APPROVAL SHEET

Title of Thesis: Low Power Context Aware Hierarchical System Design

Name of Candidate: Stanislav Bobovych
Computer Engineering PhD,
2020

Thesis and Abstract Approved: N. Banerjee
Nilanjan Banerjee
Associate Professor
Department of Computer Science and
Electrical Engineering

Thesis and Abstract Approved:
Ryan Robucci
Associate Professor
Department of Computer Science and
Electrical Engineering

Date Approved: April 9, 2020

ABSTRACT

Title of Thesis: Low Power Context Aware Hierarchical System Design

Stanislav Bobovych, Computer Engineering PhD, 2020

Thesis directed by: Nilanjan Banerjee, Associate Professor
Ryan Robucci, Associate Professor
Department of Computer Science and
Electrical Engineering

Embedded systems have come to play central roles in our lives. Smart, wearable ultra low power systems have become prevalent thanks to advances in sensing, communication, computation, and miniaturization. Though the individual subsystems have improved, designing these systems is a challenge since their compute, communication, and sensing components are tightly coupled. System engineers creating these devices must consider the whole hardware/software stack and the application at hand in order to stay within the limited energy budget of these systems. There has been a lot of research focused on each particular facet of low power systems, yet few have tried to approach the problem from the system and application integration level. The objective of the proposed research is to explore various hierarchical architectures through the development and deployment of systems for geospatial and medical applications.

Low Power Context Aware Hierarchical System Design

by

Stanislav Bobovych

Thesis submitted to the Faculty of the Graduate School
of the University of Maryland in partial fulfillment
of the requirements for the degree of
Computer Engineering PhD
2020

I dedicated this work to my loving wife, family and friends.

ACKNOWLEDGMENTS

Though many people have influenced me throughout this journey, I would like to thank the following for contributing to my success in this endeavor:

My advisers, Dr. Nilanjan Banerjee and Dr. Ryan Robucci, for their mentorship, dedication and guidance during my graduate studies.

My thesis committee, including Dr. Chintan Patel, Dr. Tim Finin, Dr. Nirmalya Roy, and those mentioned above for their time, support and advice during the preparation of this work.

Dr. Richard Allen for lending his expertise in sleep medicine and making possible our data collection at Johns Hopkins Center for Sleep.

Dr. Alex Nelson for charting the waters of wearable capacitive sensors in our research group.

Xiaoxia Zheng for lending her expertise and time in developing and fabricating capacitor sensor prototypes.

Gurashish Singh and Rebecca Baldwin for their work on initial prototypes of CapGait and RestEaZe.

Matt Law and Matt Fertig for their assistance in collecting sleep studies data.

Dr. David Lachut for being an amazing lab mate whose conversations with me helped develop and flesh out ideas.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

INTRODUCTION

Energy efficient systems have evolved from environmental monitoring to use in wearables, and it is estimated that by 2025 this market will grow to $70 billion (Dr Peter Harrop & Holland 2015). One of the major driving forces for research in embedded low power systems is that they can be used for pervasive monitoring since they are non-intrusive and appropriate for use in health care, sports and assistive technologies. Health care will be the largest segment of this emerging market due to the phenomenon of aging population in developed countries where on-the-go health monitoring and diagnosis will be a key factor in reducing cost of medicine and providing a high standard of well being. Low-level advancements in energy harvesting, sensors, and wireless technology have greatly improved the capabilities of embedded systems. It has also become clear that many optimizations can be gained from hardware/software co-development. These tightly integrated, application-specific solutions will be able to outperform non-custom implementations. However, this approach is very expensive, time consuming and requires a wide range of experts, so there is still a lot of demand for using off-the-shelf components for small-scale embedded systems. In this latter case, traditional techniques for reducing energy usage such as event-driven computation, synchronous sampling, and power-saving modes are not enough to extract the maximum utility from an embedded system. In addition, creating and developing the software architecture for embedded systems devices remains a manual, trial-and-error process that is challenging in competitive fields of commercial, industrial and research sectors.

Systems can be divided into several categories such as: industrial, home, and embedded. These categories exist on a wide power usage spectrum (see Fig 1.1). In the embedded domain, systems can be divided into three tiers: large, medium, and small. Large systems are defined as those that use a few watts. They tend to be powered by a reliable power source such as AC mains. Large systems can consume a few milliwatts to a few watts and are powered by replaceable or rechargeable batteries with potential to utilize renewable energy sources such as solar panels or micro wind turbines. Medium systems consume hundreds of microwatts to a few microwatts and can be powered with similar sources as large systems, albeit on a smaller scale. Small systems can consume in the range of nanowatts to microwatts and tend to be powered exclusively by energy harvesters without reliance on large batteries. This thesis focuses on embedded systems that fall within the range of medium to small. These systems are special-purpose, standalone, portable, low-profile, limited by power resources, have small computational and communication capabilities.

Embedded systems are constrained by energy and power. The small physical size of these systems constrains the size of the battery and hence the total energy capacity. The advances in battery capacity are evolving slowly, so it's not possible to rely on battery technology improvements over the short term of embedded system development. Power consumption must be optimized to create a practically useful system. In general, it affects the design of the power-supply in voltage regulators, system dimensions and cooling. Therefore, these systems have to be designed around the trade off between flexibility/availability and power efficiency. The relationship between energy and power is summarized by this equation 1.1. Faster execution translates to less total energy used, but may necessitate higher peak power which can increase the complexity of the power supply circuitry and heat sinking solutions.

$$E = \int P(t)dt \tag{1.1}$$

Thus, the energy consumption of a system is driven by the application requirements and is constrained by the availability of energy and total operational time of the system. Power consumption can be analyzed at different levels, with the lowest of these being transistor level and the highest being the system level. Broadly, the strategies for reducing energy utilization include: exotic electronics, decreasing supply voltage, reducing the duty cycle of system, curtailing number of duty cycles, redundant subsystems, completely shutting off system, and power supply gating (i.e., turning off unused components).

One approach to building high-availability, low-power systems is to design completely custom ASICs and sensors in hardware/software co-development for a particular application. The tightly integrated, application-specific solution will be able to outperform non-custom solutions. However, this is very expensive, time-consuming, and requires a wide range of experts, thus there is still a significant demand for using off-the-shelf components to build systems. On the other hand, duty-cycling approaches are simple to implement, but can drastically reduce availability of a system. An alternative approach is hierarchical design.

This thesis focuses on system level-design of low-power embedded systems using a hierarchical approach. A hierarchical system architecture consists of several tiers that operate at different power consumption levels. The higher power tiers can be switched off or transitioned into low-power modes while lower tiers are operating. The lower tiers minimize the use of higher power tiers through duty cycling, virtualization of services, virtualization of sensors and context-based computation. The challenge is making this transparent to the application while extending the operating time of systems. This work demonstrates the efficacy of approaching low-power embedded system design with an architecture that

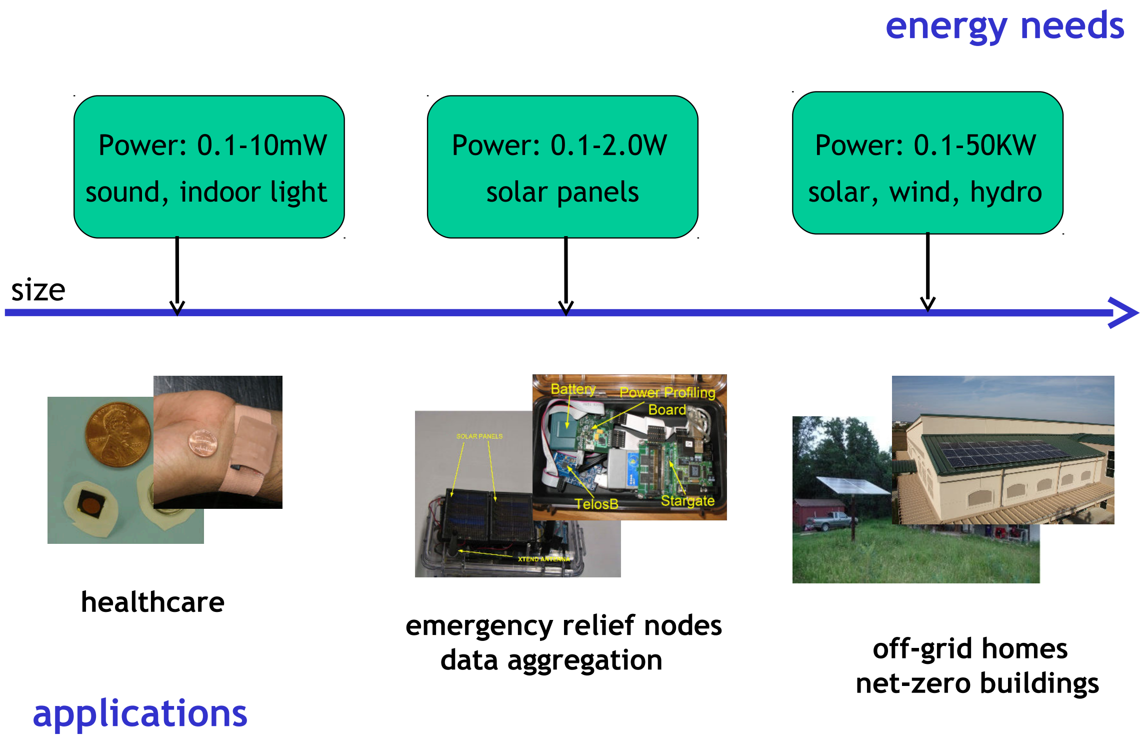

FIG. 1.1. The gamut of energy scavenging systems.

considers the hardware as time varying computation and sensing.

1.1 Contributions

- Developed a multi-tiered, multi-radio hardware architecture that supports high availability and high computational capability while surviving on energy harvested from medium-sized solar panels.
- Developed a hierarchical context-aware signal processing architecture for gait analysis. Implemented an FPGA based prototype of this system that used innovative wearable textile-based capacitive sensor arrays (CSAs).

- Developed wearable sensor hardware and software for sleep activity monitoring and gait data collection that can be used for other biomedical applications. The sleep monitoring research yielded a data set consisting of synchronized and labeled electromyograph (EMG) of the tibialis anterior muscle, 3-axis accelerometer measurements, 3-axis gyroscope measurements, capacitive data from three channels and video of the subjects. The data is labeled with respect to the subject's lying position (front, back, left side, right side), leg movement (none, small, medium, large), and foot flexion (none, tiny, medium, large). This data set will be used in research of restless leg syndrome detection, leg movement classification and sleep fragmentation.
- Developed a solar panel emulation system that has a wide operating range, is field portable and cost effective compared to contemporary systems.

1.2 Thesis organization

The organization of this thesis is as follows. Chapter 1 is the motivation for this thesis. Chapter 2 is a summary of prior findings in the realm of hierarchical system design. The common theme throughout this thesis is how to integrate processing, sensing, and harvesting to build context aware hierarchical systems that break down the signal processing algorithm into a hierarchy of processing elements. This includes approaches to building large and medium hierarchical systems and methods for evaluating them. We have explored ways to virtualize services of a high power system using a low power system through statically partitioning the functionality between a high and low power processor. Another approach we explored is determining context using a low power subsystem and then using this context to intelligently decide whether to wake up a low power processor. In addition, we have gained experience evaluating energy harvesting solutions such as solar cells which will aid us in evaluating other energy scavenging techniques. Chapter 3 describes the research methodology used by this work. Chapter 4 covers Perpetuu, a medium scale system with a large solar panel which virtualizes services by task sharing between a high and low power processor based on context and the technique of switching off devices. Chapter 5 examines RestEaZe, a sleep monitoring system which virtualizes sensors. Chapter 6 presents CapGait, a sensing system for gait analysis. It is a small-scale system which used task sharing between a processor and custom hardware. In this implementation the sensing and processing were tightly integrated and based on context. Chapter 7 present SunaPlayer, a tool for rapidly prototyping and profiling a wide scale of solar powered embedded systems. Chapter 8 presents a conclusion to this work and potential future directions of research.

Chapter 2

BACKGROUND

This thesis draws from various fields of study in embedded systems, sensing, signal processing and circuit design. This field is rapidly evolving, with what is considered low power changing very quickly as new technologies emerge to drive down energy usage and increase efficiencies. For example, in the past the microcontroller, sensors and communication were separate, discrete hardware components. Developments in integrating these components have resulted in state of the art System on Chip (SoC) with integrated radios and sensor controllers like CC26xx series from Texas Instruments (TI). Though system integration has reduced the complexity of hardware design and allows for physically smaller systems, the control, management and application software of systems remains complex. Managing the various power states of modern SoCs and microcontrollers is so complex that the use of real time operating systems with power-aware scheduling has become standard.

2.1 Embedded Low Power Systems

When discussing Embedded Low Power (ELP) systems, it is important to distinguish between power-aware and low-power. Low-power refers to designing sensors and other components to use as little energy for a given task as possible. On the other hand, power-aware means taking into account the available energy and deciding whether to perform an action or not. One approach to being power-aware is to schedule tasks based on power requirements through system level modeling (Liu *et al.* 2001). To accurately model energy usage of components, benchmarking circuitry and software have been developed to serve as sample data for simulators (Shnayder *et al.* 2004a). One of the major drivers for this research has been sensor networks where the goal is to balance communication and energy efficiency of nodes. For example, Hempstead presents a comprehensive system design which uses the techniques of power profiling, event-driven computation, hardware acceleration, common case optimization, and fine grain power management (Hempstead *et al.* 2005). Others have explored localized signal filtering and data compression (Yazicioglu *et al.* 2009) as well as event-driven, non-uniform sampling of analog data (Fesquet, Sicard, & Bidégaray-Fesquet 2010) (Roa *et al.* 2014). An orthogonal drive in ELP design has been to reduce the physical size of the system to make them wearable and integrable into every day items such as clothing (Park *et al.* 2006). This research has had an impact on design of battery powered consumer devices, with recent examples such as FitBit, Sensoria Fitness Socks, and Owlet baby monitoring smart sock. As sensing and computational technology improves, ELP design has shifted from running systems on a battery to using the battery as a backup or completely eliminating the battery altogether.

2.2 System Design

The overall goal of system design for an embedded low power system is to balance energy consumption with availability and computational capability. This can be solved with techniques that consist of three categories: energy management, signal processing and sensor design.

2.2.1 Energy Management Techniques

Energy management in mobile and sensor systems has been studied for over a decade. Approaches such as CPU Dynamic Voltage Frequency Scaling (DVFS) (Yuan & Nahrstedt 2003; Govil, Chan, & Wasserman 1995), disk spindown (Helmbold, Long, & Sherrod 1996), turning off banks of RAM (Huang, Pillai, & Shin 2003), microsleep (Huang, Pillai, & Shin 2003), Wi-Fi power saving modes (Rahmati & Zhong 2007; Pering, Raghunathan, & Want 2005), and multiple radios (Shih, Bahl, & Sinclair 2002; Pering *et al.* 2006) are applied for conserving energy on laptops and mobile phones. While such techniques can save energy on a computationally capable embedded platform, they do not eradicate its high *non-reducible* power draw incurred by the system. Previous work on Hierarchical Power Management (HPM) (Sorber *et al.* 2005b; Banerjee *et al.* 2007) resolves this issue by combining multiple platforms of varied power draw and capabilities while performing intelligent task sharing between the platforms. The applications of HPM is in developing more energy efficient laptops (Sorber *et al.* 2005b), sensor network data aggregation points (Banerjee *et al.* 2007), and more recently the design of energy-efficient desktops (Agarwal *et al.* 2009). Finally, the concept of virtualizing services of a high power tier through a lower power tier has been explored in the development of a low power sensor hub for smart phones (Shen *et al.* 2015).

2.2.2 Hierarchical Signal Processing

Hierarchical architecture can help minimize the energy consumption of the system by selectively keeping the higher power subsystem switched off when not required. The decision of which subsystem should be turned on or off can be based on simple algorithms such as peak-valley detection or more complex machine learning algorithms such as Hidden Markov Models (Sako & Kitamura 2013; Pang & Ding 2013) and higher order learning techniques (Yin *et al.* 2013; Lu 2013). Some examples of systems designed with hierarchical architecture are Somniloquy, Turducken, and UMassDieselNet (Agarwal *et al.* 2009; Sorber *et al.* 2005a; Banerjee, Corner, & Levine 2007).

2.2.3 Capacitive Sensing

Electrical capacitance is the ability of a body to store an electric charge. When talking about capacitance, a canonical parallel plate model is used where two conductive parallel plates are separated by a dielectric. The capacitance of this configuration is summarized in 2.1, where C is the capacitance in Farads, A is the area of overlap of the two plates in square meters, ε is the permitivity of the dielectric material between the plates, and d is the separation between the plates in meters.

$$C = \varepsilon \frac{A}{d} \tag{2.1}$$

This type of sensing has been applied to industrial, automotive, and healthcare applications (Ouh *et al.* 2012). For instance, capacitive sensors have been used in positioning (Ripka & Tipek 2013), humidity sensing (Gründler 2007), tilt sensing (Terzic, Nagarajah, & Alamgir 2011), pressure sensing (Palasagaram & Ramadoss 2006), MEMS-based sensing (Palasagaram & Ramadoss 2006), and gesture recognition (Nelson *et al.* 2015). The proximity sensor is a critical device with applications to robotics, industrial monitoring, mobile phone touch screens and healthcare (Lo *et al.* 2011; 2012).

In one modality, sensor hardware measures capacitance changes from a single electrode that acts as one plate of a capacitor while the other plate is an object that is virtually grounded. The measured capacitance can then be used to estimate the distance between the sensor while the target object cheaply and using comparatively less energy than other

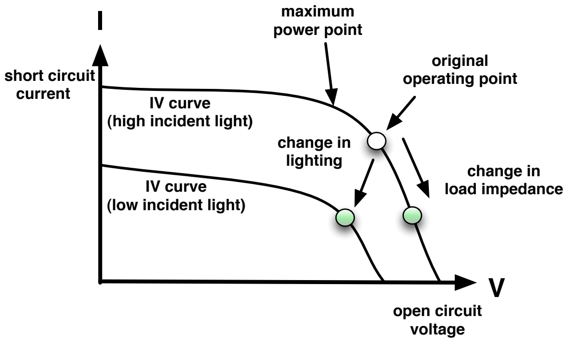

FIG. 2.1. The figure shows two IV curves for a solar cell. The IV curve is characterized by the short circuit current, open-circuit voltage, and the maximum power point. Each point on the IV curve corresponds to a load impedance. Each IV curve is a function of the ambient light intensity and temperature. The IV characteristics are highly non-linear.

sensors.

2.2.4 Energy harvesting with solar panels

Solar panels are very popular renewable energy harvesters since they strike a good balance between size, cost and energy output. Solar or Photovoltaic (PV) panels are a non-linear device that harvest energy from incident light such as natural sunlight or light from artificial sources like incandescent and fluorescent lamps. PV panels scale from small cells used for micro-harvesting indoor light, medium sized panels that power sensing and environmental monitoring systems, to large PV arrays that power buildings. Small photovoltaic cells are capable of generating over 50 $\mu\mathrm{W}/\mathrm{cm}^2$ (Gruetter 2011) and have been used in battery-less sensor nodes that use a super capacitor to store energy instead of batteries (Rahman *et al.* 2012). One obvious limitation of solar cells is that they do not generate energy when the light source is obscured or is not present.

Photovoltaic (PV) panels are made of PV modules that convert incident light energy into electric charge. A PV cell is modeled by a current-voltage or IV trace (Figure 2.1) that is characterized by an open-circuit voltage (OCV) (the voltage drop across the panel leads when there is no load, i.e., voltage limit), a short circuit current (SCC) (the current draw when the leads are connected using a shunt, i.e., current limit), and the maximum power point (MPP) (the maximum power harvested by the solar cell ($I \cdot V$), i.e., power limit). The IV characteristics of the panel between SCC and OCV is highly non-linear (National Instruments), and each point on the curve corresponds to a load impedance. The PV cell acts as a current source, and if the load current increases, the voltage across the load drops. The IV curve of a solar panel, and thus the operating point of the panel, also changes with ambient temperature and incident light intensity. The goal of a solar emulation platform is, therefore, to assure that the voltage and current supplied to an attached load is on the appropriate IV curve for a wide range of solar cells.

2.3 Applications of Low Power Systems

Low-power embedded systems have a great variety of applications such as medical, military, emergency services, commercial and space. In this work, the emergency services and medicine will be used to demonstrate various hierarchical architectures.

2.3.1 Emergency services

Natural disasters are characterized by lack of network connectivity and power grid access. While thousands of lives are lost during a natural disaster, the aftermaths of such calamities can also be severe. For instance, consider Superstorm Sandy (`http://en.wikipedia.org/wiki/Hurricane_Sandy`) or the Tokyo earthquake (`http://www.bbc.co.uk/news/world-asia-pacific-12711226`). Victims were left without power and Internet access for weeks. A critical task after a natural disaster is to rescue survivors and guide them from their present location to safe zones or relief camps. However, due to the lack of access to centralized map services like Google maps, delivering such data to survivors is unlikely.

2.3.2 Medical applications

Biomedical signals, especially those mechanical in nature, tend to be low frequency. For example, in a recent gait study (Toulouse, Badaoui, & Serviere 2015) pressure sensors were used in the heel that were sampled at 128 Hz. Another paper applied frequency analysis to Vertical Ground Reaction force data set and found that most of the information content was below 10 Hz (Karout *et al.* 2015). Event complex activity detection has been achieved by only using an accelerometer that was sampled at 51.2 Hz (Lee *et al.* 2015). Multi-sensor systems typically rely on interrupt driven sensing with fixed rates, and in the example of the eWatch, it was 20 Hz (Maurer *et al.* 2006). This work was extended to examining the accuracy and performance of non-uniform sampling techniques such as random, exponential back off, and entropy based sampling compared to uniform sampling (Krause *et al.* 2005). Restless leg syndrome detection and gait analysis are used as test cases in this work.

Restless Leg Syndrome Restless sleep is a major problem in our society affecting productivity, health and quality of life for adults and learning ability of children. The leg movements characterizing restless sleep uniquely mark sleep quality and several medical conditions. However, capturing the complexity of these leg movements in sleep requires sophisticated sensors and analytics, currently limited to sleep laboratories. Leg movements during sleep can be divided into Complex Leg Movements during Sleep (CLMS) and Periodic Leg Movements of Sleep (PLMS). CLMS are general leg movements while PLMS movement is the classic periodic dorsiflexion of the foot at the ankle produced by activation of the anterior tibialis muscle. The significance of these two movements differ. CLMS patterns reveal texture and severity for sleep-behavioral intervention; PLMS reveal possible medical conditions such as restless legs syndrome (RLS), sleep apnea, and REM behavior disorder. Both need to be measured. The standard PLMS measure uses muscle activity (EMG) from electrodes secured to the skin in a costly and limited access sleep lab. There is no simple ambulatory way to accurately record leg EMG in the home. A new method enabling simple ambulatory PLMS and CLMS measures is required for sleep characterization.

Gait In 2010 alone, 2.3 million fall injuries and 662,000 hospitalizations due to falls were reported in the United States (Center of Disease Control and Prevention 2010). Each year, one in every four adults above the age of 65 fall (Center of Disease Control and Prevention 2010). The healthcare costs associated with the aftermath of falls are severe, totaling 30 billion dollars in 2010 alone (Center of Disease Control and Prevention 2010). While fall detection methods can help provide immediate remedy to the patient, fall prediction and prevention are the only plausible ways to minimize the economic and health impact of falls. While observable factors such as improper medication can be a valid predictor of a future fall, the spatio-temporal variation in the gait of an individual can reliably predict future falls (Maki 1997; Bridenbaugh & Kressig 2010). For instance, gerontology researchers have observed that increase in stride-to-stride variability, stride speed, and spatio-temporal variation in inter-leg spacing have a high correlation to falls (Helbostad & Moe-Nilssen 2003). Unfortunately, these variations are usually small and often occur over a long period of time (Maki 1997), so capturing these variabilities solely in a clinical setting creates several difficulties (Bridenbaugh & Kressig 2010).

Chapter 3

APPLICATIONS AND DEPLOYMENTS

In order to validate various approaches to the design and architecture of low power systems, we built end-to-end systems for specific applications. These systems collected data in real settings, which we could analyze and use for controlled experiments.

We prototyped our map-server node, `perpetuu` 4.3, using a multi-tiered system that combines a PIC micro-controller subsystem with a gumstix subsystem. It houses a low power Wi-Fi radio, a high bandwidth Wi-Fi radio, and a long range XBee radio. The two tiers are powered by a common power supply board that houses the harvester circuit. The harvester circuit charges two Li-ion batteries from a medium-sized solar panel. Each sub-module was profiled for power consumption and a comprehensive system-wide power consumption model was built. The solar panel used by `perpetuu` was also profiled and modeled, resulting in an energy production model. With these models in hand, we experimented using solar irradiance data from real world locations to make sure that the system was tested against a wide operating range.

We prototyped our sleep monitoring system, RestEaZe 5.4, using textile capacitive sensors built into the ankle band and a data collection board which houses a IMU, Bluetooth module, MSP40 micro-controller, flash memory, capacitive measurement IC, and charging circuitry. Similar to Perpetuu, a system power model was created for RestEaze along with a profile for the battery. The system was used to collect data from a variety of subjects in a clinical setting. This data was used to train a model which could detect motion using low power capacitive sensors. Using this model, the collected data, and the system power model we were able to experiment with various ways to reduce power consumption and estimate the total system run time.

We prototyped our gait monitoring system 6.1 using capacitive sensor arrays built out of conductive textile sewn into an elastic bandage. The data was collected and processed using a custom system that consists of Spartan VI FPGA, Bluetooth module, and two capacitive measurement ICs. We evaluated the power consumption of our hierarchical signal processing architecture by synthesizing the FPGA design in a 180-nm CMOS process. Since the design is heirarchical, we estimated the power consumption of each hierarchy using dynamic power analysis on our synthesized ASIC cores.

Chapter 4

VIRTUALIZING SERVICES

Designing a system for dissemination of personalized map data after a natural disaster is challenging. First, the deployed system must function without the need for external power. The system should be self-sustainable and should be computationally capable to render information-rich maps to end-users. Hence, a research challenge is to balance computational capability and availability of the system with unpredictable energy supply. Secondly, during a natural disaster the topography of the scene changes frequently. For instance, a road may be flooded, a tree might have collapsed, or a bridge might have fallen. Therefore, for the map data to be reliable, the GIS data store resident on the nodes should be up-to-date. Since a backhaul link to an Internet-resident GIS server is absent after a natural disaster, it is critical that the system uses alternate mechanisms to update its GIS data store. Finally, the system should be able to disseminate maps seamlessly on common handhelds such as mobile phones without the need to install additional applications. During a natural disaster it is important that useful information be disseminated on devices that the users are comfortable using.

4.1 Perpetuu: A Heirarchical GIS Emergency Dissemination System

To address the above challenges, we present the design of `perpetuu`, a self-sustainable solar powered GIS microserver that can serve information-rich maps and directions to end users on web browsers on their common handhelds like laptops and mobile phones. The system comprises of a low-power embedded node that can survive and function under variable energy harvested from medium sized solar panels. The `perpetuu` nodes may be networked together using a long range, low-power radio to form a mesh network. At the core of `perpetuu` is a tiered system architecture that combines a low-power computationally weak micro-controller subsystem, a high-power computationally-capable micro-processor subsystem, and a virtualized software architecture to explore trade-offs in energy consumption, system availability, and computational capability. The node runs a custom map stack that can serve multi-scale maps to survivors over a Wi-Fi link. `perpetuu` uses a novel crowdsourcing approach where users can annotate maps with dangerous areas on standard browsers, and the GIS store on the `perpetuu` nodes are updated continuously using these annotations. The key contribution of `perpetuu` is the design of a self-sustainable system using hardware and software components that can serve rich map data to survivors of natural disasters. A boxed `perpetuu` node (illustrated in Figure 4.3) can be deployed from an aircraft into a disaster zone.

4.1.1 Research Contributions

The design, implementation, and evaluation of `perpetuu` presents a three-fold research contribution.

- **Hardware architecture for emergency response systems:** We present a multi-tiered, multi-radio hardware architecture that supports high availability and high computational capability while surviving on energy harvested from medium sized solar panels. High availability is critical in an emergency scenario where survivors can request data at *any* time, and computational capability is important since the system needs to generate map tiles and map layers that contain in-vitro generated routes from the user's present location to safe zones. Specifically, our system combines a micro-controller device with a micro-processor unit, and efficient power profiling and power gating circuits. The platform can support computationally intensive services like in-situ rendering of maps and high availability, at an overall low-power consumption. The architecture, therefore, is fine-tuned for an emergency response application.

- **Virtualized software architecture for emergency response systems:** We present a set of software techniques where the micro-controller subsystem virtualizes resources at the micro-processor subsystem. For instance, dynamic data from users (such as map annotations) is cached at the micro-controller node and synchronized with other nodes in the network lazily at low-power consumption. The system houses a novel custom-designed map stack that can render maps at different resolutions as a function of residual battery capacity. We combine a wakeup controller algorithm, hardware measurement-driven energy predictors, and low-power Wi-Fi sensing (to detect survivors), to minimize the use of the higher power micro-processor subsystem.
- **Prototype implementation:** We have implemented a fully functional `perpetuu` node using a low-power PIC micro-controller and a gumstix platform, augmented with a custom power supply and harvester board. The gumstix runs a fully functional map stack that can serve maps to users on their handhelds and can generate routes and directions from their present location to areas of safety. While our final goal is to build `perpetuu` as a highly optimized system-on-chip (SoC), our designed prototype serves as a proof-of-concept demonstration. We evaluate `perpetuu` using

measurements on our prototype and trace-based simulations, and demonstrate that it can operate near-perpetually while serving maps to a large number of survivors.

4.2 Design Goals

The key contribution of the `perpetuu` system is applying a multi-tiered architecture in the design of self-sustainable systems that are fine-tuned for emergency response applications. Before describing the hardware and software components of `perpetuu`, we first detail the underlying research goals that motivate the design choices of a `perpetuu` node.

Balance energy consumption with availability and computational capability: A `perpetuu` node must in-situ generate map tiles from vector and raster data, detect survivors, and render maps on a survivor's handheld device with low latency. Since survivors may request maps from a `perpetuu` node at any time, the nodes should be highly available and vigilant. Moreover, since a `perpetuu` node will be deployed in a natural disaster scene with no access to the grid power, the nodes must survive on small and variable energy generated by medium sized solar panels. Hence, it is critical that the system balances the variable energy supply with the required computational capability and availability.

Maintain up-to-date GIS store: The physical topography of a disaster scene can change quickly. For instance, a tree might have fallen on a road, an area might be flooded, or a building might have collapsed. Providing safe directions to users on maps must account for these topography changes. Hence, the map repository on the `perpetuu` nodes must be updated frequently. However, the lack of a backhaul connection to a GIS or map server makes it challenging to design a system where the underlying raster and vector data is not stale.

Portability and compatibility with off-the-shelf handhelds: `perpetuu` needs to be portable and weatherproof, since the nodes may be thrown into a post-natural disaster scene after a catastrophe. Moreover, it is important that survivors of natural disasters can use familiar handhelds like smartphones or tablets to download maps from a `perpetuu` node during the aftermath of a natural disaster without having to install a special application. This requires a `perpetuu` node to act as a GIS microserver and render maps with a Google map-like interface on web browsers over a smartphone supported wireless interface like Wi-Fi.

4.3 Motivating a Tiered Architecture

To build a hardware architecture that best fits the design goals for a `perpetuu` node described above, we profiled the power consumption of two platforms— a PIC micro-controller (`http://en.wikipedia.org/wiki/PIC_microcontroller`) and a gumstix. (`http://www.gumstix.com/`). The PIC micro-controller is a computationally weak platform, with a maximum clock speed of 8 MHz and is optimized to operate in a power regime between 5 mW and 14 mW. Conversely, the gumstix platform hosts a 600 MHz ARM cortex processor, and is computationally capable but has a power signature in the range of 1560 mW and 1700 mW, with the in-built Wi-Fi radio in transmit mode. We measured these power values using our prototype node. A GIS microserver node that in-vitro generates maps and safe routes from a user's present location to safety camps using road networks, and houses map data, requires a computationally capable platform like the gumstix. However, the system also needs to be highly vigilant—a survivor may request maps from the node at *any* time. An always-on gumstix platform in its idle mode can consume 100s of milliwatts of irreducible power (Lorch 1995), a power budget that cannot be supported by medium-sized solar panels. The PIC micro-controller, on the other hand, can maintain always-on operation while consuming 10s of milliwatts. Therefore, a system that combines both platforms and intelligently shares tasks between the two platforms will export a wide spectrum of power states and can therefore explore trade-offs in energy consumption, system availability, and computational capability. We show the potential energy savings of a tiered system using simple analysis below. While multi-tiered systems have been used in the past for building Throwboxes (Banerjee, Corner, & Levine 2007), sensor systems (Banerjee *et al.* 2007), and vigilant desktops (Agarwal *et al.* 2009), the key challenge that we address in this work is the use of the multi-tiered architecture for a specific application—map dissemination to survivors after a natural disaster. The rest of the chapter, therefore, focuses on the key design decisions made in using the tiered architecture to build a self-sustainable GIS microserver.

4.3.1 Energy Savings

In this section, we present a simple analytical study that demonstrates the advantage of using a multi-tiered system. In our analysis, we assume that there are two processing tiers — a low-power tier that consumes an average of P_l Watts when active and a higher power tier that consumes an average of P_h Watts when active. We assume a workload where the higher power tier must be awake fraction f of the time, and the average energy consumed in the boot up of the higher power tier is P_b. We conservatively assume that the boot up cost is incurred for the same fraction f of the time that the higher power tier is switched on. An alternative to the tiered system is a single tier system that uses the higher power tier alone and consumes P_h. The tiered architecture consumes lower energy than a single tiered system if

$$P_h > f \cdot (P_l + P_h + P_b) + (1 - f) \cdot (P_l) \tag{4.1}$$

Hence, for $f < \frac{P_h - P_l}{P_h + P_b}$, the tiered architecture will consume lower energy than a system that comprises of the higher power tier. Using the numbers for P_h, P_l, P_b from Table 4.1, for $f < 0.5$, the tiered architecture performs better than a single tier system. Also, the tiered architecture provides the same availability as a single tiered system. This analysis demonstrates that for reasonable f, our system can provide the same availability at lower energy consumption when compared to a single tiered system.

Subsystem	Power (W)
PIC (sleep)	0.005
PIC (active)	0.014
PIC (taking measurements)	0.020
XBee RX/idle	0.241
XBee TX	1.0
Wi-Fi (low-power) TX	0.486
Gumstix (bootup)	1.270
Gumstix + Wi-Fi (idle)	1.560
Gumstix + Wi-Fi (transmit)	1.700

FIG. 4.1. Power consumption of subsystems for the `perpetuu` node.

4.4 Hardware Architecture

The system architecture for a `perpetuu` node is illustrated in Figure 4.2. The system uses a multi-tiered, multi-radio hardware architecture that is fine tuned for a self-sustainable GIS micro-server application. The system consists of two subsystems (a low-power and a high-power subsystem) that are tightly integrated and powered using a common power supply. The low-power platform is a micro-controller (we use a PIC (`http://en.wikipedia.org/wiki/PIC_microcontroller`) for our prototype) with supporting hardware circuits, a long range low-power XBee radio, and a low-power Wi-Fi radio. The high-power subsystem is a computationally capable microprocessor-driven system (a gumstix (`http://www.gumstix.com/`) is used for our prototype) that houses a medium range high bandwidth Wi-Fi radio, and a novel custom designed map stack for serving maps to users. The high-power subsystem remains switched

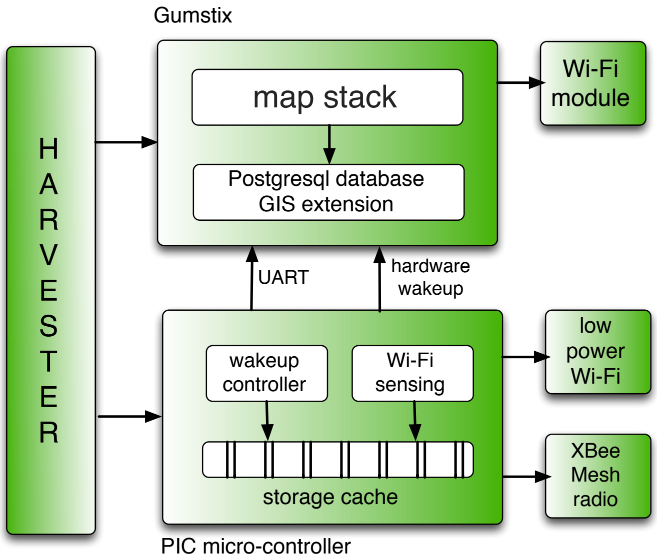

FIG. 4.2. The overall architecture of the `perpetuu` node. A `perpetuu` node consists of two subsystems; a computationally constrained low-power micro-controller tier with a long range XBee radio and a low-power Wi-Fi radio, and a computationally capable higher power micro-processor-driven subsystem with a high-power Wi-Fi radio. Our prototype uses a PIC micro-controller and a gumstix module for the low-power and high-power subsystems respectively. Software services on the low-power tier virtualizes the software services on the higher power tier (gumstix). This allows high vigilance at low energy consumption. The multi-tiered system is powered by a custom power supply that consists of a harvester that can scavenge solar energy using medium sized solar panels.

off most of the time and is woken up to serve maps only when survivors request maps and the system has sufficient residual energy. The micro-controller subsystem supports several low-power sleep and idle modes that can consume as low as 10s of microwatts of power. The low-power subsystem is frequently woken up from these low-power modes via timer interrupts to scan for Wi-Fi enabled devices (i.e., survivors) using the low-power Wi-Fi radio. The low-power subsystem, therefore, maintains high availability with low energy consumption. The low-power subsystem controls the wakeup of the higher power subsystem using an intelligent wakeup controller described in §4.5.4. The two tiers are tightly integrated using a custom designed power supply that provides power gating, wakeup controls, fine grained power measurements, and a harvester circuit board for scavenging energy from solar panels and charging a Li-ion battery. The custom board provides a serial interface for communication between the two subsystems and I/O devices. The XBee radio is used to form a network of `perpetuu` nodes. The range of the XBee radio is close to 1 kilometer when consuming 1W in transmit mode. The Wi-Fi radios, on the other hand, with a 9 dBi antenna, have a range of close to 200 m. The use of external antennas minimize the number of `perpetuu` nodes required to cover a geographic region.

The hardware architecture for `perpetuu` bears semblance to other systems built in the past (Agarwal *et al.* 2009; Sorber *et al.* 2005a; Banerjee, Corner, & Levine 2007). However, there are several hardware components in `perpetuu` that are designed to support map dissemination in a disconnected environment like a natural disaster. For instance, `perpetuu` has hardware support for solar harvesting using temperature-controlled maximum power point tracking. The maximum power point tracking (MPPT) IC senses the surrounding temperature and appropriately configures the regulator so that the harvesting circuit operates within 5% of the maximum power point of the solar panel used. `perpetuu` provides support for fine grained energy measurements and power gating. For instance, our custom power supply board can collect high frequency measurements on power harvested from solar panels, residual battery capacity, the power draw of the micro-controller, the power draw of the gumstix, the power consumed by the XBee and low-power Wi-Fi radio. We use low-power current accumulators to measure the power consumed by different subsystems. Using analog switches, power can be gated to the low-power Wi-Fi module, the XBee module, as well as the gumstix. The fine-grained energy measurements as well as the ability to power cycle peripherals and subsystems allow `perpetuu` to implement fairly sophisticated energy controllers that can balance energy harvested from the solar panels with energy utilized by the system.

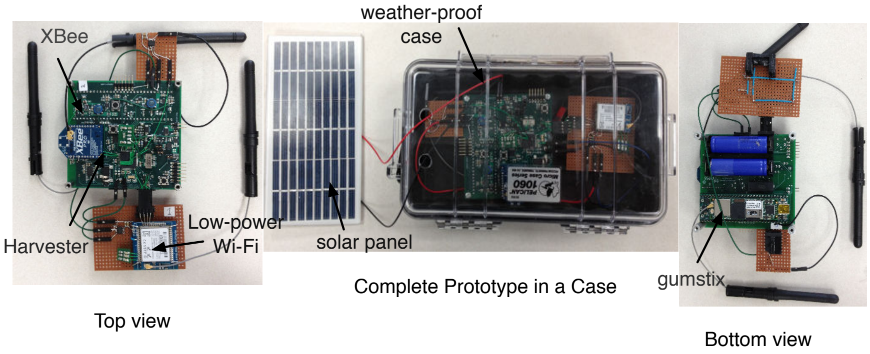

FIG. 4.3. Prototype of the `perpetuu` system. A `perpetuu` node is a multi-tiered system that combines a PIC micro-controller subsystem with a gumstix subsystem. It houses a low-power Wi-Fi radio, a high bandwidth Wi-Fi radio, and a long range XBee radio. The two tiers are powered by a common power supply board that houses the harvester circuit. The harvester circuit charges two Li-ion batteries from a medium-sized solar panel.

4.5 Software Architecture

The software architecture of `perpetuu` is driven by two design principles: (1) Virtualize the resources of the high-power subsystem at the low-power tier to maintain always-on operation and functionality at minimal energy consumption. (2) Utilize computationally intensive and energy hungry components such as the gumstix, Wi-Fi radios, and the XBee radio intelligently, to balance energy harvested from solar panels with energy used by different system services. `perpetuu` meets these design goals using the following modules.

4.5.1 Portable map stack

The goal of a `perpetuu` node is to serve maps to survivors of natural calamities. To this end, `perpetuu` houses a portable and novel custom-designed map stack on the gumstix and a minimal version of the map stack on the PIC micro-controller. The map stack is composed of software layers that supports bidirectional communication between the `perpetuu` node and the survivor's handheld device. The map stack retrieves map layers from a PostgreSQL database augmented with a postGIS extension. The map layers are sent over to a user's smartphone browser where they are merged to generate map tiles using custom written Javascript. The map stack also supports communication from a user's smartphone to a `perpetuu` node—a user can annotate maps with dangerous zones or dangerous points that are sent over to a `perpetuu` node as json-encoded vector layers. The PostgreSQL database stores four types of map layers. The *base* layer consists of raster images (compressed jpgs) for the underlying map of an area. These raster images are generated offline using Mapnik (`http://www.mapnik.org` 2011) and are accompanied by a shape file that stores the mapping of the tile boundaries and the raster images. There are three other adaptive vector layers that are generated dynamically and combined with the base layer to generate the map tiles at a survivor's handheld device: a *safe location* layer that are annotations of where the relief camps are located—first responders can program this layer on a `perpetuu` node; *crowd layers* that include annotations collected from survivors and first responders on changes in the topography of the scene, for example, a fallen bridge or a flooded street; and a *routing layer* that is generated in-situ by the map stack that displays routes from a user's present location to relief camps. The `perpetuu` node runs a mapserver, a webserver (apache), and houses several custom `php` modules that generate map layers adaptively. We next outline three novel components of the map stack that are key to making the `perpetuu` system low-power, adaptive, and usable in an emergency scenario.

Multi-resolution map dissemination: The map stack can disseminate the layers described above adaptively to a user's smartphone. This decision is made based on the residual battery capacity at the node and the predicted energy harvested. For instance, the routing layer alone can take 13 seconds to generate and transmit to a smartphone, while the base layer can take close to 2 seconds to disseminate. The base layer is a set of raster images for the map. However, the base layer does not contain any information on how to route a survivor

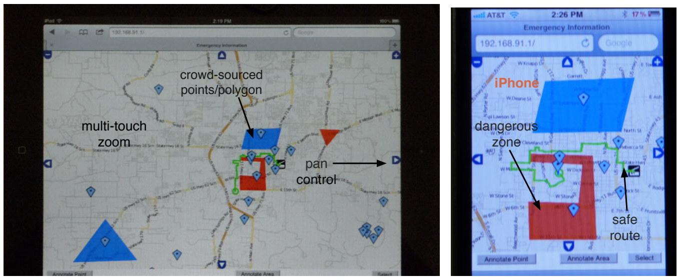

iPAD web browser

iPhone web browser

FIG. 4.4. Maps exported on an iPad and iPhone. The red box (iPhone fig) corresponds to a dangerous zone (e.g., flooded area). The green line corresponds to a route from within the dangerous zone to a relief camp outside the dangerous zone. The markers and areas annotated correspond to the crowd points annotated by users.

from his present location to a safe location. The routing layer generates the routes that can help a survivor find directions from his present location to a safe location. While the route generation is a considerable overhead, in our system we want the user to get this critical information even when the residual battery capacity is low. Hence, if the system is running out of battery, it can choose to not disseminate the base raster images, but present routes as text entries. Therefore, the intuition is to provide the bare minimum critical map information that is useful for the survivor while conserving energy by not disseminating the other layers. In the worst case, the system can disseminate only the base layer without the route information. We do not combine the map layers to create tiles at the `perpetuu` node but combine the tiles at the browser on the survivor's handheld device. This optimization offloads some of the computation to the end-user device.

Efficient route generation: The map stack generates safe paths using road networks to route users from their present location to a safe location. The present location of the user is retrieved using a GPS unit supported by most smartphone and tablets and HTML5 supported by browsers. The route generation module uses input layers (*crowd* layers) and vector layer on road segments to route survivors around potentially dangerous areas. The routing algorithm assigns high weights to road links that intersect with dangerous areas, such as a flooded street. It then uses pgRouting (`http://www.pgrouting.org/`) and the A* search algorithm to determine the shortest and safest route to a safe location. Since links that intersect with dangerous zones have high weights, they are excluded from the generated path. An example of a safe route displayed on an iPhone browser is illustrated in Figure 4.4. The red block denotes a dangerous zone and in the example it is assumed that the user's location lies within this dangerous zone. The generated route, shown as a green path, successfully routes the user from his present location to a safe location around the dangerous zone. Note that the shortest road path from the user's location to the safe location intersects the dangerous zone and is therefore not used.

Crowd layers: The third key innovation in the design of the map stack is the ability for users to annotate a map with points and areas that might be dangerous. For instance, survivors or first responders can annotate maps on their smartphones with areas that denote flooded regions or points that may denote a fallen tree or construction. This user input is stored at the gumstix database as crowd layers (dynamic layers). These crowd layers are periodically synchronized with the PIC, and lazily transmitted to other nodes using the XBee radio at the PIC. It is critical to have this ability to crowdsource information in an emergency scenario since a backhaul connection to a GIS server is non-existent, and the topography of an emergency scene may change quickly. For instance, if a street is flooded it is important that the map store on the `perpetuu` nodes are up-to-date in order to generate reliable directions for survivors. Figure 4.4 denotes example areas and points that are annotated by users as dangerous on their smartphone. Since these layers are dynamically generated at different `perpetuu` nodes, these layers are synchronized continuously between the nodes using the long range radio on the low-power subsystem. Since the crowd layers are vector layers, the data transferred is minimal, and hence the energy overhead is small. However, inter-node synchronization frequency may be an important system parameter. We have not explored how the frequency of inter-node synchronization impacts the system's performance. Our goal in this work is to focus on the performance of a single node. As future work, we plan to thoroughly evaluate the impact of inter-node synchronization frequency. A client device may connect to multiple `perpetuu` nodes if it supports Wi-Fi virtualization, i.e., the ability to connect to multiple access points (perpetuu nodes in this case) using a single Wi-Fi card. However, if it does not, we believe that the inconsistency in the crowd layers across nodes can be made minimal by higher inter-node synchronization frequency.

4.5.2 Storage cache

A major battery drain of the `perpetuu` node is the transmit power of the XBee or the mesh radio. In transmit mode, the XBee radio can consume up to $1W$. The XBee radio is used for transferring dynamic layers like the crowd layers from one `perpetuu` node to all other nodes in the network. To amortize the transmit energy cost of the XBee radio over transfers, we batch and cache files corresponding to the user annotations at the EEPROM on the PIC micro-controller. In essence, we virtualize the gumstix database at the low-power subsystem. This virtualization allows for synchronizing dynamic layers across `perpetuu` nodes without waking up the gumstix. The geospatial information and the crowd layers from the gumstix are copied to the PIC over a UART and are stored at the EEPROM. To extended the lifetime of the EEPROM, we have implemented a storage system that considers wear leveling. The EEPROM is divided into two halves. The first half is used to write data when the second half is being garbage collected. The EEPROM maintains a metadata table, called the table of contents (TOC) that stores the size of files (files store the dynamic crowd layers), the mac address of the XBee module that the data was received from, the size of the file, a crc32 checksum, and the EEPROM address where the data starts. A status byte in the TOC keeps track of whether the data has to be transmitted out of the node using the XBee radio to other nodes in the network or should be merged with the database at the gumstix. To provide wear leveling, new files are appended to the log, and a new entry is made in the table of contents. A garbage collector cleans half of the EEPROM when it is fully utilized by marking entries in the TOC as free.

4.5.3 Wi-Fi sensing

The low-power subsystem in a `perpetuu` node maintains high vigilance at low-power consumption, and wakes up the higher power subsystem when a survivor requests maps from the node. However, for this hierarchical architecture to function with low latency, the low-power subsystem must be able to detect a nearby survivor without waking up the high-power subsystem. To this end, in `perpetuu` we assume that a survivor requesting maps would use a Wi-Fi enabled handheld device to download maps from the node. A `perpetuu` node, therefore, houses two Wi-Fi radios—a low-power Wi-Fi radio that communicates with the PIC micro-controller and is used to sense Wi-Fi devices in the vicinity of the node, and a higher power Wi-Fi radio attached to the gumstix that is used for disseminating maps to survivors on their handhelds. To sense Wi-Fi activity, the low-power Wi-Fi module virtualizes the Wi-Fi module at the gumstix. It exports and advertises the same ESSID as the gumstix Wi-Fi radio. A mobile device carried by a user associates with the Wi-Fi module, and the Wi-Fi sensing service at the PIC determines that a user is in the vicinity by comparing the RSSI (received signal strength indication) in the association channel with a threshold. We have experimentally determined that a threshold of 100 dBm is robust to wireless activity from other devices in the 2.4 GHz spectrum. This simple yet robust method for detecting survivors makes the system completely asynchronous and allows it to save a substantial amount of energy by keeping the gumstix system switched off for long periods of time. If a survivor's device is detected, the `perpetuu` node runs a wakeup controller, described in the next section, to decide whether the gumstix should be awakened to serve maps to the user.

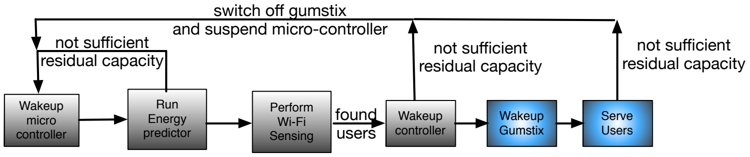

FIG. 4.5. The figure shows a flow of control as decisions are made by the `perpetuu` system to determine whether the Wi-Fi sensing module should be woken up to sense survivors, and whether the higher power gumstix module should be woken up to serve users.

4.5.4 Energy profiler and wakeup controller

The goal of a `perpetuu` node is to last near-perpetually while serving maps to as many users as possible. Near-perpetual operation requires the system to *never* run out of battery charge. So a primary decision that a `perpetuu` node has to make is the following: whether it should wake up the low-power Wi-Fi radio to sense users, and if users are nearby whether it should wake up the gumstix module to serve maps? The decision flow for `perpetuu` is illustrated in Figure 4.5. This decision is based on the residual battery capacity and the likely energy that is going to be harvested in the future. Since `perpetuu` is an online system, it uses a heuristic to make energy predictions for the next hour. It then invokes a wakeup controller that uses the predicted energy, hardware measurements on residual battery capacity, and profiles of energy consumption of different components of the system to determine whether the low-power Wi-Fi module and the gumstix should be awakened. It also determines the period of time, T_{wait} that the gumstix should remain on to serve users. T_{wait} determines the resolution of maps that can be served as well as the number of survivors that can be served with maps. We next outline our energy prediction algorithm.

Energy predictor: The energy prediction algorithm must be implemented on the low-power subsystem. Hence, it must be computationally efficient yet robust. Our energy prediction algorithm extends the Weather Conditioned Moving Average (WCMA) estimation method (Piorno *et al.* 2009), and uses both short term and long term weather conditions to predict the likely energy harvested in the next hour. In our prediction algorithm, we divide a day into 24 time slots, $\{n_0, ..., n_{23}\}$ corresponding to each hour. We keep track of the energy harvested in the last D days, $\{\vec{d_{-D}}, ..., \vec{d_0}\}$, for each of these time slots, where $\vec{d_0}$ is a vector containing the energy harvested in the time slots for the present day. Note that $\vec{d_0}$ will have entries for the first n slots, if we are predicting the energy harvested for the $(n+1)^{\text{th}}$ hour. Our algorithm estimates the energy harvested in the next hour using a combination of two prediction terms weighted using a scalar α:

$$\begin{aligned} E'\left(\vec{d_0}, n+1\right) = \quad & (1-\alpha) && \times \left(\mathrm{R}_1 \times M_D\left(\vec{d_0}, n+1\right)\right) + \\ & \alpha && \times \left(E\left(\vec{d_0}, n\right) + \Delta'\right) \end{aligned}$$

$M_D\left(\vec{d_0}, n+1\right)$ is the mean energy harvested in the past D days in the $(n+1)^{\text{th}}$ hour slot. It is is scaled by a ratio, R_1, of the present-day solar conditions to previous D days' conditions. R_1 is computed as the average energy harvested during the present day's n time slots divided by average energy harvested during the last D days for the same n time slots. The second predictive term is built by adding an expected delta, Δ', to the previous hour's measured energy harvest, $E\left(\vec{d_0}, n\right)$. The predicted delta, Δ' , is set equal to $M_D(\vec{d_0}, n+1) - M_D(\vec{d_0}, n)$. The second term is important since it takes into account recent observations of incident light. The expected delta, Δ' , captures the expected short-term trend based on the time of day. For instance, during sunrise, the previous hour energy consumption would always be smaller than the present hour energy harvested, and using the last hour energy values to estimate the next hour energy harvested can be erroneous. Thus the total energy predictor, $E'\left(\vec{d_0}, n+1\right)$, utilizes recent observations of solar irradiance in the second term and longer-term observations in the first term. From our experiments, we found that $\alpha = 0.4$ works best for weighting the two terms. Note that while we focus on predicting the energy harvested in the next hour, the scheme can be easily extended to predicting the energy harvested in the next 6, 12, or 24 hours.

Wakeup controller: The wakeup controller determines, based on the predicted energy harvested and the residual battery capacity, whether the system should wake up the low-power Wi-Fi module to sense users and the gumstix to serve clients. The goal of the wakeup controller is to avoid an *empty* battery. However, since the wakeup controller executes in online mode, it tries to avoid an empty battery in the time horizon of an hour. Note that the wakeup controller is executed every time the PIC micro-controller wakes up from sleep (at a frequency of once every 5 minutes during the day and once every 10 minutes during the night), hence if the decisions made by the wakeup controller are accurate, it can assure that the system does not *ever* encounter an empty battery. The two wakeup intervals during the day and night was determined experimentally to work best for the arrival rates and distribution used in this work. A battery is considered empty in a `perpetuu` node when the residual battery capacity is less than 20%. This is because in Li-ion batteries the amount of time taken to charge a battery from 0 to 20% (slow charging) is higher than above 20% (fast charging). Moreover, deep recharges of the Li-ion battery can minimize battery lifetime.

The wakeup controller maintains a profile of the average energy consumed by different components of the system using current accumulators: energy consumed by the gumstix node and map dissemination (E_g) at different resolutions, energy consumed by the gumstix bootup (E_b), energy consumed by the Wi-Fi sensing module at the low-power subsystem (P_{Wifi}), and the idle power consumption of the power supply board and the PIC (P_I). The algorithm uses the predicted energy harvested in the next hour, E_h, and first determines whether the energy harvested alone is sufficient to wakeup the Wi-Fi sensing module and the gumstix. Note that detecting a client is useless if the system lacks sufficient energy to wakeup the gumstix to serve maps. If the predicted energy harvested alone is sufficient, the system wakes up the Wi-Fi sensing module to look for users (our harvester can power the system directly from the solar panels without using the batteries). If the predicted energy harvested is insufficient, the wakeup controller determines if the residual battery capacity and the energy harvested together is sufficient to wakeup the Wi-Fi sensing module and the gumstix. Consequently, if the Wi-Fi sensing module senses a Wi-Fi user nearby, it wakes up the gumstix to serve maps. Another decision made by the system is to determine the amount of time that the gumstix module should remain on while serving users. The resolution of the map served as well as the number of users served is a function of this *wait* time. If more than one user is served by a single wakeup of the gumstix, the bootup energy consumed by the gumstix is amortized over a larger number of survivors. `perpetuu` always uses an upper bound on the *wait time* and the wait time never exceeds this maximum, and is a fraction of the maximum wait time. This fraction is proportional to the battery charge left in the system divided by the maximum battery capacity. The reasoning behind using this fraction follows: the time the system should keep the gumstix on to perform useful work should be proportional to the residual energy reserve of the system. If the micro-controller infers that its predictions were incorrect and the system will run out of energy it switches off the high-power devices, and transitions to a sleep mode.

4.6 System Implementation

We have designed and implemented a fully functional `perpetuu` prototype that can serve maps to users on their handhelds. For our prototype, we use a PIC18F micro-controller with a 4KB of RAM and 128 KB of flash and a gumstix module that houses an ARM Cortex processor with 512MB of RAM. The flash storage on the gumstix is augmented with a 4GB SD card to store the PostgreSQL database. We fabricated a custom

4-layer printed circuit board that houses the PIC, the harvester circuits, and the power supply with the measurement and power gating circuits. The gumstix is plugged into the PCB. The entire system is powered by a solar panel with a short circuit current rating of 320mA, an open circuit voltage rating of 17V (the nominal values are 240mA and 15V respectively), and a 2 Ah Li-ion battery. The size of the solar panel is 210mm by 184mm and the maximum power harvested from the panel is close to 3.5W. The solar panel is rated at a conversion efficiency of 10%. The software services resident at the micro-controller are light weight and work efficiently under the memory and storage constraints of the micro-controller. Figure 4.3 illustrates our prototype and the maps viewable using off-the-shelf browsers on an iPhone and iPad. A `perpetuu` node advertises an ESSID over Wi-Fi, and a user's device connects to a `perpetuu` node assuming that it is a wireless access point. The end-user can then use a browser on the handheld device and type any `url`, and the DNS server on the gumstix redirects the url request to localhost, and map layers are downloaded on the browser of the handheld device. We used an 802.11b Wi-Fi radio for our implementation. Each mesh node can store maps for an entire city. Although the Wi-Fi range of the nodes are short (access link), users can download maps for the entire or part of the city from a node. These maps can help the user route himself from a specific location to a relief camp, for instance. The maps can be viewed on any handheld device that has a web browser and an outgoing Wi-Fi interface. The end-user *does not need to install any custom application*. The maps provide a Google-map like UI, and the ability for users to annotate the map with dangerous points and dangerous areas (illustrated in Figure 4.4).

4.7 Summary

Presented in this thesis is a solar powered GIS microserver node that uses a multi-tiered hardware and software architecture for serving information-rich maps to survivors of natural disasters. `perpetuu` combines a computationally weak but low-power microcontroller device with a computationally capable gumstix platform to explore tradeoffs between energy consumed, system availability, and computational capability. Through the use of energy prediction, an intelligent wakeup controller, and a virtualized software architecture, the system can balance energy harvested from solar panels with energy consumed by different subsystems. A `perpetuu` node can serve maps near-perpetually while serving a large number of survivors on browsers of common handheld devices using a custom designed map stack.

Chapter 5

VIRTUALIZING SENSORS: POWER CONSUMPTION REDUCTION BASED ON CONTEXT

The well-recognized importance of adequate restful sleep for health and the high prevalence of disturbed sleep (Institute of Medicine 2006) has produced wide-spread recognition of the clinical need and commercial potential for a practical cost-effective system to evaluate and support restful sleep at home. The large number of existing at-home sleep recording system (MNaHS) devices attests to both the need and commercial potential but also the failure of any to meet the need. The common sleep problems include poor sleep habits, various idiosyncratic behaviors disrupting sleep, insomnia, circadian rhythms, undiagnosed restless legs syndrome, low iron status and sleep apneas. Once identified, many can be managed through guided self-help while others need professional input. Evaluation of most requires repeated *multi-night* recordings given circadian rhythm effects, marked variability of events disturbing sleep (Marie Trotti *et al.* 2009), and desired intervention evaluation. The prevalence and significance of the problem require efforts to get wide public and consumer acceptance of using a system that could accurately identify significant sleep features, support appropriate self-help and provide triage for professional evaluation.

A wide range of MNaHS have been developed from bed pads to sound systems (Nandakumar, Gollakota, & Watson 2015). None of these systems aside from wrist worn moni-

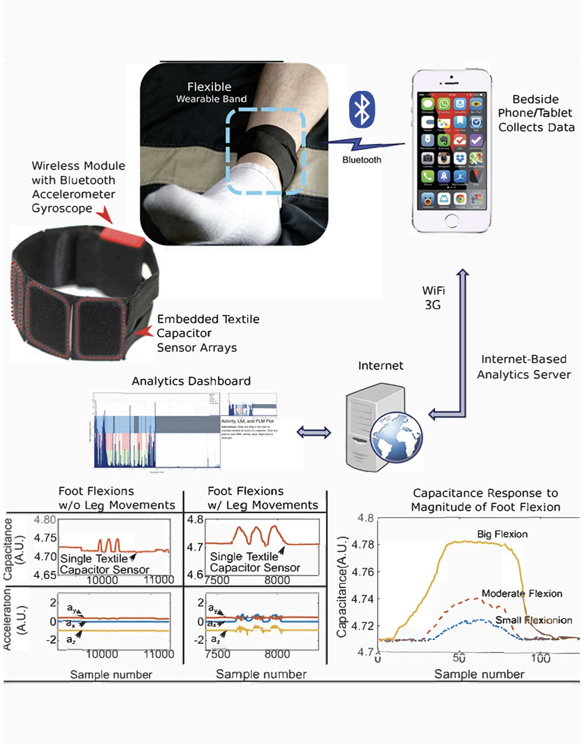

FIG. 5.1. Prototype RestEaZe system. The system can distinguish foot flexions from other general leg movements. The system can also measure the amplitude of the foot flexions.

tors have gained general commercial success for the following reasons: discomfort, sleep-room constraints, and limited accuracy (M Kelly, Strecker, & Bianchi 2012). Only wrist worn monitors have achieved wide commercial acceptance. These monitors, however, focus on “fitness for the fit” with sleep as an afterthought. Moreover, the data from existing commercial and professional wrist-worn systems are valid for group analyses but have very weak and sometimes not significant correlations across subjects, indicating that there is limited value from these existing systems in obtaining significant measures for an individual. Moreover, the user could obtain essentially the same or better information from a sleep log. The claims that these monitors could identify sleep depth or stages have been discredited (Mantua, Gravel, & Spencer 2016).

5.1 RestEaZe: Low-power Accurate Sleep Monitoring Using a Wearable Multi-sensor Ankle Band

To address the above limitations we present RestEaZe, an ankle-worn band that captures leg movements during sleep. Recording from the arm, while convenient, is inaccurate since it fails to identify features critical for sleep such as brief arousals. Moreover, arm movements are limited in drowsy waking before or during sleep and rarely show even basic sleep related movements (Gabelia *et al.* 2014). In contrast, leg movements (LM) occur commonly in sleep providing a rich panoply of characteristics to be analyzed in relation to sleep biology and clinical sleep status. This is not surprising given the relevant neurobiology involving motor control in sleep, e.g. (1) active inhibition or reduction of LM during quiet restful sleep, (2) preparatory activation of LM with arousal/waking anticipating possible movement, and (3) activation or disinhibition of some sleep-related simple, repetitive movements primarily foot dorsiflexion noted as periodic limb movements in sleep (PLMS) (Picchietti *et al.* 2007).

The goal of RestEaZe is to use leg movements to characterize sleep in the home at a cost lower than sleep labs. The RestEaZe system is illustrated in Figure 5.1. The ankle band provides measurements from four separate sensors: 3-D accelerometer, gyroscope, magnetometer, and an array of textile capacitive sensors. These sensors allow RestEaZe to use machine learning algorithms to calculate leg orientation and characterize leg movements. Figure 5.1 shows that the capacitance sensor detects foot dorsiflexion and can distinguish them from other complex leg movements using the inertial sensors. These two types of leg movements can then be used to calculate measures unique to the leg and also those previously determined from the arm such as sleep efficiency, sleep fragmentation, and brief arousals. For the system to be useful in the home and for large cohort studies, however, it is important that the ankle band can capture the relevant data at low-power consumption and can last on a single charge for several days. This work focuses on the sensor system and algorithms designed to make the system low-power and usable.

Research contributions: The design, implementation, and evaluation of the RestEaZe system presents the following novel research contributions.

- **Multi-Sensor Ankle Band for Monitoring Sleep:** We present the use of flexible textile-based capacitive sensor arrays and inertial sensors built into ankle bands for sleep quality analysis. Our sensors can detect two types of movements using a single point measurement system: (1) Foot flexions whose periodicity can be used to help diagnose restless leg syndrome, ADHD in children, and insomnia (L Marcus *et al.* 2014); and (2) Complex Leg Movements that can be used to determine texture of sleep (e.g, brief arousals). We have developed analytics that use these movements to determine restful/restless sleep, brief arousals and sleep fragmentation, and sleep efficiency. The focus of this work, however, is not on the development of the sleep measures but on the design of a low-power system development for longitudinal sleep monitoring studies in a home setting.

- **Hierarchical Sensing for low-power consumption:** We present a system that uses a machine learning algorithm on data from the low-power capacitive sensor arrays to predict when relevant movements are likely to occur, and then wakeup other energy hungry sensors such as the gyroscope, accelerometer, and magnetometer. The algorithm is low overhead and can be implemented on the micro-controller system with relevant optimizations.

- **Prototype development and evaluation:** We have prototyped a fully functional RestEaZe ankle band. We have performed a thorough evaluation both in the context of a sleep lab (on four subjects) and in-home settings on eight subjects. We have obtained relevant institutional review board (IRB) approval to perform these studies.

In our evaluation, we show that the RestEaZe band can accurately capture relevant sleep measures while lasting for more than 5 days on a single battery charge.

5.2 Need for RestEaZe

In this work, we present a novel enabling system called RestEaZe that combines the following two components: (1) textile-based non-rigid capacitive sensors arrays for measuring foot flexions (PLMS), a 9-axis IMU for measuring leg movements (CLMS), and leg position, built into a small, non-intrusive, low-cost, wireless ankle-worn unit. This system, rather than recording the standard PLMS leg muscle activity using EMG, records the very specific effect of that muscle activity, i.e. foot flexion. The system, illustrated in Figure 5.1, is the only single point measurement system that can distinguish foot flexion from other leg movements; (2) an analytic engine that uses the measured foot flexions, leg movements, and leg positions to calculate reliable PLMS and CLMS indices and features, described in the next sections. The system has an optional backend providing sophisticated analytics for relating movement features with medical conditions. This backend provides sleep indices, activity visualizations and logs for the consumer. The smartphone and ankle band unit allow the user to evaluate and possibly modify his restless sleep and optionally discuss results with his doctor. The same system with extended backend support allow the doctor to determine the nature of restless sleep, its severity and possible associated medical condition.

Another key attribute of our system (which is not part of this dissertation) is the inference of brief arousals during sleep which is a primary assessment for restful sleep producing alertness the next day. Sleep assessment examining restful sleep and health outcomes had initially focused on total sleep time. As sleep medicine advanced our understanding of the texture of sleep it became apparent that for restful, healthy sleep, the continuity of

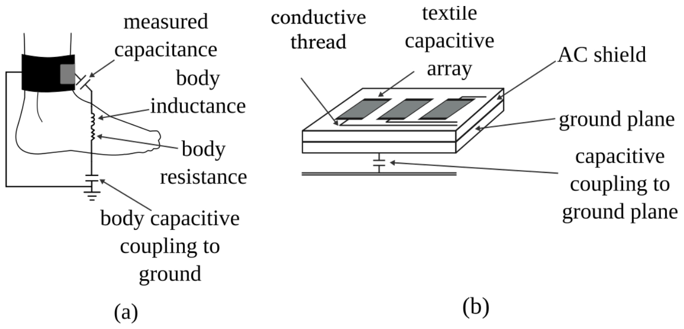

FIG. 5.2. (a) The figure illustrates the equivalent electrical circuit which allows capacitive measurements. The user is coupled into the ground which allows for more sensitive measurement of motion by the capacitive sensors. (b) The figure shows the construction of the sensor array. The bottom layer allows for coupling the user into the system ground. The middle layer shields the capacitive sensors on the top layer from the user's body directly underneath the sensors. Conductive threads connect the various layers to the capacitance sensing IC.

sleep (uninterrupted sleep) was almost as important as total sleep time. An excessive number of brief micro-arousals fragmenting sleep even without reduction in total sleep time produces decreased alertness (Stepanski *et al.* 1987; E. Martin *et al.* 2000; E Martin *et al.* 1996), impaired mood and executive function, and mental flexibility (E. Martin *et al.* 2000; E Martin *et al.* 1996). Similar results have been reported from animal studies showing frequent brief arousals without loss of total sleep time impaired maze learning (Tartar *et al.* 2005). Measuring wake and brief arousals disrupting sleep is seen as important as measuring total sleep time for evaluation of restful sleep. The rich panoply of sleep LM characteristics associated with arousal provide a putative measure of arousal and waking.

5.3 Hardware System Description

For capturing sleep quality measures, we have developed a fully functional custom sensing platform called RestEaZe. It is comprised of a multi-sensor ankle band which is depicted in Figure 5.4. The band is intended as a single point measurement system for

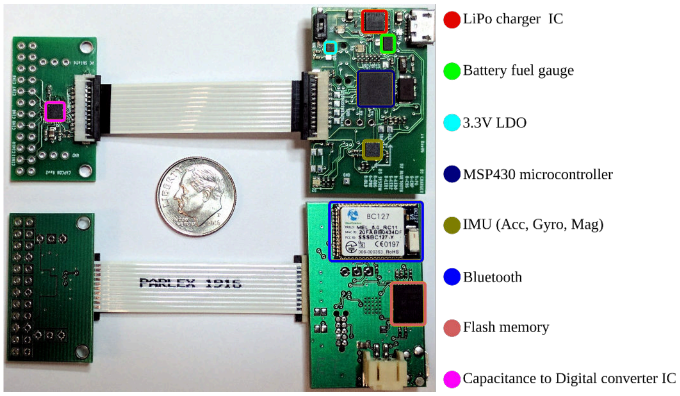

FIG. 5.3. Capsense system consists of an MSP430FR5969 micro-controller, BC127 Bluetooth 4.0/BLE module, N25Q 512 MBit flash memory, MPU9250 accelerometer/gyroscope/magnetometer (IMU), and battery charging/monitoring chips. The daughter board contains the AD7147 capacitance to digital converter chip.

capturing leg movements during sleep, and captures other attributes such as sleep position. The band has three textile capacitive sensors sewn in and are connected to a custom embedded system motherboard called Capsense using embroidered conductive fibers. The Capsense board houses a MSP430 micro-controller, inertial measurement unit (IMU) with accelerometer, gyroscope and magnetometer, Bluetooth module, flash memory and various battery charging/monitoring chips. This system connects with a flat-flex connector to a daughter board which contains the capacitance measurement circuit. The capacitive plates, made of conductive fabric, are sewn onto a non-conductive pleather band and connected to a routing layer. In order to shield the routing layer, an AC shield layer made of conductive fabric sits in between the capacitive sensor array and the routing layer. Finally, the ground layer sits beneath the routing layer. Sheets of non-conductive satin cloth are used to prevent any of the conductive layers from electrical shorting. The conductive thread network on the routing layer is made of silver-plated thread and is stitched by an embroidery machine. The embroidered traces connect to the Capsense daughter board by hand sewing. Because of wear and tear, conductive thread/layers fray and can form undesirable micro-connections.

This construction minimizes the chance of this happening thereby increasing the reliability of capacitive measurements.

The Capsense system employs a modular and sewable friendly form factor, with a hardware architecture that supports low-power sleep monitoring. The modular overview of the Capsense system hardware is illustrated in Figure 5.3. The whole system is powered by one rechargeable 400 mAh lithium polymer battery which can be charged from the on-board USB connector. When the system is connected to a charger, the charge management IC switches the system power source from the battery to USB. In this way, the system can operate normally while the battery is charging. The RestEaZe system can operate in two modes: (1) serial streaming mode; and (2) flash storage mode.

Streaming mode: In the serial streaming operating mode, the system can collect data for 20 hours on a full charge. Another device, such as a PC or smart phone has to maintain a Bluetooth connection to the ankle band to collect data streamed from the band. In addition to data packets, the band also periodically sends status packets that contain information such as battery and sensor status.

Flash mode: When the system operates in flash storage mode, instead of sending data packets over Bluetooth, the system writes data to flash memory. At the default sampling rate of 25 Hz, up to 14.6 hours of data can be stored on the 512 Mbit chip without compression. At the end of a data collection session, a `STOP` command followed by a `START_DUMP` puts the system into a mode in which data is read from the flash, assembled into data packets and sent over Bluetooth to the PC or smart phone. In addition to data packets, flash status packets are sent to help the end device track the progress of the flash dump. For example, a night's worth of data (typically eight hours), will take close to 30 minutes to offload from the band.

Capacitive sensing: A unique component of our system is an array of textile capacitive

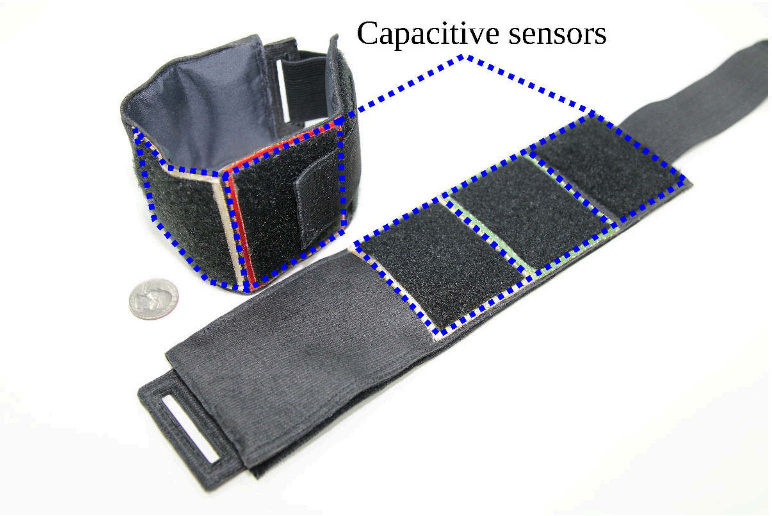

FIG. 5.4. The RestEaZe ankle band. The band comprises textile capacitive sensors built into the ankle band and a data collection board which houses a IMU, Bluetooth module, micro-controller, and charging circuitry.

sensors built into the ankle band. This sensing technique is based on the principle of parallel plate capacitor, where the conductive fabric and the user's body serve as two plates of a virtual capacitor (see Figure 5.2 for the equivalent electrical circuit for our system). Change in distance between the capacitive sensors and the user's body produces a change in capacitance that can be used to classify the type of motion or serve as a wakeup trigger. The capacitive sensors (and the sensing circuit) is ultra-low-power (consuming few micro-watts of power) and can capture any movement in the vicinity of the plates. Hence, they are used to capture both periodic leg movements and complex leg movements and can act as a low-power wakeup sensor for the more energy expensive IMU. Therefore, the capacitors in conjunction with the IMU behave as a hierarchy of sensors and the power consumption of the system can be optimized by keeping the higher power sensor asleep most of the time. The capacitive sensors, on the other hand, are active all the time and whenever they detect a movement, the control logic wakes up the IMU for data collection from the capacitive sensors and the IMU.

Module	Sleep power	Active power
micro-controller	$<$1 mW	$<$1mW
flash	110 μW	11.6 mW
capacitive sensor	7.0 μW	3.5 mW
Bluetooth	19.5 mW	51.0 mW idle, 57.9 mW TX
IMU	30 μW	26.2 mW

Table 5.1. Power consumption of individual modules that constitute the Capsense platform.

External synchronization: When collecting data in a sleep lab, we have to be able to synchronize our data with data produced by the sleep lab's polysomnograph (PSG) hardware. RestEaze outputs a pseudo-random binary waveform to a GPIO port. This port is then connected to a custom PCB board which contains a voltage divider and snaps which are compatible with standard electromyogram (EMG) electrodes. This pseudo-random pattern is recorded by PSG and the band so that we can synchronize the two systems.

Practical implementation challenges: Capacitive sensors are sensitive to various sources of noise. Additionally, interfacing a PCB to fabric has been a challenge. We tried several ways of connecting the capacitive sensors to the capacitance-to-digital (CDC) conversion IC. Originally, the CDC chip was on the main compute board that clamped onto cloth backing that had sewn conductive thread terminals. Over time, the clamps would tear up the conductive thread and the connection between the CDC and sensors would break. Also, noise from other components on the PCB affected capacitive sensor readings. In the current iteration, we designed a small and easily sewable daughter board which contains the CDC chip, and breakout through-holes for connecting capacitive channels to sensors by hand sewing the daughter board to the embroidered traces. The Capsense daughter board contains the CDC chip and breaks out 6 capacitive channels. This has been a rugged solution with some bands reliably collecting over a dozen nights of data (see Figure 5.5).

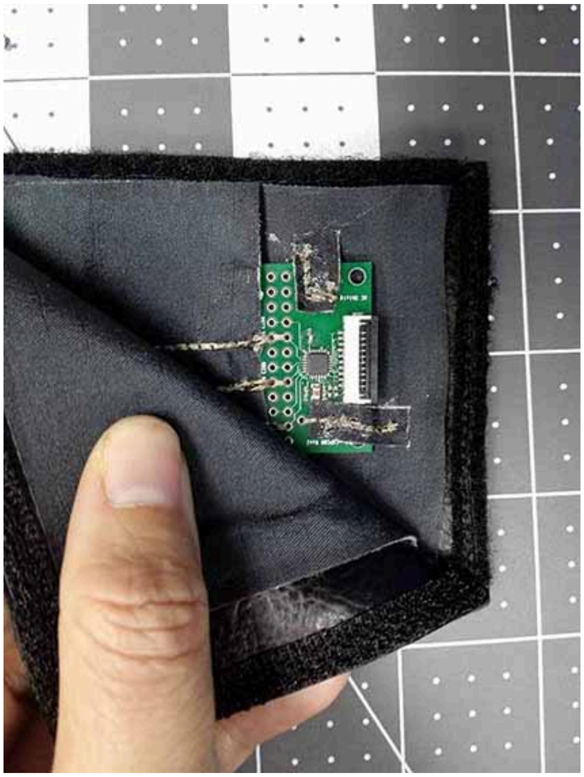

FIG. 5.5. Conductive threads connected to ground, shield and capacitive sensors connect to Capsense daughter board in the RestEaZe band.

5.4 Software System Description

We performed several software optimizations at the RestEaZe band to collect data from multiple sensors in a synchronized fashion.

Sensor reading: The IMU has an internal sampling rate of 1 kHz for the accelerometer/gyro and we sample this sensor at 25 Hz. The Capacitance-to-Digital conversion IC (CDC) as an internal sampling rate of 250 kHz and we also sample this sensor at 25 Hz. Since the sensors are sampling asynchronously, we generate a system tick every 40 ms using a combination of a built-in real time clock and timer to sample the CDC. We use a 32-kHz crystal as the reference clock for these circuits. The crystal we are using has a frequency tolerance of 20 ppm, so the worst case drift is 1.152 sec over 8 hours. The IMU is configured to produce a hardware interrupt, and we record the time of this event relative to the last system tick using capture/compare hardware built into the micro-controller.

Configuration: The system is designed to be easily configured by a text based command system. The user can connect to the system over Bluetooth and submit commands to reconfigure the platform. We leverage the MSP430 FRAM capability for persistent memory storage to store configuration information that the system uses for initialization at boot time. This flexibility allows the platform to be used for other applications.

Low power operation: The default system configuration is aimed at data collection for analysis and does not leverage low-power features of the various sensors and submodules. However, the system can be configured to put the Bluetooth, IMU, CDC and flash memory into low-power modes. A summary of energy consumption between different operating modes can be found in Figure 5.9. A hierarchical sensing algorithm we developed takes advantage of this to extend battery life.

5.5 Hierarchical sensing

We use a hierarchical sensing system to optimize the power consumption of the RestEaZe band. As shown in Table 5.1, the capacitive sensor consume an order of magnitude lower power in active mode than the IMU and 20x less power than the Bluetooth module. Hence, if it were possible to detect a onset of a movement using just the capacitive sensor array, the IMU and the Bluetooth module could be kept in low-power mode, and only awakened when an activity of importance is likely to occur. Therefore, the key challenge is to detect the onset of a complex leg movement or a periodic leg movement using the capacitors alone. Moreover, the algorithm needs to be light-weight so that it can be implemented on a micro-controller platform. To this end, we design a simple machine learning algorithm based on Support Vector Machines to predict the occurrence or absence

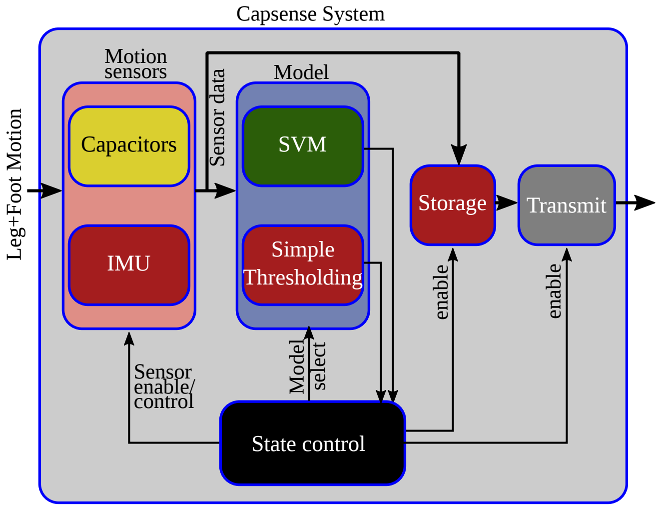

FIG. 5.6. The Capsense system uses an SVM model to detection leg motion. The SVM model determines when to wake up the IMU for leg movement detection.

of a movement using the capacitor data.

5.5.1 Machine Learning for Detecting Movement Onset

The key element in our machine learning algorithm is to determine an appropriate set of features on the capacitance data that can be used to detect a movement. We empirically determined the appropriate feature and kernel (see Figure 5.7 and Table 5.5) for our machine learning algorithm.

For our experimentation we used data collected on two pediatric patients, PED1 and PED2 (see Table 5.6). The movements in this data set was labeled by a trained technician. The ground truth for the complex leg movements was video data while the ground truth for the periodic leg movements (dorsiflexion) was EMG on the anterior tibialis. The unlabeled regions of the data were randomly partitioned and chosen as time periods with “no motion” label. This resulted in a total of 1508 distinct training data points. The predictor was tested against data from two patients, RZ0015 and RZ0016, that were not used to train the model (see Table 5.6). It successfully classified movements, and in combination with a data reconstruction algorithm the restored signals had very high correlation to the original signals. An example of this can be seen in Figure 5.8.

We then calculate several statistical features such as mean amplitude, standard deviation of amplitude, skew, kurtosis and root-mean-square amplitude in a window. Once the features were calculated, an SVM model was trained for each combination of features. The accuracy of the models for each combination can be seen in Table 5.2. Although the combination of standard deviation and kurtosis yielded the highest score in this model, since the predictor had to be implementable on a micro-controller, we have to be cognizant of other factors such as support vector size, computational complexity of the predictor and the computational complexity of calculating the feature. Hence, we chose the top feature, standard deviation to train the SVM.

In addition to the feature, the kernel used for the classification needs to be determined. Both RBF and linear kernels were tested along with varying C and gamma parameters using a grid search. The model is trained on the full development set. The scores are computed on the full evaluation set. The hyper-parameters tuned for highest precision can be found in Table 5.3. The hyper-parameters tuned for highest recall can be found in Table 5.4. The linear kernel with $C = 1000$ was chosen since it struck a good balance between accuracy and implementation complexity. Additionally, the size of the support vector was tuned by adjusting the number of features used to fit the model.

Implementation: Capacitive measurements are inserted into a six element ring buffer. Since the smallest movement we want to measure is 0.5 seconds, our sample window is half of that. Every time the ring buffer fills up, the standard deviation of the measurements is calculated and submitted to the predictor. If the predictor classifies the feature as a movement the system is switched into the mode where the IMU is turned on and sensor measurements are written to flash. While the system is in the high-power mode the

Table 5.2. Comparison of models calculated from different combinations of features: {kernel: ’linear’, C: 1000}

Avg	StdDev	Skew	Kurtosis	RMS	TP	TN	Accuracy
F	T	T	F	F	0.85	0.86	85.67
F	T	F	F	T	0.93	0.78	85.33
F	T	F	F	F	0.82	0.88	85.27
F	T	T	T	F	0.83	0.88	85.14
F	T	T	T	T	0.90	0.78	84.16
F	T	F	T	F	0.78	0.90	83.87
T	T	F	T	F	0.90	0.77	83.59
T	T	F	F	F	0.92	0.75	83.56
F	T	F	T	T	0.88	0.77	82.29
T	T	T	F	F	0.94	0.71	82.22
T	T	T	T	F	0.96	0.68	81.35
F	T	T	F	T	0.87	0.75	81.06
T	T	T	T	T	0.92	0.69	80.47
T	T	F	T	T	0.93	0.66	79.47
T	T	T	F	T	0.94	0.64	78.85
T	T	F	F	T	0.94	0.62	77.90
F	F	F	T	F	0.97	0.20	57.14
T	F	T	T	F	0.72	0.43	56.34
T	F	T	F	T	0.24	0.88	56.10
F	F	F	F	T	0.22	0.88	55.92
F	F	T	T	T	0.82	0.31	55.89
F	F	T	T	F	0.97	0.15	54.82
T	F	T	T	T	0.61	0.48	54.77
F	F	F	T	T	0.64	0.45	54.04
T	F	F	T	F	0.48	0.59	53.57
T	F	T	F	F	0.20	0.85	52.75
T	F	F	F	F	0.68	0.38	52.68
T	F	F	T	T	0.50	0.55	52.64
F	F	T	F	F	0.00	1.00	51.00
F	F	T	F	T	0.23	0.77	50.78
T	F	F	F	T	0.40	0.61	50.61

Table 5.3. Best parameters tuned for precision: {kernel: ’linear’, C: 1}

class	precision	recall	f1-score	support
no	0.91	0.96	0.93	2018
yes	0.72	0.53	0.61	401
avg / total	0.88	0.89	0.88	2419

Table 5.4. Best parameters tuned for recall: {kernel: ’rbf’, C: 100, gamma: 0.001}

class	precision	recall	f1-score	support
no	0.95	0.94	0.94	2018
yes	0.69	0.74	0.72	401
avg / total	0.91	0.90	0.90	2419

Table 5.5. Implemented SVM: {kernel: ’linear’, C: 1000}

class	precision	recall	f1-score	support
no	0.84	0.84	0.84	439
yes	0.85	0.85	0.85	467
avg / total	0.84	0.84	0.84	906

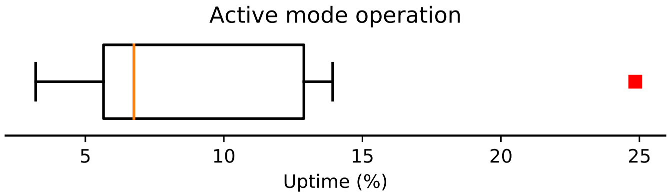

FIG. 5.7. Percentage of time spent in active mode for two subjects (0,1) of nine nights total.

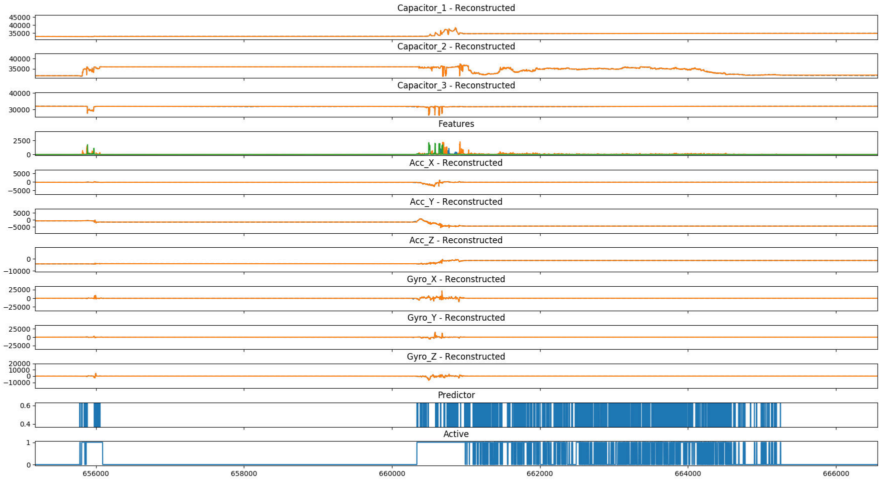

FIG. 5.8. Using motion predictor on RZ0015 data to generate a synthetic reconstruction of right leg data with a correlation of 0.99 for each IMU axis. In this instance, the IMU would be active 13.56 % of the total time.

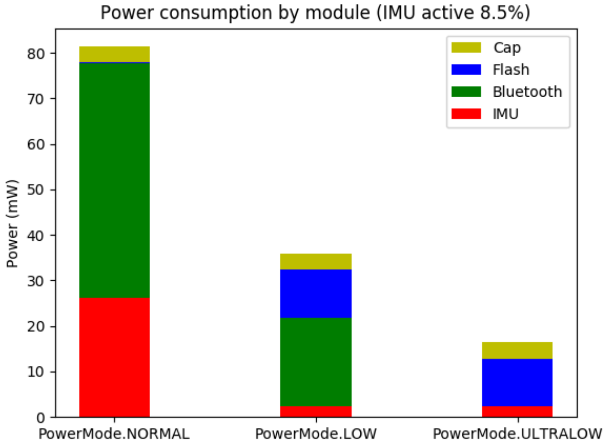

FIG. 5.9. Power consumption breakdown by different system modules in different modes of operation. PowerMode.NORMAL: IMU is always on, using Bluetooth to transmit data, conventional operation. PowerMode.LOW: IMU is idled and woken up by machine learning predictor, Bluetooth is in sleep mode, flash is used to store data. PowerMode.ULTRALOW: IMU idled and woken up by machine learning predictor, Bluetooth is turned off using a relay, flash is used to store data.

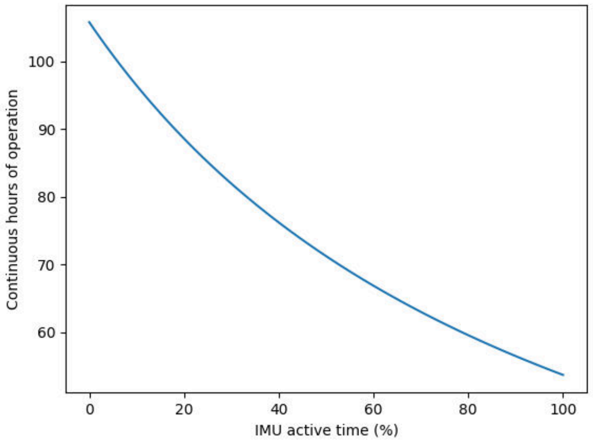

FIG. 5.10. When operating in ultra-low-power mode, the bands estimated total continuous operation decreases as the IMU is turned on more often.

accelerometer values are inserted into a six element buffer, and when the buffer is full the standard deviation of the accelerometer values is calculated. If the standard deviation is above an experimentally derived threshold, the system continues in high-power mode. The system architecture is summarized in Figure 5.6.

The support vector is 93 elements long and the predictor was implemented in C. Because the micro-controller did not contain a floating point unit, the calculations were performed using a fixed point math library provided by Texas Instruments to execute the predictor within the 40-ms window set by the required sampling rate. Algorithm 1 describes the algorithm implemented on the micro-controller.

ALGORITHM 1: Main loop for low-power operation

```
$IMU\_on$ = $FALSE$;
repeat
    if ($IMU\_on$ == $TRUE$) then
        readIMU();
    end
    readCDC();
    ring_buffer_insert_CDC_data();
    if ($IMU\_on$ == $TRUE$) then
        ring_buffer_insert_IMU_data();
        if (CDC_buffer_is_full() == $TRUE$) then
            acc_std_dev = calc_acc_std_dev();
            if ($acc\_std\_dev < 10.0$) then
                $IMU\_on$ = $FALSE$;
                power_imu_off();
            end
        end
    else
        cap_std_dev = calc_cap_std_dev();
        if (CDC_buffer_is_full() == $TRUE$) then
            $motion$ = predict_motion();
            if ($motion == TRUE$ and $cap\_std\_dev > 1.0$) then
                $IMU\_on$ = $TRUE$;
                power_imu_on();
            end
        else
            $IMU\_on$ = $FALSE$;
        end
    end
until $TRUE$;
```

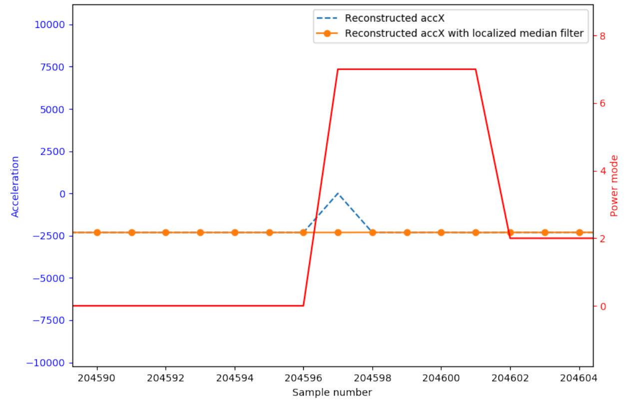

FIG. 5.11. During transitions from low-power to high-power modes, the accelerometer sometimes produces spurious data. A localized 5 point median filter is used to remove these spikes.

Data reconstruction: Since the data from the IMU is collected intermittently and our sleep quality analysis algorithm requires a uniformly sampled data stream, the sensor data must be processed to conform to this. The reconstruction algorithm is *sample and hold* with the accelerometer data filtered with a localized median filter with a window size of 5 (see Figure 5.11 for an example). The filter is used when the system transitions from low to high-power state since sometimes the accelerometer data exhibits spurious spikes. The gyroscope, however, did not show spikes, so only sample and hold is used. Over the course of five nights, a band with the regular firmware and a band with the low-power firmware were worn on the right leg. The data from the low-power band the the regular band was analyzed through our sleep quality algorithm and matched closely.

Smart phone application: For this system to be useful for non-technical users, we developed an Android application. It allows a user to pair with the bands, calibrate sensors, and start/stop data collection with minimal instruction. The main screen of the application is shown in Figure 5.12. Battery meters for each band serve to alert the user when they are on low battery. The application can be configured to switch the band between streaming

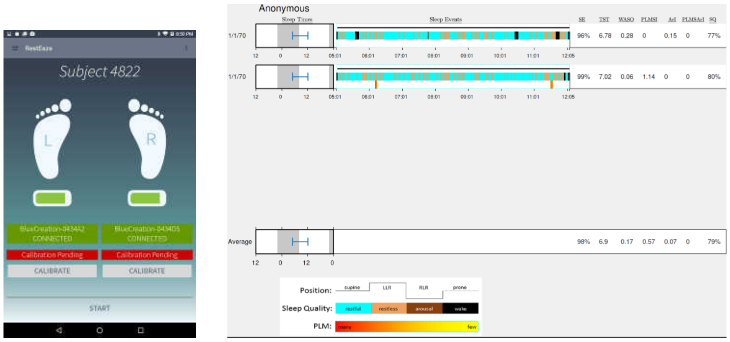

(a) (b)

FIG. 5.12. (a) The figure shows the home screen of our Android app that interfaces with the Capsense system. (b) The figure shows an output from a subject as an output of the backend analytics.

and flash storage modes. Finally, the application supports storing the collected data locally or uploading data to a backend server. These modalities of operation cover several use cases necessary for sleep research. For long term, in-home, studies to be successful it is necessary for the system to function without the guarantee of Internet access and without the possibility of assistance from a technician. Figure 5.12 (a) shows a screen shot of the smartphone application.

Extracting sleep measures: The multi-sensor data is analyzed by our custom written data analysis software to extract sleep quality measures such as the number of complex leg movements, number of periodic leg movements, number of arousals during the time, wake after sleep onset duration, and sleep efficiency. The analytics underlying the calculation of these metrics is beyond the scope of this dissertation, but we have verified the accuracy of these metrics with standard PSG (polysomnograph) tagging. Figure 5.12 (b) shows an example report generated by our backend analytics.

5.6 Evaluation

The RestEaZe system comprise the ankle bands, a smartphone application, and back-end analytics to calculate sleep quality metrics. These metrics can be used to assess the quality of sleep, texture of sleep, and can be used as an aid for diagnostics of sleep related disorders. The goal of this work is to present the hardware and software architecture of the RestEaZe system and present algorithms that make the system low-power and usable. Hence, our evaluation focuses on the following key questions.

- What is the average power consumption of RestEaZe when applied on a wide variety of subjects, i.e., subjects with sleep-related disorders, and healthy subjects that may or may not experience disrupted and fragmented sleep?
- Can the RestEaZe system, with its low-power consumption, provide the same quality of data as an always-on RestEaZe system?

5.6.1 Experimental Setup

To answer the above questions, we performed our system evaluation in two phases. In the first phase, we used RestEaZe in a sleep lab in conjunction with polysomnograph (PSG). In a PSG session, we collected a variety of other physiological data on a subject such as EEG, EKG, respiration rate, blood oxygen level, and EMG for muscle groups on the leg. Additionally, we recorded the entire sleep session using an IR camera. The data from the RestEaZe band is tightly synchronized with data from the proprietary PSG system using a pseudo-random sequence generated by the band which is fed into the PSG data collection module using an auxiliary input port. For the sleep lab studies, we recruited four subjects, 2 adult subjects and 2 pediatric subjects. The demographics and sleep related disorder experienced by each subject is described in Table 5.6 and Table 5.7. For the sleep

Table 5.6. Demographics of the subjects we evaluated RestEaZe on in the sleep lab.

Patient	Age	Sex	Disorder	Active %
PED1	10	male	ADHD	4.45
PED2	16	female	ADHD	8.64
RZ0015	60	female	insomnia	0.71
RZ0016	69	female	RLS	18.92

lab evaluation, our RestEaZe bands are always-on. The sleep lab data analysis provides us clinically valid ground truth on arousals (EEG), periodic leg movements (dorsiflexions measured by EMG on the anterior tibialis muscle), and complex leg movements (video). Using this ground truth data, we have validated the efficacy of our analytics. A trained technician scored the PSG data for event start time, event duration, body position (supine, prone, left/ride side), leg state (crossed/uncrossed), leg position (right/left on top), flexion, and complex leg movement (toe wiggling, foot jolts, etc). We ran trace-based simulations to evaluate the efficacy of our power management system. This data set constitutes the ground truth of labeled PSG and RestEaZe data for building a model. In both subjects data sets the total leg movement time constituted less than 10% of the total sleep time—pointing to the need for an effective power management strategy.

To demonstrate the efficacy of the power management system, we further performed 8 nights of in-home evaluation on 8 different subjects. The age of the subjects ranged from 20-35 and included six male and two female subjects. The subjects wore two RestEaZe bands, one on the ankle and the other just above the ankle on the same leg for a night. One of the bands had the power management algorithm implemented on the micro-controller, while the second band was always-on. The bands could not be placed on top of each other as this would interfere with capacitive sensing. Due to this arrangement the axis of the IMUs between the two bands were unlikely to line up perfectly. We discuss this limitation in the next section. The average recording session lasted 7.5 hours with a standard deviation

Table 5.7. Demographics of the subjects we evaluated RestEaZe on in-home data collections. This table shows the mapping between patient and experiment ID which is used in other figures.

Patient	Sex	Experiment ID
0	male	exp8
1	female	exp12
2	female	exp20
3	male	exp18
4	male	exp19
5	male	exp21
6	male	exp22
7	male	exp24

of 1.3 hours.

5.6.2 RestEaZe Energy consumption

We use two metrics to evaluate the energy consumption of the RestEaZe system: (1) Total uptime that RestEaZe sensors (apart from the capacitors) were on during an overnight study (**Uptime**)–this is a time when a leg movement occurs; and (2) Total energy consumption of the RestEaZe band when compared to a band that is always-on.

Figure 5.14 shows the distribution of the total **Uptime** for the eight in-home subjects. The median **Uptime** is close to 11%, while the maximum **Uptime** is 20%. Figure 5.14 show the **Uptime** per subject. The **Uptime** varies from subject to subject depending on how disruptive a subject's sleep is. For instance, subject 7 suffers from chronic insomnia, and hence the **Uptime** for 7 is 4.9 times higher than Subject 2, who is a sound sleeper. Figure 5.10 shows the number of hours that the band would last on a single charge of the 400 mAh battery as a function of the **Uptime**. With a 10% **Uptime**, the system would last for 95 hours, which is equivalent to four days of continuous operation. Figure 5.15 show the average power consumption of RestEaZe for the eight in-home subjects. The overall power

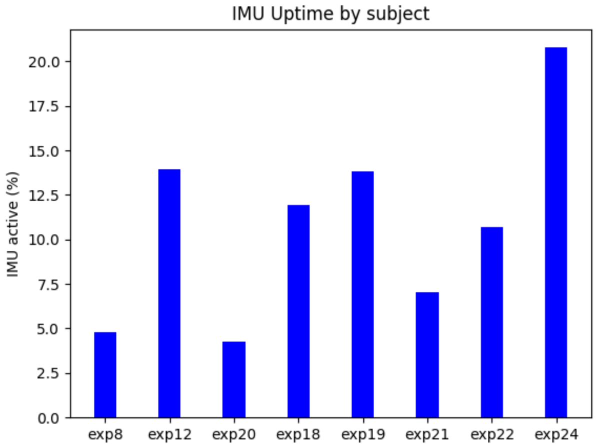

FIG. 5.13. Percentage of time spent in active mode for eight different subjects.

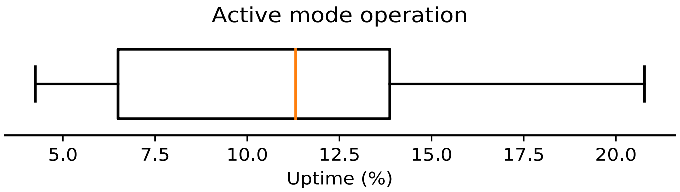

FIG. 5.14. Boxplot of time spent in active mode for eight different subjects.

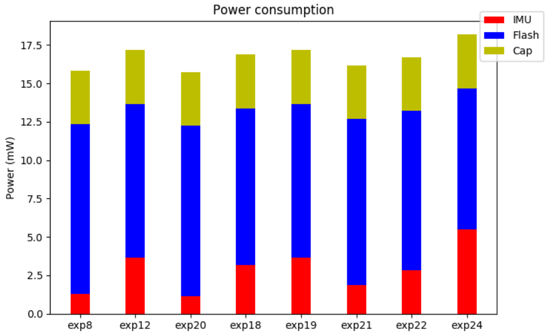

FIG. 5.15. Average power consumption by module of eight different subjects, collected in-home during normal sleep.

consumption system varies from 15.75 mW to 18.18 mW. We can draw two conclusions from the figure. First, the RestEaZe band has a lifetime that varies from 4.5x to 5.2x when compared to an always-on band. Second, we observe that writing to the flash consumes a substantial amount of this power. The system power consumption can be further optimized by caching data in memory and compressing it before it is written into the flash.

5.6.3 Data Quality Collected by RestEaZe

An important evaluation metric for RestEaZe is the quality of data reconstructed by the system when running in low-power mode. We evaluate the data quality by comparing the reconstructed data with the actual raw data from the always-on system as well as comparing the error in the sleep measures calculated by RestEaZe and the always-on system.

Correlating the raw sensor data: The low-power version of the system, ideally, should be able to capture all movements for a subject while remaining in the low-power mode as long as possible. We compare the output of the always-on system with the reconstructed data from RestEaZe and calculate the Pearson correlation coefficient for the two time series for the sensors on the band that we reconstruct the data for (accelerometer

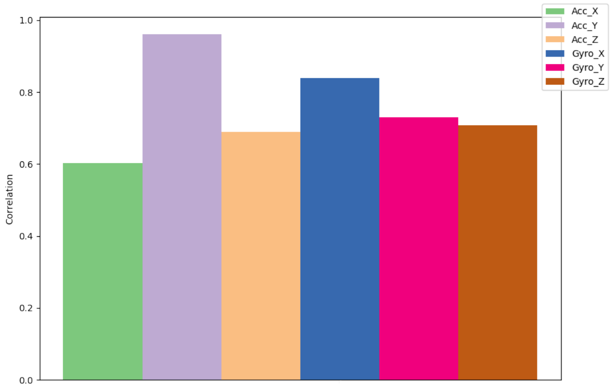

FIG. 5.16. Pearson correlation between data from two regular bands worn over night on the right leg.

and gyroscope). The resulting correlations for the eight in-home studies is shown in Figure 5.17. Although one data set has a very high correlation across all sensors, the others tend to have low correlation coefficient values. After careful scrutiny, we found that since the bands were not placed at exactly at the same position and orientation, the IMU sensors might have experienced different rotational and linear forces during movements. To validate this observation we performed a more controlled experiment in which the two regular bands were worn exactly in the same orientation overnight. The result shown in Figure 5.16 shows high correlations. In spite of this, in the next section we show that the final values of the sleep measures calculated by the RestEaZe band closely matches the always-on band.

Comparing sleep quality metrics: The RestEaZe system provides several sleep measures as output from the analysis on the leg movement data. These sleep measures are similar to the measures provided by an overnight sleep study performed using polysomnograph. We have independently validated the accuracy of the measures calculated by RestEaZe with PSG. The measures include: (1) `CLMnum:` The total number of Complex Leg Movements; (2) `ARL:` the total number of awakening and micro-arousals during

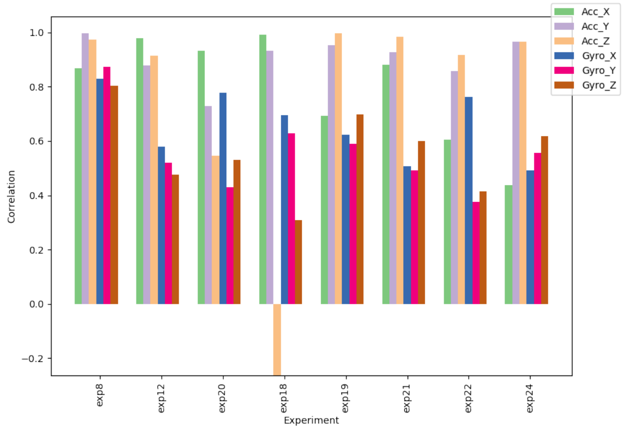

FIG. 5.17. Pearson correlation between data from regular band and low-power band.

the night; (3) `PLMnum:` the total number of periodic leg movements during the night (dorsiflexions); (4) `WASODur:` Wake after sleep onset duration; (5) `Sleep Efficiency:`; and (6) `T`ST: Total Sleep Time. To evaluate the accuracy of determining these metrics with RestEaZe, we calculate the error in the value of these metrics calculated from the always-on band and RestEaZe. Figure 5.18 shows the error distribution across all the eight subjects. We observe the median error in calculating these metrics is less than 10%, which is within range of the error when PSG is manually tagged by trained clinicians.

Summary of Results: RestEaZe collects very similar quality of leg movement data (with error in the sleep measures $<$ 10%), while having a lifetime of 5X compared to an always-on sensor.

5.7 Related Work

RestEaZe builds on previous work on sleep monitors and studies on the significance of leg movements for sleep. Here we compare and contrast RestEaZe with the most relevant literature.

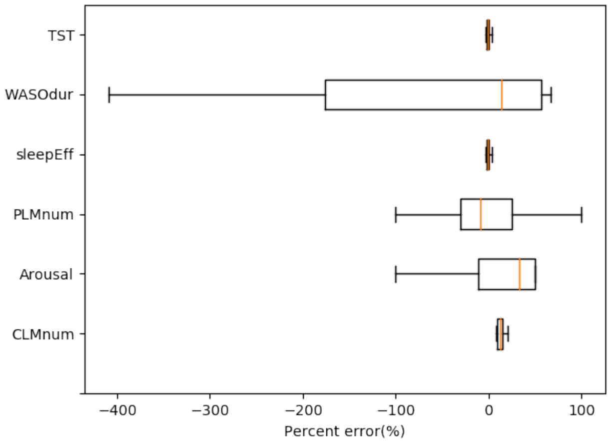

FIG. 5.18. Output of our IMU analytics script contains several parameters. This is the percent error between the regular and low-power modes.

Sleep Monitoring Devices: There are several devices that perform sleep monitoring and are listed in Table 5.8. We detail the differences and similarities between the devices in Table 5.8. Of these, only the RestEaZe device has the sensitivity to accurately identify PLMS and micro-arousals. One other system claims to measure PLMS (not micro-arousals), the PAM-RL, but is 4-5 times more expensive, difficult to use, and cannot assess or compare movement in both legs. The competitive advantages of RestEaZe are in the basic underlying RestEaZe concepts of using leg movements for determining texture of sleep and its use as an aid for diagnostics of sleep disorders. These advantages will persist over time. RestEaZe provides a convenient, accurate, low cost means for diagnosis and monitoring of awakening arousals, micro-arousals, and PLMS in the home setting.

Relevance of Periodic Leg Movements During Sleep (PLMS): High rates of PLMS indicate possible sleep apnea, RLS, or PLMD. PLM detection may also have some future relevance for those with cardiovascular risk factors, given the now well-established relation between each PLM and significant transient blood pressure and heart rate elevations (Sid-

System	Location	Leg Activity	PLM	Arousals
Phone apps	Pocket	No	No	No
Fitbit	Wrist	No	No	No
Actitrac	Wrist	No	No	No
PAM-RL	Ankle	Yes	Yes	No
RestEaZe	Ankle	Yes	Yes	Yes

Table 5.8. The table highlights the major advantages of RestEaZe over existing sleep monitoring systems.

diqui *et al.* 2007; Pennestri *et al.* 2007) that can be reduced by dopamine treatment (Bauer *et al.* 2016). But more significantly, evaluating PLMS for children addresses a currently recognized medical need and role for RestEaZe. PLMS $>$ 5 per hour in children is considered clinically significant (L Marcus *et al.* 2014; Pennestri *et al.* 2006) particularly when associated with increased arousals and poor sleep (L Marcus *et al.* 2014). PLMS occur commonly for children with attention deficit hyperactivity disorder (Sadeh, Pergamin-Hight, & Bar-Haim 2007; L Picchietti & Stevens 2007) and RLS (Picchietti *et al.* 2007). The PLMS indicate disturbed sleep exacerbating associated problems, e.g. attention deficit hyperactivity disorder (ADHD). PLMS relate to reduced brain (substantia nigra) iron (Li *et al.* 2016). PLMS in children often occur with low but normal serum ferritin indicating possible functional brain iron deficiency. Children with PLMS and low normal serum iron (<50-75 mcg/l) can be successfully treated with oral (Simakajornboon *et al.* 2003; Tilma *et al.* 2013) or even IV iron (Grim *et al.* 2013). In these studies, the critical issue was the PLMS, not the sleep complaint or a diagnosis of RLS.

5.8 Summary

This work demonstrates the design, implementation, and evaluation of the RestEaZe system, a sleep monitoring system that uses leg movements as a marker for arousals, restless sleep, and as an aid for diagnostics of sleep-related disorders such as RLS and ADHD.

We present the software and hardware architecture for RestEaZe and present a hierarchical system which uses low-power textile capacitive sensors to wakeup high-power consuming sensors like the IMU. Using evaluation on 4 subjects in a sleep lab and 8 other subjects in a home setting, we show that the RestEaZe sensor has a lifetime of 5x as long as an always-on band while providing similar quality data.

Chapter 6

CONTEXT AWARE SYSTEMS

While fall detection methods can help provide immediate remedy to the patient, fall prediction and prevention are the only plausible ways to minimize the economic and health impact of falls. While observable factors such as improper medication can be a valid predictor of a future fall, the spatio-temporal variation in the gait of an individual can reliably predict future falls (Maki 1997; Bridenbaugh & Kressig 2010). For instance, researchers in Gerontology have observed that increase in stride-to-stride variability, stride speed, and spatio-temporal variation in inter-leg spacing have a high correlation to falls (Helbostad & Moe-Nilssen 2003). Unfortunately, these variations are usually small and often occur over a long period of time (Maki 1997), so capturing these variabilities solely in a clinical setting creates several difficulties (Bridenbaugh & Kressig 2010). A typical gait analysis session costs several hundred dollars and a one time session can not detect problems that develop over time.

To reliably evaluate the risk that an individual may have a future fall, the subject's gait must be monitored continuously in a home setting. Designing systems that can continuously monitor spatio-temporal variations in an individual's gait, however, faces unique challenges. First, for long-term adoption, the system should ideally be minimally obtrusive and invisible to the user. At the same time, it should be able to accu-

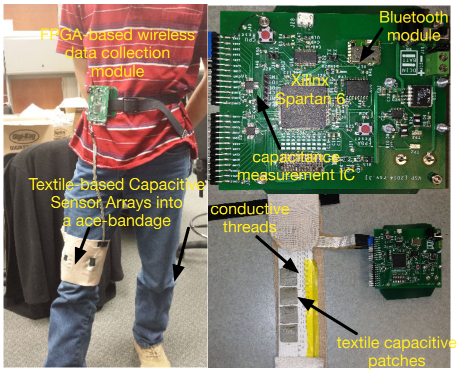

FIG. 6.1. The figure shows the prototype of our system. Capacitive Sensor Arrays built out of conductive textile are sewn into an elastic bandage. The data is collected and processed using a FPGA-based wireless module. The right subfigure shows the zoomed-in version of our custom-defined FPGA board, and the sensor array.

rately collect gait attributes such as variation in inter-leg spacing when the user is walking and changes in stride speeds over time. Finally, the system should have a long battery lifetime, minimizing the need to recharge batteries. Existing systems rely either on camera-based techniques or accelerometers built into smartphones to perform such an analysis (`http://mcroberts.nl`). While depth cameras such as the Kinect (`http://www.xbox.com/en-US/kinect`) can help collect data on an individual's gait when the camera is facing the user, it cannot reliably collect data at all possible locations. Smartphone sensors like accelerometers, on the other hand, can infer certain gait attributes such as stride speed but cannot determine more fine-grained gait features such as the variation in the distance between the legs when the user is walking. Moreover, accuracy of using the smartphone sensors is predicated on where the user carries his smartphone, i.e., whether it is in the user's pocket or in his hands.

6.1 CapGait: A Context Aware Gait Analysis System

To address the above challenges, we present a system that uses textile-based capacitive sensor arrays built into wearable ace bandages that can continuously collect data on a user's gait. Figure 6.1 illustrates the prototype of our system. The textile sensors can be built into clothing of the user and worn at all times. The system uses a novel hierarchical low-power processing subsystem to accurately detect two attributes of gait—stride speed and inter-leg spacing of a walking individual—that have a high correlation with falls (Maki 1997).

Research Contributions: The design, implementation, and evaluation of our prototype system presents three novel research contributions.

(a) **Wearable textile-based capacitive sensor arrays (CSAs) for gait analysis:** We present the use of flexible textile-based capacitive sensor arrays built into items of daily use, e.g. an ace bandage determining gait analysis for fall prediction. Our sensors measure changes in capacitance due to motion in the proximity of the sensors. An array of these capacitive sensors help cancel noise in the environment, provide high spatial resolution, and sense modality for fine-grained motion detection.

(b) **Context-aware hierarchical signal processing for low-power wearable sensors:** We present a hierarchical signal processing architecture that breaks down the signal processing algorithm for calculating the gait attributes into a hierarchy of processing stages. The low-power processing units calculate observations from a collection of capacitive sensors and extract features from the data relevant for motion detection. The more centralized processing unit aggregates features from multiple low-power processing units to determine gait attributes. The processing stages are activated based on relevant context. For instance, the low-power processing units perform feature extraction only when walking is detected and the centralized processor performs analysis only when relevant features are extracted by the lower level processors. We demonstrate that the hierarchical architecture can accurately detect gait attributes while consuming minimal energy.

(c) **Prototype development and evaluation:** We have prototyped our system using an array of rectangular capacitive textile sensors on an ace-bandage and a FPGA-based processing unit. Our evaluation using the prototype shows that the system can measure stride speeds with 85% accuracy and inter-leg spacing with an accuracy of 98%. We further evaluate the power consumption of our hierarchical signal processing architecture by synthesizing the FPGA design in a 180-nm CMOS process. We show that the power consumption of our hierarchical system is four times lower than a system where a single centralized processor is used for all the processing.

6.2 Textile-based Capacitive Sensor Arrays

To capture all attributes of an individual's gait, a continuous 3D model of the leg movements while the user is walking is required. However, capturing this 3D model is infeasible without an elaborate imaging and 3D modeling system. Ultimately, we aspire to build such a sufficient 3D model of a individual's limb movements using low-cost minimally-obtrusive wearable sensors. As a first step towards that goal, this work focuses on measuring a subset of gait attributes such as stride speed and inter-leg spacing. To capture these attributes we use textile and conductive-thread based capacitive sensor arrays (CSA). While capacitive sensing is used in multi-touch screens (Lee, Buxton, & Smith 1985) and touch sensors built of wearable fabrics (Saponas, Harrison, & Benko 2011), our approach explores more robust *proximity* sensing using wearable fabrics. To understand how these sensors work, Figure 6.2(a) diagrams an example sensor array built into an ace bandage on one leg. The effective capacitance between the plates change as the other leg moves into the vicinity of the plates. As the leg moves close to the plates it forms a set of series capacitors with the body. This also works at greater distances, enabling more-distant proximity sensing, though there are some challenges. For instance, stray capacitances exist that divert current

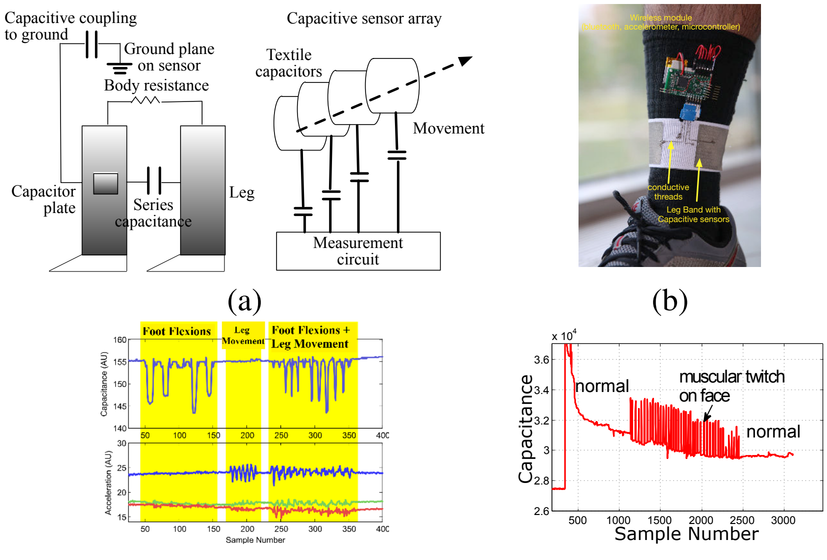

(a) (b)

(c) (d)

FIG. 6.2. (a) A diagram illustrating the CSA placed on the leg and the equivalent capacitance measurement system. (b) An application of the CSA for collecting foot flexion movements. (c) Data from the CSA and the accelerometer shows that the CSA can reliably capture foot flexions while the accelerometer can capture leg movement. (d) The sensitivity of the CSA is demonstrated using an experiment where the CSA is placed as a band aid on the face of an individual. The sensors are able to capture subtle muscular twitching on the cheek.

away from the series measurement path affecting capacitance measurement. At close proximity, the sensor-body capacitances are large and the stray capacitances can be ignored. For larger distances, we mitigate the effect of stray capacitances by using arrays of sensors and analyzing changes over time rather than instantaneous individual values. Our hierarchical signal processing algorithm builds observations by combining data from capacitor plates using an observation model (described in §6.3) to determine the approximate location, velocity, and acceleration of a body part with respect to the plates.

The use of capacitive sensor arrays (CSA) built into clothing has several advantages over conventional sensors such as pressure sensors and accelerometers used for gait analysis. First, pressure sensors and accelerometers can capture a narrow set of gait movements

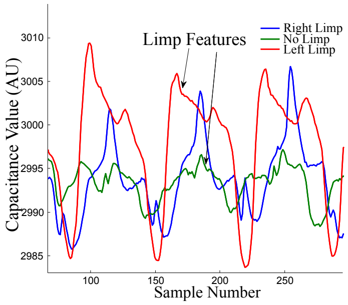

FIG. 6.3. Processed data from our sensor clearly showing the limp in the subject's movement

localized to the part of the body where they are placed. Because the CSAs are more versatile, they can capture subtle movements that occur in the proximity of the sensors. To demonstrate a scenario where the accelerometers cannot capture a specific leg movement, we present Figure 6.2 (b) and (c). Figure 6.2(b) shows a prototype band, similar to our ace bandage prototype worn around the ankle of an individual. The band consists of a capacitor sensor array and an accelerometer. Figure 6.2(c) presents data captured using the band when a subject performed three types of leg movements: (a) ankle flexions without intentionally moving the leg, (b) moving the leg without ankle flexions, and (c) moving the leg and performing ankle flexions. The figure shows that the capacitive sensors are able to reliably capture ankle flexions while the accelerometer can capture leg movements. [1] Therefore, by combining the two sensors we can reliably capture both ankle flexions and leg movements. Secondly, the CSAs have high sensitivity, are low cost (less than a cent apiece), and consume low-power. To demonstrate the sensitivity of motion capture using a CSA we present Figure 6.2(d). The figure plots the capacitance measured using textile-based capacitive sensors built into a band-aid placed on the cheek when a subject performed facial muscular twitching. The figure shows that the capacitive plate is able to reliably capture this subtle motion. Further, Figure 6.3 shows data from our ace-bandage prototype

---

[1] Ankle flexions during sleep have a high correlation with determining restless leg syndrome (Tyvaert *et al.* 2009)

worn by an individual with a limp. The figure illustrates that our sensors can capture the limp reliably, demonstrating the sensitivity and versatility of the CSAs. While in this work, we target monitoring leg movements, our overarching goal is to build a methodology for using this new type of sensor measurement on various areas of the body for a variety of body motion sensing applications. The effort requires complementary signal processing algorithm and hardware implementation framework. We next describe our signal processing framework that analyzes the data from the capacitive sensor array.

6.3 Hierarchical Signal Processing

Each textile capacitor plate measures the capacitance between one leg and the plate (depicted as the capacitor between the two legs in Figure 6.2(a)). The capacitance increases as one leg comes closer to the plate, or inversely, the capacitance value provides a measure of the location of the leg with respect to the plate. Our goal, however, is to extract other motion features such as velocity in addition to position and consequently convert these features into gait attributes such as stride speed and distance between legs. To accomplish this, we use an array of capacitor plates. Data from this array is processed using a hierarchical signal processing architecture illustrated in Figure 6.4(a). Several of the processing elements in our architecture are activated selectively to conserve computational effort and energy, and the data through communication channels between the processing stages is dynamically varied according to context to conserve power.

At the lowest level of the hierarchy, data from multiple capacitive plates are aggregated into *observations*, as shown in Figure 6.4(a). Aggregating capacitor channels into observations is important for extracting important motion attributes such as location, direction of motion, and speed of movement. These observations are described using the

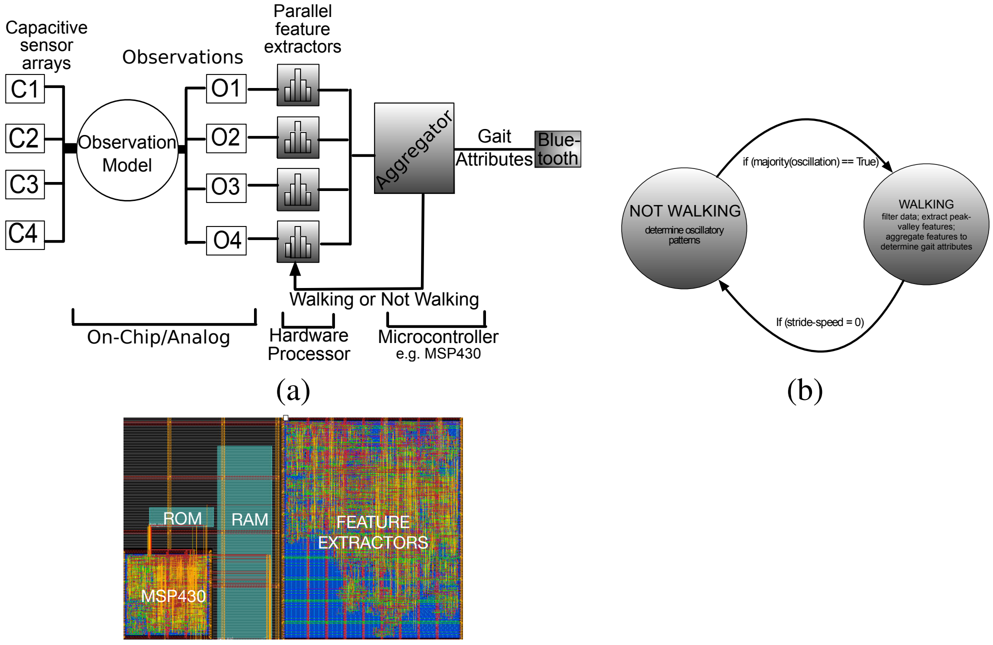

(a) (b)

(c)

FIG. 6.4. (a) Our hierarchical context-aware signal processing architecture. The processing is staged to facilitate maximal parallelism and minimal energy consumption. Higher levels of the hierarchy can be switched off till the lower stages determine interesting features. The feature extraction is optimized based on the underlying context (walking or not walking in our prototype system). (b) The processing is driven by the underlying context. If the system is in the "not walking" state, it looks for oscillatory patterns to determine if the user is walking and if the system is in the "walking state," the feature extraction and aggregation is tuned to extracting gait attributes such as stride width and stride speed. (c) Our synthesized cores place in an IC floor plan.

following observation model.

$$\begin{bmatrix} y_1 \\ ... \\ y_k \end{bmatrix} = \begin{bmatrix} W_1^1 & ... & W_n^1 \\ ... & ... & ... \\ W_1^k & ... & W_n^k \end{bmatrix} \cdot \begin{bmatrix} c_1 \\ ... \\ c_n \end{bmatrix} \tag{6.1}$$

where $y_1, ..., y_n$ are the observations, W_i's can take a value of 0, 1, or -1, and $c_1, ..., c_n$ are the outputs from the n capacitor plates. To understand how these observations work, let us assume four capacitor plates $C_1, ..., C_4$. Let us assume that there are two observations $y_1 = c_2 + c_3$ and $y_2 = c_2 - c_3$. The observation y_1 is roughly equivalent to data from a capacitor whose plate is the union of the capacitor plates c_2 and c_3. y_1 is maximized when the second leg is in the center of the two plates, so this can be used to localize a limb with respect to the plates. y_2 is calculated from the differential between two plates and can be used to track motion as the leg moves from plate c_2 to c_3 and vice-versa. These observations can be used to track location and movement of objects in the proximity of the plates. A collection of these observations help us reliably determine strides and inter-leg spacing—gait attributes whose temporal variations have high correlation with imminent falls. The observation model is implemented on the capacitance measurement IC in hardware, and in the future, we plan to implement these observations using ultra-low-power analog subthreshold circuits.

The next level of hierarchy takes data from observations and extracts features such as waveform peak and valley locations. This level is implemented as a parallel array of processing stages. The feature extraction, however, at this stage is context-aware. In our prototype, the underlying context that determines the feature extraction is whether a user is walking or not walking. As illustrated in Figure 6.4(b), if the system determines that the user is walking, the features are targeted towards collecting gait attributes; otherwise the system tries to determine when the user has started walking. Since a user in a home setting would not be walking often, the context-aware feature extraction helps us implement low-power, coarse-grained feature extraction to determine walking and more fine grained filtering and feature extraction to calculate gait attributes, leading to large energy savings.

Walking is established through detection of an ongoing oscillatory pattern. The method was motivated by a zero-crossing detection approach but modified according to the data features already available – peak-valley data. The algorithm uses the occurrence of a positive-valued peak followed by a negative-valued valley (or vice versa) as a zero crossing. For each parallel processor analyzing an observation, a binary vote for walking is cast when enough zero crossings occur in a specified time frame. This vote (if true) is sent to the highest level of the processing hierarchy and a simple majority vote determines that the context has now changed to walking. The higher level processor then triggers the low level processing units to capture features tuned to collecting gait attributes. A key design factor here is that enough sensors and measurements must be made to confidently establish the walking context, a question of *observability*. In our system the number of sensors was established through experimental trial-error-adjustment alongside design of the processing stages. We performed two optimizations to minimize computation and energy consumption when the underlying context is "not walking." First, the feature extractors do not filter the data. This saves energy in the system and filtering is not necessary for coarse feature extraction like the zero-crossing for detection of oscillations. Even if a step is missed in this process the only penalty is latency in detecting walking. Second, we reduced the sampling rate to minimize processing.

Determining stride speed and inter-leg spacing: When a state of walking is established by the system, the data from the observations undergo a filtering process. We have implemented an IIR filter (Storn 1996) that approximates a low-order windowed-integrator. This pre-filtering of the data presents the peak-valley detector with a measure of area under/above the peak/valley, and hence the size of the peak, rather than just an instantaneous maxima/minima. Peak-valley data (value and time) is extracted and transmitted. The peak valley detection unit employs a strategy to look for a series of increasing samples followed by a series of decreasing samples for peak detection and vice versa for valley detection. Rapid glitches are rejected as noise so as not to cause unnecessary data flow and wake-up signals in the system. The extracted peaks and valleys are then used by the aggregator (a higher power processor in Figure 6.4(a)) to determine stride speed and spacing between legs.

To understand how the aggregator determines stride speed and leg spacing, it is important to understand the observations derived from the capacitor data. In our prototype, we use four capacitor plates $\{c_1, c_2, c_3, c_4\}$, and derive six observations from these four sensors $y_1 = c_1 + c_2, y_2 = c_1 - c_3, y_3 = c_1 - c_4, y_4 = c_2 + c_3, y_5 = c_2 - c_4, y_6 = c_3 + c_4$. The observations that sum capacitor outputs (*summed observations*) emulate a larger capacitor plate and a peak in the data stream from the observation can be used to locate and place the object with respect to the plates. The observations that take the differential between the capacitor plates (*differential observations*) can help find the direction of motion by analyzing peaks and valleys. In our algorithm at the aggregator, depicted in Figure 6.4(a), a majority vote is used to determine a walking state. Peaks from summation channels and valleys from differential channels are counted as votes. In a detected stride, the valleys and peaks coincide in time to when the opposite leg's position is close to the capacitive sensors as the leg swings forward or backward. The time between strides is then used to determine the stride speed. Similarly, to determine the inter-leg spacing we first determine the start of a stride using the summed and differential observations described above, and then we use the average of the amplitude of the first peak (after the stride is determined) in the differential stages to calculate the inter-leg space. In our experiments, we found that the use of multiple observations improves the reliability of determining the gait attributes.

Our context-aware hierarchical architecture facilitates aggressive duty-cycling of higher power processing elements. The aggregator, for instance, can be kept switched off if the feature extractors do not see relevant features. While we demonstrate the efficacy of our system using two gait attributes that are related to falls, the system can be extended to other gait attributes and can be generally used to design low-power signal processing on data collected from an array of sensors.

Prototype Implementation: We have prototyped a fully functional sensing unit that comprises of four 1-inch by 1-inch capacitive sensor plates sewn into an ace bandage (Figure 6.1). The signal processing architecture is prototyped on a Spartan 6 FPGA. Additionally, we have designed a custom PCB that houses the FPGA module, a Bluetooth module, and a capacitance measurement IC. The feature extractor processors are built as hardware modules on the FPGA, and we use the microblaze core as the aggregator. As the microblaze is a proprietary Xilinx processor implementation, we intend to use the TI msp430 open core implementation in our final system. Additionally, since an FPGA module is not designed for low-power consumption, to demonstrate the efficacy of our hierarchical architecture, we have synthesized, placed, and routed ASIC cores for our signal processing architecture using 180-nm CMOS technology. The cores were synthesized using Cadence RTL Compiler and placed and routed using Cadence Encounter. We used the low-power synthesis and place and route options in both the tools. The synthesized cores placed in an IC flloorplan is illustrated in Figure 6.4(c).

6.4 Summary

Described in this thesis is a system that uses textile-based capacitive sensor arrays built into clothing for capturing gait attributes that have a high correlation with falls. The system uses an array of capacitive sensors to detect movement in the proximity of the capacitor plates. Our novel hierarchical and context-aware signal processing algorithm converts data from the capacitive sensor array into high-level gait attributes such as stride speed and inter-leg spacing. We have prototyped our sensor on an ace bandage and our processing unit on an FGPA, and show that the system can accurately measure gait attributes with high accuracy and low energy consumption.

Chapter 7

EVALUATING SOLAR POWERED EMBEDDED SYSTEMS

Renewable energy-driven systems are key to a self-sustainable society. The increase in the number of renewable energy-driven homes, sensors, and compute systems bear ample testimony to the growing popularity of such self-sustainable systems. Solar panels are the most popular renewable energy harvester. Solar or Photovoltaic (PV) panels are a non-linear device that harvest energy from incident light such as natural sunlight or light from artificial sources like incandescent and fluorescent lamps. PV panels scale from small panels used for micro-harvesting indoor light, medium sized panels that power sensing and environmental monitoring systems, to large PV arrays that power buildings. Irrespective of their size and the type of systems they power, evaluating and debugging solar panel-driven systems is a tedious and cumbersome process. Debugging such systems involve deploying them under realistic conditions, performing remote diagnostics, evaluating software and hardware bugs, and then repeating the development lifecycle. Rapid prototyping, hence, is difficult, if not infeasible for such systems. Therefore, in lieu of time and resources, most of the evaluation for PV panel driven systems is performed in a *single* setting. Such an evaluation lacks external validity and it is often unclear how performance would be affected when the environmental conditions change.

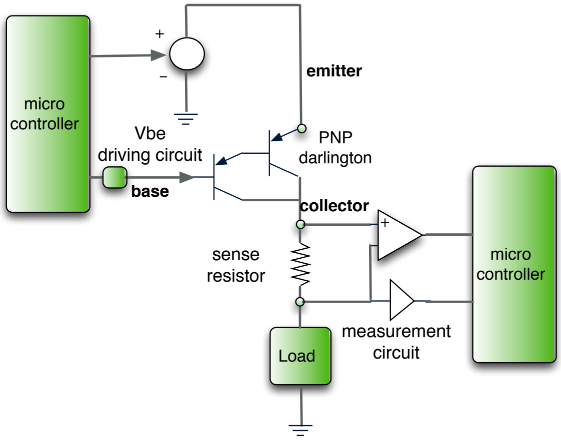

FIG. 7.1. SunaPlayer uses a non-linear analog device, a PNP darlington transistor, a multi-scale driving circuit, and a high resolution measurement circuit, to emulate the non-linear characteristics of a solar cell. The driving circuit is controlled by a state-machine based PID controller running at a micro-controller.

7.1 SunaPlayer: High-Accuracy Emulation of Solar cells

Solar panel emulation is a plausible solution to streamlining the process of evaluating solar panel driven systems. The goal of the emulation is to perform evaluations of solar-panel driven systems in the laboratory; it is not intended to replace the use of solar cells in the deployed systems. The emulation can help determine (a) which solar cells to use and (b) address bugs in the system quickly by rapid prototyping and evaluation in the lab. While there are systems (high-end power supplies for instance) available in the market that can perform solar panel emulation, they are expensive and not portable.

A solar cell emulation platform emulates the current and voltage characteristics of a solar cell, modeled by its IV curve. The goal of the emulator is to power a *load* as if it were powered by an actual solar cell. Hence, a wide range of IV traces collected from a variety of solar cells under different lighting and temperature conditions can be *replayed* in a laboratory setting to power and test the system being evaluated. Such a device can radically ease the development lifecycle of solar panel powered devices. The existing solar emulation platforms, however, are limited in several ways. For example, most emulators use linear *digital* solutions and have limited output current range (Zhang *et al.* 2011). The key challenge to emulating solar cells across a wide operating current/voltage ranges is to build accurate non-linear models for solar cells of different sizes and under different lighting and temperature conditions. Since existing system use linear devices to emulate a solar cell (a non-linear device), the emulators are fine tuned and sensitive only in a narrow operating range of output voltages and currents [1]. Secondly, most emulators rely on high performance personal computers to compute the operating current and voltage values on an IV trace of a solar cell for a given load. Therefore, the system is powered using wall sockets and is not portable. For several applications, however, portability is paramount. For instance, consider a solar panel driven mesh network where the system designer wants to determine the type and size of the solar panel he should use. For such an application, the emulator can be programmed with a set of solar panel IV traces and deployed with the mesh nodes. The emulator can then in-vitro sense the ambient light intensity and temperature and select the appropriate solar panel IV curve to emulate. For such an emulator to be practically useful, it should consume minimal power, and should be portable.

To address the above limitations, we present the design, implementation, and evaluation of the SunaPlayer, an emulation platform for solar panels spanning micro solar cells to large sized solar panels. The overall circuit design and the prototype SunaPlayer is illustrated in Figures 7.1 and 7.2. SunaPlayer uses a versatile non-linear *analog* device, a PNP darlington transistor as the current source with a diode for emulating a solar cell. The PNP darlington transistor has an IV characteristic similar to a solar cell. The SunaPlayer uses a multi-scale driving and measurement circuit, and a novel proportional-integral-differential controller to in-vitro build an accurate non-linear model of the solar cell for a given lighting

---

[1] Commercial power supplies that can be used as solar cell emulators are very expensive and are not portable (`http://sine.ni.com/nips/cds/view/p/lang/en/nid/204239`).

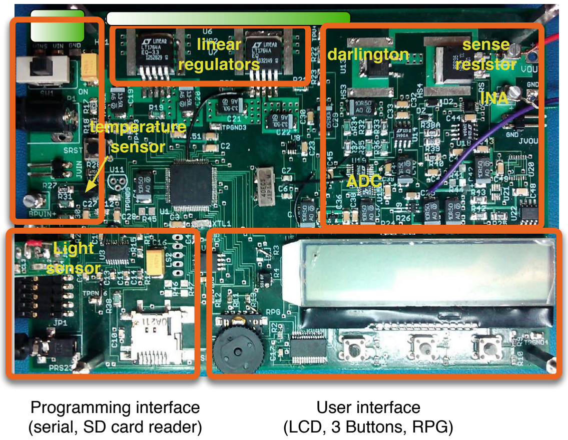

FIG. 7.2. The SunaPlayer prototype.

condition. The emulator takes IV curves as input and emulates the solar cell for a given load. Generating IV curves for solar cells is out of scope of the SunaPlayer. However, there are systems that can profile solar cells accurately (e.g. Ekho system (Hester, Scott, & Sorber 2014)). The SunaPlayer supports a wide operating range, spanning output currents from 430 μA to 1.89 A (at low voltages; four orders of magnitude) and output voltages from 0.02 V to 9.8 V (three orders of magnitude), and can emulate solar cells with an accuracy of close to 99%. The SunaPlayer does not make any assumptions on the type of load being powered by the harvester. It can power linear and non-linear loads, combination of storage devices like capacitors and batteries, or devices that are directly powered by the harvester without a storage element.

Research contributions: The design, implementation, and evaluation of the SunaPlayer presents the following novel research contributions. (1) **Novel hardware design:** Unlike previous work that use linear devices to emulate solar cells, the SunaPlayer emulates the solar cell using a high gain non-linear *analog* device, a PNP darlington transistor, whose IV characteristics are similar to a solar cell. The PNP darlington is controlled using a multi-scale driving circuit that allows the SunaPlayer to be sensitive across a wide range of output currents and voltages, across a wide range of lighting and temperature conditions. (2) **Novel software controller:** SunaPlayer uses a novel state machine based software controller that uses PID logic, caching, and device-specific optimizations to accurately control the PNP darlington. The iterative controller helps the SunaPlayer build an accurate model for the solar cell without a priori information on the load characteristics. (3) **Functional prototype:** We have prototyped a fully functional SunaPlayer device, and evaluated it in the context of two applications: emulating small, medium, and large sized solar panels, and in-vitro determination of the appropriate solar cell for a given system. We show that the SunaPlayer can emulate solar panels with IV characteristics ranging from 430 μA – 1.89 A (4 orders of magnitude) and 0.02 V — 9.8 V (3 orders of magnitude) with an accuracy of more than 99%.

7.2 Hardware Architecture

The SunaPlayer employs a hardware architecture that supports solar panel emulation for a wide range of open circuit voltages and short circuit currents. The overall SunaPlayer hardware circuit is illustrated in Figure 7.1. The primary emulation device is a PNP darlington transistor. Using a darlington for emulating solar cells is a novel approach. This non-linear analog device has a large operating range, is power efficient, and has an IV response close to that of solar cells. The device is controlled by a driving circuit. The output load voltage and current is measured using a multi-scale measurement circuit. A low-power 16MHz PIC micro-controller senses the output current and output voltage, and iteratively controls the driving circuit to output the appropriate voltage and current.

The SunaPlayer can operate in two modes: (1) offline mode; and (2) online mode. In the offline mode, the system is programmed with a set of solar panel IV traces. The SunaPlayer chooses an IV trace to emulate and automatically finds the correct operating voltage and current on the IV trace for a given load. The IV traces could correspond to solar panel profiles for different lighting and temperature conditions. They could also correspond to IV curves for solar panels of different sizes and compositions. If the solar panel traces are not collected experimentally, the system can also use IV traces from datasheets. In the online mode, the SunaPlayer senses the ambient light intensity and selects the appropriate IV curve to emulate. The online mode is useful for testing a system on a set of solar panel traces in an actual deployment. It is possible to add a third mode of operation in which the emulator uses an analytical model of the IV curve. This would require a preprocessing step that approximates a complex analytical model (one that possibly requires logarithmic and exponential function evaluations) as a set of piece wise linear functions. This mode will be implemented in the future.

The SunaPlayer can power a wide range of loads or charge a wide range of storage devices. For example, the emulator can charge a battery, a super-capacitor, or a combination of storage devices and charge controllers. The emulator can directly power a sensing or embedded system without a storage device. We next describe the hardware components of the SunaPlayer.

PNP darlington device: Existing devices for emulating solar panels assume a constant voltage source like a capacitor powering the load (Zhang *et al.* 2011), or use linear digital solutions like programmable voltage regulators (Li & Chou 2004) to emulate a solar cell. Ekho (Zhang *et al.* 2011), for instance, assumes a voltage source such as a capacitor charging the load. A storage device like a capacitor provides a deterministic output voltage across the load, and the emulator can measure this voltage to determine the output current on the solar cell's IV curve. Unfortunately, such a system cannot be used in configurations where the load is directly powered by the solar panel without a storage device or if the system uses a combination of a charge controller circuit and two or more storage devices like capacitors and batteries. The other solution is to use a linear regulator to generate the load

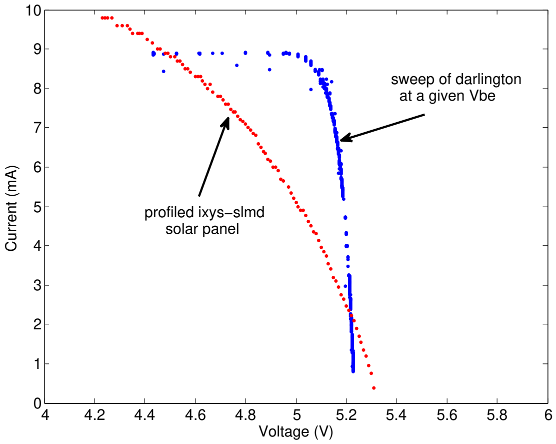

FIG. 7.3. The figure shows a part of the IV curve of a profiled solar panel at 7500 lux. The solar panel has an open circuit voltage of 5.5 V and a short circuit current of 10 mA. The IV trace is generated by varying the output load powered by the solar panel under an incandescent lamp. The blue line is the IV sweep for the PNP darlington and is generated using a set of load impedances. The figure shows that the darlington and the solar cell have similar IV characteristics, and they behave as a current source with a diode.

voltage and current. The output of the linear regulator is divided using a resistor divider, and appropriate output voltages are produced using a digital potentiometer. However, such a solution suffers from several drawbacks. First, the sensitivity of a linear device is limited for a wide range of current and voltage outputs. The digital potentiometer used to generate the output voltage usually has 1024 intervals, which provides the same resolution when emulating very small current and voltage outputs and large current and voltage outputs. Hence, the device is not sensitive to small changes in load currents and voltages or small changes in light intensity. It might be possible to get a wider operating range using multiple linear regulators. However, the power consumption, given the efficiency of most linear regulators, would be high. Moreover, with multiple PNP darlingtons the SunaPlayer can have an even wider operating range. Second, and more importantly, a linear regulator is a voltage source and may not be appropriate for emulating a current source like a solar panel. To understand why, consider the following property of a linear regulator. If the load current increases, the regulator will supply the current without any drop in the voltage. Hence, the system will have transients that do not mimic a PV cell.

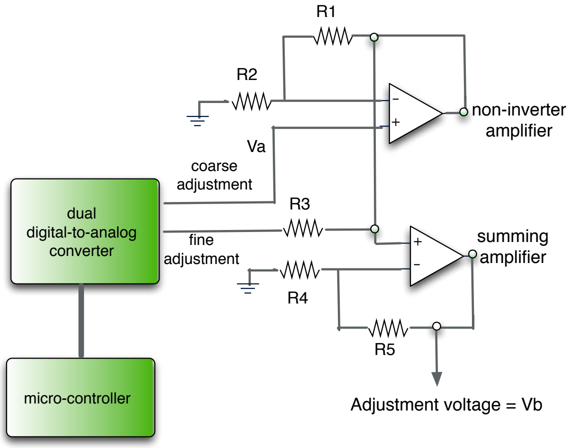

FIG. 7.4. The multi-scale circuit to drive the base voltage of the PNP darlington. The circuit uses a combination of a fine adjustment that provides the SunaPlayer with the ability to fine-grain control the changes in V_{be} and a coarse adjustment that helps generate base voltages close to $V_E(V_+)$ for small output currents.

SunaPlayer uses an analog non-linear device, a PNP darlington transistor, to emulate solar panels. A PNP darlington is a high gain non-linear device that acts as a current source (a corresponding NPN darlington transistor can also be used as a current sink). It is a pair of bipolar junction transistors (BJT) connected in a common collector configuration. This configuration creates a device with a higher current gain compared to a single BJT. The gain of the darlington can be approximated as $\beta_1 \cdot \beta_2$, where β_1 and β_2 are the individual transistor gains. The load is attached to the collector, while the load current and load voltage can be controlled using the bias voltage V_{be}—the voltage drop across the base and the emitter. The darlington overcomes the drawbacks of a linear digital device (voltage regulator) and a single BJT. First, the darlington can source a wide range of currents and voltages. By adaptive fine-grained control of the base-emitter voltage drop, the output current and the output voltage can be controlled with a fine resolution across a wide range of output current and voltage. Moreover, the BJTs allow a non-linear rescaling of the quantization error over a long range. In particular it allows better proportional error at lower currents and voltages while supporting a wide operating range. Second, since the current gain of the device is high (unlike a single BJT), the input base current required to generate a high output current is low. Hence, the system can be controlled using low-power micro-controllers. Third, for a fixed value of V_{be}, the output characteristics of a darlington is close to the output characteristics of a photocell. As illustrated in Figure 7.3, if the load current increases, the voltage will drop as it does with a photocell. Hence, the system responds more accurately to rapid load perturbations. Finally, the darlington is also an efficient non-linear analog device that can be current limited and voltage limited by appropriately setting V_E, the emitter voltage.

The darlington transistor has advantages over a linear regulator when emulating solar cells; however, accurate and fine-grained control of the device presents several technical challenges. The darlington is a non-linear device and the output characteristics of the device are difficult to model. Moreover, our system does not make any assumptions on the load characteristics; hence, in the absence of any analytical models that characterize the collector current and voltage as a function of the bias voltage (V_{be}), it is challenging to design a low latency controller that finds the operating point on a solar cell's IV curve accurately. Additionally, like any transistor, the darlington chain operates in three modes—cutoff, saturation, and active modes. Unfortunately, if the darlington is in saturation mode, it takes a long time to switch from saturation to active state, and a high saturation voltage can lead to high-power dissipation when large output currents are required. Therefore, it is important to assure that the darlington does not transition to saturation mode. We address these challenges in our design of the driving and measurement circuit and the novel state machine based controller described in §7.3.

Multi-scale driving and measurement circuit: The driving circuit accurately sets V_b (base voltage) that determines V_{be} (base-emitter drop), the bias voltage that controls the collector (load) current and voltage. Given the non-linear characteristics of the darlington, for high sensitivity at low and high current ranges, the driving circuit for the darlington must provide variable scale control for different output ranges. Moreover, the darlington is a high gain non-linear device, and hence, for high sensitivity, the driving circuit that sets V_b must provide very fine-grained changes in V_b. This will allow the SunaPlayer to respond to small changes in load impedance and light intensity. At the same time, it must be possible to set V_b close to $V_e(V_+)$, when very small output currents and voltages are desired. We solve the problem using a multi-scale driving circuit, illustrated in Figure 7.4.

The darlington's base voltage is controlled by dual 16-bit DACs with an SPI interface to the micro-controller. The DACs output a fine adjustment voltage and a coarse adjustment voltage, as illustrated in Figure 7.4. To provide coarse scale adjustments, V_a (in the figure) is fed to a non-inverting opamp stage with a gain of (1 + R_1/R_2), effectively 5 for our circuit. The fine-grained control line is combined with the coarse grained adjustment using an amplifier and a set of resistors that create a summing amplifier. Effectively, it provides a V_{adj} that is equal to $x \cdot V_f + y \cdot V_c$ where V_f is the fine grained voltage and V_c is the coarse grained line. x is small, 0.5 in our case, and y is close to 7.5. The fine adjustment voltage has a resolution of 22.9 μV and a range of 0 to 1.5 V. The coarse adjustment voltage has a resolution of 858 μV with a range of 0 to V_+. The fine adjustment range of 1.5 V allows the darlington to be biased from cutoff to high current levels by only varying the fine adjustment voltage. The coarse adjustment, on the other hand, allows the darlington base voltage to be set close to V_+, generating a small bias, and small output voltages and currents. This multi-resolution circuit allows both fine grained and coarse grained adjustment of V_b, and provides the SunaPlayer the flexibility to produce very low output currents and high output currents, and at the same time be highly sensitive to small changes in the load impedance.

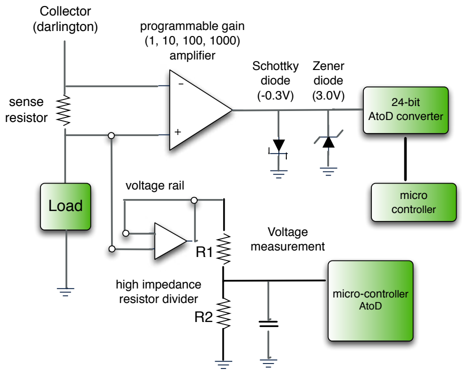

FIG. 7.5. The multi-scale circuit for measuring the output voltage and the output current. The instrumentation amplifier with gains of 1, 10, 100, and 1000 allows measuring low and high currents accurately, and provides multiple resolutions of output measurements. The diodes clamp the output voltage and allows accurate measurements using the ADC. The voltage is measured using a high impedance rail that minimally effects the load currents.

In the SunaPlayer design, the darlington's emitter voltage, V_e, is controlled by a V_+ programmable power supply ($V_e = V_+$). The value of V_+ is varied by changing the value of a digital potentiometer, programmable from the micro-controller using a SPI interface. The voltage, V_e, can be varied from 5 V to 10 V. The digital potentiometer has 256 steps and the change in V_+ from a single LSB step is a function of the voltage range. For example, if V_+ = 5 V, a least significant bit (LSB) provides a change in V_+ of 9.7 mV while for a V_+ = 10 V a LSB provides a change of 114 mV. This fine resolution allows for accurately setting V_e on the darlington. The value of V_+ depends on the solar cell being emulated. V_+ helps limit the output voltage of the emulator, similar to the voltage limiting characteristic of a solar cell. However, to avoid the darlington from transitioning from active to saturation mode, V_+ must be higher than the open-circuit voltage of the solar cell.

The load current is measured by the voltage induced across a low tolerance 1 ohm sense resistor, attached to the collector (Figure 7.5). An instrumentation amplifier, with programmable gains of 1, 10, 100, and 1000, amplifies the voltage across the sense resistor. The use of multiple gain levels allow very accurate current measurements for very low and high currents. Theoretically, the measurement circuit can measure few nanoamps of current accurately. For highly accurate measurements, the voltage output from the amplifier is clamped using diodes, so that the output voltages do not fall below -0.3 V or is above 3.0 V. The first diode is a Schottky diode that prevents the voltage from dropping below -0.3 V , preventing negative voltage drops. The second diode is a Zener diode that prevents the voltage from rising above 3.0 V. Another important consideration for our measurement circuit design is the thermal properties of the sense resistor. Under high output load currents significant power may be dissipated. A one ohm resistor with a 2 A output current dissipates 2 W, and even with a heat sink the temperature of the resistor may rise. The temperature coefficient of resistance, therefore, is taken into account in the current calculations. To measure the load voltage, a high input impedance voltage follower amplifier is used (illustrated in Figure 7.5). This allows accurate voltage measurements with very minimal loading due to the measurement circuitry. The load current sensing circuits may be calibrated using two current sources (100 mA and 200 μA) that are optionally connected to the load using a solid state switch. The primary processing unit of the SunaPlayer is a 16 MHz PIC18F87K22 micro-controller with an additional 16 MBit flash, serial I/O, and a microSD card reader slot to store IV traces. We have designed, implemented, and assembled a fully functional prototype of the SunaPlayer, illustrated in Figure 7.2. The boxes illustrate the different components of the 4-layer printed circuit board.

7.3 Software Architecture

The SunaPlayer does not make any assumptions about the characteristics of the load attached to the harvester. To power any arbitrary load, the software controller that sets the base voltage must address two key research challenges. First, the load impedance is unknown, hence it is impossible to a priori determine the operating point on the IV

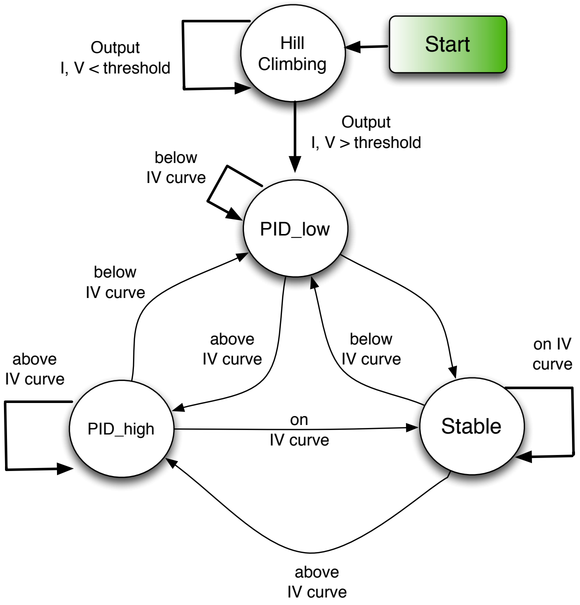

FIG. 7.6. The state machine for our software controller. The controller determines the base voltage V_b that must be applied to generate a bias voltage V_{be}. The state machine starts in a `hill climbing` state that transitions the darlington from cutoff to active mode and stabilizes the output voltage. The `PID_low` and `PID_high` states use PID controllers when the output voltage and output current values are below and above the IV curve respectively. The SunaPlayer is in a stable state when the output voltage and current is on the IV curve.

curve of the solar panel. Secondly, while the characteristics of a PNP darlington is close to a PV cell, like a PV cell, accurately modeling the non-linear relationship between the inputs (bias voltage V_{be} and light intensity) and the output current and output voltage is difficult. To address the two challenges, the SunaPlayer leverages a novel controller for the base voltage that uses a finite state machine and a proportional-integral-derivative (PID) logic that quickly converges to an operating point on the solar panel IV curve without a priori knowledge of the load impedance and the relationship between the input and output variables of the darlington.

State machine-based controller: The idea underlying the software controller is simple. The micro-controller applies a bias voltage (V_{be}) by changing the voltage on the base of the darlington. This produces an output voltage and output current at the collector (load). The output voltage and the output current is sensed using our multi-scale measurement circuit and read by the micro-controller. The micro-controller then determines the change in V_{be} (ΔV_{be}) that must be applied to the d arlington base to reduce the *error* between the operating point on the solar cell IV curve and the present output current and voltage. Since the actual operating point on the IV curve is unknown, we use a heuristic to determine the error. The solar panel IV traces are offline modeled as a set of polylines. We apply the k-means algorithm (`http://en.wikipedia.org/wiki/K-means_clustering`) on the IV trace and then use linear regression on the clustered points to determine the polylines. An example set of polylines for a solar panel IV trace is illustrated in Figure 7.7. The SunaPlayer stores the last four output voltage and current values generated by the darlington by varying V_{be} and performs a linear regression on these stored points to generate a line. The controller then finds the intersection of the modeled line and the polylines generated from the panel's IV trace. The intersection point is the *predicted* operating point for the given load. The error is calculated as the Manhattan (`http://en.wiktionary.org/wiki/Manhattan\_distance`) distance between the predicted operating point and the point corresponding to the present output voltage and current. Note that the load might be non-linear and our algorithm approximates the load's IV characteristics using a set of piecewise linear functions. The approximation of the IV trace of the solar cell as polylines and the load characteristics as a set of linear functions reduces the computational complexity of finding the *predicted error* and can be efficiently implemented on a micro-controller platform.

For every V_{be}, an error is calculated, and the system uses a proportional-integral-differential (PID) controller to iteratively reduce the error. A PID controller is a generic feedback-based controller that is used for industrial applications. If there is no a priori knowledge of the underlying process, the PID controller is considered the best performing controller (`http://en.wikipedia.org/wiki/PID\_controller`). The PID controller, in our application, determines ΔV_{be}, the change in V_{be} in the next iteration. ΔV_{be} is a function (summation) of three entities: (1) proportional error ($K_1 e$), where e is the error described above (2) integral error ($I \int e_t dt$), which is a summation of the errors in the previous iterations of the algorithm, and the (3) differential error ($D \cdot de/dt$) is the slope of the error over time. K, I, and D are constants for the proportional, integral, and differential errors. The three types of errors help faster and more accurate convergence in our application over using a purely proportional controller. The differential error, for instance, can help determine when the system is close to the actual point of operation while the integral factor helps smoothen errors caused by darlington transients. Note that the error e, described above, accounts for hardware error induced due to the measurement circuit, leakage currents of the darlington, and circuit noise.

While applying the PID approach, we ran into several practical challenges due to the non-linear characteristics of the darlington. For instance, when the output voltage and current is zero, the darlington is in cutoff mode. On increasing V_{be}, the darlington switches from cutoff to active mode, but it takes a finite amount of time before the output of the darlington stabilizes. Hence, calculating the error using the extrapolation method described above can be erroneous during this stage. Additionally, given the non-linear characteristics of the darlington, we have found that the optimal PID constants that converge to the operating point when the output voltage and current values are below the IV curve are different from the PID constants when the output voltage and current values are above the curve. This is because the increase in V_{be} required to switch the transistor on is different from the decrease in V_{be} required to proportionally switch it off.

Our controller, therefore, is a state machine that operates in four states (illustrated in Figure 7.6), depending on which region of the IV space the output current and output voltage lies. The first state is a `hill climbing` stage where the darlington switches

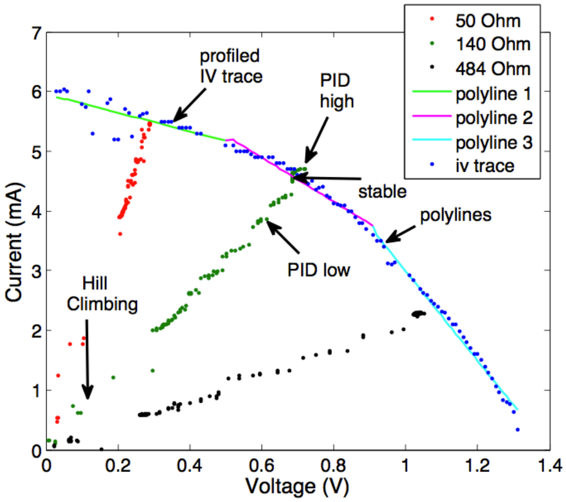

FIG. 7.7. The figure shows an example run of the state machine based controller for three load impedances of 50, 140, and 484 Ωs. The figure shows that the output current values in the `hill climbing` state is randomly spread, but it stabilizes and shows a linear trend in the `PID_low` state since the load is a linear device. The figure also shows the `PID_high` and `Stable` states.

from cutoff to active mode and its output stabilizes. For our application, we determined experimentally that the controller should transition from the `hill climbing` state to the `PID_low` state when the output voltage and output current is more than 25% of the maximum voltage and current in the IV trace. The system stays in the first PID state, `PID_low`, till the output voltage and output currents lie below the IV curve. As soon as the system reaches the IV curve, it transitions to the `Stable` state and the controller stops changing V_{be} as the operating point has been reached. If the system overshoots the IV curve, it transitions to the `PID_high` state where a different set of PID constants are used. Finally, the system can move out of the stable state if the load impedance changes or a different IV curve is emulated. The load impedance can change if subsystems like the radio module is switched on or off and a new IV curve is emulated if the ambient light intensity changes. Figure 7.7 shows an example run of the state machine controller as it transitions through the states. The load is a resistor, a linear device. As shown in the figure,

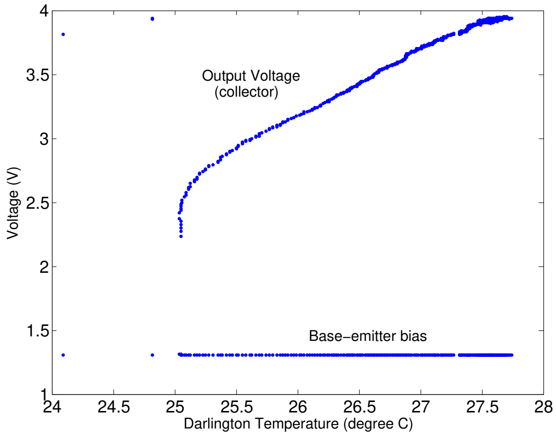

FIG. 7.8. The figure shows that the output voltage of the darlington changes by 1.5 V when its temperature varies by 2.5°C. The bias voltage (V_{be}) is kept constant. The temperature change occurs because a large 400 mA current is sourced by the darlington, causing high heat dissipation. This artifact occurs because the output gain of the darlington changes logarithmically with temperature.

in the `hill climbing` state the output voltages and currents are not stable and do not show the linear trend that is characteristic of a resistor. However, in the `PID_low` state the outputs are stable and show a linear trend and converge to an operating point on the IV curve, modeled by polylines.

Latency optimization: To improve the latency of convergence to an operating point, in the `Stable` state, the SunaPlayer caches V_{be}, the output current, and the output voltage in a lookup table (LUT). The SunaPlayer can then use these cached values to improve the latency of convergence when the desired output voltage and current is close to a cached value. When the SunaPlayer controller determines the predicted operating point on the IV curve, using the extrapolation method described above, it linearly searches the entries in the look up table to find the output voltage and output current that has the shortest Manhattan distance to the predicted value. The system then sets the bias voltage, V_{be} in the LUT entry. Caching operating points reduce the number of iterations required in the PID state to converge.

Brand	Incident Light	OCV	SCC
Cymbet	7500 lux	1.3 V	6 mA
Ixys	13000 lux	5.5 V	10 mA
Solamax	13000 lux	2.9 V	100 mA

Table 7.1. The three characteristic solar panels and ambient light intensity used to evaluate the SunaPlayer.

Practical implementation challenges: At high currents and high voltages, the thermal properties of the darlington creates a practical challenge, illustrated in Figure 7.8. When a current of above 400 mA is sourced by the darlington, the output voltage of the darlington (load voltage) changes from 2.5 to 4 V, even though the bias voltage V_{be} is kept constant at 1.3 V. A closer look shows that the temperature of the darlington changed from 25 to 28^oC during the course of the experiment. Unfortunately, the gain of the darlington changes with temperature, hence, if its temperature changes, the output current and voltage rises steadily even if the bias is constant. Therefore, for deterministic results, the temperature of the darlington must be kept constant when changing V_{be}. In our prototype, we use a processor fan on top of the darlington to dissipate the generated heat. This helps minimize the temperature fluctuation of the system.

7.4 Evaluation

We evaluate the SunaPlayer while focussing on the following questions. **(1)** How accurately does the SunaPlayer emulate solar cells? What is the range of output currents and voltages that the SunaPlayer can support? **(2)** What is the SunaPlayer's latency of convergence to an operating point on the IV curve? How sensitive is the device to changes in load impedance and ambient light intensity? **(3)** What is the energy consumption of different components of the SunaPlayer? While answering these questions, we also present micro-benchmarks on system overhead and cost, as well as associated accuracy and latency

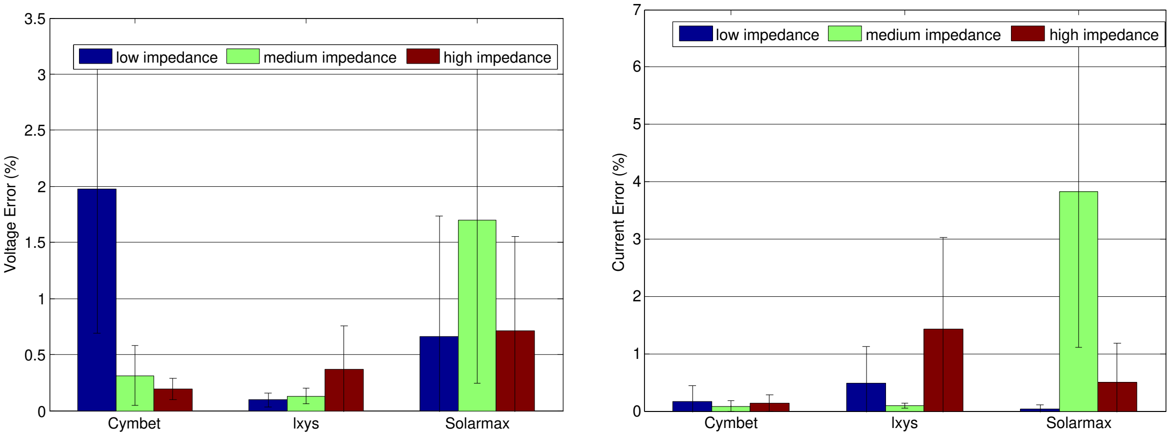

(a) (b)

FIG. 7.9. (a) The figure shows the accuracy of converging to an operating voltage on the IV curve. (b) The figure shows the accuracy of converging to an operating current on the IV curve. The impedance values for the load was varied from 45 Ω to 2.1 $K\Omega$s. The low impedance values correspond to the low voltage, high current region of the IV curve, the medium impedance is close to the maximum power point, and high impedance correspond to the high voltage, low current region of the IV trace. The average accuracy is close to 99% for both current and voltage for the wide range of solar panels and impedance values.

tradeoffs in the system.

Experimental setup: We evaluate the SunaPlayer in the context of two applications. The first application is *offline panel emulation*. For this application, the SunaPlayer replays IV traces from solar cells under low, medium, and bright lighting conditions. For the second application, *solar cell comparison*, the SunaPlayer determines the appropriate solar cell IV curve to emulate based on ambient light intensity when the system is deployed in the wild. For the first application, we generate IV traces for three solar cells under low, medium, and high ambient lighting conditions. To generate IV traces under controlled lighting, we use a programmable incandescent lamp, whose light intensity is controlled by regulating the current draw using a transformer. The incident light intensity is continuously monitored using a light meter. The operating points on the IV trace of the solar cell is generated using

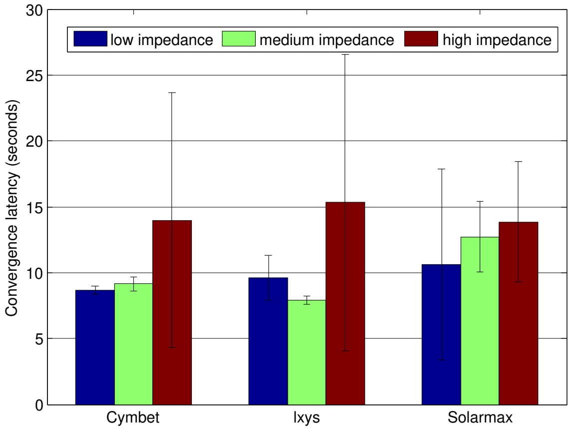

FIG. 7.10. The figure shows the latency of converging to an operating point for the SunaPlayer. The worst case latency is less than 10 seconds. This is the absolute worst case latency of convergence when the SunaPlayer starts at a cutoff state. In the common case when the devices transitions from one IV curve to another, the latency is less than 2 seconds.

a high resolution programmable load. A summary of the incident light (in lux), type of solar cell, the short circuit current, and the open circuit voltage of the solar cell under the given lighting conditions is described in Table 7.1. In our experiments, we use these three IV traces as they represent three characteristic points in the spectrum of output voltages and currents that the SunaPlayer can support. We use the programmable load and a resistor-capacitor box to emulate loads that the SunaPlayer powers. We used resistive loads for baseline evaluation of the system. We note that the SunaPlayer can support higher currents and voltages than the Solarmax and Ixys panels in Table 7.1. However, generating IV traces for very high output currents and voltages require direct sunlight that is impossible to control and regulate. Therefore, we have performed a proof of concept evaluation on synthetic operating points and verified that the SunaPlayer can emulate high output currents and voltages.

Range and accuracy: We first evaluate the output current and voltage range that the

SunaPlayer can support. The range is defined as the smallest and largest output current and output voltage that the SunaPlayer can generate with high accuracy. To determine the smallest voltage and smallest current, we set V_+ to the smallest value that our system can support (5V), and then reduce the bias voltage (V_{be}) to as small as possible so that the darlington is in active and not in cutoff mode. This bias voltage switches on the darlington by the smallest possible amount. We then try a set of load impedances and measure the voltage drop across the load and its current draw with a high resolution multimeter and our measurement circuit simultaneously. We determine the error in the measurement, and if the error is below a threshold, we assume that the system can generate the output voltage and current. To measure the largest possible current and voltage, we set V_+ to the maximum value that the SunaPlayer can support (10V). We then increase the bias voltage so that the darlington is in active mode and not in saturation and measure the largest possible output voltage and current. Through this experiment, we found that the SunaPlayer can accurately generate output voltages spanning 0.02 V to 9.8 V (at low currents), and output currents from 430 μA to 1.89 A (at an output voltage of 2.12 V). The analog-to-digital converter induces a noise of +/- 1.5% when measuring low voltages around 0.02 V, and induces an error of 0.03 V when measuring high voltages (9.8 V). It also induces a 7% error at ultra-low currents. The error is caused by an offset voltage of the instrumentation amplifier. For the higher current values, the heat dissipated by the darlington is a bottleneck. Even with the processor fan, there is a steep rise in the darlington temperature. We believe that with a passive heat sink, the range can be further improved. In spite of the above, the SunaPlayer can support output voltages that span three orders of magnitude and currents that span four orders of magnitude. The system can, therefore, emulate a wide range of solar cells under a wide range of lighting conditions.

We next evaluate the accuracy of the SunaPlayer. The accuracy of converging to an operating point on the IV curve is calculated as a duple $[(C_I - A_I)/A_I, (C_V - A_V)/A_V]$, where A_V, A_I are the operating voltage and current for the given load and C_V, C_I are the output voltage and current of the SunaPlayer when the system is in the `Stable` state (see Figure 7.6). To determine the accuracy across a range of solar panels and lighting conditions, we perform our experiments in a laboratory setting using the three IV traces described in Table 7.1. For each IV trace, the SunaPlayer powers small, medium, and large loads. These loads correspond to the [low voltage, high current], [medium voltage, medium current], and [high voltage, low current] regions of the IV trace and vary from 45 Ω to 2.1 KΩs. We perform five experimental runs for every load and every IV trace, and calculate the average accuracy and the standard deviation. Figure 7.9 plots the voltage and current convergence errors for the SunaPlayer for the three IV traces. We find that the average accuracy is high–close to 99% for most cases. The overall accuracy of converging to the correct current value is higher than the voltage values. This is because the resolution of our current measurement circuit is higher than our voltage measurement circuit. Overall, the system can accurately converge to the correct operating voltage and current for a wide range of solar panels and lighting conditions.

Latency: We next measure the latency of converging to the operating points on the three IV traces. We use the same set of experiments described above for the accuracy measurements. Figure 7.10 shows the latency of converging to an operating point in seconds. In the worst case, the SunaPlayer incurs a latency of 10 seconds to converge to an operating point from cutoff mode (0 output voltage and 0 output current). Note that this is the *absolute worst case* latency of convergence of the SunaPlayer. We show in the next section (sensitivity) that the system responds to changes in lighting conditions quickly (in less than 2 seconds)—the latency is similar to that reported in previous solar emulation platforms (Li & Chou 2004). To understand the latency of the SunaPlayer, we break down the average latency into the fraction of time spent on measuring the output voltage and current and the time taken by the state machine algorithm. The primary bottleneck of the system is in measuring the

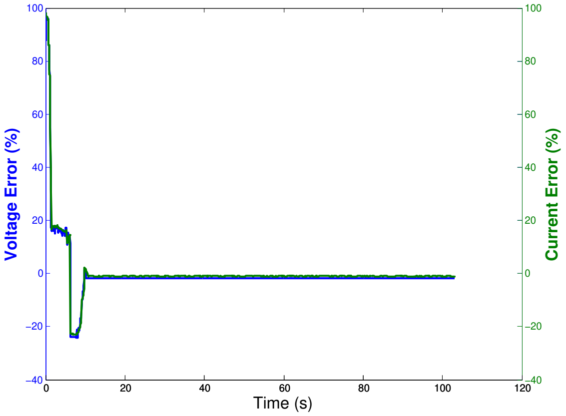

FIG. 7.11. The error in the output voltage and output current as a function of time. The error stabilizes close to 0% at 5 seconds, which is the *worst* case latency of the system. If a lower accuracy is acceptable to the application, the latency can be reduced. For example, at an error of 18%, the latency is less than a second. The negative error corresponds to values of the output voltage and current above the IV curve.

output current and voltage. It incurs close to 64% of the latency. Using a high speed instrumentation amplifier (`http://www.linear.com/product/LT1102`) and a high speed analog to digital converter will eliminate this overhead and bring down the latency to less than 720 milliseconds.

There is also an inherent tradeoff between latency and accuracy of the SunaPlayer. Since the SunaPlayer controller uses an iterative PID-based algorithm to converge to an operating point, a larger number of iterations can lead to better accuracy at the cost of higher latency. Figure 7.11 explores this tradeoff and presents the result of an experiment where we measure the time taken by the SunaPlayer to converge from cutoff to an operating point. We use the Cymbet panel IV trace for the experiment and use a load impedance of 140 Ω. From the figure, we see a diminishing return trend. Moreover, depending on the accuracy that an application desires, the latency of convergence can be reduced. For example, if an application can tolerate a 20% error, the latency can be less than a second.

Sensitivity: For accurate emulation of solar cells, the SunaPlayer must be sensitive to changes in load impedance and lighting conditions. When the load impedance changes, the operating point shifts on the IV curve. When the lighting conditions change, the operating point moves to a new IV curve. Sensitivity is important when the device is deployed in the wild to power systems. For such an application, the SunaPlayer continuously senses the ambient light intensity and chooses the appropriate IV curve to emulate. To evaluate system sensitivity, we perform two experiments. In the first experiment, we select a single IV curve to emulate. We then determine the smallest change in load impedance to which SunaPlayer is sensitive. We found that the SunaPlayer changes operating points when the load resistance changes by at least 1 Ω; hence, the system is sensitive to very small changes in load impedance. In the second experiment, we test the online mode of the SunaPlayer and sensitivity to changes in light intensity. For the experiment, we select two IV traces and *assume* that they correspond to ambient indoor light intensity and dark lighting conditions. The SunaPlayer measures the ambient light intensity using the onboard light sensor and selects the appropriate IV trace. Next, we set the lighting conditions to ambient lighting for one minute, then switch off the lights and set a dark light ambience for the next one minute, and then switch the lights on again. The result for the experiment is shown in Figure 7.12. The figure shows that the system is sensitive to changes in ambient lighting. It correctly switches from one IV trace to the other when the light intensity changes and automatically converges to a new operating point. The time to converge to the new operating point on a different IV curve is close to two seconds; therefore, the convergence latency is low.

Portability and cost: For the online mode of operation, the SunaPlayer must be portable and should consume low-power. We measured the power consumption of different components of the SunaPlayer prototype when it is sourcing zero output current and voltage and V_+ is set to 6 V. The results present a baseline power consumption of the system. The power consumption of the board without the use of the LCD screen is 148 mW (24 mA). However, the bottleneck of the present prototype is the 24-bit AtoD which

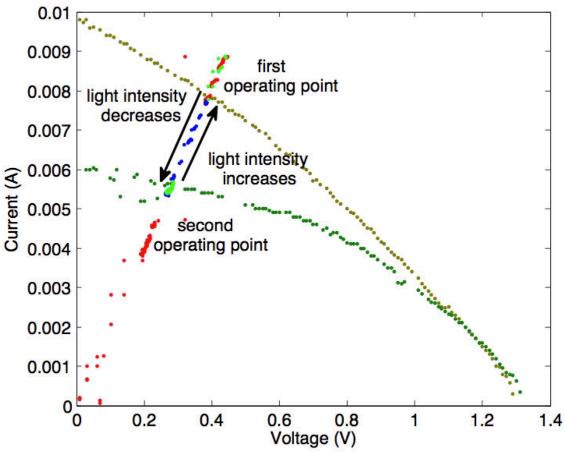

FIG. 7.12. The figure shows that the SunaPlayer is sensitive to changes in light intensity. The red dots show the convergence of the SunaPlayer to the first operating point. A change in lighting is detected by the SunaPlayer and the blue dots show the SunaPlayer transitioning to a different IV curve. Increase in the light intensity again causes the SunaPlayer to converge to the original operating point represented by the green dots.

consumes close to 72 mA. Using a different AtoD (`http://www.analog.com/en/analog-to-digital-converters/ad-converters/ad7192/products/product.html`), the consumption can be reduced to 5 mA. Therefore, the entire system can operate while consuming less than 30 mA when sourcing small currents, which is less than 30 times the quiescent current for a Vertex FGPA (`http://www.xilinx.com/univ/xupv5-lx110t.htm`). We could not compare the energy numbers with existing systems since these systems are tethered and were not designed for portable use. The embedded software on the SunaPlayer consumes only 39% program memory and 28% of the SRAM, and hence, is lightweight. Another important consideration for the SunaPlayer is its dollar cost. While an equivalent power supply (`http://sine.ni.com/nips/cds/view/p/lang/en/nid/204239`) that provides similar range and accuracy would cost around $4000, the SunaPlayer if manufactured in quantities of 10K or above would cost less than $100 a unit—an order of magnitude lower cost.

7.5 Related Work

The SunaPlayer builds on previous work on emulation platforms for solar panels and batteries and complements a suite of research efforts on renewable energy driven systems.

Emulators: There are several prior efforts on designing emulation hardware for batteries and solar cells. B# (Chou *et al.* 2003) is a battery emulator that uses an adjustable linear regulator to emulate output battery voltages. The operating range of the emulator is 1.0–4.5 V, and the system can source a maximum current of 800 mA. S# (Li & Chou 2004) uses the same architecture as B#, and extends the approach to solar panel emulation. B# reads the current draw of the system and looks up the corresponding voltage using a lookup table, and generates a control signal at the output of a digital to analog converter. While using a linear regulator simplifies the design of the control logic, it suffers from the drawbacks described in §7.2—a voltage regulator is a linear device that is not sensitive for wide ranges of output voltages and currents. Ekho (Zhang *et al.* 2011) is a solar cell emulator tuned to small currents. Ekho uses a simple design where the load is powered by a capacitor. Hence, the voltage drop across the load can be measured using the voltage drop across the capacitor, and consequently, the corresponding operating point on the IV curve (voltage and current) can be calculated. The system uses a FPGA to speed up the IV calculations. While the system has high accuracy, its operation relies on the existence of a constant voltage source to power the load. The system similar to the SunaPlayer is the National Instrument's programmable power supply (`http://sine.ni.com/nips/cds/view/p/lang/en/nid/204239`) that is capable of simulating solar panels with a wide range of operation. The power supply offers fine grained measurement resolution similar to our measurement circuit. However, the programmable power supply is a closed source system and the design is not available to the research community. Additionally, the programmable supply costs several thousands of dollars ($4,000–$5,000) and is not a portable device. In addition to emulators, simulators have also been designed for mote-class devices (Shnayder *et al.* 2004b; Titzer, Lee, & Palsberg 2005; Per 2005; Werner-Allen, Swieskowski, & Welsh 2005) and micro-harvesters (Gummeson *et al.* 2010; Clark *et al.* 2009). The simulation platforms use analytical models that are often inaccurate since they do not account for a wide variety of ambient and environmental conditions that effect the power generated. Our work also builds on research in the 1980s that use bipolar junction darlington transistor chains for profiling solar panels (Kern & Wagemann 1987) and the use of power transistors for conditioning the output of solar panels deployed in residential settings (Steigerwald, Ferraro, & Turnbull 1983). However, we use a darlington chain to emulate solar cells for a wide range of solar cells.

Renewable energy driven systems: There is a plethora of research related to renewable energy driven systems. These include medium scale systems for environmental sensing (Lin *et al.* 2005; Kansal *et al.* 2006), networking (Banerjee, Corner, & Levine 2007), large scale systems for renewable energy driven homes (Banerjee, Rollins, & Moran 2011), and micro-harvester systems for indoor environment monitoring (Yerva *et al.* 2012). The number of renewable energy driven devices have grown in the past with the goal to create an energy-independent society. There has been a focus on building harvesters for renewable energy driven devices (Brunelli *et al.* 2008), adaptive algorithms that balance energy demand with supply (Hammons 2008), and hardware platforms and programming paradigms that minimize the power consumption of devices (Banerjee, Corner, & Levine 2007). The SunaPlayer is complementary to these renewable energy driven devices. It can be used in evaluating and debugging these devices without deploying them in the wild. Moreover, it can help in evaluating the devices under a wide range of environmental conditions.

7.5.1 Evaluation on non-resistive loads

The evaluation presented in this work was performed with a variable resistive load. Resistive loads have a single current component with a proportional dependence on voltage. In contrast, common CMOS load models predict a combination of constant and linearly-dependent supply current components, which is actually easier to regulate. However, modern systems can be complex, and may even incorporate dynamic supply-voltage dependent behaviors. Analytical models for every such complex system are infeasible to produce. Luckily, power distribution subsystems for these devices typically employ batteries and supply filtering to smooth rapid supply transients. It is important that our system respond quickly enough to go unnoticed to such systems. We foresee testing the SunaPlayer with such devices and estimating the overall system sensitivity to parameters though experimental parameter perturbations. Any parameter modifications which result in dynamics significantly varied from that predicted by simpler approximate models may be investigated further. However, we do believe that the system is designed to handle these loads. For instance, the PID controller is an iterative algorithm that has been shown to model non-linear control systems well. Secondly, the system handles a non-linear load using pairwise linear functions in the `PID_low` state, hence, the convergence accuracy for the SunaPlayer should be similar to resistive loads at the cost of longer convergence latency.

7.5.2 Thermal properties

One of the side effects of a wide operating range is accounting for heat dissipation of different components. For instance, the gain of the darlington changes logarithmically with increase in temperature. While a processor fan provides a temporary remedy, a surface mount heat sink attached to the junctions of the darlington is a more robust solution. Similarly heat sinks are required for the sense resistor and power supply modules. We plan to augment our present design while provisioning for these heat sinks in the next version of the board.

7.5.3 High speed components

A primary latency and energy bottleneck is the speed of the analog-to-digital converter (ADC) and the instrumentation amplifier (INA). The present prototype does not support high speed measurements of the output current and voltage. By replacing the present AtoD and INA chips with high speed ICs, the system can be made more power and latency efficient.

7.5.4 Wider range of operation

While the SunaPlayer supports a wide range of operation, this range can be further extended with minor tweaks to the circuit. The darlington is capable of sourcing up to 20 V output and several amps of current. However, the hardware bottleneck of the system is the linear regulator used to set V_+. Our present IC supports a maximum of 10 V; hence, limiting the maximum output voltage to 9.8 V at low currents. By simply replacing the linear regulator with a wide range one can improve the range. Moreover, using a higher resolution digital potentiometer that controls the linear regulator can help fine grained control of V_+.

7.6 Summary

Presented in this thesis is the design, implementation, and evaluation of the SunaPlayer, an emulation platform for solar panels. The SunaPlayer uses a high gain analog device, a PNP darlington transistor as the current source that scales better for wider range, and a multi-resolution driving and measurement circuit to control the darlington. It uses a sophisticated state machine based PID controller that converges to an operating point on a solar cell's IV curve with high accuracy while incurring short latency. The SunaPlayer can emulate a wide range of solar panels under a range of lighting conditions, can source output voltages from 0.02 V to 9.8 V, currents from 430 μA to 1.89 A, and can power arbitrary loads such as a combination of storage devices, charge controllers, and sensing systems.

Chapter 8

CONCLUSION

This work demonstrates the use of three different hierarchical architectures: virtualized services, virtualized sensors and context aware computing in the design of low power embedded systems. In general, hierarchical system design allows for higher operational times compared to traditional embedded system designs. The choice of a particular approach is tightly coupled to the available sensing, computation and energy storage resources. The virtualized services approach is suitable in designs where two or more subsystems with a large disparity in power consumption need to be combined. In the case of the emergency mesh node application, this architecture combined with energy harvesting through solar panels enables perpetual operation of the system. The virtualized sensors approach is suitable where multiple sensors are capable of measuring the same phenomenon and have a large disparity in energy consumption and/or computational cost. In the case of the sleep monitoring system, using a low power capacitive sensor to detect motion yielded a five fold improvement in the operating time of the system. The context aware approach is suitable for systems where computation comprises the vast majority of the system's power consumption. In all of the explored architectures, the common principle is that context of the system must be measured in a low power way since that is what drives which subsystems and tiers are used. Given that solar panels are used as energy sources for low power embedded systems, the solar replayer expands the diagnostic capabilities of designers.

This work can serve as a foundation for formalization of a framework for the co-design of multi-tiered ultra-low-power embedded system hardware-software.

REFERENCES

[1] Agarwal, Y.; Hodges, S.; Chandra, R.; Scott, J.; Bahl, P.; and Gupta, R. 2009. Somniloquy: augmenting network interfaces to reduce pc energy usage. NSDI'09, 365–380. Berkeley, CA, USA: USENIX Association.

[2] Banerjee, N.; Sorber, J.; Corner, M. D.; Rollins, S.; and Ganesan, D. 2007. Triage: Balancing Energy and Quality of Service in a Microserver. In *Proceedings of ACM MobiSys*.

[3] Banerjee, N.; Corner, M. D.; and Levine, B. N. 2007. An Energy-Efficient Architecture for DTN Throwboxes. In *IEEE Infocom*.

[4] Banerjee, N.; Rollins, S.; and Moran, K. 2011. Automating energy management in green homes. In *Proceedings of the 2nd ACM SIGCOMM workshop on Home networks*, HomeNets '11, 19–24. New York, NY, USA: ACM.

[5] Bauer, A.; Cassel, W.; Benes, H.; Kesper, K.; Rye, D.; Sica, D.; W Winkelman, J.; Bauer, L.; Grieger, F.; Joeres, L.; Moran, K.; Schollmayer, E.; Whitesides, J.; C Carney, H.; S Walters, A.; Oertel, W.; and Trenkwalder, C. 2016. Rotigotine's effect on plm-associated blood pressure elevations in restless legs syndrome: An rct. 86.

[6] Bridenbaugh, S. A., and Kressig, R. W. 2010. Laboratory review: the role of gait analysis in seniors mobility and fall prevention. *Gerontology* 57(3):256–264.

[7] Brunelli, D.; Benini, L.; Moser, C.; and Thiele, L. 2008. An efficient solar energy harvester for wireless sensor nodes. In *Proceedings of the conference on Design, automation and test in Europe*, DATE '08, 104–109. New York, NY, USA: ACM.

[8] Center of Disease Control and Prevention. 2010. Falls among older adults: An overview. `http://www.cdc.gov/homeandrecreationalsafety/falls/adultfalls.html`. Accessed: 2019-08-25.

[9] Chou, P.; Park, C.; Park, J.; Pham, K.; and Liu, J. 2003. B#: a battery emulator and power profiling instrument. In *ISLPED '03*.

[10] Clark, S. S.; Gummeson, J.; Fu, K.; and Ganesan, D. 2009. Towards autonomously-powered CRFIDs. In *Workshop on Power Aware Computing and Systems (HotPower 2009)*.

[11] Dr Peter Harrop, Mr James Hayward, R. D., and Holland, G. 2015. Wearable technology 2015-2025: Technologies, markets, forecasts. Technical report, IDTechEx.

[12] E Martin, S.; M Engleman, H.; J Deary, I.; and J Douglas, N. 1996. The effect of sleep fragmentation on daytime function. 153:1328–32.

[13] E. Martin, S.; Brander, P.; J. Deary, I.; and J. Douglas, N. 2000. The effect of clustered versus regular sleep fragmentation on daytime function. 8:305–11.

[14] Fesquet, L.; Sicard, G.; and Bidégaray-Fesquet, B. 2010. Targeting ultra-low power consumption with non-uniform sampling and filtering. In *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, 3585–3588. IEEE.

[15] Gabelia, D.; Mitterling, T.; Hgl, B.; Wenning, G.; and Frauscher, B. 2014. Do periodic arm movements during sleep exist in healthy subjects? a polysomnographic study. 15.

[16] Govil, K.; Chan, E.; and Wasserman, H. 1995. Comparing algorithms for dynamic speed-setting of a low-power CPU. In *Proceedings of the First ACM International Conference on Mobile Computing and Networking (MobiCom'95)*.

[17] Grim, K.; Lee, B.; Y Sung, A.; and Kotagal, S. 2013. Treatment of childhood-onset restless legs syndrome and periodic limb movement disorder using intravenous iron sucrose. 14.

[18] Gruetter, J. 2011. Squeezing out power. `http://powerelectronics.com/energy-harvesting/squeezing-out-power`.

[19] Gründler, P. 2007. Conductivity sensors and capacitive sensors. *Chemical sensors: An introduction for scientists and engineers* 123–132.

[20] Gummeson, J.; Clark, S. S.; Fu, K.; and Ganesan, D. 2010. On the limits of effective micro-energy harvesting on mobile CRFID sensors. In *Proceedings of 8th Annual ACM/USENIX International Conference on Mobile Systems, Applications, and Services (MobiSys 2010).* To appear.

[21] Hammons, T. 2008. Integrating renewable energy sources into european grids. *International Journal of Electrical Power & Energy Systems* 30(8):462 – 475.

[22] Helbostad, J. L., and Moe-Nilssen, R. 2003. The effect of gait speed on lateral balance control during walking in healthy elderly. *Gait & posture* 18(2):27–36.

[23] Helmbold, D. P.; Long, D. D. E.; and Sherrod, B. 1996. A dynamic disk spin-down technique for mobile computing. In *Proceedings of the Second ACM International Conference on Mobile Computing and Networking (MobiCom'96).*

[24] Hempstead, M.; Tripathi, N.; Mauro, P.; Wei, G.-Y.; and Brooks, D. 2005. An ultra low power system architecture for sensor network applications. In *ACM SIGARCH Computer Architecture News*, volume 33, 208–219. IEEE Computer Society.

[25] Hester, J.; Scott, T.; and Sorber, J. 2014. Ekho: Realistic and repeatable experimentation for tiny energy-harvesting sensors. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, SenSys '14, 1–15. New York, NY, USA: ACM.

[26] Huang, H.; Pillai, P.; and Shin, K. G. 2003. Design and implementation of power-aware virtual memory. In *Proceedings of USENIX Technical Conference*.

[27] Institute of Medicine. 2006. Sleep disorders and sleep deprivation: An unmet public health problem. Washington, DC: The National Academies Press.

[28] Kansal, A.; Hsu, J.; Zahedi, S.; and Srivastava, M. B. 2006. Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems*.

[29] Karout, H.; Diab, M.; Moslem, B.; Alkhatib, R.; Corbier, C.; and Badaoui, M. E. 2015. Frequency content analysis of gaitvertical ground reaction force. In *Advances in Biomedical Engineering (ICABME), 2015 International Conference on*, 53–56. IEEE.

[30] Kern, R., and Wagemann, H.-G. 1987. Uncomplicated measurement procedure for i-v-characteristics of photovoltaic generators at remote sites. 314–318.

[31] Krause, A.; Ihmig, M.; Rankin, E.; Leong, D.; Gupta, S.; Siewiorek, D.; Smailagic, A.; Deisher, M.; and Sengupta, U. 2005. Trading off prediction accuracy and power consumption for context-aware wearable computing. In *Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on*, 20–26. IEEE.

[32] L Marcus, C.; Traylor, J.; R Gallagher, P.; Brooks, L.; Huang, J.; Koren, D.; Katz, L.; B.A. Mason, T.; and E Tapia, I. 2014. Prevalence of periodic limb movements during sleep in normal children. 37:1349–52.

[33] L Picchietti, D., and Stevens, H. 2007. Early manifestations of restless legs syndrome in childhood and adolescence. 9:770–81.

[34] Lee, S. I.; Ozsecen, M. Y.; Della Toffola, L.; Daneault, J.-F.; Puiatti, A.; Patel, S.; and Bonato, P. 2015. Activity detection in uncontrolled free-living conditions using a single accelerometer. In *Wearable and Implantable Body Sensor Networks (BSN), 2015 IEEE 12th International Conference on*, 1–6. IEEE.

[35] Lee, S.; Buxton, W.; and Smith, K. C. 1985. A multi-touch three dimensional touch-sensitive tablet. In *ACM CHI*, CHI '85, 21–25. New York, NY, USA: ACM.

[36] Li, D., and Chou, P. H. 2004. Maximizing efficiency of solar-powered systems by load matching. ISLPED '04.

[37] Li, X.; Allen, R.; J. Earley, C.; Liu, H.; E. Cruz, T.; Edden, R.; B. Barker, P.; and van zijl, P. 2016. Brain iron deficiency in idiopathic restless legs syndrome measured by quantitative magnetic susceptibility at 7 tesla. 22.

[38] Lin, K.; Hsu, J.; Zahedi, S.; Lee, D. C.; Friedman, J.; Kansal, A.; Raghunathan, V.; and Srivastava, M. B. 2005. Heliomote: Enabling long-lived sensor networks through solar energy harvesting. In *Proceedings of ACM Sensys*.

[39] Liu, J.; Chou, P. H.; Bagherzadeh, N.; and Kurdahi, F. 2001. Power-aware scheduling under timing constraints for mission-critical embedded systems. In *Proceedings of the 38th annual Design Automation Conference*, 840–845. ACM.

[40] Lo, P.-H.; Hong, C.; Lo, S.-C.; and Fang, W. 2011. Implementation of inductive proximity sensor using nanoporous anodic aluminum oxide layer. In *TRANSDUCERS*, 1871–1874. IEEE.

[41] Lo, P.-H.; Hong, C.; Tseng, S.-H.; Yeh, J.-H.; and Fang, W. 2012. Implementation of vertical-integrated dual mode inductive-capacitive proximity sensor. In *IEEE MEMS*, 640–643. IEEE.

[42] Lorch, J. 1995. A complete picture of the energy consumption of a portable computer. Technical report.

[43] Lu, T. 2013. A motion control method of intelligent wheelchair based on hand gesture recognition. In *ICIEA 2013*, 957–962. IEEE.

[44] M Kelly, J.; Strecker, R.; and Bianchi, M. 2012. Recent developments in home sleep-monitoring devices. 2012:768794.

[45] Maki, B. E. 1997. Gait changes in older adults: predictors of falls or indicators of fear. *Journal of the American geriatrics society* 45(3):313–320.

[46] Mantua, J.; Gravel, N.; and Spencer, R. 2016. Reliability of sleep measures from four personal health monitoring devices compared to research-based actigraphy and polysomnography. 16:646.

[47] Marie Trotti, L.; L Bliwise, D.; A Greer, S.; Sigurdsson, A.; Birna Gudmundsdottir, G.; Wessel, T.; M Organisak, L.; Sigthorsson, T.; Kristleifur, K.; Sigmundsson, T.; and B Rye, D. 2009. Correlates of plms variability over multiple nights and impact upon rls diagnosis. 10.

[48] Maurer, U.; Smailagic, A.; Siewiorek, D. P.; and Deisher, M. 2006. Activity recognition and monitoring using multiple sensors on different body positions. In *Wearable and Implantable Body Sensor Networks, 2006. BSN 2006. International Workshop on*, 4–pp. IEEE.

[49] Nandakumar, R.; Gollakota, S.; and Watson, N. 2015. Contactless sleep apnea detection on smartphones. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '15, 45–57. New York, NY, USA: ACM.

[50] National Instruments. Part II Photovoltaic Cell I-V Characterization Theory and LabVIEW Analysis Code. `http://www.ni.com/white-paper/7230/en`.

[51] Nelson, A.; Singh, G.; Robucci, R.; Patel, C.; and Banerjee, N. 2015. Adaptive and personalized gesture recognition using textile capacitive sensor arrays. *IEEE Transactions on Multi-Scale Computing Systems* 1(2):62–75.

[52] Ouh, H. K.; Lee, J.; Han, S.; Kim, H.; Yoon, I.; and Hong, S. 2012. A programmable mutual capacitance sensing circuit for a large-sized touch panel. In *IEEE ISCAS*, 1395–1398. IEEE.

[53] Palasagaram, J. N., and Ramadoss, R. 2006. Mems-capacitive pressure sensor fabricated using printed-circuit-processing techniques. *Sensors Journal, IEEE* 6(6):1374–1375.

[54] Pang, H., and Ding, Y. 2013. Dynamic hand gesture recognition using kinematic features based on hidden markov model. In *GCN 2012: Volume 5*, 255–262. Springer.

[55] Park, C.; Chou, P.; Bai, Y.; Matthews, R.; and Hibbs, A. 2006. An ultra-wearable, wireless, low power ecg monitoring system. In *Biomedical Circuits and Systems Conference, 2006. BioCAS 2006. IEEE*, 241–244.

[56] Pennestri, M.-H.; Whittom, S.; Adam, B.; Petit, D.; Carrier, J.; and Montplaisir, J. 2006. Plms and plmw in healthy subjects as a function of age: Prevalence and interval distribution. 29:1183–7.

[57] Pennestri, M.-H.; Montplaisir, J.; Colombo, R.; Lavigne, G.; and Lanfranchi, P. 2007. Nocturnal blood pressure changes in patients with restless legs syndrome. 68:1213–8.

[58] 2005. Sense: A wireless sensor network simulator. In Szymanski, B., and Yener, B., eds., *Advances in Pervasive Computing and Networking*, 249–267.

[59] Pering, T.; Agarwal, Y.; Gupta, R.; and Want, R. 2006. Coolspots: Reducing the power consumption of wireless mobile devices with multiple radio interfaces. In *Proceedings of Mobisys*.

[60] Pering, T.; Raghunathan, V.; and Want, R. 2005. Exploiting radio hierarchies for power-efficient wireless device discovery and connection setup. In *Proceedings of the IEEE International Conference on VLSI Design*.

[61] Picchietti, D.; Allen, R.; S Walters, A.; Davidson) Mount, J.; Myers, A.; and Ferini-Strambi, L. 2007. Restless legs syndrome: Prevalence and impact in children and adolescents the peds rest study. 120:253–66.

[62] Piorno, J.; Bergonzini, C.; Atienza, D.; and Rosing, T. 2009. Prediction and management in energy harvested wireless sensor nodes. In *Wireless VITAE 2009*, 6–10.

[63] `http://en.wikipedia.org/wiki/Hurricane_Sandy`. Hurricane sandy.

[64] `http://en.wikipedia.org/wiki/PIC_microcontroller`. PIC microcontroller.

[65] `http://www.bbc.co.uk/news/world-asia-pacific-12711226`. Japan earthquake.

[66] `http://www.gumstix.com/`. Gumstix platform.

[67] `http://www.mapnik.org`. 2011. mapnik.

[68] `http://www.analog.com/en/analog-to-digital-converters/ad-converters/ad7192/products/product.html`.

[69] `http://www.linear.com/product/LT1102`.

[70] `http://www.xilinx.com/univ/xupv5-lx110t.htm`.

[71] `http://en.wikipedia.org/wiki/K-means_clustering`.

[72] `http://en.wikipedia.org/wiki/PID\_controller`.

[73] `http://en.wiktionary.org/wiki/Manhattan\_distance`.

[74] `http://sine.ni.com/nips/cds/view/p/lang/en/nid/204239`.

[75] `http://www.pgrouting.org/`. PG Routing.

[76] Rahman, A. A.; Rashid, N. A.; Aziz, A. S.; Witjakson, G.; et al. 2012. Design of autonomous micro-solar powered energy harvesting system for self-powered batteries-less wireless sensor mote. In *Electronics Goes Green 2012+(EGG), 2012*, 1–4. IEEE.

[77] Rahmati, A., and Zhong, L. 2007. Context-for-wireless: Context-sensitive energy-efficient wireless data transfer. In *Proceedings of Mobisys*.

[78] Ripka, P., and Tipek, A. 2013. *Modern sensors handbook.* Wiley. com.

[79] Roa, G.; Le Pelleter, T.; Bonvilain, A.; Chagoya, A.; and Fesquet, L. 2014. Designing ultra-low power systems with non-uniform sampling and event-driven logic. In *Integrated Circuits and Systems Design (SBCCI), 2014 27th Symposium on*, 1–6. IEEE.

[80] Sadeh, A.; Pergamin-Hight, L.; and Bar-Haim, Y. 2007. Sleep in children with attention-deficit hyperactivity disorder: A meta-analysis of polysomnographic studies. 10:381–98.

[81] Sako, S., and Kitamura, T. 2013. Subunit modeling for japanese sign language recognition based on phonetically depend multi-stream hidden markov models. In *Universal Access in Human-Computer Interaction*. Springer. 548–555.

[82] Saponas, T. S.; Harrison, C.; and Benko, H. 2011. Pockettouch: through-fabric capacitive touch input. In *ACM UIST*, 303–308. ACM.

[83] Shen, H.; Balasubramanian, A.; LaMarca, A.; and Wetherall, D. 2015. Enhancing mobile apps to use sensor hubs without programmer effort. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '15, 227–238. New York, NY, USA: ACM.

[84] Shih, E.; Bahl, P.; and Sinclair, M. J. 2002. Wake on Wireless: An event driven energy saving strategy for battery operated devices. In *Proceedings of the Eighth ACM Conference on Mobile Computing and Networking*.

[85] Shnayder, V.; Hempstead, M.; Chen, B.-r.; Allen, G. W.; and Welsh, M. 2004a. Simulating the power consumption of large-scale sensor network applications. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, 188–200. ACM.

[86] Shnayder, V.; Hempstead, M.; Chen, B.-R.; and Welsh, M. 2004b. PowerTOSSIM: Efficient Power Simulation for TinyOS Applications. In *SenSys*.

[87] Siddiqui, F.; Strus, J.; Ming, X.; Lee, I.; Chokroverty, S.; and S Walters, A. 2007. Rise of blood pressure with periodic limb movements in sleep and wakefulness. 118:1923–30.

[88] Simakajornboon, N.; Gozal, D.; Vlasic, V.; Mack, C.; Sharon, D.; and M McGinley, B. 2003. Periodic limb movements in sleep and iron status in children. 26:735–8.

[89] Sorber, J.; Banerjee, N.; Corner, M. D.; and Rollins, S. 2005a. Turducken: Hierarchical power management for mobile devices.

[90] Sorber, J.; Banerjee, N.; Corner, M. D.; and Rollins, S. 2005b. Turducken: Hierarchical power management for mobile devices. In *Proceedings of ACM MobiSys*.

[91] Steigerwald, R.; Ferraro, A.; and Turnbull, F. G. 1983. Application of power transistors to residential and intermediate rating photovoltaic array power conditioners. *Industry Applications, IEEE Transactions on* IA-19(2):254–267.

[92] Stepanski, E.; Lamphere, J.; Roehrs, T.; Zorick, F.; and Roth, T. 1987. Experimental sleep fragmentation in normal subjects. 33:207–14.

[93] Storn, R. 1996. Differential evolution design of an iir-filter. In *Evolutionary Computation, Proceedings of IEEE International Conference on*, 268–273. IEEE.

[94] Tartar, J.; Ward, C.; Mckenna, J.; Thakkar, M.; Arrigoni, E.; McCarley, R.; Brown, R.; and Strecker, R. 2005. Hippocampal synaptic plasticity and spatial learning are impaired in a rat model of sleep fragmentation. 23:2739–2748.

[95] Terzic, E.; Nagarajah, R.; and Alamgir, M. 2011. A neural network approach to fluid quantity measurement in dynamic environments. *Mechatronics* 21(1):145–155.

[96] Tilma, J.; Tilma, K.; Norregaard, O.; and R Ostergaard, J. 2013. Early childhood-onset restless legs syndrome. symptoms and effect of oral iron treatment. 102.

[97] Titzer, B.; Lee, D.; and Palsberg, J. 2005. Avrora: scalable sensor network simulation with precise timing. In *IPSN 2005*.

[98] Toulouse, C.; Badaoui, M. E.; and Serviere, C. 2015. Gait analysis and falls estimation using the slope's variation of synchronous statistics. In *Advances in Biomedical Engineering (ICABME), 2015 International Conference on*, 293–296. IEEE.

[99] Tyvaert, L.; Houdayer, E.; Devanne, H.; Bourriez, J.; Derambure, P.; and Monaca, C. 2009. Cortical involvement in the sensory and motor symptoms of primary restless legs syndrome. *Sleep medicine* 10(10):1090–1096.

[100] Werner-Allen, G.; Swieskowski, P.; and Welsh, M. 2005. Motelab: a wireless sensor network testbed. In *Proceedings of the 4th international symposium on Information processing in sensor networks*, IPSN '05. Piscataway, NJ, USA: IEEE Press.

[101] Yazicioglu, R.; Torfs, T.; Penders, J.; Romero, I.; Kim, H.; Merken, P.; Gyselinckx, B.; Yoo, H.; and Van Hoof, C. 2009. Ultra-low-power wearable biopotential sensor nodes. In *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, 3205–3208.

[102] Yerva, L.; Campbell, B.; Bansal, A.; Schmid, T.; and Dutta, P. 2012. Grafting energy-harvesting leaves onto the sensornet tree. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, IPSN '12, 197–208. New York, NY, USA: ACM.

[103] Yin, L.; Dong, M.; Duan, Y.; Deng, W.; Zhao, K.; and Guo, J. 2013. A high-performance training-free approach for hand gesture recognition with accelerometer. *Multimedia Tools and Applications* 1–22.

[104] Yuan, W., and Nahrstedt, K. 2003. Energy-efficient soft real-time CPU scheduling for mobile multimedia systems. In *Proceedings of the Symposium on Operating Systems Principles*, 149–163.

[105] Zhang, H.; Salajegheh, M.; Fu, K.; and Sorber, J. 2011. Ekho: bridging the gap between simulation and reality in tiny energy-harvesting sensors. HotPower '11.