# eProfiler: High-Precision Power Monitoring System for IoT **Devices Featuring Extreme Dynamic Range of Operation**

**Giannis Kazdaridis** iokazdarid@uth.gr ECE Department, University of Thessaly, Greece

Ioannis Zografopoulos zografop@uth.gr ECE Department, University of Thessaly, Greece

Nikos Sidiropoulos nsidirop@uth.gr ECE Department, University of Thessaly, Greece

**Polychronis Symeonidis** posymeon@uth.gr ECE Department, University of Thessaly, Greece

**Thanasis Korakis** korakis@uth.gr Thessaly, Greece

# ABSTRACT

Modern Internet-of-Things (IoT) devices and sensor systems exhibit extreme dynamic current consumption profile, since latest microprocessors and electronics support ultra-low currents in the sleep phase, of only a few nA, while they expend several mA in the active state. Existing power meters are incapable of measuring their expenditure in order to aid the development of energy-efficient schemes. In this work we introduce the eProfiler, a novel in-situ system for measuring the power expenditure of sensing devices that illustrate a wide current range. Our meter features an autoranging shunt-resistor switch that supports ultra-fast alternations of 6ns, formed with the aid of high-speed comparators, that can successfully adapt to any sharp transition. Moreover, a high-speed Analog-to-Digital Converter (ADC) with multiple inputs is employed to simultaneously monitor the shunt-resistors, delivering a speed of 150 kSamples/s at 16 bit resolution. The proposed meter can be leveraged for the real-time, as well as the long-term monitoring of IoT devices with its cost being less than 90 euros. Furthermore, the proposed system can actuate or trace state alternations on I/O pins with a maximum delay of 2  $\mu$ s, in order to provide correlation capabilities with the obtained power measurements. The eProfiler features a wide dynamic range of 1.000.000:1 while also being able to monitor currents of a few nA. The average obtained error of the meter is 0.45 %, with a maximum error of 1.6 %.

#### CCS CONCEPTS

• Hardware → Wireless devices; • Computer systems organization  $\rightarrow$  Sensor networks; Embedded hardware.

# **KEYWORDS**

Power Consumption Monitoring, Energy-Efficiency, Shunt Resistor Switch, Low-Power Design, Sensor Networks, IoT

ENSsys '20, November 16-19, 2020, Virtual Event, Japan

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8129-1/20/11...\$15.00 https://doi.org/10.1145/3417308.3430265

ECE Department, University of

#### **ACM Reference Format:**

Giannis Kazdaridis, Ioannis Zografopoulos, Nikos Sidiropoulos, Polychronis Symeonidis, and Thanasis Korakis. 2020. eProfiler: High-Precision Power Monitoring System for IoT Devices Featuring Extreme Dynamic Range of Operation. In The 8th International Workshop on Energy Harvesting and Energy-Neutral Sensing Systems (ENSsys '20), November 16-19, 2020, Virtual Event, Japan. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/ 3417308.3430265

#### 1 INTRODUCTION

Energy-efficiency in the domain of Internet-of-Things (IoT) dominates the interest of the research community since real-world applications mainly employ battery operated sensors. Therefore, the life expectancy of a sensor network is solely dependent on the sensing node's power profile and the capacity of the battery in use. Evidently, further improvements in the sensors' power profile will increase the life duration of the network, which requires realistic feedback from in-situ power meters. To this end, several works [4, 5, 10-12, 21-24, 26, 28, 32, 37, 41, 47, 52], offer real-time monitoring to assist in the development of energy-efficient algorithms.

Recent IoT systems, integrate state-of-the-art elements such as RF radios, power regulators, sensing modules and micro-controllers, that feature improved power characteristics, as well as various operating and sleep modes. As a result, modern devices feature extremely low currents when switching to their sleep state conserving as much energy as possible. For instance, the Waspmote [49] and the eZ430-RF2500 [6] motes consume 860 nA and 690 nA respectively in their sleep state, while they draw several mA when active. Similarly, the ICARUS [15] mote exhibits a wide power profile dissipating only 22 nA in its quiescent state. Another prime example with significant power consumption deviations between its active and sleep state is the double-dip energy-harvesting system [27]. Although double-dip's consumption when active is typically in the range of few mA, in its quiescent state power consumption diminishes to only 700 nA. Identical behavior is observed by other energy-harvesting devices such as the [3, 51]. Apart from the aforementioned IoT systems as whole, their power supply sub-circuits are equally intriguing featuring wide dynamic currents that may vary from a few nA to a several mA. All the aforementioned scenarios suggest that the power profile of modern sensor nodes exceeds

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

The research leading to these results has received funding by the European Horizon 2020 Programme for research, technological development and demonstration under Grant Agreement Number 857201 (H2020 5G-VICTORI).

at least three orders of magnitude. Therefore, sophisticated tools with wide dynamic range spanning the entire spectrum of possible current draws are required to capture and characterize the power profile of IoT devices.

Commonly, power meters utilize a precise, low-impedance resistor, the so-called shunt resistor, which is placed in series, between the power supply and the Device Under Test (DUT). The voltage developed across the shunt is proportional to the current draw of the DUT according to Ohm's law. The selected shunt plays significant role in the accuracy of the entire system, thus, it must be carefully chosen. The shunt resistor role is twofold, (i) it should not affect the operation of the DUT, and (ii) provide sufficient level of detail to the next stage. Typically, a maximum voltage drop of 100 mV in a 3.3 V rail is tolerated by the DUT, therefore, the shunt must be calculated according to this limitation. For instance, for a maximum current of 100 mA, the selected resistor may not be greater than 1  $\Omega$ . This signal (1 mV-100 mV) can be processed by a pre-amplification circuit that magnifies the observed voltage drop to a full-scale analog signal to ensure accurate sampling by the ADC Integrated Circuit (IC). However, when measuring currents ranging from 1  $\mu$ A or 1 nA the selected resistor (1  $\Omega$ ) will introduce a voltage drop as low as 1  $\mu$ V or 1 nA respectively, which is infeasible to prevail over noise and parasitic thermo-voltages [35]. Therefore, the gain must be generated from the shunt resistor itself. In essence, it is necessary to adjust the shunt resistor during the measurement procedure, without interrupting the operation of the DUT, while constantly measuring even in the transition phases.

In this paper we introduce eProfiler, a high-precision power consumption meter that integrates a resistor switch able to adapt among the ranges of mA,  $\mu$ A and nA. The device supports 150 kSamples/s with 16 bit resolution. The key contributions are outlined:

- we present an innovative dynamic shunt switch able to adapt to different consumption ranges within 6 ns
- eProfiler is the first in-situ power meter able to measure currents of a few nA
- eProfiler integrates GPIO tracing and actuation to allow the correlation of the obtained data with specific DUT events
- we evaluate the performance of our meter in terms of error accuracy noting a maximum observed error of 1.6 %

The remainder of the paper is organized as follows. Section 2 reviews the related work. System architecture and implementation are described in section 3, while the system's evaluation is presented in section 4. Finally, section 5 concludes the paper.

#### 2 RELATED WORK

In this section we distinguish and present the state-of-the-art power meters. The SPOT [11] is an in-situ power meter that provides a dynamic range of 45.000:1, with an average error of 3 %. It uses a differential amplifier to boost the voltage drop across the shunt resistor and a Voltage-to-Frequency Converter (VFC) for the digitization process. The resulted frequency signal is accumulated and stored into two on-board counter ICs. Notably, this meter mainly focuses on long-term energy consumption monitoring, rather than acquiring detailed power measurements as we intend to do in this work. Moreover, the VFC IC cannot offer the desired resolution for measuring currents lower than 1  $\mu$ A. Two other power meters are the EMPIOT [4] and the ECO [37]. They both measure the voltage

drop across a shunt resistor to determine the power draw employing the INA219B and the INA226 current sense amplifiers respectively. These amplifiers integrate 12 bit and 16 bit ADC units, respectively. Apparently, the selected amplifiers feature high input offset voltage which results in significant output errors. However, the biggest pitfall of the mentioned ICs is their high input bias current which affects the measurement accuracy, discussed in detail in Section 3.3. Consequently, these amplifiers are not suitable for measuring currents in the range of  $\mu$ A and lower, but only in the mA range.

The authors in [26] present the FlockLab testbed, an observer board employed to remotely configure, measure power draw, and trace GPIO events of several attached sensor devices. For acquiring power consumption measurements, the observer features the MAX9923 current sense amplifier and a 24 bit ADC, while the Gumstix embedded device is used as the host computer. The MAX9923 is an excellent choice able to measure very low currents, since it features only 1 pA input bias current. However, the authors utilize only one shunt resistor, thus it is not feasible to accurately monitor currents ranging from a few nA to a few mA. Even with the aid of the selected 24 bit ADC, it is not possible to attain such dynamic range. The main drawback of FlockLab is the fact that it employs a Linux board, i.e., the Gumstix device, to communicate with the ADC unit and obtain the measurements. Linux being a process scheduling OS, and not a real-time system, cannot guarantee fixed intervals between the measurements, which is a significant requirement when designing power meters. In our work we also use a Linux-based device since this is the only way to support long-term monitoring capabilities along with an immediate user-interface. However, we overcome the aforementioned issue by employing a Programmable Real time Unit (PRU) integrated in the embedded system we selected to use which guarantees precise timing during the measurement acquisition.

Another distinguished work is the Nemo [52] in-situ power meter, which employs a dynamic shunt resistor switch of 4 different resistors, enabling power measurements in the range of 0.8  $\mu$ A to 202 mA (250.000:1), maintaining the average error of 1.34 %. The Nemo is the first in-situ meter to adopt a dynamic shunt switch, which uses MOSFET elements to enable or disable each shunt resistor. In addition, a single Operational Amplifier (OpAmp) is used to amplify the voltage drop, applied to the entire resistor array. This is not an appropriate strategy since the impedance presented by each MOSFET affects the overall impedance of the array and leads to imprecise results. In our implementation we use a similar dynamic shunt switch formed by 3 resistors but we employed 3 different amplifiers, each applied directly to the corresponding resistor, without any other external elements impacting on the measurement accuracy. Moreover, the Nemo employs a voltage comparator to generate an interrupt to the host MCU upon a sudden current increase or decrease, which in turn adjusts the resistance of the shunt resistor switch accordingly. Despite using the modern TI MSP430F2618 MCU, the obtained delay when switching between different ranges is roughly 7  $\mu$ s, which is significant when measuring transient phenomena. In fact, it is common that upon these transitions (i.e. wake-up from sleep state) sensor nodes may present extreme spikes in their consumption due to the fact that some electronics are initially powered (such as power capacitors, MCU, RF Radio, etc.). Notably, the Nemo discards any measurements polled

eProfiler



Figure 1: Architecture Diagram of the eProfiler

within this period, since the validity of the data obtained during this period is questionable. In our implementation we tackle this challenge using high-speed comparators directly controlling the switches which bridge the shunt resistors, achieving 6 ns speed.

## **3 SYSTEM IMPLEMENTATION**

In this section we highlight the challenges for designing a novel power meter, while we describe the proposed architecture and we detail the characteristics of our system.

# 3.1 Design Challenges

**Wide Current Range:** Any power monitoring system aiming to accurately measure the power consumption profiles of IoT sensing devices should feature a dynamic range spanning the entire spectrum of possible current draws. Hence, it must support readings from several mA all the way down to a few nA, which results in a range of at least 1.000.000:1 (considering 100 nA to 100 mA).

**High-speed & high-resolution:** The power meter should capture even the shortest in duration events to effectively record the power profile of a sensor system. Considering that the latest WiFienabled IoT nodes implement the IEEE 802.11n, that supports  $T_X$  rates up to 72 Mbps, and that a typical short packet has a size of 300 Bytes, the propagation of that frame would require roughly 33  $\mu$ s. Therefore, a sampling rate of at least 100 kHz is required.

**Long-term monitoring:** The power monitor must be able to capture power data that last for several hours or even days.

**Non-invasive operation:** The power consumption monitoring procedure should not interfere with the normal operation or affect the actual energy consumption of the DUT.

**GPIO tracing:** We expect that a high-fidelity power meter must synchronize with the DUT to allow the correlation of the obtained measurements. An efficient way to implement this, is to trace GPIO events artificially realized in the DUT. The readings of these events must be as prompt as possible.

# 3.2 System Architecture

The architecture we followed to develop the proposed meter is illustrated in Fig. 1, while the developed Printed Circuit Board

ENSsys '20, November 16-19, 2020, Virtual Event, Japan



Figure 2: The eProfiler Power Meter

(PCB) is shown in Fig. 2. Notably, this board attaches on top of the selected embedded PC. The device consists of 6 main blocks, the embedded device, the ADC unit, the current sense amplifiers, the shunt resistors switch, the comparators block and the power supply circuits. Our meter is based on the high-side current [43] sensing topology, which exhibits significant advantages over the low-side configuration featuring undesirable ground path-induced disturbances. We develop a dynamic shunt resistor switch that supports 3 different ranges (i.e., mA,  $\mu$ A, nA) in order to adapt to a wide range of currents. High-speed comparators are used to control the shunt resistor array, without the intervention of a logic unit that usually induces significant delays as denoted in [28, 52]. An important distinction of our architecture is the employment of a dedicated current amplifier for each shunt resistor and the fact that the selected ADC supports multiple inputs enabling the continuous polling of the amplifier channels. Apart from the amplifiers, the eProfiler is constantly monitoring the outputs of the comparators (labeled as SW2, SW3) on each measurement cycle. Upon completion of each experiment, a post-process script combines all the data to a final measurements file, accounting for the state of each comparator indicating which resistor/range was active at every discrete sample. The other advantage of our meter is the PRU unit integrated by the embedded device we selected which enables realtime communication with the employed ADC, supporting fixed intervals between measurement cycles. Moreover, we use the same PRU unit to trace I/O events of the DUT with ultra-low latency, to support the correlation of the obtained power measurements with specific activities of the DUT. The eProfiler also monitors the derived voltage rail that powers the DUT that can be used to infer accurate energy measurements, especially when significant power rail fluctuations occur as noted in [52].

#### 3.3 Selected Components & Characteristics

In this section we present the components we opt for detailing their characteristics and the reasons for their selection.

**BeagleBone Embedded PC:** The embedded PC we opted for, is the Linux-based BeagleBone Black Rev. C [2], or the BeagleBone Wireless that supports wireless connectivity. The BeagleBone is a low-cost, embedded platform characterized by sufficient processing power capabilities (1GHz CPU with 512MB RAM), low-power consumption and several communication interfaces. The platform is responsible for controlling the peripheral units, and implements the software framework for the energy monitoring system. Furthermore, the device features an external microSD, used to locally cache the acquired measurements, prior to the offloading process. Apparently, the only limitation in the measurement duration is the ENSsys '20, November 16-19, 2020, Virtual Event, Japan

	Туре	G. Err.	$V_{OS}$	$I_B$	CMRR	PSRR	Used by
INA139	CurS	0.10%	200 µV	10 µA	115 dB	-	-
INA219	CurS	-	$40 \ \mu V$	20 µA	120 dB	100 dB	EMPIoT [4]
INA225	CurS	0.05%	75 μV	72 µA	105 dB	140 dB	-
INA226	CurS	0.02%	2.5 μV	10 µA	140 dB	112 dB	ECO [37]
INA282	CurS	0.01%	20 µV	25 µA	140 dB	110 dB	-
INA301	CurS	0.10%	25 µV	120 µA	110 dB	140 dB	-
AD8219	CurS	0.10%	200 µV	130 µA	110 dB	110 dB	-
LTC6102	CurS	-	3.5 μV	60 pA	-	150 dB	in [32]
MAX9923	CurS	0.12%	0.1 μV	1 pA	140 dB	99 dB	FlockLab[26]
OPA2333	OpAm	-	$2 \mu V$	150 pA	130 dB	120 dB	Nemo [52]
MAX4239	OpAm	-	$0.1 \ \mu V$	1 pA	140 dB	140 dB	uCurrent[47]
				-			

**Table 1: Compelling Amplifiers and their Specifications** 

size of the card employed. We avoid using a micro-controller device, which is common in the other power meters, but opt for a Linuxbased system in order to support fast communication with the user and the ability to transfer large files. Notably, the BeagleBone is the only embedded system that incorporates a PRU unit.

**ADC unit:** The ADC we employ is the Texas Instruments (TI) ADS8332 [45] supporting the SAR architecture [48], which features 16 bits of resolution, high SNR of 91 dB and high Power Supply Rejection Ratio (PSRR) of 74 dB. The ADS8332 performs conversions at 500 kSamples/s (kSps), while it integrates 8 input pins allowing for parallel sampling with the aid of its internal multiplexer. Therefore, the supported sampling rate is split among the configured channels. Apparently, when using all the supported current ranges (mA,  $\mu$ A, nA) 3 inputs are employed, thus the maximum attained speed is 166 kSps. On the other hand, when the user decides to also monitor the voltage rail of the DUT, 4 inputs are engaged resulting in a maximum speed of 125 kSps. Of course, users can also select to monitor only 1 or 2 channels (different ranges) in order to attain higher sampling speeds (500 and 250 kSps respectively).

**Current Sense Amplifier:** The current amplifier is of vital importance since it amplifies the obtained voltage drop across the shunt resistors and feeds it to the utilized ADC. Table 1 summarizes the compelling and most used current amplifiers.

We select the MAX9923 [1] since it presents excellent characteristics in all parameters. The MAX9923 presents ultra-low Input Offset Voltage ( $V_{OS}$ ).  $V_{OS}$  defines the output voltage deviation from the ideal value. Notably, the rest of the other amplifiers feature a few or in some cases several  $\mu$ A, thus they may induce significant variations from the ideal value, while the MAX9923 features just a few nA. PSRR and Common Mode Rejection Ratio (CMRR) also affect the performance of the amplifier. They define the ability of the device to reject common-mode and power supply noise signals respectively. High CMRR and PSRR is required when a differential signal is amplified in the presence of strong or even some electromagnetic interference. Of course, the selected IC also features low gain error, which is also critical for our application.

Another crucial parameter, overlooked by other works, is the Input Bias Current ( $I_B$ ), which is drained by the amplifier itself to bias the internal circuitry of the IC. All listed amplifiers present good performance when measuring currents in the range of mA, but when it comes to amplifying signals in the range of nA, they drain comparable current to the monitored signals, which results in false readings. For instance, consider measuring a 1  $\mu$ A load current using the INA139 amplifier which features 10  $\mu$ A  $I_B$ . The INA139 amplifies the obtained signal across the resistor, but also drains 10  $\mu$ A, which get amplified and summed up with the load current, resulting in a totally inaccurate reading. Consequently, amplifiers that feature  $I_B$  of a few  $\mu$ A should only be used to measure currents in the range of mA and above, and should be avoided for lower current ranges. Even when considering the OPA2333, that features 150 pA  $I_B$ , the obtained error for measuring a 1 nA current will be 15 %. The MAX9923 and the MAX4239 are the only ICs that feature such low  $I_B$  values, specifically 1 pA. However, the MAX4239 being an OpAmp implies that its gain is controlled by external resistors, which can greatly drift due to temperature fluctuations adversely affecting the accuracy of the amplifier. Hence, it is recommended to utilize a current amplifier instead of a generic OpAmp as in [11, 52].

Another limitation of the current sense amplifiers, is the fact that they present high error rates at low differential input voltages (usually below 1 mV). We evaluated the INA139, INA282, AD8219 and LTC6102 obtaining an error of 5 - 8 % at the input voltage of 500  $\mu$ V. On the other hand, the INA255, and the MAX9923 exhibit minimum errors at the same input voltage, roughly 2 %. This finding is not mentioned in the datasheets of the respective ICs, but is highly important when monitoring differential signals in the range of  $\mu$ V.

The MAX9923 is available at 3 different amplification levels (gain), 25, 100 and 250. We select to use the 25 gain, in order to effectively cover the whole range from 100 nA to 100 mA, using 3 current amplifiers, in combination with the selected shurt resistors.

**Dynamic Shunt Resistor Switch:** The shunt resistor switch features 3 resistors targeting 3 different current ranges. The values of the resistors have been selected at 1  $\Omega$ , 100  $\Omega$  and 10 k $\Omega$  aiming for the currents ranges of 100 mA to 1 mA, 1 mA to 10  $\mu$ A and 10  $\mu$ A to 100nA respectively. We refer to these bands as mA,  $\mu$ A and nA current ranges respectively for the rest of the paper. We also note that each current amplifier is set to obtain only 1 to 100 mV input voltage signal in order to sustain high accuracy. Of course, the shunt resistors should be as precise as possible while featuring ultra-low temperature coefficients to avoid drift due to temperate deviations. Such resistors are available in a 4 terminal format [42].

The tripping between different measurement ranges is realized by two ultra-fast comparators, the ADCMP601 featuring 4 ns delay. The ADCMP601 [46] requires a reference voltage, based on which it determines whether the obtained signal is below or above the desired threshold. In order to allow remote configuration we integrate the MCP4728 [44], which is a 12 bit DAC IC featuring 4 discrete outputs. The mentioned DAC, provides a resolution of roughly 800  $\mu$ V in the scale of 0 – 3.3 V, offering wide flexibility in the provided reference threshold. In our application, we configure the transition thresholds at 1 mA and at 10  $\mu$ A. An extra NAND gate is employed in order to drive the third switch (nA range) that induces an extra delay of 1 ns. The ADCMP601 features programmable hysteresis (from 2 - 160 mV) that is once again controlled through the MCP4728 DAC. Thus, we avoid fluctuations between two current ranges when the measured load is close to the given threshold. Lastly, an extra switch is connected across the first shunt resistor - that is controlled by the Beaglebone and used to bypass the input pins of the first amplifier - in order to measure its  $V_{OS}$  for calibration purposes as in [11].

**Power Supply & Voltage Reference Circuits:** ADCs are very sensitive ICs prone to external noise, thus must be treated very carefully in the design process. Especially when engaging embedded

eProfiler



Figure 3: Various Supply Rails Employed for the Evaluation Performance of ADS8332

systems, such as the BeagleBone, that incorporate high-frequency clocks and step-down regulators that commonly feature high voltage ripple, the process of powering the ADC and current amplifier is even more challenging [11]. To this end, eProfiler embeds the Analog Devices (AD) ADM7151 [34] ultra-low noise, high PSRR linear regulator. The regulator features the extremely low noise of 1  $\mu$ V and 94 dB PSRR. The voltage of the regulator is configured at 3.45 V through externals resistors. The ADC of eProfiler operates using a 3.3V reference value, but we selected to provide an extra 150 mV headroom to the on-board electronics to avoid potential saturation issues. It is worth noting that other power supply circuits such as the TPS7A49, which is recommended for powering precision analog applications [31, 50], features 12.7  $\mu$ V noise, while a typical regulator such as the LP2985 presents 30  $\mu$ V.

Typically, ADCs require a reference voltage to convert the obtained signals to the corresponding digital value. This voltage can significantly affect the performance of the system, thus an accurate voltage level, with extremely low noise, and high stability is required. For this purpose we employ the LTC6655 that features 625 nV peak-to-peak noise, high accuracy of 0.025 % and ultra-low drift of 2 ppm/°C. Notably, the LTC6655 outperforms the best in class reference circuits [8] (such as the REF3433 or the MAX6126).

**PCB Design Considerations:** We follow several principles for the design of the PCB board in order to attain high performance measurements. More specifically, we place a guard ring around the analog front-ends to prevent high-frequency noise from the digital parts of the system to distort the sensitive analog signals [11]. Also, we placed two separate ground planes [9, 40] for the digital and the analog parts that are connected through a pair of back-to-back Schottky [39] diodes and we place bypass capacitors [7] to all used ICs as close to their supply pins as possible. Lastly, a combination of ferrite beads and decoupling capacitors is employed in every power rail, used to attenuate the high frequency noise [29, 36].

# 3.4 Software Implementation

To reliably get measurements at specific time intervals, a real-time system is required. The Beaglebone helps greatly in this regard as it features two on-board PRUs [30]. The main processor, used by the Linux host, communicates with the PRUs via a section of shared memory. One more section of shared memory exists between the two PRUs available for communication. The latter of the two memory sections achieves faster communication, requiring only 1 cycle at a 200 MHz clock frequency, while the first memory requires 3 cycles. To fully control the ADS8332 we employ both PRUs programmed in assembly language for timing consistency. The first is used to control the ADC (ADC-PRU) and to obtain the measurements, while the second serves as a clock generator (Clock-PRU), allowing the user to precisely control the sampling

rate. Actually, the ADC-PRU reads the Clock-PRU state through their shared memory pool avoiding toggling and reading of GPIO states. The selected ADC communicates through the SPI protocol, thus a custom SPI implementation was developed on the ADC-PRU. The process of acquiring one measurement of 16 bits lasts roughly 835 ns. When sampling at 500 kSps, the time interval between the measurements is 2000 ns, thus the remaining available time is 1165 ns. Within this period the acquired data are saved into the shared memory. Moreover, within the available time the desired GPIOs are monitored by the PRU, such as the outputs of the comparators (SW2, SW3), as well as the GPIOs used to trace events on the DUT. Similarly, these GPIOs are used as actuation pins. These actions require only one cycle by the PRU, while 3 additional cycles are required to write the obtained values into the memory (at 20 MHz speed). Notably, we perform GPIO measurements / actuation at every sampling cycle, thus the maximum attained delay is 2  $\mu$ s; faster than the delay presented in FlockLab, which varies from 90 to 280  $\mu$ s. Lastly, we note that an intuitive, network accessible User-Interface (UI) is developed, using the Angular framework for the front-end, and the Node.js platform for the back-end aiding the measurement process.

### 4 PERFORMANCE EVALUATION

In this section we evaluate the performance of the eProfiler.

#### 4.1 **Power Supply Evaluation**

In this section we highlight how different power sources can affect the performance of our system. We consider 3 different power rails, the output rail of the ADM7151 which is at 3.45 V, and the 3.3 V and 5 V rails of the BeagleBone. Figure 3, illustrates the measured rails with our tool. Apparently, the voltage rail supplied by the ADM7151 is quite stable featuring only 450  $\mu$ V ripple voltage, while the 3.3 V and the 5V rails present 3.3 mV and 67.5 mV, respectively. Next we power the ADS8332 by employing the different rails while measuring a fixed reference voltage. The obtained results when powering the ADC with the ADM7151 introduces a STDEV of 0.1 (the values used in calculations are in mV), while the STDEV when the 3.3 V and 5V rails are employed are 6 and 80 respectively. Notably, the deviation of the obtained reference voltage is extremely high when powering the ADC with the Beaglebone's 3.3 V and the 5 V rails. This clearly highlights the impact of a power source on the performance of measurement acquisition systems [33, 50].

#### 4.2 Comparators' Evaluation

In this experiment we evaluate the response speed of the comparators in our system. To this end, we artificially bypass the  $\mu$ A and nA ranges through external jumpers, to disable the auto-ranging feature, performing only mA range measurements. However, the two comparators remain operational, both configured to monitor ENSsys '20, November 16-19, 2020, Virtual Event, Japan



Figure 4: Comparators Evaluation & Measurement Error Performance of MAX9923 and eProfiler & RV1805 Power Draw

the output of the first current amplifier (mA range). Moreover, the BeagleBone samples their output values through the PRU unit. As external load, we engage the eZ430-RF2500 mote [6], and configure it to switch from its active state to its Low-Power Mode 0 (LPM0) state, and then to its Low-Power Mode 2 (LPM2) state. The current consumption of the mote is roughly 3.8 mA, 1 mA and 510  $\mu$ A in the given states, respectively. Thus, we set the thresholds in the comparators accordingly, to match the given currents. Notably, a 20  $\Omega$  shunt resistor was employed to better adapt to the given currents. Next we monitor this setup by sampling only the output of the first current amplifier (mA range) with a speed of 500 kSps. The Fig. 4(a) illustrates the draw of the node with the blue line, while the orange, and the green lines plot the outputs of the two comparators (signals SW2 and SW3). The comparator detects the current state of the eZ430 promptly, alternating its output accordingly.

#### 4.3 **Obtained Accuracy**

In this set of experiments we evaluate the accuracy of the acquisition system. We first measure the accuracy of the MAX9923 by varying its input voltage and recording its output with a high-end voltagemeter (Keysight 34470A). The Fig. 4(b) illustrates the obtained error in the input voltage range from 0.1 mV to 10 mV. The MAX9923 presents notable error at 0.1 mV input (6.2 %), thus we decided to avoid applying such low signals. Similarly, it attains roughly 1.89 % error at 0.5 mV, while its error at 1 mV is only 1.08 %. When measuring, higher input voltages its accuracy improves further, introducing only 0.12 % error at 10 mV input and 0.013 % at 100 mV input. By properly selecting the shunt resistors, we can always maintain the input voltage polled by the MAX9923 in the 1 mV – 100 mV range; the obtained error will never exceed the aforementioned 1.08 % regardless if measuring nA,  $\mu$ A or mA currents.

Next we proceed by repeating the same experiment while measuring the obtained results with the on-board ADC in order to evaluate the accuracy of the entire system. In this experiment we vary the input voltage from 1 mV to 100 mV as actually happens when measuring the different ranges. Fig. 4(c) plots the obtained error. Apparently, the proposed device features a maximum error of 1.6 %, obtained at 1 mV input voltage, while the error at 10 mV and 100 mV input voltages is 0.33 % and 0.018 % respectively. It is worth noting that the average error is only 0.45 %, which illustrates an outstanding performance. Notably, the eProfiler can be used to measure even lower currents, in the range of 10 nA to 50 nA, however the obtained error at these inputs is 7.8 % and 4.4 %, respectively. This happens because the ADC unit fails to monitor such low voltage values. However, the whole setup can be re-adjusted with different shunt resistor values to support different current ranges. For instance, the values of 10  $\Omega$ , 1 k $\Omega$  and 100 k $\Omega$  can be employed to support the current range from 1 nA to 1 mA, maintaining the same accuracy, as highlighted in the next experiment.

In this experiment we characterize the power consumption of an ultra-low power Real-Time-Clock (RTC), the RV1805 [38]; the corresponding results from our measurement device are illustrated in Fig. 4(d). The RV1805 dissipates only 17 nA in its time keeping mode, however to retain the stated accuracy an auto-calibrating function - lasting for several seconds - is enacted on a periodic basis. In Fig. 4(d), we demonstrate this auto-calibrating feature, lasting from second 20 to 70, alongside the corresponding power draw increase. Notably, we have configured the RV1805 to provide interrupt signals every 10 seconds, resulting in burst currents of 13.6  $\mu$ A. Our obtained results are validated by the consumption values provided in the manufacturer's datasheet, as well as by the acquired measurements leveraging the uCurrent meter [47] observing minute discrepancies (less than 0.11 % obtained error). Our setup highlights the prominence of a high-sampling tool featuring extreme dynamic range to obtain high fidelity power consumption measurements even in nano-Ampere current ranges.

Notably, the proposed meter has been successfully employed to evaluate the power draw of several devices as well as discrete modules in our previous works [13–20].

# 5 CONCLUSIONS

In this work we present eProfiler, a high-fidelity power consumption meter with long-term monitoring capabilities. The meter supports wide current ranges spanning from 100 nA to 100 mA, and features an average error of only of 0.45 % throughout its measuring range. It is the first meter to support monitoring of nA currents which are becoming increasingly popular in IoT ecosystems. The eProfiler integrates GPIO state tracing to allow correlation of the obtained results with specific events realized on the DUT. All the aforementioned features aid in the power profile characterization of IoT devices. The eProfiler provides an in-situ, economically reasonable, and precise measuring framework, especially, for energy-harvesting applications. Leveraging eProfiler the consumption profile assessment of sensor nodes, and the efficiency of energy-harvesters can be effectively evaluated, enabling realistic estimations for the lifetime of deployed IoT sensor networks. Lastly, we note that the accuracy of the eProfiler can be further improved by applying the calibration method presented in the Excalibur [25].

eProfiler

- MAX9923 Ultra-Precision High-Side Current-Sense Amplifier. [Acc. 20-Oct-2020]. https://www.maximintegrated.com/en/products/analog/amplifiers/ MAX9923.html.
- [2] BeagleBone Black Rev. C Embedded board. [Acc. 20-Oct-2020]. http://beagleboard. org/BLACK.
- [3] Bradford Campbell, Branden Ghena, and Prabal Dutta. 2014. Energy-Harvesting Thermoelectric Sensing for Unobtrusive Water and Appliance Metering. In Proceedings of the 2nd International Workshop on Energy Neutral Sensing Systems (ENSsys) (Memphis, Tennessee) (ENSsys '14).
- [4] Behnam Dezfouli, Immanuel Amirtharaj, and Chia-Chi (Chelsey) Li. 2018. EM-PIOT: An energy measurement platform for wireless IoT devices. *Journal of Network and Computer Applications* (2018).
- [5] Prabal Dutta, Mark Feldmeier, Joseph Paradiso, and David Culler. 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 (St. Louis, MO, USA) (IPSN '08).
- [6] eZ430 RF2500. [Acc. 20-Oct-2020]. https://tinyurl.com/7z7zav9.
- [7] Guy Hoover. [2010]. Consequences of Not Following the Data Sheet Bypass Recommendation.
- [8] Hubert Halloin, Pierre Prat, and Julien Brossard. 2013. Long term characterization of voltage references.
- [9] Hank Zumbahlen. [2012]. Staying Well Grounded.
- [10] A. Hergenroeder, J. Wilke, and D. Meier. 2010. Distributed Energy Measurements in WSN Testbeds with a Sensor Node Management Device (SNMD). In Proceedings of the 23th International Conference on Architecture of Computing Systems 2010.
- [11] Xiaofan Jiang, Prabal Dutta, David Culler, and Ion Stoica. 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).
- [12] G. Kazdaridis, S. Keranidis, P. Symeonidis, P. S. Dias, P. Gonçalves, B. Loureiro, P. Gjanci, and C. Petrioli. 2017. EVERUN: Enabling Power Consumption Monitoring in Underwater Networking Platforms. In Proceedings of the 11th Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization (Snowbird, Utah, USA) (WiNTECH '17).
- [13] G. Kazdaridis, S. Keranidis, P. Symeonidis, P. Tzimotoudis, I. Zographopoulos, P. Skrimponis, and T. Korakis. 2019. Evaluation of LoRa Performance in a City-Wide Testbed: Experimentation Insights and Findings. In Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (Los Cabos, Mexico) (WiNTECH '19).
- [14] G. Kazdaridis, N. Sidiropoulos, I. Zografopoulos, P. Symeonidis, and T. Korakis. 2020. The ICARUS Mote: Employing Off-Chip RTC to Attain 22 nA Sleep Current in Duty-Cycled IoT Devices. In Proceedings of the 10th International Conference on the Internet of Things (Malmö, Sweden) (IoT '20 Companion).
- [15] G. Kazdaridis, N. Sidiropoulos, I. Zografopoulos, P. Symeonidis, and T. Korakis. 2020. Nano-Things: Pushing Sleep Current Consumption to the Limits in IoT Platforms. In Proceedings of the 10th International Conference on the Internet of Things (Malmö, Sweden) (IoT '20).
- [16] G. Kazdaridis, P. Skrimponis, I. Zographopoulos, P. Symeonidis, T. Korakis, and L. Tassiulas. 2017. Enabling Asynchronous Awakenings in Wireless Sensor Networks Towards Removing Duty-Cycle Barriers. In Proceedings of the 11th Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization (Snowbird, Utah, USA) (WiNTECH '17).
- [17] G. Kazdaridis, D. Stavropoulos, V. Maglogiannis, T. Korakis, S. Lalis, and L. Tassiulas. 2014. NITOS BikesNet: Enabling Mobile Sensing Experiments through the OMF Framework in a City-Wide Environment. In Proceedings of the 15th International Conference on Mobile Data Management (IEEE MDM '14).
- [18] G. Kazdaridis, P. Symeonidis, I. Zographopoulos, T. Korakis, K. Klun, and N. Kovac. 2017. On the development of energy-efficient communications for marine monitoring deployments. In 2017 13th International Conference on Advanced Technologies, Systems and Services in Telecommunications (TELSIKS).
- [19] G. Kazdaridis, P. Tzimotoudis, N. Sidiropoulos, P. Symeonidis, and T. Korakis. 2020. The NITOS Wireless Sensor Network Testbed for Experimenting with Long-Range Technologies. In 10th International Conference on the Internet of Things Companion (Malmö, Sweden) (IoT '20 Companion).
- [20] G. Kazdaridis, I. Zographopoulos, N. Sidiropoulos, P. Symeonidis, and T. Korakis. 2019. Nano Power Draw in Duty-Cycled Wireless Sensor Networks. In Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (Los Cabos, Mexico) (WiNTECH '19).
- [21] G. Kazdaridis, I. Zographopoulos, P. Symeonidis, P. Skrimponis, T. Korakis, and L. Tassiulas. 2017. In-situ Power Consumption Meter for Sensor Networks supporting Extreme Dynamic Range. In Proceedings of the 11th Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization (Snowbird, Utah, USA) (WiNTECH '17).

- [22] S. Keranidis, G. Kazdaridis, V. Passas, T. Korakis, I. Koutsopoulos, and L. Tassiulas. 2013. Online energy consumption monitoring of wireless testbed infrastructure through the NITOS EMF framework. In Proceedings of the 8th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (Miami, Florida, USA) (WiNTECH '13).
- [23] S. Keranidis, G. Kazdaridis, V. Passas, T. Korakis, I. Koutsopoulos, and L. Tassiulas. 2014. NITOS Energy Monitoring Framework: Real Time Power Monitoring in Experimental Wireless Network Deployments. SIGMOBILE Mob. Comput. Commun. Rev. (2014).
- [24] Nordic Power Profiler Kit. [Acc. 20-Oct-2020]. https://tinyurl.com/sf22708.
- [25] C. Li and B. Dezfouli. 2019. Excalibur: An Accurate, Scalable, and Low-Cost Calibration Tool for Sensing Devices. *IEEE Sensors Journal* (2019).
- [26] R. Lim, F. Ferrari, M. Zimmerling, C. Walser, P. Sommer, and J. Beutel. 2013. Flock-Lab: 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* (Philadelphia, Pennsylvania, USA) (IPSN '13).
- [27] P. Martin, Z. Charbiwala, and M. Srivastava. 2012. DoubleDip: Leveraging Thermoelectric Harvesting for Low Power Monitoring of Sporadic Water Use. In Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (Toronto, Ontario, Canada) (SenSys '12).
- [28] Current Ranger Power Meter. [Acc. 20-Oct-2020]. https://lowpowerlab.com/guide/ currentranger/.
- [29] Michael Cobb. [2009]. Powering High Speed Analog-to-Digital Converters with Switching Power Supplies.
- [30] Molloy Derek. [2019]. Exploring BeagleBone: tools and techniques for building with embedded Linux.
- [31] TPS7A49 Low noise High-PSRR Low-dropout Voltage Regulator. [Acc. 20-Oct-2020]. https://www.ti.com/product/TPS7A49.
- [32] A. Pötsch, A. Berger, C. Leitner, and A. Springer. 2014. A power measurement system for accurate energy profiling of embedded wireless systems. In Proceedings of IEEE Emerging Technology and Factory Automation (ETFA).
- [33] Xavier Ramus. April 2016. Minimum power specifications for high-performance ADC power-supply designs.
- [34] ADM7151 Ultra Low Noise High PSRR RF Linear Regulator. [Acc. 20-Oct-2020]. https://www.analog.com/en/products/adm7151.html.
- [35] M. Riedel and M. Schwarzbach. [Acc. 20-Oct-2020]. Current measurement with auto-ranging yields 180dB range dynamics. http://www.imc-berlin.com/.
- [36] Rob Reeder. [2012]. Designing Power Supplies for High Speed ADC.
- [37] M. Rottleuthner, T. C. Schmidt, and M. Wählisch. 2019. Eco: A Hardware-Software Co-Design for InSitu Power Measurement on Low-End IoT Systems. In Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems (ENSsys'19).
- [38] RV1805-C3 Real-Time-Clock. [Acc. 20-Oct-2020]. https://tinyurl.com/y5v9fklx.
- [39] Sanjay Pithadia and Shridhar More. [2013]. Grounding in mixed-signal systems demystified, Part 1.
- [40] Sanjay Pithadia and Shridhar More. [2013]. Grounding in mixed-signal systems demystified, Part 2.
- [41] A. Schulman, T. Thapliyal, S. Katti, N. Spring, D. Levin, and P. Dutta. 2016. BattOr: Plug-and-debug energy debugging for applications on smartphones and laptops Stanford CS Tech Report.
- [42] VISHAY Shunt Current Shunts and Current-Sensing Resistors. [Acc. 20-Oct-2020]. https://tinyurl.com/yxg2mtv4.
- [43] Tim Regan, Jon Munson, Greg Zimmer and Michael Stokowski. [2005]. Current Sense Circuit Collection - Making Sense of Current.
- [44] MCP4728 12-Bit Quad Digital to Analog Converter. [Acc. 20-Oct-2020]. https: //ww1.microchip.com/downloads/en/DeviceDoc/22187E.pdf.
- [45] ADS8332 Analog to Digital Converter. [Acc. 20-Oct-2020]. http://www.ti.com/ product/ads8332.
- [46] ADCMP601 Rail to Rail TTL/CMOS Comparator. [Acc. 20-Oct-2020]. https: //www.analog.com/en/products/adcmp601.html.
- [47] uCurrent Power Meter. [Acc. 20-Oct-2020]. http://tiny.cc/dz069y.
- [48] Walt Kester. [2015]. Which ADC Architecture Is Right for Your Application?
- [49] Waspmote Device. [Acc. 20-Oct-2020]. https://tinyurl.com/a7qvx5c.
- [50] Four ways to improve your ADC system power supply rejection Texas Instruments. [Acc. 20-Oct-2020]. https://tinvurl.com/y24x820x.
- [51] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta. 2012. Grafting Energy-Harvesting Leaves onto the Sensornet Tree. In Proceedings of the 11th International Conference on Information Processing in Sensor Networks (IPSN) (Beijing, China) (IPSN '12).
- [52] Ruogu Zhou and Guoliang Xing. 2013. Nemo: A High-fidelity Noninvasive Power Meter System for Wireless Sensor Networks. In Proceedings of the 12th International Conference on Information Processing in Sensor Networks (Philadelphia, Pennsylvania, USA) (IPSN '13).