

Micro Power Meter for Energy Monitoring of Wireless Sensor Networks at Scale

Xiaofan Jiang, Prabal Dutta, David Culler, and Ion Stoica

Computer Science Division
University of California, Berkeley
Berkeley, California 94720

{fxjiang,prabal,culler,istoica}@cs.berkeley.edu

ABSTRACT

We present SPOT, a *scalable power observation tool* that enables *in situ* measurement of nodal power and energy over a dynamic range exceeding four decades or a temporal resolution of microseconds. Using SPOT, every node in a sensor network can now be instrumented, providing unparalleled visibility into the dynamic power profile of applications and system software. Power metering at every node enables previously impossible empirical evaluation of low power designs at scale. The SPOT architecture and design meet challenges unique to wireless sensor networks and other low power systems, such as orders of magnitude difference in current draws between sleep and active states, short-duration power spikes during periods of brief activity, and the need for minimum perturbation of the system under observation.

Categories and Subject Descriptors: B.8.2 [Performance and Reliability]: Performance Analysis and Design Aids

General Terms: Measurement, Performance, Design.

Keywords: Power, energy, meter, monitoring, scalable, dynamic range, embedded, wireless sensor networks

1. INTRODUCTION

Energy-efficiency has pervaded nearly every aspect of wireless sensor network (“sensornet”) research, from platform designs [11, 3, 6], to MAC layers [17, 10], to routing protocols [12, 13, 16], to applications [18, 5], and across a range of duty cycles [2, 15, 7]. What is missing is a method to empirically evaluate the energy-efficiency claims of this growing corpus of literature. On one hand, simple approximations of nodal energy usage derived from estimates of node duty cycle and communication rates [4] do not capture the low-level system power profile. On the other hand, simulators that extrapolate system macrobenchmarks – the large-scale, long-term, and system-wide behavior of a sensor network – from models based on microbenchmarks of a single node

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN’07, April 25-27, 2007, Cambridge, Massachusetts, USA.
Copyright 2007 ACM 978-1-59593-638-7/07/0004 ...\$5.00.

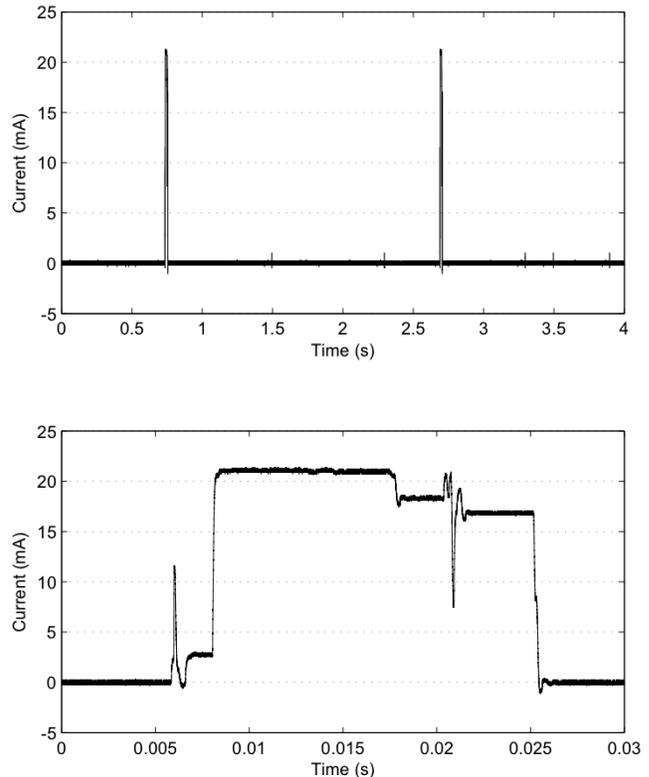


Figure 1: A typical nodal current profile consists of long periods of low-current sleep punctuated by short, periodic bursts of high-current activity.

cannot assure the accuracy of their generalizations [14].

Motivated by roboticist Rodney Brooks’ famous observation that “the real world is its own best model,” we suggest that what is now needed are empirical measurements of *in situ* energy usage *at scale* to calibrate existing models, characterize their variance, and validate their generality. However, characteristics unique to sensornets make their power and energy monitoring challenging. Figure 1 shows the current profile of a typical sensornet application. Long periods of low-current sleep are punctuated by short, periodic bursts of high-current activity.

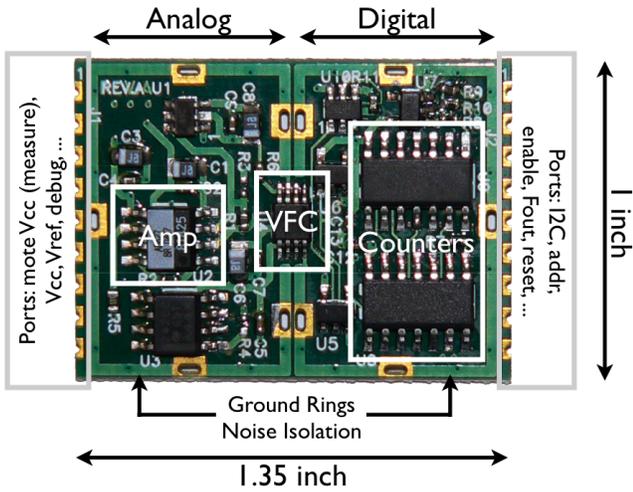


Figure 2: The scalable power observation tool (SPOT) consists of a sense resistor, amplifier, voltage-to-frequency converter, and two counters.

Since more than three orders of magnitude separate the current draw in the sleep and active states, it might seem reasonable to ignore the energy usage in the sleep state. However, because sensornets operate at very low duty cycles ranging between 0.1% to 1%, both the sleep and active states account for non-trivial fractions of the system power budget. This suggests that a metering system must have a *dynamic range that significantly exceeds three orders of magnitude* to capture sleep current with sufficient resolution.

Figure 1 shows that a node may be active for a short time, on the order of 20 ms, before returning to the sleep state. However, the node’s current draw is not constant during this 20 ms period. Rather, the current profile includes two large transients lasting tens of microseconds, four different levels, and oscillations during some level transitions. Such features in nodal current profiles are common and reflect the wide variety of system components and the sum of their various power states and state transitions. This suggests that *sampling rates approaching tens of kHz or even MHz* may be needed to faithfully capture these ephemeral features.

To prevent excessive perturbation of the system under test, the metering system should be minimally invasive: it should require low computational and storage resources from the host node, and it should be able to operate in a stand-alone manner with limited host interaction. Finally, *in situ* metering *at scale* requires small size and low cost.

Our solution to this metering problem is shown in Figure 2. This system, a *scalable power observation tool* (SPOT), enables *in situ* measurement of nodal power and energy with a dynamic range exceeding 10000:1 and a temporal resolution in the order of microseconds. SPOT accumulates current internally and exports digital I/O lines to enable/disable metering and calibration, analog lines to measure instantaneous current, and an I2C interface to read the accumulated current (energy) and reset the counter. The remainder of this paper presents more detailed design requirements and technical challenges, followed by SPOT’s design analysis and evaluation. We conclude with a discussion of the research enabled by this work and our future plans.

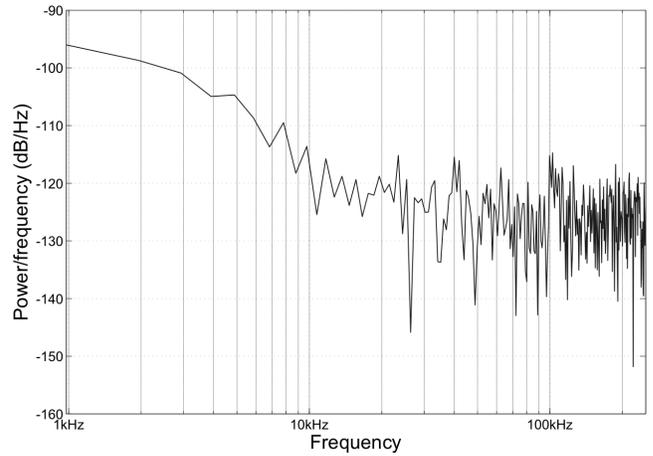


Figure 3: The power spectral density of a node current profile shown in Figure 1.

2. PROBLEM FORMULATION

This section presents the four basic requirements of our micropower meter for sensornet nodes (“motes”): dynamic range, sampling rate, perturbation, and ease-of-integration.

2.1 Dynamic Range

A sensornet node can exhibit a bewildering array of power profiles, depending on its application profile. For example, a simple *sense-and-send* application with a 0.5 Hz duty cycle might exhibit the profile shown in Figure 1. In contrast, a *sense-and-store* application may sample sensors in a few milliseconds every five minutes, buffer these sensor readings in RAM, write them to flash once a day, and attempt to upload the samples once a week. The current profile for such an application would be starkly different from that shown in Figure 1. Since SPOT is to be used in a testbed, we cannot assume a particular application profile *a priori*, implying the system needs to have a wide dynamic range spanning the entire spectrum of possible current draws.

2.2 Sampling Rate

Wireless sensor nodes are pulsing applications. Their pulse width, or the active cycle time, needs to be considered when determining the appropriate sampling rate. If the duration of an active pulse is shorter than the sampling rate of the energy meter, the energy in that pulse may be missed. In addition, the spectral content within a active pulse should also be sampled at a rate that satisfies the Nyquist rate.

While we can estimate the minimum active cycle time by observation, we cannot guarantee that the spectral content will be band-limited to a particular range. In our example, the width of an active cycle is approximately 20ms. However, it’s power spectral density (PSD) exhibits energy across the entire spectrum, and therefore it is necessary to bandlimit the signal with a low pass filter (LPF). The cutoff frequency for the LPF should be the highest frequency in the PSD that still contains significant energy. The cutoff frequency in our example is around 20kHz, as shown in Figure 3. This implies a minimum sampling rate of $20kHz \times 2 = 40kHz$ per Nyquist’s theorem.

Intuitively, we can see that to capture most of the energy

content in the $20ms$ active cycle in Figure 1, especially the oscillations at $6ms$ and $21ms$, we will need to sample at least every $0.1ms$, which implies a sampling frequency of $10kHz$.

2.3 Perturbation

Energy monitoring should not affect the actual energy consumption of the mote under study (hereby referred to as the device under test, or DUT). When the user of the energy monitor is separate from the DUT, accessing data from the energy monitor should not affect the energy consumption; when the user of the energy monitor is part of the DUT itself, then accessing the energy monitor should present a minimal perturbation to the DUT.

The energy monitoring device should not affect the measured power of the DUT. This implies that the energy monitor should be powered from a separate power source; however, if it must share the same power supply as the DUT, the point of measurement should be confined to the DUT.

The energy monitor should not require the DUT to perform extensive computation, if any. This implies that for the case when the user and the DUT are the same mote, a meter that can only provide power measurements is out of the question because to obtain an energy measurement, the DUT will need to constantly read from the meter to accumulate energy measurement. Instead, this computation should be offloaded to the energy monitor to minimize the energy and CPU usage of the DUT.

2.4 Ease-of-Integration

One of the goals for this project is for every mote on a network to be equipped with an energy meter. This suggests that the system needs to be easy to integrate, both electrically and mechanically, into a sensornet node. To be practical, the meter must be inexpensive, or perhaps comparable in cost to the sensornet nodes.

3. RELATED WORK

Many commercially available energy meters operate by measuring the voltage drop around a shunt resistor R tied to the power supply of the circuit under observation, as shown in Figure 4. The voltage drop is proportional to the current and can be multiplied by a voltage from a separate channel to obtain the true power usage. The voltage across the resistor is usually first magnified by an amplifier, and then undergoes an analog-to-digital conversion (ADC). The rest of the computations, such as dividing by the resistance and multiplying by the voltage, are usually done in the digital domain. This power computation occurs many times in a second (from a few Hz to MHz) and is summed in an integrator to obtain energy. The results are stored in registers and presented at the output either as PWMs (pulse width proportional to energy) or in digital form.

Oscilloscopes are often used to capture the power profile of a system at bench scales since their high sampling rates, often in the GHz range, can provide very fine temporal resolution. However, their high cost makes them unsuitable for large scale usage. But, even if cost were not an issue, oscilloscopes would not be suitable. First, most oscilloscopes are limited to milliamp-level vertical sensitivities, rather than the needed microamp levels. Although a pre-amplifier can be used, this adds cost and complexity. Second, they are not easy to integrate with typical sensornet nodes.

Sensornet-specific solutions have been proposed as well.

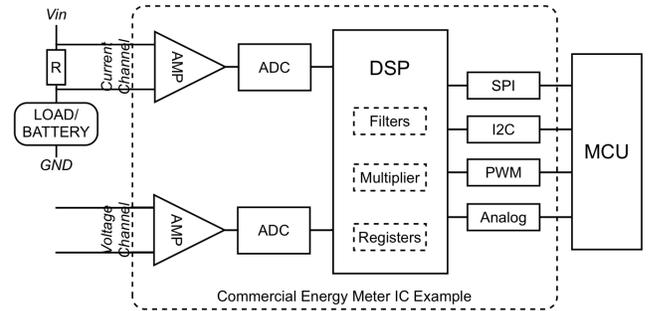


Figure 4: Architecture of a typical energy metering circuit.

For example, the DS2438 [8] is a battery monitor IC used in the Heliomote [6]. This is an 8-pin IC with a simple 1-Wire interface and integrated temperature sensing. However, the on-board ADC is only 10-bits, providing a theoretical maximum dynamic range of 1000:1, which is far below requirements for motes (10000:1). It samples at a frequency of 36.41Hz, also far below the 20kHz needed to capture interesting spikes.

Commercially available integrated circuits are not designed to meet our dynamic range and sampling requirements simultaneously. For example, battery monitoring ICs such as the Maxim Semiconductor’s BQ2019 are targeted towards long term measurement of batteries and provides low temporal resolution. Multi-function metering ICs such as the Analog Devices’ ADE7753 [1] have maximum dynamic ranges of around 1000:1, failing our vertical resolution requirement. Microchip’s MCP3906 [9], an energy measurement IC supporting the IEC 62053 international energy metering specification, offers a dynamic range of only 1000:1 and has offset currents that exceed node sleep currents by an order of magnitude, making it unsuitable for our purposes.

A very different approach, proposed by Shnayder et al., is to simulate power draw using an empirically-generated model of hardware behavior [14]. The simulator, called PowerTOSSIM, characterized its underlying model by instrumenting and profiling a single node. While PowerTOSSIM takes an important step toward providing better visibility into nodal power profiles, since its model is based on microbenchmarks of a single node taken in particular environment, it also raises several questions about the model’s generality: How representative is the node that was instrumented to calibrate PowerTOSSIM’s model? How does interaction with the physical environment shape energy usage? How do temperature variations affect leakage currents? How much variance occurs within a single node, and across different nodes, for the same operation? These questions can only be answered by instrumenting an entire network of nodes *in situ* and *at scale*.

4. SPOT ARCHITECTURE

In this section, we present the SPOT architecture and highlight some of its key features. A more detailed analysis is presented in Section 5. Our architecture consists of four stages: sensing, signal conditioning, digitization, and energy output, as shown in Figure 5. This configuration shows the DUT and the user of the energy meter to be the same mote,

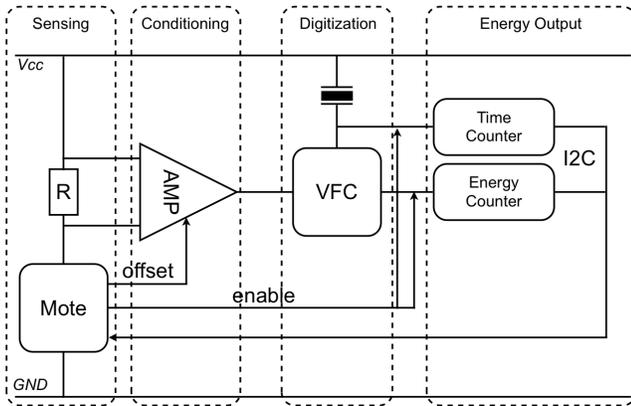


Figure 5: Architecture and Primary Components

which need not be the case. In the sensing stage, a shunt resistor is put in series with the mote (DUT), converting current to voltage. This voltage is proportional to the power consumption of the mote. In the conditioning stage, a differential amplifier and a low pass filter is used to amplify and band-limit the signal, respectively. In the digitization stage, a voltage-to-frequency converter (VFC) is used to convert voltage into a periodic wave with a frequency proportional to the input voltage, which is also proportional to the power. In the energy output stage, pulses from the output of the VFC are summed (i.e. integrated) in a counter to obtain energy measurements while a separate counter is used to keep a time-base; finally, the values of the counters are read back by the mote via I2C.

4.1 Differential Amplifier

While differential amplifier is simple in principle, it is more of an art to correctly use one in practice. Because we configured our shunt resistor at the high-side of the mote (see Section 5.1), the common-mode voltage is equal to V_{cc} (3.3V). This signifies that not only does our amplifier needs to have a high common-mode input range, we need to introduce a second voltage supply (e.g. 5.5V) to correctly biased the amplifier.

Due to variations in device sizing, amplifiers have offset voltages (i.e. output is non-zero even if input is zero) and it's different for every amplifier. This presents us with a calibration problem. To compensate, we included a calibration switch which zeros the input at startup and records the counter value (corresponding to the offset). This value is stored in the MCU and used in the calibration curve.

Because we are trying to amplify a very small signal, any noise is significant. We placed multiple RC filters (for different frequencies) around sensitive areas (e.g. voltage supplies of IC chips). Furthermore, we observed that the oscillator for the digital portion of our circuit introduces significant noise to the amplifier. To solve this problem, we carefully placed separate ground rings in the board to separate the analog portion of the circuits from the digital portion (see Figure 2).

4.2 Voltage-to-Frequency Converter

To achieve the desired dynamic range, we choose a novel approach of using a voltage-to-frequency converter as the

analog front end (AFE). It converts the signal from the analog to the digital domain. While VFC has a clocked input, the voltage to frequency conversion is entirely analog in the sense that any arbitrarily small voltage is integrated inside an analog integrator and will eventually trigger an output pulse. This essentially gives the VFC *infinite resolution, limited only by noise*. In contrast, traditional ADCs will fail to capture voltage levels smaller than the specified *bit* resolution. Please refer to Section 5.3 for details.

4.3 Energy Counter and Internal Timebase

To minimize MCU overhead, we use an internal 32-bit counter to accumulate the power readings to provide direct energy measurement. The counter is free-running at a maximum rate of 0.9MHz or 0.9 million counts per second. It overflows in roughly $\frac{2^{32}}{2^{20}} = 2^{12}$ seconds or about once an hour. This implies that the MCU will only need to read once every hour if long term energy is all the application needs. But it can read as fast as I2C bus speed allows to obtain a fine-grain power profile by differentiating the energy readings.

To find the energy consumption for a given time window, we can read the energy counter at the beginning and again at the end. For this to work, we need some way of keeping the elapsed time. This can be done in the MCU by using the MCU's timer. However, this introduces MCU overhead and the time will not be accurate because many cycles will have elapsed from `timer_fired()` to an actual read from the energy counter. Instead, we included another counter dedicated to be the time-base. This counter is triggered by the same oscillator (1MHz) that provides the base frequency to the VFC. We wired the I2C bus address such that one command can atomically capture the two counter readings at the exact same time. This method has the additional benefit of nullifying any jitter in the oscillator due to temperature.

4.4 Power Introspection

Because SPOT provides power and energy measurements without significant MCU perturbation, it is perfectly reasonable for a mote to use SPOT to monitor its own power and energy consumption, allowing the application to perform power adaptation.

5. DESIGN ISSUES AND TRADEOFFS

In this section, we revisit the SPOT architecture with an eye toward design issues and tradeoffs in the sensing, signal conditioning, digitization, and accumulation.

5.1 Sensing

As with most sensing systems, the first stage is the sensor itself. In the case of energy monitoring, the physical quantity we are trying to measure is power, which is a product of voltage and current. For the purposes of this paper, we assume that the voltage is fixed and known *a priori*, or can be measured separately. Of course, current can be measured in several ways. We chose to use a shunt resistor, which is one common method. A small resistor is placed in series with the power supply and the voltage across this resistor, which is proportional to the current, is measured.

There are several design considerations with this approach. The resistor can be placed between the mote and either the positive power supply or the negative power supply (ground). The former is called "high-side current sensing" while the

latter is called “low-side current sensing”. Low-side sensing is desirable because the differential voltage across the resistor is equal to the voltage measured, with respect to ground, at the connection between the resistor and the mote. This simplifies the amplification stage because there is no common-mode voltage. However, low-side sensing also creates a problem known as ground bounce – as the current draw fluctuates, the negative supply of the system also fluctuates. This is undesirable because many electronic components are sensitive to ground fluctuations. This is also why the resistance must be kept small since the magnitude of this fluctuation is proportional to the resistance.

High-side current sensing places the resistor between the positive supply rail and mote’s power input. By placing the resistor on the positive supply side, the voltage fluctuations are shifted from the negative supply side to the positive side. This is more desirable because most components are more resilient to fluctuations in the positive supply rail. However, this introduces common-mode voltage because the voltage on both sides of the resistor, measured with respect to ground, is non-zero. The presence of significant common-mode voltage can cause problems for amplifiers, as we discuss in Section 5.2.

We place a 1Ω resistor R between the positive supply rail (V_{cc}) and the mote’s positive supply in a high-side configuration, as shown in Figure 5. The value of 1Ω was chosen to limit the supply voltage fluctuation. Assuming a maximum current of $40mA$, typical of current mote technology, the maximum drop is $40mV$, which is reasonable.

5.2 Signal Amplification and Conditioning

The current draw of a typical mote, such as Telos [11], ranges from $2\mu A$ to $40mA$.¹ Using a 1Ω resistor, the minimum voltage needed to capture is $2\mu A \times 1\Omega = 2\mu V$, which is too small for signal processing. Therefore, we first amplify this signal using differential amplifier, as shown in Figure 5.

The gain of the differential amplifier is set such that the maximum input voltage of $40mV$, multiplied by the amplifier gain, is equal to the maximum input of the next stage and less than the maximum output of the amplifier. In this case, the input of the next stage is limited by the supply voltage, which is commonly $3.3V$ or $5V$. Therefore, a reasonable gain is $\frac{3.3V}{40mA} = 82.5$. The remainder of this section explores some of the design and implementation challenges of this stage.

5.2.1 Dynamic Range

The required dynamic range is one of the key challenges in our system. We cannot artificially increase the amplifier gain because it is limited by the maximum input range of the next stage. With a gain of 82.5 , a $2\mu V$ input translates to $0.165mV$ output, which is still quite small. At this point, we can either defer this problem to the next stage or try to alleviate it through some sort of signal processing.

One way to “decrease” the dynamic range requirement is to use a low pass filter (LPF) to lower the amplitude and widen the pulses. However, because LPFs are not energy preserving, they require software compensation and present timing difficulties during measurements. We chose to preserve the signal integrity and let the next stage deal with the dynamic range requirement.

¹ $2\mu A$ can be obtained by lowering the supply voltage below recommended values.

5.2.2 Common Mode Voltage

Amplifiers are composed of MOSFET transistors. To operate correctly, transistors need to be biased into the saturation region. The voltage level of the input signal, which is propagated through the gates of MOSFETs, needs to have a sufficient potential difference with respect to the drain and the source. The drains and sources are tied to the voltage source and ground respectively. This implies that the DC component – the common-mode voltage – of the input signal needs to be lower than the supply voltage by a small margin but higher than ground by a small margin as well. Because we chose high-side current sensing configuration, our common-mode voltage is the same as V_{cc} . Hence, we introduced a second power supply that delivers $5.5V$ to bias the amplifier.

5.2.3 Reference Voltage

The reference voltage of a differential amplifier is the baseline voltage to which input and output voltages are referenced. The reference voltage also plays a role in biasing the amplifier and must be chosen carefully. For example, for the amplifier we are using, if the common-mode voltage is around $4V$ and the supply is $5.5V$, the reference input needs to be greater than $3V$, as required by the specifications. The exact requirement on the reference voltage varies across different chips. It is desirable to experimentally determine the limits. Reference voltage also allows us to adjust or even cancel out the internal offset so that the output is always positive. The internal offset due to sizing differences and manufacturing variations may be positive or negative. Because the next stage may not cope well with negative voltages, it is desirable to offset this value so that the total offset is always positive. This does not eliminate the need for calibration because the total offset still varies across chips.

Additionally, a reference voltage can give us the flexibility to measure power in both directions (i.e. consumption and recharging). For example, if we set the reference voltage to $\frac{V_{cc}}{2}$ (see Figure 7), we can record the nominal (zero power in either direction) voltage by setting the input to zero and capturing the counter values (or rather how fast the counter increments). A reverse flow in current (e.g. recharging) will simply cause the counter to count slower than this nominal rate. This will cut the resolution in half but provide a signed rather than unsigned value for metering.

5.2.4 Input and Output Offset

Differential amplifiers consists of pairs of CMOS gates. Due to process variations, the sizing of these gates are usually not perfectly matched. As a result, there will be non-zero output even when the input is zero, called offset. There are two types of offset, input offset and output offset. Input offset is multiplied by the gain while output offset is directly added on the amplified signal. Because our application is concerned with very small voltages, even a small variation in offset will skew our results. Therefore we need to calibrate the amplifier in order to eliminate the offsets.

There are typically two ways to calibrate offsets. One method involves purely using analog circuits, but this method is expensive both in terms of cost and power. The alternative, and the method we choose, is to do it digitally. As seen in Figure 6, we added a digitally controlled switch to the resistor’s load. The switch is nominally at A, allowing normal measurement. During calibration, the mote toggles

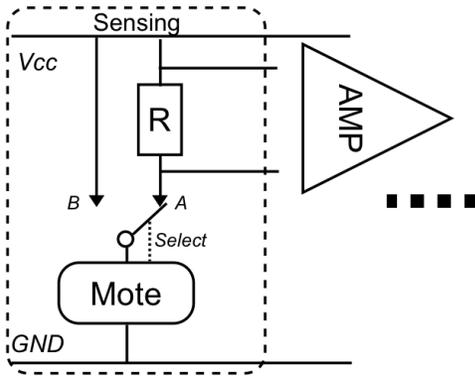


Figure 6: SPDT Switch for Calibration. In one configuration, current flow bypasses the shunt resistor, allowing the input offset to be measured. In the other configuration, current passes through the shunt resistor, allowing the mote current to be measured.

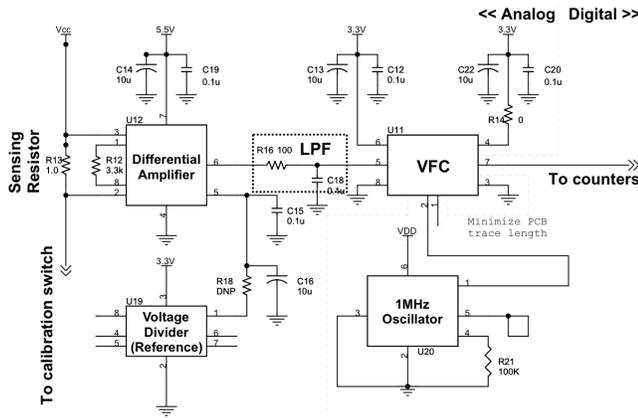


Figure 7: A simplified representation of the analog portion of schematic.

the switch line to B, which results in a zero voltage drop across the resistor. The mote then takes a reading at this configuration before returning the switch to A.

5.2.5 Effect of Noise

As discussed in Section 4.1, the differential amplifier is quite sensitive to noise. Figure 8 shows the effect of the oscillator on the amplifier before any filtering. To reduce noise, we have placed multiple filters around chip power supplies as seen in Figure 7, and separated analog circuits from digital circuits using slotted ground guard rings as seen in Figure 2. We also employ separate analog and digital ground planes which meet directly under the VFC. The guard ring slots can be used to attach separate metal shields to the analog and digital portions of the circuit, although we have not yet needed to do so.

5.2.6 Lowpass Filtering

As previously discussed in Section 2.2, energy content is present across the entire spectrum. To avoid false readings

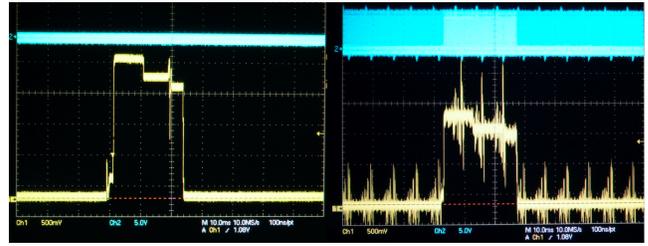


Figure 8: Noised introduced by the VFC oscillator requires analog lowpass filtering, ground guard rings, separate analog and digital ground planes, and decoupling capacitors to contain.

due to high frequency components in our signal, we cutoff the frequency content at the point where the energy content drops significantly such that it will not adversely affect our energy readings. In our system, this point occurs at 20kHz as seen in Figure 3. We place a single LPF between the differential amplifier and the VFC to achieve this purpose, as shown in Figure 7.

5.3 Digitization

Digitization is the stage that converts the analog signal to the digital domain, allowing digital signal processing (DSP) to be applied and eventually interfaced with a digital system (MCU). This is one of the central stages in most metering systems and is where some of the original signal information is lost due to the finite resolution of digital systems.

Analog-to-digital converters (ADC) are the most commonly used digitization device. ADC takes an input voltage and outputs a digital signal with a specified *bit* resolution. For example, a 12-bits ADC means that for an input range of 0-40mV, the minimum voltage level that can be captured is $\frac{40mV}{2^{12}} = 9.8\mu V$. Furthermore, 12-bits means there are 2^{12} discrete steps over the input range and any value between the steps will need to be rounded.

Our dynamic range requires at least 14-bits of resolution (Section 2.1). We considered several different designs:

- Internal ADC of the MCU is inadequate because they are usually only 12-bits. Furthermore, it incurs significant MCU overhead.
- 16-bit ADCs are slightly more expensive but not uncommon. However, one of our goals is to internally integrate power to obtain energy. If our power measurement (output of the ADC) is already in a pure digital format (e.g. parallel bus, I2C, SPI), we will need a compatible device to integrate the digital signal. This implies either a DSP or MCU. A separate DSP or MCU chip adds unnecessary complexity and cost. On the other hand, using the same mote as the integrating MCU would incur significant CPU and power overhead.
- IC chips that integrates 16-bit ADC with counters are essentially energy monitoring ICs. However, as discussed in Section 3, we have not found a single IC that satisfies all of our requirements.
- Voltage-to-frequency converter (VFC) is a low cost analog digitization device in the sense that the out-

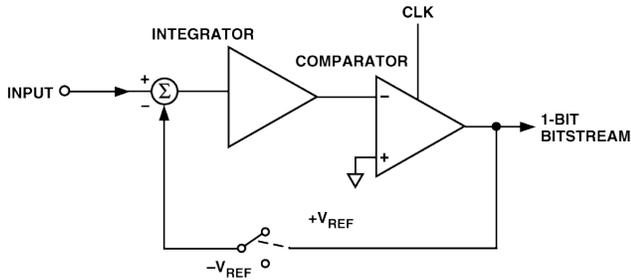


Figure 9: The architecture of a VFC.

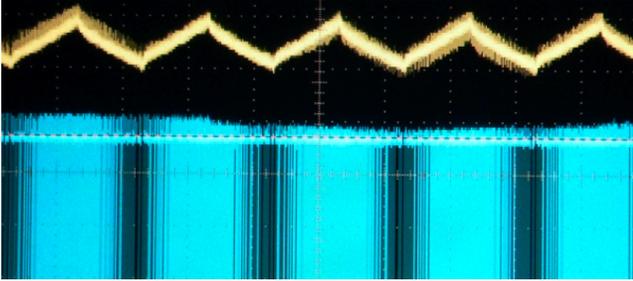


Figure 10: The VFC transfer function shows output frequency is proportional to input voltage.

put is a simple digital-compatible pulse train instead of a full fledged digital bus as in the case of ADC. We further investigate VFC below.

A VFC operates by integrating the input voltage and feeding the output of the integrator to a comparator as seen in Figure 9. As soon as the charge accumulated in the integrator exceeds the reference voltage (V_{ref}), the comparator outputs a pulse, which also acts as negative feedback to “balance” the charge inside the integrator. The net effect is a pulse train whose frequency is proportional to the input voltage. It is different from ADC because any arbitrarily small voltage can be captured via the analog integrator. The transfer function for the VFC we chose is:

$$f_{OUT} = 0.1f_{CLKIN} + 0.8(V_{IN}/V_{REF})f_{CLKIN} \quad (1)$$

where f_{CLKIN} is the 1MHz clock input and V_{ref} is 3.3V. A more graphical representation is shown in Figure 10.

While the resolution of VFC is infinite in theory since it’s analog, the actual fidelity of the signal is limited by the internal noise. The VFC we are using has a maximum peak-to-peak noise of $100\mu V$. Assuming the gain of our amplifier is around 80 and our lowest input voltage is $2\mu V$, the smallest input to the VFC is then $80 \times 2\mu V = 160\mu V$, which is barely above the noise. While this places stringent noise filtering requirement on our design, the actual signal-to-noise ratio (SNR) is not bad since the specified $100\mu V$ is the maximum peak-to-peak noise (as opposed to RMS noise).

Our choice of using VFC is also due its easily-integratable output. As mentioned before, we need to sum the output of the VFC to obtain energy and we want minimal perturbation of the DUT. ADC usually require a DSP or MCU at the output to “interpret” the digital output. If the same

mote is used to read ADC’s output, it will incur significant overhead because power readings need to be read at a fairly high frequency. To reduce complexity, perturbation, and cost, we favor VFC’s elegant pulse train output, which can be easily accumulated using simple counters (see Section 5.4 for details).

5.4 Energy Output

The last stage is energy accumulation and digital output. To provide the user with a direct energy value, we will need to integrate the power readings from the previous stage:

$$E(t) = \int_{t_0}^t P(\tau)d\tau$$

If the previous stage uses an ADC, we will need a DSP or MCU, which either incurs significant complexity/cost or perturbation, as explained in the previous section.

Fortunately, we can use simple counters to sum the pulse train output from the previous stage. Every rising or falling edge of a pulse in the pulse train increments the counter by one. Higher power consumption leads to more frequent pulses, which in turn leads to faster counter increments. The difference in counter values between t_0 and t_1 represents the energy consumption during time $t_1 - t_0$. However, counters have a finite range, represented by its bits. To prevent counters from overflowing, the MCU needs to read and reset the counters periodically. This presents overhead, but it can be minimized by using a large counter that overflows infrequently. We choose to use a 32-bit counter. From equation 1, we find that the VFC has a maximum output frequency of 0.9MHz (assuming VFC is maximum when $V_{IN} = V_{REF}$ and $f_{CLKIN} = 1MHz$). This implies that the counter will overflow roughly once every hour (see Section 4.3), which is infrequent.

Finally, for the user of SPOT (e.g. the mote) to read the counter values, we need a compatible digital output format. Parallel output is not possible because 32 bits require too many pins. Serial output such as I2C or SPI is desired. Since our data rate is low, I2C is more than sufficient.

Additionally, to relieve the MCU of time-keeping responsibilities and to precisely measure elapsed time, we included another 32-bit counter dedicated to time-keeping and is triggered by the same oscillator that drives the VFC. Please refer to Section 4.3 for preliminary descriptions. For this scheme to work, we need some way of reading the two counter chips at the exact same time. Fortunately, one valid command to the counters that we use is “capture”, which essentially takes a snapshot of the counter values and stores them in registers. Furthermore, because the addresses for the counter chips can be dynamically changed, we can wire the I2C addresses for the two chips to be the same during a “capture” command, then re-wire the addresses to be different and read back the captured values individually.

5.5 SPOT Application Programming Interface

We provide a high-level application programming interface to SPOT to make using it simple for TinyOS developers. Our API is shown in Figure 11.

6. EVALUATION

In this section, we evaluate SPOT’s accuracy as a tool for metering energy and power for a typical sensor network workload. In particular, we evaluate SPOT’s dynamic range, res-

```

// Initialize the counters and calibrate the system.
result_t init();

// Signals end of initialization phase.
event void initDone();

// Read both energy and time counters at the same time.
// 1. Configure addresses to be the same
// 2. Issue I2C command to capture counter values
// 3. Configure addresses to be different
// 4. Read from time counter
// 5. Read from energy counter
// 6. Signal readDone
command result_t readCounter();

// Signals the completion of reading the counters
// Returns the two 32-bit counter values.
event void readDone(uint32_t time, uint32_t energy);

// Enable or disable the counters from counting.
command result_t setEnable(bool enable);

// Performs calibration.
command result_t cal();

// Signals end of calibration.
event void calDone();

```

Figure 11: The SPOT application programming interface includes commands to initialize, calibrate, enable, read, and reset the meter.

olution, and stability, which represent the most challenging requirements for an embedded power meter.

6.1 Dynamic Range and Resolution

A principal requirement of our system is a dynamic range exceeding 10000:1 with current resolution of at least $2\mu A$. To evaluate SPOT’s resolution, we loaded a 3.3V power source with resistors of different value in the $M\Omega$ range to create currents of 1.08, 2.07, and $3.48\mu A$. In addition, we loaded the power source with a TelosB mote in sleep mode, which draws $9.09\mu A$, to provide an additional data point.

Table 1: Snapshots of SPOT’s uncalibrated, free-running Time and Energy counters taken approximately every 40 ms (first two columns).

Time Count	Energy Count	Counts/Sec
397369588	2535884271	161804.7903
397424011	2535893077	161806.5891
397453917	2535897916	161806.9953
397489646	2535903697	161801.3378

For each data point, we take 600 consecutive readings from SPOT. An example of the uncalibrated readings that SPOT output is shown in Table 1 in the first two columns. The first column lists counter values from the Time Counter and the second column lists counter values from the Energy Counter. The elapsed time between readings is the sum of the I2C access time (I2C is running at 100kHz) and the radio packet transmit time (one packet is transmitted for each reading).

For simplicity, we will refer to values in the table using the initials of the column heading plus the row number in subscript (starting at 1). For example, the second Energy Count is designated as EC_2 . The uncalibrated data shown in Table 1 can be used in several ways. For example, we can determine the energy the mote consumed between the 1st reading and the 4th reading simply by computing $EC_4 - EC_1$; this

energy corresponds to an elapsed time of $(TC_4 - TC_1)/10^6 = 0.12$ seconds, since the Time Counter increments at 1MHz. Since power is the derivative of energy, we can estimate the average power over an interval by simply taking the difference of the energy counts. The third column shows the power in units of counts/second. For example, C/S_2 is obtained by computing $(EC_2 - EC_1)/(TC_2 - TC_1) \times 10^6$. This value is linearly proportional to power and current, since voltage is constant.

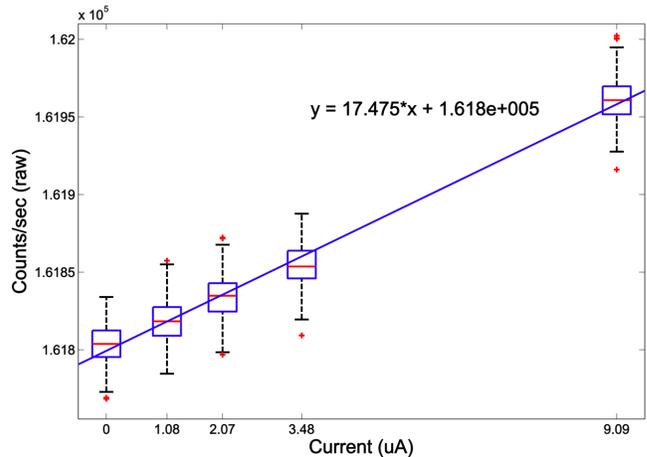


Figure 12: SPOT current resolution. SPOT can resolve currents at the microamp level.

Power values (counts/sec) for each current are plotted in Figure 12 in boxplots to show their distributions. The power for $0\mu A$ current is not zero due to the amplifier non-zero offset voltage and the VFC minimum frequency of 0.1MHz.

Notice that while values between μA boundaries do overlap, the variance is limited and the medians values sit at regular intervals from each other. This suggests that SPOT is able to resolve $2\mu A$ or smaller currents, but that it is necessary to take multiple samples. In practice, the counter reading rate is not likely to be low, which effectively averages the readings over much longer runs.

A simple linear regression of the five medians shown in Figure 12 generates a useful calibration curve for SPOT. Despite being far from optimal, since this curve is generated using data for points between 0 and $9.09\mu A$ rather than the full range extending to 45mA, we show that this curve is still useful in the next section.

Because our Vcc is 3.3V and our amplifier has a gain of around 66, we can tolerate a maximum input current of 45mA, assuming the amplifier offset does not exceed 0.3V. At the other extreme, Figure 12 shows that we can resolve currents at $1\mu A$ or even less. This means that SPOT has a dynamic range exceeding $\frac{45mA}{1\mu A} = 45000 : 1$, which surpasses our dynamic range requirement.

6.2 Long-Term Tracking Accuracy

Because motes spend the majority of their time sleeping, it is important to evaluate SPOT’s accuracy in monitoring a mote’s energy consumption during its sleep state. This is more difficult than monitoring a mote in active state because the current to be monitored is four orders of magnitude smaller.

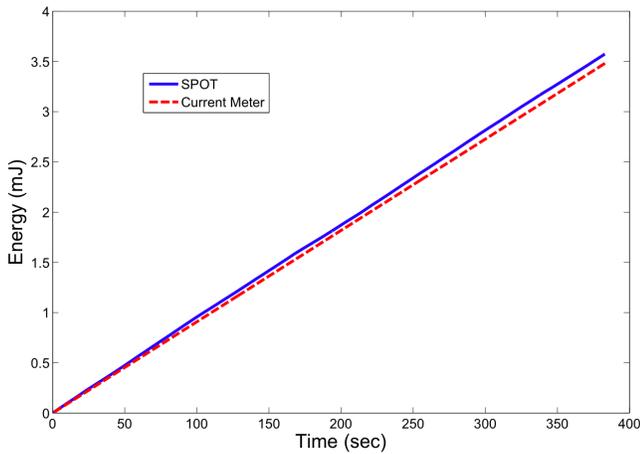


Figure 13: Energy metering of $9.09\mu A$ load over a period of time exceeding 7 minutes. The accumulated error is 0.1mJ or 3% of the actual energy usage.

The counter readings obtained from SPOT were calibrated using the equation found in Section 6.1 and compared with a reference curve obtained using an professional-grade current meter, as shown in Figure 13. The mote under observation, a TelosB [11], is drawing approximately $9\mu A$ of current.²

As seen in Figure 13, SPOT is able to closely track energy usage, but with a small drift. After 6 minutes, the error is less than 0.1mJ or 3%. The absolute error accumulates over time but the relative error should stay constant, which may be acceptable for most applications. If more accurate calibration is required, additional calibration points will be needed.

6.3 Energy and Power Tracking

Most sensornet applications are duty-cycled, waking up occasionally to take a sensor sample or send a message, and sleeping the remainder of the time. To evaluate SPOT’s tracking accuracy in monitoring motes with relatively low duty-cycles, we now consider a typical workload. The mote under observation is a TelosB mote running at slightly higher than 2Hz duty-cycle. The reference curve is generated by integrating the measured power signals collected using a high-end digital storage oscilloscope. Because the oscilloscope does not have enough dynamic range to resolve active state power and sleep power at the same time, we will focus on only the active state power in this experiment, since the previous section demonstrated SPOT’s ability to track sleep state power.

The top graph in Figure 14 shows the energy monitored by SPOT and the oscilloscope. The flat portions of the curve represent energy consumed by the mote during sleep while the sharp rises between flat steps represent energy consumed during active cycles. The change in energy during active cycles is approximately 0.4mJ while there is no observable energy change during sleep states. An interesting observation is that energy consumed during different active cycles are all slightly different. This is likely due to random backoff in the medium access control layer, which causes the variations

² $9\mu A$ represents a typical sleep state current for TelosB powered at 3.3V voltage supply.

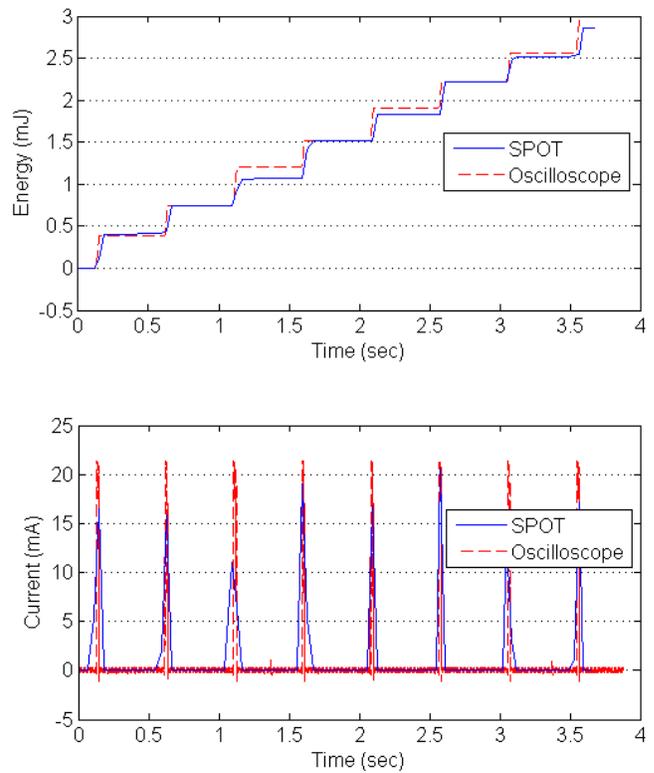


Figure 14: Energy and power tracking of duty-cycling TelosB mote using SPOT compared with a digital storage oscilloscope. The flat portions of the top graph represent energy consumed by the mote during sleep while the sharp rises between flat steps represent energy consumed during active cycles. The bottom graph shows the power draw recorded by SPOT by differencing successive energy usage readings.

in active cycle time. This also explains why SPOT does not match the oscilloscope more closely, since the SPOT and oscilloscope measurements are taken at different times.

The bottom graph of Figure 14 shows the power draw recorded by SPOT.³ Note that the observed resolution is quite low. This is because we wait for the radio on every sample. However, because SPOT samples the signal at 1MHz internally, even-though we are not reading the counters as fast, the energy measured is still quite accurate.

7. ENABLED RESEARCH

Although the SPOT system itself represents a unique, mixed-signal, hardware-software design, its true impact will be measured by the novel research it enables. In this section, we outline our near-term plans for this system and speculate on how SPOT may affect future sensornet research.

SPOT, as described in this paper, is a single-sided, leadless chip carrier (LCC) that can be treated as a modular circuit component and directly soldered to a printed circuit

³The terms “power” and “current” are used interchangeably because the voltage is constant and therefore power is directly proportional to current.

board. To be useful, the module must intercept the power line between the power supply and the node under test. We have built a Mica2/MicaZ carrier board that incorporates the SPOT module, all needed power supplies, a real-time clock, extra counters, and power line intercept, as shown in Figure 15. A Telos [11] version of the carrier board is under development. We plan to instrument an entire testbed with these modules.

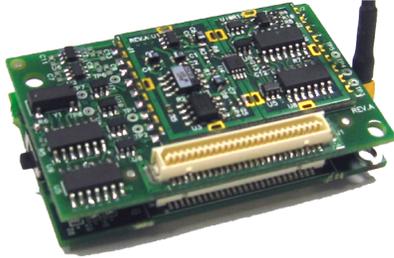


Figure 15: The SPOT module attached to its MicaZ carrier board.

Recall that state-of-the-art simulators like PowerTOSSIM base their models on microbenchmarks of a single node taken in a particular environment. An instrumented testbed will allow us to answer a number of questions about the generality of these models: How representative is the node that was instrumented to calibrate PowerTOSSIM’s model? How does interaction with the physical environment shape energy usage? How do temperature variations affect leakage currents?

Beyond the simple verification of mote power models, we believe the high impact research enabled by SPOT will be macrobenchmarking the energy-efficiency claims existing network protocols as well as implementations of similar or identical protocols from different TinyOS distributions.

We envision two variations on SPOT that would provide still greater visibility into nodal power profiles. First, it would be simple to allow the node under test to adjust its supply voltage, within bounds, and multiply, in the analog domain, the voltage and current to create a true power/energy meter. This would allow profiling nodal power draw across a range of voltages, as is common when running on batteries. Second, the instantaneous power output of the preceding system could be digitized using a high-speed ADC, collected in a deep FIFO, and read out more slowly by the node under test. This would allow captures of very fine-grained node power profiles.

8. CONCLUSION

We presented the requirements, architecture, design trade-offs, and preliminary evaluation of SPOT – an accurate and sensitive meter for monitoring wireless sensor network power profiles and energy usage patterns at scale. We showed that SPOT is able to meet or exceed challenging requirements unique to wireless sensor networks and other low power systems, as summarized in Table 2. SPOT employs a simple architecture that uses a pure analog sampling front end to provide a large dynamic range, and a dual counter energy accumulation stage to provide high temporal resolution. These features are useful for profiling a range of systems – like pagers, PDAs, and cellphones – which exhibit highly

bimodal or widely varying power profiles. SPOT will enable heretofore impossible empirical evaluations of low power designs at scale, and it will enable a new class of sensornet research in which applications can integrate the dynamic power profile of a system into the application logic.

Metric	Requirement	SPOT
Dynamic Range	> 10000 : 1	45000 : 1
Resolution	< 2 μ A	< 1 μ A
Sampling Rate	> 20kHz	Internally at 1MHz Output at I2C speed
Perturbation	Minimal	1 Ω additional load to DUT Energy measurement via I2C At least one read per hour
Integration	Easy	1.35" x 1" all-in-one
Cost	< \$25	Off-the-shelf ICs

Table 2: SPOT satisfies the power and energy metering needs of sensornet nodes *in situ* and *at scale*.

9. REFERENCES

- [1] Analog Devices. Single phase multifunction energy metering ic with di/dt input. http://www.analog.com/UploadedFiles/Data_Sheets/ADE7753.pdf, Aug. 2003.
- [2] M. Buettner, G. Yee, E. Anderson, and R. Han. X-mac: A short preamble mac protocol for duty-cycled wireless sensor networks. In *SenSys*, Nov. 2006.
- [3] P. Dutta, M. Grimmer, A. Arora, S. Bibyk, and D. Culler. Design of a wireless sensor network platform for detecting rare, random, and ephemeral events. In *IPSN/SPOTS*, Apr. 2005.
- [4] S. Gurun and C. Krintz. Full system energy estimation for sensor network gateways. In *Technical Report*, Oct. 2006.
- [5] T. He, S. Krishnamurthy, J. A. Stankovic, T. F. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh. An energy-efficient surveillance system using wireless sensor networks. In *MobiSys*, June 2004.
- [6] J. Hsu, J. Friedman, V. Raghunathan, A. Kansal, and M. Srivastava. Heliomote: Enabling self-sustained wireless sensor networks through solar energy harvesting. In *ISLPED*, Aug. 2005.
- [7] S. Kim, S. Pakzad, D. E. Culler, J. Demmel, G. Fenves, S. Glaser, and M. Turon. Health monitoring of civil infrastructures using wireless sensor networks. In *Technical Report*, Oct. 2006.
- [8] Maxim-IC. Ds2438 smart battery monitor. <http://datasheets.maxim-ic.com/en/ds/DS2438.pdf>, July 2005.
- [9] Microchip. Energy metering ic. <http://ww1.microchip.com/downloads/en/DeviceDoc/21948c.pdf>, Aug. 2005.
- [10] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *SenSys*, Nov. 2004.
- [11] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power research. In *IPSN/SPOTS*, Apr. 2005.
- [12] O. Saukh, P. Marron, A. Lachenmann, M. Gauger, D. Minder, and K. Rothermel. Generic routing metric and policies for wsns. In *EWSN*, Feb. 2006.
- [13] K. Seada, M. Zuniga, A. Helmy, and B. Krishnamachari. Energy-efficient forwarding strategies for geographic routing in lossy wireless sensor networks. In *Conference On Embedded Networked Sensor Systems*, Nov. 2004.
- [14] V. Shnayder, M. Hempstead, B. rong Chen, G. W. Allen, , and M. Welsh. Simulating the power consumption of large-scale sensor network applications. In *SenSys*, Oct. 2004.
- [15] G. Tolle, J. Polastre, R. Szewczyk, N. Turner, K. Tu, S. Burgess, D. Gay, P. Buonadonna, W. Hong, T. Dawson, and D. Culler. A macroscope in the redwoods. In *SenSys*, Nov. 2005.
- [16] Y. Xu, S. Bien, Y. Mori, J. Heidemann, and D. Estrin. Topology control protocols to conserve energy in wireless ad hoc networks. In *Technical Report 6*, Apr. 2003.
- [17] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *INFOCOM*, June 2002.
- [18] P. Zhang, C. M. Sadler, S. A. Lyon, , and M. Martonosi. Hardware design experiences in zebranet. In *SenSys 2004*, Nov. 2004.