



# EMPIOT: An energy measurement platform for wireless IoT devices

Behnam Dezfouli<sup>a,\*</sup>, Immanuel Amirtharaj<sup>a</sup>, Chia-Chi (Chelsey) Li<sup>a,b</sup>

<sup>a</sup> Internet of Things Research Lab, Department of Computer Engineering, Santa Clara University, USA

<sup>b</sup> Intel Corporation, Santa Clara, USA

## ARTICLE INFO

### Keywords:

Energy efficiency  
Testbed  
Accuracy  
Linux  
I2C driver  
802.15.4  
802.11

## ABSTRACT

Profiling and minimizing the energy consumption of resource-constrained devices is an essential step towards employing IoT in various application domains. Due to the large size and high cost of commercial energy measurement platforms, alternative solutions have been proposed by the research community. However, the three main shortcomings of existing tools are complexity, limited measurement range, and low accuracy. Specifically, these tools are not suitable for the energy measurement of new IoT devices such as those supporting the 802.11 technology. In this paper we propose EMPIOT, an accurate, low-cost, easy to build, and flexible, power measurement platform. We present the hardware and software components of this platform and study the effect of various design parameters on accuracy and overhead. In particular, we analyze the effects of driver, bus speed, input voltage, and buffering mechanism on sampling rate, measurement accuracy and processing demand. These extensive experimental studies enable us to configure the system in order to achieve its highest performance. We also propose a novel calibration technique and report the calibration parameters under various settings. Using five different IoT devices performing four types of workloads, we evaluate the performance of EMPIOT against the ground truth obtained from a high-accuracy industrial-grade power measurement tool. Our results show that, for very low-power devices that utilize 802.15.4 wireless standard, the measurement error is less than 3.5%. In addition, for 802.11-based devices that generate short and high power spikes, the error is less than 2.5%.

## 1. Introduction

The importance of low-cost and accurate energy measurement of IoT devices is justified by two important aspects: First, the percentage of energy consumed by connected devices is increasing due to the significant growth in the number of IoT devices. Gartner predicts the number IoT connected devices will surpass 20 billion by 2020 (Gartner). Second, IoT devices are mostly battery-powered or rely on energy harvesting. Therefore, it is important to measure, profile, analyze, and improve the energy efficiency of these devices to satisfy the QoS requirements of applications.

Analytical (and simulation-based) energy estimation tools multiply the time spent in each state (e.g., sleep, processing, transmission/reception) by the power consumed in that state. This approach, however, is not accurate due to the following reasons (Kaup et al., 2014; Gomez et al., 2012; Zhou and Xing, 2013; Trathnigg et al., 2008; Eriksson et al., 2009): (i) Most of the proposed models focus on simple wireless technologies such as 802.15.4 and LoRa (Wang et al., 2006; Martinez et al., 2015). However, as new technologies such as 802.11

and LTE are being adopted by IoT, it is important to profile the energy efficiency of devices using these complex technologies. Specifically, it is difficult to simulate and evaluate the properties of real-world environments (e.g., interference), physical layer, and MAC layer parameters, and study their effects on energy consumption (Li et al., 2013; Dezfouli et al., 2014, 2015). For example, while the energy consumption characteristics of 802.15.4 devices are mostly affected by MAC parameters, 802.11's physical layer parameters are diverse and significantly affect energy consumption. In addition, 802.11 offers sophisticated MAC mechanisms (such as access categories, frame aggregation, automatic power save delivery) where predicting their effect on energy consumption is considerably more difficult compared to 802.15.4's MAC. To show this complexity, for example, our testbed results reported in Fig. 1 illustrate the significant effect of 802.11 physical layer parameters on current consumption. The high variations of power consumption versus physical layer parameters justify the importance of extensive empirical measurements under various configurations in order to model and profile the energy characteristics of 802.11 devices. (ii) The energy consumption of all the operational modes of a given system-

\* Corresponding author.

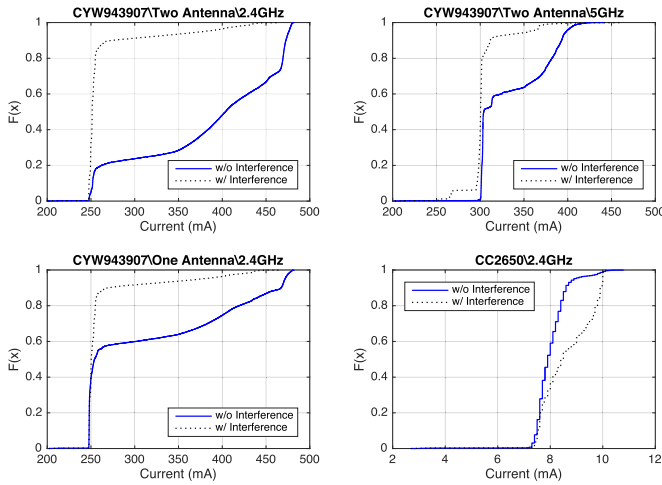
E-mail addresses: [bdezfouli@scu.edu](mailto:bdezfouli@scu.edu) (B. Dezfouli), [iamirtharaj@scu.edu](mailto:iamirtharaj@scu.edu) (I. Amirtharaj), [chelsey.li@intel.com](mailto:chelsey.li@intel.com) (C.-C. Li).

<https://doi.org/10.1016/j.jnca.2018.07.016>

Received 22 March 2018; Received in revised form 22 June 2018; Accepted 31 July 2018

Available online 14 August 2018

1084-8045/© 2018 Elsevier Ltd. All rights reserved.



**Fig. 1.** Empirical CDF of current consumption for an 802.11-based IoT device (CYW943907 (Semiconductor)) versus an 802.15.4 IoT device (CC2650).

on-chip (SoC) may not be available. For example, the datasheet may include only the average current consumption for nominal transmission power values. As another example, the energy consumption of processors depends on utilization level, frequency scaling, and I/O operations (Stathopoulos et al., 2008). (iii) IoT boards usually include a SoC and several peripheral components (e.g., ADC, sensors, memory). Therefore, even if the energy characteristics of the SoC are known, estimating the total energy consumption of the board is very challenging. (iv) Code structure, algorithms, and data structures, affect the energy consumption of the IoT system. For example, when a cache memory is present, the type of the data structures used affects cache performance when memory allocation is contiguous. (v) Analytical models usually do not take into account the cost of start-up energy. However, for example, the start-up energy cost of a sensor that is periodically woken up to collect a sample can significantly contribute to the total energy cost (Brusey et al., 2016). Despite these shortcomings, most of the research contributions on low-power wireless IoT systems rely on analytical and simulation-based energy estimation due to their simplicity.

Although the commercial energy measurement platforms provide very high accuracy, they are costly and bulky. For example, Keysight 34465A (Keysight Technologies, 2018) costs more than \$1300, includes 2 MB storage and supports maximum sampling rate 50 Ksps. Similarly, Keithley 7510 (Tektronix, 2018) costs more than \$3500, includes 2 MB storage and supports sampling rate up to 1 Msp. In addition, due to their size and cost, these devices are not viable solutions to monitor the energy consumption of a large number of IoT devices in a testbed. Furthermore, their programmability is very limited and inflexible. These shortcomings also apply to the Monsoon power meter (Monsoon Inc) (cost  $\approx$  \$800), which is widely-used by academia (Manweiler and Roy Choudhury, 2011; Serrano et al., 2015).

Due to these challenges, the research community has proposed several power measurement platforms. Based on their main shortcoming, these platforms are classified into the following categories: (i) *complex* (e.g. (Zhou and Xing, 2013; Stathopoulos et al., 2008; Jiang et al., 2007; Dutta et al., 2008; Andersen and Hansen, 2009; Naderiparizi et al., 2016)), (ii) *limited supported range* (e.g. (Trathnigg et al., 2008; Jiang et al., 2007; Haratcherev et al., 2008; Zhu and O'Connor, 2013; Hartung et al., 2016; Pötsch et al., 2017; Pötsch et al., 2014)), and (iii) *low accuracy* (e.g. (Gomez et al., 2012; Keranidis et al., 2014)). The platforms of the first category present complex circuitry, which makes the device costly and hard to build. For the second category, the power measurement platforms cannot be used with a wide variety of IoT devices due to their limited current measurement range, which is a few tens of milliamps. Specifically, the new generation of IoT devices utilizes wireless

technologies (such as 802.11) that result in current spikes as high as 700 mA. For example, Fig. 1 shows that the current consumption range of the 802.11 device is 46 times wider than that of the 802.15.4 device. Finally, if the accuracy of an energy measurement platform is low, then it cannot be used for effective study, development and debugging of IoT devices. In addition to these shortcomings, most of these platforms ignore the effect of voltage variation on energy measurement (Zhou and Xing, 2013; Trathnigg et al., 2008; Andersen and Hansen, 2009; Haratcherev et al., 2008). Furthermore, in terms of accuracy analysis, evaluations are very limited and mostly include only one IoT device type (Zhou and Xing, 2013; Jiang et al., 2007; Dutta et al., 2008; Andersen and Hansen, 2009).

In this paper we introduce EMPIOT, an accurate, low-cost, easy to build, and flexible power measurement platform. EMPIOT has two main components: a shield board, which includes a low-cost INA219 (INA219) energy monitoring chip, and a base board, which runs the controlling and data collection software. The shield board supports both current and voltage measurement within the operational range of various IoT devices. The software, written in C++, is composed of two threads for communication with the shield as well as computing energy and saving the data to a file. Despite the simplicity of its hardware, in this paper we show the effect of various design parameters on performance through extensive experimental studies. Specifically, we use two I2C drivers, namely BCM (C library for Broadcom BCM 2835) and Linux driver (Linux I2C Driver), and show that the BCM driver achieves higher speed, lower energy consumption, and predictable timing characteristics. We also evaluate the effect of input voltage on sampling rate and measurement accuracy. Our results show that reducing the shield's operational voltage increases the conversion time and lowers the sampling rate. In addition, we confirm that this platform does not require a very reliable power source to achieve high accuracy. In fact, the accuracy of current and voltage measurement is always within 100  $\mu$ A and 4 mV, respectively. Based on these results, for those IoT devices with sleep current consumption of less than 100  $\mu$ A, we propose a hybrid energy measurement model that combines empirical evaluation with analytical modeling. We also study the effect of batched and continuous file writes, as well as driver, on the energy consumption of EMPIOT. Through empirical evaluations and formulating a mathematical model, we show how the buffering mechanism used affects the power consumption of EMPIOT based on the number of processor cores.

In order to calibrate the platform, we have designed a *programmable calibration tool*, which enables us to generate currents and voltages in a wide range and precisely control the duration of each change. Using this tool, we calibrate EMPIOT for currents up to 800 mA. We study the accuracy of this platform using five different IoT devices and four types of loads. Our ground truth is a high-accuracy industrial-grade power measurement tool (Tektronix). Our results confirm that the measurement error is less than 3.5% for very low-power devices that use the 802.15.4 wireless standard and generate a peak current of 30 mA. In addition, the error is less than 2.5% for 802.11 devices when their current consumption surpasses 100 mA. We also show that neglecting voltage variations in the energy measurement process may result in up to a 0.5% increase in measurement error, especially for battery-powered 802.11-based IoT devices.

We emphasize that this paper does not intend to propose a complex power measurement tool. Rather, we study if and how the EMPIOT shield, or even an existing off-the-shelf INA219 breakout board (such as (INA219 Breakout Board)), can be used for accurate power measurement of a wide variety of IoT devices.

The rest of this paper is organized as follows: Section 2 presents the platform components and studies the effect of various parameters on design. The calibration methodology and results are presented in Section 3. Section 4 studies the accuracy of EMPIOT. We present related work in Section 5. We conclude the paper in Section 6.




Fig. 2. Components of EMPIOT.

## 2. Platform design

In this section we present the hardware and software design of EMPIOT. We also study the implications of design choices on performance through extensive empirical evaluations.

### 2.1. Hardware

EMPIOT is composed of a shield that is installed on and communicates with a Raspberry Pi (RPI), as shown in Fig. 2. The core of the shield is an INA219 (INA219 Zero-Drift), which is a low-cost ( $\approx \$2$ ) current and bus voltage monitoring chip. INA219 is available in two classes, INA219A and INA219B. We use INA219B due to its lower (0.5%) variations versus temperature. The energy draw of the chip is 1 mA, and it can operate using a 3–5.5 V supply.

INA219 measures bus voltage directly, and current is measured through digitizing the voltage across a shunt resistor. In other words, given a shunt value  $R_{shunt}$ , we can compute current using  $i = v_{shunt}/R_{shunt}$ , where  $v_{shunt}$  is the digitized shunt voltage value.

The ADC is a delta-sigma type and uses high frequency (500 KHz) to collect analog samples. After collecting samples, the ADC uses a low-pass digital filter for noise reduction, and a decimator averages the analog samples. The conversion ready bit is set automatically when a new sample is ready. The INA219 datasheet mentions that the duration of these operations for 12 and 9-bit sampling resolutions are 532–586  $\mu$ s, and 84–93  $\mu$ s, respectively. However, as we will show in Section 2.3.1, the actual sample preparation time is longer than these reported values.

The full scale voltage range supported across the shunt resistor is 40 mV. Since the ADC resolution is 12 bits, one LSB size for shunt voltage is  $\sigma_v^{shunt} = 40/(2^{12} - 1) \approx 10 \mu$ V. We have used  $R_{shunt} = 0.1 \Omega$  with 0.5% accuracy. Therefore, EMPIOT's current resolution is  $\sigma_i^{shunt} = \sigma_v^{shunt}/R_{shunt} = 10 \mu$ V/0.1  $\Omega = 100 \mu$ A for currents up to 400 mA. Depending on the maximum possible current, the power gain amplifier (PGA) can be configured to achieve the full-scale range through dividing shunt voltage by 2, 4, or 8, before digitization. Therefore, shunt voltage in four various ranges can be measured: [0, 40]mV, [0, 80]mV, [0, 160]mV, [0, 320]mV. The maximum supported bus voltage is either 16 V or 32 V, depending on the configuration applied. When configured for 16 V, the accuracy of bus voltage measurement is  $\sigma_v^{bus} = 16/(2^{12} - 1) \approx 4$ mV. Considering the nominal operational range of IoT devices, in this paper we assume current and voltage are less than 800 mA and 5.5 V, respectively.

The shield board communicates with the base board using the Inter-Integrated Circuit (I2C) bus. The minimum and maximum bus speeds supported by INA219 are 0.1 MHz and 2.5 MHz, respectively. We chose

RPI as the base board to configure the chip and collect the results because: (i) its I2C rate is fast enough to support the sample generation rate of INA219, (ii) the memory card enables power sampling for very long durations, (iii) it can be used for programming and debugging of the attached IoT device, and (iv) the existence of multiple communication technologies (Ethernet, WiFi) simplifies testbed setup and remote access.

### 2.2. Software

The pseudo-code of the software is given in Algorithm 1. At a high level, the program consists of the main function and two threads. The main function first initializes the I2C driver and configures the bus speed. The two drivers we used for this work are BCM (C library for Broadcom BCM 2835) and Linux i2c-dev (Linux I2C Driver) (version 4.4). In the rest of this paper we refer to these drivers simply as *BCM* and *Linux*. Next, INA219's configuration registers are programmed to adjust gain and resolution. The programmed gain value determines the maximum measurable current, and resolution refers to the number of bits per sample. The main function then initiates two threads: the sampler thread polls the conversion ready bit, and when set, it reads the bus voltage and shunt voltage values. In addition, the current system time (in nanoseconds) is read using the `clock_gettime()` system call. As this system call is used to add timing information to each sample, it is important to ensure its delay is negligible. To this end, we used a high sampling rate (500 MHz) logic analyzer and observed that the pin toggling delay using BCM library is 40 ns. Then, we toggled the pin right before and after the system call. This study showed that the delay of this system call is around 445ns with negligible variations; therefore, this provides a reliable mechanism to append a time stamp to each sample collected, which is later used for energy calculation.

The sampler thread also computes energy consumption through calling the `compute_energy()` function, which implements the Riemann integral approach. In addition to energy calculation, raw data are written to a file. To this end, we introduced the `sample_writer` thread to read data off a buffer and write to a file. We employed two different buffering mechanisms to implement *batched* and *continuous* file write operations: When using the *two-buffer* mechanism, the sampler thread and `sample_writer` thread use two different buffers. When a buffer is full, the sampler thread unblocks the `sample_writer` to flush that buffer, thereby entries are batched and then written to the file. Meanwhile, the sampler thread uses the other buffer. The other approach utilizes a single *circular buffer*. In this approach, as soon as a new entry is written to the buffer, the sampler thread signals the `sample_writer` thread to write that entry to the file. In this case, a mutex locks the buffer to avoid concurrent accesses. For both cases the buffers are implemented as fixed-size arrays allocated on stack memory to avoid the overhead of memory re-allocation. This also results in a higher utilization of cache memory due to the contiguous placement of array entries in random access memory.

The software supports raw data collection and energy measurement: (i) for a given time duration, (ii) after a certain number of samples were collected, and (iii) by receiving an external trigger to indicate the start and stop of measurement. In the trigger mode, falling and rising edge interrupts trigger the starting and stopping of power measurement. Using this feature, one can annotate an IoT device's code to start and stop the measurement at particular locations, thereby, providing fine-grained energy measurement of various operations such as encryption and transmission.

### 2.3. Design parameters

In this section we study the effect of various design parameters on performance in terms of sampling rate, sampling offset, accuracy, and energy consumption of the base board.

**Algorithm 1** EMPIOT's Software.

```

1 function main()
2   setup the I2C driver (BCM/Linux) and bus speed;
3   configure the INA219 gain and resolution;
4   create thread sampler;
5   create thread sample_writer;
6   return;
7 thread sampler()
8   while sampling is enabled do
9     while conversion ready == 0 do
10      check the bit;
11      read bus voltage and shunt voltage values;
12      compute_energy(prev_smp, new_smp);
13      if use two buffers then
14        if use_buffer1 == true then
15          add entry to buffer1;
16          if buffer1 is full then
17            signal the sample_writer thread;
18            use_buffer1 == false;
19          else
20            add entry to the buffer2;
21            if buffer2 is full then
22              signal the sample_writer thread;
23              use_buffer1 == true;
24        else if use a circular buffer then
25          lock the buffer;
26          add entry to the buffer;
27          unlock the buffer;
28          signal the sample_writer thread;
29 thread sample_writer()
30   if use two buffers then
31     if use_buffer1 == true then
32       flush buffer1 to the file;
33     else
34       flush buffer2 to the file;
35   else if use a circular buffer then
36     t
37     lock the buffer;
38     write entry to the file;
39     unlock the buffer;
40 function compute_energy(prev_smp, new_smp)
41   prv_power = (prv_smp.voltage) × (prv_smp.current);
42   new_power = (new_smp.voltage) × (new_smp.current);
43   time_diff = (new_smp.time) × (prv_smp.time);
44   new_energy = new_power × time_diff ;
45   tri = (new_power - prv_power) × time_diff / 2 ;
46   new_energy -= tri;
47   energy += new_energy;

```

**2.3.1. Sampling rate**

Fig. 3 shows sampling rate versus bus speed, driver, and input voltage, where “input voltage” refers to the power source of EMPIOT's shield. The main observations are as follows: First, the sampling rate is lower than the conversion rate supported by the chip. For example, for 12-bit conversion, the sampling rate is lower than the 1.8 KHz ADC conversion rate we reported in Section 2.1. Second, a lower voltage results in a lower sampling rate. Third, reducing the bus speed to as low as 200 KHz results in reducing the sampling rate. Fourth, the Linux driver affects sampling rate for 9-bit resolution. We study these observations in more details as follows.

In order to evaluate software overhead, we measured the time spent by the `sampler` thread between collecting a sample and the next polling of conversion ready bit. Using a high-speed logic analyzer, our results show that the processing overhead of the `sampler` thread is 0.46  $\mu$ s, which is negligible. To measure I2C read delay, we used a simple program that continuously reads two bytes from the shield board. After capturing I2C traffic using a logic analyzer, we compute the inter-

val between sending I2C addresses. Fig. 5 shows the results.<sup>1</sup> This figure reveals the effect of bus speed and driver on read delay. In particular, the BCM driver achieves a faster and more stable I2C performance compared to the Linux driver. In fact, for a given baud rate, Linux I2C read is at least 20  $\mu$ s slower than BCM. In order to justify the lower performance of Linux driver compared to BCM, we used `strace` to capture the system calls made by the software. Our evaluations show that when using BCM, the number of system calls is always fixed (exactly 181) and does not depend on the sampling rate. We also observed that these system calls are only made during the initialization of the BCM driver, after which the software directly communicates with the driver. In contrast, when using the Linux driver, all the I2C communication requests pass through the kernel; therefore, the number of system calls depends on the sampling rate. Fig. 5 also shows that the I2C delay of Linux driver has higher variations compared to BCM. The

<sup>1</sup> Please note that we did not report the results for bus speed 2500 KHz and input voltage 3.3 V because we observed a very unreliable I2C communication in this condition. We believe that INA219 cannot keep up with this high clock rate when the voltage is 3.3 V.

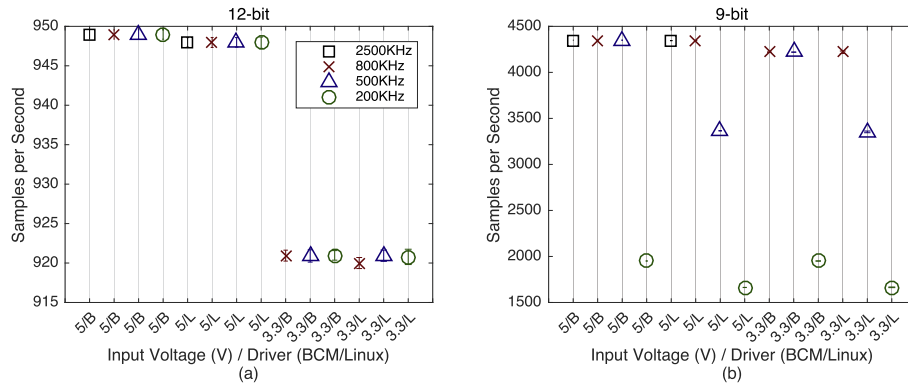


Fig. 3. Sampling rate for (a) 12-bit resolution and (b) 9-bit resolution. Error bars show 95% confidence interval.

**Table 1**  
Polling rate of conversion ready bit using BCM and Linux drivers for 12 and 9-bit resolution.

		Polling Rate			
		2500 KHz	800 KHz	500 KHz	200 KHz
12-bit	BCM	45	15	9	3
	Linux	23	9	6	2
9-bit	BCM	9	2	1	1
	Linux	4	1	1	1

average range of variations for Linux I2C read delay is 22  $\mu$ s, and for BCM this value is 4  $\mu$ s. These results indicate that the Linux driver does not achieve a stable communication rate with mission-critical and high rate sensors, such as those used in medical and industrial applications. The effect is also obvious when using 9-bit sampling (cf. Fig. 3(b)). For example, when using 500 KHz bus speed, using the Linux driver reduces the sampling rate to 3360, compared to the 4350 samples collected per second when the BCM driver is in use. From the software point of view, the overhead of the Linux driver affects the number of times the conversion ready bit is polled per sample collection round. Table 1 reports the results. Therefore, the Linux driver falls behind the sample conversion rate when the sampling rate is high.

We next analyzed the time interval between a change in input and reading the corresponding value, which is referred to as *sampling offset*. In order to identify the causes of sampling offset, we used the experiment shown in Fig. 4. To introduce quick and predictable changes in power, we have used a pin toggle using an ARM-Cortex R4 board (CWY9443907 (Cypress Semiconductor)) which sets a pin from high to low (i.e., 3.3V–0V) in 50 ns. For this experiment, the “Output Pin” in Fig. 4 is initially high (3.3V). At particular intervals, the pin is set to low (0V) and remains in this state for 2ms. We compute the sampling delay offset by measuring the interval between setting a pin to low and collecting the corresponding sample. To this end, a logic analyzer logs the status of “Output Pin” as well as I2C communications. Fig. 6 shows the measured values. Since the I2C read delay is independent of the input voltage (as Fig. 5 shows), the results of Fig. 6 indicate a longer conversion time when using a lower voltage value, i.e., 3.3V. For example, increasing conversion time by 64  $\mu$ s decreases the number of samples collected per second by 47. As the power measurement chip uses a delta-sigma ADC, we believe that the lower voltage value slows down the operation of the decimator and averaging circuitry.

2.3.2. Measurement accuracy

In this section we analyze the effect of input voltage variations on accuracy. We consider the following voltage sources to run the INA219 of EMPIOT’s shield board:




Fig. 4. In order to measure I2C read delay, the logic analyzer captures I2C communications. To measure sampling delay offset, the logic analyzer captures both “Output Pin” status and I2C communications.

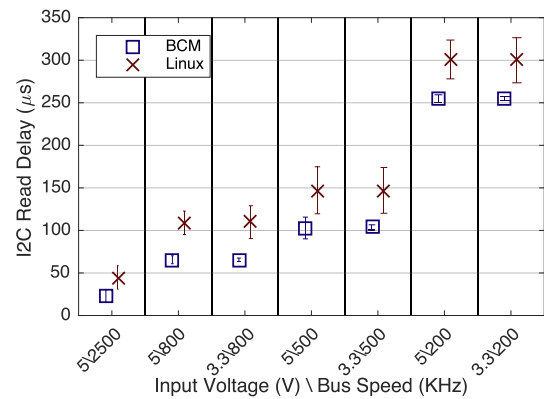


Fig. 5. I2C read delay measured as the time interval between issuing a read command and reception of requested bytes. Error bars show 95% confidence interval.

- External (5VExt): A 5V external power (Keysight Technologies).
- 5V from RPi (5VRPi): The 5V pin of RPi header.
- 5V from RPi with a load (5VRPi w/Load): The 5V pin of RPi header. In addition, the RPi’s USB port is connected to a Cypress CYW943907 IoT device. Its idle energy consumption is about 100mA, but it periodically wakes up and sends ping packets that result in up to 400mA current consumption.

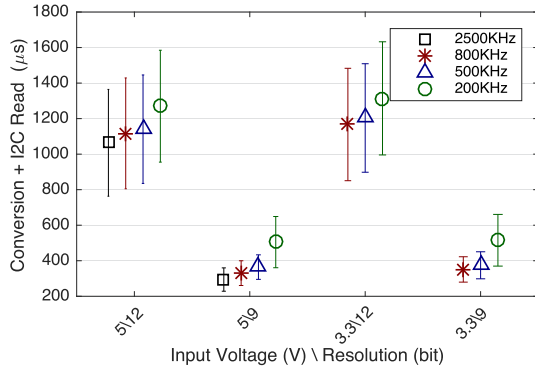


Fig. 6. Sampling offset, which reflects the interval between a change in input current and reading the corresponding value. We used the BCM driver for this experiment. Error bars show 95% confidence interval.

- 3.3 V from RPi (3.3VRPi): The 3.3 V pin of RPi header.

We have used an industrial-grade DMM (Tektronix) with a sampling rate of 500 Ksps to measure the variations of these power sources. The 95% variations of these sources are as follows: 5VExt: 2 mV, 5VRPi: 103 mV, 5VRPi w/Load: 300 mV, 3.3VRPi: 24 mV. Caused by operating system processes and EMPIOT software, we observe that RPi as a power source exhibits higher variations compared to 5VExt. Nevertheless, in order to simplify the design and reduce its cost, we are interested in confirming if RPi can be used as the source of power for EMPIOT's shield board.

In order to measure accuracy across a wide range, we used three different fixed loads: 200 µA, 5 mA and 100 mA. These loads are generated by connecting the shield's output to three different resistors. Fig. 7 presents the effect of voltage variation on the accuracy of bus voltage and current measurement. These results show the variability of measurements caused by factors such as electromagnetic interference and white noise. However, the important observation is that, irrespective to the source of power and load value, the measurement error of EMPIOT is always less than 4 mV and 0.1 mA for bus voltage and current, respectively. The reported error ranges comply with the values we previously mentioned for INA219 (cf. Section 2.1).

### 2.3.3. Overhead analysis

In this section we investigate the energy consumption of the RPi running EMPIOT's software. The two main causes of power consumption are file write and polling the shield board. In these experiments we measure the energy consumption of EMPIOT boards using either RPi3 (Raspberry Pi 3) or RPiZW (Raspberry Pi Zero with WiFi).<sup>2</sup> In addition to energy, we have logged the time of writing each sample to the file through toggling a pin and logging pin activation times. This technique avoids introducing extra software overhead to log file write instances. To remove the variations caused by Ethernet/wireless communication, we used UART to communicate with the RPi board whose energy is being measured. For example, for RPi3, when using UART instead of Ethernet, in addition to reducing power variations, current consumption is reduced by about 30 mA. In our implementation, each sample entry is 16 bytes: 8 bytes to store timestamp (i.e., the `timespec` structure), 4 bytes for bus voltage, and 4 bytes for current.

Fig. 8 shows the power consumption trace for RPi3. Vertical red bars indicate file write activities. As it can be observed, using the circular buffer mechanism results in continuously writing to the flash memory, which in turn increases the energy consumption. For example, for 9-bit resolution and BCM driver, the base power is increased by about

0.05w when using the circular buffer mechanism (compare sub-figure (a) and (c)). We have extracted a similar set of traces for RPiZW, and we observed that this increase is about 0.008w for this board.

Using the traces collected for RPi3 and RPiZW, Fig. 9 summarizes the effect of sampling resolution, driver, and buffering mechanisms on energy consumption. In order to prepare this figure, although the power consumption of the circular buffer mechanism is readily available, we need to compute the energy consumption of the two-buffer mechanism based on the energy consumption of polling and file write, as follows. The `sample_writer` thread is activated when a buffer is full. We refer to the time required to fill a buffer as  $t_b$ . After  $t_b$  seconds of buffering samples, the `sample_writer` thread is activated to write the samples to a file. It should be noted that, during file write, the device is also buffering samples. We refer to the file write duration as  $t_{wb}$ . Based on these values, we can compute the average energy consumption per second as

$$E = E_b + E_{wb} = p_b \times \frac{t_b - t_{wb}}{t_b} + p_{wb} \times \frac{t_{wb}}{t_b} \quad (1)$$

where  $p_b$  is the power consumption of buffering operation, and  $p_{wb}$  is the power consumption of writing and buffering simultaneously. The two variables  $t_b$  and  $t_{wb}$  are computed as follows. The duration of filling a buffer of size  $L_b$  samples is  $t_b = L_b/R_s$ , where  $R_s$  is the sampling rate. The `sample_writer` thread requires  $t_w = L_b \times L_s/W$  seconds to transfer  $L_b$  samples to a file, where  $W$  is the file write speed in bits per second and  $L_s$  is the length of each sample in bits. Therefore, the energy consumption of the two-buffer mechanism is

$$E = \frac{p_b \times W + L_s \times R_s \times (p_{wb} - p_b)}{W} \quad (2)$$

As Fig. 9 shows, the preferred buffering mechanisms for RPi3 and RPiZW are two-buffer and circular buffer, respectively. Since RPi3 has a multi-core processor, it schedules the `sample_writer` thread on a different core than that of the `sampler` thread. Although scheduling the `sample_writer` thread on a separate core increases processor power consumption, the speed of file write process is increased. On the other hand, RPiZW has a single core, thereby, the `sample_writer` thread requires more time to write the buffered data to the file. The longer  $t_{wb}$  duration of RPiZW increases the impact of  $p_{wb}$ . However, in the circular buffering mode the processor writes small amounts of data between polling instances, therefore, the energy consumption of circular buffer is less than the two-buffer mechanism.

Fig. 9 also reflects the higher energy consumption of the Linux I2C driver. For example, on RPi3, when using two-buffer and 12-bit resolution, using the Linux driver increases the base power to 1.38 W, as compared to the 1.26 W achieved with the BCM driver. As we discussed earlier, I2C communication through the Linux driver results in a significantly higher number of system calls, which increases processing load.

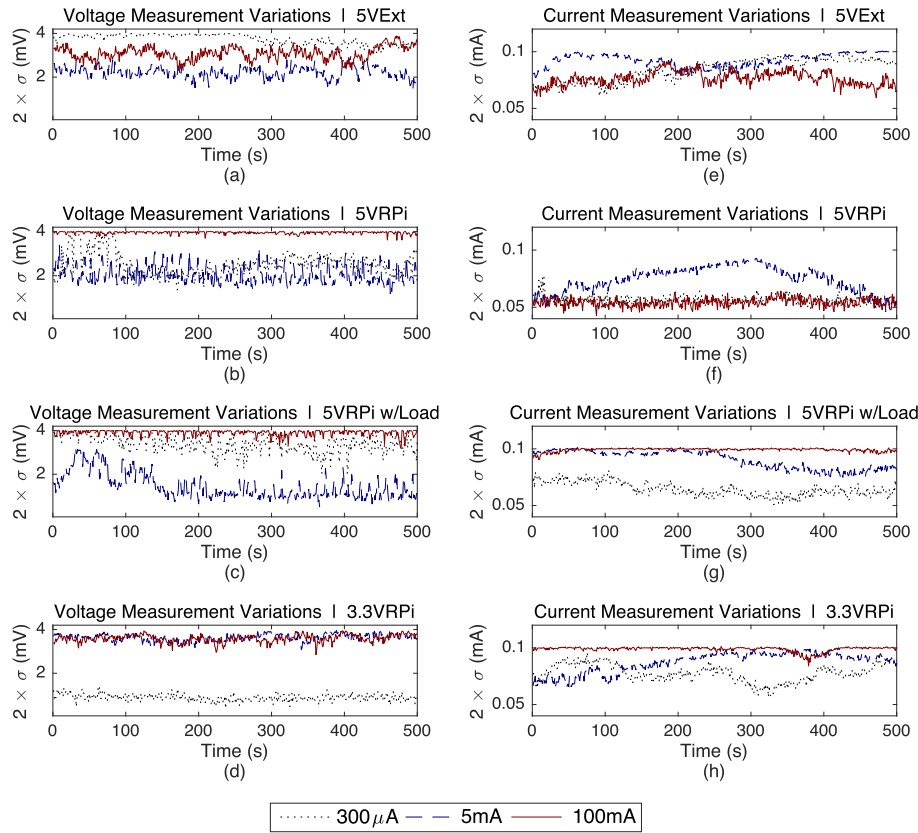
### 2.4. Overcoming the two limitations of EMPIOT

As mentioned earlier, the current measurement resolution of the proposed platform is 100 µA. However, since IoT devices support multiple low-power modes, current draws less than 100 µA cannot be measured. Our solution to this problem is to use a *hybrid energy calculation model* that computes the energy consumption of low-power modes similar to analytical energy estimation models. Specifically, we require the IoT device to inform EMPIOT right before and after each transition into a low-power mode. In this case, energy is computed as follows,

$$\sum_{\forall p_{ps}[i] \in P_{ps}} \sum_{\forall (t_s, t_e) \in P_{ps}} (t_e - t_s) \times p_{ps}[i] + \sum_{\forall S_j \in S} \Delta_j p_j \quad (3)$$

where  $P_{ps}$  is the set of all power saving modes that have a power consumption less than 100 µA,  $p_{ps}[i]$  is power saving mode  $i$ ,  $t_s$  is the start of power saving mode,  $t_e$  is the end of power saving mode,  $S$  is the set of samples collected during the normal operation (i.e., no power saving mode),  $S_j$  is a sample  $j$  collected during normal operation, and  $\Delta_j$  and  $p_j$

<sup>2</sup> The RPi3 used in this paper is based on a 900 MHz BCM2837, and the RPiZW is based on a BCM2835 SoC.



**Fig. 7.** The measurement variations of voltage (sub-figure (a) through (d)) and current (sub-figure (e) through (h)) over time for three different loads and various input sources. Sampling resolution is 12 bit. The y-axis is the 95% range of variations.

refer to the duration and power of sample  $S_j$ , respectively. Please note that during the power saving modes, power samples (i.e.,  $S_j$ ) are not used for energy measurement. This technique of capturing low power consumption is accurate and easy to implement due to two reasons: First, using a constant value as the power consumed in a low-power mode is reasonable because these modes refer to sleep states during which power consumption is stable. Second, the number of low-power states is usually limited to two or three modes. For example, CC2650 (CC2650) offers two sleep modes: *standby* and *shutdown*. During the standby mode, real-time clock (RTC) is running and the contents of RAM and CPU are retained. The device wake ups from the shutdown mode through a trigger. The energy consumed by these two modes are 1  $\mu$ A and 100 nA, respectively. For those devices that the power consumption of their sleep modes is not known *a priori*, a DMM could be used to extract the respective values, which are then hardcoded into EMPIOT. We will evaluate the effectiveness of the hybrid energy estimation technique in Section 4.

Another limitation of EMPIOT is its warm-up time. Our studies show that the first 3 to 5 samples collected after initialization are not reliable. Therefore, when using 12-bit resolution, the first 5 ms of an operation cannot be measured. This limits the minimum duration of energy measurement. This problem, however, can be simply addressed in software: instead of actually turning on and off the shield board per measurement, we can ignore the samples collected when energy measurement is inactive.

### 3. Calibration

Factors such as the resistance of the shunt resistor path, inaccuracy of shunt resistor, and ADC non-linearity cause differences between the EMPIOT's measurements and a ground truth. Thereby, calibration is an essential part of the design. To this end, the purpose of current calibra-

tion is to find function  $f_A(i_e)$  so that

$$i_a = f_A(i_e) \quad 0 \leq i_a \leq i_{\max} \quad (4)$$

where  $i_a$  is the actual current draw,  $i_e$  is the current reported by EMPIOT when the actual current is  $i_a$ , and  $i_{\max}$  is the maximum current draw of the IoT device. Similarly, the purpose of current calibration is to find function  $f_V(v_e)$  so that

$$v_a = f_V(v_e) \quad 0 \leq v_a \leq v_{\max} \quad (5)$$

where  $v_a$  is the actual voltage,  $v_e$  is the voltage reported by EMPIOT when the actual voltage is  $v_a$ , and  $v_{\max}$  is the maximum voltage supported by IoT device.

Accurate estimation of  $f_A(i_e)$  and  $f_V(v_e)$  requires a load that can generate current and voltage values across the supported measurement range of EMPIOT. Although current variations can be generated by using an IoT device, this method does not result in an accurate calibration due to the following reasons: First, the duration of a change in current draw may not be long enough to match the samples collected. In other words, due to the fast variations of current as well as the difference in the sampling offset of the DMM and EMPIOT, correlation of the samples collected by the two devices is very challenging. For example, when the current draw changes suddenly, some values may not be captured by EMPIOT due to its lower sampling rate compared to a DMM. Furthermore, DMM and EMPIOT might report the variations with different time offsets. Therefore, subtracting the pairwise values does not reflect a realistic measurement error. If the error of EMPIOT versus DMM is in fact a linear function, then the aforementioned calibration errors may prevent us from finding a linear fit or cause a linear function that its slope is higher or lower than the real value. In addition, the IoT board may not cover the supported current range of EMPIOT, which results in calibration gaps. In fact, since turning on each component (such as RF transceiver) results in a jump in energy

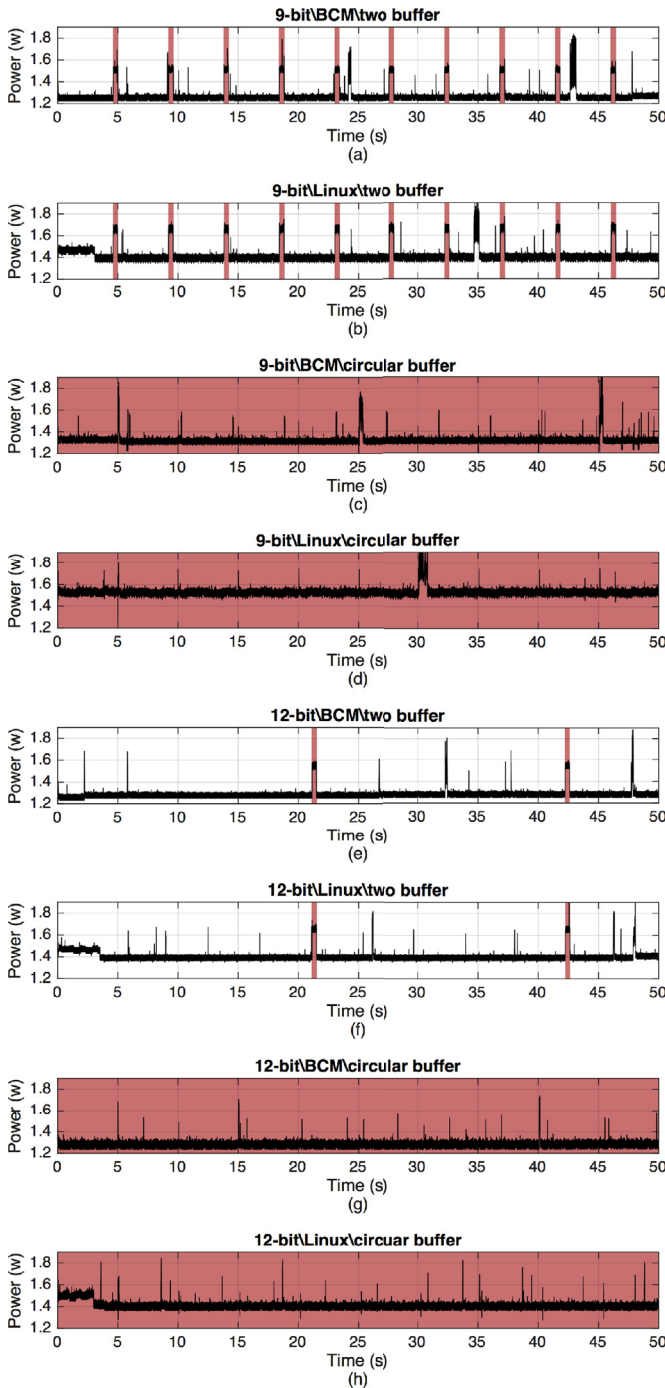


Fig. 8. The power trace of RP3 versus sampling resolution, driver, and buffering mechanisms.

consumption, it is almost impossible to generate a linear increase. Another solution is to use a potentiometer as the load. However, similar to using an IoT device, we cannot predict the transition and duration of drawing a particular current value, therefore, the measurement offsets affect calibration error. Furthermore, potentiometers usually support low current values, and the calibration for currents higher than 100 mA requires an expensive potentiometer. Due to these limitations, the existing works perform calibration either: (i) manually by using fixed resistor values (Trathnigg et al., 2008; Dutta et al., 2008; Haratcherev et al., 2008), or (ii) by using expensive equipment (Naderiparizi et al., 2016; Pötsch et al., 2017; Lim et al., 2013).

To address the aforementioned challenges and simplify the calibration process, we have designed a low-cost, accurate, and *programmable calibration tool*, which provides dynamic voltage and current ranges. This calibration tool is in fact a programmable load where its resistance and timing characteristics are controllable through a software (written in Python) running on a RPi. This software controls reconfiguration frequency and records output settling time between two consecutive configurations. Specifically, if reconfiguration frequency is  $t$ , the software records the output settling times at  $n \times t + t/2$ , where  $n \in \mathbb{Z}^+$  corresponding to all of the supported output values. Recording output settling instances enables us to correlate the measurement values of DMM and EMPIOT when the load is stable. Furthermore, this feature prevents the need for accurate time synchronization of DMM and EMPIOT.

The calibration tool uses an 8-bit digital potentiometer, AD5200 (Analog Devices AD5200), which has a maximum resistance of 10 k $\Omega$ . The maximum current output of this digital potentiometer can be expressed as  $I_{potmax} = V_{in}/R_w$ , where  $I_{potmax}$  is the maximum current supported by the digital potentiometer when it is programmed to the minimum value,  $R_w$  is the constant wiper resistance when the digital potentiometer is programmed to 0, and  $V_{in}$  is the voltage input to the calibration tool. Since the resistance range of the digital potentiometer is limited, we have added a number of resistors in parallel in order to extend the supported current range. Therefore, based on Kirchoff's law, the total maximum current of calibration tool is expressed as follows,

$$I_{max} = I_{potmax} + \sum_{j=1}^n I_j \quad \forall n \in \mathbb{Z}^+ \quad (6)$$

where  $I_{max}$  is the maximum supported current of the calibration tool,  $\sum_{j=1}^n I_j$  is the sum of the currents that  $n$  resistors can support, and  $I_j$  is the current that a resistor  $R_j$  carries. In our design, when  $V_{in} = 5V$ , the maximum current output of calibration tool is 1 A. To dynamically adjust the resistance of the calibration tool, resistors  $R_j$  are attached to a switch network that is implemented by four digital ADG1612 (Analog Devices ADG1612) switches. Therefore, although the digital potentiometer cannot handle currents higher than 20 mA, our design supports high current ranges by switching on and off the resistor paths. Specifically, a new resistor path is enabled whenever a 20 mA or a 100 mA increase in current is required, and the digital potentiometer is used to fine-tune current between range 0–20 mA. By programming the resistance value, the calibration tool can generate various current and voltage values. The calibration tool's software programs AD5200 and ADG1612 through SPI and GPIO interfaces.

The minimum current resolution supported by the calibration tool, denoted as  $I_{res}$ , depends on two parameters of the digital potentiometer: maximum resistance  $R_{max}$ , and the number of programmable bits. The equation that determines the digitally-programmed output resistance is expressed as follows,

$$R(x) = \frac{x}{2^n} \times R_{max} + R_w \quad 0 \leq x \leq 2^n \quad (7)$$

where  $x$  is the value programmed into the digital potentiometer,  $n$  is the number of bits supported by digital potentiometer,  $R_w$  is the constant wiper resistance, and  $R_{max}$  is the digital potentiometer's maximum resistance. Therefore, the minimum resolution of current output is expressed as follows,

$$I_{res} = \frac{(R(x) - R(x-1))}{R(x) \times R(x-1)} \times V_{in} \quad 0 < x \leq 2^n \quad (8)$$

Based on the components used and Equation (8), the minimum current output is 0.476 mA and the minimum current resolution is 1.82  $\mu$ A. In addition, minimum voltage output is 0.06 mV and minimum voltage resolution is 0.01  $\mu$ V. The wide range and high resolution of the calibration tool enable us to calibrate both current and voltage within the operating range of various IoT devices. It is worth mentioning that the entire calibration solution costs about \$100, including manufacturing costs. This is less than 1% of the cost compared to current commercial solutions.



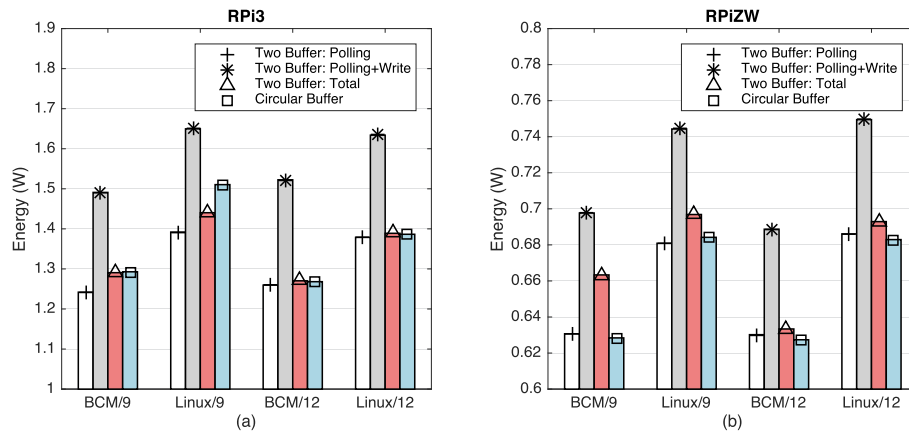


Fig. 9. The effect of sampling resolution, driver, and buffering mechanisms, on the energy consumption of RPi3 and RPiZW running EMPIOT software per second.

The calibration tool generates line interrupts to trigger the start and stop of measurements by EMPIOT and DMM. As mentioned earlier, EMPIOT can be started and stopped through line interrupts. To trigger the DMM, we have used a high-accuracy and programmable device (Tektronix), which exposes several programmable GPIO pins. We have developed a Python script using SCPI (Standard Commands for Programmable Instruments) to configure the sampling rate and enable the DMM to start and stop sampling based on the line interrupts received. This script communicates with the DMM through a TCP/IP connection and transfers the sampled data from the DMM buffer to a PC when a measurement completes.

EMPIOT’s shield includes a jumper to enable current flow from input to output (cf. Fig. 2). For calibration purposes we have removed this jumper and connected the pins to a DMM; thereby, EMPIOT and DMM measure current simultaneously. Instead of pairwise comparison of the traces collected by EMPIOT and DMM, we use the timing data logged by the programmable load. As the timing data reflects load stability instances, we can safely compare the two closest entries of the traces collected by DMM and EMPIOT.

One of the goals of this paper is to show if an off-the-shelf INA219 breakout board can be used instead of EMPIOT’s shield with the software and configuration parameters proposed in this work to achieve a high level of accuracy. To this end, in addition to reporting calibration data for EMPIOT’s shield, we have used an INA219 breakout board (Adafruit INA219 Current Sensor Breakout) to study the effect of hardware design on calibration. This breakout board has been installed on a bread board and communicates with a RPi running EMPIOT’s software.

Fig. 10 shows the measurement errors of three EMPIOT boards and three breakout boards conducted in a normal indoor temperature 25 °C. Please note that the value above each figure refers to: (i) left value: the maximum supported current configured through the programmable gain amplifier (PGA), and (ii) right value: the input voltage. As it can be observed, EMPIOT’s shield presents lower error compared to the breakout boards. Since INA219 measures current through a shunt resistor, the distance and impedance of the circuit path between the resistor and chip highly affect measurement accuracy. Therefore, the higher accuracy of EMPIOT’s shield is due to the thicker and shorter path used. In terms of input voltage, using 5 V instead of 3.3 V slightly reduces error. In addition, Fig. 10 shows that the error of EMPIOT’s shield increases linearly versus current. However, the breakout boards’ error shows a quadratic behavior for currents beyond 300 mA. For both these cases, instead of using a calibration table, we simply find the best fitted curves. Table 2 reports the calibration values for these boards.

In addition to current, the voltage measured may not reflect the actual bus voltage. We used the calibration tool to generate a variable voltage in the range of 2V–5.5V. Our results show that the error of voltage measurement is a fixed offset. In fact,

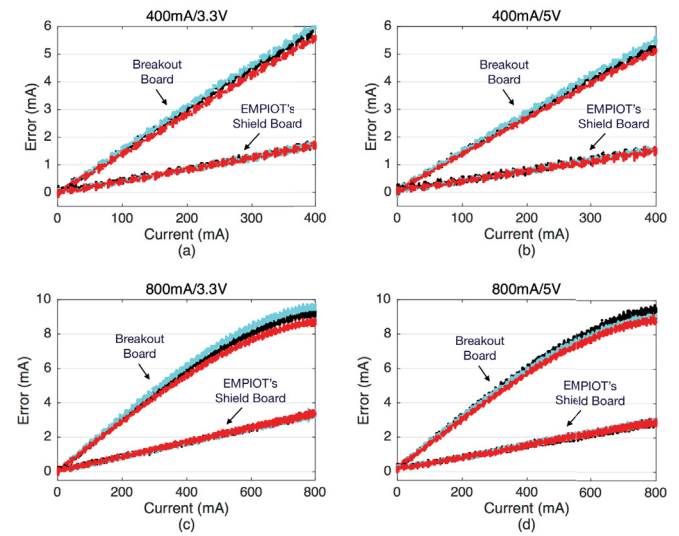


Fig. 10. Current measurement error ( $i_a - i_e$ ) versus ground truth. The value above each figure shows the maximum supported current (left side) and input voltage (right side).

EMPIOT’s voltage measurements versus DMM results in a linear function  $v_a = f_v(v_e) = v_e + 0.027$ . For the breakout board the calibration function is  $v_a = f_v(v_e) = v_e + 0.097$ .

#### 4. Performance evaluation

In this section we study the effectiveness of design parameters and calibration on accuracy when EMPIOT’s shield and the breakout board are used. Based on the results reported in Section 2, the following setting is used for both EMPIOT’s shield and the breakout board: (i) BCM driver is used, (ii) bus speed is 2500 KHz, (iii) INA219’s voltage is 5 V, and (iv) file writes are batched. Our ground truth is the energy measured by two high accuracy DMMs (Tektronix) that record current and voltage with 500 Ksps sampling rate and 18-bit resolution. A Python script programs the DMMs in trigger modes. Therefore, when an IoT device begins its operation, EMPIOT as well as the two DMMs start energy measurement. The three main operations performed by the IoT devices are *sleep*, *software encryption* and *transmission*. Please note that before each transmission the transceiver listens to the medium due to employing carrier-sense multiple access (CSMA) (Dezfouli et al., 2015). The encryption operation is used to generate a high processing load. Depending on the workload used, each node transitions between the available states every 500 ms. We introduce four types of

**Table 2**  
Current calibration functions.

EMPIOT	400mA/3.3 V	$i_c = f_A^{-1}(i_a) = 0.9957 i_a$	RMSE < 0.00005	R-Square = 1
	400 mA/5 V	$i_c = f_A^{-1}(i_a) = 0.9963 i_a$	RMSE < 0.00005	R-Square = 1
	800mA/3.3 V	$i_c = f_A^{-1}(i_a) = 0.9949 i_a$	RMSE < 0.0024	R-Square > 0.9998
	800 mA/5 V	$i_c = f_A^{-1}(i_a) = 0.9956 i_a$	RMSE < 0.0016	R-Square = 1
Breakout Board	400mA/3.3 V	$i_c = f_A^{-1}(i_a) = 0.9853 i_a$	RMSE < 0.00006	R-Square = 1
	400 mA/5 V	$i_c = f_A^{-1}(i_a) = 0.9869 i_a$	RMSE < 0.00006	R-Square = 1
	800mA/3.3 V	$i_c = f_A^{-1}(i_a) = 0.0079 i_a^2 + 0.9816 i_a$	RMSE < 0.0028	R-Square > 0.9998
	800 mA/5 V	$i_c = f_A^{-1}(i_a) = 0.0074 i_a^2 + 0.982 i_a$	RMSE < 0.0061	R-Square > 0.9994

loads:

- Workload 1: All the three operations are included,
- Workload 2: Includes send and sleep operations,
- Workload 3: Includes encryption and sleep operations,
- Workload 4: Includes encryption and send operations.

We have used five different IoT boards with various energy characteristics. The first device is a TI SensorTag CC2650 (CC2650 SimpleLink Multistandard Wireless). The CC2650 SoC includes an ARM Cortex-M3 processor and supports IEEE 802.15.4 standard. Since the energy consumption of this device in sleep mode is 1  $\mu$ A, we used this board to measure the accuracy of EMPIOT for profiling the energy of very low-power devices. Specifically, we are interested in seeing the effectiveness of the hybrid energy measurement technique proposed in Section 2.4 in terms of accuracy. We used TI-RTOS as the operating system running on this device. During the transmission state, the device sends 30-byte 802.15.4 packets as fast as possible. The small packet size introduces quick variations in power consumption.

In order to generate very fast and high temporal variations in power, we used four 802.11-based IoT devices: Avnet BCM4343W (BCM4343W), Cypress CYW43907 (Cypress CYW43907), RPiZW, and RPi3. The Avnet board includes an ARM Cortex-M4 processor and supports 802.11a/g/n. When using the power save mode, the minimum energy consumption of the board is around 10 mA, processing consumes around 40 mA, and packet transmission results in spikes up to 350 mA. The Cypress board includes an ARM Cortex-R4 processor and supports 802.11a/g/n. As the board includes other components including an Ethernet chip, the sleep power of the board is around 96 mA. The processing power is around 140 mA, and packet transmissions increase power consumption up to 400 mA. We have used Free-RTOS and WICED Studio (Cypress WICED Studio) for software development on the Avnet and Cypress devices. Furthermore, to generate higher variations in power, the software developed for these devices enables the power save mode (PS-Poll) mechanism of 802.11. At certain intervals, the radio wakes up and ping packets are transmitted as fast as possible. Since ping packets are small, this behavior results in fast and short spikes in power consumption.

For RPiZW, the base and processing power are about 130 mA and 180 mA, respectively, and 802.11 transmissions result in spikes as high as 300 mA. For the regular RPi3, the sleep and processing currents are about 280 mA and 330 mA, respectively, and 802.11 transmissions increase current consumption up to 500 mA. A program (written in C language) controls the transition between the three operations provided. Please note that for RPiZW and RPi3 we refer to the process inactivity time as sleep time.

Fig. 11 shows the energy consumption of CC2650 and BCM4343W. Although CC2650 shows a clear transition between the three states, BCM4343W presents spikes across the trace. It should be noted that in contrast with 802.15.4, 802.11 communication requires association with an access point. Therefore, when 802.11 power save mode is enabled, the device wakes up every 100 ms to receive the beacon packets generated by the access point.

Fig. 12 presents power measurement error when using EMPIOT's

shield and the breakout board. Error is computed as  $\frac{|E_x - E_a|}{E_a} \times 100$ , where  $E_a$  is the actual energy consumption measured by DMM, and  $E_x$  refers to the energy measured by either EMPIOT's shield or the breakout board. Each marker is the median of 10 experiments, where an experiment is 30 s long.<sup>3</sup> These results indicate that EMPIOT is in fact an accurate power measurement platform, even if an off-the-shelf breakout board is used instead of the shield. Comparing sub-figures (a) through (e) shows that energy measurement error is higher when measuring very small currents. Specifically, while the measurement error for 802.11-based boards is less than 2.5%, the error is less than 3.5% for SensorTag. Considering Workload 1 on SensorTag, we observed that more than 40% of variations in current are less than 1 mA. If the variation is caused by transition to the sleep mode, the hybrid power measurement model is used and the CC2650 device informs EMPIOT about its transition to and from the sleep mode by generating interrupts. However, since the current resolution of INA219 is 100  $\mu$ A, for example, EMPIOT cannot detect current variations between 3.25 mA and 3.27 mA. Therefore, as the power consumption of SensorTag is significantly lower than that of other boards, small errors in current measurement result in a more considerable effect on total measurement error. Nevertheless, these results confirm the effectiveness of the hybrid power measurement technique.

Fig. 12 also reveals the effect of calibration on accuracy. Since the measurement error of the breakout board is higher than that of EMPIOT's shield, calibration has a higher effect on accuracy. In addition, this is particularly important for 802.11-based boards because, as Fig. 10 shows, the error of the breakout board is higher and non-linear for currents higher than 300 mA.

Although our studies in Section 2.3.1 showed that EMPIOT supports around 1000 and 4000 samples per second for 12 and 9-bit resolutions, respectively, it should be noted that the actual sampling rate of the delta-sigma ADC is about 500 KHz, and the decimator averages the analog samples collected per sampling interval. For example, when 12-bit sampling is used for 802.11-based devices, the rapid power variations of these platforms are captured and averaged per millisecond, according to the sampling rate we reported in Section 2.3.1. Consequently, we observe that EMPIOT achieves high accuracy even in the presence of very fast temporal power variations.

Fig. 12 also shows that the measurement error of 9-bit resolution is slightly higher than that of 12-bit resolution. This is particularly obvious for the SensorTag measurements. As we mentioned earlier, compared with the 802.11 devices, SensorTag generates smaller variations in power, thereby a higher resolution is required to capture the changes. Compared with 12-bit, using 9-bit resolution enhances the sampling rate (cf. Fig. 3), however, the minimum detectable variation in current is increased to 780  $\mu$ A. For 802.11-based devices with a high data transmission rate, the higher sampling rate enables us to capture shorter variations in power. Therefore, the increase in sampling rate compensates the error introduced due to lower resolution.

<sup>3</sup> The 30-s experiment duration is due to the limitation of the DMM used in terms of the number of samples stored in its memory.




Fig. 11. The energy trace of SensorTag CC2650 (using 802.15.4) and Avnet BCM4343W (using 802.11) when transitioning between sleep, encryption and transmission. This figure shows the significantly higher variations of current consumption caused by 802.11 compared with 802.15.4.

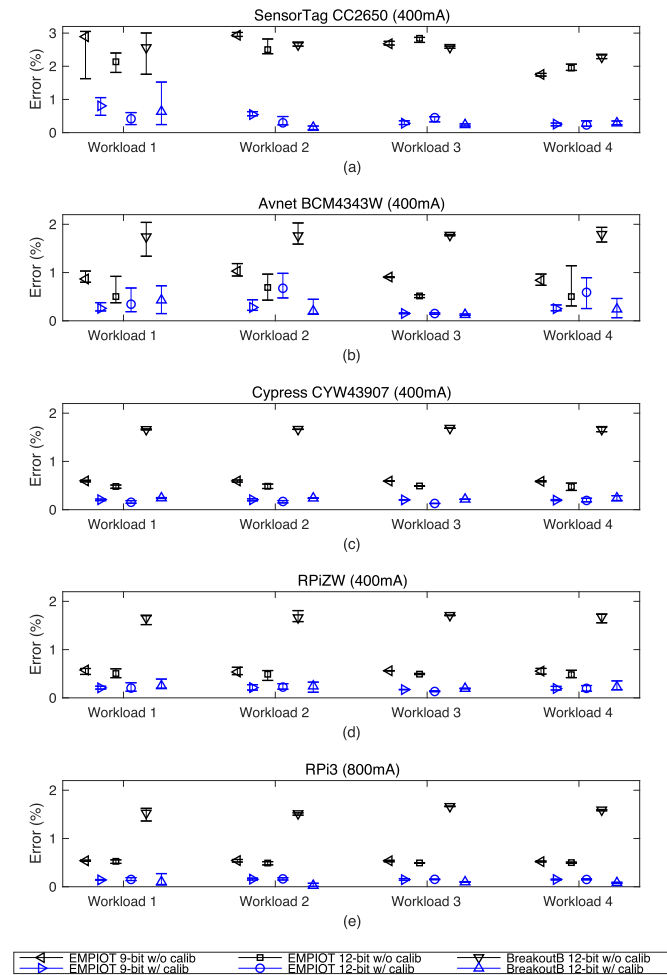


Fig. 12. The energy measurement error of EMPIOT versus the ground truth. Error bars show median, lower quartile and higher quartile. Each marker is the median of ten experiments each 30 s long.

4.1. Importance of voltage measurement

Most of the existing energy measurement platforms ignore the effect of voltage variations on energy measurement (e.g. Zhou and Xing, 2013; Trathnigg et al., 2008; Andersen and Hansen, 2009; Haratcherev et al., 2008). To verify the implication of this assumption on accuracy, we have used Workload 1 and computed energy when: (i) the average value of voltage measurements is used, and (ii) the voltage samples are used. Error is computed as the difference between these cases. Fig. 13 shows the results when a power supply and a battery are used as the sources of power for various types of IoT devices. These results indicate that neglecting voltage would result in up to a 0.45% increase in energy measurement error. The impact of voltage on energy mea-

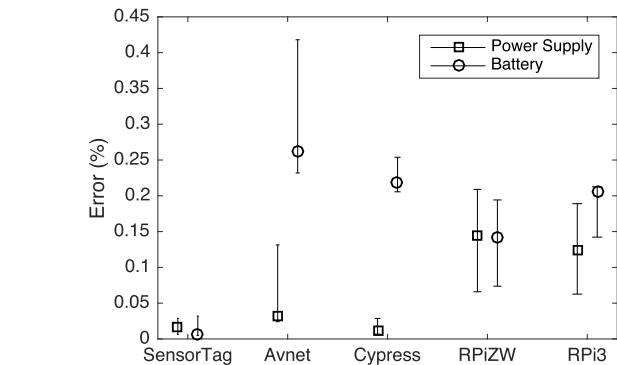


Fig. 13. The effect of ignoring voltage on the accuracy of energy measurement. Error bars show median, lower quartile and higher quartile.

surement depends on various factors including: the stability of power source, the number and intensity of sudden increases in current draw, and the electronic characteristics of the IoT device such as voltage stabilization. For example, when a battery is used, sudden variations of current result in a higher measurement error when voltage is ignored. These results also show that the Cypress and Avnet devices cause significant variations in input voltage, compared with RPi3 and RPiZW. Our studies show that these variations are caused by the operations of 802.11 transceiver, thereby highlighting the importance of including voltage for 802.11-based IoT devices. On the other hand, the energy consumption of SensorTag and its 802.15.4 transceiver does not cause any significant variation in voltage.

5. Related work

In this section we provide an overview of the existing energy measurement solutions in two domains: empirical measurement, and analytical modeling.

5.1. Empirical measurement

Based on their main shortcoming, we categorize the platforms proposed for empirical energy measurement as follows.

5.1.1. Costly and bulky

An oscilloscope is used in (Feeney and Nilsson, 2001) to analyze the energy consumption of 802.11 transceivers. Using a current probe is the approach employed in (Milenkovic et al., 2005) for power monitoring. The current probe relies on the magnetic field generated by current draw; therefore, it cannot be used to detect small currents or variations (milliamps level). The authors in (Kaup et al., 2014) added a shunt resistor to a USB cable and measured both current and bus voltage using USB1608-FSPlus (Measurement Computing Corporation.), a 16-bit ADC. The USB1608-FSPlus is controlled by a PC and the sampling rate is 1 Ksps. All of these approaches are expensive and their size prevents integration with the IoT devices of a testbed.

### 5.1.2. Complex circuit

Similar to ICs such as BQ2019 (BQ 2019), the SPOT (Jiang et al., 2007) platform uses voltage to frequency conversion and relies on the resources of the device under test to operate. In addition, the oscillator and the converter are sources of noise and error and may interfere with the device under test. More importantly, attaching SPOT to an IoT device is not plug-and-play. SPOT's measurement error is less than 15% and its resolution is 1  $\mu$ A; however, it assumes that the typical operating range of the IoT device is 5  $\mu$ A to 50 mA LEAP2 (Stathopoulos et al., 2008) is a FPGA-based platform that enables individual monitoring of various components such as processor and memory. iCount (Dutta et al., 2008) relies on the linear relationship between current and switching frequency of boost switching regulators; counting the switching cycles reflects the current drawn during an interval. A shortcoming of iCount is that it utilizes the resources of device under test and increases its energy consumption. Another limitation of this approach is that boost converters are not always available on IoT boards. For example, both CYW43907 (Cypress Semiconductor, 2982) and CC2650 (CC2650 SimpleLink Multistandard Wireless,) (used in Section 4) utilize internal voltage regulators.

Energy Bucket (Andersen and Hansen, 2009) counts the number of charges and discharges of a buffer capacitor. Although Energy Bucket can measure currents in the range 1  $\mu$ A to 100 mA, the major shortcoming is the dependency of sampling rate on capacitor value. For example, when the current drawn is very small, the inter-sampling interval would be long. In addition, the platform assumes that bus voltage is fixed, and the paper does not include accuracy analysis. Nemo (Zhou and Xing, 2013) uses a shunt resistor switch composed of a series of resistors. Each resistor is enabled or disabled based on current intensity. The voltage across the resistor switch is amplified using a differential op-amp and then digitized by a 12-bit ADC. Although this technique eliminates the need for a high-precision ADC or an adjustable amplifier, Nemo's processor must quickly react to changes in current and adjust the resistor value. Depending on the input current (which is supported in the range 1  $\mu$ A and 200 mA) the resolution of Nemo varies in the range 0.013  $\mu$ A–48  $\mu$ A. Unfortunately, Nemo does not measure bus voltage, and the evaluations are simple and do not include an IoT device with high power variations.

$\mu$ Monitor (Naderiparizi et al., 2016) proposes a power monitoring platform based on counting capacitor charging and discharge cycles. For the loads within the range of 1  $\mu$ W to 10 mW, the accuracy of  $\mu$ Monitor is almost within 10% of the results obtained from a 16-bit ADC that digitizes the voltage value over a shunt resistor. Unfortunately, the evaluations use static loads. Similar to Nemo (Zhou and Xing, 2013), Pötsch et al. (2014, 2017) propose the use of two shunt resistors (1  $\Omega$  and 100  $\Omega$ ) for measuring low and high currents up to 100 mA. Shunt resistor voltages are amplified and then sampled by a 16-bit ADC. A 32-bit microcontroller communicates with the ADC and collects the samples. The actual sampling rate and resolution of this platform have not been evaluated.

In addition to the limitations highlighted, a common shortcoming of the platforms reviewed in this section is their complex circuitry. Specifically, these platforms include various components such as ADC, op-amp, resistor series, high precision capacitors, and processor. Therefore, building these platforms is costly and time consuming. In contrast, EMPIOT is very easy to build, and the cost of a complete platform (shield plus base board) is around \$35 when a RPi3 is used as the base board, and \$15 when a RPiZ is used.

### 5.1.3. Limited range

PowerBench (Haratcherev et al., 2008) is capable of providing a 5 KHz sampling rate and 30  $\mu$ A resolution. PowerBench assumes that the maximum current consumption of the host device is 65 mA. This platform samples the amplified voltage (51x) across a shunt resistor using a 12-bit ADC. The authors in (Trathnigg et al., 2008) assume energy profiling requires current measurement only. The error of this platform

is less than 5% as long as current is higher than 20  $\mu$ A. However, the maximum supported current is 35 mA. iWEEP\_HW (Zhu and O'Connor, 2013) employs a multi-layer architecture to measure the power consumption of processor, transceiver, and sensors using separate ADC channels. This platform uses a PIC18 processor and a 10-bit ADC which measures voltage across shunt resistors. iWEEP\_HW can measure currents up to 40 mA with a maximum sampling rate of 150 KHz. PotatoScope (Hartung et al., 2016) is a microcontroller-based oscilloscope that focuses on reliable energy measurement in outdoor environments and in the presence of significant temperature variation. PotatoScope includes an ARM Cortex-M3 processor that is attached to a 12-bit ADC to sample current and voltage. The voltage across a 0.47  $\Omega$  shunt resistor is amplified by 200x. Since the ADC's reference voltage is 2.5 V, the maximum measurable current is 26.6 mA.

The main shortcoming of these platforms is their limited current measurement range, which is less than 100 mA. In this paper we showed that the new generation of IoT devices rely on high data rate technologies that increase temporal power consumption as high as 700 mA. Therefore, none of the platforms studied in this section can be used with these IoT devices. The performance evaluation results presented in Section 4 confirm the effectiveness of EMPIOT for measuring the energy of 802.11 IoT devices.

### 5.1.4. Low accuracy

Energino (Gomez et al., 2012) uses the 10-bit ADC of Arduino boards. The ADC supports an input voltage of 0–5 V, therefore, the LSB is 4.88 mV and the current measurement resolution is 25 mA. Energino uses a Hall-effect current sensor with sensitivity 185 mV/A to measure currents up to 5 A. The actual sampling rate is 1 KHz. NITOS (Keranidis et al., 2014) uses the ATmega2560 micro-controller with a 10-bit ADC. Since the Arduino's ADC cannot measure millivolt level variations in voltage, a voltage amplifier has been used. In order to maximize sampling rate, NITOS: (i) uses ADC free running mode, (ii) increases the ADC's prescaler clock from 125 KHz to 1 MHz, and (iii) uses interrupt service routine to get ADC values. Although these enhancements result in up to 63 KHz sampling rate, the accuracy of ADC is reduced by 11% due to the higher clock used. Furthermore, the current measurement resolution of the platform is 25 mA.

## 5.2. Analytical estimation

In addition to empirical measurement, analytical modeling can be employed to estimate power. In (Wang et al., 2006) the authors model the energy consumption of wireless communications for single and multi-hop networks. However, since the authors rely on a simple CSMA protocol, the model would not be valid in the presence of interference and collision. Furthermore, the energy consumption of the processor and other components such as sensors, has not been taken into account. Similar to this work, unfortunately, most of the existing models only consider the energy consumption of wireless communication (Du et al., 2010; Li and Mohapatra, 2007). In particular, due to its prevalence, 802.15.4 is the primary technology modeled.

In addition to communication aspects, models have been proposed for modeling the energy consumption of processors (Bazzaz et al., 2013; Konstantakos et al., 2008). After accurate energy profiling of processor and peripherals, the energy cost of each instruction is used to calculate the energy of applications.

The authors in (Martinez et al., 2015) propose a framework for modeling the power consumption of various components on an IoT device. To model the power consumption of networking, they take into account the effect of transmission power, re-transmission, and spreading factor, as used by wireless technologies such as LoRa. Processing power has been modeled based on the average power consumed per arithmetic instruction and the time complexity of the algorithm used. To apply the models to a system, however, it is essential to first run experiments and fit the models into the empirical results collected.

Although analytical approaches are very useful during the system design phase, empirical energy measurement is inevitable. In addition to the points we mentioned in Section 1, the following points further justify the importance of empirical evaluation. First, accurate and thorough energy profiling is necessary to formulate analytical energy estimation models. For example, the energy consumption of the board might be significantly affected by environmental interference or the start-up power of components. Second, empirical long term energy analysis is necessary to identify the hidden sources of energy depletion (Brusey et al., 2016). For example, some studies show the significant effect of link breakage and server outage on energy usage. Third, various types of batteries are affected differently by load and environmental factors.

## 6. Conclusion

The EMPIOT platform presented in this paper enables low-cost, programmable, and accurate energy measurement of a wide range of IoT devices. After explaining the hardware and software components of this platform, we evaluated the effect of various parameters on performance. Our studies show that: (i) the BCM driver results in a higher sampling rate and lower energy consumption, (ii) lowering the shield's operational voltage slightly reduces the sampling rate, (iii) voltage variations do not affect accuracy, and (iv) the effect of two-buffer or circular buffering mechanisms on power consumption of the base board depends on the number of processor cores. In terms of calibration, we developed a programmable load that enables us to generate a wide range of load values and calibrate the platform without requiring precise timing control.

For IoT devices whose their minimum energy consumption is more than 100  $\mu$ A, EMPIOT can be simply connected to the input power supply of the device. If the minimum energy consumption is less than 100  $\mu$ A, EMPIOT uses a hybrid energy estimation model to take into account the time spent in sleep modes. To this end, extra connections are required to inform EMPIOT before and after every sleep duration. We evaluated the performance of EMPIOT using five different IoT devices and four types of workloads. Our results confirm that energy measurement error is less than 3.5% for IoT devices that utilize either 802.15.4 or 802.11 wireless standards. Beyond the scope of this paper, our study on the I2C performance and energy efficiency of Linux provides insights into designing Linux-based IoT systems.

Some of the future work avenues are as follows: Although the energy consumption of simple wireless technologies such as 802.15.4 and LoRa have been thoroughly studied and modeled, in addition to 802.11b/n standards, newer technologies, such as 802.11ac, NB-IoT, and mMTC, are being used for IoT applications. The EMPIOT platform provides a low-cost and scalable solution to deploy testbeds and profile the energy consumption of these complex wireless technologies. Another area of future work is to port EMPIOT to non-Linux-based systems and measure its sampling rate and overhead when the operating system of the base board is a RTOS. Not only for a wide variety of IoT devices, the measurement range of EMPIOT makes it a suitable energy measurement platform for other applications, such as measuring and evaluating the effect of energy efficiency techniques proposed for cloud computing platforms (Aldawsari et al., 2015; Baker et al., 2017).

## Acknowledgment

This work has been partially supported by a research grant from Cypress Semiconductor Corporation (Grant No. CYP001).

## References

34465A. Keysight Technologies, 34465A (612 digit). Available from: <https://literature.cdn.keysight.com/litweb/pdf/5991-1983EN.pdf>.  
 Aldawsari, B., Baker, T., England, D., 2015. Trusted energy efficient cloud-based services brokerage platform. *Int. J. Intell. Comput. Res.* 6, 630–639.

ADG1611. Analog Devices, Quad SPST Switches. Available from: [http://www.analog.com/media/en/technical-documentation/data-sheets/ADG1611\\_1612\\_1613.pdf](http://www.analog.com/media/en/technical-documentation/data-sheets/ADG1611_1612_1613.pdf).  
 AD5200. Analog Devices, 256-Position and 33-Position Digital Potentiometers. Available from: [http://www.analog.com/media/en/technical-documentation/data-sheets/AD5200\\_5201.pdf](http://www.analog.com/media/en/technical-documentation/data-sheets/AD5200_5201.pdf).  
 Andersen, J., Hansen, M.T., 2009. Energy Bucket: a tool for power profiling and debugging of sensor nodes. In: *Proceedings of Third International Conference on Sensor Technologies and Applications (SENSORCOMM'09)*. IEEE, pp. 132–138.  
 Baker, T., Asim, M., Tawfik, H., Aldawsari, B., Buyya, R., 2017. An energy-aware service composition algorithm for multiple cloud-based iot applications. *J. Netw. Comput. Appl.* 89, 96–108.  
 Bazzaz, M., Salehi, M., Ejlali, A., 2013. An accurate instruction-level energy estimation model and tool for embedded systems. *IEEE Trans. Instru. Measure.* 62 (7), 1927–1934.  
 BCM4343W.BCM4343W: 802.11b/g/n WLAN, Bluetooth and BLE SoC Module. Available from: [https://products.avnet.com/opsdata/d120001/medias/docus/138/AES-BCM4343W-M1-G\\_data\\_sheet\\_v2\\_3.pdf](https://products.avnet.com/opsdata/d120001/medias/docus/138/AES-BCM4343W-M1-G_data_sheet_v2_3.pdf).  
 BQ2019.BQ2019: Advanced Battery Monitor. Available from: <https://www.ti.com/lit/ds/slus465e/slus465e.pdf>.  
 Brusey, J., Kemp, J., Gaura, E., Wilkins, R., Allen, M., 2016. Energy profiling in practical sensor networks: identifying hidden consumers. *IEEE Sensor. J.* 16 (15), 6072–6080.  
 C library for Broadcom BCM 2835. Available from: <http://www.airspayce.com/mikem/bcm2835/>.  
 CC2650. Texas Instruments, CC2650 SimpleLink Multistandard Wireless MCU. Available from: <https://www.ti.com/lit/ds/swrs158b/swrs158b.pdf>.  
 CYW943907AEVALIF. Cypress Semiconductor, CYW943907AEVALIF Evaluation Kit. Available from: <http://www.cypress.com/documentation/development-kitsboards/cyw943907aevalif-evaluation-kit>.  
 CYW43907. Cypress Semiconductor, CYW43907: IEEE 802.11 a/b/g/n SoC with an Embedded Applications Processor. Available from: <http://www.cypress.com/file/298236/download>.  
 DMM7510. 7.1/2 Digit Graphical Sampling Multimeter. Available form: <https://www.tek.com/tektronix-and-keithley-digital-multimeter/dmm7510>.  
 Dezfouli, B., Radi, M., Whitehouse, K., Razak, S.A., Tan, H.-P., 2014. Cama: efficient modeling of the capture effect for low-power wireless networks. *ACM Trans. Sens. Netw.* 11 (1), 20.  
 Dezfouli, B., Radi, M., Razak, S.A., Hwee-Pink, T., Bakar, K.A., 2015. Modeling low-power wireless communications. *J. Netw. Comput. Appl.* 51, 102–126.  
 Du, W., Mieveville, F., Navarro, D., 2010. Modeling energy consumption of wireless sensor networks by systemc. In: *Fifth International Conference on Systems and Networks Communications (ICSNC)*. IEEE, pp. 94–98.  
 Dutta, P., Feldmeier, M., Paradiso, J., Culler, D., 2008. Energy metering for free: augmenting switching regulators for real-time monitoring. In: *Proceedings of the 7th International Conference on Information Processing in Sensor Networks (IPSN'08)*, pp. 283–294.  
 Eriksson, J., Österlind, F., Finne, N., Dunkels, A., Tsiftes, N., Voigt, T., 2009. Accurate network-scale power profiling for sensor network simulators. In: *Proceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN'09)*, vol. 9. Springer, pp. 312–326.  
 Feeney, L.M., Nilsson, M., 2001. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In: *Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'01)*, pp. 1548–1557.  
 Gartner. IoT Units Installed Base by Category (Millions of Units). Available from: <https://www.gartner.com/newsroom/id/3598917>.  
 Gomez, K., Riggio, R., Rasheed, T., Miorandi, D., Granelli, F., 2012. Energinio: a hardware and software solution for energy consumption monitoring. In: *Proceedings of the International Workshop on Wireless Network Measurements (WiOpt'12)*, pp. 311–317.  
 Haratcherev, I., Halkes, G., Parker, T., 2008. PowerBench: a scalable testbed infrastructure for benchmarking power consumption. In: *Proceedings of the International Workshop on Sensor Network Engineering (IWSNE'08)*, pp. 37–44.  
 Hartung, R., Kulau, U., Wolf, L., 2016. Distributed energy measurement in WSNs for outdoor applications. In: *Proceedings of the 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON'16)*.  
 INA219 Breakout Board. Adafruit Industries. Available from: <https://learn.adafruit.com/adafruit-ina219-current-sensor-breakout/downloads>.  
 INA219. Texas Instruments, INA219 Zero-Drift, Bidirectional Current/Power Monitor With I2C Interface. Available from: <http://www.ti.com/lit/ds/symmlink/ina219.pdf>.  
 Jiang, X., Dutta, P., Culler, D., Stoica, I., 2007. Micro power meter for energy monitoring of wireless sensor networks at scale. In: *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN'07)*, p. 186.  
 Kaup, F., Gottschling, P., Hausheer, D., 2014. PowerPi: measuring and modeling the power consumption of the Raspberry Pi. In: *Proceedings of the 39th Annual IEEE Conference on Local Computer Networks (LCN'14)*, pp. 236–243.  
 Keranidis, S., Kazdaridis, G., Passas, V., Korakis, T., Koutsopoulos, I., Tassioulas, L., 2014. NITOS energy monitoring framework: real time power monitoring in experimental wireless network deployments. *SIGMOBILE Mob. Comput. Commun. Rev.* 18 (1), 64–74.  
 Konstantakos, V., Chatzigeorgiou, A., Nikolaidis, S., Laopoulos, T., 2008. Energy consumption estimation in embedded systems. *IEEE Trans. Instru. Measure.* 57 (4), 797–804.  
 Li, J., Mohapatra, P., 2007. Analytical modeling and mitigation techniques for the energy hole problem in sensor networks. *Pervasive Mob. Comput.* 3 (3), 233–254.

- Li, Q., Martins, M., Gnawali, O., Fonseca, R., 2013. On the effectiveness of energy metering on every node. In: Proceedings of the IEEE International Conference on Distributed Computing in Sensor Systems (DCoSS'13), pp. 231–240.
- Lim, R., Ferrari, F., Zimmerling, M., Walser, C., Sommer, P., Beutel, J., 2013. FlockLab: a testbed for distributed, synchronized tracing and profiling of wireless embedded systems. In: Proceedings of the 12th International Conference on Information Processing in Sensor Networks (IPSN'13), p. 153.
- Linux I2C Driver. Available from: <https://www.kernel.org/doc/Documentation/i2c/dev-interface>.
- Manweiler, J., Roy Choudhury, R., 2011. Avoiding the rush hours: wifi energy management via traffic isolation. In: Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services. ACM, pp. 253–266.
- Martinez, B., Monton, M., Vilajosana, I., Prades, J.D., 2015. The power of models: modeling power consumption for iot devices. IEEE Sensor. J. 15 (10), 5777–5789.
- Measurement Computing Corporation. USB-1608FS-Plus data acquisition device. Available from: <http://www.mccdaq.com/pdfs/manuals/USB-1608FS-Plus.pdf>.
- Milenkovic, A., Milenkovic, M., Jovanov, E., Hite, D., Raskovic, D., 2005. An environment for runtime power monitoring of wireless sensor network platforms. In: Proceedings of the Thirty-seventh Southeastern Symposium on System Theory (SSST'05), pp. 406–410.
- Monsoon Inc. Monsoon Power Monitor. Available from: <https://www.msoon.com/LabEquipment/PowerMonitor/>.
- Naderiparizi, S., Parks, A.N., Parizi, F.S., Smith, J.R., May 2016.  $\mu$ Monitor: in-situ energy monitoring with microwatt power consumption. In: Proceedings of the IEEE International Conference on RFID (RFID'16). IEEE, pp. 1–8.
- Pötsch, A., Berger, A., Leitner, C., Springer, A., 2014. A power measurement system for accurate energy profiling of embedded wireless systems. In: Proceedings of the 19th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'14), pp. 1–4.
- Pötsch, A., Berger, A., Springer, A., 2017. Efficient analysis of power consumption behaviour of embedded wireless iot systems. In: Proceedings of the Instrumentation and Measurement Technology Conference (I2MTC), pp. 1–6.
- Serrano, P., Garcia-Saavedra, A., Bianchi, G., Banchs, A., Azcorra, A., 2015. Per-frame energy consumption in 802.11 devices and its implication on modeling and design. IEEE/ACM Trans. Netw. 23 (4), 1243–1256.
- Stathopoulos, T., McIntire, D., Kaiser, W.J., 2008. The energy endoscope: real-time detailed energy accounting for wireless sensor nodes. In: Proceedings of International Conference on Information Processing in Sensor Networks (IPSN'08), pp. 383–394.
- strace: Linux syscall tracer. Available from: <https://strace.io>.
- TI-RTOS: Real-Time Operating System (RTOS). Available from: <http://www.ti.com/tool/TI-RTOS>.
- Trathnigg, T., Jürgen, M., Weiss, R., 2008. A low-cost energy measurement setup and improving the accuracy of energy simulators for wireless sensor networks. In: Proceedings of the Workshop on Real-world Wireless Sensor Networks, pp. 31–35.
- U8000. Keysight Technologies, U8000 Series, Single Output DC Power Supplies. Available from: <https://literature.cdn.keysight.com/litweb/pdf/5989-7182EN.pdf?id=1460471>.
- Wang, Q., Hempstead, M., Yang, W., 2006. A realistic power consumption model for wireless sensor network devices. In: Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks (SECON), vol. 1. IEEE, pp. 286–295.
- WICED Studio. Cypress, Wireless Internet Connectivity for Embedded Devices. Available from: <http://www.cypress.com/products/wiced-software>.
- Zhou, R., Xing, G., 2013. Nemo: a high-fidelity noninvasive power meter system for wireless sensor networks. In: Proceedings of the ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN'13), pp. 141–152.
- Zhu, N., O'Connor, I., 2013. Energy measurements and evaluations on high data rate and ultra low power wsn node. In: 10th IEEE International Conference on Networking, Sensing and Control (ICNSC), pp. 232–236.

**Behnam Dezfouli** is an Assistant Professor at the Department of Computer Engineering, Santa Clara University (SCU), US. Before joining SCU, he was a Postdoctoral Research Scientist and a Visiting Assistant Professor at the University of Iowa, US, during 2015–2016. Behnam Dezfouli served as a Postdoctoral Research Fellow and a Research Fellow for University Technology Malaysia, Malaysia, and Institute for Infocomm Research, Singapore, during 2014 and 2012, respectively. He received his PhD in Computer Science from University Technology Malaysia in 2014. His research interests are wireless and software-defined networks, security, and cyber-physical systems. He is currently an associated editor of IEEE Communications Letters journal. He also has served as the TPC member of several conferences.