured devices.

6.6 Parameter Tuning

Many factors affect Monjolo’s accuracy, measurement range, and overhead. The current transformer’s turns ratio, rectification scheme, and capacitance values can all be tuned to optimize for different desired characteristics.

An obvious limitation of the Monjolo design is the minimum load requirement. As Figure 8 shows, many household loads are below the 17 W minimum observed in Figure 5. This minimum is due to the LTC3588 minimum voltage threshold. In order for the input boost converter to operate, the rectified wave must meet a minimum voltage. The minimum primary load is therefore fully configurable in turns ratio. By selecting a higher turns ratio, we can cause low wattages to generate higher voltages than at the 1:500 ratio that we use to evaluate Monjolo.

Figure 18 shows the range of wattages from 1.15 W to 20 W generating wakeups by means of a 1:50 turns ratio. This 1.15 W minimum could be decreased further but already represents the ability to meter all of the devices listed in Figure 8. The alarm clock draws 1.37 W – and has a power factor of 0.49 – meaning the apparent power of 2.8 W meets the necessary threshold.

There are two reasons why a user might desire a lower turns ratio like the one on which Monjolo was evaluated. First, a lower turns ratio will generate fewer wakeups and therefore cause a lower power overhead. Second, Figure 5 shows the quantization error increasing with increasing wattage. Shifting the transfer function to lower wattages will be at the expense of relatively higher quantiza-

Component	Cost	Component	Cost
MSP430G2152	\$1.72	CC2420	\$4.20
LTC3588	\$3.58	Harvesting Coil	\$5.02
FRAM	\$1.05	Storage Caps	\$4.60
Other Components	\$1.17	PCB	\$0.97
Packaging	\$5.23		
Total			\$27.54

Table 2: Per unit costs of Monjolo by major component (in quantities of 1000). A majority of the cost comes from the energy harvester, coil, and capacitors. Very little cost is spent on processing due to the limited Monjolo application complexity.

Device	Cost	Device	Cost
Monjolo	\$27.54	iMeter Solo	\$39.99
Kill-A-Watt	\$29.95	Watts up?	\$95.95
Kill-A-Watt Wireless	\$34.99	Watts up? .Net	\$235.95

Table 3: Cost comparison of plug load power meters. Monjolo is more affordable even than meters with smaller feature sets.




Figure 18: Monjolo activations over a range of low power loads with a 1:50 current transformer turns ratio. By changing the current transformer front-end, Monjolo can operate with primary loads as low as 1.15 W. This demonstrates that the minimum load for operation is not a fixed limitation but rather a tunable parameter of the Monjolo design.

tion error at high wattages. This inaccuracy will be counter-acted in some regards by the fact that a higher turns ratio will generate more wakeups and therefore provide higher granularity.

Turns ratio is not the only way to adjust the nominal activation rate or power overhead. Switching from half-wave to full-wave rectification provides higher granularity with higher power overhead, but does not affect runtime. Increasing the storage capacitance decreases wakeup frequency but increases runtime without affecting power overhead.

Monjolo’s sample rate as observed by the data aggregator is tunable by adjusting R_{timer} and C_{timer} . We set this rate at 0.2 Hz, but it could be adjusted upward for higher fidelity data collection. This parameter is also software tunable as each sensor can tune the V_{timer} level at which it transmits a packet. For comparison, many of the commercial meters we evaluate sample at 1 Hz, although some sample slower and the Kill-A-Watt Wireless samples once every two minutes.

All of these parameters are user-selectable. Whether optimizing for accuracy, measurable load range, or power overhead, Monjolo is configurable for target load power.



Figure 19: Monjolo sensor in plug-load form. The case is open so that the female socket is visible on the left and the male plug is visible on the right. The energy-harvesting power supply is visible in the lower right PCB. The Epic Core is attached to the opposite side of the same PCB.

6.7 Monjolo Realizations

After evaluating the trade-offs of the Monjolo design space, we select a particular design point to realize Monjolo in a more convenient form factor for real-world deployments at scale. Figure 19 shows a plug-load version in its case and Figure 20 shows the split-core version that can be clipped into a circuit breaker panel to unobtrusively measure a circuit. Both designs use full-wave rectification, a 500 μF storage capacitor, and a 2.2 μF timer capacitor. The plug-load version has a current transformer turns ratio of 1:67 and the split-core version has a 1:3000 ratio. These two realizations allow us to cover nearly all loads present in a building either directly, or aggregated at the circuit level.

7. DISCUSSION

In this section we explore the limitations of Monjolo, present possible solutions, and offer ideas for future improvements.

7.1 Limitations

While Monjolo is effective in many different situations, it does have a number of limitations. Perhaps the most significant limitation is that it fundamentally lacks the ability to measure fine-grained data or accurately measure low power factors.

7.1.1 Lack of Fine-grained Data

Due to the nature of Monjolo’s power metering scheme, it is unable to collect precise, high-resolution data about the load. Traditional power meters record voltage, current, wattage, power factor and other values at relatively high frequencies, often 0.5-1 Hz. While this makes Monjolo unsuitable for scientific or revenue-grade metering, it remains well suited to typical home or long-term measurements, in which high resolution data are unnecessary, but efficiency, cost, and deployability are paramount.

7.1.2 Dependence on Power Factor

Because Monjolo harvests from any current in the wire to which it is connected, it measures apparent rather than real power. We offer two ideas on how to correct for this discrepancy if needed. First, as described in Section 6.2.2, a global meter that is capable of measuring the real power of an entire circuit may be used in tandem with Monjolo’s time-in-mode reporting to derive the real power of each individual plug load. Alternately, if Monjolo could estimate the power factor of the load, it could use that information to derive the real power component of the apparent power.




Figure 20: Monjolo sensor in split-core form. (a) The split-core current transformer with the energy harvesting power supply attached to the bottom. (b) The Monjolo node installed in a circuit breaker panel where it is easily clipped onto the wire running to a breaker to measure an entire circuit. A clip-on, energy-harvesting energy meter dramatically changes the panel-level energy metering landscape, enabling unprecedented energy visibility and installation simplicity.

7.2 Early Shutdown

As described in Section 5.2.1, fixing the duration of each Monjolo activation is crucial for accurately estimating the rate. Besides disabling the input and letting the low voltage threshold disable the output regulator, one way to assure a consistent discharge is to have the microcontroller time its runtime and shut itself off after a fixed period. This would allow the software more control over wakeup duration and timing. Common energy harvesting ICs, including the LTC3588, do not support this type of self-shutdown operation, however. Another approach might use a digitally-controlled, non-volatile, analog threshold.

7.3 Best of Both Worlds

Monjolo offers a new entry in the design space of energy meters. Typical energy meters (Table 1) trade standby power for increased metering capabilities and sampling rates by using active metering circuitry and an AC-DC power supply. Monjolo eliminates standby power at the cost of metering fidelity by using a harvesting scheme that can only operate when a load is present and drawing power. Perhaps, by combining these techniques, a new meter could achieve the best of both worlds. If a typical active energy meter added an inductive coil to its AC front-end, it could potentially power gate itself when no load was present. When current started flowing to the load it could reactivate itself, and use traditional energy metering techniques to take and report high accuracy measurements.

8. CONCLUSION

The need for energy submetering in buildings is crucial and growing, but current approaches that employ expensive circuitry, require intrusive installations, and draw unnecessary standby power leave much to be desired. Against this backdrop, we offer Monjolo: a power-proportional, energy-harvesting approach to energy-metering. With Monjolo, the energy-harvesting power supply is the sensor, removing the need for bulky AC-DC power supplies and active measurement circuitry. Metering is accomplished by counting and timing Monjolo activations, which are proportional to the load’s power draw. The simplicity of this design makes it well suited to both a plug-load meter and a split-core design that can easily clip onto a wire in a circuit breaker panel. Both of these configurations lower the bar to making ubiquitous submetering a reality.

9. ACKNOWLEDGMENTS

Special thanks to Shane DeMeulenaere, Nathan Heidt, and Sarah Spitzer for implementing a Monjolo proof-of-concept that demonstrated the basic principle, Andrew Robinson for his work on the data aggregator hardware and drivers, the students of Lab 11 for enduring days of toggling lights and clicking relays during calibration testing, the anonymous reviewers for their insightful feedback, Anthony Rowe for shepherding this paper, and Rodrigo Fonseca for suggesting the Monjolo name. This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA. This material is based upon work partially supported by the US Agency for International Development, the National Science Foundation under grants CNS-0964120 and CNS-1111541, an NSF Graduate Fellowship, and generous gifts from Intel, Qualcomm, and Texas Instruments.

10. REFERENCES

- [1] Cooper Power Systems Energy Harvesting Power Supply. http://www.cooperindustries.com/content/dam/public/powersystems/resources/library/1150_GridServer/B115009057.pdf.
- [2] EnOcean Application Note 308: Current Sensor – Power Line Monitoring for Energy Demand Control. http://www.enocean.com/fileadmin/redaktion/pdf/app_notes/AN308_CURRENT_SENSOR_Jan09.pdf.
- [3] Kill-A-Watt Wireless. <http://www.p3international.com/products/consumer/p4220.html>.
- [4] PLM-1LP. http://www.epd.com/power_meters.html#PLM-1.
- [5] Raspberry Pi. <http://www.raspberrypi.org/>.
- [6] Smarthome iMeter Solo. <http://www.smarthome.com/2423A1/iMeter-Solo-INSTEON-Power-Meter-Plug-In/p.aspx>.
- [7] Watts Up? .Net. <https://www.wattsupmeters.com/secure/products.php?pn=0&wai=0&spec=1>.
- [8] Analog Devices. ADE7753 Single-Phase Multifunction Metering IC. http://www.analog.com/static/imported-files/data_sheets/ADE7753.pdf.
- [9] Analog Devices. ADE7757 Energy Metering IC. http://www.analog.com/static/imported-files/data_sheets/ADE7757.pdf.
- [10] CR Magnetics. CR2550 Series Low Cost Remote Current Indicator. <http://www.crmagnetics.com/Products/Assets/ProductPDFs/2550.pdf>.
- [11] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler. Energy metering for free: Augmenting switching regulators for real-time monitoring. In *Information Processing in Sensor Networks*, pages 283–294, 2008.
- [12] 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.
- [13] V. Gupta, A. Kandhalu, and R. Rajkumar. Energy harvesting from electromagnetic energy radiating from AC power lines. In *In proceedings of HotEmNets 2010*, June 2010.
- [14] 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.
- [15] D. Jung and A. Savvides. Estimating building consumption breakdowns using on/off state sensing and incremental sub-meter deployment. In *SenSys’10: In Proceedings of the Eighth ACM Conference on Embedded Networked Sensor Systems*, Nov. 2010.
- [16] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso. A platform for ubiquitous sensor deployment in occupational and domestic environments. In *IPSN ’07: Proceedings of the 6th international conference on Information processing in sensor networks*, Cambridge, Massachusetts, Apr. 2007.
- [17] Linear Technology. LTC3588 piezoelectric energy harvesting power supply. <http://cds.linear.com/docs/en/datasheet/35881fa.pdf>.
- [18] National Science and Technology Council–Committee on Technology. Federal research and development agenda for net-zero energy, high-performance green buildings, Oct. 2008.
- [19] National Science and Technology Council–Committee on Technology. Submetering of building energy and water usage: Analysis and recommendations of the subcommittee on buildings technology research and development, Oct. 2011.
- [20] J. Ortiz, Y. Noh, G. Saldanha, D. Su, and D. Culler. Towards real-time, fine-grained energy analytics in buildings through mobile phones. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 42–44. ACM, 2012.
- [21] I. Paprotny, Q. Xu, W. Chan, R. White, and P. Wright. Electromechanical energy scavenging from current-carrying conductors. 2012.
- [22] J. E. Petersen, V. Shunturov, K. Janda, G. Platt, and K. Weinberger. Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives. *International Journal of Sustainability in Higher Education*, 8(1):16–33, 2007.
- [23] N. B. Priyantha, A. Kansal, M. Goraczko, and F. Zhao. Tiny web services: design and implementation of interoperable and evolvable sensor networks. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 253–266. ACM, 2008.
- [24] T. Schmid, D. Culler, and P. Dutta. Meter any wire, anywhere by virtualizing the voltage channel. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pages 25–30. ACM, 2010.
- [25] T. Weng, B. Balaji, S. Dutta, R. Gupta, and Y. Agarwal. Managing plug-loads for demand response within buildings. In *Proceedings of the Third ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 13–18. ACM, 2011.
- [26] T. Wu and M. Srivastava. Low-cost appliance state sensing for energy disaggregation. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys ’12, pages 53–55, New York, NY, USA, 2012. ACM.
- [27] Q. Xu, I. Paprotny, M. Seidel, R. White, and P. Wright. Stick-on piezoelectromagnetic ac current monitoring of circuit breaker panels. 2013.