

Assessing the impact of sensor-based task scheduling on battery lifetime in IoT devices

Original

Assessing the impact of sensor-based task scheduling on battery lifetime in IoT devices / Chen, Yukai; Wang, Wenlong; Jahier Pagliari, Daniele; Macii, Enrico; Poncino, Massimo. - In: IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT. - ISSN 0018-9456. - ELETTRONICO. - 70:(2021), pp. 1-15. [10.1109/TIM.2021.3088498]

Availability:

This version is available at: 11583/2913541 since: 2021-09-07T14:47:56Z

Publisher:

IEEE

Published

DOI:10.1109/TIM.2021.3088498

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Assessing the Impact of Sensor-based Task Scheduling on Battery Lifetime in IoT Devices

Yukai Chen, *Member, IEEE*, Wenlong Wang, *Member, IEEE*, Daniele Jahier Pagliari, *Member, IEEE*, Enrico Macii, *Fellow, IEEE*, Massimo Poncino, *Fellow, IEEE*

Abstract—A well-known system-level strategy to reduce the energy consumption of microprocessors or microcontrollers is to organize the scheduling of the executed tasks so that it is aware of the main battery non-idealities. In the IoT domain, devices rely on simpler microcontrollers; workloads are less rich and, batteries are typically sized to guarantee lifetimes of more extensive orders of magnitude (e.g., days, as opposed to hours). Load current magnitudes in these IoT devices are therefore relatively small compared to other more powerful devices, and they hardly trigger the conditions that emphasize the battery non-idealities. In this work, we carry out a measurement-based assessment about *whether task scheduling is really relevant to extend the lifetime* of IoT devices. We run experiments both on a physical commercial IoT device hosting four sensors, an MCU, and a wireless radio, as well as on a “synthetic” device emulated with a programmable load generator. We used both secondary lithium-ion and primary alkaline batteries to explore the impact of battery chemistries further. Results show that the impact of different schedules is essentially irrelevant, with a maximum difference of only 3.98% in battery lifetime between the best and worst schedules.

Index Terms—Power Management, Task Scheduling, IoT, Battery-powered Device, Multi-sensor Device, Measurement Scheduling.

I. INTRODUCTION

The requirement of highly energy-efficient operations has made battery-powered sensor-based devices become the ground in which Dynamic Power Management (DPM) strategies have found their most diversified applications [1]–[6]. Some of these solutions explicitly take into consideration the fact that devices are battery-powered; this substantially amounts to realize that there is a substantial difference between the power *consumed* by the device and that is actually *drawn* by the battery [2], [4], [6]. Two fundamental factors result in this decoupling. First, a DC-DC converter is typically placed between the battery and the load, whose conversion efficiency depends on the operating input and output power levels; Secondly, the battery cannot provide any arbitrary amount of power with the same efficiency; this is the so-called “rated capacity” effect, which represents the fact that the effective capacity of a battery (i.e., its available energy) decreases as the discharge current increases.

The former factor can be mostly taken into account at design time by choosing the best devices (batteries and converters)

Y.K. Chen, W.L. Wang, D. Pagliari, and M. Poncino are with the Department of Control and Computer Engineering, Politecnico di Torino, Italy. (E-mail: yukai.chen@polito.it, wenlong.wang@polito.it, daniele.jahier@polito.it, massimo.poncino@polito.it). E. Macii is with the Interuniversity Department of Regional and Urban Studies and Planning, Politecnico di Torino, Italy. (E-mail: enrico.macii@polito.it)

that can provide the lowest possible losses. The latter aspect, however, pertains exclusively to the operations of the load. Existing solutions do account for this problem by focusing on the scheduling of the computational tasks executed on the hardware: of the generic applications [1], [4], of the sensing tasks in a multi-sensor device [7], [8], of the transmission tasks of the sensor nodes [9], [10], or globally for the tasks executed in the various nodes of a wireless sensor network [3].

These solutions generally try to reshape the current profile of the tasks in order to address the two above non-idealities, i.e., to match power levels of battery and load to reduce converters’ losses and reduce the rated capacity effect. This can be achieved typically by scheduling tasks so that the load profile avoids significant current variations and reduces the peaks, compatibly with the dependencies of the tasks. While these works provide general scheduling guidelines that can reduce the battery drain, they have two main shortcomings.

First and foremost, most of these works focus only on computational or communication-related tasks that only involve the processor [4], [11] and neglect the impact sensing tasks, which may have non-negligible power demands [2]. Typical Internet of Things (IoT) devices are often equipped with several sensors [8], [12], and their aggregate power demand tends to be more significant than that of the other tasks. At the same time, however, sensing tasks offer additional degrees of freedom in designing the task schedule. Sensing operations are usually *independent of each other*; thus, they can be freely scheduled without dependence constraints. As such, they can be fully serialized (to minimize peak currents) or fully parallel (to maximize idle intervals), with all the intermediate partial combined serial/parallel options. These two factors (their power demand and the number of additional scheduling options they offer) are the apparent motivation of why sensing tasks should not be neglected in determining the optimal battery-aware task scheduling.

Neglecting the sensing task has, however, a second and more subtle implication. Different sensors require quite diverse supply voltages, implying that each sensor will be powered by a distinct DC-DC converter, resulting in different efficiency and power losses. While this could be simply handled using a constant efficiency, the reality is that DC-DC converters’ efficiency is positively affected by the output current (in this case, the sensors’ currents); ignoring this variability will usually lead to over-optimistic results.

A second issue with previous works is that they tend to overrate the *recovery effect* in batteries, i.e., adding rest periods during discharge might prolong battery operation time

(obviously, net of the inserted time intervals). Many methods propose strategies based on stretching idle times among tasks to exploit the recovery effect (e.g., [11], [13], [14]). However, the magnitude of the recovery effect is strongly dependent on the battery chemistry; recent works have shown that the recovery effect is virtually absent [15] in typical lithium-ion batteries. Therefore, an IoT device powered by a lithium-ion battery playing with idle times will hardly benefit. While in primary (non-rechargeable) batteries, the effect of recovery can be more evident, its impact will still be limited with respect to that achieved by the scheduling of the computational tasks.

The work of [7] partially address these shortcomings by taking into account all types of tasks executed in the device as well as the conversion efficiency of the converters, while virtually neglecting the recovery effect; it calculates battery-aware scheduling of tasks through a heuristic exploration, showing that a battery-aware task scheduling can prolong battery lifetime as much as 20% with respect to a battery- and converter-agnostic one. However, that work relies on results obtained from simulation, without an experimental validation based on measurements. The issue of measurement-based validation is particularly relevant because battery models tend to be accurate for discharge rates (i.e. C-rates. A 1C rate for a battery with an X mAh capacity corresponds to a current of X mA) close around the nominal capacity (i.e., $0.2C \dots 2C$).

Conversely, the typically low duty cycles of IoT devices and the lifetime requirements of days or weeks are such that the average power consumption of the whole device ends up being relatively small, e.g., below the mA scale. As a result, it might be that the levels of power are not large enough to trigger the non-idealities of batteries appreciably, especially for some battery chemistries. For instance, an advanced lithium-ion battery can handle a discharge rate in mA range without actually experiencing a measurable rated capacity effect. Moreover, this low-level discharge current strongly affects the conversion efficiency of the IoT device's converters.

In this work, we present an experimental validation of existing task scheduling methods by carrying out measurements to estimate the lifetime of an IoT multi-sensor device when powered by the lithium-ion battery and the primary battery, with the objective to assess the actual impact of different scheduling of the tasks executed by the device. In particular, our experimental measurements implements the tasks scheduling solution proposed in [7], including sensing, computation, and communication tasks. The scheduling exploits the freedom offered by mutual independence of sensing tasks and the power/speed trade-off of the Micro Controller Unit (MCU); moreover, it considers DC-DC converter efficiency resulting from each module's power level. The conclusion drawn from our experiments is that in a realistic scenario *task scheduling hardly impacts the lifetime* of the device, both in the case of lithium-ion and of the primary batteries.

The paper is organized as follows. Section II provides some background on the three main battery non-idealities that are the subject of our study and reviews the current state-of-the-art devices scheduling works; Section III describes the characterization of a generic multi-sensor IoT device and how the scheduling of the tasks executed on it can impact the power

consumption. The experimental analysis is carried out in Section IV, we conduct two different sets of measurements to investigate the influence of task scheduling on the IoT device lifetime. Finally, Section V summarizes the main observations that can be drawn from our experiments.

II. BACKGROUND

A. The Non-Idealities of IoT Device Batteries

IoT devices are primarily sensor-based electronic devices. These IoT device nodes are typically deployed by dispersing them over a wide physical area, and it is impossible to rely on the utility grid. Therefore, batteries have become the most widely used energy storage devices in IoT applications. Although the self-powered IoT devices by energy harvesting are proposed for designing power-autonomous green IoT devices, the current solution for power supply in the IoT field is still using batteries as energy storage devices. According to the different IoT application scenarios, designers choose either a primary or rechargeable battery for their IoT devices [16]. The authors of [17] reports that batteries used for low-power sensory and IoT devices are mainly very low-cost primary batteries due to low cost and self-discharge, and high capacity. For instance, Netatmo and X-sense offer smart smoke alarms as IoT devices powered by primary batteries used to detect and monitor toxic gas in real-time by mobile phone; Petsafe, Xiaomi, and Dogness produce smart pet feeders accept primary batteries as secondary power supply. However, with the development of green IoT applications and the cost reduction of rechargeable batteries, rechargeable batteries will be more widely used in the future. Whatever the battery type and chemistry, a battery is far from being an ideal voltage source; the power it can provide is strongly dependent on the charge and discharge currents profile. In this work, we are fundamentally interested in the discharge process, although most of the discussion below can apply to the battery charge.

1) *Rated Capacity Effect*: The most commonly acknowledged non-idealities of battery is the *rated capacity* effect, which expresses the change in battery capacity at different discharge rates. When a battery is continuously discharged, the available battery capacity decreases as the discharge current increases.




Figure 1. Primary battery capacity dependence on the discharge current [18]

Figure 1 shows the relation between battery capacity and discharge current of the alkaline battery E91 [18] by Energizer. The dependence shown in Figure 1 indicates that as the discharge current increases from 25mA to 500mA, the corresponding available capacity decreases from 3,000mAh to 1,700mAh (For this battery, the datasheet does not provide a nominal capacity). This implies that for the largest discharge currents, the E91 behaves approximately as a half-capacity battery with respect to the case of the smallest current.

In the case of rechargeable batteries, the rated capacity in datasheets is usually described as shown in Figure 2, i.e., as voltage vs. capacity curve, parameterized with respect to discharge current. Figure 2 refers to the lithium-ion battery NCR18650B [19] by PANASONIC. The usable capacities are obtained when the battery voltage reaches the cut-off voltage value (2.5V in this case). We can immediately notice how the sensitivity to current is quite modest and shows only a small improvement for a 0.2C current (i.e., 3,200mA*0.2 = 640mA). Comparing Figure 1 and Figure 2 indicates that the rated capacity effect is much more significant in the primary battery than it in the lithium-ion battery.

Discharge Characteristics (by rate of discharge)




Figure 2. Lithium-ion battery voltage vs. discharge current curves [19]

2) *Capacity Dependence on Discharge Variation*: The plots in the previous sections do refer to a constant current discharge. In a real scenario, load currents will exhibit significant variations in the time domain but also in the frequency domain. Generally speaking, a battery is more efficient if there is no variation in the load profile. Therefore, a constant current is a favorable case for measurements, and any variation will degrade the equivalent capacity of the battery. This information is seldom provided by datasheets but has been described by many research works [20]–[22].

As an example of the variation over time, consider a periodic square wave alternating with a given duty cycle D between two current values I_1 (D) and I_2 ($1-D$); the average value of the current will be $I_{avg} = D * I_1 + (1-D) * I_2$. Figure 3 shows the SOC traces during battery discharge of Panasonic lithium-ion UR16650ZT battery with an 800mA constant discharge profile and two square wave discharge profiles of period cycle 2,000s, 50% duty cycle, average value 800mA, two

different swing values ($\pm 300mA$ and $\pm 500mA$) [20]. The SOC traces indicate a difference in discharge time between a constant discharge at 800mA with respect to the square wave discharge at $I_{avg} = 800mA$. The magnitude depends on the difference between I_1 and I_2 : the larger, the more evident the difference will be. As shown in the figure, the square wave's discharge time with $\pm 500mA$ swing values has a more significant difference to the constant discharge time than the square wave with $\pm 300mA$ swing values.




Figure 3. Capacity dependence on discharge variation: constant discharge vs. periodic square wave discharge [20]

3) *Recovery Effect*: A third effect is the so-called *recovery effect*. Compared to the previous two, this one is potentially beneficial from the battery duration perspective. It states that if the battery stays in periods of rest (zero or very little discharge) for a sufficient duration, capacity lost during discharge can be rejuvenated to a certain extent. Figure 4 shows the measurement results of battery discharge voltage traces from three different load profiles, and one is a continuous discharge scenario; the other two are intermittent discharge with different idle periods cases. The curves reveal that the voltage rises when the load is reduced, as shown by the intermittent discharge curve, and the overall discharge time extends. Furthermore, the longer period of the rest, the more voltage recovered, as illustrated by two intermittent discharge profiles shown in Figure 4.




Figure 4. Battery discharge voltage profiles of continuous and intermittent discharge by 250mA discharge rate with Energizer E91 primary battery.

The work of [23] investigates the capacity recovery effect concerning discharge current, Depth of Discharge (DOD), and recovery break time; the measurements show that the recovery

effect of the battery is positively correlated with these three parameters, and it is possible to build a capacity recovery battery model based on these three parameters. On the other hand, the experimental results in [15] reveal that the recovery effect almost does not exist in the Alkaline, Nickel-Metal Hydride (NiMH) and lithium-ion battery chemistries, which are commonly used as power supplies for IoT applications.

B. Energy-efficient Power Management Techniques

Several different techniques for prolonging the lifetime of battery-powered sensor-based devices have been proposed in the literature, focusing on different application contexts.

A large part of these existing works corresponds to the energy optimization and lifetime enhancement for the Wireless Sensor Networks (WSN) [5], [8], [10], [24]–[27]. The WSN is constructed by a certain amount of sensor nodes and gateway nodes. Each node equips with its energy storage devices. In most cases, the energy storage components of nodes are not identical. Not all the energy storage components are replaceable or rechargeable, making nodes have a different lifetime. When a node finishes its energy storage device, it is considered “dead”. If a certain number of nodes lose functionality, the whole network cannot operate, which tells the end of the WSN lifetime. The works related to the WSN usually focus on the whole network lifetime instead of focusing on one single device node. For this reason, the works related to WSN do not care about techniques used for prolonging the lifetime from the perspective of battery. They usually use high-level techniques to optimize protocol or network configuration [28], [29].

The existing works related to WSN have one consensus that sensors’ power consumption is not comparable to the transmission power consumption [10], [30], [31]. There are works related to WSN focus on the scheduling of sensors while they schedule the sensors in the network scale, instead of scheduling multiple sensors within one node [7], [24], [26], [29] propose a similar idea to schedule the different sensors in the whole WSN. [26] illustrates that each sensor has its coverage area in the network; there are overlaps among the sensor coverage area; it proposes using the overlapped area to activate or deactivate some sensor nodes in order to achieve an optimal network lifetime. Similar to these works, the work [27] focuses on the gateway in the network.

The possibility of harnessing battery recovery effect in IoT and WSN applications to reduce battery size or extend battery life is investigated. These works have assumed the existence of this recovery effect and proposed different power management techniques in the form of power supply architectures (e.g., multiple battery setup) and communication protocols (e.g., burst mode transmission) in order to exploit it. The work of [13] demonstrates the recovery effect can prolong the battery lifetime by up to 45% by modifying discharging profile and inactive periods rather than constantly discharging, while this work only study the AAA NiMH battery chemistry. [6] proposes to use multi-battery for each node. There is always one battery works and the others in the rest state. While the work does not mention the battery’s chemistry, and they use a straightforward linear battery model to prove

the proposed method. Leveraging the recovery effect during the communication period to prolong the system lifetime is studied in [14]. The work [9] devise the optimal transmission task scheduling by considering the recovery effect. However, the work [15] declares that the lithium-ion battery’s recovery effect is negligible based on their measurements. Another critical point related to the battery in the existing works is that the used battery model’s accuracy is not enough for conducting accurate power estimation. The work [32] try to estimate the lifetime of WSN by adapting an analytical kinetic battery model, while this model cannot account for capacity dependence on load current variation, and the used model is only limited to the NiMH battery.

Besides the WSN applications, in the traditional real-time embedded system applications, there are existing works [1], [3], [4], [7] target on the scheduling computation tasks to derive the battery-friendly discharge current profile by considering the non-ideal battery discharge characteristics. The work [33] proposes to deploy the super-capacitor and battery simultaneously in the wireless sensor nodes to construct a hybrid energy storage system to extend the lifetime. In case of not convenient for replacing or charging batteries in the WSN, [31] study to include the energy harvesting devices to optimize the lifetime by leveraging environmental resources convert to electrical energy.

III. BATTERY-AWARE TASK SCHEDULING

Battery-aware scheduling of the various tasks run by a IoT device must carefully balance the main non-ideal discharge characteristics of batteries described in Section II; the dependence on the current magnitude, on current variance, and on periods of rest.

In order to better understand the interplay of these three effects, we define a straightforward model of our target IoT device. We assume that the device hosts:

- N_S sensors (each one represented by a corresponding task), each one taking a time $T_{s,i}$ to complete sensing and requiring a power $P_{s,i}$. The idle power consumption of each sensor is $P_{s,i,off}$.
- An MCU with an idle power consumption of $P_{c,off}$, and N_p power states associated to different voltage/frequency points; at a given voltage/frequency point $i = (1, \dots, N_p)$, the MCU takes $T_{c,i}$ to execute its tasks and requires a power $P_{c,i}$. Power states are sorted in increasing values of power.
- A wireless transceiver implementing some given protocol. For simplicity we assume that each message has a fixed size and transmission requires always the same T_{tx} and T_{rx} , consuming respectively a power P_{tx} and P_{rx} , and the power consumption in idle state is $P_{x,off}$.

There is only one simple (yet realistic) dependence rule: sensing must occur before processing, and transmission must follow processing. Therefore, there is total freedom in the scheduling of the sensing tasks. Notice that the MCU is indeed active during the sensing tasks as it has to drive the collection of the sensed data through its interfaces.

Finally, the global execution model consists of a duty-cycled workload of a period of T_{cycle} . The scheduling determines the

execution sequence of the tasks, which results in an “active interval” T_{active} ; while in the remaining part (“idle interval”) of the cycle $T_{idle} = T_{cycle} - T_{active}$, all the modules of the device are put into the idle state.

Figure 5 shows two possible schedules of a system hosting four sensors. The heights of the bars denote power consumption, and the widths indicate the time; hence, each bar’s area is the energy. It is important to emphasize that the two plots are identical from the point of view of the energy consumption: the total areas of the various “boxes” is the same in both cases. What makes them different for the battery are the above-mentioned battery non-idealities.




Figure 5. Two extreme cases of task scheduling.

These plots describe two scenarios that privilege one battery property over the others. The schedule on the top features all the sensing in series, followed by data processing and transmission. This schedule has the smallest peak power (so it will suffer less from the rated capacity effect), but it will have the shortest idle interval (so it will exploit less recovery effect). Notice that sensing tasks are executed in decreasing order of power consumption not by chance: it is a commonly accepted principle that this is best for the battery as the most larger currents are drawn in the maximal charge conditions, i.e., at the beginning of the cycle [1]–[3].

Conversely, in the bottom plot, all sensing tasks are executed in parallel at the beginning of the period. This scheduling will stress the battery more, yet for a shorter time and maximize the idle interval. Even if we assume that battery recovery is modest, we have to consider that states that P_{off} is likely to be 2 or 3 orders of magnitude smaller than active power; therefore, there will be a limit in compress the activity period, after which an increase in the idle period will not be convenient anymore. It is therefore possible that an intermediate schedule in which only a subset of the sensor tasks is executed in parallel will result in the best balance between the two effects.

Given that the degrees of freedom essentially lies in the scheduling of the sensing tasks and that, therefore, the number of options is limited (having more than ten sensors is quite unusual), the work of [7] proposes a heuristic exploration tool to determine the schedule that yields the most extended battery lifetime. The method of simulation-based exploration for the

most energy-efficient task scheduling proposed in [7] adopts the circuit equivalent battery model in the simulation. Figure 6 illustrates the structure of the circuit equivalent battery model, and it can account for the capacity dependencies on the current magnitude and dynamics. All the model parameters can be extracted from the battery’s datasheet, and the model can be applied to both lithium-ion batteries and primary batteries.




Figure 6. Circuit equivalent battery model for the simulation exploration [20].

The full battery charge stored capacity can be modelled by a full-capacity capacitor (C). This is simply obtained by converting the nominal battery capacity C_{nom} (in Ah) using Equation 1, where $1V$ is the initial voltage across the capacitor that defines a fully charged battery [34].

$$C = \frac{C_{nom} \times 3600}{1V} \quad (1)$$

The voltage generator $V_{oc}(SOC)$ of the model describes the relation between SOC and voltage. $R(SOC)$ represents the battery internal resistance. The methodology to extract these two quantities is described in [35]. It needs two voltage versus capacity curves as shown in Figure 2. Two discharge current (I_1 and I_2) curves are fitted to build two functions $V_{batt}^{I_1}(SOC)$ and $V_{batt}^{I_2}(SOC)$. The parameters $V_{oc}(SOC)$ and $R(SOC)$ are derived by solving the equations as follows:

$$\begin{cases} V_{oc}(SOC) = R(SOC) \times I_{batt}^{I_1} + V_{batt}^{I_1}(SOC) \\ V_{oc}(SOC) = R(SOC) \times I_{batt}^{I_2} + V_{batt}^{I_2}(SOC) \end{cases} \quad (2)$$

$$R(SOC) = \frac{V_{batt}^{I_1}(SOC) - V_{batt}^{I_2}(SOC)}{I_{batt}^{I_2} - I_{batt}^{I_1}} \quad (3)$$

$$V_{oc}(SOC) = V_{batt}^{I_1}(SOC) + R(SOC) \times I_{batt}^{I_1} \quad (4)$$

The two pairs of RC groups in series of the battery model account for the battery’s sensitivity to the load dynamics. The first RC group (parallel between $R_S(SOC)$ and $C_S(SOC)$) defines the short-time constant $\tau_S = R_S \times C_S$ of the battery voltage response due to the double layer capacity effects; the second RC group (parallel between $R_L(SOC)$ and $C_L(SOC)$) defines the long-time constant $\tau_L = R_L \times C_L$ of the battery voltage response due to the characteristic diffusion effects in the electrolyte. The work [35] provides the detailed steps to extract these RC groups’ quantities from the datasheet if it provides the voltage trace of a pulse current. There are existing works compute these parameters by conducting pulse current test. Since it is too difficult to obtain the parameters of these two RC groups and they do not affect the long-term simulation results, this circuit equivalent model is frequently used with only the $R(SOC)$ in the right part for the simulation does not focus on the instantaneous simulation results.

Two voltage generators $V_{lost}(I_{batt})$ and $V_{lost}(f_{load})$ on the left side part of the battery model account for the rated capacity effect $C(I_{batt})$ and the effect of load current variation $C(f_{load})$ as described in the section II; these two voltage generators bring a voltage drop at the SOC node, thus affect the battery SOC. $V_{lost}(I_{batt})$ is derived by computing, at each simulation time step Δt , the following equation:

$$\Delta SOC(I_{batt}) = \frac{I_{Batt} \times \Delta t}{C(I_{batt})} - \frac{I_{Batt} \times \Delta t}{C_{nom}} \quad (5)$$

where $C(I_{Batt})$ is the relationship between capacity and battery current can be derived from datasheet as described in [36] and C_{nom} is the nominal capacity. The effect of the discharge variation is not an instantaneous quantity; therefore, the model uses the Short Time Fourier Transform (STFT) to compute load frequency components in each time interval window. $V_{lost}(f_{load})$ is obtained by evaluating Equation 6 at each Δt .

$$\Delta SOC(f_{load}) = \sum_{i=1}^{N_{FFT}} \left(\frac{I_{Batt}(i) \times \Delta t}{C(f_{load})} - \frac{I_{Batt}(i) \times \Delta t}{C_{nom}} \right) \quad (6)$$

N_{FFT} is length of timing window in STFT; $I_{Batt}(i)$, $i = 1, \dots, N_{FFT}$ is a string of current values within a timing window; $C(f_{load})$ is the relation between capacity and load frequency. The method proposed in [36] shows that using the information in Figure 2 can derive the relation between discharge energy and current and the relation between discharge time and current. Then the relation between discharge power and current is computed based on the previous two relations. Using the power and current relation and the energy and current relation can draw *Ragone plot*. The diagonals in the *Ragone plane* indicate the discharge time; the inverse of each discharge time represents a frequency. Thus the relation between energy and frequency is extracted; after converting energy to capacity, the $C(f_{load})$ is computed.

In our work, we will use the result of such exploration to compare, through measurements, the alleged optimal energy-efficient schedule against the two extremes of Figure 5.

IV. EXPERIMENTAL RESULTS

A. Experiments on an Industrial IoT Device

1) *Device Description*: Our first validation is run on an industry-level multi-sensor IoT device prototype shown in Figure 7 which has been manufactured as a final deliverable of one EU-funded project and is currently being commercialized by one of the industrial project partners.

The minimum board supply voltage of this industrial IoT device is 3.5V. It hosts four different sensor modules, one low-power MCU, and one wireless transceiver. The detailed specifications of each module are given as follows:

- A low-power MCU by NXP (LPC54114 [37]), based on ARM CORTEX-M4 architecture. The MCU controls the activation and de-activation of the sensors according to the scheduling. It also collects the data from each sensor and performs a linear regression of the sensed data [38]. The data to be transmitted are simply the trend (i.e., the slope of the extrapolated line) and the last received




Figure 7. The IoT multi-sensor device prototype: (A) ARM CORTEX-M4 architecture MCU module; (B) Wireless transceiver module; (C) Air quality sensor module; (D) Air quality poisonous gas sensor module; (E) Noise sensor module; (F) Humidity and temperature sensor module.

datum. In this way, the receiver (not considered here) can reconstruct the data set as a set of piece-wise segments.

- A wireless transceiver Texas Instruments (TI) CC2538 [39], compliant to the IEEE 802.15 standard [40]. The radio transmit gathered data at the end of every period.
- Four sensor modules by Mikroelektronika provided as ClickTM boards; each board includes the sensor and other electronic components (e.g., DC-DC converter and analog to digital converter).
 - Humidity and temperature sensor module [41], it contains a HDC1000 sensor [42] by TI. It measures relative humidity 0-100% with an accuracy of 3% and temperature in a range [-20°C, 85°C] with a 0.2°C accuracy.
 - Air quality sensor module [43], which embeds the SGP30 sensor [44] from Sensirion; it provides detailed information of the air quality parameters by measuring various types of gases and outputs them in the form of two complementary air quality readings: Total Volatile Organic Compounds (TVOC) and CO_2 .
 - Noise sensor [45] module carries the SPG0410HR5H-B [46] surface mount silicon microphone from SiSonic to detect the noise in the environment.
 - Air quality poisonous gas sensor module [47] includes an MQ135 [48] sensor for detecting ammonia, nitrogen oxides benzene, smoke, and other poisonous gases.

Table I lists the rated voltage and current of the six modules (MCU, transmitter, and four sensors) of the test IoT device.

It is important to emphasize that the sensor modules, as discussed above, include the sensors and other electronic components; therefore, the sensor module's power consumption is larger than that of the sensor itself. For completeness of the information, Table II summarizes the electrical characteristics of the sensors hosted by the sensor modules.

Notice how the whole module power consumption is much

Table I
MODULES ELECTRICAL PARAMETERS OF INDUSTRIAL IOT DEVICE

Module	Voltage (V)	I (mA)
LPC54114 (MCU) [37]	3.3	6.0
CC2538 (Transmitter) [39]	3.3	7.0
Temperature and Humidity (Sensor) [41]	3.3	6.5
Air Quality (Sensor) [43]	3.3	55.5
Noise (Sensor) [45]	3.3	11.0
Poisonous Gases (Sensor) [47]	5.0	56.4

Table II
ELECTRICAL PARAMETERS OF SENSOR IN EACH SENSOR MODULE.

Sensor	Sensor Module	Voltage (V)	I (mA)
HDC1000 [42]	Temperature and Humidity	2.7 to 5.5	0.18
SGP30 [44]	Air Quality	1.62 - 1.98	48.2
SPQ0410HR5H-B [46]	Noise	1.5 - 3.6	0.12
MQ-135 [48]	Poisonous Gases	5.0	≤ 160

larger than one sensor alone in some cases. For instance, an HDC1000 sensor's current is only 0.18mA, while the corresponding module is rated as a 6.5mA current load. In this specific case, the difference is due to the presence of LEDs, voltage regulators, and some essential electronic components such as amplifiers, pull up/down resistors, capacitors, etc. As these are inherently part of the specific sensor in order to make it "usable", in the evaluation of the schedule by using the method of [7], we adopt the power consumption values of sensor modules, i.e., those in Table I.

2) *Battery Description*: We selected the rechargeable lithium-ion battery NCR18650B [19] from PANASONIC and the primary AA battery E91 [18] from Energizer in our experiments. To achieve accurate measurement results for the lithium-ion battery, we adopted a battery gauge (BQ27z561 [49]) produced by TI to record and monitor battery discharge characteristics.

The battery gauge has its independent power source and is equipped with an Evaluation Module (EVM) with all essential peripherals for measurement and communication. The EVM provides interfaces to connect the lithium-ion battery and the load device to the gauge. All captured data are transmitted to the PC via a Micro USB interface and can be visualized by a software called BqStudio, which TI provides for interacting with the gauge. The gauge implements the Impedance Track algorithm developed by TI. It consists of two main phases, the learning phase and the usage phase. Before the learning phase, the chosen battery's chemistry file provided by TI should be selected and loaded in BqStudio. TI provides a large number of chemistry files corresponding to different batteries. The chemistry file typically consists of the battery's basic electrical parameters and their corresponding lookup tables with specific parameters (SOC, T(temperature)), such as nominal capacity, cut-off voltage, OCV[SOC, T] table, and internal resistance[SOC, T] table. Moreover, some parameters related to the charging and discharging process to guarantee safety (e.g., battery voltage thresholds for charging and discharging phases, maximum allowable current of the gauge) should be

uploaded to the internal flash memory through BqStudio.

In the learning phase, the battery is first discharged from the predefined SOC_1 to SOC_2 ; then inserted a rest interval; after the rest, the battery is charged to SOC_1 , relaxed again; finally discharged again from SOC_1 to SOC_2 . The battery impedance is learned during the learning phase ($R[SOC, T] = (V-OCV[SOC, T])/I$), and the maximum battery capacity is computed ($Q_{max} = \text{Total Passed Charge}/(SOC_1 - SOC_2)$). In the usage phase, the gauge builds a battery model according to the data provided in the chemistry file and data captured in the learning phase, and it runs the simulations with this battery model and uses the new measured data to calibrate the derived battery model. We set the SOC_1 and SOC_2 to be 99% and 1% in the learning phase, respectively; the discharge current was 0.2C, and the charging current is 1C. The rest intervals are 5 hours after the discharge phase and 2 hours after the charge phase; these are the values suggested by the gauge user manual.

As the battery gauge is not designed for AA batteries, the HP 34401A digital multimeter is used to measure the primary AA battery discharging characteristics. According to the rated AA battery's nominal voltage, a set of three E91 batteries is used for supplying the IoT multi-sensor device to satisfy the minimum board supply voltage.

Table III shows the key electrical data of the two batteries. Figure 8 shows the measurement experiments under the lithium-ion battery scenario.

Table III
MANUFACTURER'S PARAMETERS OF THE EMPLOYED BATTERIES.

Parameters	NCR18650B	E91
Rated Capacity	3,200mAh	3,000mAh-25mA 2,500mAh-100mA 2,000mAh-250mA
Nominal Voltage	3.6V	1.5V
Cut-off Voltage	2.5V	Varies w.r.t. discharge conditions




Figure 8. (A) IoT multi-sensor device prototype; (B) Panasonic NCR18650B lithium-ion battery; (C) Texas Instruments battery gauge BQ27z561.

3) *Workload Description and Task Scheduling*: The workload executed by the device is structured as a periodic sequence of the following major tasks as described in Section III: sensing from the four sensors, processing, and transmission of the processed data to a sink node. Note that the operations of the sink node are irrelevant for the purpose of our analysis. Although there are only four sensors in the IoT device prototype, it is impractical to carry out experimental measurements on all the different schedules. We decided to use the simulation-based design space exploration methodology proposed in [7] as indicated in Section III to obtain the optimal and worst energy-efficient schedules. Notice that this task requires building battery models of the two kinds of batteries described above and the various modules of the IoT device as listed in III because the method of [7] is based on simulation. The batteries are modelled by the circuit equivalent battery as indicated in Figure 6. All other modules are considered as loads of system; therefore, they are modelled by the voltage and current drain with the values as listed in Table I. Further detailed modelling and simulation explanation please refer [7].

We then analyzed the impact of different schedules on battery operation time by measuring the difference between the optimal and worst schedules, which can also be seen as the upper bound of the impact of different schedules on battery lifetime. In this practical IoT device case, the heuristic exploration proved to be marginally helpful. It yields the fully serial schedule as the best (longest operation time) one, with all the tasks in parallel as the worst one. This differs from the results shown in [7] because different battery and load current profiles were used.

Notice that full serial/parallel refer to the sensing tasks only, as there are precedence constraints such that all sensing tasks should precede the processing tasks, which must be completed before the transmission task. The two schedules are conceptually illustrated in Figure 5. Notice that sensing tasks follow a monotonically decreasing order in the full serial schedule (from most to least power demanding). We then apply these two schedules to the device and run the measurements.

From the point of view of the battery, the experiment is designed as follows. The device is turned on with a fully charged battery; it executes the corresponding schedule (serial or parallel) for a predefined time so that all the schedules have the same operation time. After one execution of the experiment, the lithium-ion battery is recharged to 100% SOC and set to rest for 12 hours, while the AA battery pack is replaced with the new one. After the rest period or the replacement, a new run of the different schedule is executed at the same given operation time.

Notice that it is crucial to guarantee that the batteries have the same initial voltage before starting each discharge cycle for the test. For the lithium-ion battery, even if we use the same battery cell and recharge it to 100% SOC after one schedule, there is still a tiny difference in the initial voltage; thus, we should discharge it to a predefined initial voltage each time before the experiments start. For the AA battery, the fresh AA battery typically has a slightly different initial voltage due to the manufacturing variations; we thus slightly discharge all the pack's batteries before the experiments to the

predefined initial voltage. In this way, we can guarantee that the battery's initial voltage remains consistent across different and successive runs. In the following experiments, the initial lithium-ion battery voltage is set as 4,100mV, and the initial AA battery pack voltage is set as 4,400mV.

In order to make the captured data stable and give the air quality and poisonous sensor modules enough time to warm up, we set each sensor module sampling period to 30 seconds.

Regarding the experimental environment's choice, we decided to perform the measurements in a general laboratory environment. As the power consumption of the sensor modules is independent of the measured quantity's magnitude, for our analysis (assessing power consumption for the different temporal sequences of operations), there is no difference in measuring indoor or outdoor quantities. It is also worth mentioning that our experiments were taken in the laboratory throughout the day. During normal working hour operations, with doors and windows opening and closing, a variable number of persons, and so on, these factors lead the laboratory environment is not very different from outdoor ones.

B. Measurements Results of the Industrial IoT Device

We use the battery voltage to measure the effect of different schedules on the device lifetime. Specifically, all schedules are applied for the same amount of operation time. At the end of this time, we select the schedule with the highest battery residual voltage (say, V_{max}) as the reference value, that is, the energy-optimal one. Then, on the voltage waveforms of the other schedules (they have all final voltages $< V_{max}$), we search the point in time where their voltages reached V_{max} . The difference in the time is the "benefit" obtained with the energy-optimal schedule.

1) *Measurements Results of Lithium-ion Battery*: The device operations consist of multiple cycles of duration: $T_{cycle} = 3$ minutes. Each cycle consists of sensing (30 seconds for each sensor), MCU operations (2 seconds), and transmission (3 seconds). Sensors are in an active mode for the same time so that the data transfer in the parallel schedule can be synchronized. MCU operations include transfer from the sensors and delivery to the wireless transmitter. The transmitter time includes the transceiver's startup time and the transmission overhead (modulation, protocol, encoding, etc.). The two schedules have therefore duty cycles of $(30 \times 4 + 2 + 3)/180 = 69\%$ (serial) and $(30 + 2 + 3)/180 = 19\%$ (parallel).

Figure 9 shows the comparison between the full serial and full parallel schedules for the 3-minute cycle described above with 1,440 minutes (24 hours) operation time. After the device operates for 24 hours, the residual voltage of the full serial schedule is slightly higher than that of the full parallel one (see the last cycle in Figure 9), which matches the simulation results [7]. The difference in discharge time between the two schedules can be seen by observing the last few discharge cycles. To quantify the difference, we select the battery voltage in the last cycle of the full serial task schedule as the reference value (3,739mV) and search at what cycle the voltage trace of the full parallel schedule reaches this voltage. This occurs nine cycles (27 minutes) before that of the serial schedule

as indicated by the red arrows in Figure 9, corresponding to a 1.875% (27/1,440) difference in discharge time after 24 hours. Notice that when the battery voltage reaches 3,739mV, the battery SOC is dropped to approximately 54%.




Figure 9. Lithium-ion battery results of 3-minute cycle and 1,440 minutes operation time: Voltage traces of last 15 cycles from 1,395 to 1,440 minutes.

To verify our experimental results' generality, we first repeated the measurements above with the same battery but with three different operation times (1,320 minutes, 1,488 minutes, 1,740 minutes) to estimate the influence of operation length. Then we adopted two more Panasonic NCR18650B cells to repeat the above experiments to strengthen the generality of our conclusion. Notice that enlarging operation time is not unrestricted; we should guarantee that the battery voltage is always higher than the minimum board supply voltage. Twelve dots indicate twelve measurement results (one colour for the same battery, and each colour includes four different operation times) corresponding to the 3-minute cycle period category shown in Figure 10. The cluster of these dots indicates that the difference between optimal and worst schedules keeps stable with different batteries and operation times; the rows belong to the 3-minute cycle period category in the Table IV illustrates that all the differences are marginal.

Notice that this magnitude is consistent with the difference of simulation results reported in [7] when a "robust" lithium-ion battery is used. More considerable differences were reported only when less powerful batteries are used.




Figure 10. Lithium-ion battery practical board measurement results with 3, 10, 30, and 60 minutes cycle period using 3 different batteries.

2) *Impact of Cycle Period Length:* The results of the previous section refer to a 3-minute cycle period. We ran another

set of experiments with a 10-minute, 30-minute, and 60-minute cycle and different operation times to assess the cycle period's impact. As we did in the 3-minute case, we adopt three different batteries in the measurements. For 10-minute cycle periods cases, we conduct three different operation times as indicated in Table IV. Figure 10 shows that the cluster of the dots becomes flatter than in 3-minute scenarios, which tells that all differences have minor variance compared to the ones of the 3-minute cycle period cases.

Table IV
MEASUREMENT RESULTS STATISTICS OF PRACTICAL BOARD WITH LITHIUM-ION BATTERY.

Cycle Period	Operation Time (mins)	Maximum Difference	Minimal Difference	Average Difference
3 mins	1,320	2.05%	1.82%	1.89%
	1,440	2.08%	1.88%	1.94%
	1,488	2.02%	1.81%	1.88%
	1,740	1.90%	1.72%	1.78%
10 mins	4,960	0.20%	0.20%	0.20%
	5,160	0.19%	0.19%	0.19%
	5,440	0.18%	0.18%	0.18%
30 mins	33,000	0.09%	0.09%	0.09%
60 mins	78,000	0.08%	0.08%	0.08%

Because the overall measurement becomes extremely time-consuming with 30-minute and 60-minute cycle periods, it becomes impractical to perform multiple measurements on each cell as we did when cycle periods are 3 and 10 minutes. Hence, we only performed one measurement on each cell with the same predefined operation time as indicated in Table IV to compare the difference between optimal and worst schedules. The cluster of the dots in Figure 10 and the quantitative values in Table IV illustrate that the difference between optimal and worst schedules is diminishing as the length of the cycle period increased, and the difference keeps stable with using different batteries. The 3-minute and 10-minute cycle length results also indicate that various operation time does not affect the difference between two schedules. This is due to the fact that, since the execution times of all tasks are the same and only the cycle period changes, the duty cycles tend to become similar, and therefore the difference fades away. As prolonging the cycle period virtually increases the time spent by the device in an idle state, we can infer that the device's static power is marginal, and increasing the cycle period will basically make the impact of the schedule virtually irrelevant.

Again, this is consistent with the conclusion drawn in [7], where it was observed by simulation that, as the cycle period increases, the impact of the task scheduling on battery discharge time becomes negligible.

3) *Measurements Results of Primary AA battery:* This experiment aims at assessing whether the recovery effect, which can be significant in primary batteries, yields different considerations with respect to the previous section.

For every single measurement in the experiments using primary batteries, we need a battery group that includes three primary batteries connected in series to satisfy the board power supply. Therefore, two battery groups are required to accomplish one comparison between the full serial and the full parallel schedules. We conducted 18 comparisons with 36

battery groups (108 batteries) for four different cycle periods (3-minute, 10-minute, 30-minute, and 60-minute). Figure 11 shows the difference between optimal and worst schedules under four different cycle periods. Notice that each point represents a comparison. We performed twice the measurements for each operation time. The two points of the 30-minute and 60-minute cycle have the same operation time as shown in Table V. The points belong to the 3-minute and 10-minute cycle periods have four and three different operation times, respectively, as indicated in Table V.




Figure 11. Primary battery practical board measurement results using with 3, 10, 30, and 60 minutes cycle period.

Figure 11 shows the overall results of different cycle periods, and they indicate that the difference between the two schedules in the primary battery cases does not decrease monotonically as the cycle length gets, which is different as illustrated by lithium-ion battery cases shown in Figure 10. When the cycle period is 10 minutes, it has the most significant difference among other cycles. Simultaneously, the difference in 30-minute and 60-minute cycle cases diminishes with the cycle increases as lithium-ion battery cases.

Table V
MEASUREMENT RESULTS STATISTICS OF PRACTICAL BOARD WITH
PRIMARY BATTERY.

Cycle Period	Operation Time (mins)	Maximum Difference	Minimal Difference	Average Difference
3 mins	900	2.33%	2.00%	2.17%
	918	2.29%	1.96%	2.12%
	960	2.19%	2.19%	2.19%
	1,120	2.06%	2.06%	2.06%
10 mins	4,140	3.86%	3.62%	3.74%
	4,200	3.81%	3.81%	3.81%
	4,460	3.98%	3.76%	3.87%
30 mins	18,000	1.83%	1.67%	1.75%
60 mins	48,000	1.13%	1.13%	1.13%

For the experiments with a 3-minute cycle period, the idle time in the full parallel schedule is about 2.5 times longer than the full serial one so that the battery can exploit recovery further. This results in the full parallel schedule being better than the serial one, as shown in Figure 12, it indicates that the difference of discharge times is only 2.19% when the operation time is 16 hours. The numerical results of other operation times under a 3-minute cycle period scenario are listed in Table V; all the differences are less than 2.35%.

In the case of the 10-minute cycle period, the situation is, however, reversed as illustrated by Figure 12 and Figure 13.

As the idle state intervals are similar for the two schedules, and both allow a good amount of recovery between two consecutive cycles, the full serial schedule now outperforms the full parallel one due to its relatively low maximum current magnitude. The cluster of green colour points in Figure 11 shows that the variance of the difference between different batteries and operation times is insignificant; Table V indicates its maximum difference is 3.98%, while the minimum difference is 3.62%.

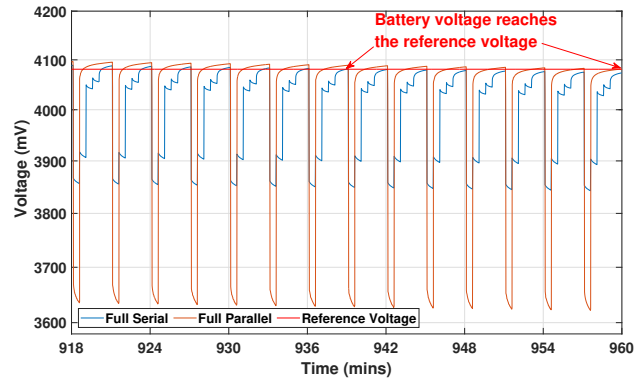


Figure 12. Primary battery with 3-minute cycle and 960 minutes operation time: Voltage traces of last 14 cycles from 918 to 960 minutes.




Figure 13. Primary battery with 10-minute cycle and 4,200 minutes operation time: Voltage traces of last 20 cycles from 4,000 to 4,200 minutes.

The 30-minute and 60-minute cycle period cases have a relatively smaller difference than the short cycle cases, which are around 1.75% and 1.13% as indicated in Table V respectively. The full serial schedule is better than the full parallel one. This is because the long enough idle interval for both schedules can fully exploit the recovery effect and the full serial schedule has less rated capacity effect. Furthermore, with the growth of the operation duration, the low duty cycle leads the relative difference between the two schedules to diminish because most of the time in each cycle is in an idle state. It is equivalent to saying that the schedule has only a tiny portion in each cycle affecting the operation time; therefore, the impact of different schedules on the overall operation time fades away.

It is emphasized again that this is consistent with the results obtained from [7]: as the cycle length grows, the effect of

different schedules on the battery discharge time becomes negligible. Concerning the magnitude of the difference between schedules, this experiment is consistent with the results of the lithium-ion battery's measurements. However, the recovery effect is not negligible and can reverse the two schedules' relative performance (with parallel being better than serial). Generally speaking, we can say that the recovery effect is the dominating effect for the AA primary battery when the cycle period is short, and the rated capacity effect is the dominating effect when the cycle period is extended.

C. Experiments on a Synthetic IoT Device

The main limitation of the practical IoT device used in the previous section is that the design space exploration of the task schedule is limited due to minor differences in the power consumption among the four sensor modules, as shown in Table I. In particular, the air quality and poisonous gas sensor modules have similar power consumptions, and the same applies to the temperature and noise sensor modules. In order to allow a better design space exploration, we designed a synthetic multi-sensor device by using a programmable load generator, which can emulate a device equipped with sensors that have more diverse load power values.

1) *Synthetic Multi-sensor IoT Device Configuration*: This experiment assumes that the multi-sensor IoT device we used is equipped with six different sensor modules. Table VI shows the power consumption of the six emulated sensor modules; we assume the power supply of each module is fixed as 3.3V, and the minimum board supply voltage of the device is 3.5V.

Table VI
POWER CONSUMPTION OF EACH MODULE IN THE SYNTHETIC DEVICE.

Module	Active Power (mW)	Idle Power (mW)
Sensor 1	300	12.00
Sensor 2	250	10.00
Sensor 3	200	8.00
Sensor 4	150	6.00
Sensor 5	100	4.00
Sensor 6	50	1.50
MCU	30-20-10	0.75
Transmitter	30	0.75

The operations are the same as our practical IoT device. The same lithium-ion battery NCR18650B and primary AA battery E91 are adopted in the synthetic scenarios. We also used the method of [7] to explore the optimal energy-efficient schedule and obtained the same results as the previous prototype IoT device, the full serial task schedule (best) being more energy-efficient than the full parallel task (worst). Therefore, the full serial and full parallel task schedules are deployed in this experiment. Additionally, we also selected one intermediate schedule. The details of three task schedules deployed in synthetic IoT device are as follows:

- Full serial schedule: all sensors are activated one after another in non-increasing order of power consumption. Each sensor has a 15-seconds sampling period. After all the sensors finish sampling, MCU is activated for 3 seconds to manipulate the captured data. Then the transmitter starts to transfer data and finishes data transmission

after 2 seconds. As soon as the transmitter finishes the operation, the device goes to the idle state until the next cycle starts.

- Full parallel task schedule: instead of activating the sensors sequentially, all the sensors are activated simultaneously by a 15-seconds sampling period at the beginning of each cycle. The following procedure is identical to the full serial policy.
- Intermediate task schedule: the six sensors are divided into three groups. The sensors in each group collect data simultaneously. Groups are activated one after another in serial. After all groups finish sampling, the rest procedure is identical to the previous two schedules.

The three deployed task schedule are sketched in Figure 14.




Figure 14. Three task schedules deployed in the experimental measurement.

2) *Synthetic Device Measurements Setup*: The constructed measurement environment is shown in Figure 15. We used a programmable electronic load (RIGOL DL3021) [50], which generates the power consumption trace corresponding to the desired task schedule (component B). The battery gauge (component F) measures the lithium-ion battery discharge characteristics, and the multimeter (component A) measures the AA battery discharge characteristics. The component C shown in the Figure 15 is programmable linear DC power supply [51], it is used for recharging lithium-ion battery.




Figure 15. Synthetic device measurement environment: (A) HP 34401A digital multimeter; (B) RIGOL programmable electronic load DL3021; (C) RIGOL programmable linear DC power supply DP711; (D) Texas Instrument BqStudio software interface; (E) Panasonic NCR18650b lithium-ion battery; (F) Texas Instruments battery gauge BQ27Z561.

The measurement procedure is the same as in previous experiments. The fully charged battery is connected with the programmable load generator. We calibrate the initial battery voltage before each task schedule is executed as described before; then, the three schedules are executed for the same predefined operation time to compare the difference of battery voltage traces. We adopted three lithium-ion batteries and 162 primary batteries (each measurement need three cells to satisfy the board minimal operation voltage, and each comparison need three runs with three schedules shown in Figure 14) in the following measurement with different cycle periods, and we also defined various operation time as we did in previous experiments to investigate the influence of cycle length, operation time, and different batteries.

D. Measurements Results of Synthetic IoT Device

1) *Measurements Results of Lithium-ion Battery:* Figure 16 graphically shows the voltage discharge traces of 5-minute cycle and 5,280 minutes operation time obtained by the measurements. As before, we measured energy efficiency in terms of the residual battery voltage. After discharging 5,280 minutes, the full serial schedule exhibits the largest residual battery voltage, which is 3,789mV; the battery SOC is around 55% at this point; followed by the intermediate one and the full parallel being the worst.




Figure 16. Lithium-ion battery with 5-minute cycle and 5,280 minutes operation time: Voltage traces of last 40 cycles from 5,080 to 5,280 minutes.

As we did the measurements for the industrial IoT device, we conducted multiple measurements by considering three different lithium-ion batteries, four distinct cycle periods, and numerous operation time as illustrated in Figure 17 and Table VII.

Figure 17 shows that all the measurement results obtained using a synthetic device are entirely consistent with those on the practical device, despite the more considerable variance of the sensor's current. The difference between the worst schedule and the optimal schedule is monotonically decreasing as the cycle period increases, as indicated by the cluster of points in Figure 17 becomes flatter as the cycle period increases. Note that the full serial schedule is always the energy-optimal one as shown in Figure 16, even if by a small amount. Table VII exhibits that the maximum difference among all these different conditions for the synthetic IoT




Figure 17. Lithium-ion battery synthetic experiment results with 5, 10, 30, and 60 minutes cycle period using 3 different batteries.

Table VII
MEASUREMENT RESULTS STATISTICS OF SYNTHETIC DEVICE WITH LITHIUM-ION BATTERY.

Cycle Period	Operation Time (mins)	Maximum Difference	Minimal Difference	Average Difference
5 mins	4,950	3.43%	3.33%	3.37%
	5,120	3.52%	3.32%	3.42%
	5,280	3.22%	3.12%	3.19%
	5,500	3.27%	3.18%	3.21%
10 mins	9,000	2.22%	2.11%	2.19%
	9,200	2.31%	2.09%	2.20%
	9,400	2.28%	2.17%	2.21%
30 mins	38,250	1.02%	0.94%	0.97%
60 mins	84,000	0.79%	0.71%	0.74%

device with lithium-ion battery is 3.52%, and the minimal difference is only 0.71%.

2) *Measurements Results of Primary AA Battery:* The battery setup is the same as the experiment on the practical device: three AA batteries are connected in series, and the discharge conditions are identical to the case of lithium-ion batteries. We applied then the same three schedules discussed in the previous section.

The battery voltage traces of three task schedules with a 5-minute cycle and 3,180 minutes operation time are depicted in Figure 18. The plots show that the residual battery voltages of the three different schedule are almost indistinguishable towards the end of each cycle, and their difference becomes negligible. The difference between full serial and full parallel schedules is only 1.57%. We conducted the measurements with other operation times as shown in Table VIII, and for each operation time, we measured twice. All the measurement results are plotted by the red points in Figure 19. The rows belong to the 5-minute cycle period category in Table VIII indicate the maximum difference between full serial and full parallel schedules is 1.68%, and the minimum difference is 1.45% under different operation time conditions.

Notice that the ranking of the schedules is now reversed: as in the experiment using the industrial IoT prototype device, the full serial schedule has the highest power consumption, the full parallel schedule is the optimal energy-efficient one.

When switching to the experiments in which a 10-minute cycle is used, the experiment results show that the full serial schedule achieves the highest residual voltage, while the full parallel schedule achieves the lowest residual voltage, as




Figure 18. Primary battery with 5-minute cycle and 3,180 minutes operation time: Voltage traces of last 12 cycles from 3,120 to 3,180 minute.




Figure 19. Primary battery synthetic experiment results using with 5, 10, 30, and 60 minutes cycle period.

Table VIII
MEASUREMENT RESULTS STATISTICS OF SYNTHETIC DEVICE WITH
PRIMARY BATTERY.

Cycle Period	Operation Time (mins)	Maximum Difference	Minimal Difference	Average Difference
5 mins	2,970	1.68%	1.52%	1.60%
	3,060	1.63%	1.47%	1.55%
	3,180	1.57%	1.57%	1.57%
	3,450	1.59%	1.45%	1.52%
10 mins	6,000	2.83%	2.67%	2.75%
	6,300	2.70%	2.54%	2.62%
	6,600	2.88%	2.73%	2.80%
30 mins	18,000	1.17%	1.17%	1.17%
60 mins	36,000	1.00%	0.83%	0.92%

shown in Figure 20, Table VIII lists the results of other operation time under the 10-minute cycle period, indicating the green dots cluster in Figure 19 has maximum difference is 2.88% and the minimal difference 2.54%.

The results of the 10-minute cycle period conditions are the same as the practical industrial IoT device's measurement results: the optimal schedule is flipped between full parallel and full serial. We analyze the reasons for this above. The idle time becomes large enough to allow all the three schedules to exploit the recovery effect further; therefore, the rated capacity effect becomes the dominant effect under 10-minute cycle scenarios. Specifically, the shorter cycle period has a limited length of the idle period; the schedule with a relatively long idle period gives the battery a long time to exploit its recovery effect, thus achieving high residual voltage. On the other hand,

when the idle period increases to longer enough in the long cycle cases, all the schedules have adequate time to leverage the recovery effect, then the rated capacity effect plays a pivotal role in affecting the residual voltage.




Figure 20. Primary battery with 10-minute cycle and 6,600 minutes operation time: Voltage traces of last 20 cycles from 6,400 to 6,600 minute.

From the relative position of the dots cluster shown in Figure 19, we observe that the results are similar to the results of practical industrial IoT device with primary battery shown in Figure 11. The different batteries and operation durations do not affect the difference between optimal and worst schedules. The difference under the 10-minute scenario has the largest value among all cycle periods. The optimal schedule under 30-minute and 60-minute cases both are full serial schedule as 10-minute scenario, and the differences between the two schedules with 30-minute and 60-minute cycle period are pretty small (around 1.0%) because of the super low duty cycle and long operation time as described in practical IoT device case.

V. CONCLUSIONS


As one of the possible knobs to reduce the energy demand of embedded devices used for IoT applications, some researchers have focused on designing a battery-aware power consumption profile by acting on scheduling the tasks executed by the device. In this work, the results of our measurements reveal that this knob is not practical for such small-scale devices. The impact of different task schedules in terms of the device lifetime is negligible; this irrelevance is a consequence of the relative low-power consumption of the workload, which hardly triggers the non-ideal characteristics of the battery, and is due to two main facts: firstly, to the nature of IoT workloads, which perform relatively simple tasks and typically executed periodically with quite low duty cycles; secondly, as these devices are designed for a lifetime in the order of days or weeks, batteries are sized to support these workloads for a long period of times. Quantitatively, the maximum difference between the schedule with the longest and shortest lifetimes was 3.98% for primary batteries and 3.52% for rechargeable batteries, which are in the range of error of the measurement instrumentation. It is worth mentioning that, despite the slight differences, as the theory suggests, in most configurations, a schedule where all the tasks are executed serially yields the

most extended lifetime, as this profile diminishes the battery rated capacity effect. However, for primary batteries, where the recovery effect is more pronounced than in rechargeable ones, some configurations of the load and the amount of available idle time can sometimes reverse this finding.

REFERENCES

- [1] P. Chowdhury and C. Chakrabarti, "Static task-scheduling algorithms for battery-powered dvs systems," *IEEE Transactions on very large scale integration (VLSI) systems*, vol. 13, no. 2, pp. 226–237, 2005.
- [2] S. Hong, D. Kim, and J.-e. Kim, "Battery aware real time task scheduling in wireless sensor networks," in *11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA'05)*. IEEE, 2005, pp. 269–272.
- [3] Y. Xiang, X. Xin, and W. Wang, "Priority-based low-power task scheduling for wireless sensor network," in *2009 international symposium on autonomous decentralized systems*. IEEE, 2009, pp. 1–5.
- [4] K. Wei, W. Zhang, Y. Yang, G. Song, and Z. Zhang, "Battery-aware task scheduling for energy efficient mobile devices," *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, vol. 97, no. 9, pp. 1971–1974, 2014.
- [5] R. Khajuria and S. Gupta, "Energy optimization and lifetime enhancement techniques in wireless sensor networks: A survey," in *International Conference on Computing, Communication & Automation*. IEEE, 2015, pp. 396–402.
- [6] R. Jin, Z. Che, Z. Wang, M. Zhu, and L. Wang, "Battery optimal scheduling based on energy balance in wireless sensor networks," *IET Wireless Sensor Systems*, vol. 5, no. 6, pp. 277–282, 2015.
- [7] Y. Chen, D. Pagliari, E. Macii, and M. Poncino, "Battery-aware design exploration of scheduling policies for multi-sensor devices," in *Proceedings of the 2018 on Great Lakes Symposium on VLSI*, 2018, pp. 201–206.
- [8] J. Lin, W. Xiao, F. L. Lewis, and L. Xie, "Energy-efficient distributed adaptive multisensor scheduling for target tracking in wireless sensor networks," *IEEE Transactions on Instrumentation and Measurement*, vol. 58, no. 6, pp. 1886–1896, 2008.
- [9] H. Li, C. Yi, and Y. Li, "Battery-friendly packet transmission algorithms for wireless sensor networks," *IEEE sensors journal*, vol. 13, no. 10, pp. 3548–3557, 2013.
- [10] P. Györke and B. Pataki, "Energy-aware measurement scheduling in wsns used in aal applications," *IEEE Transactions on Instrumentation and Measurement*, vol. 62, no. 5, pp. 1318–1325, 2013.
- [11] V. Mahima and A. Chitra, "Battery recovery based lifetime enhancement (brle) algorithm for wireless sensor network," *Wireless Personal Communications*, vol. 97, no. 4, pp. 6541–6557, 2017.
- [12] P. Schonle, G. Rovere, F. Glaser, J. Bossler, N. Brun, X. Han, T. Burger, S. Fateh, Q. Wang, L. Benini *et al.*, "A multi-sensor and parallel processing SoC for wearable and implantable telemetry systems," in *ESSCIRC 2017*. IEEE, 2017, pp. 215–218.
- [13] C.-K. Chau, F. Qin, S. Sayed, M. H. Wahab, and Y. Yang, "Harnessing battery recovery effect in wireless sensor networks: Experiments and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 7, pp. 1222–1232, 2010.
- [14] F. Sun, C. Yi, and Y. Li, "Battery-friendly scheduling policy in MAC layer for WBAN data packets transmission," *IET Communications*, vol. 11, no. 9, pp. 1423–1430, 2017.
- [15] S. Narayanaswamy, S. Schlueter, S. Steinhorst, M. Lukasiewicz, S. Chakraborty, and H. E. Hoster, "On battery recovery effect in wireless sensor nodes," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 21, no. 4, p. 60, 2016.
- [16] Y. Shim, H. Park, and W. Shin, "Joint time allocation for wireless energy harvesting decode-and-forward relay based IoT networks with rechargeable and non-rechargeable batteries," *IEEE Internet of Things Journal*, 2020.
- [17] M. Radfar, A. Nakhlestani, H. Le Viet, and A. Desai, "Battery management technique to reduce standby energy consumption in ultra-low power IoT and sensory applications," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 67, no. 1, pp. 336–345, 2019.
- [18] Energizer, "ENERGIZER E91 AA battery datasheet," <https://data.energizer.com/pdfs/e91.pdf>, 2019.
- [19] PANASONIC, "NCR18650B Li-Ion 18650 cell datasheet," <https://www.batteryspace.com/prod-specs/NCR18650B.pdf>, 2019.
- [20] Y. Chen, E. Macii, and M. Poncino, "A circuit-equivalent battery model accounting for the dependency on load frequency," in *Proceedings of the Conference on Design, Automation & Test in Europe*. European Design and Automation Association, 2017, pp. 1177–1182.
- [21] D. Baek, Y. Chen, A. Bocca, L. Bottaccioli, S. Di Cataldo, D. J. Pagliari, E. Patti, G. Urgese, N. Chang *et al.*, "Battery-aware operation range estimation for terrestrial and aerial electric vehicles," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 6, pp. 5471–5482, 2019.
- [22] Y. Chen, S. Vinco, D. J. Pagliari, P. Montuschi, E. Macii, and M. Poncino, "Modeling and simulation of Cyber-Physical Electrical Energy Systems with SystemC-AMS," *IEEE Transactions on Sustainable Computing*, vol. 5, no. 4, pp. 552–567, 2020.
- [23] C. Maurer, W. Commerell, A. Hintennach, and A. Jossen, "Capacity recovery effect in lithium sulfur batteries for electric vehicles," *World Electric Vehicle Journal*, vol. 9, no. 2, p. 34, 2018.
- [24] S. Chand, B. Kumar *et al.*, "Maximising network lifetime for target coverage problem in wireless sensor networks," *IET Wireless Sensor Systems*, vol. 6, no. 6, pp. 192–197, 2016.
- [25] P. Györke and B. Pataki, "Application of energy-harvesting in wireless sensor networks using predictive scheduling," in *2012 IEEE International Instrumentation and Measurement Technology Conference Proceedings*. IEEE, 2012, pp. 582–587.
- [26] B. Shah, A. Abbas, G. Ali, F. Iqbal, A. M. Khattak, O. Alfandi, and K.-I. Kim, "Guaranteed lifetime protocol for IoT based wireless sensor networks with multiple constraints," *Ad Hoc Networks*, p. 102158, 2020.
- [27] S. A. Karthikeya, R. Narayanan *et al.*, "Power-aware gateway connectivity in battery-powered dynamic IoT networks," *Computer Networks*, vol. 130, pp. 81–93, 2018.
- [28] R. Yan, H. Sun, and Y. Qian, "Energy-aware sensor node design with its application in wireless sensor networks," *IEEE Transactions on instrumentation and measurement*, vol. 62, no. 5, pp. 1183–1191, 2013.
- [29] Q. Wang, W. Yan, and Y. Shen, "n-person card game approach for solving set k-cover problem in wireless sensor networks," *IEEE Transactions on Instrumentation and Measurement*, vol. 61, pp. 1522–1535, 2012.
- [30] F. Pianegiani, M. Hu, A. Boni, and D. Petri, "Energy-efficient signal classification in ad hoc wireless sensor networks," *IEEE Transactions on Instrumentation and Measurement*, vol. 57, no. 1, pp. 190–196, 2007.
- [31] J. Gutiérrez and e. a. Villa-Medina, "Automated irrigation system using a wireless sensor network and GPRS module," *IEEE Transactions on instrumentation and measurement*, vol. 63, no. 1, pp. 166–176, 2013.
- [32] L. M. Rodrigues, C. Montez, G. Budke, F. Vasques, and P. Portugal, "Estimating the lifetime of wireless sensor network nodes through the use of embedded analytical battery models," *Journal of Sensor and Actuator Networks*, vol. 6, no. 2, p. 8, 2017.
- [33] M. T. Penella and M. Gasulla, "Runtime extension of low-power wireless sensor nodes using hybrid-storage units," *IEEE Transactions on Instrumentation and Measurement*, vol. 59, no. 4, pp. 857–865, 2010.
- [34] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino, and R. Scarsi, "Discrete-time battery models for system-level low-power design," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 9, no. 5, pp. 630–640, 2001.
- [35] M. Petricca, D. Shin, A. Bocca, A. Macii, E. Macii, and M. Poncino, "An automated framework for generating variable-accuracy battery models from datasheet information," in *International Symposium on Low Power Electronics and Design (ISLPED)*. IEEE, 2013, pp. 365–370.
- [36] Y. Chen, E. Macii, and M. Poncino, "Frequency domain characterization of batteries for the design of energy storage subsystems," in *2016 IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC)*. IEEE, 2016, pp. 1–6.
- [37] NXP Semiconductors, "LPC5411x MCU datasheet," <https://www.nxp.com/docs/en/data-sheet/LPC5411X.pdf>, 2019.
- [38] T. Istomin, A. L. Murphy, G. P. Picco, and U. Raza, "Data prediction+ synchronous transmissions= ultra-low power wireless sensor networks," in *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, 2016, pp. 83–95.
- [39] Texas Instruments, "CC2538 powerful wireless microcontroller system datasheet," <http://www.ti.com/lit/ds/symlink/cc2538.pdf>, 2019.
- [40] J. A. Gutierrez, E. H. Callaway, and R. Barrett, *IEEE 802.15.4 Low-Rate Wireless Personal Area Networks: Enabling Wireless Sensor Networks*. USA: IEEE Standards Office, 2003.
- [41] Mikroelektronika, "HDC1000 Click sensor module information," <https://www.mikroe.com/hdc1000-click>, 2019.
- [42] Texas Instruments, "HDC1000 sensor datasheet," <http://www.ti.com/lit/ds/symlink/hdc1000.pdf>, 2019.
- [43] Mikroelektronika, "Air Quality 4 Click sensor module information," <https://www.mikroe.com/air-quality-4-click>, 2019.
- [44] Sensirion, "SGP30 sensor datasheet," <https://www.sensirion.com/en/environmental-sensors/gas-sensors/sgp30/>, 2019.
- [45] Mikroelektronika, "MIC Click sensor module information," <https://www.mikroe.com/mic-click>, 2019.

- [46] Knowles Acoustics, “SPG0410HR5H-B sensor datasheet,” <https://download.mikroe.com/documents/datasheets/spq0410hr5h-b.pdf>, 2019.
- [47] Mikroelektronika, “Air Quality Click sensor module information,” <https://www.mikroe.com/air-quality-click>, 2019.
- [48] “Mq135 sensor datasheet,” <https://download.mikroe.com/documents/datasheets/mq135.pdf>, 2019.
- [49] Texas Instruments, “BQ27Z561 impedance track battery gauge datasheet,” <http://www.ti.com/lit/ds/symlink/bq27z561.pdf>, 2019.
- [50] RIGOL, “RIGOL programmable DC electronic load,” <https://www.rigol.eu/products/dc-load.html>, 2019.
- [51] —, “RIGOL programmable linear DC power supply,” <https://www.rigolna.com/products/dc-power-loads/dp700/>, 2019.



Yukai Chen (M’15) received the M.Sc. degree in Computer Engineering and the Ph.D. degree in Computer and Control Engineering from Politecnico di Torino, Turin, Italy, in 2014 and 2018, where he is currently a Postdoctoral research fellow. His research interests include computer-aided design in non-functional properties (power, temperature, reliability, and cost) optimization from embedded systems to large-scale cyber-physical systems.



Wenlong Wang (M’20) received the M.Sc. degree in Computer Engineering from the Department of Control and Computer Engineering at Politecnico di Torino, Turin, Italy, in 2018, where he is currently a senior research fellow. His current research interest includes computer-aided design for large-scale cyber-physical energy systems and small-scale IoT devices, with particular emphasis on the power and energy modeling, simulation and optimization.



Daniele Jahier Pagliari (M’15) received the M.Sc. and Ph.D. degrees in Computer Engineering from Politecnico di Torino, Torino, Italy, in 2014 and 2018, respectively. He is currently an Assistant Professor in the same institution. His research interests are in the computer-aided design and optimization of digital systems, with a particular focus on energy-efficiency aspects and on emerging applications, such as machine learning at the edge.



Enrico Macii (SM’02-F’05) is a Full Professor of Computer Engineering with the Politecnico di Torino, Torino, Italy. He holds a Ph.D. degree in Computer Engineering from the Politecnico di Torino. His research interests are design automation for electronic digital circuits and systems, with a particular emphasis on low-power consumption aspects, energy efficiency, sustainable urban mobility, clean and intelligent manufacturing. He is a Fellow of the IEEE.



Massimo Poncino (SM’12-F’18) is a Full Professor of Computer Engineering with the Politecnico di Torino, Torino, Italy. His current research interests include several aspects of design automation of digital systems, with emphasis on the modeling and optimization of energy-efficient systems. He received a Ph.D. in computer engineering and a Dr.Eng. in Electrical Engineering from Politecnico di Torino. He is a Fellow of the IEEE.