μMonitor: In-situ Energy Monitoring with Microwatt Power Consumption

Saman Naderiparizi[1], Aaron N. Parks[1], Farshid Salemi Parizi[1], Joshua R. Smith[2,1]
[1]Electrical Engineering Department and [2]Computer Science and Engineering Department
University of Washington, Seattle, USA-98195

***Abstract*—Knowledge of energy flow in a microwatt-class energy harvesting system is essential to reliable deployment and scheduling of sensing, computation, communication, and actuation tasks. However, existing techniques for monitoring energy flow fail to meet the basic requirements for in-situ realtime monitoring systems by failing to be efficient and failing to perform accurately across a wide dynamic range. The proposed system, μMonitor, makes use of a highly power-optimized "Coulomb counting" implementation to achieve less than 1.7 microampere current draw, 94% efficiency in-situ, and high energy flow measurement accuracy across four orders of magnitude.**

I. Introduction

Any battery operated system, such as a modern low-power connected device, has strong motivations for measuring its own energy consumption. In laptops and smartphones, battery life is estimated based on power consumption trends and used to inform the user of the estimated remaining time before a charge is required.

Energy harvesting systems, which derive their minuscule power requirements from sources of energy in their environment, have an even greater motivation for monitoring energy flow. These systems extract energy from ambient sources (both intentional and unintentional), such as temperature gradients, vibration, and radio signals, store this energy temporarily using a battery or capacitor, and then expend it to do tasks involving sensing, computation, communication, or actuation.

In these systems, energy monitoring is essential for reliable performance of even the most basic tasks, as the system must be aware of the available energy in order to successfully deploy and complete tasks. For instance, transmitting a packet over a radio link is an energy-atomic task, meaning the amount of energy required to complete the entire task must be available before the task begins (or else the task will fail entirely). If an energy harvesting system wishes to transmit a packet, it must first accumulate the energy required to do so, and therefore must have knowledge both of the energy required to send the packet and of the energy accumulated on its storage element. Additionally, if it wishes to schedule such an energy-atomic task to occur at a certain rate or at a certain future time, it must also be aware of the rate at which charge is being accumulated in order to predict future energy availability.

A. Power monitoring considerations

A method for monitoring energy flow in these ultra low-power systems should be judged by two main considerations: accuracy across dynamic range, and efficiency:

1) Accuracy across dynamic range: The target platform may have a variation in consumed power across many orders of magnitude, for example, from nanowatts to milliwatts. Additionally, the target platform may spend an overwhelming majority of its time in a low power state, making it essential to have accurate measurement of sleep mode power consumption when estimating charge state. Incoming harvested power may also vary by several orders of magnitude for many harvesting sources. Monitoring accurately across this entire range is critical to achieve low integral error, but presents significant design challenges.

2) Efficiency: An in-situ monitoring system must not detract significantly from the available power, as the target platform's power budget is already very constrained. Simple, low parts count designs with minimal computational requirements will help achieve this goal.

B. Conventional power measurement methods

Conventional methods such as those listed in [1], while effective in some scenarios, do not sufficiently address the above criteria in the context of an energy harvesting system. Below we list several standard approaches to monitoring energy flow and enumerate the flaw with each in the context of ultra low power or energy harvesting systems:

1) Series sense resistor: A serial resistor in the power path of interest, across which voltage is monitored. This method is confined to a very narrow dynamic range, as developed voltage across the resistor must be large enough to be accurately measured but not so large as to interfere with system operation or dissipate significant power. Some work partially overcomes this issue with meticulously designed highly sensitive preamplification around a low-value sense resistor [2], while others use dynamic selection of sense resistor values [3].

2) Voltage monitoring: Monitoring the voltage across the charge reservoir can give an estimate of how much charge is left. For example, a fully charged "12 volt" battery will have an open-circuit voltage of 12.6V, while a fully discharged one will read 10.5V. However, this method has poor accuracy due to the large number of variable involved in determining battery voltage. Plus, in scenarios where input power is low (in the order of microwatts), the voltage changes on the charge reservoir can be as low as single digit millivolts. Accurately recording these minuscule voltage changes requires a high resolution and low-noise digitizer, which is not available on a power limited platform.

3) Current clamp metering: Monitoring the charge flow through a conductor by monitoring the induced magnetic field is a technique that is well known in high-power systems, but hasn't been proven in microwatt class systems to the best of our knowledge.

C. μMonitor approach

μMonitor approaches energy measurement from a different perspective, by applying a in-situ charge-to-pulse converter, known as a coulomb counter, which emits a pulse for each pre-measured unit of charge that moves across it. Pulses can be counted by a logic system in order to estimate energy availability. The μMonitor turns the WISPCam [4] (or any other energy scavenging system) into an energy-aware device that can make appropriate decisions based on its remaining energy as well as on the measured value of incoming harvested power.

While some existing work has demonstrated current monitoring using a Coulomb counter in the context of wireless sensor nodes [5], the system characterized in that work is a bench-top measurement unit and does not provide the realtime, in-situ measurement capability that is the goal of μMonitor.

We show that μMonitor allows precision monitoring across more than a three orders of magnitude dynamic range, and carefully select circuit parameters for minimal wasted power.

In this paper, we first present the details on designing the μMonitor. Then, we present some benchtop experiments that characterize the μMonitor. Then, we apply μMonitor in-situ in a radio frequency harvesting system, the WISPCam [4], to demonstrate how accurate this method can be in the face of large input power fluctuations. Finally, we introduce some potential applications that can be enabled by μMonitor.

II. Design

In this section we introduce the steps we took in getting to the final revision of our design. We will also explain in detail the design choices involved.

A. Coulomb counting

The key technique behind μMonitor is to shuttle charge across a known value capacitor as it flows from one circuit block to another. In our system, we monitor charge flow from the energy harvesting circuit to the system's charge reservoir (battery or supercapacitor).

A switch will first be closed allowing charge sourced from the energy harvester to be stored on this known capacitor. When adequate charge is stored to reach some upper voltage threshold V_{high}, the first switch is opened and a second switch closed allowing this charge to flow to the charge reservoir (for example, to be stored on a battery). The second switch opens yet again when some lower threshold V_{low} is attained. By knowledge of V_{high} and V_{low} and the capacitor size, energy delivered to the charge reservoir with each switching event may be determined as shown in Equation 10. By counting how many times the capacitor discharged into the charge reservoir, total energy delivered into the charge reservoir can be estimated.

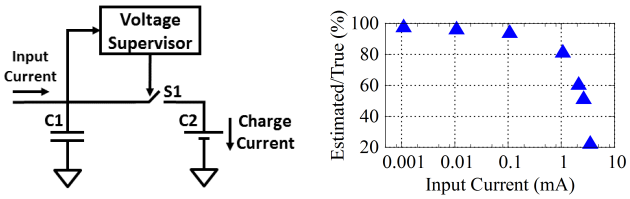

(a) Version 1 hardware (b) Estimated/actual energy vs input current

Fig. 1: Figure 1a shows Iteration 1 of the hardware, as analyzed in simulation. Figure 1b shows the estimated/actual energy limitation of this first iteration design as the input current increases.

$$E = \frac{1}{2}C(V_{high}^2 - V_{low}^2) \tag{1}$$

This charge-shuttling circuit topology in itself is not novel, and is known colloquially as a coulomb counter. However, the application of a coulomb counter to in-situ power monitoring with high dynamic range and with very low power overhead has not been previously demonstrated.

B. Topology selection, in three iterations

Three iterations were carried out during the circuit design phase of μMonitor. We walk through the design choices behind each iteration in order to better establish the reasoning behind our final design.

1) Iteration 1: Single switch coulomb counter: To implement the μMonitor Coulomb counter, a voltage supervisory circuit can control a switch based on the V_{high} and V_{low} thresholds. Figure 1a shows a block diagram for the first attempt at such a design. This achieves the goal of monitoring accurately when input power is low. However, this first circuit iteration does not perform accurately when the input current (constant current charging C1) is a non-negligible fraction of the charge current (transient current flowing into the charge reservoir, modeled here as C2). Nor does this circuit accurately measure current when the duty cycle of switch S1 is non-negligible (the ON time is a noticeable fraction of the period). Both these cases produce non-negligible charge flow directly from the input to C2, effectively bypassing the monitoring capacitor C1 entirely. Figure 1b shows the simulated decrease in estimated/actual energy with increasing input power.

While this first implementation is desirable in its simplicity, because of these performance limitations, this early implementation was abandoned.

2) Iteration 2: Isolating current paths: To overcome the flaws of the Iteration 1 design, we must isolate the incoming energy storage stage from energy delivered to the charge reservoir such that no charge can bypass the coulomb counter. To that end, the block diagram in Figure 2 was designed and evaluated. The charge current here can be measured through $C1'$ which is practically a very large current since the series resistance of a ceramic capacitor is very small.

But discharging a capacitor into another capacitor is not an energy-neutral operation, since some energy will waste as heat and electromagnetic radiation. Assuming the input current is negligible comparing to the maximum charge current in

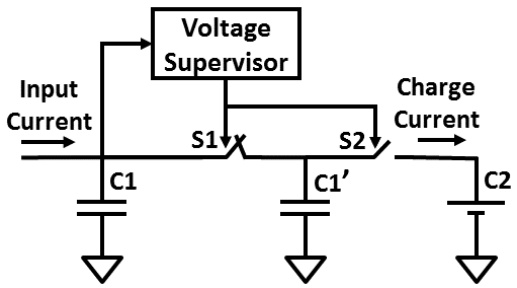

Fig. 2: The second version of the μMonitor circuit design, which isolates the input storage capacitor from the large charge reservoir

Figure 1a, then the wastage energy and its characteristics will be as below:

Let $V_1 = V_2 + \Delta V + \frac{\Delta Q}{C_1}$ where ΔV is the difference between V_1 and V_2 when switch S1 opens and ΔQ is the charge injected to C1 in every charging cycle. So right before S1 closes we have:

$$\begin{aligned} E_1 &= \frac{1}{2}C_1V_1^2 = \frac{1}{2}C_1(V_2+\Delta V+\frac{\Delta Q}{C_1})^2 \\ &= \frac{1}{2}C_1V_2^2 + \frac{1}{2}C_1\Delta V^2 + \frac{1}{2}\frac{\Delta Q^2}{C_1} + C_1V_2\Delta V + V_2\Delta Q + \Delta V\Delta Q \end{aligned} \tag{2}$$

$$E_2 = \frac{1}{2}C_2V_2^2 \tag{3}$$

Right after S1 opens again:

$$E_1^* = \frac{1}{2}C_1V_1^{*2} = \frac{1}{2}C_1V_2^2 + \frac{1}{2}C_1\Delta V^2 + C_1V_2\Delta V \tag{4}$$

$$E_2^* = \frac{1}{2}C_2V_2^{*2} = \frac{1}{2}C_2V_2^2 + \frac{1}{2}\frac{\Delta Q^2}{C_2} + V_2\Delta Q \tag{5}$$

So the efficiency of delivering charge from C1 to C2 can be calculated as below:

$$\Delta E_1 = E_1 - E_1^* = \frac{1}{2}\frac{\Delta Q^2}{C_1} + V_2\Delta Q + \Delta V\Delta Q \tag{6}$$

$$\Delta E_2 = E_2 - E_2^* = \frac{1}{2}\frac{\Delta Q^2}{C_2} + V_2\Delta Q \tag{7}$$

$$\eta = \frac{\Delta E_2}{\Delta E_1} = \frac{\frac{\Delta Q}{C_2} + 2V_2}{\frac{\Delta Q}{C_1} + 2V_2 + 2\Delta V} \tag{8}$$

Since in our work always C2 is far more larger than C1 we can simplify Equation 8 to below:

$$\eta = \frac{\Delta E_2}{\Delta E_1} \approx \frac{2V_2}{\frac{\Delta Q}{C_1} + 2V_2 + 2\Delta V} \tag{9}$$

Based on Equation 9, as ΔV becomes larger or C1 gets more charge during charging period, the efficiency will decrease. It is important to note that, not only smaller efficiency will lead into having lower estimated/actual energy, but also it will result in larger total wastage energy overhead due to charge monitoring.

3) Iteration 3: Final μMonitor design: To solve the problems described above, the final iteration of the μMonitor design makes use of a system that picks thresholds dynamically on-the-fly. The idea is to keep changing the voltage threshold as the battery/supercapacitor is being charged. At each moment in the time we make sure that the C1′ voltage charges up to maximum αV_{C2}. Figure 3 shows the block diagram for this circuit. We set R1 and R2 such that:

$$\alpha = \frac{R1+R2}{R2} \tag{10}$$

By having $C2'$ in the circuit shown in Figure 3, we will make sure that after S2 opens $C1'$ has discharged to $\beta V2$. Based on conservation of charge (when connecting two capacitors), β can be found by the below equation:

$$\beta = \frac{C1'\alpha + C2'}{C1' + C2'} \tag{11}$$

Then, while S2 is open, $C2'$ will discharge its accumulated charge to C2. So in theory the overall efficiency of this topology will be dictated by the efficiency of discharging $C1'$ into $C2'$ (we will call it η_1) and the efficiency of discharging $C2'$ to C2 (we will call it η_1). So the overall efficiency (η_{total}) can be found as shown below:

$$\eta_1 = \frac{\beta^2 - 1}{\alpha^2 - \beta^2}\frac{C2'}{C1'} \tag{12}$$

Also using Equation 9 with ΔV equal to zero:

$$\eta_2 = \frac{2V_2}{(\beta-1)V_2 + 2V_2} = \frac{2}{\beta+1} \tag{13}$$

$$\eta_{total} = \eta_1 \times \eta_2 = \frac{2(\beta^2-1)}{(\beta+1)(\alpha^2-\beta^2)}\frac{C2'}{C1'} \tag{14}$$

A remarkable feature of this design is that η_{total} is independent of V_2 and is only related to α. In the other words, as C2 is being charged, the efficiency of delivering charge will remain constant and, consequently, accuracy of incoming power estimation will not be affected much as the voltage on C2 changes. Just to give the reader a notion of what is the value for a typical overall efficiency, assuming α is 1.05, and $C1'$ is equal to $C2'$ then η_{total} will be more than 95%

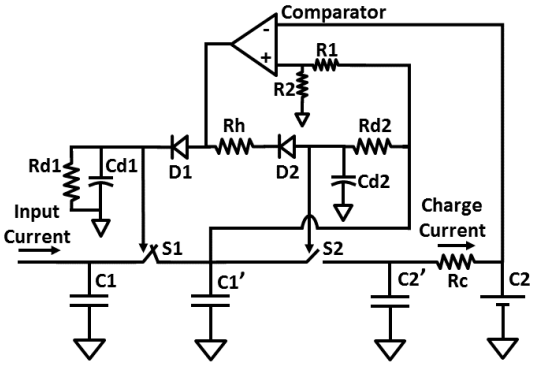

Fig. 3: Final revision of the proposed hardware

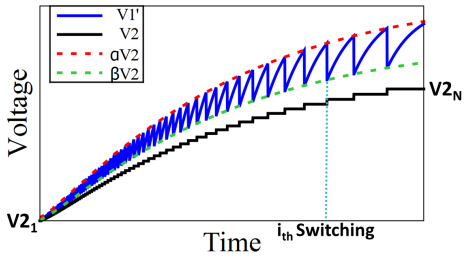

Fig. 4: Qualitative curves that show the operation of dynamic thresholding as C2 is getting charged up. The bottom (βV2) and top (αV2) thresholds are shown with dashed lines.

C. Switching considerations

Another challenge that we faced and addressed in this work is the potential shoot-through current introduced due to the fact that S1 and S2 may conduct at the same time. Basically, it is required that S1 and S2 do not conduct at the same time as this will cause some charge transfer from C1 to C2 without being monitored by $C1'$, resulting in reduced estimated/actual delivered energy. Since the duration of this shoot-through time interval will differ from one switch to another due to topology and process variations, it is important to address this problem fully.

Here are the two problem intervals that need consideration:

- **Comparator output going from low to high:** the voltage threshold has been reached and $C1'$ should be discharged to $C2'$. In this case, S1 should open fully first and then S2 closes. According to Figure 3, S1 will open instantly and S2 will close after a delay which is proportional to Cd2×Rd2;
- **Comparator output going from high to low:** $C1'$ has been discharged to $C2'$ and $C1'$ needs to get charged up again from C1. In this case, S2 should disconnect first before S1 closes. According to Figure 3 S2 will open with a delay which is proportional to (Rh||Rd2)×Cd2 and S1 will close with a delay which is proportional to Rd1×Cd1.

It is important to mention here that the comparator has a single threshold, so in theory and assuming that comparator does not have any hysteresis, it will go low immediately after it goes high. The reason is that right after closing S2, the $C1'$ voltage will drop a bit and comparator output will go low. This can cause a serious problem, if S2 reopens before $C2'$ and $C1'$ reach equilibrium. To make sure that $C2'$ and $C1'$ have enough time to reach equilibrium, we induced a delay from the time the comparator output goes low to the time S2 opens. The purpose of Rh is to provide a known deterministic delay period.

D. Determining delivered energy

Here we present a mathematical model with which we can estimate the delivered energy to C2 in the final μMonitor system of Figure 3. Figure 4 shows the voltage of C2 and $C1'$ as we are accumulating charge on C2.

The key question here is how to estimate the amount of delivered energy when knowing only $V2_1$, $V2_N$, and N (total number of switching events). The solution is detailed below:

$$\begin{aligned} E_{discharged} &= \sum_{i=1}^{N} \frac{1}{2} C1'(\alpha^2 {V2_i}^2 - \beta^2 {V2_i}^2) \\ &= \frac{1}{2} C1'(\alpha^2 - \beta^2) \sum_{i=1}^{N} {V2_i}^2 \end{aligned} \tag{15}$$

Based on conservation of charge when discharging $C2'$ into C2, the equation for $V2_i$ will be as below:

$$V2_i \approx (1 + \frac{C2'\beta}{C2}) V2_{i-1} \rightarrow V2_i \approx (1 + \frac{C2'\beta}{C2})^{i-1} V2_1$$

Since C2 is much larger than $C2'\beta$ we can employ Taylor linear approximation:

$$\begin{aligned} &{V2_i}^2 \approx (1 + 2(i-1)\frac{C2'\beta}{C2}) {V2_1}^2 \\ &\rightarrow {V2_N}^2 \approx (1 + 2(N-1)\frac{C2'\beta}{C2}) {V2_1}^2 \end{aligned} \tag{16}$$

Now we substitute the Equation 16 into Equation 15 to get the final result:

$$\begin{aligned} E_{discharged} &= \frac{1}{2} C1'(\alpha^2 - \beta^2) \sum_{i=1}^{N} (1 + 2(i-1)\frac{C2'\beta}{C2}) {V2_1}^2 \\ &= \frac{1}{2} C1'(\alpha^2 - \beta^2) {V2_1}^2 \sum_{i=1}^{N} 1 + 2(i-1)\frac{C2'\beta}{C2} \\ &= \frac{1}{2} C1'(\alpha^2 - \beta^2) {V2_1}^2 (N + N(N-1)\frac{C2'\beta}{C2}) \\ &= \frac{N}{2} C1'(\alpha^2 - \beta^2) V2_1 V2_N \end{aligned} \tag{17}$$

Finally, we know that based on Equation 14, to get a more accurate estimate of $E_{delivered}$ we should take η_{total} into consideration as well. So an accurate estimation of $E_{delivered}$ will be as below:

$$E_{delivered} = \eta_{total} \times E_{discharged} = N(\beta - 1) \times C2' \times V2_N \times V2_1 \quad (18)$$

The above equation has been driven with the assumption that capacitance of a ceramic capacitor is constant over different bias voltages. In fact, this is not a valid assumption. It turns out that there is over a 300% variation in capacitance of a 150uF ceramic capacitor when its voltage goes from zero to 4.5V. It is crucial to take this voltage dependence into account to achieve a more accurate delivered energy estimation.

To address the voltage dependence of capacitance, we used the fact that C2, the charge reservoir, is far larger than the other capacitors in the system, and therefore has a quasi-static voltage on the timescale of individual switching events.

We initially measure the capacitance of the ceramic capacitors for a few evenly spaced voltages to create a look up table. Then, periodically, we will sample the voltage on C2 and will reset the switching counter to zero. At each voltage we can get the approximate capacitance by applying piece-wise linear interpolation over the capacitance look up table. In our case, we can assume a fixed switching count as that within which the charge reservoir voltage will not change significantly. Let's call this N_p. For simplicity let's assume that we had total $K \times N_p$ switching events. So, using Equation 18 the delivered energy estimation will be as follows:

$$E_{delivered} = N_p \sum_{i=1}^{K} (\beta_i - 1) C2'_i \times V2_{i+1} \times V2_i \quad (19)$$

This compensates for the voltage dependence of the measurement capacitors, allowing μMonitor to be accurate over the entire charging cycle of the charge reservoir.

III. Prototype Implementation

μMonitor has been implemented fully using off-the-shelf components. μMonitor prototype is shown in Figure 5. This is built such that it can be combined into any designs by easily inserting that in series with their charge reservoir. An ultra-low-power TLV3691 comparator [6] is employed to achieve the dynamic threshold. ADG801 and ADG802 [7] are the low-power complement switches that deployed on the board. R1 and R2 are selected such that we achieve $\alpha = 1.056$. Also, we pick Rd2 and Cd2 such that S1 closes about 50us after comparator output goes low and we pick Rd1 and Cd1 such that S1 closes about 130us after comparator output goes high.

An an ultra-low-power MSP430FR5969 microcontroller [8] (MCU) was used to count the number of switching as well as calculating how much energy has been stored. When the system has been powered initially, the MCU will start a timer and will go to one of its lowest power modes. At this mode the CPU is off, and all high frequency clocks are also powered down and the MCU's quiescent current is about 0.5uA. The only modules that are active in this low-power mode are VLO, which is a very low-power low frequency oscillator, and the timer. Timer is being clocked externally from the output of comparator to count the number of switching.

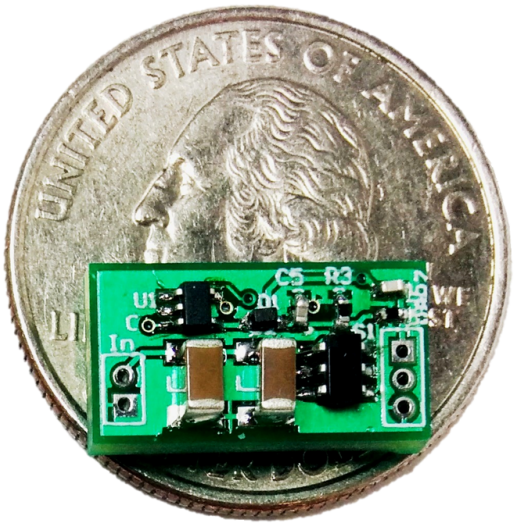

Fig. 5: Photo of the μMonitor prototype

We also have initialized an interrupt routine which will fire whenever switching count reaches 1000. Since, C2 is typically well over 1000 times the ceramic caps and α is 1.056, when the interrupt fires we can ensure that voltage across C2 has not changed more than 100mV. After the interrupt fired, the entire MCU will wake up and will sample the the voltage of C2, clear its timer and will go back to sleep again.

Finally, Rh, is the hold resistor which will hold $C1'$ and $C2'$ connected for a period of time after comparator output goes low. Rc, is the current limiter resistor that should be picked based on the maximum permitted charge current of charge reservoir which is usually mentioned in the datasheet.

IV. Results

In this section we first explain our measurement and experimental setup, and then walk the reader through each experiment conducted and its results. We used two different charge reservoirs, a 470mF supercap [9] and a 750mAh Nickel-Metal-Hydride (NiMH) battery. Figure 8 shows the accuracy and efficiency of our system as a function of the charging power.

All the ground truth measurements were collected using two 16-bit resolution National Instruments DAQs, one for monitoring incoming power and the other to monitor the charge power. The MCU was also running for the duration of each experiment, loaded with the firmware required to perform energy measurement using μMonitor, and was obtaining its power from the charge reservoir to ensure that its power draw is reflected in the reported efficiency of μMonitor. The measurement setup is shown in Figure 6 In each experiment the baseline for comparison is mostly estimated/actual energy (accuracy) and efficiency, which can be found in the below equations:

$$\text{Efficiency} = \frac{\text{Delivered Energy}}{\text{Input Energy}} \quad (20)$$

$$\text{Accuracy} = \frac{\text{Estimated Energy}}{\text{Delivered Energy}} \quad (21)$$

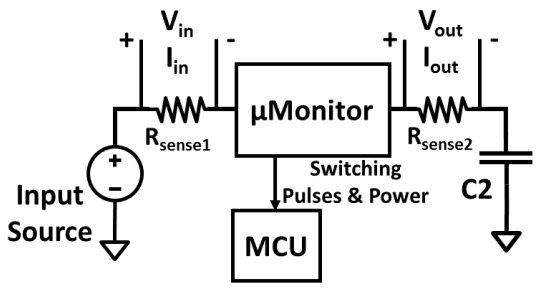

Fig. 6: Measurement setup for μMonitor evaluation

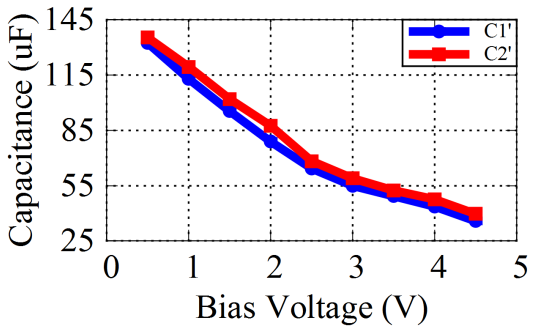

Fig. 7: Capacitance change of the deployed C1′ and C2′ versus bias voltage

A. Capacitance change versus bias voltage

As discussed before, one of the most important challenges that we should take into consideration is the fact that capacity of ceramic capacitors has a large variation over different bias voltages. So, before populating the ceramic capacitors on μMonitor we carefully measured their capacitance over different bias voltages to use that later as our reference for delivered power estimation. Figure 7 shows these result for C1′ and C2′ that has the nominal capacitance of 150uF. We can see that there is over 400% variation in capacitance as bias voltage goes from about zero to 4.5V.

B. Dynamic Threshold

The first experiment is to evaluate our system in comparison with the circuit shown in Figure 2. The target in this experiment is to see how much benefit we might get from dynamically setting the threshold instead of having it fixed. We first simulated the fixed-threshold circuit 2 and set the upper and lower threshold at 3.25V and 3.2V, respectively, letting C2 charge from 2V to 3.1V. We also charged C2 from 2V to 3.1V using the dynamic-threshold μMonitor implementation. Clearly, μMonitor outperformed the fixed threshold method in both accuracy and efficiency, as expected. The accuracy numbers for μMonitor were all between 98.5% and 88.3%, also the efficiency numbers were all between 94% and 95%.

C. Performance over long duration

In this experiment we connected a 470mF supercapacitor to the output of μMonitor, and the input was connected to a voltage source with a 15KΩ series resistor. We changed average input power by adjusting the input voltage source. During each experiment, the μMonitor charged its supercapacitor from approximately 2V to 3.1V. In a second experiment, we connected a 750mAh 2.5V Nickel-Metal-Hydride battery to the

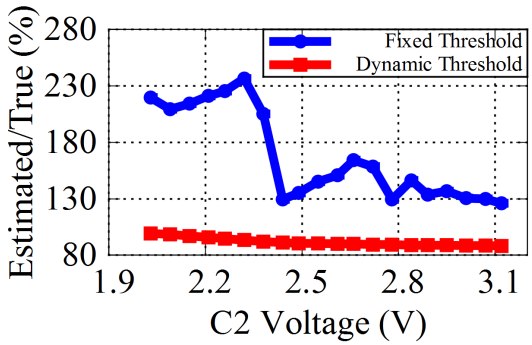

(a) Estimated/actual energy curve

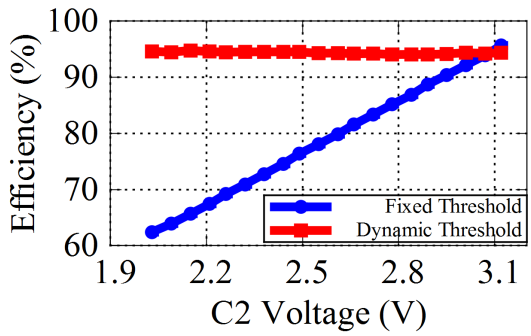

(b) Efficiency curve

Fig. 8: Figure 8a shows the estimated/actual energy comparison and Figure 8b shows the efficiency comparison between dynamically setting a threshold or have it fixed at 3.225V

μMonitor and did the same experiment. These trials were all done for the moderate input powers, from tens of micro-watts to less than a few milliwatts. All of the results are presented in Figure 9.

Based on this experiment we observed fairly constant accuracy when charging the battery, 96% to 97%, and a slightly variable accuracy for the supercapacitor case, from 91% to 102%. The reason for accuracy variation when charging a supercapacitor is that, since the supercapacitor's voltage is changing the instantaneous input power will change, and this will result in our system having variable accuracy. However, this variation is still within 10% of the actual delivered power, an acceptable accuracy overall for many use cases.

D. Dynamic range evaluation

In this experiment the target was to evaluate how our system operates with a very wide dynamic range of charging powers. We changed the input to μMonitor such that its corresponding average charge power varies from about 500nW all the way to about 15mW; power levels below 500nW were introducing significant errors in ground truth measurements, and a current limit on charge current set the upper value of 15mW. Figure 10 reflects these results.

Based on these results, we achieved an acceptable accuracy from about 500nW up to about 15mW, which is more than 4 orders of magnitude in dynamic range. Efficiency will drop as we go to lower power, which is expected since the leakage power overhead introduced by μMonitor will become increasingly significant compared to the input power as input power decreases. For the supercapacitor case we did not go

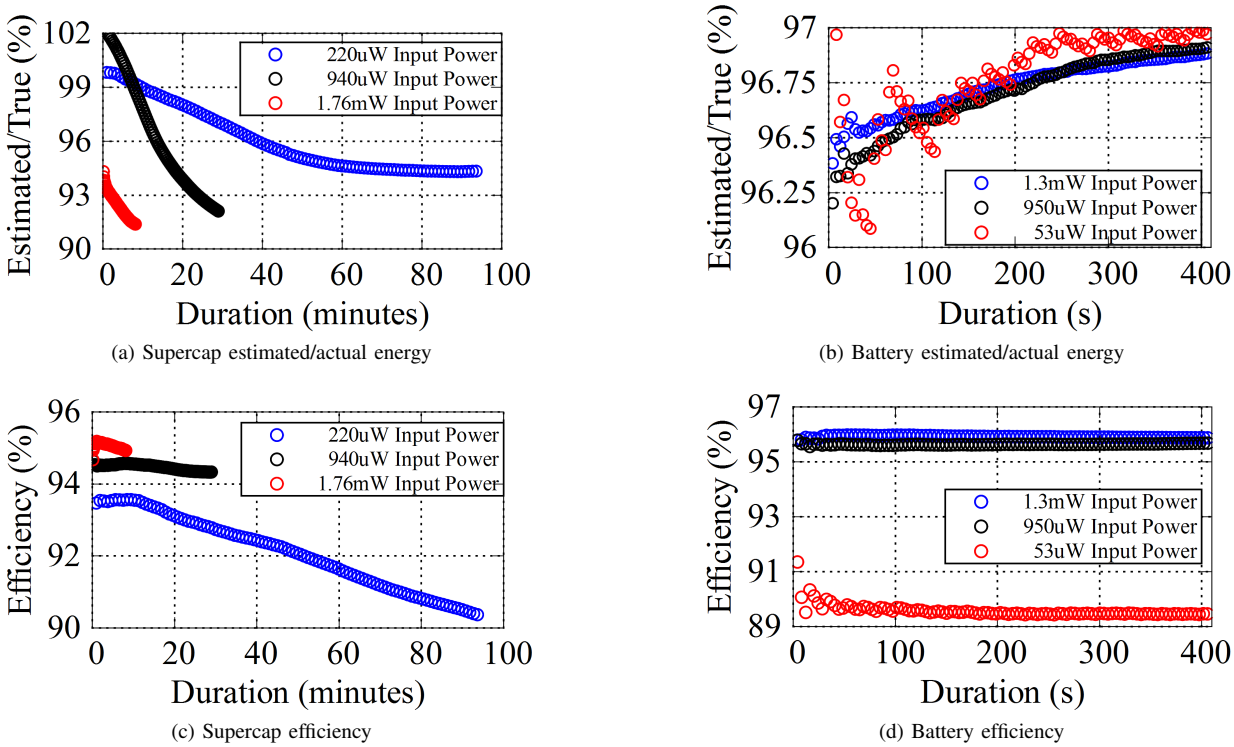

(a) Supercap estimated/actual energy

(b) Battery estimated/actual energy

(c) Supercap efficiency

(d) Battery efficiency

Fig. 9: Estimated/actual energy and efficiency of our system during long periods of time and for a 480mF supercap and a 2.5V 750mAh battery.

below 9uW charge power due to the large leakage power of the supercapacitor itself, which turned out to be about 8uW at 2.5V.

E. Monitoring harvested energy from an RFID reader

To characterize μMonitor as part of a real-world energy harvesting target platform, we connected the output of the harvester in a battery-free (RF harvesting) RFID camera [4] to the μMonitor input, and used a 470mF supercapacitor as the charge reservoir. An Impinj Speedway R420 RFID reader emitting 4W EIRP was used to power and communicate with the RFID camera platform. Our setup was placed in 12 different positions between 2ft and 23ft from the RFID reader antenna. Throughout each experiment we were logging input and output energy for μMonitor. Figure 11 shows the accuracy and efficiency results of this experiment.

Based on the results, we got an estimated energy which is within 6% of the actual delivered energy in the wild scenarios. As expected from the benchmark experiments, decreasing input power will result in decreasing efficiency, that is why efficiency is lower at longer distances.

F. Quiescent current of μMonitor

Finally, to measure the quiescent current of the μMonitor, we disconnected the charge reservoir from the μMonitor and, using a Source Measure Unit, we found the input current at which the μMonitor will be in equilibrium; meaning all of the DC voltage values of the μMonitor will remain constant. We did the same experiment for different DC bias voltages, and table I shows these quiescent current results.

We did this experiment only for the 2V to 3.5V range since this is the operating range for our MCU. Below this range would have excluded MCU power consumption from the result and above this range would have destroyed the MCU.

TABLE I: Quiescent current of μMonitor vs different bias voltages

Bias Voltage(V)	2	2.2	2.5	2.7	3	3.3	3.5
Quiescent Current(μA)	1.24	1.28	1.4	1.44	1.54	1.61	1.68

V. Applications

μMonitor can provide any low-power system with energy-awareness. It enables systems to make more intelligent operation decisions based on estimates of remaining energy. As an example, WISPCam [9] is a battery-free RFID camera that can take a picture and perform some basic image processing tasks [10] once it has accumulated enough energy to perform a single image capturing operation. But typically, WISPCam-like systems must harvest most of the time, and image (or video) capture will be triggered by an external sensor condition rather than an energy level condition. For example, in surveillance applications, WISPCam harvests energy most of the time, and uses a motion sensor to capture images whenever some motion is happening in the environment. So if WISPCam be aware of its stored energy, it can adjust its frame rate such that it can have pictures that are evenly spaced for a desired period of time.

As another example, consider NFC. The NFC-WISP is a wirelessly powered NFC tag data logger [11]. One of the applications of this tag is logging cold chain data. The user charges the tag for a period of time, and then attaches the tag on a temperature/shake sensitive product. With μMonitor, systems like NFC-WISP can plan energy expenditures, uniformly

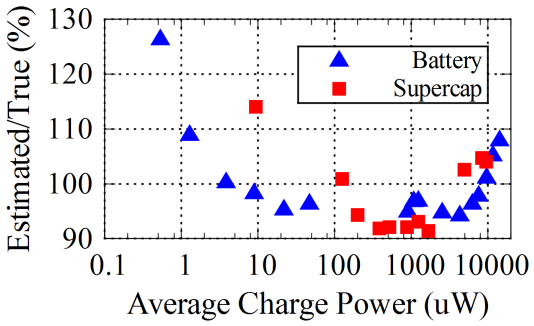

(a) Estimated/actual energy curve

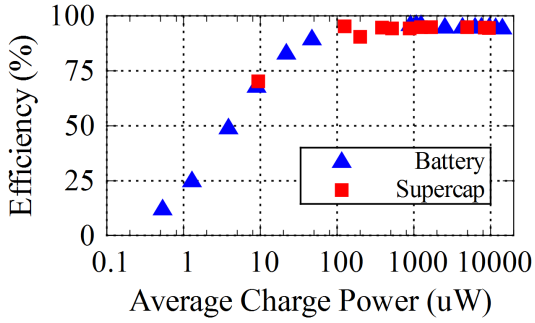

(b) Efficiency curve

Fig. 10: Figure 10a shows the accuracy (estimated/actual energy) of charge monitoring and Figure 10b shows the efficiency of μMonitor for a very wide range of charging power.

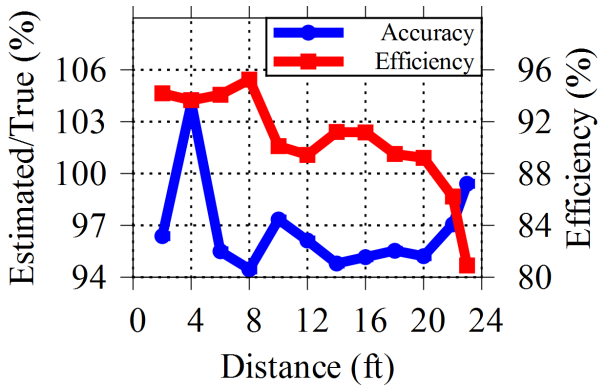

Fig. 11: μMonitor is connected to a WISPCam which has been placed at 12 different distances from an RFID reader that emits 4W EIRP. The estimated/actual energy and efficiency is shown in the graph.

spreading the samples out to cover a target time period, instead of collecting short bursts of data and then running out of energy for the most of chain duration. When more energy is available, the samples can be planned more densely.

VI. Conclusion

This paper presented the μMonitor, a low-power, in-situ, accurate power monitoring system. We first addressed the limitations of current power monitoring systems and why they are not suitable for ultra-low-power and energy scavenging devices. Leveraging the "Coulomb counting" technique, we developed a power monitoring prototype from off-the-shelf components. Also, we presented a mathematical model which firstly simplifies the energy estimation calculation with negligible introduced error, and secondly is appropriate for implementing on ultra-low-power MCUs without having floating point operation support.

The μMonitor was evaluated by computer simulation as well as benchmark experiments and wild scenarios. In all of the situations the μMonitor operated over four orders of magnitude range of charging power and estimated the delivered energy with high accuracy. Furthermore, the total quiescent current overhead of the μMonitor was less than 1.7μA up to a 3.5V operating voltage.

VII. Acknowledgements

This work is funded by Intel Science and Technology Center for Pervasive Computing (ISTC-PC), NSF award number CNS 1305072, and a Google Faculty Research Award. We would also like to thank our reviewers for providing helpful comments.

References

[1] A. Hergenroder and J. Furthmuller, "On energy measurement methods in wireless networks," in *Communications (ICC), 2012 IEEE International Conference on*, pp. 6268–6272, June 2012.

[2] X. Jiang, P. Dutta, D. Culler, and I. Stoica, "Micro power meter for energy monitoring of wireless sensor networks at scale," in *Information Processing in Sensor Networks, 2007. IPSN 2007. 6th International Symposium on*, pp. 186–195, April 2007.

[3] R. Zhou and G. Xing, "Nemo: A high-fidelity noninvasive power meter system for wireless sensor networks," in *Information Processing in Sensor Networks (IPSN), 2013 ACM/IEEE International Conference on*, pp. 141–152, April 2013.

[4] S. Naderiparizi, A. Parks, Z. Kapetanovic, and B. Ransford, "Wispcam: A battery-free rfid camera," in *RFID (RFID), 2015 IEEE International Conference on*, IEEE, 2015.

[5] J. Andersen and M. Hansen, "Energy bucket: A tool for power profiling and debugging of sensor nodes," in *Sensor Technologies and Applications, 2009. SENSORCOMM '09. Third International Conference on*, pp. 132–138, June 2009.

[6] "Tlv3691 comparator datasheet." http://www.ti.com/lit/ds/symlink/tlv3691.pdf. Accessed January 2016.

[7] "Adg80x datesheet." http://www.analog.com/media/en/technical-documentation/data-sheets/ADG801_802.pdf. Accessed January 2016.

[8] "Msp430fr5969 datasheet." http://www.ti.com/lit/ds/symlink/msp430fr5969.pdf. Accessed January 2016.

[9] "Murata 470mf supercap datasheet." http://www.murata.com/. Accessed January 2016.

[10] S. Naderiparizi, Y. Zhao, J. Youngquist, A. P. Sample, and J. R. Smith, "Self-localizing battery-free cameras," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '15, (New York, NY, USA), pp. 445–449, ACM, 2015.

[11] Y. Zhao, J. Smith, and A. Sample, "Nfc-wisp: A sensing and computationally enhanced near-field rfid platform," in *RFID (RFID), 2015 IEEE International Conference on*, pp. 174–181, April 2015.