

# Profiling and Improving the Duty-Cycling Performance of Linux-based IoT Devices

Immanuel Amirtharaj, Tai Groot, and Behnam Dezfouli\*

Internet of Things Research Lab, Department of Computer Engineering, Santa Clara University, USA  
iamirtharaj@scu.edu, agroot@scu.edu, bdezfouli@scu.edu

**Abstract**—Minimizing the energy consumption of Linux-based devices is an essential step towards their wide deployment in various IoT scenarios. Energy saving methods such as duty-cycling aim to address this constraint by limiting the amount of time the device is powered on. In this work we study and improve the amount of time a Linux-based IoT device is powered on to accomplish its tasks. We analyze the processes of system boot up and shutdown on two platforms, the Raspberry Pi 3 and Raspberry Pi Zero Wireless, and enhance duty-cycling performance by identifying and disabling time-consuming or unnecessary units initialized in the userspace. We also study whether SD card speed and SD card capacity utilization affect boot up duration and energy consumption. In addition, we propose Pallex, a parallel execution framework built on top of the systemd init system to run a user application concurrently with userspace initialization. We validate the performance impact of Pallex when applied to various IoT application scenarios: (i) capturing an image, (ii) capturing and encrypting an image, (iii) capturing and classifying an image using the the k-nearest neighbor algorithm, and (iv) capturing images and sending them to a cloud server. Our results show that system lifetime is increased by 18.3%, 16.8%, 13.9% and 30.2%, for these application scenarios, respectively.

**Index Terms**—Energy efficiency; boot up; shutdown; edge and fog computing; userspace optimization; machine learning.

## I. INTRODUCTION

As more Internet of Things (IoT) devices are deployed every day, the demand for energy efficient and low-power devices is increasing at a fast pace. These devices are increasingly deployed in remote regions or harsh environments where the only reliable source of power is a battery or the power provided by an energy harvesting solution. In addition, even for those devices that are connected to the power grid, it is essential to reduce their energy footprint given the increasing cost of energy and the efforts towards reducing carbon emissions [1], [2]. To this end, various solutions have been proposed, such as energy-efficient hardware design [3]–[5], low-overhead operating systems (OS) [6]–[9], and low-power networking stacks [10]–[13]. The importance of energy-efficient design has also been demonstrated by mining the questions raised by developers in forums [14].

With the increases in processing power and memory capacity for embedded systems, new process-intensive IoT applications are being introduced. Such applications benefit from edge and fog computing to improve responsiveness and minimize the overhead of data exchange with cloud platforms [15]. For example, image processing through artificial intelligence on IoT devices presents several benefits including faster decision-making and lower reaction time to environmental changes, less wireless interference with nearby devices, and improved security due to eliminating the need to upload raw images to a cloud service [16]. According to ABI Research [17], 90% of the data generated by edge devices is being processed locally. In summary, local storage and processing can help satisfy the stringent latency requirements of mission-critical applications [11], reduces network utilization, reduces the processing overhead of resource-constraint devices, enhances security, and enables the system to continue its operation even in the presence of intermittent network connectivity [18].

Among the most-popular OSs for IoT edge devices is Linux. According to the Eclipse IoT survey report [19], 71% of IoT developers rely on this OS. Existing low-cost devices such as the Raspberry Pi 3 (RPi3), Raspberry Pi Zero Wireless (RPiZW), Arduino Yun, and Beaglebone Black support this OS. Despite low-level software and hardware improvements targeted at lowering their energy consumption, widespread use of these devices in IoT contexts requires employing user-level energy conservation techniques. Unfortunately, although there are valuable studies on the power measurement and modeling of this device [20]–[23], less attention has been paid to improving energy efficiency.

One of the most effective approaches towards improving energy efficiency is *duty cycling*. This approach has been widely employed by the wireless sensor network community to achieve a long node lifetime, in some cases up to a few years [11], [12], [24]. In a duty-cycled application, a device powers on, performs its intended task, and then powers off for a specified interval. Without relying on duty cycling, the lifetime of a system using an RPi3 and a 2400mAh battery is around 6 hours, assuming 400mA and 5V for current and voltage consumption.

The major burden of a duty-cycled Linux-based device is its boot up time. Specifically, the system software and hardware must be loaded and initialized before running user applications. Compared to real-time OSs such as

\*Corresponding Author

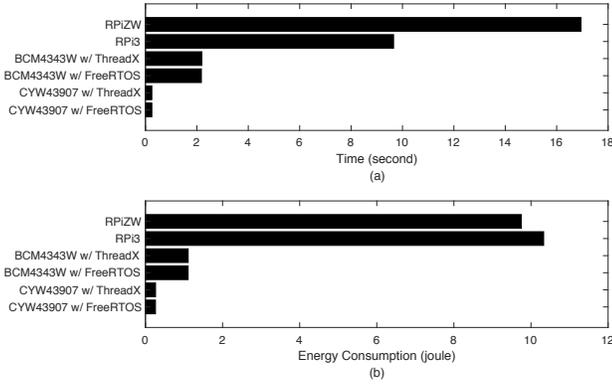


Fig. 1: The (a) boot up duration and (b) energy consumption of the popular IoT operating systems on different hardware platforms.

FreeRTOS [9] and ThreadX [8], the boot up time of Linux is several orders of magnitude longer. To show this, we have measured the boot up time of various hardware platforms when using Linux, FreeRTOS and ThreadX. The results are presented in Figure 1. These results indicate that for Linux-based systems a significant amount of energy is consumed during the boot up process at the beginning of each duty cycle.

There are two types of energy optimization techniques commonly applied to reduce the energy consumption of Linux: *general optimization* [25]–[29] and *application-specific optimization* [30], [31]. General optimization refers to the improvement of OS code to run faster and consume less energy, agnostic to the application type. Sample techniques belonging to this category include improving the filesystem and compression methods of unpacking the kernel during the bootloader phase [25], running boot up scripts in parallel and disabling kernel print statements [26], and saving a copy of the boot image to a file for reuse on subsequent boot ups [27]. Applying general optimization techniques, however, requires a deep understanding of the boot up process and OS. In addition, hardware initialization may pose some challenges. For example, assume that a system image is taken after the hardware devices have been initialized. When this image is reloaded during the next duty cycle, some hardware devices might not be initialized, and therefore the user application cannot function properly. Addressing this limitation requires deep system knowledge and thorough testing to ensure system reliability. Additionally, applying security updates requires image regeneration. Even if the image can be updated, reliability would be a major concern. If images are generated on-the-fly after updates, the system might fail to boot up if the image is corrupted. Enabling fallback images requires an overhead for image verification and uses at least twice the storage space.

Application-specific optimization, on the other hand, refers to either OS improvement or the removal of unnecessary components, depending on application requirements. For example, by customizing the bootloader and thinning

the kernel of unnecessary modules, the authors of [30] decreased the boot up time on an embedded Android device by 65%. In [31], the authors optimized a Linux-based smart television and reduced its boot up time by five seconds. However, since they defined boot up time as the interval between power on and the time instance at which the user can interact with the device, they delayed the initialization of certain components until after the home screen of the television was loaded. Unfortunately, the existing application-specific optimization techniques do not provide simple and universally applicable guidelines for IoT scenarios. When Linux-based devices are used in IoT applications, it is desirable to tailor the system based on application requirements. Even the Linux distributions released for IoT boards are preloaded with unnecessary services and initialize hardware devices not required by specific IoT application scenarios. For example, a device running an image classification algorithm may not require sound utilities, remote login, service discovery daemons, time synchronization, or all the wireless technologies offered by the device.

In this paper we focus on the userspace level and application-specific optimization to improve the performance of duty-cycled Linux systems. In other words, the goal of this paper is to profile and improve the energy consumed by the boot up and shutdown phases of the RPi3 (Raspberry Pi 3, based on the quad-core BCM2837 SoC) and the RPiZW (Raspberry Pi Zero w/ Wireless, based on a single-core BCM2835 SoC) in such a way that minimal work is required to tailor a standardized installation image to specific IoT applications. In particular, the contributions of this paper are as follows: First, we overview the boot up process and present the implementation of a testbed to measure the duration and energy consumption of this process. In order to reveal the effect of loading services on system performance, we profile the start and end of loading units that require more than 10ms to initialize. By categorizing system units into multiple classes, we show that customizing the set of active units, which we refer to as the *unit configuration*, is a very effective approach towards improving duty-cycling performance. For example, unit configuration reduces the energy consumption of the boot up process using a RPi3 by 43.62% for an application that only requires communication with the camera interface. In addition, through profiling the resource utilization of the RPi3 and RPiZW in terms of processing, memory and I/O, we further analyze the overhead of initializing units and show the possibility of running user application processes while userspace initialization is still in progress. Second, in the context of flash memory, even when up to 95% of its capacity is utilized, our results show no effect on the boot up duration or energy consumption. However, using faster flash memory always results in a slightly lower boot up duration (around 1.5%) and energy consumption (around 2.5%). Third, we investigate two shutdown approaches, *graceful* and *forced*, and evaluate the impact of unit configuration on their performance. Our results

TABLE I: Key abbreviations and notations

$E_{btl}$	Energy of the bootloader phase
$E_{knl}$	Energy of the kernel phase
$E_{usi}$	Energy of the userspace initialization phase
$E_{sdn}$	Energy of the shutdown phase
EU	Essential Units
GPIO	General Purpose Input/Output
NRS	Networking Related Services
$P_{btl}$	Bootloader phase
$P_{knl}$	Kernel phase
$P_{usi}$	Userspace initialization phase
RPi	Raspberry Pi 3 or Raspberry Pi Zero w/ Wireless
RPi3	Raspberry Pi 3
RSL	Raspbian Stretch Lite
RPiZW	Raspberry Pi Zero w/ Wireless
SDC	Secure Digital Card
SoC	System on a Chip
$T_{btl}$	Duration of the bootloader phase
$T_{knl}$	Duration of the kernel phase
$T_{usi}$	Duration of the userspace initialization phase
$T_{sdn}$	Duration of the shutdown phase
UART	Universal Asynchronous Receiver Transmitter

confirm that using unit configuration reduces the energy consumption of the graceful and forced shutdown by up to 43.87% and 57.42%, respectively. Additionally, the benefits and risks of both categories are highlighted when used for duty-cycled systems. Fourth, we propose *Pallex*, a parallel execution framework to execute a user application while userspace initialization is still in progress. Using Pallex, a user application is split into several stages that execute at different points of the userspace initialization phase. Our evaluations considering different user application scenarios show that in terms of lifetime, Pallex improves the duty-cycling performance of the RPi3 and RPiZW by 30.2% and 9%, respectively. Although the power consumption of the RPi3 (quad-core) is higher than that of the RPiZW (single-core), the RPi3 achieves a longer lifetime due to its significantly shorter processing duration.

We have chosen to conduct our research using Raspbian Stretch Lite (RSL) on both the RPi3 and RPiZW because around 43% of Linux-based IoT systems rely on Raspbian [19]. RSL is a popular release packaged without a desktop environment, is advertised as a minimalist distribution, has a long development cycle, and is officially supported by the Raspberry Pi Foundation, thereby making it an ideal candidate for deploying energy-efficient IoT applications. However, it is worth noting that the results of this paper translate over to other Linux distributions supporting `systemd` such as Debian, Arch Linux, and Kali.

The rest of this paper is organized as follows: In Section II, we provide an in-depth description of the Linux boot up process. We present the components of our testbed and experimentation methodology in Section III. In Section IV, we profile the Linux boot up process and measure the effect of unit configuration and flash memory on the duration and energy consumption of boot up. The Linux shutdown phase is studied in Section V. The operation and performance evaluation of Pallex is presented Section VI. Section VIII overviews the related work. We conclude our findings in Section IX.

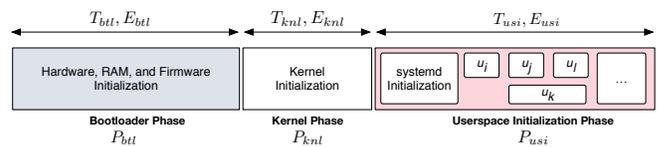


Fig. 2: The sequence of operations during the boot up process.  $u_i, u_j, u_k$  and  $u_l$  refer to userspace units.

## II. BOOT UP PROCESS

As demonstrated in Figure 2, the Linux boot up process consists of three main phases: (i) the bootloader phase ( $P_{btl}$ ), (ii) the kernel phase ( $P_{knl}$ ), and (iii) the userspace initialization phase ( $P_{usi}$ ). Each of these phases present unique opportunities to optimize boot time. In this paper we focus on  $P_{usi}$  because it enables the user to simply and effectively implement duty cycling to conserve energy.

### A. Bootloader Phase

In the Pi's boot up process, a two-stage bootloader<sup>1</sup> prepares the hardware to load the kernel. First, while the ARM processor is off, the GPU is powered up and initializes itself by executing a first stage bootloader that is burned into the SoC's ROM [32]. This stage instructs the GPU to power on the Secure Digital Card (SDC) and read a file called `bootcode.bin` from the first partition of the SDC. The execution of `bootcode.bin` enables the on-board SDRAM and loads `start.elf`, which contains firmware for the GPU. After reading in system configuration parameters from `config.txt`, the GPU loads the kernel image (`kernel.img`) along with kernel parameters (`cmdline.txt`) into the shared RAM allocated to the ARM processor. Lastly, the second stage powers on the ARM processor by triggering the reset signal [33]. Now the system is running the Linux kernel. In this paper we refer to the duration and energy consumption of this phase as  $T_{btl}$  and  $E_{btl}$ , respectively. Also, we refer to the ARM processor as "processor".

### B. Kernel Phase

The Linux kernel handles all OS-related processes such as memory management, process scheduling, driver initialization, and overall system control. The kernel is initialized in two steps: The first step occurs when the kernel is loaded into memory and decompressed. Basic memory management is enabled during this stage as well. The next major step the kernel takes is launching an `init` process to run and transition to  $P_{usi}$ . In this paper we refer to the duration and energy consumption of the kernel phase as  $T_{knl}$  and  $E_{knl}$ , respectively.

### C. Userspace Initialization Phase

The final phase of the boot up process is userspace initialization. In this paper we refer to the duration and

<sup>1</sup>Until October 2012, the RPi platform used a three stage bootloader, with an additional file, `loader.bin`, executed by the GPU between the `bootcode.bin` and `start.elf` stages.

energy consumption of this phase as  $T_{usi}$  and  $E_{usi}$ , respectively. During this phase, the units that are activated (a.k.a., initialized) in the userspace as well as the daemons that run in the background during active mode are activated by an `init` process. Most modern Linux distributions (including RSL) use `systemd` [34]–[36] as their initialization system. The predecessor to `systemd` was System V `init` (`sysvinit`), which traces its origins back to the original commercial Unix system. Compared to `sysvinit`, `systemd` offers advantages such as calendar-based job timers, a more unified API, and backward compatibility with `sysvinit`.

`systemd` is the first daemon to start during boot up and the last to exit during shutdown. In addition to operating processes and services, `systemd` is capable of triggering filesystem mounts, monitoring network sockets and running timers. Each of these capabilities is described by a set of configurations files, termed `unit` files. Unit types include service units, which manage background services; mount units, to mount filesystems; and target units: to group and control other units. There are other unit types as well, but the details of their implementation reside outside the scope of this paper. In order to manage the dependencies and ordering of unit activation, unit files employ the syntax provided by `systemd`. When `systemd` is initialized, it first loads the unit configurations and determines the boot up goal. Based on the specified hierarchy of dependencies, `systemd` activates the units in order to reach the target goal. On RSL, the `rc.local` file is loaded and executed by `systemd` after all the services have been initialized. Therefore, the Raspberry Pi Foundation recommends general consumers launch user processes through this file to ensure all necessary hardware and software components are initialized. However, performance improvements are easily achieved through manually resolving dependencies and writing custom unit files, as we will demonstrate in Section IV.

One important feature `systemd` provides is the ability to activate units in parallel. This feature can save time compared to initializing units sequentially, even on a single-core or single-threaded board, as some units require time for hardware initialization [37]. Finally, `systemd` enables developers to customize rules for automatically starting, reloading, and killing services.

### III. TESTBED OVERVIEW

Figure 3 shows the architecture of the testbed. It consists primarily of two hardware components: a *master* and a *minion*. The master is composed of a RPi3 and uses a shield board [38] for energy measurement. The minion refers to the device under test. We used two different minion boards in this paper, a RPi3 and a RPiZW, where both run the March 2018 release of RSL.

The master runs two programs: (i) a *control program* that is responsible for enabling and disabling the input power to the minion, and (ii) an *energy measurement program* that controls the EMPIOT shield to measure the energy consumption of the minion. Algorithm 1 shows

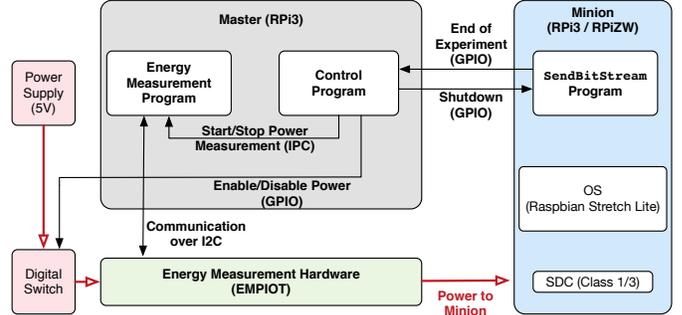


Fig. 3: The schematic of the testbed used for measuring the duration and energy consumption of the boot up phases.

the pseudo-code of the control program running on the master device. This program is used to control power supply to the minion and measure the duration and energy consumption of the boot up process by sending commands to the energy measurement program. First, the BCM [39] and WiringPi [40] libraries are initialized. Then, the energy measurement program is initialized to listen on a socket and receive commands from the control program using inter-process communication (IPC). At the beginning of each experiment, the control program communicates with the digital switch through GPIO pins controlled by the BCM library to turn on the minion. At the same time, the control program communicates with the energy measurement program through the socket to start power measurement. After the start of the boot up process, the control program uses the WiringPi library to detect and decode a message sent by the minion when the experiment reaches the desired completion point. The end of the operation depends on the experimentation scenario and refers to the cases such as the end of  $P_{usi}$  or the completion of a user application. The control program logs the duration and energy consumption of this operation for each experiment. Figure 4 shows the states of the minion during one single experiment from the instance the minion is powered on until the instance the minion sends the signal to the master after it has completed the desired set of operations.

It should be noted that the minion cannot use Ethernet, WiFi, or Bluetooth to inform the master about the end of its operation because their relevant units might not be activated in some scenarios, as we will see later in this paper. Furthermore, it is not possible to use any of these mechanisms when the performance of  $P_{btl}$  and  $P_{knl}$  is being measured. To address this challenge, we send a message over a GPIO pin from the minion to the master. When the minion completes its operation, it runs a program that generates a simple bit pattern to notify the master. The minion generates a bit pattern, instead of a simple rising or falling edge signal, because of the GPIO voltage variations during the boot up process. Therefore, the generated bit pattern avoids the master from reporting false positives. This approach enables us to measure the performance of  $P_{btl}$  and  $P_{knl}$  because `systemd` is initialized immediately after the kernel phase,

and GPIOs can be used as soon as `systemd` is initialized. The pattern generation program, named `SendBitStream`, is a unit that is activated by `systemd` once the targeted state (based on the experiment type) has been reached. Specifically, to measure  $T_{btl} + T_{knl}$  (and  $E_{btl} + E_{knl}$ ), this unit is called immediately after `systemd` initialization. This is achieved by creating a new unit without enforcing any dependencies. Similarly, to measure  $T_{btl} + T_{knl} + T_{usi}$ , this unit is called when all the required units have been activated. In order to measure boot up duration and energy until a particular unit has been activated, this unit is activated when its dependencies have been resolved. We used `systemd-analyze` to extract  $T_{knl}$ . By using this value, we can compute  $T_{btl}$  as well. A similar approach is used to measure the energy consumption of these phases.

The existing COTS energy measurement tools are either costly or do not offer the necessary features to build a fully controllable energy measurement testbed. For example, Keithley 7510 [41], includes only 2MB of storage and costs more than \$3500. The Monsoon tool, which has been widely used by the academia, interfaces with the board under test using a USB connection. Therefore, accurately controlling the start and stop of the measurement interval is not possible. On the other hand, the tools proposed by the research community reveal the following shortcomings: (i) complexity and the need to modify the board under test (e.g., [42]–[47]), (ii) offering a limited measurement range (e.g., [42], [48]–[52]), and (iii) low accuracy (e.g., [53], [54]). Due to these concerns, to build the testbed required for the experiments of this paper, we used the EMPIOT tool proposed in [38].

EMPIOT is a shield board that is installed on top of an RPi3 and communicates with the energy measurement program using I2C, as Figure 3 shows. For the experiments of this paper, we have configured the board to measure current and voltage values up to 1A and 5V, respectively. For current measurement, the board utilizes the INA219 chip to measure the voltage across a  $0.1\Omega$  shunt resistor. This voltage is then converted to current using Ohm’s law. The voltage is directly measured using the internal ADC of INA219. The EMPIOT tool is capable of supersampling approximately 500,000 samples per second (sps), which are then averaged and streamed at 1Ksps. The current and voltage resolution of this platform are  $100\mu\text{A}$  and  $4\text{mV}$ , respectively, when the 12-bit resolution mode is configured. Once a new sample is ready, the shield sets the conversion ready bit, which is read by the energy measurement program running on the master. The sample is then read off the shield board using I2C communication. The energy measurement program uses two buffers. Once a buffer is filled, an execution thread is activated to write the collected samples to a file, while the second buffer is being filled up. This program also converts the raw power samples collected during the measurement interval to energy using Riemann integration.

Figure 5 shows the actual testbed used. In addition to the master node and the two minion boards, this testbed includes a gateway, which is a RPi3 board. Since for most

---

### Algorithm 1: Master’s control program

---

```

1 function startExperiments()
2   /*to control the power switch through a GPIO pin and
3     receive the bit pattern from the minion on another
4     GPIO pin */
5   setup the BCM and WiringPi GPIO libraries;
6   /*setup the energy measurement program and make it
7     ready to measure power */
8   initialize the energy measurement energy;
9   for i = 0; i < 50; i++ do
10    cut power to the minion;
11    apply power to the minion and start time and energy
12    measurement;
13    wait for minion to transmit bit pattern;
14    record duration and energy consumption;

```

---

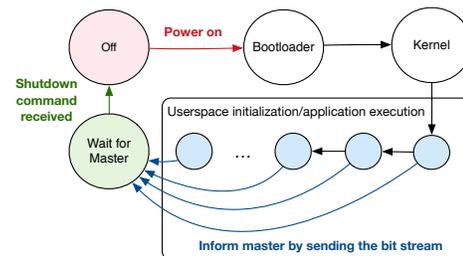


Fig. 4: The state machine of the minion device during an experiment. The minion executes the `SendBitStream` program to send a message to the master after reaching the level of completion.

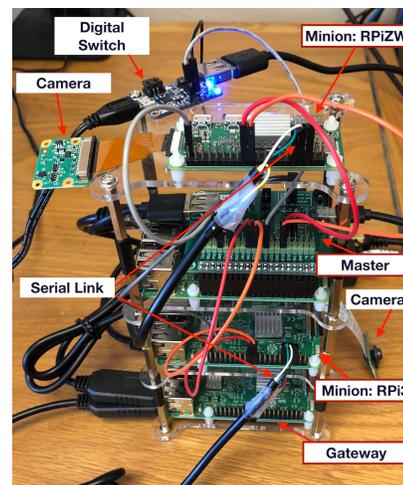


Fig. 5: The hardware components of the testbed.

of the scenarios, both the wired and wireless communication interfaces are disabled, we need to use the serial port to communicate with the minions and configure the tests. To this end, the gateway, which is always connected to our wired network, enables us to use UART and communicate with the minions. For the rest of this paper we refer to a minion board simply as "board", which is the device under test.

Unless otherwise mentioned, results are averaged over 50 experiments, with the error bars representing the 95% confidence interval of the mean.

#### IV. PROFILING AND ENHANCEMENT OF THE LINUX BOOT UP PROCESS

In this section, we first study the activation time of units during  $P_{usi}$ . We then profile system resource utilization in terms of processing, memory and I/O. In addition, we evaluate the effect of SDC speed and capacity usage on the duration and energy consumption of the three phases of the boot up process. Finally, we show how customizing userspace units can be employed to improve boot up duration and energy consumption.

##### A. A Deeper Look into the Userspace Initialization Phase

The userspace initialization phase activates a variety of units supporting different functionalities of the system. A summary of these units is available in Appendix A. In the beginning of this phase, `systemd` is initialized and it loads unit configuration files to determine which units must be activated. It then creates a dependency tree to determine the ordering of unit dependency resolution. In order to improve efficiency, units can be initialized in parallel with respect to their dependency relationships.

Figures 6(a) and (b) show the start and activation duration of units during  $P_{usi}$  for the RPi3 and RPiZW, respectively. We have used the `systemd-analyze blame` utility to extract these data. Although all of the units are enabled for these experiments, we disabled WiFi and Ethernet connectivity to extract the activation duration of units without them being affected by external factors such as communication with a WiFi access point. Also, to focus on the units that significantly contribute to  $T_{usi}$ , these figures do not display units that require less than 10ms to be activated. Please note that the x-axis of both Figures 6(a) and (b) start at  $t = 6s$  because  $T_{btll} = 3.65s$  and  $T_{knl} = 2.85s$  for both boards. It must also be noted that, since  $energy = power \times duration$ , the execution duration of each unit does not necessarily reflect its energy consumption. We will study the energy consumption of unit activation in the subsequent sections. Additionally, we will show that a long unit activation duration does not prevent us from running a user application concurrently with the unit activation.

For the RPi3, our results show that the system units imposing the highest overhead are `dev-mmcblk0p2.device`, `networking.service`, `hciuart.service`, and `systemd-resolved.service`. Except the first unit, which is responsible for bringing the root partition into the scope of `systemd`, the rest are networking-related services. Networking services generally require a longer activation duration compared to other units because they require a combination of initializing hardware and networking utilities. The `hciuart.service` is responsible for the initialization of Host Controller Interface (HCI) to provide a uniform interface for accessing Bluetooth hardware capabilities. All USB-Bluetooth adapters operate with a HCI interface over the USB link. During its initialization, HCI creates read and write communication threads, establishes a connection to the Bluetooth transceiver, and reads device buffer sizes.

Enabling real-time Bluetooth communication requires the processor to frequently context-switch to monitor UART communications. Because `bluetooth.service` depends on `hciuart.service`, it finishes initialization after `hciuart.service`. An interesting behavior of these services is that they both complete their activation after `rc.local`. This means that user applications that rely on Bluetooth cannot be started using `rc.local`. Instead, these applications require a `systemd` unit that is scheduled to be activated after the completion of `bluetooth.service`.

On the RPiZW, we observe that  $T_{usi}$  is about 17s, which is 10s longer than that of the RPi3. The same units mentioned for the RPi3 consume most of the userspace initialization time on the RPiZW as well. Since the RPiZW's processor only has one core, compared to the RPi3's quad-core processor, the difference in duration is expected. A multi-core processor can parallelize unit activations across multiple cores, while a single-core processor must context-switch more frequently between tasks. This also explains why different sets of units are presented in Figure 6(a) and (b). Furthermore, other units including `systemd-login.service`, `console-setup.service`, `alsa-restore.service`, and `systemd-user-sessions.service` reportedly require a longer activation duration (at least 1s) on the RPiZW.

**Threats to Validity.** Our analysis of unit activation duration revealed one significant shortcoming about `systemd-analyze blame`. The initialization duration calculated by this utility is not completely accurate if the dependency tree is not precisely configured. For example, if `systemd` attempts to activate a fast unit that depends on a longer unit which has not been activated yet, the faster unit cannot complete its activation until the dependency has been resolved. More specifically, if the activation duration of longer and shorter units are 800ms and 10ms, respectively, then `systemd-analyze blame` reports 810ms as the initialization duration of the faster unit. For example, for `sudo.service`, `systemd-analyze blame` falsely reports a long initialization duration because `systemd` attempts to activate it before the filesystem the service depends on is mounted. Therefore, the values generated by this utility may be longer than the actual activation duration of each unit, but these values are not shorter. To cover all the units that contribute significantly to the boot up process, we have included only those units that require more than 10ms according to `systemd-analyze blame`. In Figure 6, we include Type I errors to avoid excluding any units that require more than 10ms. Later, to eliminate Type I errors, we identify and study the impact of units that significantly affect the boot up process.

##### B. Customizing Userspace Initialization

In this section, we analyze the effect of disabling optional units on the performance of the userspace initialization phase. This is referred to as *unit configuration* in this paper. The `systemctl` [55] utility is used to implement unit configuration.

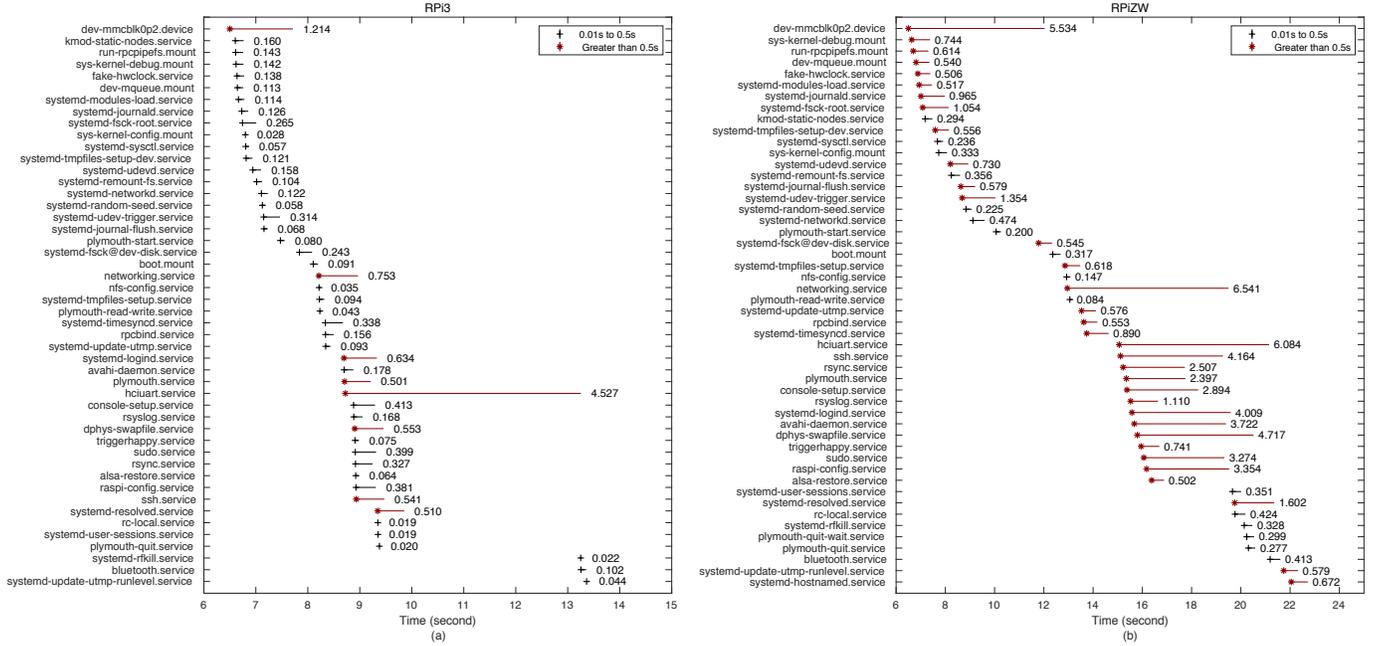


Fig. 6: The starting time and initialization duration of units during  $P_{usi}$  for (a) RPi3 and (b) RPiZW. Only the units with activation duration longer than 10ms are included in this figure. Results are averaged over 50 experiments.

Among the units activated during  $P_{usi}$ , some of them are essential to maintain stable system operation. For example, units that are responsible for mounting the file systems and loading kernel modules cannot be safely disabled. Another example, nearly all user applications will work even if `systemd-random-seed.service` is disabled. However, disabling this service is a security risk, because it is critical for maintaining higher entropy for the secure generation of random numbers used in encryption algorithms. A complete list of these units, which are referred to as *Essential Units* (EU) in this paper, can be found in Appendix A-A. The next category includes a significant number of services and is referred to as *Networking-Related Services* (NRS) in this paper. Appendix A-B overviews these services. These services are by far the most variable in terms of activation duration because they often rely on external dependencies (e.g., association with an access point, communicating with a DHCP server, etc.) and/or initializing physical hardware such as the WiFi and Bluetooth transceivers. Due to the significant effect of NRS on boot up performance, we study the effect of following unit configurations on  $P_{usi}$ :

- *EU*. Refers to the case where only the essential units (cf. Appendix A-A) are enabled. For EU configuration, we detect the end of userspace initialization when `rc.local` runs because it is the last unit file that is executed.
- *MMS*. Refers to `dphys-swapfile.service`. For configuration EU w/ MMS, we detect the end of userspace initialization when `rc.local` runs because it is the last unit file that is executed.
- *NET1*. Refers to `networking.service`. For configuration EU w/ NET1, we detect the end of userspace initialization when `rc.local` runs because it is the last unit file that is executed.

- *NET2*. Refers to `networking.service` and `sshd.service`. For configuration EU w/ NET2, we detect the end of userspace initialization when `rc.local` runs because it is the last unit file that is executed.
- *NET3*. Refers to `bluetoothd.service` and `hciuart.service`. For configuration EU w/ NET3, since `bluetooth.service` is the last service that is executed, we detect the end of userspace initialization when this service completes its initialization.
- *ALLU*. Refers to the case where all units are enabled. For this configuration, since `bluetooth.service` is the last service that completes, we detect the end of userspace initialization when this service is activated. For configuration ALLU w/o NET3, we detect the end of userspace initialization when `rc.local` runs because it is the last unit file that is executed.

In addition to the classifications detailed above, since the actual association of a RPi with an access point introduces more variations due to the control messages exchanged between the two parties, we report the results separately for the cases where a WiFi connection is established. Figures 7(a) and (b) show the impact of unit configuration on both the RPi3 and RPiZW, respectively, in terms of  $T_{usi}$  and  $E_{usi}$ .

These figures clearly show the benefits of unit configuration to enhance the performance of boot up phase. For example, applying unit configuration EU reduces the energy consumption by 43.62% compared to ALLU, for the RPi3 board. These results also reveal that WiFi significantly affects  $T_{usi}$  and  $E_{usi}$ . Specifically, for configuration EU w/NET1, enabling WiFi increases  $T_{usi}$  by around 2s and 4s, for the RPi3 and RPiZW, respectively. Besides, for this configuration, enabling WiFi increases  $E_{usi}$  by 3 and 3.5 joules for the RPi3 and RPiZW, respectively.

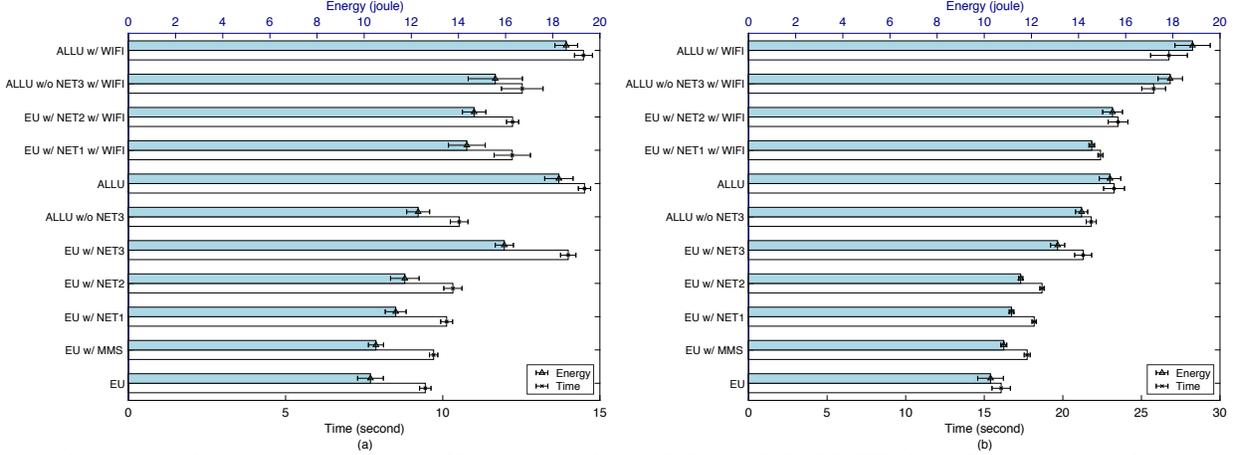


Fig. 7: The duration and energy consumption of booting up the (a) RPi3 and (b) RPiZW for different unit configurations. The lower and upper axes show  $T_{btl} + T_{knl} + T_{usi}$  and  $E_{btl} + E_{knl} + E_{usi}$ , respectively.

These results also reveal the higher effect of `hci-uart.service` and `bluetooth.service` on the RPi3. Since activating these services do not significantly utilize the processing resources of the RPi3 (as we will show in Section IV-C), the waste of processing resources, and hence energy consumption, is higher than that of the RPiZW. Therefore, EU w/ NET3 compared to EU results in a 54.8% energy increase on RPi3, compared to the 26.92% increase on RPiZW. In addition, we can observe that on the RPi3, which can initialize the WiFi and NET3 services concurrently, the duration and energy consumption of the "ALLU" and "ALLU w/ WiFi" configurations are almost equal. In contrast, for the RPiZW, the duration and energy consumption of the "ALLU w/ WiFi" configuration are higher than "ALLU" because the single-core processor needs to interleave the tasks of initializing WiFi and NET3 services. We will further study this behaviour in the next section.

It must be noted that *disabling* units does not necessarily prevent their activation by `systemd`. More specifically, units might be activated when: (i) other units relying on them are activated, or (ii) in the case of external event hooks such as attaching a device. For example, `alsa-restore.service` is automatically activated even if it has been disabled. However, as its main function is to initialize the onboard sound card, it is not required by most IoT applications. Therefore, it is worth *masking* this service using `systemctl`. This is performed by pointing the unit file to the special device `/dev/null` so that the dependency tree can mark it as resolved for dependants without actually running it.

The exact increase of duration and energy when using the networking services depends on factors such as interference, channel congestion, and the load of the access point during the association process. For example, the duration of WPA authentication and IP allocation increases as the current load of the access point is intensified. Since these external dependencies are outside the scope of the RPi performance, unnecessary services must be disabled or careful attention must be paid to link quality and access

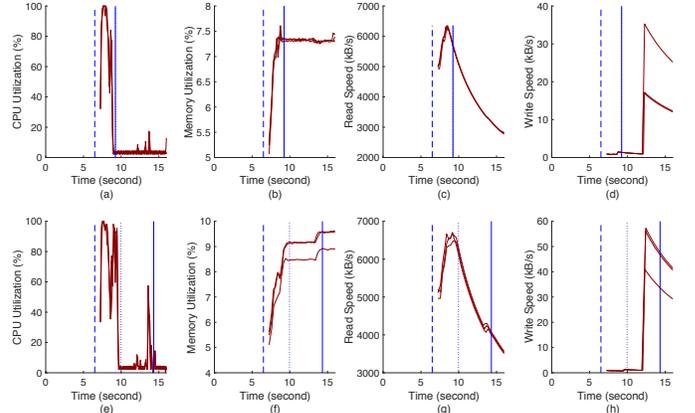


Fig. 8: The resource utilization of the RPi3 during userspace initialization for three different experiments. Sub-figures (a)-(d) represent unit configuration EU, and sub-figures (e)-(h) represent unit configuration ALLU. Dashed lines indicate the start of  $P_{usi}$ , dotted lines indicate the instance `rc.local` is activated, and solid lines indicate the end of  $P_{usi}$ .

point load to achieve a desirable performance.

### C. Resource Utilization During Userspace Initialization

In this section, we study system resource utilization during  $P_{usi}$ . To this end, we measured processor utilization, memory utilization, and SDC's I/O speed on the RPi3 and RPiZW. Resource monitoring is performed by a shell program that starts as soon as `systemd` is initialized. For this reason we noticed that resource monitoring is not available for 0.75s and 1.2s after the completion of  $P_{knl}$  on the RPi3 and RPiZW, respectively. To record processor utilization, we wrote a `gawk` [56] script to read values from `/proc/stat` and calculate the current percentage of processor utilization across all cores with high granularity and low overhead. We read directly from `/proc/meminfo` to determine memory utilization, and used the `iostat` [57] utility for collecting SDC's I/O utilization. Figures 8 and 9 show the results for three trials.

Comparing the two figures indicates the significantly higher processor utilization of the RPiZW compared to the

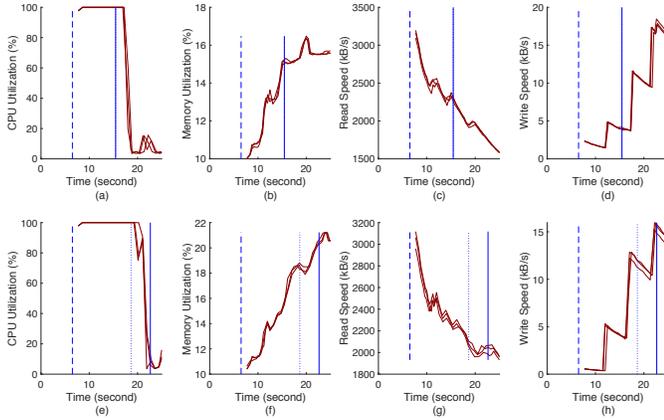


Fig. 9: The resource utilization of the RPiZW during userspace initialization for 3 different experiments. Sub-figures (a)-(d) represent unit configuration EU, and sub-figures (e)-(h) represent unit configuration ALLU. Dashed lines indicate the start of  $P_{usi}$ , dotted lines indicate the instance `rc.local` is activated, and solid lines indicate the end of  $P_{usi}$ .

RPi3. For the RPi3, we notice that processor utilization drops to less than 5% as soon as `rc.local` is invoked, regardless of the unit configuration applied (compare Figure 8(a) and (e)). For the RPiZW, however, when the unit configuration ALLU is applied, processor utilization drops to around 5% after the end of userspace initialization, which is the completion of NET3 services (compare Figures 9(a) and (e)). We justify this behavior by referring back to Figure 6. While the RPi3 is almost finished with userspace initialization (except the NET3 services) at the time `rc.local` is activated, RPiZW needs to complete the activation of multiple units. Specifically, as Figure 6 shows, a considerable number of units with loading time longer than 10ms are being activated around  $t = 20$ .

For RPi3, enabling all services increases memory utilization from around 7% to 10%. For RPiZW, the increase is from around 16% to 21%. These results indicate that even for the unit configuration ALLU, more than 900MB and 400MB of RAM is available on the RPi3 and RPiZW, respectively. In terms of I/O, since the RPi3 initializes more units in parallel, its I/O speed is almost double than that of the RPiZW. Although the processor utilization of the RPiZW is around 100% throughout  $P_{usi}$ , we can still benefit from concurrent execution of user applications within this phase if they mostly rely on peripheral initialization and I/O operations. A sample IoT application that satisfies this requirement is capturing a photo using a camera. For example, the RPi camera module [58] connects over a Camera Serial Interface (CSI) to the GPU on the RPi. While a picture is being captured and processed, the processor can switch to other tasks, as there are several steps performed by the camera that are independent of the RPi’s processor. In particular, the camera has a physical requirement of exposure time. Next, the camera needs to process the image. According to the manufacturer of the camera module’s image sensor (IMX219PQ [59]), the camera has Lens Shading Correction (LSC) functionality,

which means the image undergoes some processing on the camera module before it is sent over CSI. After this step, the image must be sent over the CSI interface as a series of Bayer frames to the RPi’s GPU. Next, the GPU’s VideoCore firmware assembles the image. Finally, the processor can receive the assembled image from the GPU. In Section VI we will show the performance improvement of this application running during  $P_{usi}$ .

#### D. Profiling the Time and Energy Consumption of Bootloader and Kernel Phases

In this section, we study the duration and energy consumption of bootloader, kernel, and userspace initialization phases versus the properties of the SDC used. In order to measure the effect of SDC speed on performance, we used two 32GB Sandisk SDCs: (i) a UHS (Ultra High Speed) class 1, and (ii) a UHS class 3. Note that UHS 1 and UHS 3 refer to minimum write speed of 10 MB/sec and 30 MB/sec, respectively. In addition to SDC speed, we are interested in measuring the effect of SDC capacity utilization on boot up performance. Modern SDCs are implemented with the NAND technology. Compared to the NOR technology, NAND offers lower power consumption, lower cost per bit, higher density and faster write speed. However, as the disk fills up, the write performance starts to degrade. In order to measure the effect of SDC utilization on boot up performance, we fill 5%, 50%, and 95% of the SDC capacity. To fill the SDC with data, we used the `dd` [60] utility which allows us to write blocks of memory to the SDC.

Tables II and III demonstrate the results averaged over 50 trials for each configuration. The following unit configurations are used for these measurements: (i) EU, and (ii) ALLU w/o NET3 w/o WiFi). We excluded NET3 due to the high variations caused by these services. In terms of SDC capacity usage, these results show no effect on boot up performance. However, the use of a faster SDC results in a slight reduction in energy consumption. For example, for the ALLU w/o NET3 configuration on the RPi3, on average the faster SDC reduces energy by 2.5%, compared to the slow SDC.

According to these results, regardless of the SDC type and capacity usage,  $T_{btl}$  and  $T_{knl}$  are similar for the two boards. However, the energy consumption of these phases is higher on the RPi3 than on the RPiZW. Compared to  $P_{usi}$ , which is process-intensive,  $P_{btl}$  and  $P_{knl}$  do not benefit from the higher processing power of RPi3’s SoC because their main operation is to load the kernel and initialize hardware components. These are tasks that are mostly synchronous and not easily threaded across cores. Therefore, since the RPi3 has a more complex and powerful SoC, more resources are wasted on this board during the first two phases. In contrast,  $E_{usi}$  of the RPi3 is actually lower than that of the RPiZW. During  $P_{usi}$ , the RPi3 runs a larger number of processes in parallel, and therefore requires less energy to reach the boot up target by reducing the total amount of time spent in this phase.

TABLE II: The duration and energy consumption of the bootloader phase, kernel phase, and userspace initialization phase for RPi3. The left and right values in each cell show duration (second) and energy consumption (joule).

		EU			ALLU w/o NET3 w/o WiFi		
		Bootloader+Kernel ( $T_{btl} + T_{knl}, E_{btl} + E_{knl}$ )	Userspace ( $T_{usi}, E_{usi}$ )	Total	Bootloader+Kernel ( $T_{btl} + T_{knl}, E_{btl} + E_{knl}$ )	Userspace ( $T_{usi}, E_{usi}$ )	Total
		Slow SDC	5%	6.5, 4.51	2.94, 5.77	9.44, 10.28	6.5, 4.51
50%	6.5, 4.51		3.23, 5.95	9.73, 10.46	6.5, 4.51	3.72, 7.32	10.22, 11.74
95%	6.5, 4.51		3.23, 5.89	9.73, 10.4	6.5, 4.51	3.72, 7.28	10.22, 11.79
AVG	6.5, 4.51		3.13, 5.87	9.63, 10.38	6.5, 4.51	3.83, 8.91	10.33, 11.92
Fast SDC	5%	6.5, 4.51	2.94, 5.6	9.44, 10.11	6.5, 4.51	3.74, 7.13	10.24, 11.64
	50%	6.5, 4.51	2.97, 5.68	9.46, 10.19	6.5, 4.51	3.66, 7.07	10.16, 11.58
	95%	6.5, 4.51	2.96, 5.67	9.46, 10.18	6.5, 4.51	3.71, 7.2	10.21, 11.71
	AVG	6.5, 4.51	2.95, 5.65	9.45, 10.16	6.5, 4.51	3.7, 7.133	10.2, 11.64

TABLE III: The duration and energy consumption of the bootloader phase, kernel phase, and userspace initialization phase for RPiZW. The left and right values in each cell show duration (second) and energy consumption (joule).

		EU			ALLU w/o NET3 w/o WiFi		
		Bootloader+Kernel ( $T_{btl} + T_{knl}, E_{btl} + E_{knl}$ )	Userspace ( $T_{usi}, E_{usi}$ )	Total	Bootloader+Kernel ( $T_{btl} + T_{knl}, E_{btl} + E_{knl}$ )	Userspace ( $T_{usi}, E_{usi}$ )	Total
		Slow SDC	5%	6.5, 3.26	9.73, 7.02	16.23, 10.28	6.5, 3.26
50%	6.5, 3.26		9.43, 6.34	15.93, 9.6	6.5, 3.26	15.4, 10.97	21.9, 14.23
95%	6.5, 3.26		9.8, 6.57	16.3, 9.83	6.5, 3.26	15.47, 10.94	21.97, 14.2
AVG	6.5, 3.26		9.65, 6.64	16.15, 9.9	6.5, 3.26	15.45, 10.94	21.92, 14.2
Fast SDC	5%	6.5, 3.26	9.47, 6.56	15.97, 9.82	6.5, 3.26	15.3, 10.86	21.8, 14.12
	50%	6.5, 3.26	9.41, 6.31	15.91, 9.57	6.5, 3.26	15.39, 10.92	21.89, 14.18
	95%	6.5, 3.26	9.47, 6.35	15.97, 9.61	6.5, 3.26	15.34, 10.92	21.84, 14.18
	AVG	6.5, 3.26	9.45, 6.41	15.95, 9.67	6.5, 3.26	15.34, 10.9	21.84, 14.16

## V. PROFILING AND ENHANCEMENT OF THE LINUX SHUTDOWN PHASE

Throughout the scope of this paper, special attention is paid to system boot up phases rather than the shutdown phase. This is because the time and energy required for boot up are higher than that required for the shutdown, and the gains are therefore more significant. However, there are also several ways through which shutdown time and energy consumption can be reduced. The naive approach is to cut the power to the RPi as soon as the user application is completed. Unfortunately, cutting the power improperly might result in corrupting data blocks on SDC. If these blocks also happen to coincide with the sectors necessary for boot up or important files in the `rootfs`, the device could be rendered unrecoverable. In order to address this problem, all of the SDC's partitions can be mounted as read-only to guarantee there would be no file operations when a power loss event occurs. However, a read-only solution is complicated and not feasible when the user application needs to store or analyze large quantities of dynamic data such as running a machine learning algorithm. Furthermore, since the filesystems are by default read-only, updating the device becomes a lengthy and challenging process, requiring virtual root filesystems mounted to RAM disks and multiple remount operations on the SDC. In this case, if large amounts of data must be stored for processing or before transmission, external storage is required. This solution, however, introduces extra power consumption and might cancel out any gains achieved by a faster shutdown. Additionally, if no external device is used, workarounds must be implemented for system logging and other system functionalities, which rely on a writeable

filesystem. This analysis is outside the scope of this paper.

The next approach is to make only the boot partition write-protected. However, this solution does not necessarily prevent data corruption, because flash partitions are not actually separated as they would be on a traditional hard drive. SDCs use a Flash Transition Layer (FTL) to map virtual file blocks to their actual location in the storage [61]. Many SDCs are preloaded with a wear-leveling firmware which uses the FTL to re-arrange data blocks, often mixing across partition lines (transparently to the RPi) to enhance block device lifetime. In this case, data corruption can occur if the power is cut while read-only data is being migrated during the wear-leveling operation. Therefore, mixing read-only and writeable partitions does not guarantee protection against improper shutdowns. The only way to guarantee it is safe to cut the device power is to ensure that no operations are being performed on the SDC.

In the rest of this section, we present and evaluate two suitable approaches, *graceful shutdown* and *forced shutdown*, to power off a duty-cycled IoT system. Furthermore, we assess the tradeoffs between system reliability and energy consumption. Please note that  $T_{sdn}$  and  $E_{sdn}$  refer to the duration and energy consumption of the shutdown phase.

### A. Graceful Shutdown

During a graceful shutdown, `systemd` sends a shutdown signal to all of the running processes. After these processes exit and the network interfaces are brought down, the filesystems are unmounted and power to the device is safely cut. The amount of time required to unmount the

filesystems is almost fixed and beyond the control of the user. However, the lower the number of running processes which must return an exit code before `shutdown.target` is reached, the faster `systemd` can finalize the shutdown; therefore, removing extraneous units expedites the shutdown phase. This behavior is best exemplified in Figure 10. These figures represent the time and energy consumption of the RPi3 and RPiZW for various unit configurations.

When WiFi is disconnected, using unit configuration EU reduces energy consumption by 43.9% and 57.4%, on the RPi3 and RPiZW, respectively, compared to unit configuration ALLU. When WiFi is connected, using unit configuration EU reduces energy consumption by 37.3% and 48.85% for the RPi3 and RPiZW, respectively, compared to ALLU. These results also show the significant effect of WiFi communication on  $E_{sdn}$ . During the shutdown phase, the system invokes `ifdown` to ensure all connected networks are brought down properly and then powered off. This process takes around 0.8s when `avahi-daemon.service` is disabled. When this service is enabled, the duration varies depending on the network speed and configuration. In our testbed, we noticed a delay of up to 12s. For example, unit configuration EU w/ NET1 w/ WiFi increases energy consumption by 96.4% and 60%, for the RPi3 and RPiZW, respectively, compared to EU w/ NET1.

### B. Forced Shutdown

Not all IoT applications require a clean shutdown to remain functional. There are two alternative commands, `systemctl halt --force` (equivalent to the traditional `halt` command) and `systemctl halt --force --force` (equivalent to the traditional `halt -f` command), which result in shorter shutdown phases. In this paper we refer to these approaches as *forced shutdown* and *forced-forced shutdown*. Both of these approaches skip the steps performed by `shutdown` to notify the running processes of the impending shutdown and wait for them to exit gracefully. Therefore, steps such as recording the shutdown event and any `STDOUT` or `STDERR` output that are typically printed to a log file during the shutdown phase are skipped by these commands. In the case of `systemctl halt --force`, this may be acceptable, as the logging of system shutdown events is not often critical to IoT applications, and any necessary shutdown logs may be generated manually by the user application. `systemctl halt --force` also properly disconnects the networking interfaces by calling `ifdown` on all the connected interfaces.

Figure 11 shows the impact of using forced shutdown on both time and energy consumption. Comparing this figure against Figure 10 demonstrates the performance benefits of killing processes rather than gracefully terminating them. Although this mechanism results in lower energy consumption for both platforms, the effects are more apparent on the RPiZW. When gracefully shutting down, the RPi3 provides the system processes with more resources to finalize their operations and terminate correctly, resulting

in a shorter shutdown phase. On the other hand, when forced shutdown is used, running processes are merely killed, which does not require significant system resources. However, the OS still waits for the networking interfaces to be brought down before continuing the shutdown phase. As a result, the extra power provided by the RPi3’s processor is wasted, making the RPiZW more energy efficient during the shutdown phase.

The second approach, `systemctl halt --force --force`, is faster than calling `systemctl halt --force` because the processes are not killed. Instead, the processes are simply abandoned as the processor cores are stopped. This command physically halts the processor and cuts power almost immediately without bringing down network interfaces or unmounting filesystems. Eliminating these steps reduces the shutdown duration of the RPi3 and RPiZW to less than 200ms and 400ms, respectively, and the energy consumption to less than 700mJ and 350mJ, respectively, as Figure 12 shows. Unfortunately, according to the documentation [55], using this command may cause data corruption. However, some approaches exist to minimize the risk. For example, all processes required or started by the user application can be killed manually first, and the `sync` command must be run to commit unsaved buffers to the SDC. For preventing issues related to wear-leveling, increasing the percentage of unused storage on SDC reduces the frequency of moving critical data blocks by the wear-leveling algorithm. However, since the chance of data corruption is not fully eliminated, it is up to the user to calculate the risk involved in shutting down the system using this mechanism repeatedly (or across many devices). These calculations are outside the scope of this paper. Additionally, for both forced-shutdown approaches, the user must consider running some processes (such as `fake-hwclock save` and `systemd-random-seed save`) manually to ensure system integrity and security.

It is worth mentioning that the `halt` system call for the ARM architecture automatically calls `machine_power_off()` to power off the board rather than entering the traditional *halted state* where the board stays powered on after the processor is powered off. Other architectures may require additional flags or even different commands in order to achieve a similar result.

## VI. PARALLELIZING APPLICATION PROCESSES WITH USERSPACE INITIALIZATION

As we showed in the previous sections, the userspace initialization phase does not fully utilize the system resources. In this section, we propose *Pallex*, a parallel execution framework to run user applications during the userspace initialization phase. After presenting the implementation of Pallex and providing guidelines for applying this framework, we evaluate its performance considering various IoT application scenarios.

### A. Pallex

The basic idea of Pallex is to divide a user application into *stages* and run each stage based on the set of available

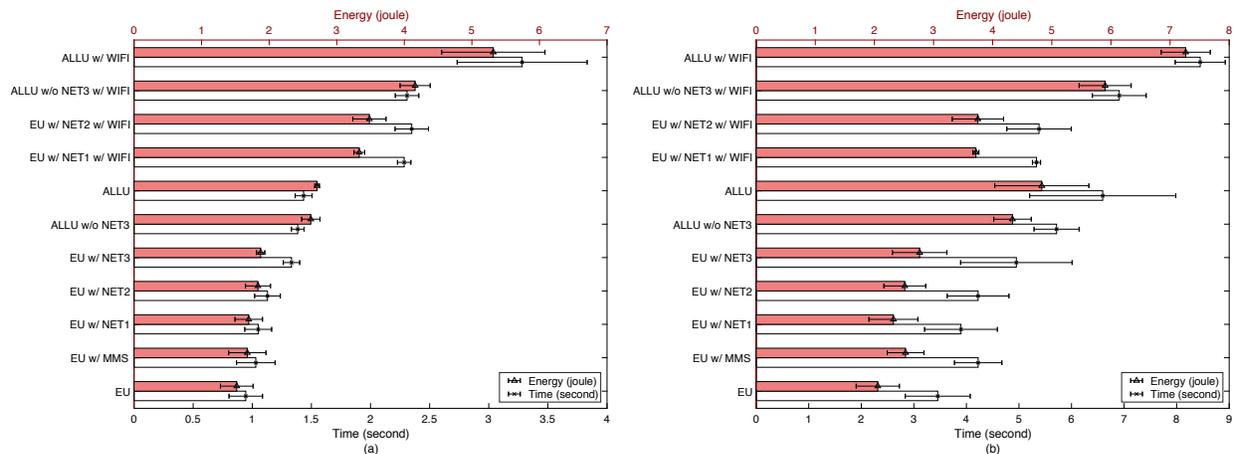


Fig. 10: The duration and energy consumption of *graceful shutdown* for the (a) RPi3 and (b) RPiZW when various unit configurations are applied. The lower and upper axes show  $T_{sdn}$  and  $E_{sdn}$ , respectively.

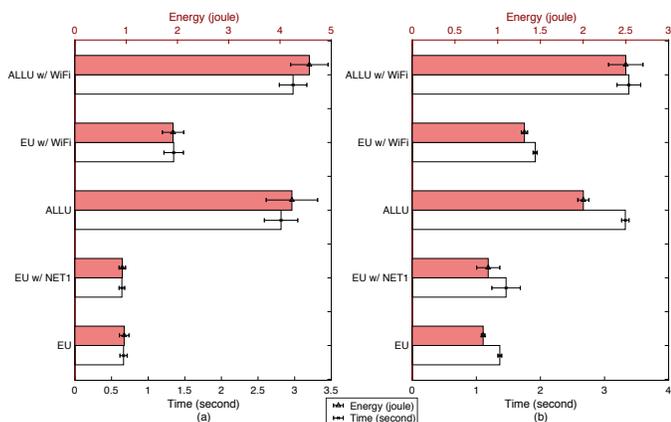


Fig. 11: The duration and energy consumption of *forced shutdown* for the (a) RPi3 and (b) RPiZW when various unit configurations are applied. The lower and upper axes show  $T_{sdn}$  and  $E_{sdn}$ , respectively.

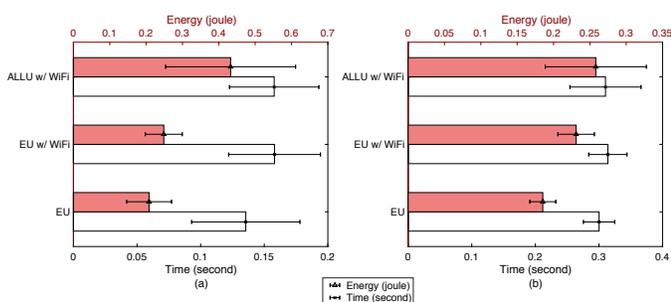


Fig. 12: The duration and energy consumption of *forced-forced shutdown* for the (a) RPi3 and (b) RPiZW when various unit configurations are applied. The lower and upper axes show  $T_{sdn}$  and  $E_{sdn}$ , respectively.

units and the completion of prerequisite stages. For a given user application, we break the code into a stage set  $S = \{s_i, s_j, \dots\}$ . Each stage  $s_i$  has two types of dependencies:

$$D(s_i) = \{s_j, s_k, \dots\} \quad (1)$$

and

$$D'(s_i) = \{u_j, u_k, \dots\} \quad (2)$$

where  $D'(s_i)$  is called the *stage dependency set* and refers to the set of user application's stages that must be completed before starting stage  $s_i$ , and  $D(s_i)$  is called *unit dependency set* and refers to the set of units on which stage  $s_i$  depends on. Each stage is invoked by `systemd` when the dependency sets specified in the stage's unit file are resolved. Therefore, for `systemd` to build the dependency tree and run the stages in an orderly manner, each stage's dependency sets must be specified in its corresponding unit file's `Requires` section.

Since the stages of execution might be implemented as independent processes, they may not share the same address space. Therefore, a mechanism is required to share data across processes. To this end, Pallex utilizes the mechanism offered by *Unix Domain Socket* (UDS) [62]. UDS is a data communication endpoint for exchanging data between processes executing on the same host OS, and provides a standard method for implementing Inter-Process Communication (IPC) on a Unix-based system. Since UDS handles communication within the Linux kernel, it is initialized in the kernel space as well. This indicates that we can use them as reliable means of communication during  $P_{usi}$ . When a stage  $s_i$  completes its execution and needs to transfer data to another stage  $s_j$ , the kernel blocks  $s_i$  until  $s_j$  is ready [63]. During this time, processor resources are used for the activation of units that are required by stage  $s_j$ . When the dependencies of  $s_j$  are resolved and it requests the shared data,  $s_i$  completes the data transfer and then exits.

There are at least three other IPC methods we could have used, including TCP/UDP sockets, POSIX message queues, and writing to a file. The reasons that we did not use these methods are as follows. First, using a TCP/UDP connection over localhost requires waiting for `networking.service` to load, which reduces the level of concurrency for those stages that do not need the networking capability. Besides, assume that a user application only uploads the results to the cloud when abnormal behavior is detected. In this case, relying on the `networking.service` unnecessarily increases the energy consumption of the

boot up process as we demonstrated in Section IV-B). Second, writing to a file increases the I/O overhead and affects the boot up time and message sharing delay because of the SDC activity during the userspace initialization phase. Lastly, we did not use POSIX message queues because they are very low-level and require careful configuration. Although both UDS and POSIX message queues are available almost concurrently, for the latter, it is necessary to configure the size and number of messages and their accepted/waiting status to make sure that enough buffer is available. Although message queues offer various features, UDS is standard, easy to use, and fast. UDS allow us to send data to the socket before the receiving program is even started with little to no configuration. Therefore, the sending program can be completely divorced from the receiver regarding dependencies. Also, UDS is agnostic towards payload size and can pause the sending program until the FIFO queue of data has begun to move and there is memory available to continue sending.

We further clarify the operation of Pallex through the scenario presented in Figure 13. This example assumes that the user application is composed of two processes (or threads), and the first process depends on the data generated by the second process to complete its task. Also, each process can be broken into a set of sequentially running stages. For the stages we assume that,

$$\begin{aligned} D(s_i) &= \emptyset, \quad D(s_j) = \{s_i\}, \quad D(s_k) = \{s_i, s_j\}, \\ D(s_m) &= \emptyset, \quad D(s_n) = \{s_m\}, \quad D(s_l) = \{s_i, s_j, s_k, s_m, s_n\} \end{aligned}$$

Each stage also depends on a set of units. Figure 13 shows the instances where the stage dependency set and unit dependency set of each stage are resolved. Since  $D(s_i) = \{\}$  and  $D(s_m) = \{\}$ , both stages can start at time  $t_1$  at which point their unit dependency sets are resolved. Therefore, these two stages run in parallel while the userspace units are being activated. At time  $t_2$ , stage  $s_i$  completes. However,  $s_j$  cannot be started because its unit dependency set is resolved at  $t_3$ . Therefore,  $s_i$  is blocked and waits until stage  $s_j$  is started and is ready to receive the data generated by  $s_i$ . Sharing the messages generated by a stage  $s_i$  with the next stages is denoted as  $M_{s_i}$ . This figure also shows that, since  $D(s_l) = \{s_i, s_j, s_k, s_m, s_n\}$ , stage  $s_l$  cannot be started before the completion of  $s_n$ . Once started, stage  $s_l$  reads  $M_{s_n}$  and  $M_{s_k}$ , the messages shared by two stages  $s_n$  and  $s_k$ . At time  $t_9$ , both the unit and stage dependency sets of  $s_l$  are resolved, and this stage has all the resources necessary to complete its operation. The user application finishes at time  $t_{10}$ . At this point, the system enters the shutdown phase.

The performance improvement achieved by Pallex highly depends on the number of and dependencies between the stages of a user application. Specifically, to minimize the waiting time of each stage and reduce energy consumption, it is important to break a user application into stages with small unit dependency sets. To this end, system developers must identify the main tasks performed by the application carefully, before restructuring them as

stages. Each stage must resolve all its dependencies before beginning execution. Therefore, each stage must comprise a set of instructions that once started, can be completed without relying on the completion of any other unit or stage. However, finding the right decomposition might not be a straightforward task due to the large number of units activated during  $P_{usi}$ . Besides, little to no improvement is observed by minimizing the waiting time of stages on units requiring execution time of only a few milliseconds. Due to the sampling-processing-sending nature of IoT applications, we can narrow down the list of essential units to simplify the task of decomposition, as follows.

- GPIO. These interfaces are initialized by the GPU. Therefore, as soon as `systemd` finishes its initialization (0.75s for the RPi3 and 1.2s for the RPiZW), user applications can use the GPIO pins. In addition to enabling the execution of user applications at the beginning of  $P_{usi}$ , using GPIOs provides a faster communication interface without the high overhead of a full networking stack, especially when a small amount of data must be communicated between nearby devices.
- I2C, SPI, CSI. The drivers for I2C (Inter-Integrated Circuit), SPI (Serial Peripheral Interface) and CSI (Camera Serial Interface) are loaded as kernel modules during  $P_{knl}$ . Therefore, the unit dependency of user applications relying on these components is resolved at the beginning of  $P_{usi}$ .
- Bluetooth. Bluetooth depends on both `hciuart.service` and `bluetooth.service`. As Section IV showed, the activation of these services finishes after `rc.local`. Therefore, user application stages that rely on Bluetooth must be started by creating a `systemd` service that starts after `bluetooth.service` instead of `rc.local`. If an application includes tasks that do not depend on these services, then running those tasks as stages that start before the completion of these services can result in a considerable performance improvement, especially due to their long activation duration.
- WiFi. Stages that rely on WiFi must be initialized after the activation of `network.target`. Another option is to start the stage after `network-online.target`, which is invoked once the network is *connected* as opposed to *available*.
- Ethernet. Similar to WiFi, the status of the Ethernet interface can be derived from `network.target`. A difference between the WiFi and Ethernet interfaces is that the speed of initializing Ethernet is faster because it does not perform the authentication and association process required for WiFi. Therefore, using Ethernet results in a shorter duty cycle. However, since most IoT applications rely on wireless communications, we mainly focus on WiFi in this paper.

In Appendix A, we present an overview of units to provide the users with guidelines regarding the impact of each unit on each application scenario.

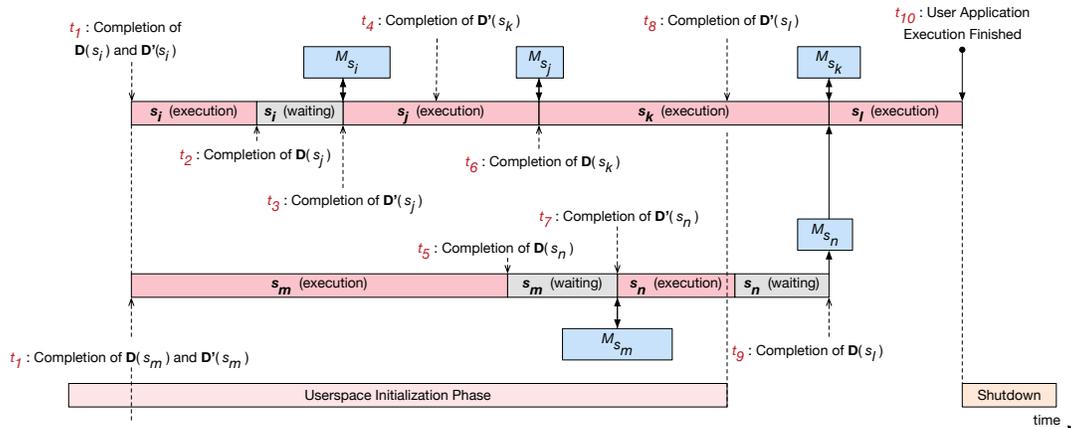


Fig. 13: An example of Pallex. The stage set of user application includes six stages,  $S = \{s_i, s_j, s_k, s_l, s_m, s_n\}$ . A new stage is started as soon as its stage dependency set and unit dependency set are satisfied. Stages share their messages using the mechanism offered by Unix Domain Socket (UDS).

### B. User Application Scenarios for Evaluating Pallex

In this section, we evaluate the performance of Pallex when applied to various types of user applications implemented on the RPi3 and RPiZW. Since using the RPi3 or RPiZW is justified when the application at hand cannot be accomplished using resource-constrained devices (such as those employing ARM Cortex-M or R processor), our user application scenarios include heavy operations like image capture, encryption, and classification. However, it should be noted that we omit image classification using RPiZW due to the high overhead caused by running the machine learning algorithm on this platform. We explain these scenarios in the following subsections.

1) *Scenario 1: Image Capture (IC)*: For this scenario we used a camera module [58] to capture an image. The application stage set includes only one stage,  $P = \{s_{cap}\}$ , where  $D(s_{cap}) = \emptyset$  and  $D'(s_{cap}) = \emptyset$ . Therefore, we use unit configuration EU for this scenario. The camera module is capable of capturing one picture per second (each around 2.5MB) and uses CSI to communicate with the RPi. After capturing the image, the user application saves the image as a JPEG file. The stage  $s_{cap}$  is written as a Python program that calls the picamera library [64] (written in Python) to capture the image.

2) *Scenario 2: Image Capture+Encryption (IC&E)*: We extend "Scenario 1" by encrypting the captured image using the AES-256 encryption algorithm. The application stage set includes one stage,  $P = \{s_{cap+enc}\}$ , where  $D(s_{cap+enc}) = \emptyset$  and  $D'(s_{cap+enc}) = \emptyset$ . Please note that since encryption depends on capture, we do not decompose the application into two stages. This scenario uses Unit configuration EU. The stage  $s_{cap+enc}$  is written as a Python program that calls the picamera library to capture the image, then calls the GNU Privacy Guard (gpg) utility (written in C) to encrypt the image.

3) *Scenario 3: Image Capture+Classification (IC&C)*: In this scenario, we capture an image and classify it using a pre-trained K-Nearest Neighbors (KNN) algorithm [65]. Since image capture and loading the model are

independent, we decompose the application into three stages: image capture ( $s_{cap}$ ), loading the KNN model ( $s_{load}$ ), and performing classification ( $s_{clas}$ ). Therefore,  $P = \{s_{cap}, s_{load}, s_{clas}\}$ , where  $D(s_{cap}) = \emptyset$ ,  $D'(s_{cap}) = \emptyset$ ,  $D(s_{load}) = \emptyset$ ,  $D'(s_{load}) = \emptyset$ ,  $D(s_{clas}) = \{s_{cap}, s_{load}\}$  and  $D'(s_{cap}) = \emptyset$ . Please note that  $s_{cap}$  and  $s_{load}$  are two concurrently running processes. Consequently, the two heavy stages, i.e., image capture and loading the model, start and run concurrently.  $s_{clas}$  depends on the completion of  $s_{cap}$  and  $s_{load}$  to perform its operation. This scenario uses Unit configuration EU. The implementation of the  $s_{cap}$  stage is as before. The stage  $s_{load}$  is implemented in a Python program that loads scikit-learn's KNN model (written in Cython) [66]. The stage  $s_{clas}$  is implemented in a Python program that uses scikit-learn's KNN algorithm to classify the image.

4) *Scenario 4: Image Capture+Upload (IC&U)*: In this scenario, we capture an image and transmit it to a cloud server through WiFi communication with an access point. The application stage set includes two stages,  $P = \{s_{cap}, s_{upl}\}$ , where  $D(s_{cap}) = \emptyset$ ,  $D'(s_{cap}) = \emptyset$ ,  $D(s_{upl}) = \{s_{cap}\}$  and  $D'(s_{upl}) = \{u_{net}\}$ , where  $u_{net}$  refers to the `network-online.target` unit. Please note that the service configuration we used for this scenario is EU w/NET1. Because networking services require a relatively long duration to initialize, we are interested in evaluating Pallex in scenarios where  $T_{usi}$  is long. For the RPi3 and RPiZW, we vary the number of images that are captured and uploaded from 1 to 3. For the RPi3, activating `networking.service` is approximately 3.5s, which means that we can capture a maximum of 3 images in parallel with this service activation without negatively impacting the duration of userspace initialization. Although activating `networking.service` on the RPiZW is longer than 7.4s, we cap the maximum number of images at 3 to present a fair comparison across the two boards. These sub-scenarios are referred to as IC&U $_x$ , where  $x$  refers to the number of images captured and uploaded. The implementation of the  $s_{cap}$  stage is as

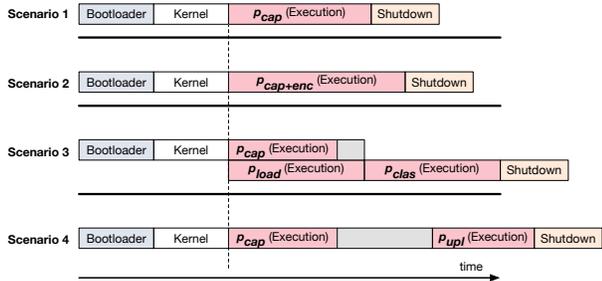


Fig. 14: The four scenarios used to measure the effect of Pallex on duration and energy consumption. We have used three variations of Scenario 4 to evaluate Pallex when using WiFi for uploading data.

before. The stage  $s_{upl}$  is implemented as a Python program that calls the `curl` utility (written in C) to upload the image.

Figure 14 shows a summary of the operations of these applications versus time.

**Threats to Validity.** As the authors in [67] showed, the choice of programming language affects the energy consumption of user application depending on the task being performed. This must be taken into account when the measurements are repeated using different programming languages than those we used.

### C. Results

Considering the user application scenarios given in Section VI-B, we present the performance measurement results of Pallex in this section. It must be noted that Pallex is compared against baseline scenarios that employ unit configuration to prevent the activation of unnecessary units. In the baseline scenarios, referred to as "EU" in the figures, the EU configuration is used, and the user application is started when `rc.local` is loaded.

Regarding duration and energy measurement, we consider the interval between the instance the RPi is powered on until the completion of the user application program. Please note that merely measuring the impact of Pallex on the duration of completing a user application does not reflect performance in terms of energy efficiency and lifetime. Since the OS dynamically adjusts the operating frequency of the processor cores based on load [68] (a.k.a., dynamic frequency scaling), it is critical to verify that the additional system load during  $P_{usi}$  does not eliminate the energy saving achieved by reducing the duration.

Figure 15 and 16 show the performance improvements achieved using Pallex for RPi3 and RPiZW during the boot up process, respectively. In these figures, the black and gray bars represent  $P_{btl}$  and  $P_{knl}$ , respectively. It must be noted that, since Pallex affects system performance after  $P_{knl}$ , both  $E_{btl}$  and  $E_{knl}$  are fixed irrespective to the RPi board used. From the userspace processing point of view, for the IC scenario, energy consumption is reduced by 24.7% and 10.77% for the RPi3 and RPiZW, respectively. For the IC&E scenario, we observe 27.22% and 5.98% improvement in terms of duration and 21.54%

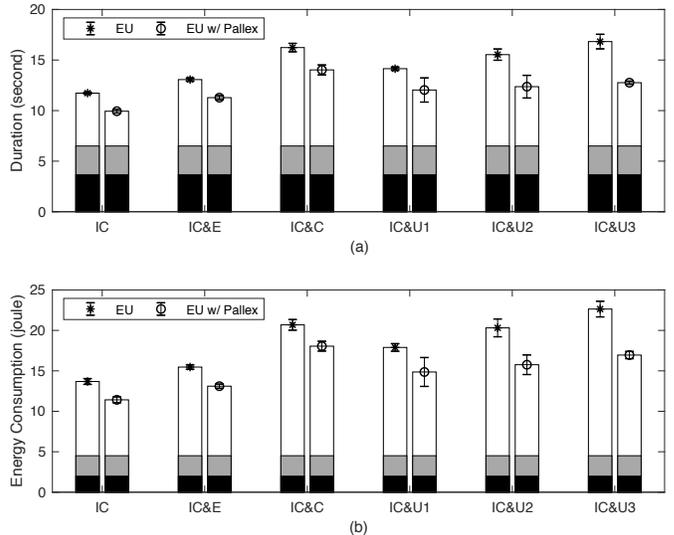


Fig. 15: The duration and energy consumption of Pallex versus normal launching of user applications. The device used is a RPi3. Black, grey, and white bars show the bootloader phase, kernel phase, and the completion of user application.

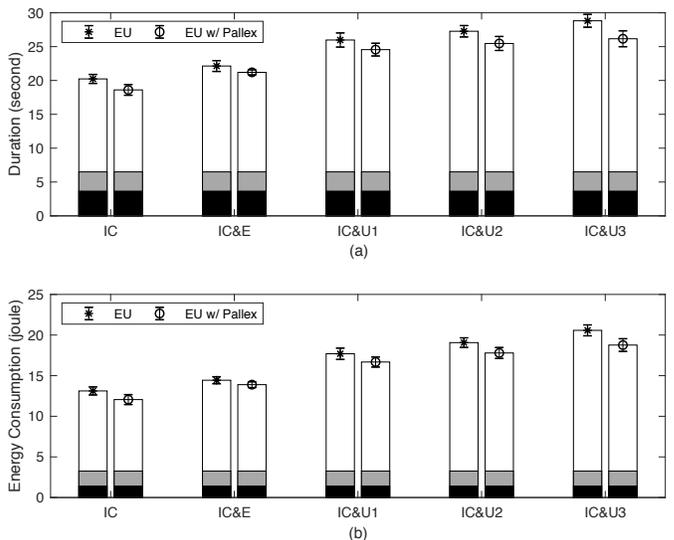


Fig. 16: The duration and energy consumption of Pallex versus normal launching of user applications. The device used is a RPiZW. Black, grey, and white bars show the bootloader phase, kernel phase, and the completion of user application.

and 4.89% in terms of energy for the RPi3 and RPiZW, respectively. Both of these scenarios include a single stage that is executed as soon as `systemd` is ready. It must be noted that the user application is not the only process running after the completion of `systemd` initialization. As explained in Section IV-B, the units that are essential for maintaining system integrity and stability are being activated during this duration as well.

Each IC&C scenario is composed of three stages, where two stages run in parallel and during  $P_{usi}$ . For this scenario we observe a 22.68% improvement in terms of duration and 16.29% in terms of energy for the RPi3.

The IC&U scenario is composed of two stages, where the first stage runs during  $P_{usi}$ . As we showed in the

previous sections, activating the `networking.service` is a lengthy process due to the initialization of hardware and networking utilities as well as association with the access point. Therefore, this scenario can significantly benefit from capturing images while the activation of networking services is in progress. In addition, more improvement is observed as the number of captured images increases: When one image is captured, we find 27.57% and 7.3% improvement in terms of userspace processing duration for the RPi3 and RPiZW, respectively. This results in an energy improvement of 22.67% for the RPi3 and 7.04% for the RPiZW. These improvements are increased to 39.36% and 11.89% for these two boards when three images are captured. Regarding energy improvement, we observe 31.35% and 10.45% improvement for the RPi3 and RPiZW, respectively, when three images are captured.

Comparing the two hardware platforms, we can observe that when Pallex is applied, the RPi3 shows a higher performance improvement compared to the RPiZW. These results are consistent with our observations in Section IV-B, indicating that the RPi3 platform consumes less energy than the RPiZW in a duty-cycling capacity despite drawing more current. The RPi3 parallelizes userspace initialization processes across multiple cores, resulting in a shorter duty-cycle duration. The impact of shortening the processing duration is more significant than the impact of the difference in current consumption across platforms.

To measure the impact of Pallex on the lifetime of duty-cycled systems, we compute system lifetime as follows:

$$\begin{aligned} \text{lifetime} &= \frac{E_{bat}}{(E_{btl} + E_{knl} + E_{usr} + E_{sdn}) \times N} \\ &= \frac{3600 \times 2400 \times 10^{-3} \times 5}{(E_{btl} + E_{knl} + E_{user} + E_{sdn}) \times N} \end{aligned} \quad (3)$$

where  $E_{btl}$ ,  $E_{knl}$ ,  $E_{user}$ , and  $E_{sdn}$  are the energy consumption of the bootloader phase, kernel phase, user application, and shutdown phase.  $E_{bat}$  is the available energy of the battery, and  $N$  is the number of cycles per hour. For the shutdown phase, we used the *forced shutdown* mechanism detailed in Section V, as it is faster than the `shutdown` command without sacrificing reliability. We also assume the capacity of the battery is 2400mAh and its voltage is 5V. Figures 17 and 18 show system lifetime versus the number of cycles per hour for the RPi3 and RPiZW, respectively.

The maximum improvement in the lifetime for the RPi3 is 30.16% for scenario IC&U3, and the minimum is 13.89% for scenario IC&C. Similarly, for the RPiZW, the maximum improvement is 9.01% for scenario IC&U3, and the minimum is 3.74% for scenario IC&C. For the networking-dependent scenarios, we notice that lifetime improvement increases versus the number of images captured and transmitted. Specifically, for the RPi3, we notice an improvement of 18.06% for IC&U1, 25.74% for IC&U2, and 30.16% for IC&U3. For the RPiZW, we notice an improvement of 5.67% for IC&U1, 6.7% for IC&U2, and 9.01% for IC&U3. For the networking-independent scenarios, IC&C achieves the highest increase in the lifetime

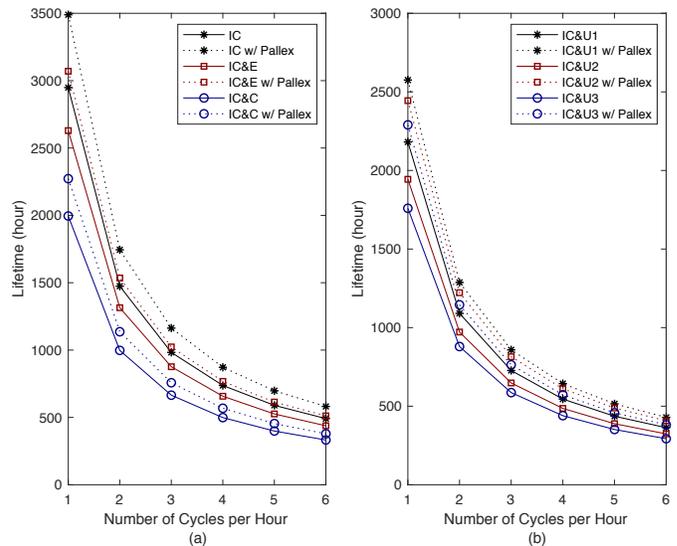


Fig. 17: Lifetime of the RPi3 for different user application scenarios. Sub-figure (a) shows networking-independent scenarios, and sub-figure (b) shows networking-dependent scenarios.

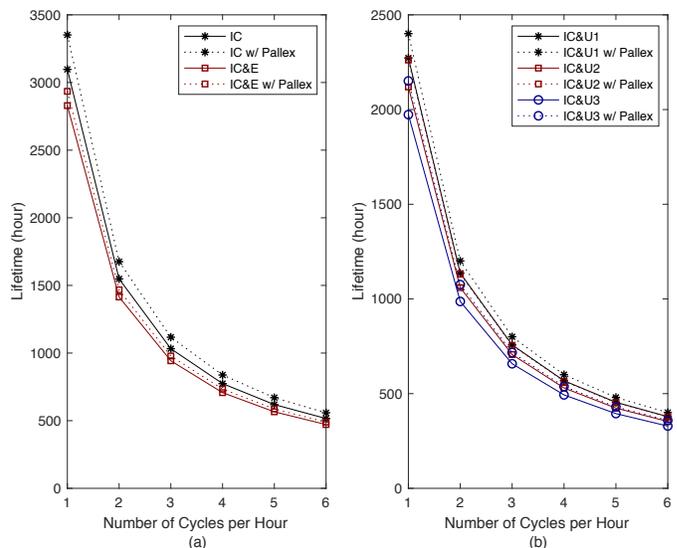


Fig. 18: Lifetime of the RPiZW for different user application scenarios. Sub-figure (a) shows networking-independent scenarios, and sub-figure (b) shows networking-dependent scenarios.

(18.33%) on the RPi3. This phenomenon is attributed to the higher processing demand and concurrent execution of  $s_{cap}$  and  $s_{load}$  stages.

The improvements in the lifetime, in particular, reduce the cost of energy harvesting systems as well as system maintenance. For example, when the amount of energy consumed per hour is reduced, a lower-cost energy harvesting system (e.g., smaller solar panels, smaller batteries) can be used as it is provided with more time to harvest and store energy. For fully battery powered systems, increasing lifetime reduces the frequency of system maintenance to replace or recharge the batteries.

## VII. THREATS TO VALIDITY

In addition to the discussions regarding threats to validity presented in the previous sections, we highlight additional internal and external threats in this section.

**Internal Validity.** The power measurement tool used for the experiments of this paper presents an accuracy of 2.5%, as proved in [38], when used in the room temperature of 23°C. Nevertheless, it should be noted that both the accuracy of the power measurement tool as well as the energy consumption of the board vary with temperature. Specifically, for the power measurement tools without a self-calibration feature, the impact of temperature on analog-to-digital converter (ADC) accuracy must be carefully taken into account.

**External Validity.** As mentioned in Sections III and IV-D, the experiments of this paper have been conducted using two different hardware platforms, one OS, and two different SDCs. Future work could address the applicability of these observations using different hardware platforms and software environments. Regarding OS, we have employed the March 2018 release of RSL. However, it must be noted that the results might be different for the new releases of this OS. A study similar to that performed in [69] can be very helpful to reveal the impact of OS maintenance and updates on performance.

**Reliability.** The testbed architecture and experimentation methodology have been presented in Section III in details. This enables the extension of this work to other systems and scenarios while establishing a clear baseline for comparison with our work.

## VIII. RELATED WORK

In this section, we overview the existing works relevant to energy measurement and modeling, boot up improvement, and energy efficient software development.

### A. Energy Measurement and Modeling

Although most of the existing works on the energy measurement and modeling of low-power systems focus on the communication cost [70], [71], there exist a few studies that model the energy consumption of the whole system. The authors in [72] present a systematic analysis of energy distribution in IoT devices that rely on energy harvesting. Although the accuracy of the proposed models have been verified against empirical evaluations, the proposed models are only applicable to IoT devices that utilize micro-controllers (e.g., Cortex-M4) and simple OSs (e.g., RTOS). In contrast, the RPi3 board includes a Cortex-A53 processor and the operation of its OS is significantly more complex than RTOSs such as FreeRTOS [9] and ThreadX [8]. PowerPi [20] proposes a power estimation model for the RPi Model B. The power estimation model is composed of the power consumption of a processor, an Ethernet NIC, and an external WiFi transceiver. Although the proposed models present a realistic power consumption of the board versus processor and communication load, the models do

not provide insights regarding the performance of duty-cycled systems regarding the boot up overhead. Besides, applying these models requires the pre-measurement of processor utilization and networking load, which might not be feasible for complex tasks. An extension of the PowerPi modeling approach has been proposed in [73] for various single-board computers. The authors in [74] model the power consumption of ARM-based architecture, using regression analysis to develop statistical power models. Attention is given to both the processor itself and the system as a whole, i.e. wireless radios, in different scenarios, such as browsing the web or multimedia benchmarking. The paper concludes that hardware changes such as increasing the size of the L2 cache can increase power efficiency, but does not go into detail concerning improvements that can be made in software to increase efficiency when modifying the hardware is not feasible.

### B. Boot-up Improvement

In this section, in addition to userspace optimization, we also study several other techniques for improving boot up, namely, bootloader optimization, kernel space optimization, and suspend optimization. Variants of these techniques can be utilized in conjunction with our work to further enhance boot up time and energy consumption.

1) *Userspace Optimization:* The authors in [28] identify several areas of improvement for boot up time on Debian, a Linux operating system, which at the time used `sysvinit`. Although Raspbian uses `systemd` instead of `sysvinit`, its backwards compatibility allows for some of these optimizations to translate over to devices using `systemd` as their `init` system. The first area of optimization involves substituting lighter default applications. For example, boot up time was improved by two seconds by using `dash` instead of `bash`, the default shell for Debian. Rewriting slow shell scripts to use internal functions instead of external functions has also been proven to improve the performance of boot up time.

Boot scripts can also be executed in parallel. This is achieved through setting the `CONCURRENCY` option in each boot script. Reordering boot scripts can also help programs complete faster. In [31], the authors identified areas during the boot up process of a smart TV when processor utilization is low. They used this information to optimize boot script execution by interleaving I/O-intensive programs with processor-intensive programs. While this can be an effective solution, maintaining system stability can be troublesome, as `sysvinit` does not track dependencies. Therefore, it is the user's responsibility to ensure any changes do not compromise system stability. Compounding the issue, dependencies between boot scripts are often not well documented and require thorough research and an understanding of both the Linux kernel and `sysvinit` before modification is possible. By comparison, `systemd`, a modern `init` system, tracks dependencies and automatically generates a tree structure to determine which startup tasks can be executed in parallel. Therefore, the feature set

of `systemd` makes boot script reordering a much simpler task.

In addition to improving boot up time, we studied several methods for shortening the duration until *user interaction time*: the time it takes until the user is able to interact with the application. In [28] the authors are able to gain an improvement of two seconds by performing certain tasks, such as setting up the network and hardware clock, in the background. In [30], [31], the authors optimize boot up time on Android smartphones and televisions by deferring long-running services until after the log-in screen is displayed to the user. However, this approach is not directly applicable to the IoT devices employing duty cycling.

*Application XIP* (Application Execute in Place) [29] is another technique for improving application start up speed. *Application XIP* allows for user applications to be executed directly from the filesystem instead of first being loaded into RAM. When a program is executed, the kernel program loader maps text segments for applications directly, resulting in a reduction in RAM footprint, faster first invocation, and reduced power consumption of flash vs RAM. However, *Application XIP* requires compiling the kernel with support for a filesystem (such as CRAMFS [75]) that supports storing files in contiguous blocks of memory. Additionally, there are hardware requirements for XIP such as random access storage that is directly accessible by the processor, which the RPi platform does not offer.

One major complexity of userspace initialization can be attributed to the evolution of the `init` system’s responsibilities. Modern `init` systems often manage process scheduling, I/O scheduling, and memory scheduling, among other responsibilities. In [76], the authors addressed this issue by simplifying the existing `init` scheme to improve an Android device’s boot up time. Their contribution involves separating the `init` scheme into two booting modes: normal boot, which executes all tasks during boot up time, and quick boot, which executes only mandatory tasks before user interaction time. However, duty-cycled IoT systems do not typically require separate modes for mandatory services and extra services, as an optimized system only runs the software required to complete its task. When device features, such as applying system updates, are not utilized during every duty-cycle, they are easily loaded after the boot up is completed. Furthermore, modern Linux kernels support dynamic kernel module loading using the `modprobe` utility.

2) *Suspend-related Improvements*: Another approach to improve boot up time is the usage of a suspend or hibernation mechanism. In [77], the standard *hibernate-resume* approach is used to optimize the boot up time of a digital camera. During the *system suspend* phase, current system information such as processor registers, I/O map information, and runtime variables (both global and local), are stored in RAM before the device powers off. When the system receives certain power events, the *resume* operation is started. The previous state of the device is

restored from RAM, resulting in a faster initialization time than a regular boot up. It is important to note that the RPi platform’s hardware does not support hibernation or suspension, as the SoCs do not support modern power management features [33].

In [27] and [78], methods for *snapshot boot techniques* are discussed with respect to boot up time improvement on a device running a Linux-based OS. These methods utilize a snapshot image created at boot up time and stored on a disk or in reserved flash memory. For subsequent boot ups, the device loads the suspended image instead of stepping through the standard kernel initialization process. As opposed to a standard *hibernate-resume* operation, this snapshot image is generated only once and then reused. A major disadvantage of snapshot images is the considerable amount time required for image generation, verification, and storage. Also, storing a full image in addition to the original system scripts and binaries necessary for image generation requires a considerable amount of disk space, which may not be available on many IoT platforms. If the image is corrupted (due to power loss, for examples), the device may become unrecoverable if a fallback image is unavailable. To mitigate these issues, the authors in [79] proposed a dual-image system, where some essential services and initialization are kept in the first image, and less critical functions are loaded by the second image, referred to as the *essential-snapshot-image* and *supplementary-snapshot-image*, respectively.

3) *Bootloader Optimization*: The bootloader is responsible for initializing the hardware and loading the kernel. In most cases, the kernel is compressed to save space and the bootloader must decompress it before use. The authors in [25] analyzed the time required by the bootloader and compared the performance of several root filesystems for fast boot up time. Decompressing the kernel with `gzip-checksum` on their system resulted in a bootloader time of 16s. Conversely, using `gzip-nochecksum` required only 12s, with a reduction of 3.8s. While the actual decompression of the image took 2.79s, verifying the checksum required an extra second on their device. In addition, the cost of calculating the checksum for the RAM disk was 2.78s. They found that storing the image in a decompressed format can circumvent this process. These results are noteworthy in the realm of embedded Linux, where processors may be slower than the RPi’s processor and the onboard flash is more responsive than the SDC used by the RPi. However, most general-purpose Linux systems can decompress the kernel faster than the storage medium can read the uncompressed alternative. Therefore, it is important to balance the speed of the processor and the I/O read speed before seeking performance gains from an uncompressed kernel.

Similar to *Application XIP*, *Kernel XIP* (Kernel Execute-In-Place) [29] is a popular technique used for optimizing bootloader initialization. In a typical boot sequence, the kernel is decompressed either during or just after it is loaded into memory. XIP enables direct execution of kernel instructions directly from ROM or

flash memory. However, this method requires the kernel be stored in an uncompressed format, thereby requiring more storage space. The RPi platform does not support `Kernel XIP` due to hardware constraints.

4) *Kernel Space Optimization*: In [25], the authors show the effect of removing unnecessary kernel modules and bundling multiple modules into a single module on decreasing kernel loading time. In another work [29], the authors evaluate the effect of disabling kernel print statements to prevent bottlenecks caused by streaming messages to the console. The `quiet` option in the kernel configuration changes the logging level of print statements to 4, which suppresses the output of regular (non-emergency) messages.

### C. Energy-Efficient Software Development

The importance of energy efficient software development has been attracting a lot of attention from both industry and academia [67], [80], [81]. The importance, in particular, is revealed by [14], which studies the popularity of energy-efficient software design by developers. Specifically, for mobile application development, the literature proposes various approaches and tools to study and improve energy consumption. For example, [82] combines empirical power measurement with statistical modeling to offer energy consumption information at the source line level. The authors in [83] propose a technique that first determines the execution trace of the code and then applies a set of cost functions to estimate the energy consumption of the execution path. The authors in [69] study the impact of software maintenance and update on energy consumption. Although these solutions do not directly address the energy consumption and overhead of duty-cycled systems from the boot up and shutdown point of views, they are complementary to our work and can be used to enhance the performance of these systems further. For example, developers can first use these techniques to improve the energy utilization of an user application and identify the energy bugs. Then Pallex can be employed to parallelize the execution of the user application's stages with the userspace initialization phase.

## IX. CONCLUSION

The increasing number of Linux-based IoT devices used for edge and fog computing necessitates the adoption of duty-cycling mechanisms to reduce the energy consumption of these devices. To this end, profiling and improving the operation of userspace initialization offers techniques that can be easily adopted by users. In this work, we presented a thorough study of the Linux boot up process, in particular the effects of unit activation on the duration and energy consumption of the boot up and shutdown phases. We showed that although some units cannot be disabled without compromising system stability, duty-cycling performance is significantly enhanced by application-specific unit configuration. Our studies also showed that there is no effect on boot up performance when up to 95% of the SDC

capacity is utilized. However, using a faster SDC results in a slightly shorter, more energy efficient duty-cycle. After analyzing the resource utilization of the RPi3 and RPiZW during the boot up process, we showed that user applications can be executed in parallel with the userspace initialization phase to reduce the energy consumption of each duty-cycle. To this end, we proposed Pallex, a parallel execution framework which relies on `systemd` and Unix Domain Sockets to break a user application into multiple phases. Our evaluations show up to a 31% reduction in energy consumption and up to a 30% enhancement in lifetime when Pallex is applied to various IoT application scenarios. Our studies also reveal the trade-off between processing power and current consumption. Although the current draw of the RPiZW is lower than that of the RPi3, this paper confirms the RPi3 platform is more suitable for duty-cycled applications. This is because the RPi3 parallelizes userspace initialization and user application processes across multiple cores, resulting in lower energy consumption by shortening processing duration.

Some potential areas of future work are as follows: Although the studies of this paper revealed the significant effects of quad-core and single-core SoCs on duty-cycling performance, extending these observations and profiling the performance of other COTS Linux-based boards is of interest. From the shutdown point of view, developing a model to predict the probability of SDC corruption based on factors such as capacity, I/O rate, and duty-cycling frequency enables users to choose the best shutdown mechanism available without compromising system reliability. Regarding Pallex, although we have provided guidelines to simplify its applicability to other IoT scenarios, it would be helpful to develop a program to analyze user application code and break it into stages based on the tasks it performs and the dependencies of those tasks. Finally, unit configuration and Pallex can be integrated with bootloader and kernel-level optimization mechanisms to further enhance performance.

## ACKNOWLEDGMENT

This research has been partially supported by the Santa Clara Valley Water District research grant SCWD02. This project involves the development of a flood monitoring system where Linux-based wireless systems, which rely on solar or battery power, capture images for analysis using machine learning to classify and report the debris carried by rivers and streams.

## APPENDIX A

### LIST OF UNITS IN RASPBIAN STRETCH LITE (RSL)

In this section, we present an overview of the units available in RSL. All units except those in the EU category can be disabled if they are not necessary for the application scenario being considered.

### A. Essential Units (EU)

- **boot.mount**: This unit helps `systemd` resolve dependency trees for units that depend on mounting `/boot` before activation.
- **dev-mmcblk0p2.device**: This unit brings the root partition on the SDC into the scope of `systemd` so that units that require the root partition's mount to finish before activation can resolve their dependencies properly.
- **dev-mqueue.mount**: This unit informs `systemd` when the POSIX message queues for internal system messages is ready.
- **kmod.service**: This service contains `modprobe`, which is used for loading and unloading kernel modules.
- **kmod-static-nodes.service**: This service creates a list of required static modules for the loaded kernel.
- **run-rpc\_pipefs.mount**: This unit directs `systemd` on how to mount the RAM-based `pipefs`, which is used every time a process is forked or a pipe ("|") is used.
- **sys-kernel-debug.mount**: Similar to `dev-mmcblk0p2.device`, this unit helps `systemd` resolve dependencies correctly. The actual mounting of `debugfs` occurs within `udev`, which is the daemon that detects hardware changes.
- **sys-kernel-config.mount**: This unit prevents the system from reaching `sysinit.target` until the kernel configuration parameters are fully loaded into the kernel from the Configuration File System (`configfs`).
- **systemd-fsck.service** and **systemd-fsck-root.service**: These services run `fsck` on each partition to ensure file system consistency. This is an important step, and does not run every time unless there are problems detected on the SDC.
- **systemd-journald.service**: Many programs rely on `journald` for logging output, including the kernel (through `kmsg`). Therefore, it should not be disabled. However, in order to speed up its initialization, it may be useful to lower the size limit of the journal logs since a dependency, `systemd-journal-flush.service`, must rotate this log file on initialization.
- **systemd-modules-load.service**: This service starts early in the userspace initialization phase to load static kernel modules.
- **systemd-remount-fs.service**: In the beginning of the userspace initialization phase, this service mounts the necessary API filesystems for the kernel (such as `/proc`, `/sys`, or `/dev`) to a RAM disk.
- **systemd-random-seed.service**: This service loads the random seed and saves it at shutdown to enable the device to generate a new value when the system restarts.
- **systemd-sysctl.service**: By loading kernel configurations, this service enables `systemctl` to perform as expected.
- **systemd-udevd.service**: This service initializes `udev`, a daemon that listens to kernel `uevents` and matches them against specified rules, to run scripts. For example, it can load drivers when a new device is attached, or

mount a USB drive when it is plugged in.

- **systemd-udev-trigger.service**: Devices plugged in before the system is powered on might not generate the kernel messages necessary for `udev` to discover them. This service probes and detects devices that `udev` would not normally discover.
- **sudo.service**: This service clears cached sudo privilege escalations to enforce user re-authentication after every reboot.
- **systemd-tmpfiles-setup.service** and **systemd-tmpfiles-setup-dev.service**: Mount `/tmp` and delete the old files. These services also create any files that are specified by user-provided configuration.
- **systemd-rfkill.service**: This service restores the `rfkill` state at the beginning of userspace initialization to ensure it matches the status saved before shutdown. Therefore, if the wireless peripherals (typically WiFi or Bluetooth) had `rfkill` preventing their use before shutdown, they will remain disabled on reboot.
- **systemd-update-utmp.service** and **systemd-update-utmp-runlevel.service**: Record and manage the system uptime, the logged-in users, and users' log-in method (such as `ssh`, `tty`, and serial port).
- **systemd-user-sessions.service**: This service enables user log-in and denies log-in attempts after the shutdown signal has been sent. If a device in the field does not require log-in capabilities, this service does not need to be started automatically. For example, it can be started by a helper program when a GPIO pin is pulled high.
- **systemd-logind.service**: This service is responsible for tasks such as user session management, processor usage quotas, and device access management.

### B. Networking-related Services (NRS)

- **avahi-daemon.service**: This service enables programs to discover and publish services and hosts running on a local network. Note that this service can significantly slow the speed of the `ifdown` command and therefore the shutdown process if not completely uninstalled. Unless necessary for the user application, `avahi-daemon` should be uninstalled.
- **bluetoothd.service**: Daemon for controlling the Bluetooth interface. `Bluez`, `bluetoothctl`, and many other Bluetooth-related utilities communicate through this daemon.
- **dhcpcd.service**: The daemon responsible for managing the DHCP protocol on all targeted network interfaces. This service can be disabled if the device does not require a network connection or is guaranteed a static IP address. The duration of IP allocation depends on external factors including link quality and the load of access point.
- **hciuart.service**: This is responsible for initializing the HCI bluetooth interface. HCI stands for "Host-to-Controller-Interface" and it is controlled over a serial UART interface.

- **networking.service**: Completes the configuration of WiFi and Ethernet interfaces based on the settings available in the `/etc/network/interfaces` configuration file.
- **nfs-config.service**: This service, along with **nfs-common.service**, loads configuration details applicable to Network File Systems (NFS).
- **rsyncd.service**: Daemon that listens on port 873 for incoming **rsync** file transfer requests. **rsync** is used for efficiently transferring and synchronizing files across computer systems.
- **rpcbind.service**: This service accepts requests for Remote Procedure Calls (RPC) and binds them to TCP ports for access and control.
- **sshd.service**: This service belongs to the **OpenSSH** package. It runs in the background to listen for and accept or deny incoming **ssh** connections according to a user-defined configuration file.
- **systemd-hostnamed.service**: This service can be used to control the hostname and related metadata by user programs.
- **systemd-networkd.service**: Brings up the system’s network manager and provides it with discovered networks, both physical and virtual.
- **systemd-resolved.service**: Provides local DNS resolution for namespaces such as **localhost** or those added by the user to overlay DNS provided by an external source.
- **systemd-timesyncd.service**: Used for time synchronization across the network.

### C. Memory Management Services

- **dphys-swapfile.service**: This service initializes, mounts, unmounts, and deletes swap files on the SDC. If the available RAM is enough for the user application, then disabling this service results in a performance enhancement in terms of faster boot up time and prolonged SDC lifetime. This service is usually required when the user application involves loading large machine learning models and data sets.

### D. I/O-related Services

- **alsa-utils.service**: Represents the tools relating to the Advanced Linux Sound Architecture (ALSA).
- **alsa-restore.service**: Initializes and restores the last state of the RPi’s onboard soundcard.

### E. Miscellaneous Units

- **fake-hwclock.service**: This service saves the current time to a file at shutdown and loads it at boot up time. Without this service, the RPi is unaware of the current epoch time until it establishes a network connection. An incorrect epoch value may cause some files to appear as if they are edited in the future.
- **plymouth.service**: Provides a flicker-free graphical boot up process. Other related services in-

- clude **plymouth-quit.service**, **plymouth-quit-wait.service**, **plymouth-start.service**, and **plymouth-read-write.service**.
- **raspi-config.service**: This service loads configuration changes made by the user such as processor governance, display overscan, and filesystem partition expansions, and applies them on reboot.
- **rsyslog.service**: Tools for log processing and conversion.
- **console-setup.service**: Configures the fonts, screen resolution, keyboard layout, etc., for virtual tty terminals.

## REFERENCES

- [1] P. Delforge, “Slashing energy use in computers and monitors while protecting our wallets, health, and planet.” Natural Resources Defense Council, 2016.
- [2] S. Chu and A. Majumdar, “Opportunities and challenges for a sustainable energy future,” *nature*, vol. 488, no. 7411, p. 294, 2012.
- [3] Z. Wang, Y. Liu, Y. Sun, Y. Li, D. Zhang, and H. Yang, “An energy-efficient heterogeneous dual-core processor for Internet of Things,” in *IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2015, pp. 2301–2304.
- [4] M. Ueki, K. Takeuchi, T. Yamamoto, A. Tanabe, N. Ikarashi, M. Saitoh, T. Nagumo, H. Sunamura, M. Narihiro, K. Uejima *et al.*, “Low-power embedded ram technology for iot applications,” in *Symposium on VLSI Technology*. IEEE, 2015, pp. T108–T109.
- [5] Cypress Semiconductor. CYW43907: IEEE 802.11 a/b/g/n SoC with an Embedded Applications Processor. [Online]. Available: <http://www.cypress.com/file/298236/download>
- [6] E. Baccelli, O. Hahm, M. Gunes, M. Wahlisch, and T. C. Schmidt, “RIOT OS: Towards an OS for the Internet of Things,” in *IEEE Conference on Computer Communications Workshops (INFOCOM Workshops)*. IEEE, 2013, pp. 79–80.
- [7] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer *et al.*, “TinyOS: An operating system for sensor networks,” in *Ambient intelligence*. Springer, 2005, pp. 115–148.
- [8] (2018) ThreadX RTOS Real-Time Operating System. [Online]. Available: <https://rtos.com/solutions/threadx/real-time-operating-system/>
- [9] (2018) The FreeRTOS Kernel. [Online]. Available: <https://www.freertos.org>
- [10] C. E. Jones, K. M. Sivalingam, P. Agrawal, and J. C. Chen, “A survey of energy efficient network protocols for wireless networks,” *wireless networks*, vol. 7, no. 4, pp. 343–358, 2001.
- [11] B. Dezfouli, M. Radi, and O. Chipara, “Rewimo: a real-time and reliable low-power wireless mobile network,” *ACM Transactions on Sensor Networks (TOSN)*, vol. 13, no. 3, p. 17, 2017.
- [12] B. Dezfouli, M. Radi, K. Whitehouse, S. A. Razak, and T. Hwee-Pink, “DICSA: Distributed and concurrent link scheduling algorithm for data gathering in wireless sensor networks,” *Ad Hoc Networks*, vol. 25, pp. 54–71, 2015.
- [13] S. Zoican and M. Vochin, “LwIP stack protocol for embedded sensors network,” in *9th International Conference on Communications*. IEEE, 2012, pp. 221–224.
- [14] G. Pinto, F. Castor, and Y. D. Liu, “Mining questions about software energy consumption,” in *Proceedings of the 11th Working Conference on Mining Software Repositories*. ACM, 2014, pp. 22–31.
- [15] J. Hong, Y.-G. Hong, and J.-S. Youn, “Problem Statement of IoT Integrated with Edge Computing,” 2018. [Online]. Available: <https://tools.ietf.org/html/draft-hong-iot-edge-computing-00>
- [16] E. Griffiths, S. Assana, and K. Whitehouse, “Privacy-preserving image processing with binocular thermal cameras,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technologies*, vol. 1, no. 4, pp. 133:1–133:25, 2018.
- [17] R. Kelly, “Internet of things data to top 1.6 zettabytes by 2022,” *Campus Technology*, vol. 9, pp. 1536–1233, 2016.

- [18] M. Chiang and T. Zhang, "Fog and IoT: An overview of research opportunities," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 854–864, 2016.
- [19] Eclipse Foundation, "Key Trends from the IoT Developer Survey 2018." [Online]. Available: <https://blogs.eclipse.org/post/benjamin-cabe/key-trends-iot-developer-survey-2018>
- [20] F. Kaup, P. Gottschling, and D. Hausheer, "PowerPi: Measuring and modeling the power consumption of the Raspberry Pi," in *39th Conference on Local Computer Networks (LCN)*. IEEE, 2014, pp. 236–243.
- [21] R. Morabito, "Virtualization on internet of things edge devices with container technologies: a performance evaluation," *IEEE Access*, vol. 5, pp. 8835–8850, 2017.
- [22] V. Vujovic and M. Maksimovic, "Raspberry Pi as a Wireless Sensor node: Performances and constraints," in *37th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE, 2014, pp. 1013–1018.
- [23] R. Fisher, L. Ledwaba, G. Hancke, and C. Kruger, "Open hardware: A role to play in wireless sensor networks?" *Sensors*, vol. 15, no. 3, pp. 6818–6844, 2015.
- [24] B. Dezfouli, M. Radi, S. A. Razak, T. Hwee-Pink, and K. A. Bakar, "Modeling low-power wireless communications," *Journal of Network and Computer Applications*, vol. 51, pp. 102–126, 2015.
- [25] K. H. Chung, M. S. Choi, and K. S. Ahn, "A study on the packaging for fast boot-up time in the embedded Linux," in *13th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE, 2007, pp. 89–94.
- [26] C. Villegas. (2006) Improve the Debian boot process. [Online]. Available: <http://bootdebian.blogspot.com>
- [27] H. Kaminaga, "Improving linux startup time using software resume (and other techniques)," in *Linux Symposium*, 2006, p. 17.
- [28] C. Villegas and P. Reinholdtsen, "State-of the-art in the boot process." Google Summer of Code, 2006. [Online]. Available: <https://pdfs.semanticscholar.org/a171/696ddb41ba8aad53cdccb6aba1c4547aa80e.pdf>
- [29] T. R. Bird, "Methods to Improve Bootup Time in Linux," in *Proceedings of the Linux Symposium*, 2004.
- [30] G. Singh, K. Bipin, and R. Dhawan, "Optimizing the boot time of android on embedded system," in *15th International Symposium on Consumer Electronics (ISCE)*. IEEE, 2011, pp. 503–508.
- [31] H. Jo, H. Kim, J. Jeong, J. Lee, and S. Maeng, "Optimizing the startup time of embedded systems: a case study of digital tv," *IEEE Transactions on Consumer Electronics*, vol. 55, no. 4, 2009.
- [32] E. Upton, J. Duntemann, R. Roberts, B. Everard, and T. Mamtora, *Learning Computer Architecture with Raspberry Pi*. John Wiley & Sons, 2016.
- [33] "BCM2835 ARM Peripherals." [Online]. Available: <https://www.raspberrypi.org/app/uploads/2012/02/BCM2835-ARM-Peripherals.pdf>
- [34] (2018) systemd System and Service Manager. [Online]. Available: <https://www.freedesktop.org/wiki/Software/systemd/>
- [35] J. Gorauskas, "Managing Services in Linux: Past, Present and Future," *Linux J.*, vol. 2015, no. 251, 2015.
- [36] "systemd.unit," 2018. [Online]. Available: <https://www.freedesktop.org/software/systemd/man/systemd.unit.html>
- [37] R. Love, *Linux Kernel Development*, 3rd ed. Addison-Wesley, 2010.
- [38] B. Dezfouli, I. Amirharaj, and C.-C. Li, "EMPIOT: An energy measurement platform for wireless IoT devices," *Journal of Network and Computer Applications*, vol. 121, pp. 135 – 148, 2018.
- [39] "C library for Broadcom BCM 2835." [Online]. Available: <http://www.airspayce.com/mikem/bcm2835/>
- [40] "Wiring Pi: GPIO Interface library for the Raspberry Pi." [Online]. Available: <http://wiringpi.com>
- [41] Tektronix. DMM7510 7½ Digit Graphical Sampling Multimeter. [Online]. Available: <https://www.tek.com/tektronix-and-keithley-digital-multimeter/dmm7510>
- [42] X. Jiang, P. Dutta, D. Culler, and I. Stoica, "Micro Power Meter for Energy Monitoring of Wireless Sensor Networks at Scale," 2007, p. 186.
- [43] T. Stathopoulos, D. McIntire, and W. J. Kaiser, "The energy endoscope: Real-time detailed energy accounting for wireless sensor nodes," *Proceedings of International Conference on Information Processing in Sensor Networks (IPSN'08)*, pp. 383–394, 2008.
- [44] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler, "Energy metering for free: Augmenting switching regulators for real-time monitoring," in *Proceedings of the 7th International Conference on Information Processing in Sensor Networks (IPSN'08)*, 2008, pp. 283–294.
- [45] J. Andersen and M. T. Hansen, "Energy Bucket: A Tool for Power Profiling and Debugging of Sensor Nodes," in *Proceedings of Third International Conference on Sensor Technologies and Applications (SENSORCOMM'09)*. IEEE, 2009, pp. 132–138.
- [46] R. Zhou and G. Xing, "Nemo: A high-fidelity noninvasive power meter system for wireless sensor networks," *Proceedings of the ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN'13)*, pp. 141–152, 2013.
- [47] S. Naderiparizi, A. N. Parks, F. S. Parizi, and J. R. Smith, "μMonitor: In-situ energy monitoring with microwatt power consumption," in *Proceedings of the IEEE International Conference on RFID (RFID'16)*. IEEE, may 2016, pp. 1–8.
- [48] I. Haratcherev, G. Halkes, and T. Parker, "PowerBench: A Scalable Testbed Infrastructure for Benchmarking Power Consumption," in *Proceedings of the International Workshop on Sensor Network Engineering (IWSNE'08)*, 2008, pp. 37–44.
- [49] T. Trathnigg, M. Jürgen, and R. Weiss, "A low-cost energy measurement setup and improving the accuracy of energy simulators for wireless sensor networks," in *Proceedings of the workshop on Real-world wireless sensor networks*, 2008, pp. 31–35.
- [50] N. Zhu and I. O'Connor, "Energy measurements and evaluations on high data rate and ultra low power wsn node," in *10th IEEE International Conference on Networking, Sensing and Control (ICNSC)*, 2013, pp. 232–236.
- [51] R. Hartung, U. Kulau, and L. Wolf, "Distributed energy measurement in WSNs for outdoor applications," 2016, pp. 1–9.
- [52] A. Pötsch, A. Berger, and A. Springer, "Efficient analysis of power consumption behaviour of embedded wireless iot systems," in *Proceedings of the Instrumentation and Measurement Technology Conference (I2MTC)*, 2017, pp. 1–6.
- [53] K. Gomez, R. Riggio, T. Rasheed, D. Miorandi, and F. Granelli, "Energino: a Hardware and Software Solution for Energy Consumption Monitoring," in *Proceedings of the International Workshop on Wireless Network Measurements (WiOpt'12)*, 2012, pp. 311 – 317.
- [54] S. Keranidis, G. Kazdaridis, V. Passas, T. Korakis, I. Koutsopoulos, and L. Tassioulas, "NITOS Energy Monitoring Framework: Real Time Power Monitoring in Experimental Wireless Network Deployments," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 18, no. 1, pp. 64–74, 2014.
- [55] "systemctl," 2017. [Online]. Available: <https://www.freedesktop.org/software/systemd/man/systemctl.html>
- [56] "The GNU Awk User's Guide." [Online]. Available: <https://www.gnu.org/software/gawk/manual/gawk.html>
- [57] S. Godard, "iostat," 2018. [Online]. Available: <http://man7.org/linux/man-pages/man1/iostat.1.html>
- [58] (2018) Camera Module (v2). [Online]. Available: <https://www.raspberrypi.org/documentation/hardware/camera/>
- [59] (2018) IMX219PQ: Diagonal 4.6mm 8.08M-Effective Pixel Color CMOS Image Sensor. [Online]. Available: [https://www.sony-semicon.co.jp/products\\_en/new\\_pro/april\\_2014/imx219\\_e.html](https://www.sony-semicon.co.jp/products_en/new_pro/april_2014/imx219_e.html)
- [60] S. K. Paul Rubin, David MacKenzie, "dd: convert and copy a file," 2018. [Online]. Available: <http://man7.org/linux/man-pages/man1/dd.1.html>
- [61] A. Gupta, Y. Kim, and B. Urgaonkar, "DFTL: A Flash Translation Layer Employing Demand-based Selective Caching of Page-level Address Mappings," in *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2009, pp. 229–240.
- [62] "Linux Programmer's Manual, Socket." 2018. [Online]. Available: <http://man7.org/linux/man-pages/man2/socket.2.html>

- [63] D. P. Bovet and M. Cesati, *Understanding the Linux Kernel: from I/O ports to process management*. "O'Reilly Media, Inc.", 2005.
- [64] D. Jones. Python interface to the Raspberry Pi camera module. [Online]. Available: <https://picamera.readthedocs.io/en/release-1.13>
- [65] N. S. Altman, "An introduction to kernel and nearest-neighbor nonparametric regression," *The American Statistician*, vol. 46, no. 3, pp. 175–185, 1992.
- [66] "scikit-learn: machine learning in python." [Online]. Available: <https://scikit-learn.org/>
- [67] S. Georgiou, M. Kechagia, P. Louridas, and D. Spinellis, "What are your programming language's energy-delay implications?" in *Proceedings of the 15th International Conference on Mining Software Repositories (MSR)*. ACM, 2018, pp. 303–313.
- [68] D. Brodowski, N. Golde, R. J. Wysocki, and V. Kumar, "CPU frequency and voltage scaling code in the Linux (TM) kernel." [Online]. Available: <https://www.kernel.org/doc/Documentation/cpu-freq/governors.txt>
- [69] A. Hindle, "Green mining: a methodology of relating software change and configuration to power consumption," *Empirical Software Engineering*, vol. 20, no. 2, pp. 374–409, 2015.
- [70] X. Vilajosana, Q. Wang, F. Chraim, T. Watteyne, T. Chang, and K. S. Pister, "A realistic energy consumption model for TSCH networks," *IEEE Sensors Journal*, vol. 14, no. 2, pp. 482–489, 2014.
- [71] Q. Wang, M. Hempstead, and W. Yang, "A realistic power consumption model for wireless sensor network devices," in *3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks (SECON'06)*, vol. 1. IEEE, 2006, pp. 286–295.
- [72] B. Martinez, M. Monton, I. Vilajosana, and J. D. Prades, "the power of models: Modeling power consumption for iot devices," *IEEE Sensors Journal*, vol. 15, no. 10, pp. 5777–5789, 2015.
- [73] F. Kaup, S. Hacker, E. Mentzendorff, C. Meurisch, and D. Hausheer, "The progress of the energy-efficiency of single-board computers," *Tech. Rep. NetSys-TR-2018-01*, 2018.
- [74] J. Nunez-Yanez and G. Lore, "Enabling accurate modeling of power and energy consumption in an ARM-based System-on-Chip," *Microprocessors and Microsystems*, vol. 37, pp. 319–332, 2013.
- [75] Cramfs: cram a filesystem onto a small ROM. [Online]. Available: <https://www.kernel.org/doc/Documentation/filesystems/cramfs.txt>
- [76] G. Lim, J. young Hwang, K. Park, and S.-B. Suh, "Enhancing init scheme for improving bootup time in mobile devices," *2015 Eighth International Conference on Mobile Computing and Ubiquitous Networking (ICMU)*, pp. 149–154, 2015.
- [77] C. Park, K. Kim, Y. Jang, and K. Hyun, "Linux bootup time reduction for digital still camera," in *Linux Symposium*, 2006, p. 231.
- [78] I. Joe and S. C. Lee, "Bootup time improvement for embedded linux using snapshot images created on boot time," in *The 2nd International Conference on Next Generation Information Technology (ICNIT)*. IEEE, 2011, pp. 193–196.
- [79] K. Baik, S. Kim, S. Woo, and J. Choi, "Boosting up embedded linux device: experience on linux-based smartphone," in *proceedings of the Linux Symposium*, 2010, pp. 9–18.
- [80] G. Procaccianti, H. Fernández, and P. Lago, "Empirical evaluation of two best practices for energy-efficient software development," *Journal of Systems and Software*, vol. 117, pp. 185–198, 2016.
- [81] G. Pinto, K. Liu, F. Castor, and Y. D. Liu, "A comprehensive study on the energy efficiency of java's thread-safe collections," in *IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2016, pp. 20–31.
- [82] D. Li, S. Hao, W. G. Halfond, and R. Govindan, "Calculating source line level energy information for android applications," in *Proceedings of the International Symposium on Software Testing and Analysis*. ACM, 2013, pp. 78–89.
- [83] S. Hao, D. Li, W. G. Halfond, and R. Govindan, "Estimating android applications' cpu energy usage via bytecode profiling," in *Proceedings of the First International Workshop on Green and Sustainable Software*. IEEE, 2012, pp. 1–7.