A Case for Application Driven Design of Energy Harvesting Sensor Systems

Neal Jackson

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2022-268
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2022/EECS-2022-268.html

December 16, 2022

A Case for Application Driven Design of Energy Harvesting Sensor Systems

by

Neal Schadewald Jackson

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Electrical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Associate Professor Prabal Dutta, Chair
Professor Kristofer Pister, Professor Stefano Schiavon

Fall 2022

A Case for Application Driven Design of Energy Harvesting Sensor Systems

Abstract

A Case for Application Driven Design of Energy Harvesting Sensor Systems

by

Neal Schadewald Jackson

Doctor of Philosophy in Electrical Engineering

University of California, Berkeley

Associate Professor Prabal Dutta, Chair

Any problem becomes tractable with enough power. Since the inception of wireless sensor networks, researchers have searched for ways to do more with less. Integrated circuits and sensors have continued to shrink in size, cost, and active and quiescent power. This has resulted in sensors with increasing computational power and longer lifetimes. By comparison, however, the options for and quantity of power available for a wireless sensor has stagnated. The energy density of non-rechargeable batteries as well as photovoltaic efficiency have approached a plateau. As a result, wireless sensors are constrained in either lifetime or power.

Technology improvements like integrated circuits that perform efficient maximum power tracking and voltage boosting, near-threshold computing, the advent of non-volatile memory technologies, and the rapid improvement of supercapacitor technology has enabled the development of sensors that can operate entirely on harvested power without batteries, and without a finite lifetime. But the lack of a reliable power source necessarily results in a system that is fundamentally unreliable. Despite this significant flaw, researchers have pursued this design archetype tirelessly, producing an impressive corpus of methods, systems, and solutions that attempt to improve batteryless design. Proponents of batteryless systems are convinced that batteries are a threat to the future of wireless sensing, and that batteryless sensing is the only way forward. Despite this, batteryless sensing has not seen widespread adoption by industry. There is a rift of design understanding between those who value reliability, and those who do not.

The core argument of this dissertation is that there is not a single design dogma, be it batteryless or battery-powered, that can provide a solution for *all* applications. Instead, the correct design process must involve a balancing act of the inclusion and sizing of energy harvesting, rechargeable, and non-rechargable energy storage to meet the goals of the application. This design space is large and difficult to navigate, resulting in many system designers defaulting to following a predetermined design template archetype instead of fully reasoning about their application and its requirements. In this dissertation, we develop a design framework for

energy harvesting systems that provides reasoned guidance for the inclusion and sizing of various power supply elements. In particular, we develop analytical and simulation tools to size rechargeable energy capacity in a more reasoned way than current heuristics and arbitrary methods.

To develop this design framework, this dissertation explores previous wireless sensor applications, identifying the appropriateness of different approaches qualitatively and quantitatively. We explore the system-level effects of harvester size and rechargeable and non-rechargeable energy capacity on wireless sensor application performance. To determine rechargeable energy capacity selection and sizing, we develop a novel heuristic for determining the minimum sufficient capacity for a sensor workload and expected energy income. We verify this heuristic through the use of a custom wireless sensor energy state simulator to estimate energy utilization and system performance. To identify technology options for energy capacity, we quantitatively compare energy buffer types and reevaluate the many qualitative claims made against rechargeable batteries by batteryless proponents, concluding that many of them are without merit. Finally, we utilize the design framework developed within this dissertation, including the heuristics and simulation tool, to design and implement wireless sensor systems to address two real indoor sensing applications that achieve long-lived operation with consistent and reliable sensing.

To all those who think they cannot:
"Sucking at something is the first step towards being sorta good at something"
— Jake the Dog

Acknowledgments

At the end of my undergraduate career, I was convinced pursuing a PhD was not for me. At the time, I did not think I had the drive or intellect to be successful in a graduate program. It took the support of many people to not only change my mind, but also get me through the last six years. I am so thankful and glad to have them in my life.

My parents, Sue Schadewald and Brad Jackson have always stressed the value of education. They have always supported me in my endeavors, scholastic or otherwise. My mom is one of the first people I will call for advice. She has helped me navigate my relationships and growing into being an adult with adult responsibilities. My dad is the first person I go to for help. He has driven with me across the country twice, first for an internship in Oregon, and second for my program in Berkeley. I thank my parents for the joy I find in building, fixing, and improving things, my sense of humor, and providing the example of love and selflessness that I strive to emulate. My younger sister Sadie Jackson has almost beat me to being the first Dr. Jackson, despite being three years younger. She is someone I look up to and I thank her for being one of my best friends.

My new family, my partner Taryn Farber, and our pets Aspen and Bella, have been a constant source of comfort and support. Having been through her own graduate struggles, Taryn understands and empathizes when I am struggling. Aspen and Bella do not, but they are always up for giving affection and love whether I need it or not. My childhood best friends Matt Peat and Trent Delehanty remain best friends today, and have offered a lifetime of advice and support. I thank my college friends Jason LaFave and Claude Christensen for being the family away from home then, and still best friends now.

Finally, I thank my friends and colleagues of Lab11, who, through their work, culture, and camaraderie set the benchmark for what I envisioned a research lab group should be. When visiting different potential graduate schools and lab groups, Lab11 was the standard by which I judged them. Nowhere else could compare. I have Josh Adkins to thank for being my brilliant and aloof roommate and co-founder of the Lab11 Berkeley Annex. It was a rare and wonderful journey being the first of Prabal's Berkeley students, setting up a new lab, doing class projects, and discovering the wonders of the Bay Area. Josh's drive to build things and make a difference kick-started my own enthusiasm. Will Huang and Jean-Luc Watson have been some of my best friends through graduate school. I thank Will for his tireless support, as he has undoubtedly helped me more than anyone with my research, writing, and fighting for bounty in the bayou. I thank Jean-Luc for keeping lab culture alive with game nights and Utah "scones." I also thank Branden Ghena, the unofficial "Director of Undergraduate Researchers" for introducing me to the lab and being the most approachable and patient graduate student I have ever met. His example led me to be a more patient and joyful teacher. I thank Meghan Clark and Pat Pannuto for being stand in advisors as well as friends. I thank Noah Klugman and Lane Powell for their kindness, honesty, randomly dropping by my apartment, and for being my dog's best friends. I thank Brad Campbell for the big beers and Deepika Natarajan for being a consistent friend since 470. I thank Thomas Zachariah for being my TA at Michigan, as well as a travel buddy in Hong Kong. Lastly,

I thank Andreas Biri, Matt Podolsky, Shishir Patil, Alvin Tan, and Tess Despres for their feedback and support. Getting through graduate school was only possible with your support.

I still remember my last October as an undergraduate when Professor Prabal Dutta persuaded me to pursue a PhD. He knew that I had what it took to be successful, even though I was doubting myself. I have him to thank, not only for persuading me to try, but also for his commitment to seeing me through all the way to graduation. Lab11 set the standard for a creative and productive research group, and Prabal set the standard for a kind, empathetic, and inspiring advisor. He is someone that will fight for his students, in academia as well as when they get caught on CCTV "improving" his faculty portrait in the Cory Hall photo case. I also thank the other members of my committee: Professors Kris Pister and Stefano Schiavon for their insight and suggestions to make this dissertation the best version I could have created. I am a better engineer, researcher, and person having had their guidance.

This work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, and with support from the Building Technology Urban Systems group at Lawrence Berkeley National Laboratory (LBNL). Additionally, this material is based upon work supported by the National Science Foundation under grant numbers CNS-1824277 and DGE-1106400, and NSF/Intel CPS Security under grant CNS-1822332.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The past two decades have witnessed the proliferation of low power, autonomous, and wirelessly connected devices. The development and deployment of these devices were inspired by the vision of ubiquitous computing: computers that are pervasive in our world, providing information about and intuitive control over our environments in such a way that they disappear into the background [1]. While initially inspired by this vision, today's internet of things (IoT) has failed to fully achieve it. There are many problems standing in the way of this vision, but one of the most pressing and fundamental problems is power. Low power embedded systems have experienced exponential increases in computational and power efficiency in the last two decades. Across the board, components used to build wireless systems have increased in power efficiency by one to two orders of magnitude. Our ability to process and infer meaning from vast amounts of collected data has also improved by orders of magnitude, with approachable and powerful frameworks for developing effective machine learning models for classification and detection. However, the methods of collecting and allocating energy to power wireless sensing systems have not improved at the same rate. So, while today's small wireless sensing devices have become smaller, more efficient, and more capable of powerful and complex sensing and data processing, their capabilities are still limited by the power and energy available to them.

1.1 The Power and Energy Dilemma

Perhaps the most important step in the wireless sensor design process is determining how to allocate or capture the energy to power it. Nowadays, there are many options for energy storage: capacitors, supercapacitors, batteries. There are even more sources for energy harvesting, including photovoltaic panels, thermoelectric and piezoelectric generators, and microbial fuel cells [2–5]. This results in a vast system-level design space that can be difficult to navigate. To reduce the design space problem to something comprehensible and manageable, it is reasonable to consider existing industrial and research designs as templates for a design. We may look to modern commercial sensors for inspiration, like the Google Nest temperature

sensor [6], or to a modern research device like the Flicker platform [7]. While these two examples use similar sensors, microcontrollers, and radios, and are designed for a similar sensing purpose, they exhibit disparate differences in their power supply design. These two different designs represent a fork in the road for wireless system design.

Like most modern commercial sensors, the Nest temperature sensor utilizes a non-rechargeable battery for an energy source [6]. Batteries have been the preferred method of powering sensors for the last twenty years, and for good reason: they are simple to use and provide reliable and predictable power. However, a battery provides a finite lifetime and battery or sensor replacement is inevitable. Wireless sensor lifetime is a first order concern, as the cost of frequent maintenance and battery replacement is clearly untenable as the number of wireless devices grows to trillions of devices. The act of maintaining what should be an invisible sensor renders it visible due to the annoyance of changing a battery. Changing a single battery is not usually an issue. Still, it is not uncommon for people to simply ignore a dying device like a smoke detector, opting to remove the battery instead of replace it when low. Changing the battery of hundreds or thousands of devices within a building is a significant and costly undertaking.

Recognizing the problem with short, finite lifetimes, researchers have forged ahead with designs that aspire for immortality. For a decade, many researchers have abandoned batteries completely and instead built systems that harvest energy to power themselves. The Flicker platform is an example of a batteryless sensing platform [7]. Sensors built with this batteryless technique possess indefinite lifetimes. The cost of immortality is steep, however. These systems generally require significantly more software and hardware design complexity, and operate unpredictably and without any quality of service guarantees. Due to this, energy harvesting systems without batteries have seem limited adoption by industry, despite significant excitement and a large corpus of work by researchers.

1.2 The Difficulty of Gathering Data

In reality, neither option on its own, battery or energy harvesting, is an entirely satisfactory solution for many sensing applications. Modern data gathering applications often simultaneously demand longevity and consistent operation. It is often not worth the cost of deployment and frequent maintenance if sensors only last a short while, or if the data they produce is not ultimately useful for the end goal of the data gathering application.

Traditional environmental sensing applications like wildlife monitoring and tracking require consistent environmental and location data gathering over seasons and years [8, 9]. Asset tracking applications have to be able to reliably locate and track commodities throughout an entire supply chain, over weeks or months in a warehouse, or on store shelves [10]. Infrastructure monitoring applications must consistently and reliably detect vibrational anomalies or corrosion throughout the multi-decade lifespan of a bridge or a building [11, 12]. All of these example applications require consistent and reliable data gathering over long periods of time.

New applications are made feasible with new and improved technology, but they will be still be limited by the longevity supplied by a non-rechargeable battery or the consistency from energy harvesting options. As an example, person detection and room occupancy counting is a useful metric for building lighting and HVAC control. Having a reliable occupancy measurement can increase efficiency as the building management system can provide light and comfortable heating tailored to where people exist within the building. Capturing, manipulating, and transmitting image data requires significant computational, memory, and networking resources. Even with advances in embedded computing and CMOS imaging, these resources are generally costly in power compared to the traditional, simpler one-dimensional data collection of traditional environmental sensors [8, 13]. If designed with a battery, such an imaging system would possess a short lifetime. If designed to be energy harvesting without a battery, the sensor can provide no guarantees for when or how many images are captured, potentially failing to perform the designed purpose: reliably detecting when and how many people are present in a space.

This choice between a battery or batteryless energy harvesting solution is a false dichotomy. This choice is predicated on the assumption that a sensor design must be exclusively one or the other: battery-based or energy harvesting. Many recent energy harvesting systems are built with the assumption that batteries cannot be trusted, both non-rechargeable and rechargeable [2, 7, 14–20]. This belief has led many designers to consider batteryless energy harvesting designs as the default option for any new application. By defaulting to one particular design archetype, designers are forced to make seemingly arbitrary and ill-considered design decisions regarding the typing and sizing of components within their power supply in order to achieve minimally feasible batteryless designs. Instead of thinking of wireless sensor power supply design as adhering to a predetermined class or pattern of design, an alternate and more successful design process can consider an application's requirements and which design options are appropriate to satisfy those requirements.

1.3 Thesis Statement

An analytical and simulation driven design framework provides better guidance for system-level power supply design than current energy harvesting intuition. We can use design points determined through design space analysis and simulation to build systems that provide higher availability, lower latency, and long lifetimes, all without the need for complex techniques for managing intermittency.

1.4 Contributions of this Dissertation

This dissertation attempts to address system-level power supply design in the context of real application requirements. Various parts of this dissertation have been published at ENSsys'18 [21] and IPSN'19 [22]. Content from these publications are included throughout

the entire document, including Chapter 3, Chapter 4, and Chapter 5. This work was developed in collaboration with Joshua Adkins and Prabal Dutta.

We begin with an examination of existing power supply designs and methods for providing energy to sensors. In Chapter 2, we discuss traditional battery-based sensor design as well as the history of energy harvesting in wireless sensor research. We explore the design and trade-offs of batteryless system design, a newer class of energy harvesting sensors that forgo the use of batteries as energy storage in favor of capacitors and supercapacitors. We identify the strengths and weaknesses of both battery-based and batteryless design points within the context of real applications from research and industry. From this exploration, we identify several common high-level application requirements that drive wireless sensor design.

In Chapter 3, we utilize these requirements to develop system-level design constraints and heuristics. We consider and compare the efficacy of batteries and energy harvesting and identify inflection conditions where either option excels over the other. We identify the benefits of a rarely used hybrid architecture that employs both options and consider the performance and lifetime implications of a system that utilizes it. Next, we define design constraints regarding energy income and energy buffer size that partition the design space into several regions to reason about the necessity of batteryless and intermittent techniques. We finish this chapter with a deep dive into the effect of rechargeable energy capacity on the performance of an energy harvesting system. From this exploration, we utilize a simple model of energy capacity to develop novel and easy to use design heuristics. These heuristics determine the minimum sufficient sizing for energy capacity for energy harvesting systems.

We expand upon the simple model and heuristics developed in Chapter 3 and build an wireless sensor energy state simulation in Chapter 4. We use this simulation to verify our capacity sizing heuristics, consider additional design variables, and examine the performance of the hybrid battery and energy harvesting architecture explored in Chapter 3. From the results of our energy simulation, we identify the benefits in energy capture and system performance when rechargeable energy capacity is increased. From these results, we conclude that energy harvesting system performance can be increased by 1.4–2$\times$ if capacity is increased several orders of magnitude more than is commonly offered by the capacitor or supercapacitor energy buffers commonly utilized by batteryless systems.

In Chapter 5, we explore the design space of rechargeable energy capacity, identifying the classes and options for charge and energy buffers. We compare capacitors, supercapacitors, and batteries. We examine the common arguments made against the use of rechargeable batteries by batteryless researchers. We find that many of these arguments are unsubstantiated when examined quantitatively and conclude that small rechargeable batteries are the superior method for providing an energy dense buffer for low power applications.

We consider the heuristics, tools, and conclusions from Chapters 3 and 4 to implement sensors for two wireless sensing applications in Chapter 6. These applications include illuminance sensing for automated lighting control applications as well as image-based person detection and counting for accurate space occupancy measurements, both of which require consistent and long-lived sensing. Both sensor designs utilize the hybrid architecture identified in Chapter 3, and we utilize the heuristics and simulation tools developed in Chapters 3 and 4

to size their power supply components to achieve high energy capture and system performance. We implement and utilize our designs to evaluate our energy simulation developed in Chapter 4. We evaluate an existing batteryless image sensing platform, quantify its performance using simulation, and identify design changes that could be made to improve performance. From these design changes, we build a new indoor image sensing platform that can capture image data with high consistency and availability and still offer a multi-year lifetime.

Finally, in Chapter 7, we summarize the contributions of this dissertation and identify several primitives and architectural changes that would enable more efficient and autonomous energy harvesting wireless sensors.

Chapter 2

Background

System researchers have been developing wireless sensors for over twenty years. Sensor systems research is uniquely application focused, with many seminal works involving real applications and deployments [8, 9, 23]. A lot of other research has been focused on developing and utilizing new and improved technology for wireless sensor designs. Over the past two decades, microcontrollers have vastly improved in processing speed and capability, sensors have continued to shrink in size, and wireless communications have increased throughput and range. Across the board, all wireless sensor components have also substantially increased their energy efficiency. However, the availability of power and energy remain limiting factors for wireless sensor designs.

The energy density of non-rechargeable batteries has not improved at the same rate as as other wireless sensor components. While newer, more efficient sensors are able to do more with the limited energy available to them, battery-based sensors are still limited in lifetime and have the potential to produce substantial battery waste. For longer deployments, or in applications where the size and weight of a non-rechargeable battery was untenable, researchers have also developed energy harvesting solutions for wireless sensors. Energy harvesting sensors do not have concretely limited lifetimes, but they are inherently limited by the availability and consistency of harvestable power. Like battery technology, the efficiency of energy harvesting methods have not increased at the same rate as other sensor technology, and specialized energy harvesting power management ICs are already highly optimized to extract energy from many sources [24–26], with few exceptions [5, 27]. System designers are left with two unsatisfying options. Depending on an application's lifetime, maintenance, and quality of service requirements, batteries may offer insufficient longevity, and energy harvesting may allow insufficient availability. This chapter seeks to explore the myriad of traditional and contemporary methods of powering sensors and their trade-offs with the goal of identifying common application requirements and how various power supply designs satisfy those requirements.

2.1 Methods to Power Wireless Sensors

All wireless sensors require electrical energy to function. Most simply and most commonly, energy can be provisioned in a finite energy storage such as a battery, also known as a primary cell. In situations where there is harvestable energy, it can be captured and stored in rechargeable storage, or a secondary cell. The next sections describe prior work in both research and industry regarding preallocation and energy harvesting wireless sensor design.

Energy Preallocation

Primary-cell batteries are the preferred method of powering sensors for both academic experimentation and commercial and industrial applications. The Telos family of motes, originally designed in 2004 [28], are still the wireless platform of choice for some modern research projects [29, 30]. The Hamilton mote is a more modern example that seeks to provide a cost-effective and longer-lived alternative to older motes [31]. Besides research platforms, the majority of commercial smart home sensors, like those offered by Ecobee, Honeywell, Lutron, Nest, Phillips, among many others, all opt to use non-rechargeable batteries as their source of energy [6, 32–35]. Industrial offerings from Emerson, GE, Honeywell, and others mostly utilize non-rechargeable power cells in their wireless sensors [36–38]. The use of primary-cells is popular in commercial and industrial sensing because they enable sensors with predictable lifetimes that are easy to design, simple to program, and reliable to operate. However, a finite energy storage provides a finite lifetime, meaning battery replacement is inevitable. To achieve a long lifetime, sensor designers must often sacrifice constraints on size to accommodate a larger battery. However, advances in energy efficiency and battery longevity have resulted in reasonably sized commercial and industrial sensors that can last up to 10 years without battery maintenance [34, 36, 38].

Energy Harvesting

Instead of preallocating energy, a system can utilize external sources of energy. A system that harvests energy is not as limited in lifetime as a primary-only system, and depending on the durability and longevity of its harvester and energy storage, can persist indefinitely. However, the quantity and consistency of external energy can vary widely. If the energy is predictable or reliable, an energy harvesting system can operate reliably. If the source is unpredictable and insufficient to power a system continuously, the system may operate inconsistently if it does not have sufficient energy storage to buffer in times of insufficient income. Capacitors, supercapacitors or batteries are all options to buffer energy when external energy is variable.

The majority of energy harvesting wireless sensors depend on unreliable external sources of energy. They utilize photovoltaic, thermoelectric, piezoelectric, ambient RF harvesting, or other methods to scavenge energy from their environment. Parameters like harvester size or surface area, impact the amplitude of power delivered to the sensor, while the size and capacity of the rechargeable storage determines how much energy can be buffered.

Energy harvesting sensors have largely been developed for use in environments with plentiful harvestable energy. Most examples of energy harvesting research devices are deployed outdoors with photovoltaics [39–44]. This is also true of many commercial and industrial sensor systems like weather and air quality monitoring stations [45], trail cameras [46], and traffic cameras [47]. Besides photovoltaic harvesting, industrial products have also leveraged temperature gradients and vibration to power sensors. The Perpetua Power Puck and Tile are thermoelectric generators designed for high temperature monitoring of critical infrastructure like steam pipes, hot tanks, and other equipment in high temperature environments [48]. Perpetua harvesters are compatible with the same Emerson, GE, and Honeywell sensor systems as a drop in replacement for the default non-rechargeable power supplies. Other companies, like ReVibe and Kinergizer, are utilizing piezoelectric harvesting for industrial vibrational monitoring [49, 50].

There have been fewer research, commercial, and industrial sensors developed for indoor environments, or other environments that lack light, a large temperature gradient, or perpetual motion and vibration. This is because it is more difficult to design a lower power system to match a lower power income, even with newer and more efficient technology. Despite these difficulties, researchers and industry have developed sensors for applications with limited access to harvestable energy. Researchers have followed two distinctly different approaches for indoor energy harvesting.

Some of the first attempts at bringing energy harvesting indoors consisted of power supply architectures similar to those of their outdoor counterparts, in that they utilized rechargeable batteries for energy storage. The EnHANTs sensor used an indoor photovoltaic panel to charge an intentionally oversized nickel-metal hydride (NiMH) battery, with plans to eventually use a thin-film battery [51, 52]. DoubleDip utilized thermoelectric harvesting to charge a lithium-manganese battery [53]. The batteries used by EnHANTs and DoubleDip are examples of older technology that had numerous limitations. Early rechargeable batteries offered short cycle lifetimes and low charge and discharge current capabilities. The cycle lifetime limit results in systems with limited longevity, even though they utilize energy harvesting.

Simultaneously, researchers at University of Washington and Intel Research Seattle began experimenting with computational RFIDs (CRFIDs) to create battery-free sensors [54]. These researchers noted the limited storage of non-rechargeable batteries, and the limited cycle lifetime of rechargeable batteries. Instead of batteries, these systems utilize small capacitor-based energy buffers that are able to store just enough energy to complete a small atomic task, be it operating a sensor, transmitting a packet via RFID backscatter, or performing some amount of computation. While the longevity of CRFIDs are not limited to to the lifetime of a battery, they require the proximity of an RFID reader to provide power and bidirectional communication. Since the development of CRFIDs, and as technology has continued to improve in efficiency, researchers have extended the technique beyond RFID to systems with active radios and other harvesting methods, including photovoltaic, piezoelectric, and thermoelectric harvesting [2, 7, 18, 55, 56]. The majority of modern wireless sensor platforms build by researchers are batteryless, with many convinced that batteryless designs are the future for wireless sensor power supply design [7, 17, 18, 57–59]. Despite the excitement

around batteryless designs, they exhibit some serious detracting qualities that may limit feasibility and adoption for many applications. The next section explores the design and development of batteryless systems in more detail and describes the benefits and trade-offs of the design.

2.2 Batteryless Energy Harvesting

The batteryless, or intermittent, sensor movement has abandoned batteries (both rechargeable and non-rechargeable) under the assumption that current battery technology has too many detracting qualities to be suitable for energy-harvesting wireless sensors. Most notably, they argue that rechargeable batteries provide insufficient lifetimes to build long-lasting deployments [2, 7, 15–20]. Batteryless systems instead utilize capacitors, some types of which offer functionally infinite lifetimes [60]. However, capacitors and supercapacitors provide much less energy storage compared to batteries. Due to this, a batteryless design results in two major drawbacks: at any given time, a batteryless system is limited to the energy provided by the short discharge cycle of its capacitor bank, and the availability of the system is determined by the consistency and intensity of energy income.

A batteryless system can technically perform the operations of a wireless sensor but it may not perform them well. When these systems are harvesting enough power to turn on and operate, they can only perform operations that require less energy than their capacitor storage can hold. Often a batteryless system's income power is intermittent, resulting in a system that operates intermittently. When harvestable energy is unavailable, batteryless systems quickly deplete their small energy stores, and lacking any future income, they power off and lose volatile state, potentially in the middle of an important operation and for an extended and unknown period of time. The intermittent reality of many harvesting sources necessitates careful management of energy and detailed and thorough software optimization to ensure that *any* operations can be completed.

Despite these complications, the concept of an immortal wireless sensor is tantalizing, and is the driving motivation behind batteryless systems research. This has resulted in a wealth of batteryless systems research, the majority of which is focused on developing software solutions that manage and preserve volatile state across power failures. State preservation has the potential to extend the runtime of sensor workloads beyond that of a single capacitor buffer discharge cycle, and is useful for building more general and capable systems. Batteryless devices are also difficult to debug and develop software for, as in addition to software bugs, energy is no longer guaranteed at any point during execution. To this end, researchers have also built tools to recreate energy conditions to help diagnose and fix intermittent energy bugs. In addition to software solutions, researchers have designed hardware platforms that maximize individual component availability, as well as platforms that dynamically tune capacitance to meet individual operation energy requirements. The next sections discuss various system designs, tools, and software and hardware techniques developed to alleviate the drawbacks to batteryless systems.

Forward Progress and State Retention

The type and amount of work that is possible within the energy envelope of a single discharge cycle of capacitor or supercapacitor buffer is severely limited. Batteryless systems are generally only able to perform a few seconds or less of computation, or send a single packet before depleting their capacitor storage and powering off, even with modestly sized supercapacitors [2, 7, 18, 56]. This section discusses the techniques developed by batteryless researchers to perform interesting sensing under unpredictable power loss.

One-shot Intermittency

One of the simplest methods employed by batteryless systems does not bother to attempt to retain state across power failures and reboots. One-shot intermittent designs instead just allocate enough capacitance to turn on and perform a simple predefined task, like sending a packet. This method is reminiscent of the simple reply behavior of RFID tags, upon which the design of early intermittent systems is based on [54]. These systems may act as a simple beacon [61, 62], or a sensor [2, 3, 55, 63]. For some platforms, most notably the Monjolo family of devices, the rate of harvesting is used as the sensor itself. Every wake-up and transmit event corresponds to an amount of energy harvested, and can be used to quantify the harvested phenomenon [3, 55, 63]. This method of coupling harvesting and sensing is the only way to ensure that interesting sensor data is observed by a batteryless sensor. Otherwise, when a batteryless sensor is built to sense a phenomena independent of the harvesting method, it is inevitable that changes in the sensed phenomena will be missed because the sensor is offline. Even when sensing and harvesting is coupled, a one-shot batteryless design often does not have the energy required to support message retransmission in the case of packet loss, or generally support any network reliability mechanisms. Any observation may be lost upon packet reception failure, and the next time the sensor wakes up it will only have enough energy to be concerned with transmitting its most recent observation. This one-shot method is relatively simple compared to other approaches to batteryless software, but it still requires tedious cold-start software optimization and capacitance tuning to allow the device to power on and complete its workload within the constraints of its tiny energy storage. This severely limits the generality of systems developed with this approach. Performing any sensing or computing outside of the hardware's intended use case is often not possible. Like all batteryless sensors, it is also impossible to distinguish between sensor failure and a lack of energy.

Checkpointing

Given the limitations of the one-shot technique, researchers have developed tools and techniques for ensuring forward progress across system reboots to enable longer and more complex workloads than a single capacitor bank can support. One method of enabling forward progress is checkpointing, where important volatile state is saved at predetermined and run time locations in code prior to a power outage, and restored upon rebooting. Many methods have been proposed, using both static and dynamic methods. Mementos utilizes a modified LLVM compiler, as well as a run-time library, to automatically place checkpointing triggers within a program [64]. During run-time, at these trigger points, Mementos measures the system voltage, and if low enough, portions of volatile state are written into non-volatile flash before turning off. The introduction of new non-volatile technologies like Ferroelectric RAM (FRAM) allows more energy- and time-efficient checkpointing, while simultaneously simplifying the state retention logic. Hibernus does not use predefined trigger points, and instead utilizes a hardware interrupt to detect a low system voltage threshold, and immediately copies volatile state to FRAM [65]. This state is subsequently restored on a rising voltage interrupt. Checkpointing is an effective method for enabling forward progress. However, software development with checkpointing can be very difficult to develop and debug, and writing and restoring state can become prohibitively expensive in time and energy as the amount of volatile state increases.

Batteryless Debugging

For batteryless systems, it is very difficult to properly and correctly develop software, especially when non-volatile memory like FRAM is involved. One-shot system software is difficult to tune and optimize to complete all within the energy provided by capacitor storage. For checkpointing systems, if checkpoints are improperly placed such that atomic blocks are not completed entirely before power loss, and yet their partial state is recorded in non-volatile memory, it can result in inconsistencies between volatile and non-volatile state. Software development and debugging for batteryless systems must also consider energy state. To this end, researchers have developed hardware-software debugging tools that emulate and replay energy state based on energy trace captures. Ekho is an energy emulator that recreates and repeats energy harvesting conditions for low power energy harvesting devices [66]. Another tool, the energy-interference-free debugger (EDB), also replays energy state, specifically for intermittent systems [67]. EDB augments normal interactive debugging tools like GDB to allow real time playback of energy state, in addition to interactive program debugging. While these tools are useful and necessary for ensuring correctness in batteryless development, they do introduce substantial complexity on top of an already complex embedded software debugging and development toolchain.

Safe Forward Progress

While conceptually simple, checkpointing is not a perfect solution to saving state, especially as that state begins to increase in footprint with increasing application complexity. There is an upper bound to the size of checkpoints, where the energy required to save and read state to and from non-volatile memory exceeds that offered by various capacitor storage sizes. The advent of FRAM (as well as other non-volatile RAM technologies) results in methods that can avoid explicitly saving and reading back the majority of application state. State that would normally reside in volatile RAM can now be placed within non-volatile memory, allowing this state to persist across power outages. While this seems like a straightforward solution to preserving state, it is complicated by the shrinking boundary between volatile and non-volatile state. Often both SRAM and FRAM are memory mapped and accessed the same way. Writes and reads to either are indistinguishable to the programmer and compiler. Code is often executed multiple times through the course of intermittent execution, and if volatile and non-volatile state are not carefully managed and accessed by the program, re-execution of various code sections can result in consistency violations [68]. The aftermentioned debugging tools can help to detect and fix intermittency bugs during runtime testing. However, many researchers have developed solutions to avoid these bugs by design. Task-based programming frameworks and models like Chain [69] and Alpaca [68] ensure forward progress and prevent intermittent consistency bugs. The task model of execution splits a software application into idempotent code blocks that read inputs from other tasks, and upon completion write their own results to channels of non-volatile memory. This has the effect of increasing task throughput in addition to ensuring safe forward progress.

The Wrong Abstraction

While effective in ensuring safe forward progress, all of these techniques introduce significant software and compiler complexity. They also necessarily limit programming scope to defined frameworks and concepts to ensure various constraints and guarantees of safety. Batteryless programmers must learn and adopt new programming models that incorporate potentially unfamiliar concepts like tasks, channels, and flows [16, 68, 69]. Once programmers have mastered the new patterns and methods of batteryless programming, intermittency is largely hidden from them. This is by design and one of the main goals of these programming models and frameworks. However, the abstraction of continuous progress is more harmful for application developers than helpful. This abstraction provides the illusion of continuous computing progress in processor time, while robbing programmers and users of any meaningful notion of progress in wall clock time. Continuous operation of a batteryless systems is often impossible, and software methods for forward progress only guarantee *eventual* task completion. While these techniques may make it easier to program batteryless systems, the abstraction of continuous forward progress is misleading regarding the timeliness and availability of systems that employ them. These techniques do not address the needs of any application that requires availability or timeliness of data. System designers and programmers are left without a way

to reason about and quantify the actual performance of applications built on top of these abstractions that mask inherent system unreliability.

Timeliness and Availability

While convenient during programming, the abstraction of continuous operation without a notion of real performance does not result in responsive and available batteryless systems. The task-based models ensure safety and improve performance over checkpointing-based approaches, but they do not consider or provide solutions for increasing batteryless system responsiveness and availability, two qualities that are especially important for wireless sensing. Timeliness and availability are goals that are generally at odds with the reality of batteryless sensing, as no guarantees can be made about energy income and thus the timeliness of event completion or system availability. A batteryless system's intermittent execution can span unknown periods of time, where it iteratively works on a task whose results may be stale and no longer useful by the time it finishes. There is also no way to ensure that a batteryless system is awake to witness a given event. Despite this, researchers have developed both software and hardware solutions in attempts to improve batteryless responsiveness and availability.

Software Timeliness and Availability

Mayfly [16] and InK [70] are task-based software frameworks like Chain and Alpaca. In addition to ensuring safe forward progress, they also attempt to maximize system timeliness and availability respectively. Mayfly associates tasks and the data they produce with deadlines, and maintains time through power failures. After a datum's deadline passes, it is discarded and energy is not used to further process or transmit it. This approach prioritizes fresh data, and allocates more energy towards completing tasks that can be completed within their deadline. This enables applications that depend on timely data, such as activity recognition. Alpaca and Chain are unsuitable for such applications, as they are unable to distinguish between expired irrelevant data and new fresh data. While it increases the timeliness of a system, Mayfly does not increase the amount of reported data by a system or its availability, as it cannot increase the energy available to the system. By design, it must sacrifice stale, potentially important data in order to report fresh data. It is uniquely suited for applications that are only concerned with close to real time data.

InK prioritizes availability by incorporating timers, interrupts, and event handling into a batteryless context. The runtime utilizes sleep states and low power interrupts to persist for longer periods and increase the likelihood of detecting and capturing events. It also utilizes an ultra-low power timer subsystem to ensure timekeeping across power failures. With these improvements, a system using Ink is able to capture 14x the number of events over the same system using Alpaca. While this seems like a dramatic increase, Alpaca was not designed with event detection workloads in mind. It performs abysmally when applied to such tasks. With controlled solar harvesting conditions, a system running on Alpaca was only able to

capture less than 5% of events. Comparatively, the system with Ink was only able to capture slightly over half of the events even when in the same light conditions [70]. The authors did not compare the performance of these runtimes with that of a traditionally programmed and continuously powered embedded system, so there is no baseline comparison.

While these approaches increase the timeliness and availability of a system, it is *impossible* for a batteryless system without continuous energy income to ever approach 100% event detection or data capture and transmission reliability. There will be always be periods of time where there is no available energy, a batteryless system will remain off, and important events will remain undetected and important data will be discarded as it grows stale.

Hardware Hysteresis Management

Beyond software techniques, researchers have also explored building specialized hardware solutions to increase system availability. Batteryless system operation is generally tied to a full swing capacitor charge/discharge hysteresis. Once enough energy is captured, a batteryless system will begin executing its workload until it has exhausted its stored energy. The general rule of thumb for batteryless system capacitor sizing is that it must hold enough energy to support the most energy intensive atomic task of the system, such as operating a sensor, performing some quanta of computation, or transmitting a single radio packet [7, 18, 56, 59]. This represents the smallest amount of energy storage that results in a feasible system with checkpointing or other forward progress mechanism. All batteryless systems operate within the bounds of capacitor storage voltage thresholds. The energy stored within a capacitor is directly related to the voltage across its electrodes.

$$E = \frac{1}{2}C(V_{max}^2 - V_{min}^2) \tag{2.1}$$

The upper hysteresis threshold V_{max}, the point at which the device is charged and turns on, is the voltage at which a capacitor is full, and the lower threshold V_{min}, the point at which the system turns off, is the minimum operating voltage of the components in the system. Smaller capacitors can charge to an upper threshold and turn on faster, but store less energy. Larger capacitors store more energy, but charge slower. Hysteresis management techniques attempt to combine different sized capacitors to optimize the charging time and available energy for specific tasks. Managing hysteresis also allows for more platform generality, as the capacitor storage can be optimized for various operations and workloads. Notable examples of platforms that utilize this technique are the Flicker and Capybara systems [7, 18].

The Flicker platform practices federated energy storage, where each module (sensor, radio, microcontroller) has its own built in capacitor that is tuned to the task it is expected to perform [7, 15]. This has the effect of allowing various components to charge their storage quickly and begin operating before other parts of the system turn on or are needed. This also isolates power failure to independent components. Federated storage provides marginal improvements with regard to component availability and energy harvesting efficiency. A federated approach provides an average increase of 6.7% for microcontroller uptime (from

64.2% to 70.9%) and a 1.5% increase for radio uptime (from 3.8% to 5.24%) over a single duty cycle. Likewise, a federated approach provides at most a 10% improvement in energy harvested over a duty cycle [15]. A federated approach incurs overhead involving additional hardware complexity in the form of voltage monitoring circuitry as well as more complex software to manage and make use of the additional voltage information. The additional hardware increases the size, cost, and energy overhead of the system.

The Capybara platform also manages capacitor hysteresis to increase system responsiveness and reactivity. However instead of using federated capacitor storage, it utilizes a central capacitor bank that can be dynamically resized [18]. This is a more flexible approach than federated energy, where each hardware peripheral is designed with a specific amount of capacitance to support a predetermined task. Capybara can resize its capacitor bank on the fly to match the energy required by an arbitrary task. This results in the lowest possible cold start and capacitor recharge times to support a given operation. Under controlled conditions with consistent harvestable power, Capybara is able to detect 2-4x of events over a system with statically allocated capacitance. However, this improvement only detects 40-70% of possible events, even under optimistic harvesting conditions. To support this dynamic resizing, Capybara introduces an extremely complex and costly power system design that occupies significant board area. The design includes five mixed capacitor and supercapacitor banks, four state-retaining switches for connecting the capacitor banks to the load and harvester, and a custom power distribution circuit to charge the configuration of capacitors.

These techniques are based on the assumption that a full swing capacitor hysteresis with full power downs are necessary and unavoidable. Batteryless systems are often not designed to support a low power sleep state with volatile state retention. Sleeping is a difficult operation for systems with relatively little energy storage, as it is difficult to proactively retain enough energy to enter and sustain a sleeping state when future availability of energy is uncertain. Additionally, energy must be micromanaged on a time scale of milliseconds to seconds. The precision timing of entering sleep state at the correct capacitor voltage level is often not worth the engineering effort to implement when compared to allowing the system power off naturally. Thus, batteryless systems are not often optimized for sleeping. For example, Capybara is by design an inefficient sleeper but has very low leakage when powered off. Its power supply has a significant "power overhead of the power system" that limits the effectiveness of any low power state besides being completely off [18]. The decision to design around constant reboots incurs significant cost in energy and time. Without significant software optimization, the initialization of a processor, radio, and other peripherals is an expensive operation. If capacitor storage is allowed to leak (over the course of hours without any harvestable energy) the voltage may drop to a point that a boost converter must cold start. Cold start is also very expensive, as harvesting ICs and boost converters are optimized for steady state operation and not cold start efficiency. For example, the TI BQ25505 harvesting boost charger is capable of operating at >90% efficiency at voltages higher than 1.5 V [25]. However, its worst case cold start efficiency is 5%, and requires at least 600 mV input and 400 mW instantaneous power to exit cold start. Over the course of operation, a batteryless system will experience many repeated cold starts, resulting in wasted energy and time.

Redundancy for Statistical Availability

Beyond software and hardware techniques for forward progress, timeliness, availability, and reactivity, researchers have also explored approaches that utilize cooperation within a network of batteryless devices. Most notably, the Coalesced Intermittent Sensor (CIS) is an abstraction of a group of intermittent sensors that combines the availability of each sensor to increase the availability of the group as a whole [20]. The CIS approach leverages the theoretical uniform distribution of uptimes of an ensemble of sensors, each of which is intermittently operating, to maximize the probability that at least one sensor is on and able to detect a phenomena at any given time. While this approach may work in controlled environments, it is dependent on several assumptions that are difficult to justify when considering real world conditions.

All nodes within a CIS ensemble are assumed to have a priori knowledge about the size of the deployment and are sufficiently co-located to experience the same harvesting conditions and sense the same phenomena. These assumptions are unrealistic when considering real-world deployment scenarios. Without a mechanism to autonomously and dynamically update network node count, it is onerous for system maintainers to physically reprogram and update each node when a new member is added to the CIS. If all nodes need to be co-located for the CIS algorithm to function, it raises the question as to why this problem could not be solved with a single node with a proportionally larger solar panel and energy storage, instead of the 20-50 nodes it would require to achieve high availability in low light levels.

The evaluation of the CIS greatly overestimates the light levels available to an indoor sensor, and assumes that node on-times are uniformly distributed. In a real deployment, with sensors that are not co-located, light levels can differ by at least 10x, and indoors levels are usually much lower than 300 lux [2]. The evaluation of the CIS sensor shows less than 50% availability when exposed to 300 lux. The only scenario in which the CIS approaches 100% availability is when it is exposed to an LED array providing a constant 1000 lux, an unrealistic environment for many indoor sensors. The CIS approach, like all batteryless approaches, is dependent the existence of any harvestable energy. In most real-world situations, periods of no energy are common, such as nighttime. These periods result in 0% availability, regardless of the amount of node redundancy. For infrequent events with high light levels, the on-times for nodes within the CIS are extended as they are exposed to more energy allowing longer sleep. Nodes that are simultaneously sleeping may then trigger on the same event, essentially synchronizing their on and off times. This results in duplicate event detection and missed subsequent unique events, and temporarily breaks the assumption of uniformly distributed on-times. Beyond the questionable assumptions made by the CIS approach, the event-detection applications that are appropriate for a CIS are limited to short and burst events that can be captured in one intermittent duty cycle.

The Batteryless Fallacy

The aforementioned programming models, runtime systems, debugging tools, hardware platforms, and device redundancy schemes represent over a decade of successive and increasing

complexity and sunk cost in batteryless design. The researchers who have developed these solutions would almost definitely argue that this body of work is justified and necessary for the usability and functionality of batteryless systems, and wireless system design in general. Beyond the ability for a system to safely progress across reboots, the gains in performance provided by various software and hardware techniques are marginal. The value of these improvements is reduced when considering the baseline and maximum theoretical performance of a batteryless system, as well as the hardware and software complexity that is necessary to achieve a feasible system. Due to the nature of batteryless systems, the maximum achievable performance is entirely limited by the instantaneous availability of energy, which is often low and unpredictable in real world applications. Despite the simultaneous high complexity and poor performance of batteryless systems, proponents of the design believe that they are the future of wireless sensor design [17]. This claim ignores the significant downsides to batteryless design, all stemming from an insufficient energy capacity that is unable to support sensor operation beyond a single duty cycle activation, and only when external energy is available. This limitation is **fundamental and inescapable**.

No matter what software framework is used, however the resizable capacitor bank is configured, or how many sensors are deployed within a coalesced sensor, a batteryless system is dead and unresponsive whenever there is insufficient energy: at nighttime, when stationary, without a nearby RFID transmitter, or exposed to an insufficient temperature gradient. To truly address this problem, batteryless systems builders are left waiting for dramatic and transformative technology improvements. Either capacitor energy density must increase by 100-1000x, or system-wide power efficiency must similarly improve to render capacitor-based storage appropriate for many applications [71]. This has resulted in an industry that continues to focus on building battery-based sensors, which are simple to build and operate reliably and predictably. With the efficiency of current processor and radio technology, many applications can persist for a decade or more on just battery power [36–38]. There are few serious attempts by industry to utilize a purely batteryless design for an application, despite more than a decade of batteryless development.

Batteryless Designs in Industry

There are some examples of industry sensor solutions that utilize supercapacitors instead of rechargeable batteries for energy harvesting applications. They deploy supercapacitors that are often large and able to support hours and days of operation, and are often paired with a backup non-rechargeable battery to ensure operation. Pressac sensors utilize EnOcean-based harvesters with a supercapacitor alongside a non-rechargeable battery to ensure a minimum reliable lifetime [72, 73]. The Davis weather and air quality stations utilize a supercapacitor to harvest daylight and include a non-rechargeable backup battery [45].

There are fewer examples of purely batteryless products and solutions from industry. Many new companies are attempting to build batteryless products, including Williot, Perpetua, and Everactive, among others [10, 48, 74]. The Williot Pixel asset tracking tag was originally meant to be an RFID-like batteryless peel and stick tag. These batteryless tags require a nearby

Application	Lifetime	Workload Power	Harvesting Power	Size/Weight	Reliable
GDI [8]	9 months	300 µW	—	25 cm^3	yes
ZebraNet [9]	1 year	400 mW	300–1000 mW	1.15 kg	yes
Williot [10]	∞ [a]	— [c]	10–100 µW	0.25 cm^3	no
Williot Battery [10]	4 years	∼ 20 nW [d]	—	2.8 cm^3	yes
CapBand [58]	∞ [a]	260 µW	∼ 100–1000 µW	∼ 20 cm^3	no
FaceBit [71]	11 days	286 µW	70–700 µW	10 cm^3	yes
Monjolo [55]	∞ [a]	1–27 mW	1–70 mW	∼ 80 cm^3	no
Kingdom [11]	6 months	∼ 200 µW [d]	—	— [e]	yes
Empire [11]	∞ [a]	— [c]	<50 µW	— [e]	no
Cathodic [12]	3–10 years	— [c]	480 µW	— [e]	yes [f]
M^3 [75]	— [b]	40 nW	20 nW–10 µW	1 mm^3	yes

[a] Indefinite capacitor cycle lifetime. [b] Lifetime limited by cycle lifetime of battery.
[c] Average workload power assumed equal to harvesting. [d] Average workload power estimated.
[e] Size not mentioned or prioritized. [f] Reliable because harvesting income is reliable.

Table 2.1: An assortment of seminal and recent wireless sensor applications and deployments. This selection is by no means complete, but provides a reasonable selection of sensors with varying power supply architectures and intended sensing applications. Wherever possible, figures are taken or calculated directly from the cited work. Some figures are not mentioned and are estimated instead.

reader for power and to communicate. For many customer applications, having sufficient reader infrastructure is likely cost-prohibitive. To address this, Williot has now introduced a battery-powered tag that performs active BLE beaconing and is compatible with any BLE device. Perpetua and Everactive are developing batteryless sensors for industrial monitoring, taking advantage of sensing modalities that are simultaneously sources of harvestable energy, like thermal gradients or vibration. These applications represent a unique opportunity for a batteryless design, as the sources of energy are reliable and correspond to valuable sensed data, resulting in a batteryless system with high availability.

There is a valid argument that batteryless design is suitable for various niche applications, like those within environments with reliable sources of harvestable energy. However, not all applications are feasible with batteryless sensor designs, limiting industry adoption. In Section 2.3, we explore applications in the context of their requirements and how different energy provisioning techniques address them.

2.3 Wireless Sensor Applications and Their Requirements

Real applications have requirements for availability, latency, and reliability. For many, these requirements may be at odds with what can currently be provided by batteryless systems,

and in some cases energy harvesting methods as a whole. There is a clear niche for batteryless systems for applications that can tolerate relaxed requirements, or are designed for reliable energy income, however it is disingenuous to argue that batteryless systems represent the future for all sensing applications [17]. The next section explores several example applications from the literature and industry, with qualitative and high level quantitative feasibility analysis. These example applications and their requirements are summarized in Table 2.1.

Application Requirements

While the popularity and excitement around energy harvesting and batteryless sensing has led to progressively better batteryless systems, the focus on batteryless design has often ignored the requirements of real world applications, and the limitations of the batteryless archetype. Constraints like sensor size, lifetime, the availability of harvestable energy, and the reliability and consistency of an application all contribute to the end design of a wireless sensor, and many applications will have requirements that make a batteryless design an unsuitable choice for many real applications requiring some level of quality of service. Instead of starting with the assumption that a design must be batteryless, a design should start from application requirements. Given an application, a set of questions must be answered to define requirements and design parameters. An appropriate set of questions to help define the power supply requirements may resemble the following:

1. **What is the application lifetime requirement?**

   This question primarily defines the required lifetime of the wireless sensors driving the application. Secondarily, it also defines the acceptability of occasional power supply maintenance.

2. **What level of availability is required by the application?**

   In this question, availability refers to the general ability of the sensor to remain online and maintain sensing at the rate required by an application. In this sense, availability is also analogous to reliability. Availability can be defined depending on the application reporting frequency and distribution. For applications that require measurements on a uniform sensing period, availability is defined by the percentage of intervals that the sensor was able to measure, or how often it was able to measure on schedule. For applications that are event driven, availability can be defined by the percentage of events captured by the sensor compared to the baseline distribution of events.

3. **What are the average sensor power requirements to drive the application?**

   The power requirements of a sensor generally define application limitations for the size of the sensor, the feasibility of various energy harvesting methods, as well as the maximum lifetime on battery power. For opportunistic batteryless sensors, average power has a different meaning, as such sensors use any energy that is instantaneously available, and the average power they utilize varies with supply. There is still a limitation, as there is

a minimum amount of power and energy required for a batteryless system to turn on and successfully activate.

4. **What are the limits of area, volume, weight, and cost for each sensor?**

   The limits of size, weight, and cost of a sensor directly affects the amount of energy storage that can be allocated, or the available surface area for an energy harvester. Many applications have various limitations on the size of individual sensors as to not be burdensome, difficult to deploy, or unaesthetic.

In the following sections, several example applications from prior work are presented and their requirements placed within the context of the above questions. While this is by no means an exhaustive list, the examples discussed form a subset that provides broad coverage of applications that can and cannot be supported by a batteryless design. Some of these applications were developed before the introduction of batteryless sensors, however they are analyzed as if the technology and ideas existed to enable a batteryless option. In Chapter 3, these questions form the basis of formalizing a set of high-level application constraint equations to determine power supply design. These constraints are utilized to develop new heuristics for sensor power supply design.

Monitoring and Tracking

Classic wireless sensor applications generally monitor a phenomena, or track the location of objects and people. These applications require regular measurement and reporting to provide clear and complete data of the phenomena being measured or the location and path of the object, person, or animal being tracked.

Habitat Monitoring

Perhaps the most notable seminal work in wireless sensor networks research is the deployment of 32 wireless sensors for habitat monitoring on Great Duck Island [8]. These sensors were built and deployed to measure the environmental conditions inside and around seabird nesting burrows, as well as their occupancy during breeding season. This subsection attempts to explore this application's requirements and identify the suitability of different power supply architectures to meet these requirements.

This habitat monitoring application has a necessary lifetime of 9 months, a time period that encapsulates the average 7 month sea bird breeding period. The goal of the study is to measure burrow conditions and usage patterns consistently over a 24–72 hour cycle. The sensor must be able to measure occupancy consistently every 5–10 minutes overnight and in the early morning, when it is most likely that birds are leaving and entering the burrow. It is acceptable to measure general environmental conditions within and outside the burrows every hour. Thus, the in-burrow sensor must be able to sense occupancy and environmental conditions at regular intervals, without interruption.

To perform this task, based on technology available when this work was published, the burrow sensors required on average ∼300 µW. This power requirement is low enough to be powered by various energy harvesting methods, however there are not many options for harvesting within a sea bird burrow. Solar and thermal differential harvesting are not ideal. The inclusion of these harvesting methods would potentially make the burrows unappealing to sea birds. For the same reason, the mote must also occupy a small volume. The mote is based on the Mica mote which occupies about 25 cm^3 [76]. Energy harvesting would also substantially increase the deployment effort and cost. External to the burrows, the gateway infrastructure for the deployment has plentiful solar energy to harvest from.

The sensor designers chose to utilize non-rechargeable batteries as a source of energy, and tailored the size of the batteries as well as the sensor workload in order to support a 9 month lifetime. Based on the nature of this application, the power supply design would be unlikely to change, even given modern technology. The application requires reliable and periodic measurement, and is unable to utilize harvestable energy for in-burrow motes.

Wildlife Tracking

Zebranet also represents a seminal work in wireless sensor design, involving the development and deployment of 30 tracking collars on zebras under study [9]. Researchers were primarily interested in tracking the location and migration of zebras in an environment with no cellular or other broadcast communications. The collars include a GPS system, a short and long-range wireless transceiver, and a processor.

The wildlife biologists working with the system designers were interested in capturing GPS position every 3 minutes for at least 1 year of operation with no human maintenance or intervention. The biologists understandably want to limit the amount of tranquilizing and re-collaring of the zebras as much as possible. Like the habitat monitoring application, Zebranet also requires continuous and reliable operation to collect periodic GPS updates. Due to the lack of options for long range networking, the ZebraNet collars utilize peer-to-peer networking and mobile base stations for eventual data collection from the monitored zebras.

Conservatively, ZebraNet collars require an average of 405 mW to perform GPS sampling and occasional peer-to-peer and base station networking. ZebraNet collars are relatively high power, even compared to other contemporary sensors like the Great Duck Island burrow sensors. GPS and long range radio technology had high power requirements at the time the collars were designed. The authors note that to power a collar for just five days would require a 13.5 A h non-rechargeable battery, weighing 4 pounds. Not only would a non-rechargeable battery be unable to provide sufficient energy to achieve the application lifetime goals, it would also be too large and heavy to put on a zebra's neck. Instead, the collar designers opted to use solar energy harvesting with a lithium polymer battery, resulting in a significantly smaller collar, weighing only 2.54 lbs total. This allows the collar to harvest the copious amounts of sunlight during the day, and power the zebra localization workload all day and night. With a batteryless design, the collars would be unable to function overnight or in the early morning, limiting zebra localization and behavior data to daytime. Given modern technology,

the collar *could* utilize thermal gradient harvesting in addition to solar harvesting. This would allow harvesting overnight, as the gradient between zebra skin and ambient nighttime temperatures would be likely be sufficient to generate energy. However, the relative power provided via thermal harvesting is several orders of magnitude less than harvesting direct sunlight, and much less than is required by the workload. Even with additional harvesting sources, a batteryless collar would have to necessarily degrade the consistency and frequency of GPS sampling and communication with peers or the base station.

Asset Tracking

In many industries, it is important but incredibly difficult to monitor and track various products and commodities through a supply chain. Asset tracking sensors have the potential to provide finer granularity tracking and monitoring, which provide clearer estimates for production and delivery, as well help to increase efficiency by better tailoring supply to demand. The lifetime of an asset consists of the amount of time it takes an item to traverse a supply chain end-to-end, from production to end retailer and user, potentially spending an unknown amount of time in a warehouse or on retail shelves. Tracking sensors must be small, cheap, and easily disposable or recyclable. Their lifetime must resemble the maximum shelf lifetime of any given product, which can vary widely.

Generally, these sensors are simple, and only need to report their existence occasionally at each step of production or delivery, or respond to queries when inventory is taken. This means their power requirements are modest, as they do not need to consistently and periodically transmit. Asset tracking is a uniquely situated application that can benefit from a batteryless design. The infrequency of location reporting coupled with the requirement for cheap and easily disposable tags results in a design space that favors a batteryless approach. As long as each step of a supply chain has infrastructure to wirelessly power and communicate with tags, batteryless asset tags can provide accurate and reliable tracking updates.

Several recent companies are developing cheap batteryless asset tags, including Williot and Jeeva Wireless [10, 77]. The Williot battery-free IoT Pixel is a peel and stick tag with an ARM Cortex M0+ and harvesting and transmission antenna. The battery-free Pixel relies on RFID-like transmitter infrastructure to supply wireless power, and tags report their existence via active BLE (Bluetooth Low Energy) broadcasts which can be received by any consumer handheld device. For applications where RFID infrastructure is inadequate, Williot also offers a battery-assisted tag with a battery lifetime of four years. Jeeva Wireless is also developing low power RF-harvesting solutions for asset tracking. However, instead of BLE transmissions, they rely on bidirectional backscatter communication with a wireless router [77].

Asset tags must be physically small and lightweight. Arguably, the most convenient form factor resembles that of a traditional "dumb" tracking label: a peel and stick tracker with an integrated harvester and radio. This form factor does not allow much space for a traditional non-rechargeable battery, but can supply enough surface area to support an energy harvesting RF antenna, an active antenna, and a limited amount of capacitor-based energy storage. At most, the minimum atomic operation of an asset tracker must be sending sending a single

transmission indicating the tag's existence. This amount of energy is easily supplied by a bank of capacitors [2, 55, 63]. Beyond size, these asset trackers must also be inexpensive, and the cost difference between a non-rechargeable battery and a bank of ceramic or tantalum capacitors is substantial at scale. The Williot battery-free Pixel occupies a sticker form factor of 2.8x4.4 cm and 0.2 mm thick and is projected to cost between \$0.10–0.70 USD. The battery-assisted Pixel is a larger and thicker sticker that is 3.6x6.0 cm and 1.3 mm thick and is projected to cost \$1–2.1 USD. Williot does not publish the power or energy requirements for their tags. We can still estimate the average power requirements for the battery-assisted Pixel given the volume and lifetime of the battery-assisted Pixel and the volumetric energy density of lithium batteries. The Williot battery-assisted tag occupies 2.8 cm^3, has a purported lifetime of four years, and the energy density of a lithium primary is 0.8 Wh/cm^3 [78]. This suggests that the battery-assisted Pixel only requires 17.8 nW on average. This figure is very low for an actively broadcasting system, and it would not be surprising that a four year lifetime is an overestimate for the battery-assisted Pixel. We would expect an average power requirement multiple orders of magnitude more for active BLE beaconing. For example, a modern low power SoC like the nRF52 series requires 15 µW to beacon every four seconds [79]. A tag like those that Jeeva is developing utilizes backscatter and will require very little power to transmit data, usually much less than a µW. This average power is well below the capabilities for directed RF harvesting from a device similar to an RFID reader, which can typically supply hundreds of µW at a minimum at reasonable ranges [54].

In environments with plentiful and reliable harvestable energy, like in the case of asset trackers within a factory or warehouse with widely deployed RFID readers, the batteryless design point allows for simple, small, and inexpensive wireless devices that require essentially no maintenance. However, when a batteryless asset tracker is out of range due to theft, accidental loss, or other reasons, it is impossible to locate them as they have no energy or way to communicate. The value of batteryless trackers is strictly limited to specific locations and areas where the tracker is expected to be. It is likely that such events are rare, and for most asset tracking use cases with plentiful infrastructure, a batteryless tracker is an ideal solution. In situations where infrastructure is insufficient to power batteryless tags, non-rechargeable batteries may be necessary. Williot backtracked from originally offering only battery-free tags to developing and offering battery-assisted Pixels to address this reality.

Wearables

Many human-centric applications are only possible when a sensor is co-located on the human body. Often, wearables are the best and sometimes only way to accurately measure or interact with humans. These wearable sensors are uniquely situated when considering their power supply design. For a wearable, a human is naturally always present and able to provide light maintenance or a recharge. This allows wearables to have shorter one-charge lifetimes, use smaller batteries, and achieve smaller form factors. Wearable maintenance is less costly when humans are always present and are already conditioned to maintain their devices. Humans have generally accepted the cost of nightly smartphone recharging, making daily or weekly

wearable charging also acceptable. However, not all populations are as diligent or willing to utilize wearable sensors if they have to repeatedly remove and charge them. Particularly, elderly populations have usability issues with discomfort and difficulties with taking wearables off to charge them and putting them back on [80]. For these populations, a wearable that does not need to be charged would be beneficial. A batteryless wearable would not have a charging constraint, however the expectation of constant and immediate wearable human interactivity demands a baseline of availability that may not be feasible with a batteryless design. It is hard to imagine a user being willing to adapt their behavior to accommodate a sporadically operating device. This section examines several recent batteryless wearable research applications in the context of the aforementioned trade-offs.

Gesture Recognition

CapBand is a batteryless gesture recognition wearable that measures capacitive changes in skin deformations to identify hand gestures [58]. While the main contribution of this work is the sensing and processing of wrist capacitance measurements, the authors also claim that the low power nature of the measurement technique makes this appropriate for a batteryless design. Capband utilizes a small solar panel for power, and a supercapacitor for energy storage. Gesture recognition is an application that requires high availability to accurately capture hand gestures. CapBand's capacitive measurement needs to run continuously to capture enough data to predict gestures, on the order of 20 measurements per second for its neural network inference. At this rate, CapBand requires an average 260 µW. This does not include the energy required to transmit the measurements. The measurement power alone allows operation only outdoors during the daytime, and indoors with very bright lighting. User frustration is likely to arise in insufficient lighting.

Wearables must optimize for size and weight. While CapBand fits within a wrist form factor, the sensing and energy management circuitry is quite large compared to a normal wristwatch or smartwatch. There is some room for miniaturization by using different IC packages and combining modules into one PCB. However, the majority of volume and surface area of CapBand is actually devoted to the solar panel and supercapacitor. A small rechargeable 40 mA h could power the aforementioned 20 Hz measurement for 27 days, and would take less than half of the volume of CapBand's designed energy harvesting frontend.

CapBand's supercapacitor is sized to support an amount of sensing and data transmission while also minimizing the amount of time it takes to charge to a usable voltage. It is not designed to support a single atomic operation. CapBand is not designed to power off frequently. Instead, it sleeps when it does not have sufficient power. A more energy dense storage, as well as one that has a more stable voltage, like a rechargeable battery, would also allow for continuous operation. A battery would also allow CapBand to operate overnight and in low light environments assuming it is charged regularly. The authors do not give consideration to a battery-based design, beyond claiming one would need frequent recharging. We believe the decision to use a batteryless design was informed more by the popularity and excitement for battery-free technology, rather than a careful design exploration that considered all options.

Smart Face Mask

Beyond traditional wrist-worn wearables, researchers are also developing ways to measure the efficacy of face masks and the health of the user. The FaceBit platform is a smart mask platform that is affixed to normal masks and measures non-contact heart rate, respiratory rate, and mask fit quality and wear time [71]. FaceBit features a hybrid batteryless energy harvesting design with a backup primary battery. This hybrid design is intended to extend the lifetime beyond that of a primary-only design. The current FaceBit design cannot operate in a fully batteryless fashion, as the reporting frequency and computational demands of the application requires more power and availability than possible with energy harvesting alone.

The FaceBit prototype can persist for 11 days before exhausting its small primary battery. Its workload consists of periodically sampling a barometer sensor once a second to detect if it is being worn. If it is being worn, it measures heart rate and respiratory rate every minute, and reports results every 2 minutes. This workload requires high availability and consistent periodicity, which is difficult to achieve with a batteryless design. The FaceBit workload requires an average power of 286 µW. While the current version relies mostly on a primary battery, the authors have explored the use of energy harvesting methods to power FaceBit. The authors explore multiple energy harvesting methods, including a photovoltaic cell, a thermoelectric generator (TEG), as well as a triboelectric nanogenerator (TENG). The photovoltaic cell must be placed external to the mask, with power leads running into the mask, potentially degrading the mask fit. The authors also test the use of a TEG placed against the inside mask surface, with the goal of harvesting from the temperature gradient from inside the mask to the surface of the mask. The TENG converts the small kinetic forces caused by breathing into electricity, and is placed within the mask.

Among the energy harvesting methods tested, only the photovoltaic in outdoor light can provide sufficient power to power the device. In worst-case indoor conditions, where dim artificial lighting only provides on the order of 10 µW/cm^2 [2] the selected photovoltaic panel for FaceBit can only provide 71 µW, assuming a 17% efficiency. The authors also explore the use of TEG and TENG generation. Thermoelectric harvesting provides a temporary burst of energy on the order of 50 µW, but quickly dissipates as the inside of the mask and the inside mask surface reach a temperature equilibrium. The TENG can provide consistent power if the mask is worn, but generally less than 2 µW when the mask wearer is talking. No power data is provided for when the wearer is just breathing. The authors believe that breath-based harvesting (Either TEG or TENG) are the most promising avenues for future development, despite their low power output. They also claim that a 10x improvement in power efficiency of the SoC (including the processor and radio) would enable a batteryless version of the FaceBit. This is not possible given the capability of the TEG and TENG methods explored. It would require closer to a 100x system-wide power improvement to enable a continuously operating batteryless FaceBit.

FaceBit must be sufficiently small and unobtrusive as to fit in an N95 or surgical mask without bothering the wearer. It must also be lightweight enough to be held securely with a magnet affixed to the outside of the mask. The choice of non-rechargeable and rechargeable

energy storage, as well as the harvesting method is the largest driver for the size, weight, and cost of a device like FaceBit. The choice to use any energy harvesting, including the tested solar panel, TEG, or TENG requires significant surface area and added cost. The use of tantalum capacitors as energy storage is also one of the least energy dense options.

The authors do not give a rationale for the decision to utilize capacitors as energy harvesting storage. While each capacitor is individually small, FaceBit requires multiple capacitors to build up enough energy storage to successfully complete its most energy intensive operations. Since the capacitor bank has very limited energy storage, the utility of having energy harvesting at all is severely limited. FaceBit can only capture enough energy to perform a few operations before depleting its storage. Even in cases of sufficient harvesting, FaceBit can only utilize a small percentage of any available energy because it fills its buffer quickly and any remaining harvestable energy is wasted. The authors do not examine how their energy storage capacity and chosen energy harvesting methods affect energy capture and the lifetime of FaceBit's primary cell.

The three capacitors that comprise FaceBit's energy storage occupy nearly 1.5x the volume of the primary battery itself, at significantly lower energy density. An alternate design could more than double its lifetime by replacing energy harvesting and the capacitor bank with additional battery storage. Alternatively, the primary battery and energy harvesting components could be replaced by a single rechargeable battery and a charging connection. A rechargeable battery is less energy dense than a primary cell and be unable to offer a similar lifetime. However, a rechargeable would reduce battery waste. Both options would result in a smaller, less cumbersome design. The removal of energy harvesting and related circuitry would also reduce the cost of the unit.

Intra-Body Power Transfer

SkinnyPower is a technology that enables wireless power transfer through human skin to enable more reliable batteryless wearable sensors [59]. Small batteryless sensors can be placed on body parts that are unable to support batteries either due to their volume or weight. SkinnyPower targets sensors for fingers, ears, or mouths, and provides power via battery-powered transmitters placed elsewhere. SkinnyPower represents an unique approach to batteryless systems by exploiting their potential for miniaturization, as well as addressing their weaknesses in availability by providing a reliable source of untethered power.

The SkinnyPower prototype consists of a wrist-worn battery-powered power transmitter and a finger-worn batteryless sensor that operates on the transmitted power. The authors target a smartwatch-sized form factor for the power transmitter while the finger-worn sensor requires a surface area of 2.5x2.5 cm^2, about the size of a smart watch underside. The wrist-worn transmitter powers the batteryless sensor by transmitting an alternating current through the skin of the user's hand. With this prototype, SkinnyPower is able to provide average power on the order of 1 mW at an efficiency of 14.5%. This is enough power to operate a batteryless sensor system continuously. Assuming a 100 mAh rechargeable lithium polymer battery powering the transmitter, a SkinnyPower system could provide 1 mW to a

batteryless finger sensor for 54 hours. Given that SkinnyPower can provide continuous power, it is possible to operate sensors like an accelerometer continuously with high availability. The authors evaluate the system with an accelerometer and compare its measured accelerometer with that of a motion capture system. They confirm that the system is able to continuously and accurately report acceleration measurements at 1 Hz.

Through the use of intra-body power transfer, the SkinnyPower approach is able to separate a sensing system's power supply from its sensing system. This allows small and lightweight batteryless sensors to be placed on body locations where it would be burdensome to place batteries, while at the same time providing the high sensing availability of a battery-powered system.

Amalgamated Harvesting and Sensing

There is an application niche where sensed phenomena is correlated with a method of harvesting. As mentioned previously in Section 2.2, there are one-shot batteryless sensors that do not directly sense any phenomena, and instead simply transmit packets when their capacitor storage is charged and full. The rate of packet transmission by these simple sensors can be used to directly estimate the rate of harvesting, and the intensity of a phenomena. There are other applications that take this primitive one step further, and integrate active sensing for more direct and accurate measurements. Applications that harvest the same phenomena that they sense are well situated for batteryless designs. The inherent unreliability of a batteryless supply is not as impactful, as whenever there is energy to be harvested corresponds to when it is important to be sensing. While energy harvesting methods are limited, many valuable applications are still feasible.

Power Metering

Usually, the proximity to existing powered infrastructure presents relatively easy solutions for powering power meters. For some use cases, this proximity is not enough to justify the cost of deployment. It is generally burdensome and costly to instrument a building's circuit panel, as traditional submetering equipment must directly connect to a circuit, which requires an electrician install. Many plug-load power meters are bulky and often block nearby outlets, and they require substantial standby power to operate. The Monjolo power meter [55] instead utilizes a current transformer wrapped around one path of an AC power line to harvest energy. The Monjolo operates on a one-shot batteryless design principle where the quantity and rate of packets can be used to estimate the current through an instrumented AC line.

The availability of a Monjolo power meter is tied to the availability of the instrumented load. Unless a load current is constant, a Monjolo sensor is unable to operate continuously, and when no load current is flowing, it will remain off and not send packets. When a Monjolo is offline for an extended period, it can be assumed that the measured load is also offline and drawing no power. The zero-load assumption breaks down when considering real-world situations where operation and packet reception are not reliable. If a Monjolo breaks, or its

packets are not reliably received for an extended period, its failure is impossible to distinguish from the sensor measuring no load. Additionally, any missed packets has an impact on the reconstructed power rate estimation resulting in underestimation of power. To account for this, Monjolo employs a monotonically increasing counter sent with each packet such that high level energy aggregation remains accurate. However, this technique does not correct the loss in temporal power resolution from lost packets.

Because the Monjolo sensor operation is directly tied to its energy income, its power requirements vary depending on the measured load. At a minimum, it requires a load of at least 17 W to harvest enough power to turn on and send a packet. This corresponds to a coil output power of less than 1 mW. At some point the induced power from the magnetically coupled load is enough to power the Monjolo indefinitely. This point represents a maximum possible measurement, as the rate of packets is limited by how fast the sensor's software and hardware can queue and send packets when running continuously. For the Monjolo, this maximum corresponds to a measured load of 480 W, and a packet frequency of 8.5 Hz.

Regarding size, the Monjolo platform must fit easily within a circuit breaker panel. Most of the volume of the Monjolo is dedicated to its coil harvester. The 500 µF storage capacitance and circuitry is small and amortized in comparison. This capacitance is tuned to support a single Monjolo activation, allowing the sensor to wake up and transmit a packet. However, with consistently powered or high current loads, the Monjolo harvester can provide significant energy, and a larger or more dense energy storage option could allow for more captured energy. This would potentially support continuous operation and more advanced active sensing to directly measure current and voltage of the load without the need for power estimation from packet rate, and the limitations of upper and lower bounds of measurement.

In fact, the Monjolo harvesting method and split-coil form factor has directly influenced further research and industry solutions. The Triumvi system is a direct improvement over the Monjolo design, directly measuring voltage through capacitive coupling, and the current transformer to harvest energy and measure current [4]. Like Monjolo, Triumvi is designed with a batteryless power supply, but supports an optional rechargeable battery for reliable, short term instrumentation. Triumvi also supports charge sharing between devices, allowing a sensor harvesting from and measuring a high load to share its harvesting energy with a sensor measuring a smaller load. This increases the overall reliability and uptime of the entire deployment. Vizi Metering is commercializing the Monjolo and Triumvi concepts with some improvements [81]. Instead of utilizing the one-shot activation and packet rate counting method of the Monjolo, the Vizi meter directly measures and reports current and voltage like Triumvi. This direct measurement is much higher accuracy than estimation through packet counting, which has no phase information or voltage measurement. Unlike both Monjolo and Triumvi, the Vizi meter is not batteryless. It uses a small rechargeable battery as its energy storage. The inclusion of a rechargeable battery allows the Vizi meter to run continuously. Continuous operation allows the meter to utilize reliable transport mechanisms to ensure all measurements are received successfully. This preserves measurement temporal resolution, which was potentially lost when Monjolo and Triumvi packets were not received. Like Triumvi, the Vizi meter also supports daisy-chained charge sharing. This daisy-chained connection

also allows communication between sensors, primarily to synchronize measurements. This is especially useful for measuring three-phase or split-phase circuits. Even though the Monjolo principle of operation is intuitive and a seemingly appropriate fit for a batteryless design, the limitations of measurement accuracy, resolution, range, and reliability are not appropriate for an industrial sensor. It was necessary for Vizi to abandon the batteryless aspect of Monjolo and Triumvi to achieve their availability, accuracy, and resolution goals.

Archaeological Site Monitoring

There are actually few real-world deployments of long-term batteryless sensors. One of these is a 3.5 year deployment in the Mithræum of the Circus Maximus to measure ambient temperature, humidity, and structural vibrations [11]. Having data regarding the environmental conditions of the site helps restorers identify the necessary techniques for preservation. The Mithræum deployment consists of three different generations of devices, starting with battery-powered sensors and evolving into a deployment of batteryless sensors. This work is an interesting and unique analysis of the trade-offs of battery-based and batteryless design.

The Mithræum is completely underground and only accessible through spiral staircases, and entry is strictly regulated to protect the site. Specifically, the installation of vibration sensors requires difficult climbing on potentially fragile vaults, putting not only the site in danger of damage, but also the safety of people installing sensors. This makes periodic sensor network maintenance extremely costly. The original lifetime goal of the battery-based sensor deployment was six months, however continuous sensing beyond that time was desired. The battery-based design was built with off the shelf processors, radios, and sensor modules. Through the course of the 3.5 year deployment, the battery deployment experiences frequent outages, an average of 9 a year. Due to these outages, the authors turn to batteryless designs for their environment and vibration sensors, producing custom application-specific hardware and software designs.

The authors do not explicitly mention many quantitative details of their designs including power and energy requirements for different operations, making it difficult to gleam the average application power requirements. However, the off the shelf components used for the first deployment, along with the described workload provide some insight, enough to make an educated guess. The first deployment utilized a Libelium Waspmote and an XBee 868LP sub-GHz radio [82, 83]. Combined, these two components require a sleep current of 34.7 µA at 3.3 V. The Waspmote is powered by an ATmega1281 and has an active current of 17 mA and frequency of 14.7456 MHz. The Xbee sub- GHz radio requires 45 mA when active and can transmit at 80 kb/s. The temperature and humidity sensor only requires 600 µA when measuring, while the accelerometer requires 18 mA when active. The temperature sensor workload consists of measuring temperature and humidity every 20 minutes, and transmitting an average and standard deviation every hour. Similarly, the vibration sensors capture a one minute burst of accelerometer data every 30 minutes. Every hour the acceleration data is summarized via an FFT, the fundamental frequency and spectral density is computed, and the results are compressed and sent. With this information, we can estimate an average

power draw for each system by making a few assumptions. For simplicity, we assume every packet sent is an arbitrary 200 bytes. We also conservatively assume that it requires 100 ms to measure the temperature and humidity sensor and one second to summarize the collected data for the hour. For the vibration sensor, we assume it requires five seconds of computational time every hour to calculate an FFT and other statistics.

In an hour, the temperature and humidity sensor will spend 300 ms measuring its sensor, one second calculating summary statistics, and the amount of time it requires the radio to transmit a 200 byte packet over the 80 kb/s link. The rest of the time is spent sleeping. This corresponds to an average of 191 µW to power the temperature sensor workload. Likewise, every hour the vibration sensor will spend two seconds measuring its sensor, five seconds computing its statistics, and the same transmit time as the other sensor. This corresponds to an average of 244 µW to power the vibration sensor workload.

The battery-based sensor used in the first "Kingdom" deployment utilizes six C cell alkaline batteries in a 2x3 configuration, two in series and three in parallel. This 2x3 configuration of 6 C-type batteries provides a maximum of 16 Ah at a nominal voltage of 4.2V, or 67.2 Wh. These batteries should provide more than sufficient energy storage for both types of sensor to last six months. Despite this, the amount and frequency of failures experienced by the real deployment points to multiple issues that were not predicted at design time. The authors lay most of the blame of sensor failures on batteries. They claim that the high humidity in the underground chambers of the Mithræum caused the frequent battery failures, which occasionally included visible degradation in the form of leaking battery acid. The authors claim that different primary battery types, such as lithium, also exhibit issues. To solve this problem, the authors elect to change their sensor design to be batteryless. The authors do not consider improving their battery-based design with any efforts to mitigate humidity with waterproof enclosures, humidity absorbers, or conformal coatings and epoxy. The authors were primarily interested in developing a deployment of batteryless sensors, and were motivated to do so when they encountered initial difficulties with batteries in humid environments.

The batteryless temperature and humidity sensors harvest energy via a TEG connected to a sublayer of Roman concrete. Likewise, the batteryless vibration sensor utilizes piezoelectric harvesting connected to the columns of the Mithræum. The first batteryless attempt uses the same COTS components as the battery-based sensor but connected to these harvesters. The COTS components were not designed to operate in a batteryless context and performed poorly. The vibration sensor was rarely able to wake up and was unable to capture useful and relevant data. The second batteryless attempt was to build a custom application-specific hardware design with more efficient components. The deployment of these new sensors was named "Empire". For the vibration sensor, they also designed a vibration-triggered activation. Rather than sense vibration occasionally and potentially miss important vibration events, the redesigned sensors began sensing when a sufficient vibration occurs and triggers a wake up. This aided in coupling the sensor's sensing modality with its harvesting, resulting in capturing more relevant data. While vibration sensing was designed to capture rare events, which is more amenable to the batteryless design, temperature sensing was not as successful, as it requires regular measurements throughout the day. The Empire temperature sensors

were only able to capture 22% of the quantity of data provided by the battery-based Kingdom sensors. The authors do not comment on the temporal distribution of the collected data, but we assume it was likely clustered around certain times of day when the temperature differential between the inside of the Mithræum and the outside was sufficient to generate power. The authors note that over two years, the batteryless Empire deployment has required no maintenance beyond a cellular failure unrelated to the sensors themselves.

This work represents an interesting design exploration and highlights a few major benefits of batteryless systems. Most notably, capacitor-based power supplies are more durable in the face of high humidity. Also, when harvesting and sensing are properly coupled, as was the case with the Empire vibration sensor, a sample of interesting vibration events can be captured. However, this work also exemplifies their downsides. Compared to the battery deployment, which was built with COTS, the batteryless deployment required a custom, more complex, and application-specific batteryless hardware design. This custom design required significantly more design time and effort. While some relevant vibration events are able to be captured by the batteryless design, there still exists the probability of not capturing events if they are are close in time or the sensor has not been able to instantaneously capture enough energy from a vibration to sense and report it. The batteryless deployment was also unable to capture the same amount of temperature and humidity data as the battery deployment, and did so with much less consistency. Batteryless designs are poorly suited for applications that require that high data fidelity or that data be captured at regular intervals.

Cathodic Protection Monitoring

Beyond the coupling of sensing to harvesting, sometimes harvestable power is reliable. Reliable power can be produced from a consistent temperature differential between a hot steam pipe and the ambient air, microbial activity within soil or wastewater fuel cells, or the voltage potential produced from a sacrificial anode in a home water heater, a boat, or a bridge. In these situations, reliable harvestable energy results in a reliable system, even when batteryless. This allows a batteryless system to maintain a regular and periodic sensing schedule. These applications are clever and convenient for batteryless systems, as a reliable power source allows them to avoid the uncertainty of an unreliable income.

Researchers have built proof-of-concept sensors that monitor and harvest from the corrosion of cathodic protection systems [12]. A corrosion protection system is sometimes employed on critical infrastructure, especially in aqueous environments, to prevent corrosion. This protection usually consists of bars of a sacrificial metal, usually of elements that are more easily corroded, like aluminum or magnesium. The critical metal structure is protected from corrosion, as the sacrificial metal corrodes first. It is important to monitor the state of a cathodic protection system, because as soon as it is depleted, the metal structure that was being protected will begin to corrode.

The chemical process of corrosion is the same process as a galvanic battery. A cathodic protection system is not a *good* battery. It provides a very low peak power, and has very high equivalent series resistance. However, a cathodic protection system can be used to

slowly charge an actual battery, a bank of capacitors, or a supercapacitor. This stored energy can then be used to monitor the health of the protection system, powering periodic voltage measurements between the sacrificial anode and the protected cathode, and reporting these readings via radio transmissions. As long as a sacrificial cathode exists and is performing its function, a sensor will be able to scavenge energy from it. Such a system must persist for the life of the cathode or infrastructure, which in the case of a water heater can be between 1-3 years for the cathode, and around 10 years for the heater itself [12]. While this lifetime could be supported by a non-rechargeable battery, it is both more elegant and less wasteful to employ an energy harvesting solution. Since the degradation of a sacrificial anode is slow, over the course of multiple years in the case of a water heater, such a sensor does not need to report frequently. A measurement once a day is sufficient to track the health of the sacrificial anode and predict its failure.

The authors propose a fully energy harvesting solution, meaning that the average power of their application must be energy neutral, and fully dependent on income from the weak corrosive battery. Via multiple experiments with a small corrosion cell consisting of a steel bucket and a magnesium rod, the researchers have discovered a maximum power point at around 300 mV and 1.6 mA for a peak power of 480 µW. Even from a small example corrosion cell based on a water heater, this amount of average power is more than enough to power a low duty-cycle sensor system. In many cases, the infrastructure being monitored will be much larger, with more surface area resulting in more harvestable power. The size of the designed system is not of particular concern, as it connects to potentially huge infrastructure. Any sensor built will be easily dwarfed by a water heater or a bridge.

Since the energy income from a cathodic protection system is stable and reliable up until the point where it fails, a batteryless design is an attractive choice. The constant trickle of income power allows a system to preserve volatile state over long sleep periods without the need for intermittent techniques to ensure forward progress. The left over energy is banked and eventually used to monitor and report the health of the sacrificial anode. The authors decide to use a parallel array of three 1 F supercapacitors to store harvested energy. This bank of supercapacitors is sized such that it provides enough energy storage for the system to run continuously and measure the voltage between the anode and cathode every 4 hours on average. After measuring, the system transmits the results via a LoRa radio, which has relatively high power and energy costs. The authors avoid tantalum and ceramic capacitors used in past designs, as their energy density is insufficient to support LoRa operations within a reasonable form factor. Alternatively, such a system could utilize a rechargeable battery. A single small rechargeable battery is more energy dense and significantly less costly than the bank of three large supercapacitors used. This would result in a smaller, lighter, and cheaper system. The authors' concerns of battery cycle lifetime, charging characteristics, and voltage requirements, are actually inconsequential, especially for this application. They do not provide any citations for any of these assertions. Given the application's infrequent periodicity, cycle depth is shallow, and average discharge and charge rates are low, resulting in less cell damage and an extended cycle lifetime [84, 85]. For most lithium-based batteries, this would result in more than the 10 year lifetime of the instrumented water heater, never mind the 1-3

year lifetime of the sacrificial anode. The authors' issue with voltage levels is easily addressed with a buck or boost regulator, depending on the battery chemistry. These trade-offs between capacitors, supercapacitors, and batteries are further explored in Chapter 5.

Smart Dust and the Future

Beyond applications that were possible ten years ago, as well as ones that are possible today, it is important to look forward to the requirements of future applications. The continuing trend in computing and wireless sensors is pushing the boundaries on size and power. Researchers have begun prototyping a new class of millimeter-scale computers and sensors, known colloquially as "smart dust" [86].

While some developed millimeter-scale systems are one-off demonstrations of the feasibility of the device class, some have been developed for specific applications. In particular, millimeter-scale sensors are uniquely appropriate for injectible or implantible medical applications. Millimeter-scale devices have been proposed as syringe-implantable medical sensors to to measure and monitor intraocular pressure for glaucoma diagnosis and management [87], heart arrhythmia [88], glucose levels for diabetes treatment [89], and cellular temperature to monitor the metabolism of cancer cells [90]. General purpose, modular, millimeter-scale sensing platforms have also been proposed for traditional sensing applications [75, 91]

The lifetime goals for many of these applications are to operate perpetually based on available energy [75, 91]. For many of the medical applications, lifetimes of a few years to decades are appropriate to avoid frequent sensor re-injection and re-implantation [88]. Typically, these applications require high availability to support reliable and periodic measurements. High quality, consistent measurements are necessary to establish reliable trends for the basis of diagnoses or treatments. Likewise, non-medical smart dust applications also require high availability. A millimeter-scale image sensor application utilizes always-on motion detection as opposed to opportunistic or periodic image sensing [91]. This approach attempts to maximize the value of captured images and minimize energy used to capture non-interesting images.

Despite the superior energy density of primary lithium batteries, at the scale of a millimeter cubed, the energy provided is insufficient to offer an adequate lifetime for these systems [75]. To achieve perpetual lifetimes, these systems utilize harvesting methods including photovoltaic, active RF, intra-body transfer, or optical harvesting [75, 89, 90]. Due to their size, these systems have limited volume for energy storage, and likewise limited surface area available for energy harvesting. This results in aggressive system-wide power limitations to achieve adequate lifetimes with high availability. Unlike the previously mentioned devices and platforms, most of which require on the order of hundreds of µW to mW of average power to drive their application, smart dust devices require tens or hundreds of nW to operate [75, 88, 90, 91]. Given the low energy density of capacitors and supercapacitors, they provide insufficient energy storage to allow continuous and perpetual operation, making batteryless approaches unsuitable for many smart dust applications. Instead, to achieve sufficient energy storage, most of the aforementioned millimeter-scale systems utilize small rechargeable thin-film lithium batteries to allow operation between periods of harvestable energy. These

batteries are sized to provide a few days of operation without any income, corresponding to capacities on the order of 1-5 µAh.

Not all applications require the extreme size constraints of smart dust. However, they can benefit from the advancements in power efficiency developed by smart dust researchers. Within the next decade it is conceivable that microcontrollers, sensors, and radio technologies will be developed with the power improvements originally meant to enable millimeter-scale computing. These sensors can remain at the scale of traditional sensor motes with relatively vast energy storage and energy harvesting capability, but with one or two orders of magnitude improvement in active and idle power. This will result in sensors that can easily achieve lifetimes of multiple decades, even with small primary batteries. With average system power on the order of nW, batteries will be limited more by self-discharge than the actual system load. For energy harvesting systems with sufficient surface area, an average income power on the order of tens or hundreds of nW to support perpetual operation is an easy task even in dim environments. Beyond lifetime extension, these sensors will be capable of doing more with the energy they have. Future systems will perform advanced sensing and processing that are currently infeasible due to power and energy constraints. Some of the advancements that have led to the power efficiency of millimeter-scale systems are starting to be commercialized. In particular, the development of high performance subthreshold and near-threshold computing has resulted in an order of magnitude improvement in active power over traditional CMOS supply voltage techniques [92, 93]. Ambiq Micro has led the commercialization of this technology, starting with ultra low power real time clocks, and more recently developing microcontrollers like the Apollo4 that requires only 5 µA/MHz. This represents an order of magnitude improvement over other modern commercial offerings, like the Nordic nRF5340, which requires 56 µA/MHz [94].

2.4 Summary

Traditional battery-powered wireless sensors cannot provide long lifetimes without sacrificing size, and upon inevitable battery replacement or sensor retirement, result in significant battery waste. Modern energy harvesting approaches forgo batteries to achieve indeterminate lifetimes, but they sacrifice software and hardware design simplicity, provide no guarantees of availability, and capture significantly less harvestable energy than is available.

Either option is a not entirely satisfactory solution for many classes of applications. Many applications have quality of service requirements that make a batteryless approach less amenable than a traditional battery-based one. However, in Section 2.3, we have described several niches of applications where batteryless designs are an appropriate design choice. They excel in environments with robust and widespread infrastructure for power transfer or a reliable source of harvested energy. They are a reasonable option in applications where the sensing modality can simultaneously be harvested from. Outside of these situations, a batteryless-based application must relax its requirements for data consistency or availability. For applications in adversarial environments where any data is better than no data, a

batteryless design may be appropriate. For any other application without the convenience of infrastructure or other reliable harvesting source, any quality of service requirements regarding the frequency, consistency, and availability of collected and reported data necessitate a different design path. Additionally, any application that cannot justify the design cost of the extra complexity that a batteryless design requires will require an alternate design approach.

In the next chapter, we make the argument that the distinction between different power supply classes is a less useful way of thinking about the design space. This thinking delineates based on the end result of a design, like whether it utilizes energy harvesting, any non-rechargeable energy storage, or the type and size of rechargeable energy storage. This leads system designers to begin their design by choosing a power supply type, instead of choosing a power supply for their application. Since batteryless systems are popular and exciting, this has resulted in many researchers making the decision to utilize a batteryless design without a proper analysis of what their end application actually requires. Instead, it is more useful to reason about an application's requirements and consider the options for a power supply. A proper design must determine the quantity of type of energy storage (both rechargeable and non-rechargeable) and energy harvesting required to meet an application's requirements. In the next chapter, we consider the myriad of system-level design options for a wireless power supply with a focus on determining heuristics for the sizing and inclusion of energy storage and energy harvesting techniques.

Chapter 3

Developing System-Level Power Supply Design Heuristics

Wireless power supplies are the result of a multi-dimensional design space exploration. While certain design decisions are straightforward, like the choice to use a certain type of harvester to suit an application environment, some design points are more difficult to determine. Harvesting and storage technologies are often difficult to directly compare, and it is sometimes difficult to determine the sizing of these elements in order to satisfy the constraints of an application. Many times, the type and size of harvesters, batteries, or capacitors are chosen arbitrarily [9, 75, 95], or are chosen to make the design minimally feasible [2, 7, 11, 55].

The goal of this chapter is to provide high-level guidance for navigating the design space of wireless sensor power supplies. With the common application requirements from Section 2.3 in mind, this chapter provides guidance for designers as to when energy harvesting is a beneficial technique for energy income instead of, or in addition to, a non-rechargeable storage. We also explore how prospective income and workloads drive component types and sizing, and if energy-harvesting is utilized, what software and hardware techniques are required to ensure a feasible design, or what is required to capture sufficient energy to achieve acceptable system availability. To limit scope of this chapter to be manageable, we focus exclusively on indoor photovoltaic energy harvesting. Occupied indoor environments are the focus of a significant amount of prior work, and for good reason: most applications aim to improve the lives of people and are necessarily present in the spaces they occupy [7, 28, 31, 52, 53, 57, 58, 96–98]. Indoor environments are lit by both natural and artificial light, and may occasionally get direct or indirect sunlight. Under these conditions, photovoltaic harvesting still offers an order of magnitude more energy than other methods. However, indoor environments still present a significantly more challenging and low power design space compared to the outdoor ones. While we focus on indoor photovoltaic harvesting in this chapter, the conclusions drawn from the following analyses are still applicable to other environments and harvesting methods.

3.1 Energy Income

A wireless sensor requires a source of electrical energy to function. Energy-harvesting has the potential to supply energy indefinitely. However, harvested energy is often variable and unpredictable, and instantaneous power delivery is limited. Conversely, primary-only systems have a limited amount of energy, but can use that energy at any rate they please, within the power limits of the battery. Because of these differences, it can be difficult to directly compare their performance and determine at what scale and in which conditions harvesting, energy preallocation, or a hybrid of the two, is the preferable power source strategy for any given application. This section begins with a description of constraints for properly sizing components to ensure a sufficient lifetime and energy income. Next, it provides a method of comparison of energy capture potential between harvesting and preallocation methods. Finally, the constraints and relations defined and explored in the following section can be used to estimate proper harvester and primary battery selection to achieve application requirements.

Sizing

Regardless of how a sensor is powered, it must have sufficient energy at the right times to continuously and reliably operate. Assuming that the average power required to drive an application is known at design time, it is relatively straightforward to determine the correct size of a battery to achieve an expected lifetime, or a photovoltaic panel to achieve an average power income. Usually battery capacity is expressed as accumulated charge in terms of mAh. If the nominal voltage of the battery is known, it can be used to estimate lifetime given a battery's capacity and intended workload.

$$E_b = UC \tag{3.1}$$

$$\bar{P}_b = \frac{UC}{T} \tag{3.2}$$

Equation (3.1) describes the energy contained in a battery with a charge capacity C and a nominal voltage U. For example, a 1 Ah battery with a nominal voltage of 3 V would provide 3 Wh of energy. Assuming a desired lifetime T, Equation (3.2) describes the maximum achievable average power supplied by such a battery over that lifetime. Likewise, the average power provided by a photovoltaic of area A is described by Equation (3.3).

$$\bar{P}_h = \eta \bar{E}_e A \tag{3.3}$$

Where η denotes the efficiency of the photovoltaic, and is often between 15 and 20%. The average energetic spectral irradiance, $\bar{E}_e$, is generally between 10 and 100 $\mu W\,cm^{-2}$ for indoor conditions [2, 99]. Given an average workload power $\bar{P}_w$ and a desired lifetime T, Equations (3.4) and (3.5) describe the necessary battery capacity and solar cell size.

$$C \geq \frac{\bar{P}_w T}{U} \tag{3.4}$$

$$A \geq \frac{\eta \bar{E}_e}{\bar{P}_w} \tag{3.5}$$

For some applications, it is obvious when utilizing energy harvesting is a better design choice over a battery. Such is the case with Zebranet collars [9]. Powering the collar would require a large and heavy battery, beyond the weight constraints of application, and would only power the system for 5 days, far below the lifetime goal of a year or more of operation. In other cases, it is much more difficult to determine which method is appropriate. This is especially apparent when available harvestable energy is limited, like indoor environments.

Harvesting can Provide More Energy than Preallocation

When potential harvestable energy is limited, it is difficult to determine whether energy preallocation or harvesting is the most suitable technique. It is very difficult to directly compare these methods, as the power and energy supplied by either option is dependent on many factors, including the variance and magnitude of available harvestable energy, the size of the battery or harvester, and the intended sensor workload. Prior work has attempted to approach a comparison of preallocated energy storage and energy harvesting by comparing the expected average power supplied by either option. They simplify the problem by considering a theoretical cubic sensor of volume $V = L^3$, with the assumption that the entire cubic sensor volume is dominated by a lithium primary battery, or the entire area of one $A = L^2$ face is dominated by a photovoltaic panel [2]. The authors compare the power provided by the battery cube, with that of the photovoltaic square, over the shelf-life of the battery.

For this analysis, Equation (3.1) is not useful for direct comparisons. A more appropriate method utilizes volumetric energy density. Battery energy capacity (E_b) can be estimated based on the volume of the battery (V) and its volumetric energy density (ρ). While ρ varies depending on the specific battery chemistry, size, and packaging, lithium primary batteries usually provide on the order of $800\,\mathrm{mWh/cm^3}$ [78]. The maximum average power ($\bar{P}_b$) provided by the battery can be calculated with a desired lifetime (T).

$$E_b = \rho V \tag{3.6}$$

$$\bar{P}_b = \frac{\rho V}{T} \tag{3.7}$$

When considering Equations (3.6) and (3.7), the authors use a more conservative $653\,\mathrm{mWh/cm^3}$ for ρ, and assume a 7 year lifetime based on shelf-life. They use this lifetime in order to calculate the maximum power that a primary cell can provide over this period [2]. This assumption for battery lifetime is very conservative, as it considers a primary battery entirely empty or unusable after its reported shelf-life. Assuming a static lifetime also ignores the impact of a sensor's workload power requirements on battery longevity. Primary battery

shelf-life is often poorly defined and understood. While some take shelf-life to mean a battery is expired and not usable after this time, shelf-life actually represents a manufacturer guarantee that a cell will retain a majority of its capacity over the shelf-life period. A 10 year shelf-life is common for lithium metal primary batteries [100–102]. Energizer claims that their lithium metal primary batteries experience a self-discharge rate of 1% per year at room temperature and humidity. This corresponds to 90% remaining capacity at the end of their listed shelf-life period [102]. A lithium battery can hardly be considered empty after its shelf-life has expired. A more accurate estimation for battery lifetime is:

$$T = \frac{\rho L^3}{\bar{P_w} + P_l} \tag{3.8}$$

Where $\bar{P_w}$ is the average power required to drive a sensor workload, and P_l is the self-discharge power of the battery. Any aging effects other than self-discharge are unpredictable and we do not consider them here. With this new definition of battery lifetime, the comparison of power is not as useful a metric, as workload power provided by the battery (P_w) is independent of the battery capacity. The power provided by a battery *is* technically limited by its maximum rated load current, which is often related to its capacity. However, this limit is not usually relevant when considering low power operation.

Conversely, the power provided by an energy harvester is unrelated to the requisite workload power at all, and is incomparable with that of an on-demand preallocated power source. Instead of power, the energy provided by a primary cell and the energy captured by a similarly sized photovoltaic over the course of the lifetime of a primary cell are directly comparable. The energy captured by a photovoltaic over a time period T is represented by Equation (3.9). The authors of [2] assume the lower end of irradiance to provide a conservative estimate, but ignore photovoltaic efficiency η. This results in an overestimation of harvestable power and energy. This analysis assumes a conservative photovoltaic efficiency of 17%. Equation (3.9) is a more realistic representation of captured photovoltaic energy.

$$E_h = \eta \bar{E_e} A T \tag{3.9}$$

With Equations (3.6) and (3.9), it is possible to calculate the energy contained in a battery of size L^3, the lifetime of that battery given an average workload power P_w, and the energy captured by a photovoltaic panel of size L^2 over the battery lifetime. This is visually presented in Figure 3.1, an updated energy-harvesting reality check [2]. This figure compares the preallocated energy provided by a battery with that of the potential harvestable solar energy over the same time period. Both battery and harvester energy are driven by the same dimension L. The energy offered by a battery is represented by the blue line, and increases with L^3. This energy, coupled with power requirements of different workloads, define the lifetime of of the battery (T). The orange and red lines correspond to the energy captured by a photovoltaic panel of area L^2 over the lifetime of the battery under different workloads. The two different workloads are represented by the orange (average 25 µW) and red (100 µW) lines. The extents of average indoor photovoltaic harvesting are represented by dashed (10 µW/cm^2)

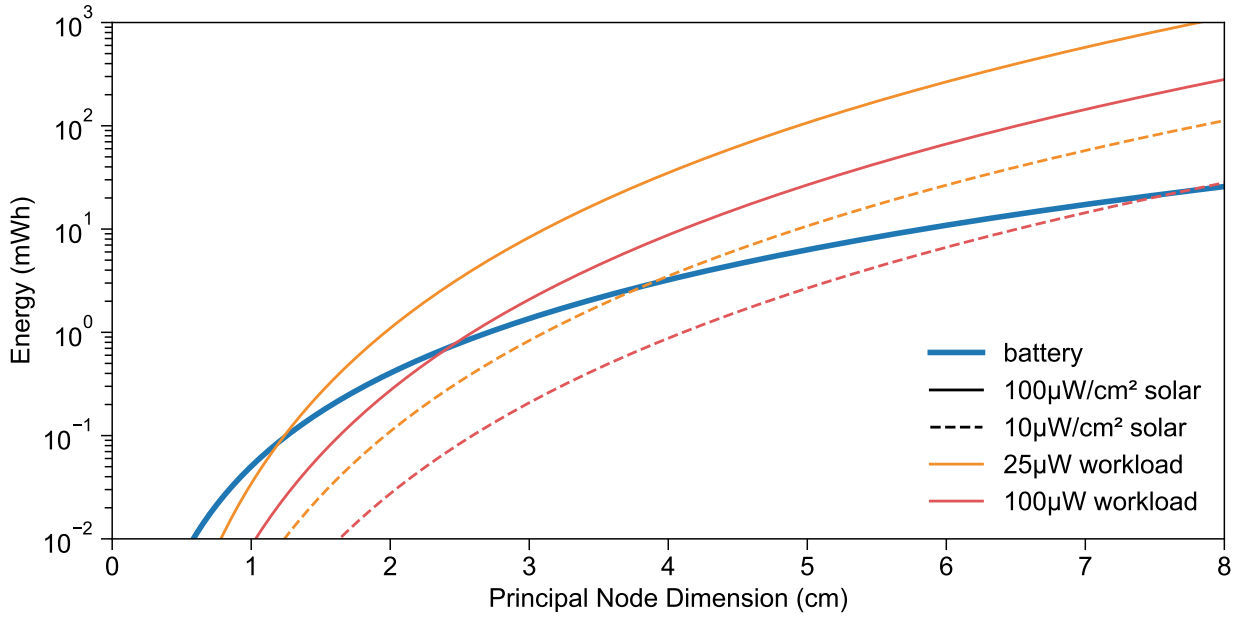

Figure 3.1: A comparison of preallocated energy and captured energy. Note the logarithmic y-axis scale. This figure compares the energy offered by a cubic battery with that of potential harvestable energy captured by a square photovoltaic over the lifetime of the battery. At a sufficient size and in sufficient harvesting conditions, while powering an appropriate workload, solar energy-harvesting can provide more energy over the same time frame as a lithium battery.

and solid (100 µW/cm^2) lines. The point at which the harvesting lines (orange, red) cross the battery line (blue) indicate the size at which a solar panel of size L^2 will harvest the same amount of energy provided by a battery of size L^3 over the lifetime (T) of the battery. These crossing points also indicate the size at which harvesting collects sufficient average power to drive its intended workload, approaching energy neutral operation. For appropriate workloads and lighting conditions, the crossing points suggest that a sensor with a driving dimension larger than 4 cm will be able to harvest more energy than is contained in a similarly sized battery for all but the most taxing income and workload scenarios.

The same analysis can be done for designs that scale down in both size and power. Millimeter-scale systems like the Michigan Micro Mote occupies 1 mm^2 and requires on the order of 100s of nW to power its workload [75]. Figure 3.2 is a reconfigured analysis, tuned for nW workloads and millimeter-scale systems. The same dark blue line represents the energy preallocated with a lithium battery of size L^2. The green lines represent a 100 nW average workload, and the teal lines represent 25 nW workload. This figure considers the same lighting conditions as the previous figure. Like for µW systems, there are sensor sizes where workloads and lighting conditions result in a sensor that will be able to collect more energy over time with energy harvesting than can be preallocated with a battery. This reaffirms the

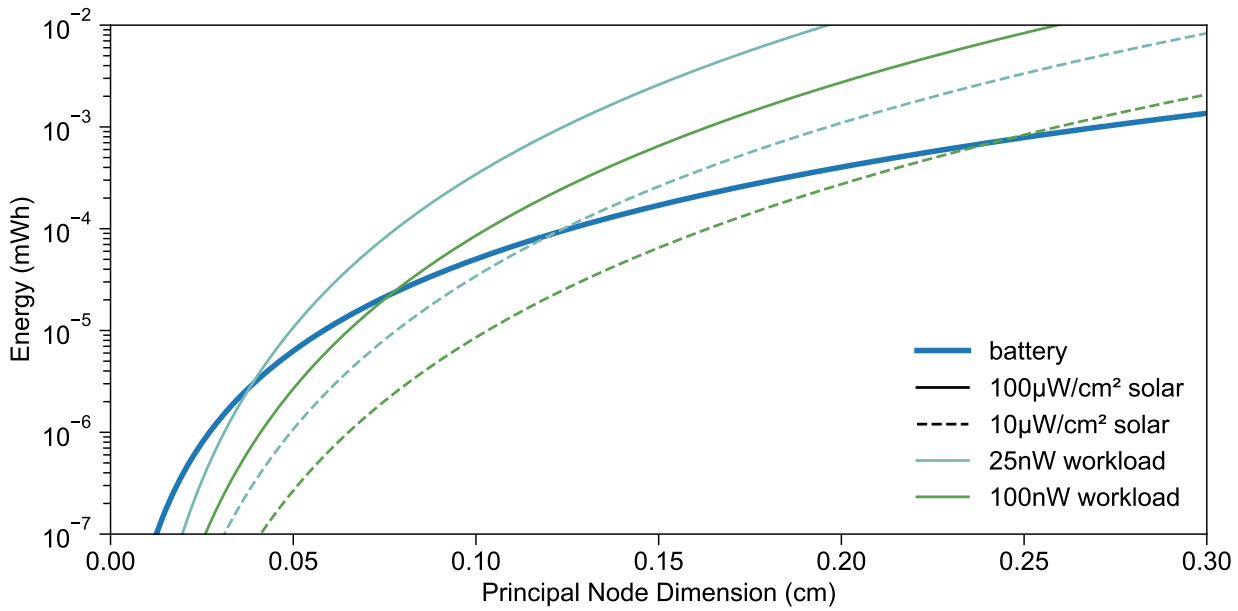

Figure 3.2: A comparison of preallocated energy and captured energy for nW applications. This figure is identical to Figure 3.1 except that it considers the nW workloads that characterize millimeter-scale systems. A principle node dimension on the order of 1-2 mm is generally sufficient for energy harvesting to collect more energy over a battery, in all but the worst case: a heavier workload with low harvesting potential.

design decisions made by the designers of many of the millimeter-scale systems, who chose to utilize energy harvesting to prolong the lifetime beyond that of a non-rechargeable battery.

A Hybrid System Maximizes Energy Capture and Reliability

At the crossing points in Figures 3.1 and 3.2, the benefits of energy harvesting are compounded when considering that energy capture will continue indefinitely, beyond that of the lifetime of a battery. However, this energy is not guaranteed to be supplied at times when the system requires it, even when harvesting supplies enough on average to match that of workload. This reality has the potential to result in unexpected system outages and system resets. To ensure that a system always has sufficient energy to operate, a hybrid system can utilize energy harvesting with a backup preallocated energy storage. This section considers the same cubic sensor as before, but combines the energy preallocated by the battery and captured by the photovoltaic panel. A hybrid system can operate even in the absence any ambient or accumulated harvestable energy. A hybrid system will possess a finite lifetime, but will have guaranteed reliability during its lifetime. This lifetime primarily depends on the disparity between its intended workload and the availability of harvestable energy. This disparity can be expressed by Equation (3.10). If the average power supplied by a photovoltaic harvester

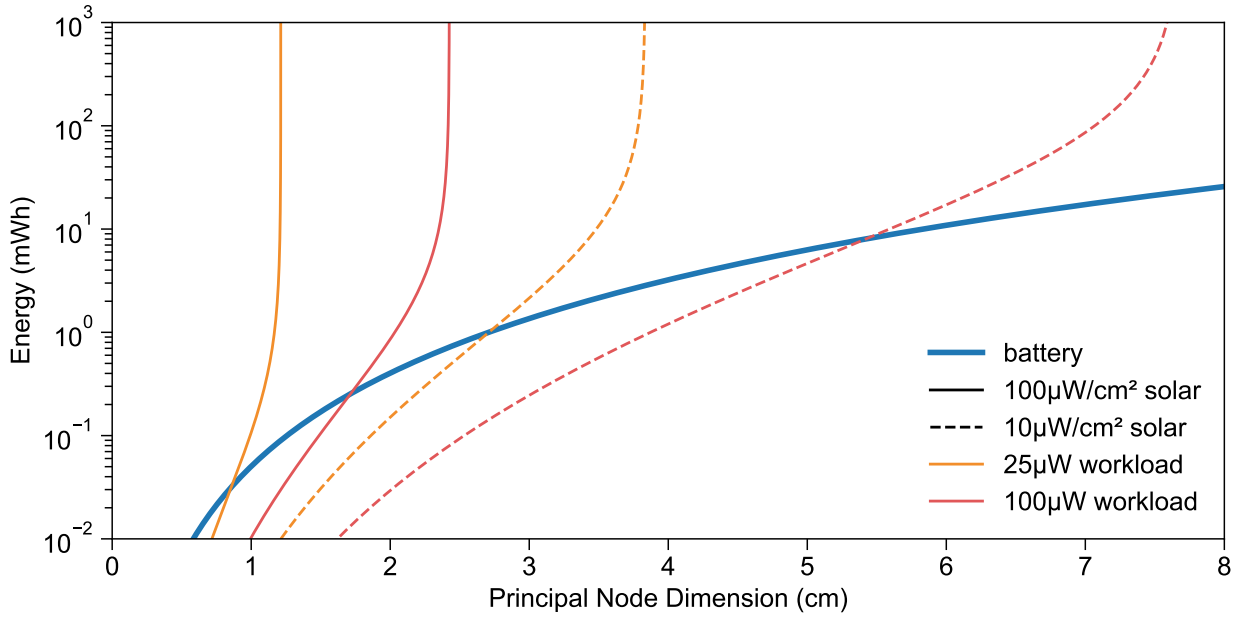

Figure 3.3: The energy captured by a hybrid system utilizing both harvesting and backup preallocated energy. This figure uses the same scale and line types as Figure 3.1. The addition of energy harvesting to a primary cell system has a compounding effect on lifetime and harvested energy. More harvested energy results in a prolongation of the lifetime of the primary cell. Subsequently, this lifetime extension results in an increase in harvested energy. Because of the increased energy and battery lifetime, the crossing points now shift to the left in the figure, allowing a reduction in volume and area required for a battery and harvester, respectively.

is sufficient to drive the intended workload, the total power drain, P_t, on the battery is equal to its self-discharge power, P_l. If the photovoltaic does not harvest sufficient power for the workload, the battery is drained at an average rate of the difference, plus the battery self-discharge. Given this new definition of sensor power, the lifetime of its battery now resembles Equation (3.11).

$$\bar{P}_t = \begin{cases} P_l & \text{if } \bar{P}_w < \bar{P}_h \\ \bar{P}_w - \bar{P}_h + P_l & \text{if } \bar{P}_w \geq \bar{P}_h \end{cases} \tag{3.10}$$

$$T = \frac{\rho L^3}{\bar{P}_t} \tag{3.11}$$

If the amount of average harvestable power is greater than required for a sensor workload, the lifetime of a hybrid system is compounded. The additional energy captured lengthens the lifetime of the system, and the longer lifetime allows for additional energy capture. Figure 3.3 explores this relationship visually, using the same workloads, lighting conditions, and visual

scale and line types as Figure 3.1. The additional energy capture results in the battery line crossing points shifting to the left, meaning that in a hybrid system a smaller harvester and battery are required to capture the same amount of energy as that provided by a battery when compared to a harvester-only system. For a hybrid system, these crossing points represent a design point that is able to harvest the same amount of energy as it has preallocated, in effect providing twice the energy as a battery-only system of the same dimensions. As the size of the system increases beyond the size at which the harvesting lines and battery line cross, the energy supplied by harvesting begins to approach a vertical asymptote. This represents the size at which the average income from harvesting is sufficient to power the workload on its own. When harvesting can sustain the system indefinitely, the backup primary battery is infrequently or never utilized, and its lifetime is limited by its self discharge. These locations where the harvesting energy lines approach a vertical asymptote represent the ideal sizing for a hybrid system given an intended workload and lighting conditions. This ideal sizing results in a harvester large enough to capture enough average power to results in energy neutrality, where the system harvests enough energy to sustain its operation.

Limitations

The above analysis is a simplified view of a complex reality. The analysis assumes that a sensor's workload and income can be treated as static averages. While it can be argued that on a timescale of hours to days, a sensor's workload power can be consolidated into an average for many applications. The same assumption cannot easily be made for energy income. For example, lighting conditions are often both diurnal and seasonal, meaning that a daily or even weekly average income from a photovoltaic will vary significantly over the course of a day or year depending on the angle and intensity of sunlight, or the occupancy and activation of artificial lighting.

For simplicity's sake, the analysis of this section also assumes the whole sensor volume is available for energy storage, and a whole face is available for energy harvesting. This assumption breaks down when considering sensors and PCB elements that require substantial volume, such as a PIR sensor, or a sensor that cannot be covered by a harvesting element, such as an image or light sensor. However, the relationship to energy capacity and harvesting capability are directly and linearly related to volume and area respectively. The general conclusions should hold, even if some percentage of the volume and area of a sensor are occupied by elements other than energy storage or harvesting.

The above analysis also does not consider the effect of rechargeable energy capacity on energy income from harvesting. The analysis assumes an ideal rechargeable energy storage with limitless storage. In reality, energy storage elements are far from limitless or ideal. A small energy capacity will fill up often, requiring shunting of any additional harvested energy. By shunting, or wasting this energy, a system with insufficient capacity will effectively reduce the average energy income power available to it. Sections 3.2 and 3.3 explore the effect of capacity in further detail.

The analysis of this section also does not consider the effects of miniaturization, especially as components approach the scale of millimeters and their packaging begins to dominate their volume. As batteries reduce in size, the proportion of volume available for energy storage versus the volume taken by packaging decreases, resulting in less energy density. This analysis assumes a static energy density for lithium primary cells, when in reality the density would also decrease as the size of the battery decreases. Similarly, harvesting elements experience less surface area for the actual harvester versus area required to affix and provide electrical connections at smaller scales. System designers seeking to build extremely small devices will have to consider additional factors when considering element sizing. Despite these limitations, the analysis of this chapter has enough basis to serve as a guiding rule-of-thumb when considering methods for energy income as well as component sizing requirements for various applications.

3.2 Harvesting Feasibility and Intermittency

From the previous section, it is obvious that the availability and magnitude of harvestable energy directly impacts the energy income of an energy harvesting system. Income limits which and how many operations it can complete and how fast or frequently it can complete them. What is perhaps not as obvious is that the rechargeable energy capacity used to capture harvested energy also impacts the overall energy income and ultimately the operation of the system. The availability and magnitude of harvested power, the size of an energy buffer, and the intended sensor workload determine whether a design is minimally feasible, and if so, whether the design will be reliant on software or hardware intermittent techniques to ensure proper operation. The necessity of the techniques mentioned in Section 2.2 depend solely on the energy income of a system and how it relates to the design's efficiency when idle as well as its energy capacity. This section seeks to illustrate the relationship of income and storage on the resultant sensor operating regimes.

Design Regimes

Figure 3.4 represents an illustration of a wireless sensor design framework. This framework splits the design space into four main regimes: Always on, Infeasible, State retention required, and State retention not required. Additionally, the figure illustrates the conditions in which hysteresis management techniques are helpful or less helpful. The following sections describe these regimes and their constraints in more detail.

Always On

If the energy harvester reliably supplies a sensor with more power than the maximum instantaneous power it will ever draw, then the power supply does not need significant energy

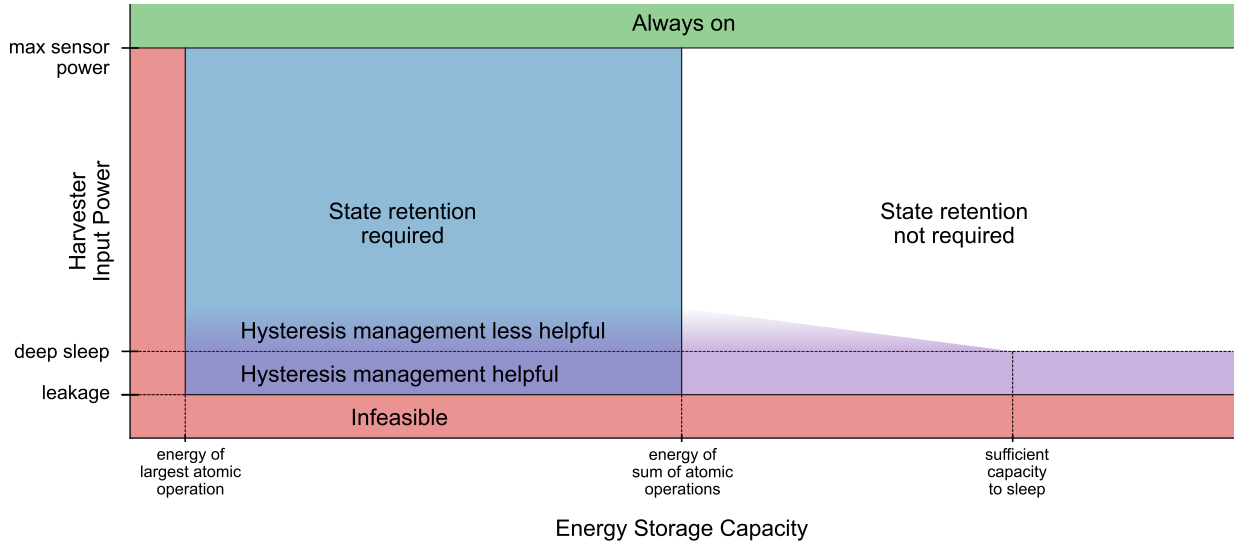

Figure 3.4: Design space for energy-harvesting sensors based on their energy income (assumed constantant for this analysis), energy storage capacity, and workload. Workload is represented by the set of atomic operations required by an application, as well as the deep sleep and leakage power. The plot breaks into four regions: Always On or effectively powered, Infeasible due to lack of energy storage or leakage higher than harvesting rate, feasible but Requires State Retention to make forward progress, and enough energy storage so that State Retention is Not Required. Additionally, sensors which have high power when they enter deep sleep before depleting their energy buffer may benefit from Hysteresis Management techniques. This benefit diminishes with lower sleep currents and higher harvesting potential.

capacity to remain operational. This design point constraint is defined by Equation (3.12).

$$\max P_w \leq \bar{P_h} \tag{3.12}$$

Where $\max P_w$ represents the maximum instantaneous power required by a sensor workload, and $\bar{P_h}$ is the average harvesting income. Additionally, P_h must also have minimum variance ($\text{Var}(P_h) \approx 0$), on the order of normal power ripples handled by power supply bypass capacitors. If the harvesting power has significant variance or frequent outages, then a sensor must have some ability to buffer energy to use when its instantaneous operating power exceeds that of the harvester input power. This green region of Figure 3.4 represents design points that have sufficient income power to be Always on.

Infeasible

If the energy harvester supplies less average power than the system self discharge and leakage, the system will very infrequently be able to charge its energy buffer. Leakage can exist in a system via component quiescent currents, energy storage self-discharge, and parasitic current

paths that are not eliminated when the system is idle or powered off. The constraint on minimum harvesting income and system leakage is described by Equation (3.13).

$$\bar{P_h} > P_l \tag{3.13}$$

Where P_l is the total average system leakage in a powered off state. Similarly, if the energy buffer capacity is less than the energy required to perform a workload's largest atomic operation, then that operation will not have enough energy to complete. The total energy required for a sensor to complete an iteration of its workload can be represented by a the finite multiset of the energy required for n atomic operations. This multiset is described by X in Equation (3.14)

$$X = \{e_i\}_{i=0}^{n} \tag{3.14}$$

Each individual e_i represents the energy required for a single atomic operation such as sampling a sensor, sending a radio packet, a processor power on reset, and performing a checkpoint. The energy capacity ($E_{capacity}$) of a system must be equal to or greater than the largest atomic operation ($\max X$). This constraint is described by Equation (3.15).

$$E_{capacity} \geq \max X \tag{3.15}$$

Designs that do not satisfy the constraints of Equations (3.13) and (3.15) have insufficient income to ever power on, or insufficient energy capacity to perform operations and are therefore infeasible. The red regions of Figure 3.4 represent design points that are Infeasible.

State Retention Required

If the energy capacity of the energy buffer is sufficient to perform the largest atomic operation as in Equation (3.15), but not enough to complete the entire multiset of atomic operations in the workload, E, then a mechanism for state preservation is required to ensure forward progress over power loss and reboots. A design requires state preservation techniques if Equation (3.16) is satisfied.

$$\max X \leq E_{capacity} < \sum_{e \in X} e \tag{3.16}$$

In Figure 3.4 the blue region labelled State retention required encompasses designs that satisfy Equation (3.16) and require state retention techniques to ensure proper operation.

No State Retention Required

A sensor that has a large enough energy buffer to support all of the atomic operations of its workload, and has sufficient harvesting capability to charge this buffer does not require state retention methods. All of its workload can be completed with a full energy storage element, without requiring a reboot and recharge. This constraint is defined by Equation (3.17).

$$E_{capacity} \geq \sum_{e \in X} e \tag{3.17}$$

Systems satisfying Equation (3.17) exist in the white region of Figure 3.4. In this region, systems can avoid the complexity of state retention software frameworks, and the energy that would previously be devoted to performing checkpoints or writing and restoring data to non-volatile memory can be used toward the workload instead. However, a system that satisfies Equation (3.17) may still experience intermittent operation depending on the variability of its energy income and size of its storage. Such systems would occupy the white region of Figure 3.4 labeled State retention not required.

Hysteresis Management

Hysteresis management is useful for sensors that frequently reboot and must cold start, recharging their storage element from, or close to, empty. These techniques are particularly beneficial for frequently restarting systems that have large energy buffers that require significant charge and time to provide a sufficient and stable system voltage. Often, intermittent and batteryless systems are designed to turn off when there is no harvestable energy. Because of this, they have not optimized the power required to maintain a low power, state retaining mode. Under these conditions, hysteresis management techniques, such as reconfigurable capacity [18] and federated energy [7] can decrease the time and energy required to cold start, increasing sensor performance. These methods can decrease cold start time by reducing the capacity that must be charged to achieve a cold start and partial system restart.

While a system's cold start length and frequency are not well represented within the metrics considered by Figure 3.4, the constraints between a design's deep sleep power, energy income, and energy capacity are. A design with average harvestable income ($\bar{P}_h$) less than what it takes to sustain a low power sleep state (P_s) is better off turning off instead of sleeping. This constraint is defined by Equation (3.18).

$$P_l < \bar{P}_h \leq P_s \tag{3.18}$$

Sensors within this regime will benefit from hysteresis techniques to reduce the amount of charge and time it takes for them to cold start. They will occupy the purple region in Figure 3.4 labelled Hysteresis management helpful.

Sensors that are able to harvest more power than it takes to maintain a deep sleep state will still benefit from hysteresis management if they frequently must cold start, or they have a lengthy and energy intensive cold start. The utility of hysteresis management is diminished when the ratio of harvester power to deep sleep power increases and when the energy capacity of a system increases. More energy and the capacity to hold it reduces the frequency of cold starts, and capacity has an averaging effect on energy income, providing power when income is insufficient, and storing in situations of excess. Sensor designs within this region represent a gradient of diminishing returns on the utility of hysteresis techniques. These designs occupy the space in Figure 3.4 above the horizontal line for deep sleep power, extending to the point of sufficient energy capacity for sleeping. This region is represented by a gradient, as the regime is not as clearly delineated as others as it depends on factors other than harvesting input power and energy capacity.

Designs that would not benefit from hysteresis management represent designs with sufficient energy income and capacity. With enough income and capacity, sensors can willfully enter low power sleep states for long periods, avoiding powering off and subsequent cold starts. These sleep states can be managed by the sensor itself, instead of being imposed by the variability of energy income. When sleeping, deep sleep power becomes analogous to leakage power, and hysteresis management techniques will not improve recharge times. While it is relatively straightforward to determine component sizing to achieve proper energy income, as described in Section 3.1, it is more difficult to determine the proper energy capacity to avoid power loss and frequent cold starts even with sufficient income. Section 3.3 attempts to identify and quantify the proper energy capacity that maximizes energy capture and minimizes cold starts.

Limitations

This framework on feasibility and design regimes makes a few simplifying assumptions to describe a complex design space. Similar to Section 3.1, The framework assumes a constant average energy harvesting income, when in reality income is often highly variable. In practice, a sensor platform defines the regions and limits of the plot, such as the leakage and deep sleep boundaries. The platform then occupies a vertical line which represents the the range of harvester input powers it might experience fixed at its designed capacity. By ignoring variability, the plot also fails to illustrate key benefits of increased capacity under varying energy incomes and workloads. Intuitively, a large capacity can store energy in times of excess and supply that energy in times of drought. This balancing out of energy income effectively increases the average power supplied by the energy harvester, up to the theoretical maximum available to be harvested. The extent of this impact is completely dependent on the variability of the energy income and workload of the sensor. In the following Section 3.3, the performance impact as related to capacity is examined more closely.

This framework also does not consider the impact of a backup energy store. A backup energy store can be viewed as the ability to inject additional energy to the system at arbitrary times, eliminating the need for state retention when there is very low harvesting potential. A backup energy store could also contribute in more subtle ways. It could allow a system to avoid the energy and complexity of state retention by providing just enough energy for a deep sleep mode with state retention rather than a full power down when the system depletes its stored energy. It could also cold start energy buffer charging to eliminate the need for reconfigurable or federated power supplies, or to increase the efficiency of the energy-harvesting front-end at low voltages. While the use of a backup energy store does constrain the sensor to a finite lifetime, as discussed in Section 3.1, energy-harvesting can substantially extend these lifetimes under certain harvesting conditions.

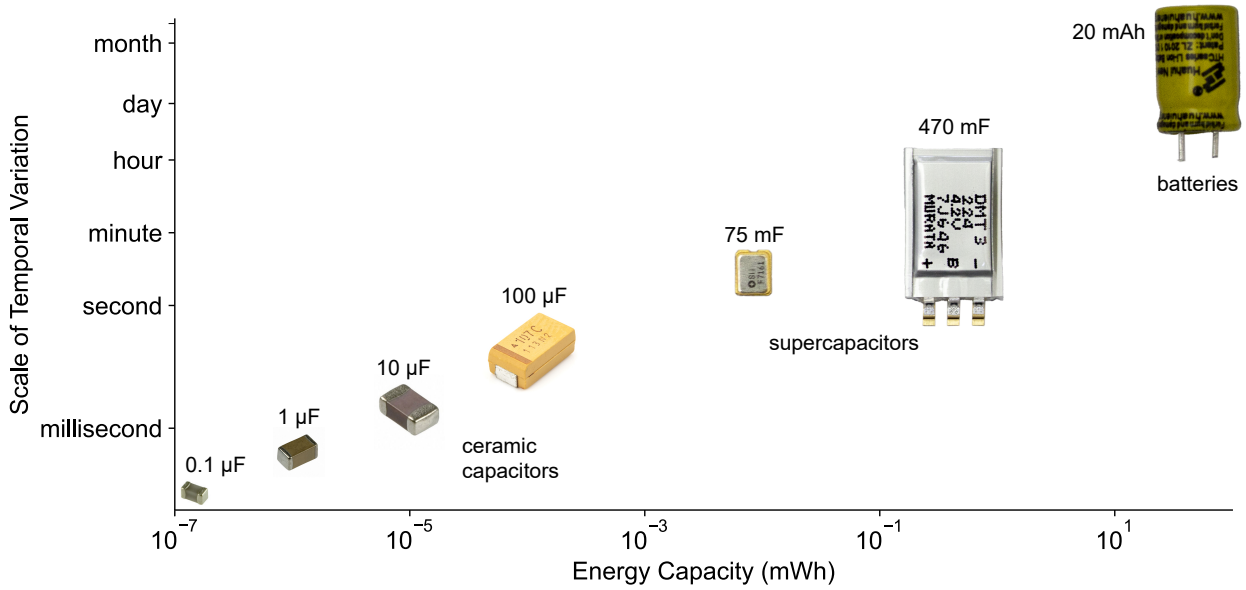

Figure 3.5: The continuum of the efficacy of energy storage technologies for averaging and filtering power at different temporal scales. Small energy buffers filter out high frequency power supply noise and spikes, while larger energy capacities filter out power variance on larger temporal scales.

3.3 A Case for Capacity

There are currently many approaches to determining the energy capacity for energy harvesting sensor designs. Many batteryless systems that expect frequent restarts size their energy storage to meet minimum design feasibility and to minimize charging hysteresis. They size their energy storage to support the largest atomic operation required by their application, as described by Equation (3.15). For example, the sensors developed for the aforementioned deployment at Mithræum of Circus Maximus have energy storage sized to provide enough instantaneous energy to ensure the completion of a few atomic operations and subsequent forward progress [11]. Other platforms also adopt this approach for sizing: Both the previously mentioned cathodic protection system and Camaroptera size their supercapacitor bank to support sending a single atomic LoRa radio packet [12, 56], SkinnyPower sizes its capacitor to provide enough energy for the boot cycle of its SoC [59], and the Flicker platform's federated capacitors are all sized to support predefined atomic operations for its peripherals [7]. Similarly, one-shot batteryless platforms size their energy buffer to support the entirety of their small workload [2, 3, 55, 63].

Other more traditional energy harvesting systems size their energy buffer to support an arbitrary amount of runtime without harvestable energy. The Pible platform uses a supercapacitor that allows 2 hours of BLE advertizing at a 100ms period, but employs a dynamic power management technique to scale back advertising frequency when available

Irradiance Trace	Total Days	Average Irradiance	90^{th} Percentile Daily	10^{th} Percentile Daily
EnHANTS A	393	15.1	25.0	5.2
EnHANTS B	375	14.9	26.0	0.80
EnHANTS C	310	746	1610	176
EnHANTS D	326	97.5	256.5	24.8

Table 3.1: Summary statistics for the indoor photovoltaic irradiance traces from the EnHANTs dataset [99]. Irradiance is expressed in units of $\mu W/cm^2$.

energy is low [57]. The ZebraNet collar has enough battery energy capacity to support five days of operation without harvestable energy [9]. This approach is also taken by millimeter-scale systems. The M^3 platform has allocated enough energy storage to persist for just over two days in sleep mode [75]. The energy capacity allocated by all of these examples is arbitrarily chosen. In cases like the ZebraNet collar and the M^3 platform, physical size constraints likely drive the limits of energy capacity. However, no consideration or analysis is given regarding the consequence of the chosen energy capacity. Like energy income or the size of preallocated energy, the sizing of capacity of an energy buffer must be considered in a principled manner.

The proper capacity of a charge buffer for a given purpose exists on a continuum of how it averages or filters temporal power variation. We illustrate this continuum in Figure 3.5. In electrical circuits, the buffer provided by capacitors generally acts as a low pass filter of incoming and outgoing power, averaging out noise, ripples, and general supply variation. As a capacitor or other type of charge buffer increases in capacity, it smooths out larger and larger temporal power variations. For example, smaller capacitors are generally used for filtering out noise, slightly larger capacitors are used to stabilize and smooth out power supply ripple, and large capacitors can be used to provide a rush of instantaneous power to circuits when needed. The capacity of a supercapacitor energy buffer increases the timescale of averaging variance in power income to hours to days, depending on its capacitance. Rechargeable batteries offer even greater energy density and can smooth out energy income variation over the course of weeks, months, and seasons.

By making arbitrary choices regarding the size of energy capacity, system designers may be impacting the efficiency of their energy harvesting and ultimately the performance of their system. This section seeks to identify the relationships between energy capacity, sensor workload, and potential energy income. By understanding these relationships, future system designers can determine suitable energy capacity analytically instead of arbitrarily.

Defining an Energy Harvesting Income

Conversely, photovoltaic energy harvesting can differ significantly over the course of a day, a week, or even a year. To accurately consider the effects of capacity on energy capture, the full trace of income power must be considered, not simply an average. For a photovoltaic, this potential income power is dependent on the time series of energetic spectral irradiance, E_e, the photovoltaic efficiency η, and panel area A. The irradiance trace E_e is described by

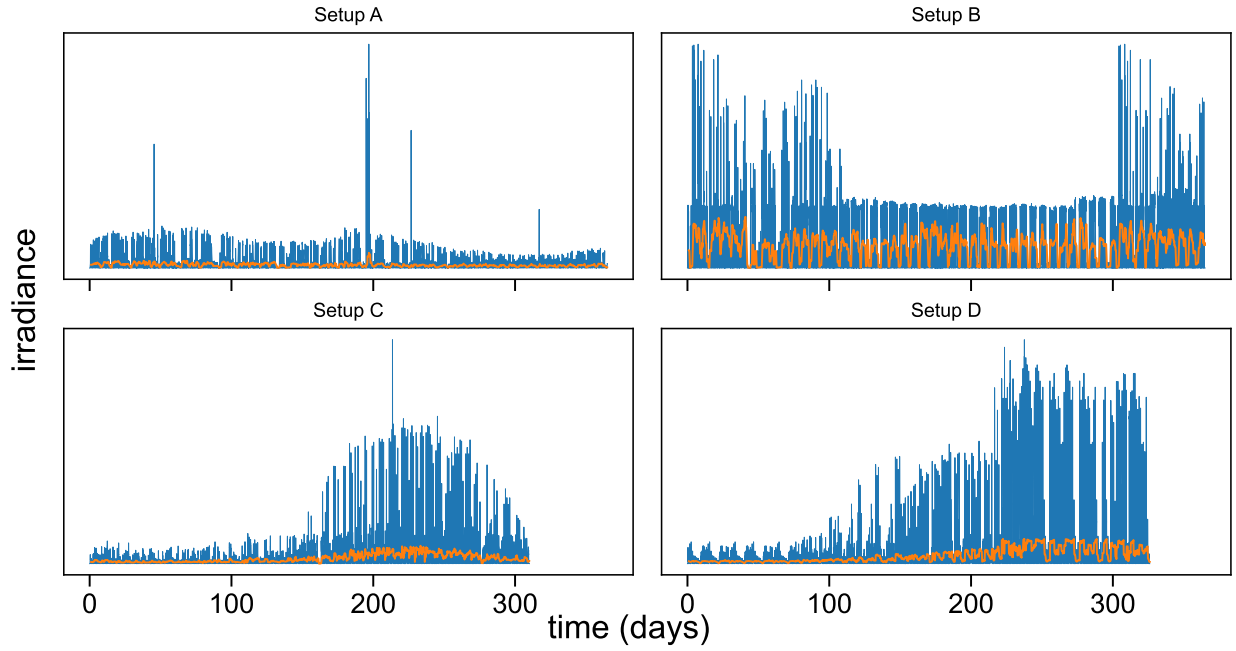

Figure 3.6: Each of the four irradiance traces from the EnHANTs dataset represent different lighting and environmental conditions. Here, each trace is min-max normalized to illustrate their individual temporal variance. Setup A, B, and D measure irradiance in a student's office. Setup A is located near a south facing window and receives some sunlight. Setup B is located at a bookshelf, and is largely occluded from sunlight. Setup D is located near a west facing window and receives significant sunlight during part of the year. Setup C measures irradiance in a conference room and gets significant sunlight from a north facing window. The blue line is the raw irradiance for each trace, while the orange line is a moving average of the irradiance with a day-length window.

Equation (3.19), and the resulting trace of power available to a harvesting system based on this irradiance is given by Equation (3.20).

$$E_e = (e_t)_{t=0}^{T} \tag{3.19}$$

$$\begin{aligned} p_t &= (\eta e_t A)_{e_t \in E_e} \\ P_h &= (p_t)_{t=0}^{T} \end{aligned} \tag{3.20}$$

The temporal variance of energy income directly impacts the amount of capacity required. For example, natural light varies significantly both diurnally as well as seasonally over the course of a year. Energy capacity must be large enough to capture energy during daytime and summer, so that it can supply it during nighttime and winter, respectively. It is important to utilize a realistic energy income trace for P_h to properly capture the effect of temporal

variance. This section utilizes the irradiance traces from the EnHANTS indoor irradiance dataset [99]. The EnHANTs dataset remains the most complete and extensive dataset for indoor light irradiance traces, capturing over a year of data in several environments. The four longest running traces are summarized in Table 3.1 and Figure 3.6. In Figure 3.6 all irradiance traces have been min-max normalized to show the characteristics of each trace. Artificial light accounts for the majority of income for Setup A and B. The irradiance for Setup B is very clearly almost all artificial light, as it consists of a pattern of five weekdays of artificial lighting followed by two weekend days with very little light when humans are not present. Conversely, traces C and D have significant income from sunlight that follows a seasonal pattern throughout the year. This indicates that A and B mostly follow a diurnal pattern, while C and D are simultaneously diurnal and seasonal. These traces exhibit different variances, which will impact the necessary capacity required to support a given workload. To isolate the impact of the variance of energy income, as well as the magnitude of the income, we synthesize income traces with arbitrary average power from mean-scaled EnHANTs traces. The traces are divided by their mean to achieve a mean of one, and then scaled by the multiplication of any average power. This allows us to generate synthetic traces with different intensities while maintaining the characteristics of each trace, resulting in an analysis with higher granularity.

Defining a Workload

A sensor workload can be represented as the time series of instantaneous powers over a time period T, as described by Equation (3.21).

$$P_w = (w_t)_{t=0}^{T} \tag{3.21}$$

The time series of required power for a system can be measured or it can be estimated if the sensor behavior is predictable and the power and duration of operations are known. A simplifying assumption is to reduce the workload trace to an average workload power, $\bar{W} = \frac{1}{T}\sum_{t=0}^{T} w_t$. For many workloads, a time scale of a few hours to a day is sufficient to calculate an average workload power assuming periodic behavior occurring within the time period, or for an application that reacts to events, the probability distribution of those events is well described by that time period.

Determining Capacity

To determine the ideal amount of energy capacity for a given sensor application, the relationship between energy capacity and energy income must be understood. As mentioned previously, energy capacity acts similar to a moving average over energy income. With sufficient capacity, energy is stored in times of excess harvestable energy, and provided in times of insufficient energy. At some point, a design can allocate enough capacity to effectively treat its energy income as an average, even if it is highly variable. While the operation is *similar* to a moving

average, it is not exactly, as real energy storage has limits when full or empty. This section seeks to explore this relationship with the goal of developing guidance and methods for making design decisions regarding energy capacity.

The ideal energy capacity for a given application depends on the expected harvesting environment for the application, as well as the intended sensor workload of the application. Energy harvesting income and a sensor's workload can both be represented by a time sequence of instantaneous power across a period T. The sequences for workload and energy income are defined by Equations (3.20) and (3.21), respectively. Given these traces of workload and income power, the trace of cumulative energy stored by an energy buffer ($Y = (y_t)_{t=0}^{T}$) is defined by Equation (3.22). The initial condition y_0 assumes the buffer starts empty, which is realistic for capacitor-sized energy buffers, but a very conservative assumption for larger buffers.

$$
\begin{aligned}
y_0 &= \begin{cases} 0 & \text{if } p_t - w_t < 0 \\ (p_t - w_t)\Delta t & \text{if } p_t - w_t > 0 \\ E_{capacity} & \text{if } (p_t - w_t)\Delta t > E_{capacity} \end{cases} \\
y_t &= \begin{cases} 0 & \text{if } y_{t-1} + (p_t - w_t)\Delta t < 0 \\ y_{t-1} + (p_t - w_t)\Delta t & \text{if } y_{t-1} + (p_t - w_t)\Delta t > 0 \\ E_{capacity} & \text{if } y_{t-1} + (p_t - w_t)\Delta t > E_{capacity} \end{cases}
\end{aligned}
\tag{3.22}
$$

This relation is essentially a capped sum over the difference of instantaneous income and workload powers. The amount of energy stored in the buffer is limited to values between zero and the maximum capacity of the buffer. We assume the energy buffer starts empty. While systems with large energy capacities are normally deployed with fully or partially charged buffers, for later analysis using these relations, starting empty will produce a more accurate estimation of required energy capacity over long periods of time. For simplicity, the Δt between each element of P_h and P_w can be assumed to be one second. This simplifies future calculations for energy. Also, as mentioned previously, w_t is replaced by the average of the workload sequence $\bar{P_w}$.

The key metric of interest given the trace of stored energy, $Y = (y_t)_{t=0}^{T}$, is the amount of available energy income that was captured by the energy buffer. Given that the average income power is equal to the average workload power for this analysis, capturing all of the available energy also indicates enough energy was captured to run the workload continuously. Insufficient capacity will fill up prematurely and be unable to capture any additional energy. Capacity, if too small, can severely limit the total energy captured by an energy harvesting system. The trace of instantaneous energy captured by an energy buffer is represented by the sequence $Z = (z_t)_{t=0}^{T}$, and each element z_t is defined by Equation (3.23). The total energy captured by an energy harvesting system over a time period T is represented by the sum of the elements in Z, defined in Equation (3.24). At any given time, the energy captured depends on whether the energy storage is full or not. If the storage is not full, the entire $p_t \delta t$ is captured. If nearly full, only a part of the available energy is captured.

$$z_t = \begin{cases} p_t \Delta t & \text{if } y_t + (p_t - w_t)\Delta t \leq E_{capacity} \\ E_{capacity} - y_t & \text{if } y_t + (p_t - w_t)\Delta t > E_{capacity} \end{cases} \tag{3.23}$$

$$E_{captured} = \sum_{t=0}^{T} z_t \tag{3.24}$$

Due to the conditional and iterative nature of the elements of the capacity sequence (Y) and captured sequence (Z), there is no way to directly solve for $E_{capacity}$. Instead, we can use Equations (3.22) and (3.23) to perform a parameter sweep of average workload power, income power magnitude, and capacity. To simplify the sweep, we assume that the sensor workload has been tailored to match the expected income, such that the average workload power is equivalent to the average energy income ($\bar{P}_w = \bar{P}_h$). This is a conservative assumption, as it is usually beneficial to tailor the workload to require less average power than its income. This provides a margin to ensure performance even if energy income is worse than expected. However, for this analysis a conservative estimate is ideal. We explore the impact of considering income margins in Section 3.3.

To determine a minimum sufficient capacity, we are essentially sweeping two variables: capacity and average power. One of the sweeps of capacity is illustrated in Figure 3.7. This figure has a fixed income and workload power of 50 µW, but capacity is varied from capacity offered by a ceramic capacitor to that of a small rechargeable battery. As capacity available to the system increases, the amount of harvestable energy also increases. There are two instances where the percentage of energy captured increases rapidly. The first occurs around a capacity of 1 mWh, which represents sufficient capacity to overcome diurnal variation. All trace types experience a substantial increase at this point, as they all exhibit diurnal variation. The second rapid increase, which occurs around 100 mWh, represents the sufficient capacity to overcome seasonal variation, which Setups C and D exhibit. Setup A also has some seasonal variation, but to a lesser extent. Setup B does not exhibit the same second rapid increase as it does not vary substantially between seasons. For this specific income power and workload, Setup C and D require on the order of 100 mWh to capture all the available energy to power a 50 µW workload continuously. Setup A and B require on the order of 20-40 mWh.

The amount of capacity that is sufficient to capture 100% varies not only on the income variation but also on the relative magnitude of the workload and income power. As the sensor's power requirements and income increase, its capacity requirement to achieve the same performance also increases. The parameter sweep resulting from varying both capacity and power results in a multitude of curves like those shown in Figure 3.7. For each of these curves, we can identify the points at which capacity is sufficient for 100% energy capture for a given power. Figure 3.8 shows the minimum capacity required to capture all of the available harvestable energy versus the magnitude of income and workload power. This is also the necessary amount of energy required (when averaged over the period of the trace) to power the workload continuously, because the average workload power is equal to the average harvestable power. From this analysis, we find that the relationship between minimum sufficient capacity

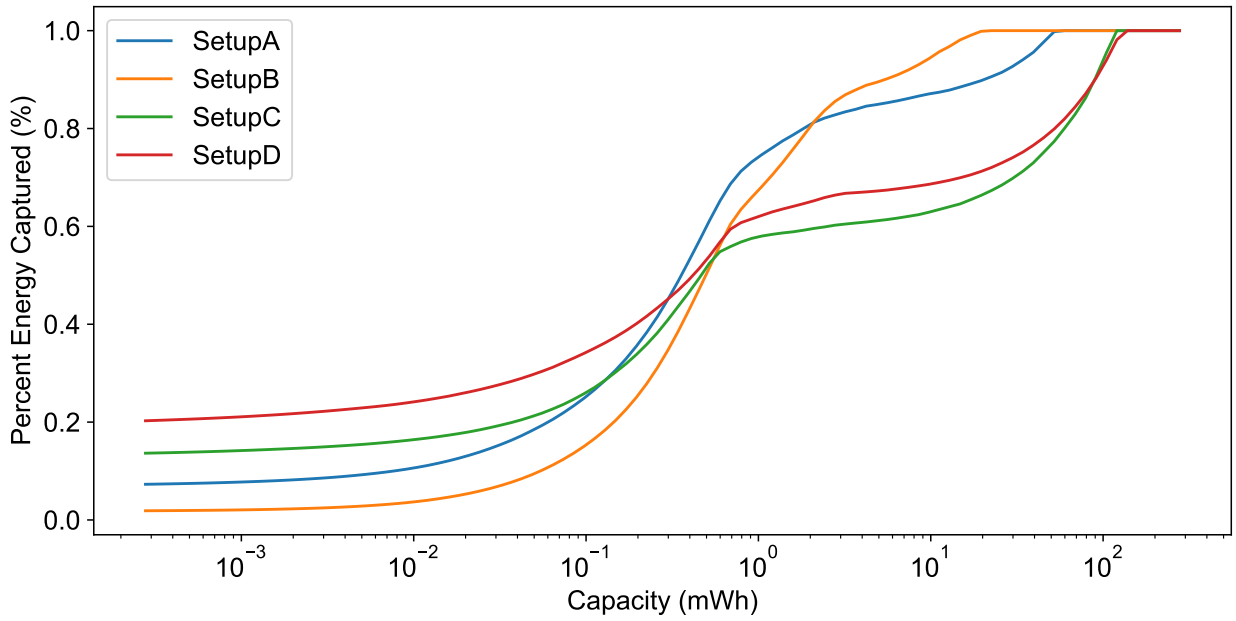

Figure 3.7: The percentage of energy captured to satisfy a workload vs a sweep of capacity. This figure assumes a 50 µW average income and workload. Note the x-axis log scale. Capacity ranges from the order of energy capacity offered by a small 100 µF capacitor, to the capacity offered by a small 100 mAh battery.

and the magnitude of workload and income power is linear, assuming the same energy income variability. Traces synthesized from Setup C and D have nearly identical minimum capacity requirements across a sweep of workload power. These traces also require the more energy capacity than traces A and B as they have more seasonal variability. Out of all the income traces, those synthesized from Setup B require an order of magnitude less energy capacity compared to Setup C and D. Traces synthesized from Setup A require amounts of energy capacity in between that of B and C or D. Energy harvesting applications that can expect an amount of natural light should size their rechargeable capacity according to either the Setup C or D lines. Such an application would have to scale its capacity by at least 2.3×10^3 times the expected magnitude of the income and workload power. For indoor applications that expect to harvest the majority of light from artificial sources, capacity can be sized according to the Setup A and B lines, to be at least 4×10^2 to 1×10^3 times the magnitude of income and workload power. This result is a significantly more reasoned heuristic for determining the minimum sufficient rechargeable energy capacity than other methods, like sizing capacity to support a single atomic operation in a workload, or sizing to support arbitrary amounts of operating time.

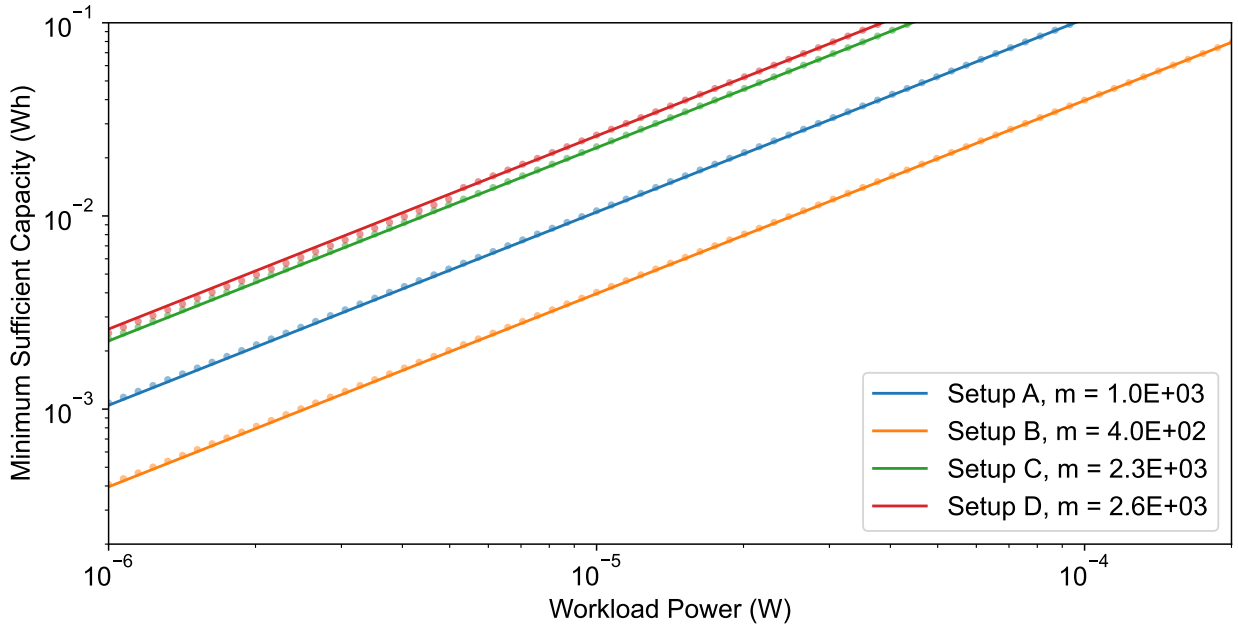

Figure 3.8: The minimum sufficient capacity to support an average workload power, assuming an identical average income power. Note the x- and y-axis log scale. The average workload and income power are set equal; however the variability of income power is determined by the synthetic EnHANTs traces. The minimum sufficient capacity follows a linear trend with increasing income and workload power. Even though each line appears parallel, lines that are higher up on the y-axis actually have a larger slope due to the log-log scale.

The Impact of Income Safety Margin on Sufficient Capacity

Throughout this section, we have assumed that the average workload power was equivalent to the average income power. This is a very conservative assumption, and engineers rarely design systems that operate on the edge of feasibility. Usually, systems are over-provisioned to allow for a margin of error, usually between 20-50%. For energy harvesting, this would represent altering a workload to require less than the expected available harvestable power, increasing the size of the harvester, or deploying in a location with more ambient harvestable energy. In Figure 3.9, we model the impact of an income safety margin on the required capacity to capture sufficient energy. We utilize the Setup D EnHANTs trace as it represents the worst case seasonal light distribution out of all of the traces. Figure 3.9 presents required capacity for multiple sweeps of workload power under different safety margins. The different sweeps consider different safety margins of income power, from no margin to a 300% margin. No margin indicates that the average workload power is equal to the average income power, while an 300% margin indicates the income power is 4× the amount required by the workload, on average[1]. As expected, when a workload requires less energy than is available to harvest,

[1]A safety margin is calculated as a percent difference, $100 \times \frac{\text{Income}-\text{Workload}}{\text{Workload}}$

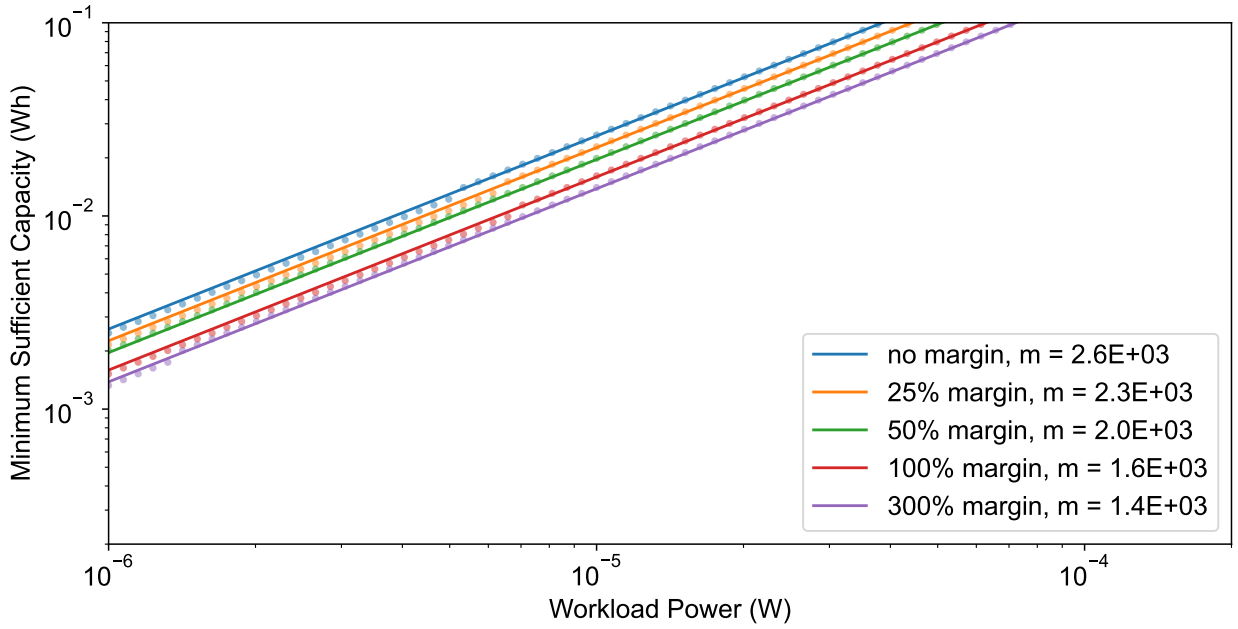

Figure 3.9: The impact of workload margin on the minimum sufficient capacity. Traces synthesized from Setup D are considered, and are scaled by an income power. Six different workload margins are considered. The zero margin corresponds to the workload power being equal to the average income power. The 0.2 margin represents the workload power requiring 80% the power provided by the average income, for example. As the margin increases, the ratio of the average workload power to the income power decreases. With a higher margin there is less energy that must be captured to satisfy the workload and the minimum sufficient capacity decreases.

it requires less energy capacity to harvest that energy. With a 25% margin, where the income provides 1.2× the average required workload power, a design that can expect an energy income with a similar distribution to the Setup D trace will require a capacity scaled by 2.3×10^3 multiplied by the average workload power.

Limitations

The analytical model presented in Equation (3.22) is simple and assumes a perfect energy buffer. In reality, energy storage technologies are imperfect. They feature less than perfect efficiency; they leak and have parasitic internal impedance. The efficiency of an energy harvesting front end is also more complex than presented in this section. Beyond the efficiency of a harvester, the DC-DC regulation of a harvesting IC does not provide 100% efficiency. Any further system voltage regulation has similar limits to efficiency. All of these imperfections result in energy loss and prevent a perfect capture of all available harvestable energy. Workload is also assumed to be a constant average power. In reality, a workload can be variable depending on

an application's reporting requirements or the probability distribution of detectable events. This section makes the argument that the variability of a workload is sufficiently less than the variability of income, to the point that workload can be treated as an average in comparison. This assumption is reasonable for many workloads that do not exhibit long term variation but does not apply to all possible sensor workloads. Additionally, this analysis did not consider the effect of a backup non-rechargeable energy store. As seen in Section 3.1, the addition of a non-rechargeable backup can result in significant increases in the lifetime of a system while maintaining operation with high availability.

3.4 Summary

In this chapter, we present analytical methods for considering the design space of an energy harvesting wireless sensor power supply. We present heuristics for determining the proper sizing of non-rechargeable battery storage and photovoltaic harvesters, and we provide a framework for considering the use of either option for powering sensors, or a hybrid solution of both preallocation and harvesting. For energy harvesting sensors, they may require batteryless and intermittent techniques to ensure feasibility and forward progress. The requirement for intermittent techniques is solely dependent on energy income and rechargeable energy storage, and they are avoidable by properly sizing a harvester and energy storage. We also present a simple model for energy capacity and analysis that indicates that the size of rechargeable capacity in energy harvesting systems directly impacts the amount of energy that is actually captured by a system. The end result of the analysis of this chapter is a reasoned heuristic for determining appropriate rechargeable capacity sizing for a system based on its income and workload that is superior to arbitrary methods used by other platform designs.

The next chapter seeks to expand upon the iterative model of capacity presented in Section 3.3 by developing an energy simulation for wireless sensors. It will seek to address the three previously mentioned limitations by modelling energy income and storage more realistically, modelling sensor workload more dynamically, and consider the impact of non-rechargeable backup energy.

Chapter 4

A Simulation-based Exploration of Capacity

In the previous chapter, we presented a simple model for energy capacity and explored the effect of capacity on the energy captured by a system. We utilized synthesized energy harvesting traces from an existing irradiance trace dataset, however we simplified the analysis by assuming a static average workload power. This chapter will explore expanding our model to an energy state simulator for wireless sensors. This simulator will utilize several different dynamic workloads based on benchmarks of real hardware, as well as considering a wireless sensor state machine that better models the behavior of a real sensor. The rest of the chapter is an analysis of the performance of these simulated wireless sensor workloads under measured harvesting conditions.

4.1 Upgrading Our Model

In Section 3.3, we argued that a sensor workload, which is often periodic or event driven, does not vary significantly beyond a few hours to a day of operation when compared to the variability of energy harvesting income. To verify this assumption and expand our analysis of capacity, we seek to to explore the dynamic effects that energy capacity and backup storage have on sensor performance in the face of workload variability. We use representative energy harvesting traces, measurements of real hardware, and synthesized dynamic workloads in a new model based on a simple wireless sensor state machine. We use this model to simulate the behavior of energy harvesting sensors. From the input traces of workload and energy harvesting income, our model produces estimates of sensor energy utilization, availability, responsiveness, and lifetime.

Irradiance Trace	Total Days	Average Power ($\mu W/cm^2$)	90^{th} Percentile Daily Power ($\mu W/cm^2$)	10^{th} Percentile Daily Power ($\mu W/cm^2$)
EnHANTs A	394	15.1	25.0	5.2
EnHANTs D	311	97.4	256.5	24.8

(a) Indoor photovoltaic irradiance traces

Workload Class	Energy per Event (uJ)	Average Period	Average Power (µW) [a]
Periodic	586	10 s	58.6
		30 s	24.5
		60 s	14.7
		120 s	9.8
Reactive	86	3.4 s [b]	25.3
		6.8 s [b]	17.6
		13.6 s [b]	11.3
Long-Running	93,300	2 weeks [c]	5.1

(b) Representative workloads

[a] Average power includes an average 5 µW idle power, measured in Section 6.1.
[b] Event times are based on a Poisson distribution for each hour of the day and drawn every second. The distribution is parameterized by collected entryway data then scaled.
[c] Event time is based on a uniform distribution and drawn every second.

Table 4.1: Representative harvesting conditions and workloads. To evaluate different energy harvesting storage techniques, we define a set of energy harvesting conditions and workloads that are representative of common sensing applications. We choose two real irradiance traces with different magnitudes and distributions of available energy. These traces are summarized in Table 4.1a. We define three workloads: periodic, reactive, and long-running, and we characterize those workloads for different event frequencies. The energy used for each event is measured on our reference hardware described in Section 6.1. Statistics for the three workloads are described in Table 4.1b.

Energy Harvesting Income

As in Section 3.3, we utilize the EnHANTs irradiance trace dataset as energy income for our simulation. Unlike our previous analysis, we do not use synthesized traces, and instead utilize the traces directly. We choose to narrow our focus to two of the traces that we think most accurately represent indoor lighting conditions: Setup A and Setup D. As shown in Figure 3.6, Setup A represents a location that receives some sunlight but is mostly lit by artificial sources while Setup D represents a location near a window that receives significant light during part of the year. Beyond Figure 3.6, these traces are summarized in Table 4.1. On average, Setup A provides an average of 15.1 $\mu W/cm^2$ which is on the lower end of expected irradiance in indoor environments. Setup D represents the higher end of indoor irradiance with an average 97.4 $\mu W/cm^2$.

Representative Hardware

We limit our analysis to the effects of capacity, independent of the differences of energy intensity or efficiency in device component selection. To do so, we define an example photovoltaic energy harvesting sensor platform that utilizes currently available commercial components. These components are listed in Table 4.2. We choose new components in an attempt to better represent prior energy storage designs and give them the benefit of the improvements that have occurred in recent years. We take benchmark measurements of various tasks performed by this representative platform, such as the amount of time and energy required to sample a sensor or send a Bluetooth Low Energy (BLE) packet. These benchmarks are used to generate energy utilization metrics for our representative workloads shown in Table 4.1. The physical size of the solar panel used by this sensor is assumed to be 10.9 cm^2 and the volume of the sensor node is similar to prior work like the Hamilton sensor [103].

Representative Workloads

We find that sensing workloads generally fall into three categories: (i) periodic sense-and-send [8], (ii) reactive event detection [11, 63], and (iii) infrequent, long-running, high-energy events [104]. All of these workloads are characterized by long periods in which a sensor is inactive, punctuated by active events, which may be periodic or drawn from a probability distribution. These active events may correspond to collecting a measurement from a sensor, performing some computation, and transmitting the results. We choose a representative workload for each of the three categories to use in our model. We characterize our "periodic sense-and-send" workload as periodically sampling a light and three channel color sensor and sending a BLE advertisement containing the data. Our "reactive workload" consists of sending a BLE advertisement upon motion detection of the main entrance of a university building, and we linearly scale the frequency of these events to create synthetic traces with varying amounts of usage. We treat these workloads as atomic. When a measurement is taken or a packet is sent in response to a periodic schedule or a detected event, the energy

Component	Function	Active Power	Idle Power
Nordic NRF52840	Processor	56 µA/MHz	940 nA [a]
	Radio	5.2 mA @ 0 dbm	— [a]
Ambiq AB1815-T3	Real time clock	55 nA	N/A [b]
ST Micro LIS2DW12	Accelerometer	1 uA @ 12.5 Hz	50 nA
Maxim MAX44009	Light sensor	650 nA	N/A [b]
Intersil ISL29125	Color sensor	56 µA	500 nA
Silicon Labs SI7021	Humidty sensor	1.5 µA @ 1 Hz	60 nA
TE Connectivity MS5637	Pressure sensor	0.6 - 5 µA @ 1 Hz	10 nA
Panasonic EKMB11011	PIR Occupancy	100 µA	1 uA

[a] Sleep current for both processor and radio. [b] No shutdown or idle mode.

Table 4.2: The components used by our representative hardware. Benchmarks of the processor, radio, and sensors presented here are used to establish our representative workloads used by our wireless sensor energy simulation. These components are among the lowest power options available, and are even 2-4x lower power than those used on relatively recent systems [7, 18, 31].

spent on those operations must be spent instantaneously and atomically. This means a sensor must currently have sufficient energy stored to complete the operation, otherwise the event is considered "missed" and counted against the availability of the system. Finally, our "infrequent expensive" workload is a contiguous task that is representative of an over-the-air firmware update, which is randomly executed with an average occurrence rate of once every two weeks. We optimistically assume this long running task can be interrupted and resumed at any point during execution and that any checkpointing does not incur any overhead.

4.2 An Energy Model for Wireless Sensors

We use the previously discussed indoor irradiance traces, generalized workloads, and hardware characterizations to model the behavior of sensors using different types and sizes of energy storage. We develop an open source[1] iterative simulator that allows parameterization of various system characteristics, including regulator efficiency, solar harvester size and efficiency, energy storage capacity, leakage, equivalent series resistance (ESR), and charge-discharge efficiency. A subset of these parameters are summarized in Table 4.3.

The simulation of our model operates as a second-by-second calculation of the energy entering and exiting a device, similar to the model presented in Section 3.3. At every step, the simulation calculates the net energy gain or loss of the system based on its current state, the instantaneous harvestable energy, and available stored energy. Occasionally, the model performs a workload event based on either a periodic schedule (in the case of a sense-and-send

[1] https://github.com/lab11/permamote/tree/master/simulator

Config Type	Parameter	Description
Device	`operating_voltage`	Output voltage of the power subsystem
	`boost_efficiency`	Efficiency of the boost converter
	`frontend_efficiency`	Efficiency of the harvesting frontend
Secondary	`capacity`	Capacity of secondary in joules or mAh
	`esr`	Equivalent series resistance in ohms
	`leakage_constant`	Factor for capacity dependent leakage
	`{max, min}_hyst`	Secondary capacity upper/lower hysteresis
Primary	`capacity`	Capacity of primary in mAh
	`leakage_percent`	Percent capacity leakage per year
Harvester (Solar)	`area`	Area of solar harvester in cm^2
	`efficiency`	Efficiency of solar panel

Table 4.3: Simulation configuration parameters. A subset of available configuration options for the sensor energy simulation. Simulated sensors can be configured to use secondary storage and an energy harvester, a primary-cell, or both. A secondary-cell can be configured with a hysteresis, with a lower bound set to `min_hyst` and an upper bound of `max_hyst`.

workload) or from a random distribution (reactive event detection or a high-energy event). For our modeling, workload schedules are generated from values listed in Table 4.1. This simulation is performed for the entirety of an input irradiance trace, which constitutes about a year of data. During a simulation, metrics such as energy utilization, the fraction of completed events versus expected events, and event time to completion are collected, and if applicable, the primary-cell lifetime is estimated from a calculation of the net difference in energy in the system from the start to the end of the simulation, if negative.

During simulation, modeled devices can be online or offline and idle or performing work. These states are shown in Figure 4.1. A device's state transitions from top to bottom of this figure and vice versa depending on the energy state of the secondary storage. If the secondary-cell energy state drops below `min_hyst`, the state of the system moves to the upper half (offline) of this diagram. The state of the system moves downward (online) if the state of charge of the secondary reaches the `max_hyst` limit. Secondary charging hysteresis limits are defined by parameters described in Table 4.3. A device's state can also move to the right or left of the state machine depending on whether a workload event is scheduled, or the prescribed workload has been completed. A new workload event is counted as failed if the device is not in the Online Idle state when it begins. In the case of the "atomic" sense-and-send and reactive workloads, the modeled device will only begin to perform an expected workload event if it has enough energy to perform the event in entirety. If the workload is not atomic, the device will only begin scheduled work if it has enough energy to make the configured minimum amount of progress. We make an assumption that the duration of atomic events are less than the one second simulation step. We also assume that a modeled sensor has perfect, zero-energy progress latching and can go to sleep at any point after an active event. If there

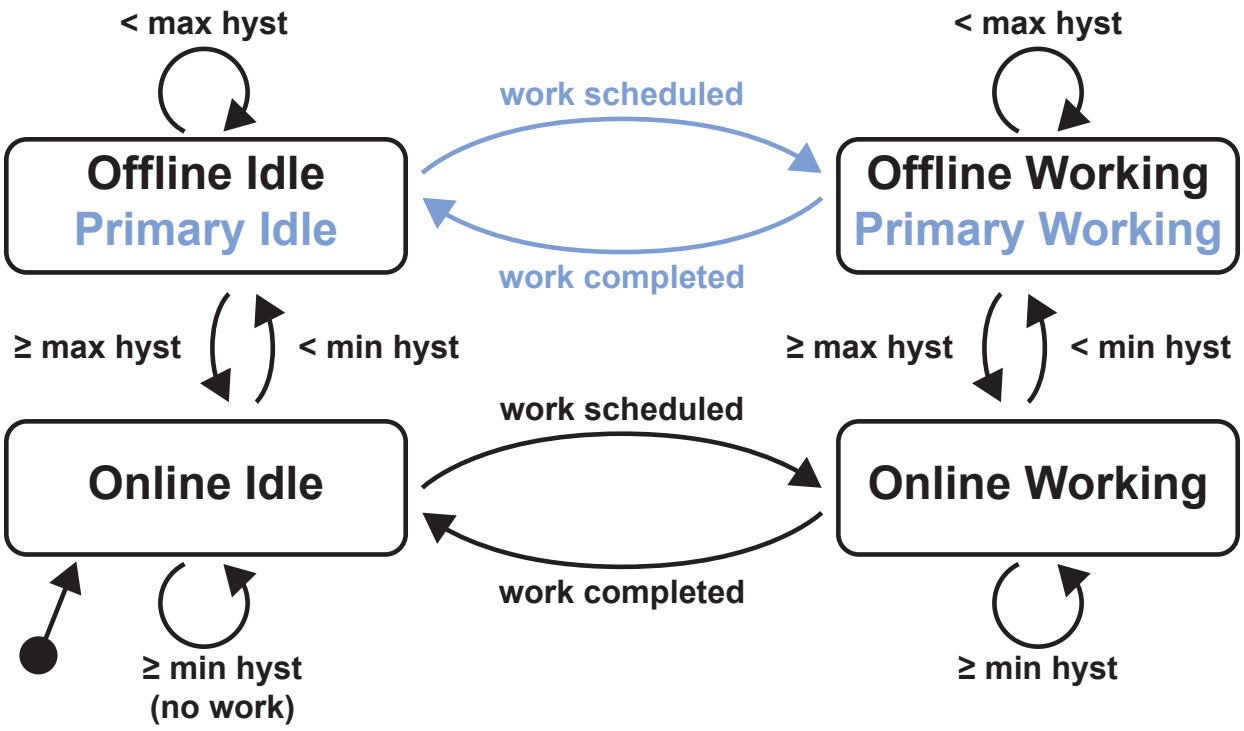

Figure 4.1: Model state machine. A modeled device can be in one of four states: Offline Idle, Online Idle, Online Working, and Offline Working. When a device is Offline Idle, it has run out of energy and is off. If a device is Online Idle, it is on and in deep sleep, ready to perform work if triggered. If triggered, a device moves to Online Working, where it performs a portion of a work event. If a workload is atomic, workload events *must* be completed in one Online Working step, without any transitions to an offline state. Offline Working means that while working on a non-atomic task, the device ran out of energy, checkpointed, and is waiting to harvest more and resume its task. For devices configured with a primary-cell, Offline Idle and Offline Working become Primary Idle and Primary Working respectively. In these states, outgoing energy is charged against the primary-cell and the device remains online and able to perform work for the life of the primary-cell.

is energy remaining after performing a task, the modeled sensor will attempt to spend the rest of the simulation period in the Online Idle state.

If the simulated device is configured with a backup primary-cell, offline states transform to "primary" states in which the device remains on and able to perform work, but charges energy usage to the primary storage instead of the secondary. During these periods, the secondary cell continues charging from harvested energy. Upon reaching the upper charging hysteresis limit, the device returns to an online state, using energy stored in the secondary. If the primary storage is depleted, the simulation ends early.

4.3 Capacity Increases Energy Capture and Availability

Using our simulator, we can confirm that the conclusions from Section 3.3 regarding the relationship between energy capacity and energy capture. An increase in sensor energy capacity results in an increase in the capture of energy during periods of abundant harvestable energy. This additional energy has the effect of increasing the availability of the sensor, as the buffer serves to average out times of energy insufficiency and excess. Higher The next few subsections explore the effect of capacity on energy utilization and sensor availability.

Ambient Energy Utilization

Ambient energy is underutilized when it is not used to support the specified sensing application. This may happen for two reasons: 1) the secondary energy storage is full but energy is still available for harvesting, and 2) the sensor performs work based on its energy state rather than its application goals. The first scenario is common for energy harvesting systems presented in prior work that charge up a capacitor and wait for an event before sensing and sending, failing to capture all energy that may have been harvested while their capacitor was full [11, 63]. For an example of the second scenario, consider opportunistic batteryless systems that transmit a packet every time their energy storage capacitor is full rather than saving this energy for use during periods of lower harvesting potential [7, 18]. Another example of this are sensors in which the harvesting rate is proportional to the sensed phenomena [55].

To explore ambient energy utilization as a function of storage size, we use our simulation of the charge-discharge patterns of idealized energy stores under the harvesting conditions and workloads described in Section 4.1. We perform multiple simulation runs while sweeping energy storage capacity and measuring the amount of harvestable energy captured and used by the intended workload. We desire to isolate the effect of capacity from the type of storage, so the idealized energy storage is assumed to have no leakage or ESR. This modeling is primarily accounting for our first classification of wasted energy, since our workload definitions *do not* perform tasks in response to available energy; instead they attempt to maximize the success rates for the specified sensing tasks. The results of this modeling are shown in Figure 4.2. The x-axis of Figure 4.2 is split by energy capacities possible with capacitors, supercapacitors, and batteries. The upper extents of capacity for capacitors represents ten

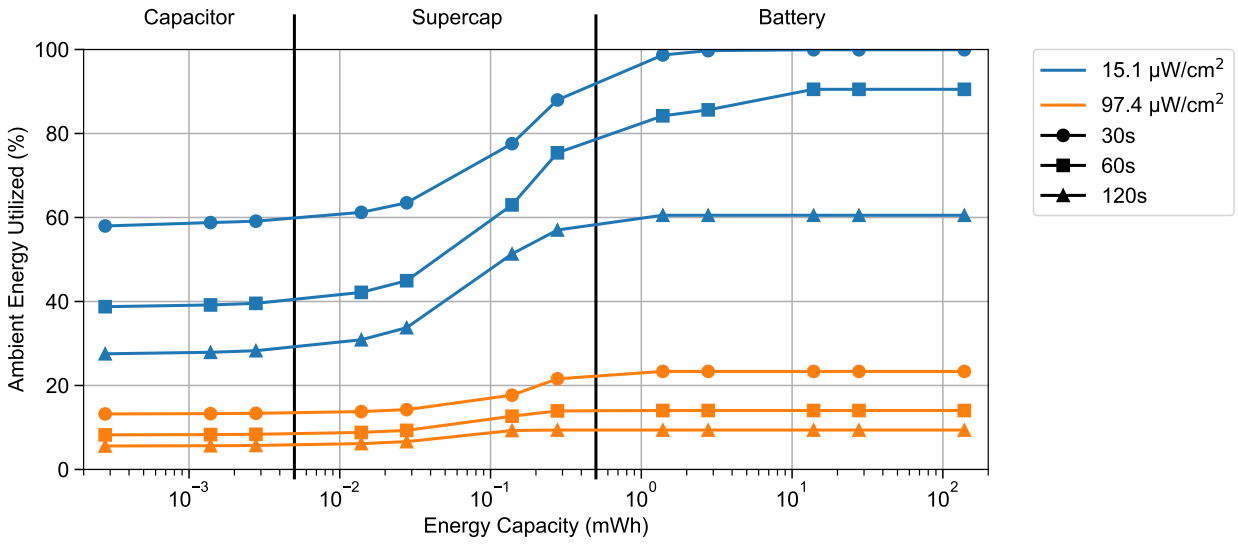

(a) Periodic application

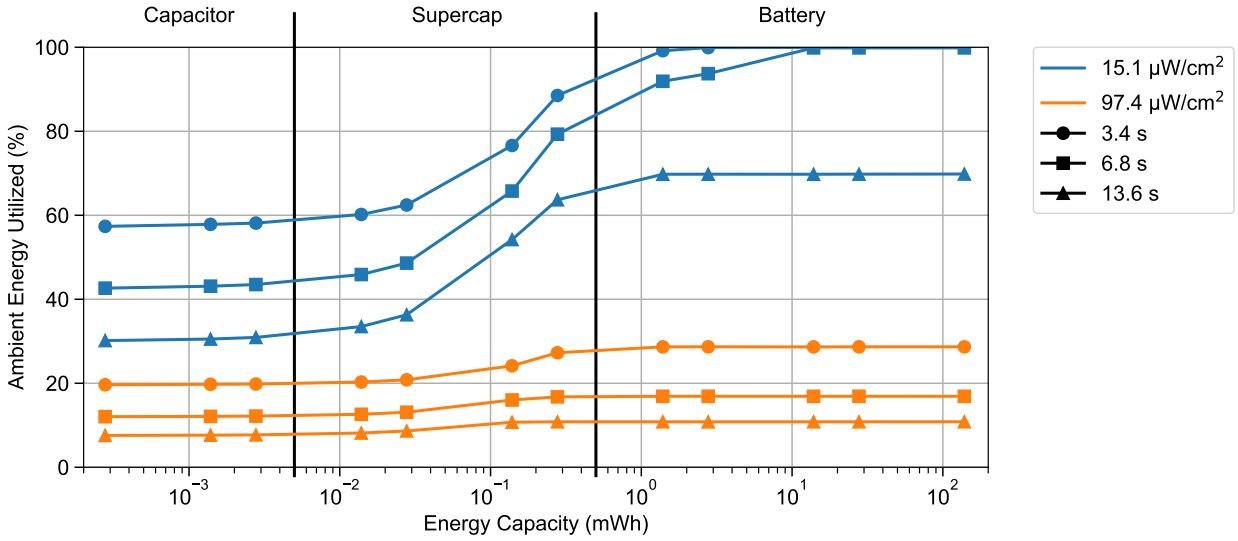

(b) Reactive application

Figure 4.2: Ambient energy utilization as a function of idealized secondary storage capacity for different harvesting scenarios and workloads. The harvesting scenarios and workloads are described in Table 4.1. Figure 4.2a represents the energy utilized by a periodic sense and send application, while Figure 4.2b is the energy utilized by a event-driven application. Despite these two workloads exhibiting different event distributions and variance, the overall energy utilization follows the same trend with energy capacity. As energy storage increases, the harvestable energy used in the application also increases. Some scenarios, such as the periodic 30 s, 15.1 µW/cm^2 case, reach 100% utilization at sufficient secondary capacities indicating that all of the available energy is captured and may not be enough to meet the application's requirements. Generally, for both workloads and irradiance traces, from the smallest to largest capacity simulated, we see a 1.4-2.3x increase in utilized energy.

100 µF tantalum capacitors [105]. For supercapacitors, the extent is represented by one large 220 mF supercapacitor [106]. These delineations are not hard partitions, as capacitors and supercapacitors have a wide range of capacities and the capacities offered can overlap in some configurations.

Regardless of the amount of available energy or the workload considered, energy utilization increases as buffer energy capacity increases. For scenarios with low harvesting potential and high power workloads, a sensor with at least 1–10 mWh of storage can accomplish 100% utilization. In cases of high harvesting potential and low power workloads, utilization often stops increasing before reaching 100%. This can be attributed to a small fraction of the available energy being sufficient to fully support the sensing task. For the low light, 15 µW/cm^2, environment, a few of the workloads are eventually able to harvest 100% of the available light as capacity increases. This indicates that the sensor performance is bounded by available harvestable energy, which is not enough to sustain the workload. Conversely, in the high light 97 µW/cm^2 environment, regardless of the workload, the simulated sensors are not harvesting all or even the majority of the available energy. This indicates they are harvesting as much as needed to power their workload.

In most cases, energy utilization for a workload is maximized with the energy capacity offered by small batteries or very large supercapacitors (>1 F). A system configured with ceramic or tantalum capacitors will be unable to fully utilize available harvestable energy, resulting not only in unreliable energy, but also less energy. A capacitor-based batteryless system will simply harvest less overall energy than a supercapacitor or battery counterpart.

Workload Availability

The amount of captured energy directly impacts the capability of an energy harvesting sensor. With more captured energy, a system is better able to complete its workload consistently and with high reliability. We utilize our simulation to also measure the ability of varying energy capacity sizes to meet the simulated sensing tasks under various harvesting intensities. The results are presented in Figure 4.3. Similar to the results of Figure 4.3, we see marked increases in performance when the energy capacity is increased to 1–10 mWh of capacity. Simulations experience 100% availability for all but the most energy constrained scenarios with high power workloads and ≥2 mWh of energy storage. Figure 4.4 is another perspective of availability, and examines sweeping the period of the periodic workload with four decades of energy capacity. These simulated results indicate that even with gratuitous energy harvesting potentials and infrequent workloads, sensors with energy capacity on the order of that offered by capacitors or supercapacitors will experience low availability compared to a sensor designed with sufficient energy capacity. This is because they are unable to capture and utilize enough harvestable energy to allow continuous sensing through periods with little to no harvestable energy, like nighttime or winter.

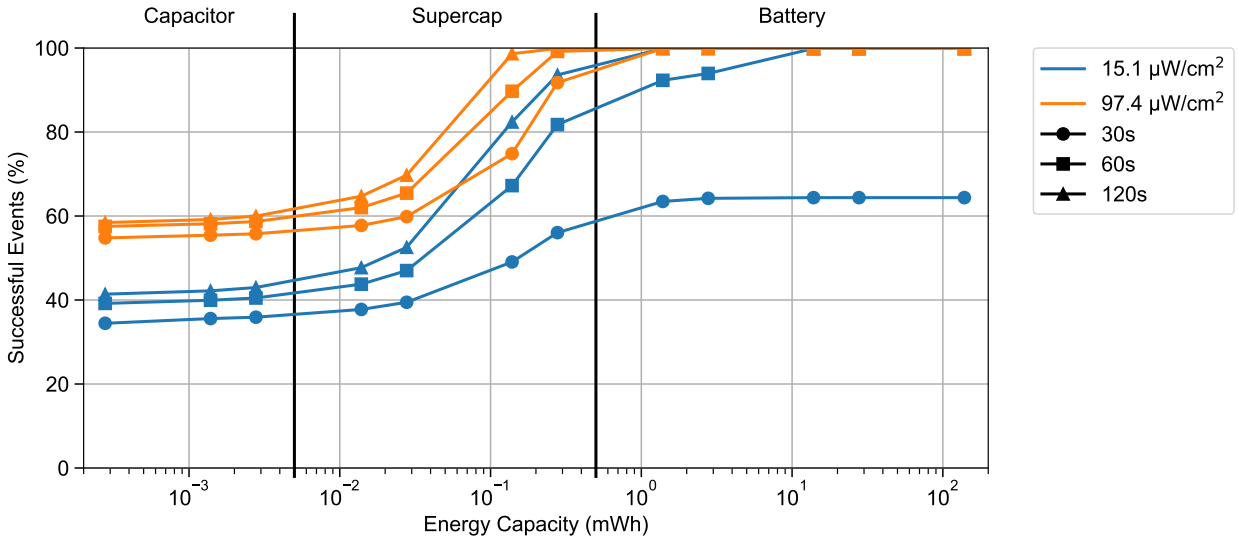

(a) Periodic application

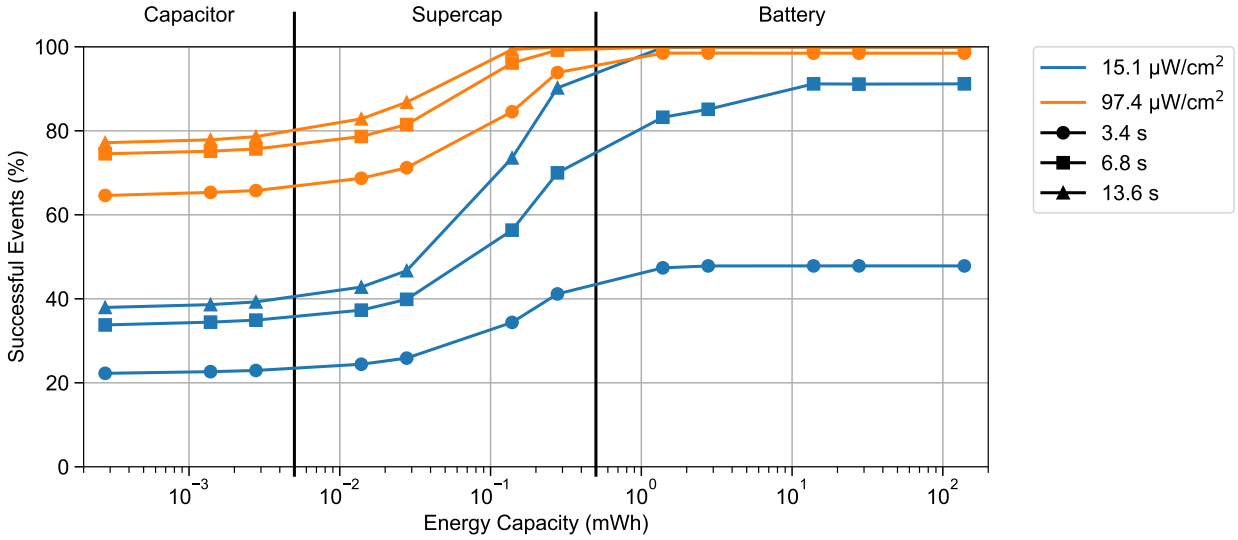

(b) Reactive application

Figure 4.3: Workload availability for different harvesting scenarios, workloads, and idealized secondary storage sizes. We define availability as the percentage of successfully completed events. If a periodic or random event occurs and the sensor does not have sufficient energy to complete the event, that event is not completed and is considered missed. As expected, workload availability follows a similar trend as energy utilization, improving with increased secondary energy storage and energy capture. For both periodic and reactive workloads, from the smallest to largest capacity simulated, we see a 1.4-2.7x improvement in availability. In cases where average harvester power is sufficient to power the workload completely, the sensor achieves 100% availability on the workload when configured with sufficient capacity.

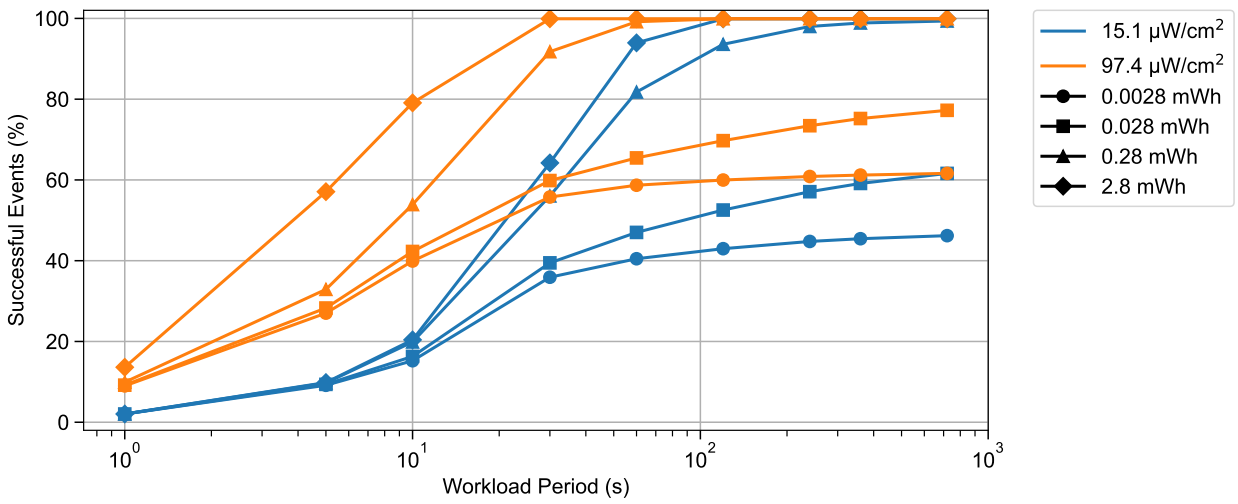

Figure 4.4: The performance of workloads with different periodicity with four decades of energy capacity. We investigate the period at which different secondary storage sizes meet a specific availability, showing that even with infrequent periodic workloads, small amounts of secondary storage have low availability while larger secondary stores approach 100% availability.

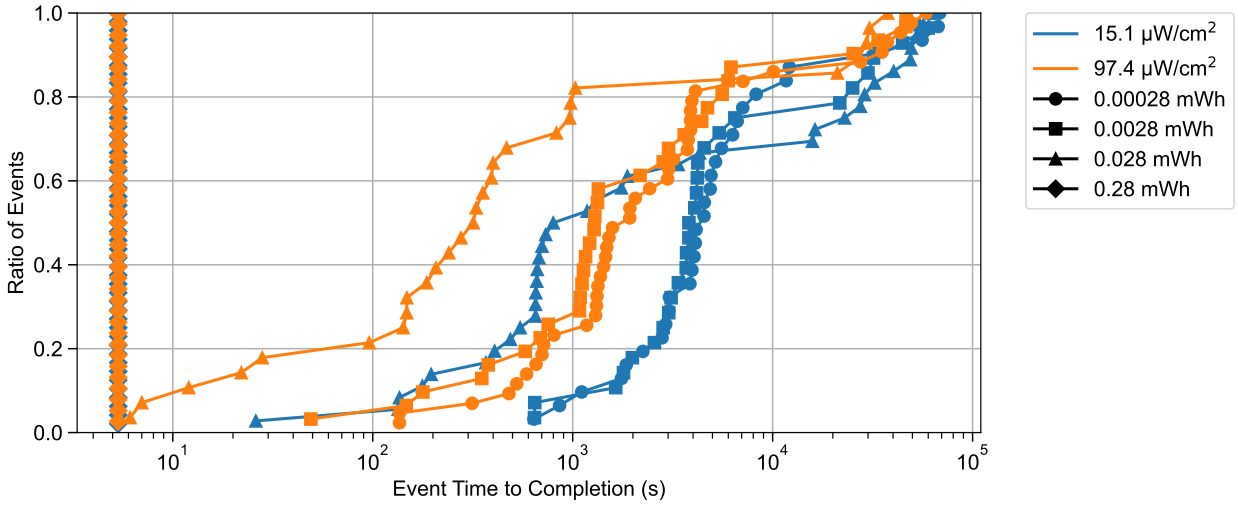

Figure 4.5: A CDF of time to completion of a long-running, high energy event. In this workload, events are not atomic, and can be paused and resumed based on available energy. With secondary capacities that are large relative to the workload (which takes 93 mJ) we see immediate completion. However, performing the event on smaller secondary capacities can take between three hours and a day to complete even for scenarios with large amounts of harvestable ambient energy.

Workload Timeliness

Finally, we analyze the ability of different storage configurations to perform a long-running, contiguous, higher energy task. The characteristics of the long-running task are presented in Table 4.1b. We This long-running task is an example of a medium-sized, 50 kB over-the-air code update. Figure 4.5 is a CDF of the time-to-completion for this update with different configurations of energy capacity, workload, and harvesting conditions. Nearly all of the configurations with 0.28 mWh of energy storage can complete the task in the minimum amount of time (5 s). In comparison, even reducing the energy storage size to match the energy required for an update causes significant latency. Many of the updates for smaller energy storage configurations do not complete for 1000-10,000 s. This aligns with the amount of time a sensor may sit idle overnight waiting for power from daytime to continue and eventually complete its update.

4.4 Availability Requires Backup

Increasing secondary capacity greatly improves availability across all workloads. However, some environmental conditions and workloads do not reach 100% availability regardless of the size of the secondary store. Some workloads simply require more energy than is available to harvest, and cannot achieve high availability because they already capture all the energy they can. For systems that can theoretically capture enough energy on average, they still may experience periods where they run out of stored energy. Some of the results in Figure 4.3 that appear to achieve near 100% availability actually miss between 0.1 and 2% of events because they did not initially start with enough energy stored. In both of these cases, a backup energy store can greatly increase availability. However, the introduction of a non-rechargeable backup means a system now has a finite lifetime. While the second case may be solved with a sufficiently large secondary-cell that is deployed fully charged, the diminishing returns of increasing secondary size and unpredictability of some reactive applications makes it difficult to rely solely on harvested power.

We argue that 100% availability is a significant improvement over even low failure rates with respect to availability and simplicity due to the lack of intermittency. As we discuss in Section 2.3, many types of applications must provide high availability to be successful. For applications that are human facing, research suggests that any unavailability results in frustration and unwillingness to adopt automated solutions [107–109]. To use energy harvesting sensors for human-facing applications, ones that control safety-critical systems, or applications that simply require consistent and periodic sensing, inherent unavailability is intolerable. Beyond the availability of the normal sensor workload, it is also important to be able to detect and identify the failures in a sensor deployment. Batteryless system failures are difficult to detect because there is no method for differentiating between lack of energy and an actual fault. While scheduled communication of current energy state may help, this would be difficult for systems that only store enough energy to perform a single

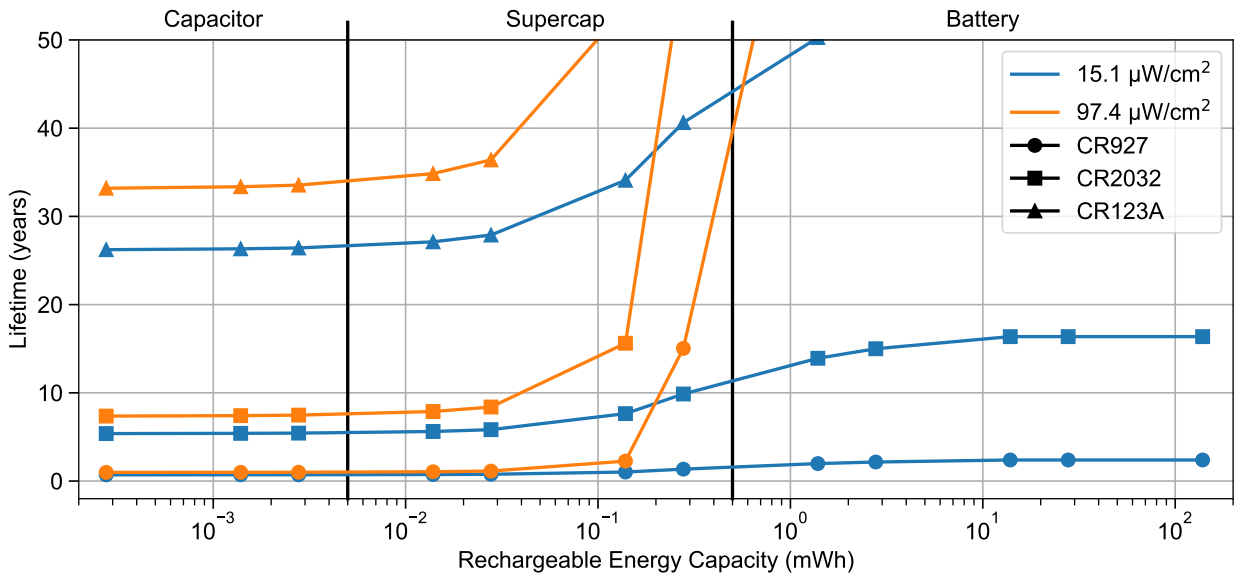

(a) Periodic Application

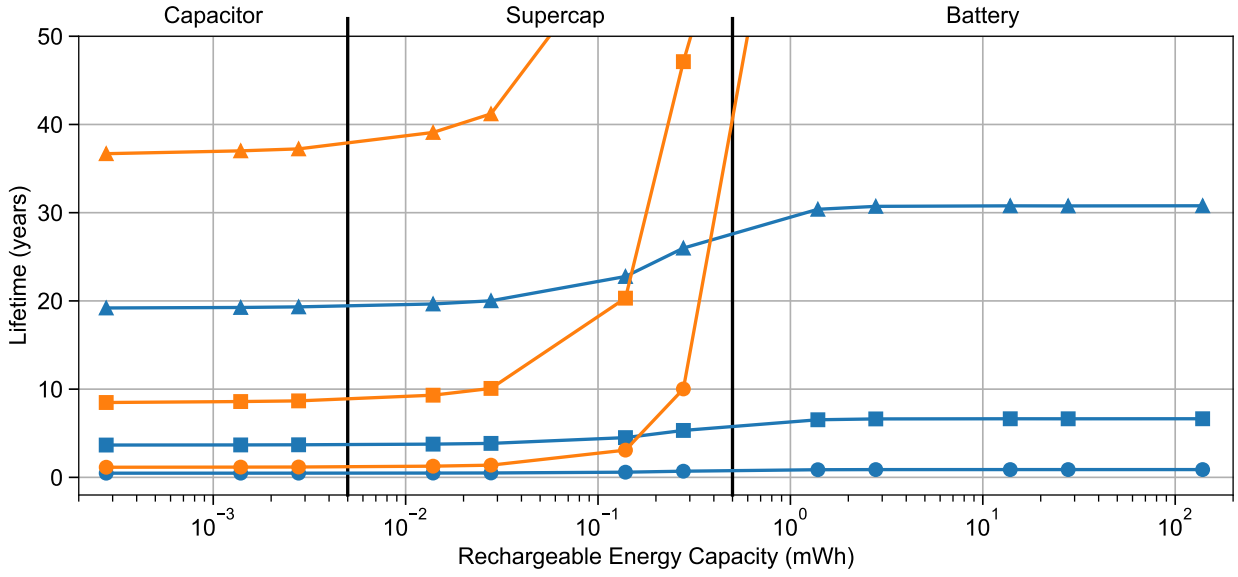

(b) Reactive Application

Figure 4.6: Estimated lifetime when varying secondary energy capacity for different harvesting scenarios and backup energy storage sizes. The periodic application's period is 30 s and the reactive application events are scaled to represent a maximum of 2000 events per hour during the peak hour. The backup sizes correspond to those found in common coin cell batteries: 90 mWh, 720 mWh, 4500 mWh for the CR927, CR2032, and CR123A respectively. As the ability to capture more harvested energy increases, the sensors lifetime increases. In some scenarios, expected lifetime becomes unbounded as the device is able to subsist entirely on harvested energy.

operation [2, 4, 7, 18]. The only way to ensure there is *always* energy available to issue a heartbeat message or report issues, a backup non-rechargeable source of energy is required.

Lifetime of a Backup Energy Store

To achieve 100% availability and avoid frequent power and state loss, designs can utilize a backup energy store. In this section we simulate the workloads of energy harvesting sensors with backup non-rechargeable energy storage. In instances where the rechargeable source is depleted, the system can operate from the backup, masking the effects of variable energy income. We estimate the lifetime of the backup based on the discharge rate resulting from the simulated workload. We consider a sensor's lifetime to be completed when its backup energy store is depleted. We assume the backup energy store is a primary battery, as they offer very low self-discharge, long shelf life, and superior energy density.

An analysis of the reliable lifetime of sensors configured with energy harvesting and a backup energy store is shown in Figure 4.6. We choose several backup energy stores with energy equivalent to those found in several types of common primary-cells. We see that with energy harvesting and a sufficiently large secondary energy store, nodes can achieve 100% reliable lifetimes that exceed what we can reasonably predict, especially for harvesting scenarios that exceed the average power of the application. In these scenarios, the inclusion of a backup energy store is critical to ensure availability in uncharacteristically adverse conditions. Even for conditions with limited energy availability we still observe significant lifetime improvements due to energy harvesting. We identify a 2-4x increase in lifetime estimates from the smallest to largest capacity simulated, if we only consider bounded results. In some cases, lifetime estimates grow exponentially as rechargeable energy capacity is increased, indicating that with sufficient capacity and income, a workload almost never needs to utilize the energy in its non-rechargeable backup. These results confirm the initial analysis performed in Section 3.1, which showed substantial lifetime improvements when considering hybrid energy harvesting systems with backup preallocated energy. We emphasize that these lifetime estimations are just estimations, and while we do model the 1 %/year leakage typical of coin cells, we do not consider the unknown physical or chemical degradation that would be experienced over decades of use or in adverse environments.

4.5 Summary

In this chapter, we present an upgraded energy harvesting and storage model and use it to simulate the behavior and common workloads of wireless sensors. This simulation improves upon the simple analysis performed in Section 3.3 by considering a dynamic workload, more realistic energy income and energy storage characteristics, and the effects of preallocated backup energy storage. We use the results from multiple simulation runs to illustrate the impact of energy capacity on the performance of wireless sensors. Increasing energy capacity allows a system to increase its energy utilization. In turn, increased energy utilization

allows a wireless sensor to more successfully complete its workload, regardless of the type or distribution of workload activity. When considering the workloads presented in Section 4.1 and the results of Figures 4.2 and 4.3, they will require an energy buffer capacity on the order of 1–10 mWh to achieve sufficient energy capture and workload availability.

Despite the improvements granted by increased capacity, a wireless energy harvesting sensor will still be limited by the magnitude of available harvestable energy. In cases where harvestable energy is insufficient to power an intended application workload, it can be beneficial to include a backup non-rechargeable source of energy to ensure high availability. While this inclusion results in a wireless sensor with a finite lifetime, energy harvesting with sufficient rechargeable energy capacity can significantly increase the system lifetime.

In the next chapter, we explore and compare the options for rechargeable energy capacity for wireless sensors, including capacitors, supercapacitors, and batteries. Capacitors and reasonably sized supercapacitors are unable to provide the energy capacity predicted by our workload simulations. However, batteries are a promising option to provide sufficient energy capacity and density to achieve the energy utilization and workload availability gains identified in this chapter.

Chapter 5

A Quantitative Evaluation of Energy Storage

From the previous two chapters, we have identified the value of properly sizing rechargeable energy capacity in an energy harvesting design to increase energy capture and system availability. Many batteryless sensor designs have chosen to utilize capacitors in their designs, which severely limit energy capacity and the performance of the system. Many modern applications are attempting to push the energy bounds of sensing, computation and networking, and batteryless systems have begun incorporating supercapacitors to provide the necessary energy capacity to make certain applications feasible [56]. These designers have eschewed batteries as an option, despite their vastly superior energy density, dismissing them qualitatively as bulky [2, 7, 15–20, 56, 59, 70, 110], inefficient [2, 19, 56, 59], expensive [7, 15–17, 20, 70], short-lived [2, 7, 16–20, 56, 59, 110, 111]. temperature-sensitive [7, 16–19], and dangerous [7, 16, 17, 20, 70, 110].

In this chapter, we reexamine each of these arguments with respect to modern capacitor, supercapacitor, and battery technology. While the early battery design of one or two decades ago were certainly guilty of many of the charges levied against them, new battery electrode materials and improved lithium-ion manufacturing processes have produced appealing options for miniature energy storage elements that do not possess the undesirable qualities of older battery technology. New battery technology paired with newly available low power energy harvesting and battery management ICs [24, 25] enables the design of high-capacity and performant energy harvesting systems. Despite improvements, batteries still underperform in some metrics compared to supercapacitors, but these deficiencies are likely inconsequential for many of the contemplated wireless sensor applications, and the substantial increase in energy capacity outweighs these detracting trade-offs.

Technology	Capacitance / Capacity	Energy Capacity (Wh)	Voltage (V)	Volume (mm^3)	Energy Density (Wh/L)	Power Density[b] (W/L)	Temperature (Charge/Discharge °C)	ESR (Ω)	Self-Discharge (nA)	Cycle Life (Cycles)[j] 100% DoD	Cycle Life (Cycles)[j] 10% DoD	Cost (USD) US[m]	Cost (USD) China
MLCC	47 µF [112]	0.000 000 259[a]	6.3	8.19	0.032	6 060 000[c]	-55 - 125	0.001-0.1[f]	<10[h]	∞[k]	∞[k]	0.16	0.03
	100 µF [113]	0.000 001 39[a]	10	20.0	0.069	6 250 000[c]	-55 - 125	0.001-0.1[f]	<10[h]	∞[k]	∞[k]	0.31	0.04
Tantalum	100 µF [105]	0.000 000 551[a]	6.3	18.6	0.030	2 670 000[c]	-55 - 125	0.2[f]	<10[h]	∞[k]	∞[k]	0.28	0.17
	220 µF [105]	0.000 003 06[a]	10	91.0	0.034	1 370 000[c]	-55 - 125	0.07[f]	<10[h]	∞[k]	∞[k]	0.37	0.16
Supercapacitor	7.5 mF [114]	0.000 070 4[a]	2.6	7.2	0.980	4690	-30 - 70[d]	25[f]	<10[i]	>10000	—	2.42	—
	33 mF [115]	0.000 139[a]	5.5	870	0.159	17 400	-20 - 70[d]	0.25[f]	<5[i]	—	—	8.65	—
	100 mF [116]	0.000 420[a]	5.5	1130	0.372	33.5	-25 - 70[d]	100[f]	<10[i]	—	—	1.10	—
	470 mF [106]	0.001 15[a]	4.2	1029	1.17	16 500	-40 - 70[d]	0.13[f]	<1000	100000+/4 yr [117][l]	—[l]	5.06	1.00
Li-ion	11 mAh [118]	0.0407	3.7	191	213	96.9	0 - 40/-20 - 60[f]	1[g]	120-400 [119]	500	10000+ [120, 121]	4.00	—
	40 mAh [122]	0.148	3.7	1010	147	44.4	0 - 40/-20 - 60[f]	3	120-400 [119]	500	10000+ [120, 121]	1.62	—
	80 mAh [118]	0.296	3.7	1010	295	147	0 - 40/-20 - 60[f]	2[g]	120-400 [119]	500	10000+ [120, 121]	7.00	—
LiPo	40 mAh [123]	0.148	3.7	660	224	224	0 - 40/-20 - 60[f]	1.5	120-400 [119]	300	10000+ [120, 121]	4.50	0.51
LTO	1.8 mAh [124]	0.0043	2.4	88.0	48.9	489	-35 - 70[f]	8	<300[g]	7000	10000+ [125]	—	1.58
	5 mAh [124]	0.012	2.4	200	60.0	600	-35 - 70[f]	2	<300[g]	7000	10000+ [125]	—	1.58
	15 mAh [124]	0.036	2.4	496	72.6	726	-35 - 70[f]	0.6	<300[g]	7000	10000+ [125]	—	1.58
	20 mAh [124, 126]	0.0480	2.4	682	70.4	1410	-35 - 70[f]	0.55[g]	<300[g]	7000	10000+ [125]	6.75	1.38
	125 mAh [124]	0.300	2.4	2360	127	1270	-35 - 70[f]	0.18	<300[g]	7000	10000+ [125]	—	1.58
LFP	70 mAh [127]	0.224	3.2	1570	143	1430	-20 - 75[f]	0.53	160 [128]	2000	30000+ [84, 85, 129]	—	1.38
	100 mAh [127]	0.320	3.2	1960	163	1630	-20 - 75[f]	0.33	160 [128]	2000	30000+ [84, 85, 129]	—	1.38
Solid State	0.7 mAh [130]	0.002 73	3.9	145	18.8	134.2	-20 - 60[f]	100	<3	4000	—	30.00	—
	0.1 mAh [131]	0.000 150	1.50	14.5	10.3	2.07	-20 - 80[f]	<200	<80	1000	—	9.31	—

[a] Energy capacity calculated with rated maximum voltage. [b] Calculated with battery and capacitor effective power.
[c] Effective power for capacitors is slightly nebulous. Here the higher end of ESR is assumed. [d] Supercapacitors experience higher ESR at lower temperatures and higher leakage at higher temperatures [117].
[e] Lithium batteries experience higher ESR, higher leakage, lower capacity and shorter lifetimes at temperature extremes. [f] ESR is frequency dependent, supercapacitors are usually rated at 1 kHz.
[g] Empirically tested. [h] Both tested and calculated from insulation resistance after absorption period. [i] Specification after 24 h of charging. [j] For batteries, measured to 80% original rated capacity.
[k] Capacitor derating is not considered. With proper design principals these should be nearly infinite. [l] Supercapacitors are time rather than cycle limited. Assumes 3V, 20 °C. No DoD dependence mentioned [117].
[m] Prices are based on cheapest available equivalent part in quantities of 100.

Table 5.1: A comparison of miniature energy storage technologies. Data is based on specific components and their datasheets, and components are chosen for each category based on their inclusion in platforms described by the literature. Some technologies are rapidly evolving, such as supercapacitors and batteries. Other citations point to general characteristics of the storage technologies not specified by their datasheets. For most applications, lithium-based batteries provide much higher energy density without reasonably impacting sensor lifetime, cost, or function. The minority of sensing applications, such as those operating at extreme temperatures, may require capacitors or active heating and cooling.

5.1 Energy Storage Technology

We summarize the notable characteristics of various examples of capacitors, supercapacitors, and batteries in Table 5.1. Many of the capacitors and supercapacitors in this table were chosen based on their inclusion in some contemporary batteryless platforms, including the Solar Monjolo [63], Flicker [7], Capybara [18], and Camaroptera [56]. The selected batteries represent some of the smallest lithium-based cells that are commercially available. This section serves to explain the characteristics of various possible energy storage technologies for low power energy harvesting applications, as a prelude to deeper dives on particular aspects and comparisons.

Ceramic and Tantalum Capacitors

Along with chip resistors, the multilayer ceramic capacitor (MLCC) is the most widely used passive component in modern electronics. MLCCs are essentially a parallel connection of many cearamic plate capacitors, packaged in a small form factor [132]. Traditionally, MLCCs are used for two applications: resonant circuits and filters and power supply bypass and decoupling [132]. With sufficient capacitance, MLCC capacitors meant for power supply decoupling can act as the sole energy storage in a system [2, 7, 63]. However, the energy storage and density of MLCC components is limited, and is often only enough to support a single small operation, like measuring and transmitting sensor results over a radio. Similarly, tantalum capacitors are often used for power supply bypass and decoupling. Tantalum capacitors consist of a tantalum metal anode and a solid electrolyte as a cathode, separated by a solid dielectric [133]. They generally provide more capacitance per volume than MLCCs, but are polar and have worse efficiency. Tantalum capacitors also do not exhibit any noticeable aging effects, whereas MLCCs experience slight capacitance change over long periods of time [134]. Compared to batteries and supercapacitors, MLCC and tantalum capacitors provide infinitely longer lifetimes and higher power density, but are severely limited in energy capacity and density.

Supercapacitors

Electrochemical double layer capacitors (EDLC) are the most common type of supercapacitor, and are utilized in many batteryless platforms [18, 56, 57, 135]. Supercapacitors generally consist of electrodes separated by a liquid electrolyte, like batteries, however energy accumulation is through electrostatic interaction instead of chemical reactions [135, 136]. EDLCs achieve significantly higher energy density than MLCC and tantalum capacitors due to their large effective surface area and very small charge separation distances. Supercapacitors are also durable and long-lived, capable of millions of charge-discharge cycles [135]. supercapacitors excel in high power, high cycling applications, where large amounts of charge must be stored and provided quickly, at a high frequency. However, they offer lower power densities due

to higher internal resistance compared to MLCC and tantalum capacitors [136], and lower energy density than batteries [106, 126].

Li-ion

Lithium-ion batteries encompass all batteries that utilize lithium ions for charge transfer. These batteries are well known for their superior energy density, and are used in small wireless and portable consumer electronics as well as in large multi-cell configurations in electric vehicles. Like all batteries, Li-ion batteries consist of two oppositely charged electrodes separated by an electrolyte. The electrodes consist of a negatively charged anode and positively charged cathode. During discharge, Li^- ions move from the anode through the electrolyte to the cathode, and during charge they move from the cathode to the anode. Lithium polymer (LiPo) batteries, most commonly used in consumer electronics, utilize a dry or gel electrolyte, while coin cells or cylindrical Li-ion cells utilize a liquid electrolyte.

There are many different types of li-ion batteries, utilizing different anode, cathode, and electrolyte materials. Most modern Li-ion cells primarily utilize graphite as an anode material, and lithium nickel manganese cobalt oxide ($LiNi_{0.33}Mn_{0.33}Co_{0.33}O_2$ or NMC for short) as a cathode material [137]. These cells offer a nominal voltage of 3.7 V, high energy density, and long lifetimes. The choice of NMC provides several benefits over earlier designs that utilize $LiCoO_2$ (LCO) or $LiMn_2O_4$ (LMO) for cathode materials. LCO lithium batteries are notorious for their thermal instability and fast capacity fade at high current rates or deep cycling [137]. Cobalt is also a toxic and expensive metal to produce. LMO lithium batteries are cheaper, have higher power density, and have significantly better thermal stability than LCO cells, but have lower energy density and still have poor cycling stability, especially at higher temperatures [137]. The combination of nickel, manganese, and cobalt in NMC cathodes results in higher structural stability, longer lifetimes, and less reliance on expensive transition metals like cobalt [137].

Batteries that utilize LCO cathodes have the propensity to ignite or explode due to thermal runaway when stressed thermally, mechanically, or electrically [138]. This is primarily due to the cathodes releasing oxygen at high temperatures, which start an exothermic reaction with the organic parts of the cell [137, 138]. Like LCO, the newer NMC cathode material is still susceptible to thermal runaway with abuse, however the peak magnitude of self-heating on runaway is an order of magnitude less than that of LCO, and onset is delayed and begins at a higher temperature [138]. Batteries with LMO cathodes are safer than both LCO and NMC, but have poor cycle cycling lifetime which has limited the commercialization of the technology.

Lithium Iron Phosphate

Lithium iron phosphate ($LiFePO_4$, or LFP), is a more recently commercialized cathode material and stands to offer similar safety and power capability to LMO while offering long lifetimes similar to NMC [137, 139]. Batteries that utilize LFP as a cathode material possess

a lower nominal voltage (3.2 V), and lower energy density than LCO and NMC batteries. However, LFP cathodes offer higher cycle stability and lifetime, have lower thermal sensitivity, and are cheaper to produce than cobalt and manganese-based cathodes [137–139]. LFP cells are also very safe compared to NMC and LCO cells. The only thermal runaway experienced by LFP batteries is dominated by reactions between the electrolyte and graphite anode, which decomposes at high temperatures [138].

Lithium Titanate Oxide

Perhaps the most promising replacement for graphite anodes is lithium titanate oxide ($Li_4Ti_5O_4$). Lithium titanate oxide, abbreviated LTO, is usually paired with an LMO cathode, and sometimes an NMC cathode [137, 140, 141]. LTO offers superior thermal stability, high discharge/charge rates, and longer lifetimes compared to graphite. These improvements come at a cost of the more expensive titanium compound, and a lower energy density and nominal voltage (2.4 V) [137, 141]. Cells that incorporate LTO anodes are also extremely safe compared to those that use graphite [137, 140, 141]. Unlike graphite, LTO anodes do not produce Li dendrites after considerable cycling [137]. This reduces the risk of inadvertent internal shorts and thermal runaway. LTO anodes also remain stable and do not break down at high temperatures [140].

Solid-state

All cylindrical or pouch Li-ion batteries employ a liquid or gel electrolyte of lithium salts dissolved in an organic, non-aqueous solvent [138]. Liquid electrolyte requires specific packaging to prevent leakage, and an internal separator to prevent shorts between the electrodes. This limits the miniaturization of Li-ion batteries, as this packaging is increasingly difficult to manufacture at smaller sizes. This results in smaller batteries that have uncharacteristically high internal resistance and lower energy density compared to larger cells with the same chemistry. The organic solvent used in liquid and gel electrolytes also can react exothermically with any oxygen released upon the breakdown of electrode material and cause thermal runaway. In addition to developing new anode and cathode materials to mitigate this, researchers and companies have also begun experimenting with replacing the non-solid electrolyte non-reactive alternatives. Solid-state batteries are safer, due to the use of non-reactive solid electrolytes, come in smaller packages, offer longer lifetimes, and have very low self-discharge [142]. However, solid-state batteries generally have limited energy and power density and high internal resistance compared to aqueous Li-ion cells. Solid-state batteries are also currently expensive to manufacture, and their commercial viability has been limited [142].

Summary

Capacitors possess superior power density and have functionally infinite lifetimes, but can only store minute amounts of energy. Supercapacitors provide one to two orders of magnitude

more energy density than MLCC and tantalum capacitors, at the cost of reduced power density, and shorter lifetimes. By contrast, batteries offer one to two orders of magnitude more energy density than supercapacitors. However, batteries do have limited lifetimes based on cycle counts, are more temperature sensitive, and some can be dangerous if mishandled.

In the rest of this chapter, we seek to compare these technologies in more depth. The following sections directly analyze each of the high-level qualitative arguments that "batteryless" platform designers have leveled against batteries. Namely, that they are bulky, inefficient, expensive, short-lived, temperature-sensitive, and dangerous.

5.2 Volume and Density

Arguments that deride batteries as bulky are likely directly comparing the size of a small battery to that of single tantalum or ceramic capacitor, without considering energy and power density. Modern commercially available miniature batteries are comparable in volume to many supercapacitors, and even a banked combination of ceramic and tantalum capacitors, while offering substantially more energy density and an acceptable power density. In this section, we explore the "bulkiness" of capacitors, supercapacitors, and batteries in the context of energy and power density. Here, we are considering volumetric density (Wh/L and W/L) instead of specific energy and power (Wh/kg and W/kg, respectively) to better compare the volume of these energy storage options. The energy density and power density of various capacitors, supercapacitors, and batteries are compared in Figure 5.1.

Energy Density

Energy density should be the primary consideration for energy harvesting power supply design to maximize energy capacity while simultaneously minimizing volume. The energy stored in a capacitor is calculated in one of two ways, illustrated by Equations (5.1) and (5.2).

$$E_{total_{cap}} = \frac{1}{2}CV^2 \tag{5.1}$$

$$E_{eff_{cap}} = \frac{1}{2}C(V^2 - V_{min}^2) \tag{5.2}$$

Where E_{total} and E_{eff} are the total and effective energy stored in a capacitor, respectively. These amounts are defined by the capacitor's capacitance C, and the applied voltage V, and the minimum voltage V_{min}. Usually V_{min} represents the minimum supply voltage for a system. Often this is the minimum open circuit voltage to cold start a boost regulator, between 400 mV to 600 mV [24, 25, 143]. For simplicity, we use E_{total} to determine energy capacity and density. For most capacitors, the unusable energy represented by $E_{eff_{cap}}$ is negligible compared to the total energy.

The energy stored in a battery cannot be directly calculated, and depends on many variables. It can be estimated by considering the charge capacity Q, often denoted in terms of Ah, and the nominal voltage of the cell V_{nom} This estimation is described by Equation (5.3).

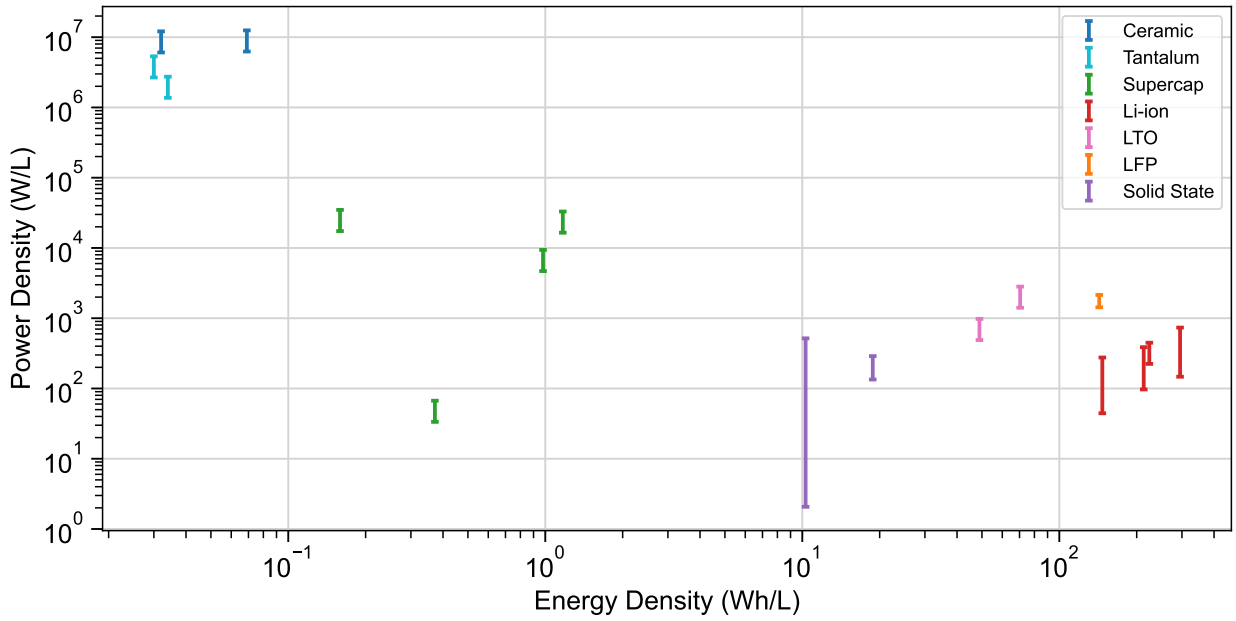

Figure 5.1: Ragone plot for components listed in Table 5.1, in log-log scale. The upper cap represents the maximum power that a storage element can provide while the lower cap represents the effective, continuous power it can provide. A ragone plot directly compares power and energy density for different devices. Ceramic and tantalum capacitors are very power dense, but provide abysmal energy density. Batteries provide superior energy density, are less power dense. Supercapacitors exist between these two extremes. Even though batteries do not provide comparable power density to either capacitors or supercapacitors, they can still provide sufficient power for common wireless sensor workloads, like operating a short or long range radio.

$$E_{bat} \approx QV_{nom} \tag{5.3}$$

Nominal voltage represents an average of the battery's voltage curve over the course of a charge/discharge cycle. The nominal voltage and capacity of a battery are provided by the manufacturer and are commonly included in a datasheet.

A selection of capacitor, supercapacitor, and battery energy capacities and densities are summarized in Table 5.1, and their energy density is compared in Figure 5.1. Among this selection, small batteries are 50-1000x more energy dense than supercapacitors and three to five orders of magnitude more dense than ceramic and tantalum capacitors. Li-ion and LiPo batteries are the most energy dense among all options. Solid-state batteries are an order of magnitude less energy dense than those with aqueous or gel electrolytes, but still an order of magnitude more dense than supercapacitors.

Several of these capacitors, supercapacitors and batteries are shown visually in Figure 5.2. The capacitor and supercapacitor configurations are based on examples from batteryless

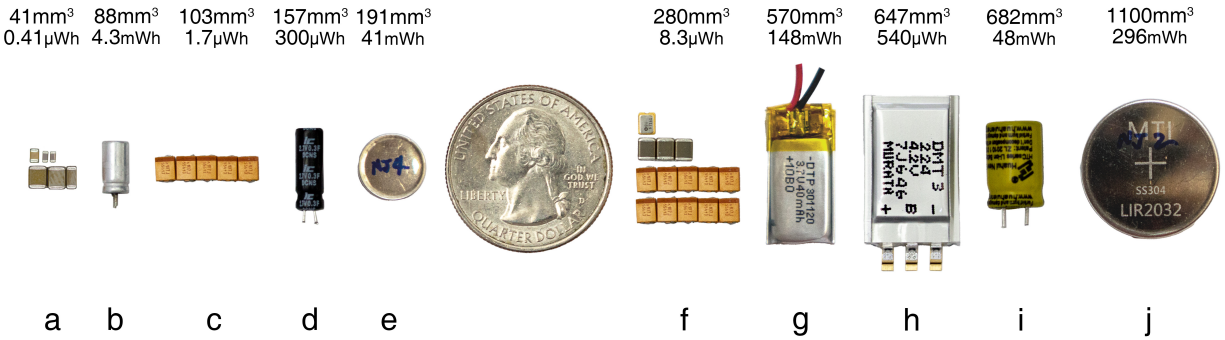

Figure 5.2: A size comparison of energy storage methods including capacitors, supercapacitors, and batteries. They are ordered left to right, by their (total) volume. Total volumes and energy storage are listed above the respective device. Configuration **(a)**, **(c)** and **(f)** represent the energy storage configurations used in the Flicker platfrom with BLE and several sensors [7], the Solar Monjolo [63] and the Capybara temperature monitor and alarm [18], which have total capacitances and energy capacities of 119 µF (0.41 µWh at 5 V), 500 µF (1.7 µWh at 5 V) and 8.8 mF (8.3 µWh at 2.6 V), respectively. Capacitors **(d)** [144] and **(h)** [106] are large supercapacitors available on the Capybara platform and have the capacitances and energy capacities of 300 mF (300 µWh at 2.7 V) and 220 mF (540 µWh at 4.2 V) respectively. Devices **(b)** and **(i)** are small LTO battery cells with 1.8 mAh (4.3 mWh at 2.4 V) and 20 mAh (48 mWh at 2.4 V) capacity respectively [124]. Devices **(e)** and **(j)** are small prototype Li-ion coin cells with 11 mAh (41 mWh at 3.7 V) and 80 mAh (296 mWh at 3.7 V) respectively [118]. Device **g** is a traditional Lithium Polymer pouch cell with 40 mAh (148 mWh at 3.7 V) [145]. The LTO battery **(b)** and the Li-ion coin cell **(e)** are among the smallest of all configurations of energy storage presented here and also provide one to two orders of magnitude more energy capacity compared to **(f)**, the largest supercapacitor presented.

platform designs described in the literature [7, 18, 63]. The batteries shown in Figure 5.2 are as small as 88 mm^3, and resemble small through-hole capacitors and coin cells. Battery **(b)** is smaller in volume than many of the capacitor configurations presented in the literature, only outdone by systems like Flicker **(a)** which utilizes only a few ceramic capacitors [7]. This small LTO battery offers 8x more energy storage and 60x more energy density than the largest supercapacitor presented (**h**). When considering the size of other components in the system, most notably the harvester (solar panel, thermocouple, RFID antenna, or piezoelectric device), the combination of ICs, and large sensors like a PIR motion sensor, the size of small rechargeable batteries is inconsequential. The Michigan Micro Mote is one of the smallest energy harvesting systems ever built, occupying a volume on the order of a single ceramic capacitor [75]. Despite its small size, it utilizes a thin film solid-state lithium battery for energy storage due to its superior energy density over any capacitor or supercapacitor option. When one considers energy density instead of purely volume, capacitors and supercapacitors are significantly more bulky than batteries.

Power Density

In addition to energy density, power density is also an important metric to consider for a design. Common wireless sensor workloads are characterized by very low sleep currents punctuated by infrequent pulses of high current, usually a radio transmission. Energy storage must provide sufficient peak power to drive these short pulses, but largely, these applications require low continuous power. The maximum peak power is largely dependent on equivalent series resistance (ESR) of the storage device, represented by an internal series resistance to the capacitor or battery cell. Internal resistance for both supercapacitors and batteries is temperature and age dependent. Both storage elements experience increased ESR at temperature extremes, and experience increased ESR as they age. The internal resistance of capacitors and supercapacitors is also frequency dependent, and usually reported for 1 kHz. This value is generally related inversely with frequency for frequencies below the capacitor's self-resonance [146]. For the relatively low frequency of charge/discharge cycles characteristic of energy harvesting devices, actual ESR is likely higher than reported for supercapacitors.

Internal capacitor, supercapacitor, and battery resistance incurs a voltage drop over this resistance which is especially noticeable and detrimental during high current loads. Ceramic and tantalum capacitors have negligible ESR and incomparably high power density, so we focus on comparing the power density of supercapacitors and batteries. There are two different metrics for quantifying power output of supercapacitors and batteries. The first is effective power P_{eff}, and represents the maximum sustainable continuous power that can be provided. The second is peak power P_{max} and represents the maximum possible current that can be provided in short pulses. These metrics are defined differently for supercapacitors and batteries. For a supercapacitor, power capabilities are defined by Equations (5.4) and (5.5) [147].

$$P_{eff_{sc}} = \frac{1}{8}\frac{V^2}{R_i} \tag{5.4}$$

$$P_{max_{sc}} = \frac{1}{4}\frac{V^2}{R_i} \tag{5.5}$$

Where V is the voltage applied to the capacitor, and R_i is the internal resistance, or ESR. For batteries, these metrics are defined by Equations (5.6) and (5.7).

$$P_{eff_{bat}} = I_{cont}V_{nom} \tag{5.6}$$

$$P_{max_{bat}} = I_{max}V_{nom} \tag{5.7}$$

Where I_{cont} and I_{max} are the battery's rated continuous and peak pulsed current respectively. These metrics are provided by the battery manufacturer and often included in the battery datasheet. The continuous and peak currents are often defined in terms of the C-rate, or a proportion of the rated capacity (in units of Ah). For example, the 1C rate of a 100 mAh battery is 100 mA, and the 2C rate is 200 mA. Continuously charging and discharging a battery beyond its specified rate will result in cycles that deliver less energy than rated, and eventually damage to the cell resulting in capacity fading. For sake of comparison, we use

P_{eff} in Table 5.1 for batteries and supercapacitors, as it is a more conservative measure of the energy storage power capability.

Among the supercapacitors and batteries featured in Table 5.1, supercapacitors provide 10-400x higher power density over batteries. However, one supercapacitor outlier provides the least power density of all options. Like energy density, power density is also directly compared in Figure 5.1. Despite their superior power density, most applications do not require the higher power density afforded by supercapacitors. It is hard to imagine a *low power* energy harvesting application that must source more than a few to a few hundred mA at 3 V infrequently, never mind continuously. Solid-state batteries are more power limited than other battery types. Despite their low energy capacity they are able to supply high C-rates, between 15-50C. This corresponds to currents on the order of 10 mA. Conventional Li-ion and LiPo cells can generally source 1C continuously. The smallest Li-ion cell listed in Table 5.1 can provide 11 mA continuously, and the largest can provide 80 mA. The LiPo cell presented can supply 40 mA continuously. Small LTO and LFP cells are capable of very high C-rates, often between 20-40C for LTOs, and 10-20C for LFP [124, 126, 127]. The smaller LTO battery listed in Table 5.1 is able to source 18 mA, while the larger LTO and LFP cells can source between 400-700 mA continuously. This power capability is more than sufficient for the majority of wireless sensor applications, either indoors with PANs, or outdoors, with cellular or LPWANs [148, 149].

Another consideration for power capability and ESR is the effect of voltage drops during high power loads. If a voltage drop is sufficiently large, it can drop the supply to a level unusable by a voltage regulator or CMOS logic. This could effectively render part of the power storage unusable in an unpredictable manner. This effect is particularly detrimental for supercapacitors. While ESR is comparable between batteries and supercapacitors, the stability of provided voltage is not. Batteries provide a stable voltage curve centered around their nominal voltage, and voltage drops due to ESR can still result in a usable voltage even when almost empty. Supercapacitors, on the other hand, experience a (approximately) linear decrease in voltage with current until empty. This means that, depending on the load (in intensity and frequency), a large load and subsequent voltage drop could occur that causes a system brown-out, potentially corrupting state, or an unexpected reset.

5.3 Efficiency

At a high level, the efficiency of an energy storage element can be defined as the actual proportion of stored energy that is used to perform a desired task or application. Batteries have been derided as "inefficient" by batteryless platform designers, however both supercapacitor and battery technology exhibit the same two phenomena that causes inefficiency, and generally to the same extent. The first phenomena is power dissipation over the internal resistance of the device. The second is self-discharge or leakage, which is represented by an internal parallel resistance to the capacitor or battery. This section seeks to quantitatively compare these two sources of inefficiency for supercapacitors and batteries.

Internal Resistance

The power dissipated over the internal resistance, or ESR, of an energy storage element can be calculated with Ohm's law, reproduced in Equation (5.8).

$$P_i = R_i I^2 \tag{5.8}$$

Due to the squared relationship of current to power, high current loads, such as a radio transmission or long computation, are the primary cause of ESR power losses. The total power required to drive a load, including losses over internal resistance, is represented by Equation (5.9).

$$\begin{aligned} P_{total} &= P_l + P_i \\ P_{total} &= IV + R_i I^2 \end{aligned} \tag{5.9}$$

Where P_l, I and V are the power, current, and voltage required to drive an intended load. Power inefficiency due to ESR can be costly when considering high current loads.

With a subset of the capacitors and batteries mentioned in Table 5.1, an 8 mA BLE transmission from a steady 3 V would incur less than 0.03% in resistive loss from a tantalum capacitor, a 2.1% loss from the 1.8 mAh LTO battery, and 6.3% loss from the 7.5 mF supercapacitor. A 130 mA LoRa transmission would incur a 0.43%, 25%, and 52% overhead, respectively [149]. These selections represent some of the worst performing examples in terms of ESR. For most supercapacitor and battery product lines, as they get smaller, capacitance and capacity decreases while ESR increases. When the surface area between cell electrodes decreases, this limits the flux of ions traveling between electrodes. This is manifested as an increased internal resistance. This has traditionally been an issue for small Li-ion batteries. However, the recent commercialization of new manufacturing processes and new cathode and anode materials that increase electrode surface area has led to small form factors with lower ESR, on the order of 1-2 Ω [118]. Current solid-state battery technology still struggles with ESR, with values at or above 100 Ω [130, 131]. In Table 5.1, there are several batteries and supercapacitors that exhibit much lower ESR, on the order of 1 Ω or less, and are thus more efficient with high current loads. The 20 mAh LTO battery is in the same product line as the 1.8 mAh cell, but has 15x less ESR. This cell would only incur a 0.15% and 2.3% loss on the respective BLE and LoRa transmissions. Generally, there are options for both supercapacitors and batteries that offer comparable ESR efficiency.

Self-Discharge

In addition to an internal series resistance, capacitors and batteries both feature a non-ideal parallel resistance that causes a continuous self-discharge. The self-discharge of supercapacitors and batteries is generally dependent on their size and temperature. Larger capacity/capacitance batteries and capacitors from the same product line will exhibit more self discharge than smaller ones.

The batteries selected in Table 5.1 typically exhibit less than 500 nA self-discharge in standard environmental conditions. Solid-state batteries exhibit even less self discharge, at similar rates to capacitors. Notably, many types of Li-ion and LiPo batteries also require additional protection circuitry to prevent deep discharges and charging/discharging at high currents. This additional circuitry incurs additional self-discharge. However, extremely efficient options exist to manage small batteries [24, 25, 150]. The LTC4071 battery charger and protection IC requires only 550 nA when in operation, and features a very low disconnect current (<1 nA) to extend shelf-life [150]. This amount of self-discharge is negligible when considering the average sleep currents of common sensors and MCU options. Generally, sleep currents will still be dominated by memory retention and low frequency clock operation, on the order of a few µA.

While the self discharge of capacitors and supercapacitors appears lower than batteries, this is partially due to how it is measured and rated. However, the self discharge of supercapacitors is actually is comparable to that of batteries when considering short term post-charge behavior. Supercapacitors exhibit a pronounced, non-ideal phenomena known as dielectric absorption [151]. Dielectric absorption represents a decreasing exponential decay of the supercapacitor voltage immediately after charging. After hours or days, self-discharge is dominated by a linear leakage current. The magnitude of the decay depends on the initial voltage, temperature, and duration of the charge [152]. Datasheets generally specify supercapacitor leakage after a subsequent 24 hour constant voltage charge and 1 hour open circuit period. This is sufficient time for the contribution of dielectric absorption to be rendered negligible. After this time, only the internal parallel leakage resistance is a factor.

In the short term, however, the discharge due to dialectric absorption dominates. For example, after an hour of constant voltage charging, the 33 mF BestCap supercapacitor listed in Table 5.1 experiences an average of 300 nA self-discharge over a 3 hour window. Considering the common use case of supercapacitors in batteryless systems, where any captured energy is immediately used whenever it is available, this short term cycle use case means that energy is never stored for an extended period of time, and dialectric absorption is a primary factor in self-discharge for batteryless systems. While we have not characterized the self-discharge of each individual supercapacitor listed in Table 5.1, this phenomena is inherent to electrochemical capacitor technology.

Generally, one can expect the short-term self-discharge of supercapacitors to be comparable with batteries. In either case, self-discharge is not sufficient to warrant discounting either option for a low power design. When considering the impact of self-discharge with that of ESR loss, there are suitable options from both technologies that would provide satisfactory performance.

5.4 Expense

Like the bulkiness argument, the claim that batteries are expensive is only valid when directly comparing the cost of a battery to that of single ceramic or tantalum capacitor. This view is

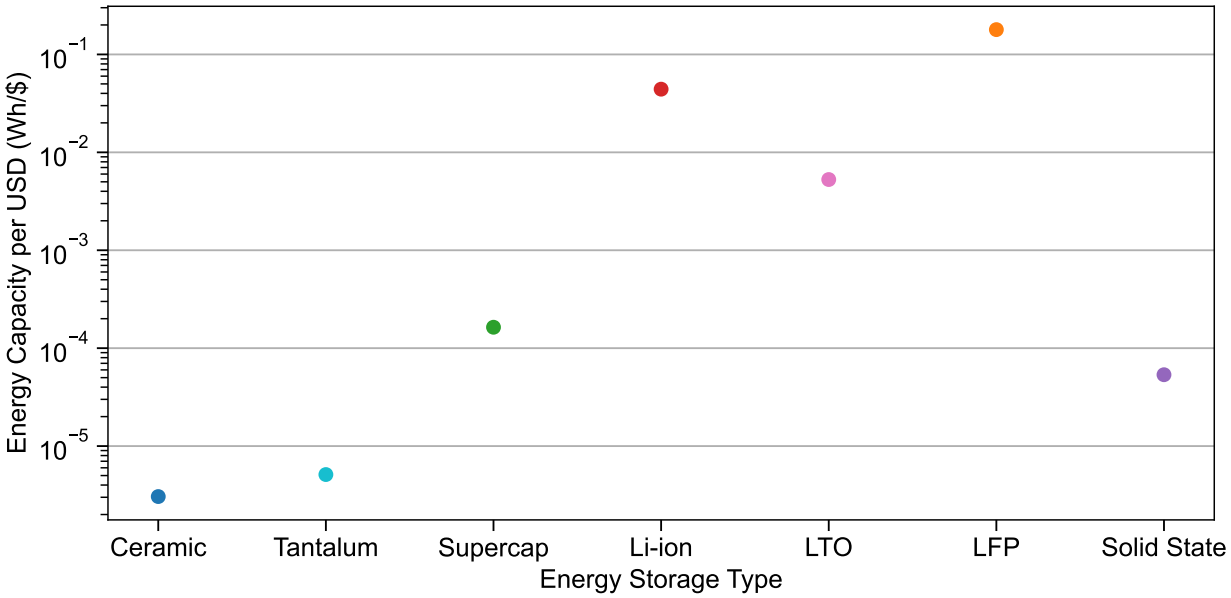

Figure 5.3: Average energy capacity for various technologies selected in Table 5.1, normalized by price in USD. Wherever possible, component costs were determined by their price in the United States. In regards to energy capacity, batteries offer 3-5 orders of magnitude more energy capacity per dollar than ceramic and tantalum capacitors. Batteries also offer 2-3 orders of magnitude more capacity than supercapacitors.

reductionist, and this section seeks to explore this argument in more detail.

Most batteries are actually comparatively cheap when considering energy capacity and density. A single battery provides significantly more energy storage per dollar than any capacitor, supercapacitor, or banked configuration. Figure 5.3 illustrates this by comparing the energy per dollar one can expect from capacitors, supercapacitors, and batteries. Solid-state batteries are still costly, as they are a newly commercialized technology and their manufacture is still expensive. Despite this, solid-state batteries offer a similar magnitude of energy capacity per dollar to supercapacitors. Other battery types offer several orders of magnitude more energy capacity per dollar than capacitors and supercapacitors.

The expensiveness of capacitors is accentuated when considering that some designs require multiple parallel capacitors or supecapacitors to build up enough energy capacity to enable required atomic operations. For example, the Capybara design is configured with 14 of the 7.5 mF Seiko supercapacitors, costing $2.42 USD, for a total of $33.88 [18]. The small 1.8 mAh battery offers 4.4x the energy capacity of the combined 14 parallel supercapacitors, for only $1.25. It would require 61 of the Seiko supercapacitors to provide the same amount of energy capacity as the 1.8 mAh battery. Those 61 supercapacitors would cost $148.

Besides the Seiko supercapacitor, several of the selected supercapacitors in Table 5.1 are as expensive or more expensive than the selected batteries [106, 114–116]. Small, 2-50 mAh LTO batteries can be purchased for $6.75 USD each from US distributors and $1.25 USD

each from Chinese manufacturers, even in small quantities [124, 126].

Also, when considering the cost of other components in a wireless energy harvesting system, like an SoC [148], harvester [153] and sensors [154], which each cost around $5 USD each in low quantities, a battery will not constitute the driving cost. Batteries are by far the most cost effective option for providing energy capacity. The argument that they are too expensive for a sensor design does not consider the comparable cost of supercapacitors, and entirely ignores the benefits of the superior energy density of batteries.

5.5 Lifetime

Batteries are often considered short-lived due to their limited cycle lifetimes, especially when compared to ceramic and tantalum capacitors. Ceramic and tantalum capacitors lifetimes are estimated to be thousands to millions of years with proper voltage derating and when used at room temperature [60]. The lifetime of capacitors, supercapacitors, and batteries generally refers to the lifetime before the device experiences parametric failure. Parametric failure is when the device is no longer within specification, usually when rated capacitance/capacity is lower than 20% its original value, or when ESR or other parasitics are sufficiently higher than rated. This section serves as an effort to quantify and compare battery lifetime with capacitors and supercapacitors, and to investigate methods for elongating their cycle lifetimes. While modern battery technology will never compete with the longevity of capacitors, for low power energy harvesting, cycle lifetime is unlikely to be the limiting factor on the lifetime of a battery, never mind the system as a whole.

While ceramic and tantalum capacitors can potentially last a million years or more [60], these lifetime estimates do not hold for supercapacitors, which are often rated in thousands of hours at a specified voltage and temperature (often 65-85 °C) [106, 115]. This lifetime is further influenced by the cycling characteristic and intensity of the workload [155]. For low power wireless sensors in room temperatures, with the sporadic cycling rate of an intermittent system, one can still expect lifetimes of one hundred thousand to one million hours (around 10-100 years) [155]. While this is a long lifetime, it is by no means infinite.

Conversely, the cycle lifetime of a battery is generally the number of full cycles at a rated continuous discharge/charge current (usually 0.5C or 1C) before the battery's capacity diminishes to 80% of its original rated capacity. Generally, NMC lithium batteries offer a cycle lifetime between 3000 and 5000 cycles [139, 156], compared to 300-700 cycles for LCO and LMO batteries. For small form factor batteries, this cycle lifetime is less, and for those in Table 5.1, they offer only 300-500 cycles [118, 123] at a 0.5 C charge/discharge rate. However, new battery chemestries like LTO and LFP batteries offer between 2000 and 7000 cycles at 0.5 C [84, 85, 124–127, 129, 139]. Solid-state batteries also offer long cycle lifetimes, between 1000 and 4000 cycles for commercially available cells [130, 131].

These rated cycle lifetimes represent the cycle life at 100% depth-of-discharge (DoD), however battery cycle lifetime is heavily influenced by the rate of discharge, as well as the depth of discharge. The reduction of battery DoD to 10% exponentially reduces cycle

capacity loss, resulting in potential lifetimes of greater than 10,000 cycles before reaching 20% capacity degradation with LFP cells [84, 85]. Similarly, LTO cells are estimated to sustain approximately 15,000 cycles before reaching 20% capacity degradation [157]. These cycle estimations for LTO cells hold true even for relatively high temperatures (50-60°C) [85, 157]. LTO chemistries can be expected to survive one thousand cycles at 100% DoD at 55°C [158]. Life cycle expectations for both 100% and 10% DoD are summarized in Table 5.1.

On its face, several thousand cycles does not seem like a lot compared to the lifetime estimates of capacitors and supercapacitors. However, for batteries, this cycle lifetime amounts to a significant amount energy. This energy also represents a significant amount of time when considering the energy capacity of batteries, and the expected charge/discharge behavior of energy harvesting wireless sensors. The comparatively vast energy capacity of batteries means that each cycle represents a significant amount of energy. This amount of energy can drive a low power workload for an extended period of time. For example, the representative periodic workload described in Table 4.1 requires an average power of 58.6 µW. For the smallest 1.8 mAh battery, a single discharge cycle represents just over 3 days of continuous operation. The rated 7000 cycle lifetime of this cell represents 117 years, assuming an identical charge/discharge cycle with 100% depth of discharge.

For wireless sensor workloads, battery cycle lifetime is not going to constitute the main source of failure in designs that incorporate them. Instead, limited shelf-life and poor battery management are likely to be the driving forces behind usable battery lifetimes. For a rechargeable battery, shelf-life represents the amount of time a battery can sit uncharged before depleting itself and experiencing capacity degradation. Batteries must also be properly charged and discharged within rated current limits, and should not be overvolted or undervolted. It is conceivable that energy harvesting systems may be deployed in areas with little available harvestable energy, and be unable to charge their batteries for long periods of time. The shelf-life for lithium-based batteries is approximately one decade, however, LTO chemistries have been shown to exhibit no long-term damage when undervolted, even to zero volts [159]. This is a significant improvement over traditional technologies like LCO, LMO, and NMC which suffer capacity degradation if undervolted due to long storage without charging. Regardless of the type of battery, designs should utilize the myriad of small battery charger and management ICs that exist to properly manage battery state and minimize shelf-life effects [24, 25, 150].

The lifetime of an energy harvesting system is also not solely dependent on the lifetime of its power supply. Most notably, some components exhibit significant long-term calibration drift. For example, each year a humidity sensor [154] expects a quarter of a percent relative humidity drift, while an oscillator [160] expects 3 ppm drift. There is also the question of relevancy in the face of decades of future progress in networking, processor efficiency, and MEMS sensors. At some point, the wireless sensor platforms built today will be obsolete and less useful, regardless of their theoretical lifetimes. Sensors do not need to last indefinitely, they need to last long enough and provide enough benefit to justify their original deployment. Replacement and renovation is inevitable for any infrastructure. Conservatively, rechargeable batteries have the capability to last between 10 and 20 years when managed properly. While not cycle limited, supercapacitors have similar lifetime estimates on the order of decades.

When considering this reality, only systems that employ ceramic and tantalum capacitors can claim to support indefinite lifetimes.

5.6 Temperature Sensitivity

In addition to lifetime considerations, batteries are also notorious for their temperature sensitivity. As mentioned previously, in temperature extremes, batteries exhibit higher ESR, self discharge, and have accelerated capacity degradation. Out of the many arguments against batteries, temperature sensitivity is the most coherent. This section seeks to quantify and compare the temperature sensitivity of capacitors, supercapacitors, and batteries and identify the applications for which it has a significant impact.

Ceramic and tantalum capacitors are very temperature resistant. Of the few selected in Table 5.1, both types are rated for -55 to +125 °C [105, 112, 113]. The ceramic and tantalum capacitors both exhibit approximately 10% capacitance difference at -55 and 85 °C. Extreme temperatures can reduce operational lifetimes, however. Tantalum capacitor lifetime is reduced from 4000 years at 45 °C to just 14 years at 85 °C [60].

Supercapacitors are also rated for extreme temperatures, and those listed in Table 5.1 can withstand temperatures between -40 to -20 °C on the low end to up to 70 °C. Supercapacitors generally experience increased ESR with lower temperatures, and lifetime limits with higher temperatures [106, 115, 117, 155]. Capacitance is generally affected by low and high temperatures. For the 33 mF BestCap, at -40 °C, ESR can increase to 20x its rated value. Similarly for the 470 mF Murata capacitor, it can experience 8x its rated ESR at the same temperature. All supercapacitors have a rated lifetime in terms of hours at a specific temperature. The BestCap can withstand 1000 hours at 70 °C, and the Murata cell is rated for 1000 hours at 85 °C before experiencing a capacitance degradation of 20%. For supercapacitors, the effects of extreme temperatures is actually quite severe. At cold temperatures, supercapacitors will be unable to supply current at a sufficient voltage, and at high temperatures their rated lifetime can decrease from hundreds of thousands of hours to just 1000, or from approximately 11 years to 40 days. Even a moderately high temperature of 45 °C can reduce some supercapacitor operational lifetimes from 14 years to 3 [155].

Like supercapacitors, batteries experience adverse effects at extreme temperatures. At low temperatures, lithium-based batteries exhibit a reduced energy capacity, support lower charge and discharge rates, exhibit higher ESR, and can experience accelerated shelf-life and cycle aging [161]. For a LFP battery at -20 °C, its capacity is reduced to approximately 60% of its original value. When this battery is stored at -20 °C for 17 days, it experiences a further 10% capacity degradation to 50%, and a 16x increase in ESR to 8 mΩ from 0.5 mΩ [161]. If charging rates are exceeded at cold temperatures, or the battery is simply cycled at low temperatures, additional battery capacity degradation can occur. The same LFP battery as above experiences a 12% degradation after only 12 cycles at -20 °C [161]. Likewise, at higher temperatures, lithium batteries also experience accelerated aging and increased internal resistance [162]. At extreme temperatures, some types of lithium batteries can

experience thermal runaway and cause explosions or fires. These effects are rarely described and quantified in battery datasheets. Instead, manufacturers simply provide a range of operational temperatures.

Li-ion and LiPo batteries are particularly sensitive to temperature. Of those listed in Table 5.1, they are generally only rated for 0-40 °C for charging, and -20-60 °C for discharging. We would expect them to perform worse than the metrics listed previously for the $LiFeMnPO_4$ battery. Generally, LTO and LFP batteries perform better in extreme temperatures than traditional Li-ion and LiPo cells. Some datasheets and authors report operating LTO batteries successfully as low as -30 °C and as high as 75 °C [124, 127, 163]. From their datasheet, the selected LTO cells in Table 5.1 exhibit a 10% capacity degradation at -20 °C, and a 20% degradation at -30 °C. At high temperatures, these cells only experience a 2% degradation at 75 °C, and less than 1% at 60 °C. LFP cells perform worse than LTO cells at lower temperatures. According to the datasheet of the cells in Table 5.1, they experience 40% capacity degradation at -20 °C, which is in agreement with the previous section [127, 161]. Solid-state batteries also perform relatively well in extreme temperatures due to their solid electrolyte, and there are commercial options rated for -20 to 80 °C. However, at -20 °C, effective capacity is reduced to just 20% of original. For this solid-state battery, effective capacity actually increases at 80 °C to 120% [131]. No information is given on the effect of temperature on ESR and long-term cycle lifetime for these cells, but related work indicates these metrics would still be severely impacted by extreme temperatures. [85, 128].

Regardless of type, lithium batteries should be kept as close to room temperature as possible, and in extreme environments like high altitudes, power plants, outer space, or close to the earth's mantle, batteries will require active heating or cooling to maintain an operational lifetime. For example, the Mars Curiosity and Perseverance rovers utilize a Radioisotope Thermoelectric Generator (RTG) to both provide heating and constant thermocouple harvesting to charge a bank of two Li-ion batteries [164]. Battery packs with included (non-RTG) heaters are also commercially available for cubesat applications [165].

Batteries are not as robust as capacitors and supercapacitors when used in extreme temperatures, and will require temperature control. However, supercapacitors also require active temperature management to maintain long lifetimes and normal performance at high and low temperatures. When considering common use cases for wireless sensors, the efficient and operational range of both batteries and supercapacitors encompasses all applications located in spaces that humans occupy, and likely will handle most of the range of temperatures outdoors in many locations.

5.7 Safety

Besides temperature sensitivity, many types of lithium metal batteries are also notorious for their propensity to burn and explode under mechanical, electrical, and thermal stress. Capacitors and supercapacitors do not exhibit same caliber of danger, and at most will "pop" if exposed to electrical stress like overcharge or a large inverted charge. For both batteries

and supercapacitors, electrical stress is rare for low power energy harvesting applications, especially if using a battery protection and voltage management circuit. However, in some extreme environments, mechanical and thermal stress can be an issue. If left uncharged for a long period, battery undervoltage will occur.

Batteries with LCO and NMC cathodes and graphite anodes are prone to fiery explosions if mishandled. However, batteries that use other electrode materials, like LFP cathodes or LTO anodes, are considered much safer [137, 138, 140, 166]. Compared to LCO and NMC, LFP has enhanced safety and stability [166]. The structure of the LFP cathode is more thermally and chemically stable. The LFP cathode material has been shown to not release oxygen during thermal runaway, even when fully decomposed at high temperatures. The P-O bond within the PO_4^{3-} ion is stronger than that of the Co-O bond in the CoO_2^{-}, so that when abused, oxygen is released slowly or not at all [137]. This results in a significantly safer cell, where any thermal runaway is dominated by anode and electrolyte reactions at extreme temperatures [138].

LTO anodes provide similar benefits over graphite anodes. Graphite is prone to expansion and contraction during charging and discharging, respectively, which causes internal damage to the cell, resulting in capacity fading and in rare cases, short circuit conditions [141]. This effect is further exacerbated at extreme temperatures. Titanate oxide is considered a zero-strain material, meaning it experiences very little change in its chemical structure during charging, discharging, or temperature changes, unlike graphite. When heated to high temperatures, graphite can break down and react exothermically with the electrolyte, releasing flammable hydrocarbons. By contrast, at the same high temperatures, LTO anodes do not produce heat or release any gasses [140]. Beyond the anode chemical stability, the high potential of LTO anodes prevents lithium dendrites from forming upon deep cycling or after many cycles [137]. Lithium dendrites are the main cause for internal shorts in batteries with graphite anodes. These features of LTO anodes greatly lower the risk of internal short circuits, thermal runaway, and explosions. The solid electrolyte in solid-state batteries also substantially improves safety. A solid electrolyte is not flammable, and since it is solid, it is unable to interact explosively with a decaying cathode or anode [142].

While there is still slight potential for danger under abuse conditions for LTO, LFP, and solid-state batteries, the danger is substantially less than traditional graphite and cobalt-based lithium batteries. For LTO batteries, the danger of thermal runaway and explosion should be equivalent with that of a supercapacitor. Solid-state batteries should be even safer than supercapacitors, as EDLC supercapacitors employ an aqueous electrolyte. The battery manufacturer for the LTO and LFP cells listed in Table 5.1 states in the datasheet that no battery protection circuit is necessary to ensure safety [124, 126]. These cells also feature a capacitor-like notched cap that fails first in the case of internal pressure and gas release. Despite their safety, battery management is still important to limit damage to any battery and maintain capacity and performance. Not all lithium batteries are equivalent, and the claim that *all* batteries are dangerous disregards the many safety improvements made with newer electrode and electrolyte materials.

5.8 Summary

Among the various arguments that batteryless platform designers have levied against batteries, very few are relevant in the face of improved battery technology. Out of all addressed here, temperature sensitivity is the only real limitation of batteries. This complicates their usage in extreme environments, like space, but does not disqualify them in the vast majority of applications. Supercapacitors also exhibit poor performance and lifetime at extreme temperatures. Batteries cannot compete against some supercapacitors on power density and cycling performance, but for low power applications, these metrics are inconsequential. In reality, batteries are by far the most energy dense rechargeable storage device available and they provide many orders of magnitude more energy capacity per dollar than any capacitor or supercapacitor. New battery technology provides comparable efficiency, lifetime, and safety to many modern supercapacitors.

The qualitative arguments made by batteryless proponents are the same ones made in the 2010s, when miniature battery technology was poorly suited for long-lived, low power applications. Two decades later, these arguments have not been revisited or verified with quantitative evidence, and continue to find purchase to further justify batteryless design decisions even when the facts suggest otherwise. The choice to use a (super)capacitor has become the default, often selected to provide just enough energy storage to allow an application to function, but far from enough to function well, as Chapters 3 and 4 quantitatively illustrate. This lack of energy capacity has forced researchers to develop complex software and hardware solutions and shoehorn capacitors and supercapacitors into applications they are poorly suited for. There is a far superior and simpler solution for the majority of energy harvesting sensing applications: **use a battery**.

Chapter 6

Implementation and Evaluation of Capacity Sizing

This chapter considers two example sensing applications and the design and implementation of wireless sensor systems to achieve the goals of these applications. We explore the system design for these applications within the context of the conclusions of Chapters 3 and 4, which suggest that energy capacity is highly important for energy harvesting wireless sensor performance. The first application that we consider is the measurement of fine-grained workplane illuminance. Illuminance measurements at the desk or computer monitor can be used to inform the operation of a lighting control system to balance artificial and natural light for lighting efficiency and occupant comfort. The goals of this application are to measure illuminance at high granularity with high availability and provide a lifetime of at least a decade. This application is relatively simple and it is used to validate the simulation and conclusions presented in Chapters 3 and 4. The second application is image-based human occupancy detection and counting. Human occupancy measurement can help inform building energy management and climate control. The goal of this application is to accurately and consistently detect the presence of humans within view while providing a lifetime of several years. This application considers an existing batteryless image sensing solution and reconsiders their design in the context of the conclusions of previous chapters. We propose designs for each of these applications, and detail their implementation and evaluation. Utilizing a hybrid energy harvesting architecture with sufficient rechargeable energy buffer and a backup non-rechargeable energy storage, our proposed sensor designs are able to provide significantly higher availability than batteryless designs, as well as much longer lifetimes than designs that solely use non-rechargeable batteries.

6.1 Measuring Workplane Illuminance

The advent of LED lighting has significantly reduced electricity consumption in residential and commercial sectors. However, residential and commercial lighting still consumes 5% of the

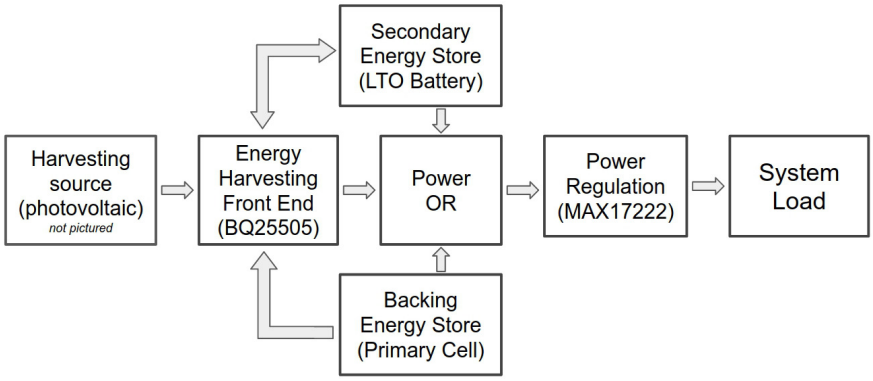

(a) Harvesting and storage architecture

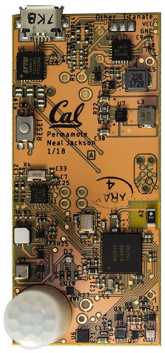

(b) Hardware

Figure 6.1: The Permamote power supply architecture is informed by the findings in Chapters 3 to 5. An LTO battery is recharged by a solar panel. When the battery is depleted, a primary-cell powers the system, providing reliability and avoiding intermittency.

total U.S. electrical consumption [167]. Beyond the utilization of LED lighting, the technique of daylighting, or lighting buildings with natural light, can further reduce the amount of electricity consumed by buildings. Since the intensity of daylight can be unpredictable as it depends on weather, it can be difficult to achieve consistent lighting with natural light alone. Artificial lighting can be used to augment insufficient natural light, but it requires fine-grained measurement to provide feedback to control and maintain a set point in a space. In particular, workplane illuminance for commercial buildings is important for occupant comfort and productivity, but is difficult to maintain with existing lighting control systems. Modern lighting control systems that perform both measurement and control are generally limited to large zones of measurement and influence. The cost of instrumentation and automation is often too exorbitant to justify fine-grained sensing. This often results in inequitable and sometimes uncomfortable lighting for occupants. Wireless sensing could provide a solution for fine-grained workplane measurement for daylighting applications, assuming the sensor does not require frequent maintenance and provides high availability and consistent measurements for the lighting system feedback loop. For this specific example, our application goals are to provide at least a ten year lifetime with consistently high availability. This section details the realization of a design to meet these requirements, and utilizes this design to verify the results of the simulation and design conclusions detailed in earlier chapters.

Design and Implementation

We design and implement a prototype sensor named Permamote to perform workplane illuminance sensing based on the application requirements described previously and the design principals discussed in Chapters 4 and 5. Permamote integrates a processor, BLE/802.15.4 radio, and various environmental, lighting, and a passive infrared (PIR) occupancy sensor. The

components used in Permamote are the same ones that we used to develop our representative hardware and workloads for our simulation. These components are listed in Table 4.2. A picture and system diagram of Permamote is shown in Figure 6.1. All hardware and software for the platform is open source[1].

Energy Harvesting and Storage.

Some of the primary goals of Permamote are to provide workplane illuminance measurements with high availability for a long lifetime of greater than ten years. Given the results of our simulation, a design that relies solely on energy preallocation is unlikely to achieve a sufficiently long lifetime. Permamote is powered by an energy harvesting front end that capitalizes on the benefits of rechargeable and non-rechargeable energy capacity. It utilizes the TI BQ25505 energy harvesting IC, which harvests energy while monitoring both rechargeable and backup energy stores, switching between them at user-configurable voltage thresholds [25].

We utilize the hueristics developed in Section 3.3 to determine the required sizing for our energy harvesting and energy storage. We select a 10.9 cm^2 amorphous silicon photovoltaic panel as our energy harvester to fit within our form factor. Assuming the lower end of efficiency (10%) and the upper bound of indoor irradiance, this panel can provide an average of 100 µW. On average, this income power provides more than sufficient average power to support various frequencies of the sense and send workload. We can consult Sections 3.3 and 4.3 to estimate the minimum required capacity for the sense and send workload. Considering the average power required by the sense and send workload with a 30 second period (24.5 µW), the minimum sufficient capacity should be 1.4×10^3 times the average workload power according to our capacity sizing heuristic. This is if we assume an income distribution similar to the EnHANTs Setup D and an income margin of 300% of the expected workload power[2]. Given this heuristic scaling factor, the capacity of the rechargeable energy storage should be on the order of 34 mWh. This agrees with results from our simulation in Figure 4.3 that suggests energy capacity on the order of 1–10 mWh will be sufficient to achieve near perfect availability in conditions like Setup D captures. The simulation suggests slightly less capacity is required compared to the heuristic estimate. This is because the simulations assume the energy buffer starts fully charged. Most systems, especially those that rely on rechargeable batteries, are deployed fully or partially charged. In our simulation, this provides an influx of energy into the system and it does not need to capture as much harvested energy to power its workload, and thus requires less energy capacity. If simulated for more time, over a longer period than the energy income trace provides, the capacity determined by simulation results may not be sufficient to continue an energy neutral operation. The heuristic for energy capacity sizing provides a safer, more conservative estimate.

Given this analysis, we select a 20 mAh (48 mWh) LTO battery to achieve this energy capacity [124, 126]. As described in Section 5.5, we configure the voltage thresholds of the BQ25505 to derate the usable capacity of this battery to increase the apparent cycle lifetime

[1] https://github.com/lab11/permamote/tree/master/hardware/permamote

[2] 100 µW represents a 300% margin over the 24.5 µW workload

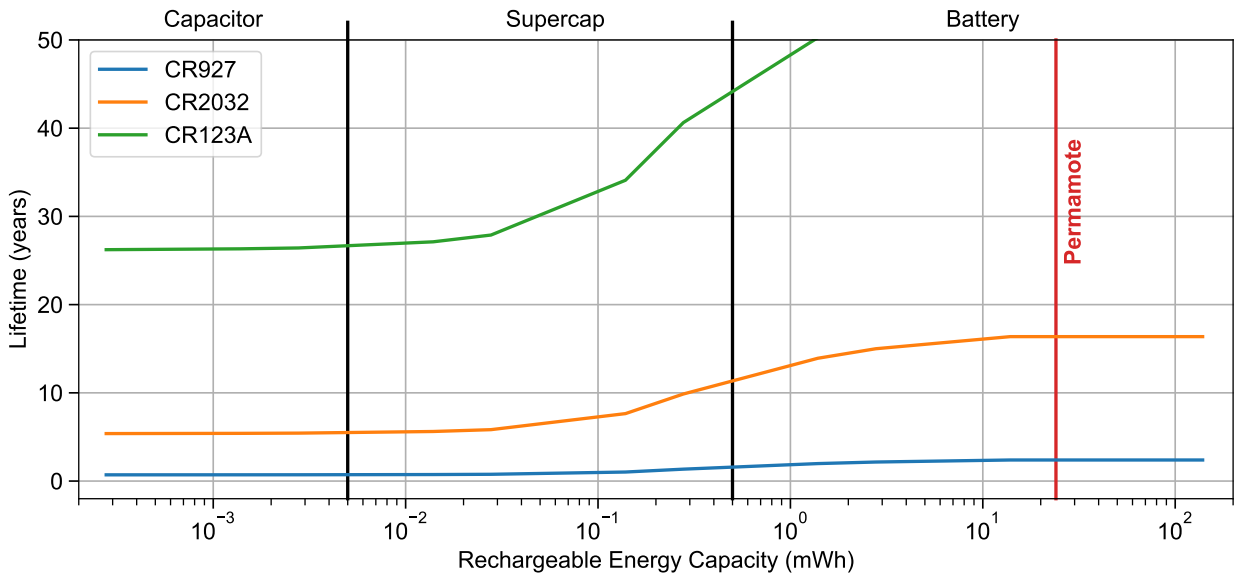

Figure 6.2: Lifetime estimation of the Permamote sense and send workload given different rechargeable buffer sizes and different primary cell sizes. This figure utilizes the low irradiance Setup A environment (15 µW/cm^2) and the 30 second sense and send workload period (24.5 µW). This is a different presentation of Figure 4.6 that identifies the rechargeable capacity of Permamote with a red vertical line.

of the battery. The resulting energy storage provides 24 mWh of energy storage, more than the capacity required to achieve the reliability and energy utilization improvements of the workloads that were simulated in Section 4.4 and Figure 4.6.

In some cases, the available harvestable power may be lower than expected and insufficient to operate in an energy neutral fashion. This justifies the addition of a reliable, non-rechargeable backup energy source. Figure 6.2 is a recreation of Figure 4.6 presenting the estimated lifetime of Permamote with its configured 24 µWh rechargeable storage and various non-rechargeable backup energy storage options. We assume a 30 second sense and send workload period. From this simulation result, Permamote requires at least one CR2032 coin cell to achieve the application lifetime goals in the worst case energy harvesting potential. Thus, Permamote is designed to be configured with either one or two CR2032 coin cells or a CR123A cell. The output of the active battery is boosted by a MAX17222 regulator, which features high conversion efficiency (>90%) at low output currents and operates down to 400 mV [143].

Processor, Radio and Sensor Selection

The components used in the implementation of Permamote are the same that were benchmarked for our simulation configuration in Chapter 4. These components are summarized in Table 4.2. We note our choice of the Nordic NRF52840 MCU over the more commonly

used MSP430FR series because of its higher active power efficiency while offering comparable sleep currents. Specifically, it only draws 56 µA/MHz compared to over 100 uA/MHz for the MSP430. The MSP430 is a common choice for batteryless systems because of its integrated non-volatile FRAM. Unlike batteryless systems, Permamote is designed with sufficient rechargeable capacity and backup energy and is intended to never power off and lose volatile state. This eliminates any reliance on state retention techniques and the need for the non-volatile FRAM present on the MSP430FR series chips. While slightly more efficient processors and radios exist than those found in the NRF52840, we value the simplicity of the SoC-based design as well the platform's capability of using either BLE or 802.15.4 over its 2.4 GHz radio.

Simulation Evaluation using Real Systems

To evaluate our simulation and the benefits of a capacity-focused design, we perform a three-month-long deployment in a partially sunlit room using i) a primary-cell only system [97], ii) a batteryless, capacitor-only system [63], and iii) Permamote, our system that features both a secondary and primary-cell. We model these systems over the same period using estimated irradiance from Permamote illuminance masurements. We compare the performance and lifetime of these three systems with the predictions generated by a simulation of their workload. We also compare the availability of Permamote to the batteryless system.

We analyze the deployment of the systems and compare their behavior to our model's predictions: i) ten CR2032 primary-cell (720 mWh) only devices, ii) an batteryless system configured with just 500µF of capacitance (about 0.36 µWh at 2.2 V), and iii) Permamote, configured with a 20 mAh (48 mWh) secondary-cell, half of which is usable, and a CR2032 backup. The primary-cell only device performs environmental sensing over BLE every second. The batteryless system sends a beacon as soon as its capacitor bank is full. When its energy is depleted, it powers off and charges again. Permamote is running the "sense and send" workload that we described in Section 4.1, and sends illuminance measurements every second. This workload stresses the model and requires more charge and discharge cycles. We use Permamote illuminance readings to estimate irradiance using the same scaling factors used by Yerva et al. [2], and use these traces as model input for the energy harvesting sensors.

Primary-Cell Only

We measure and model the workload of the primary-cell system and produce estimates for lifetime. The primary cell system requires on average 480 µW to measure and beacon every second. Our model predicts the platform lifetime to be 58 days. The average lifetime of the devices in our ten device deployment is 61 days, which is on par with our simulation estimate.

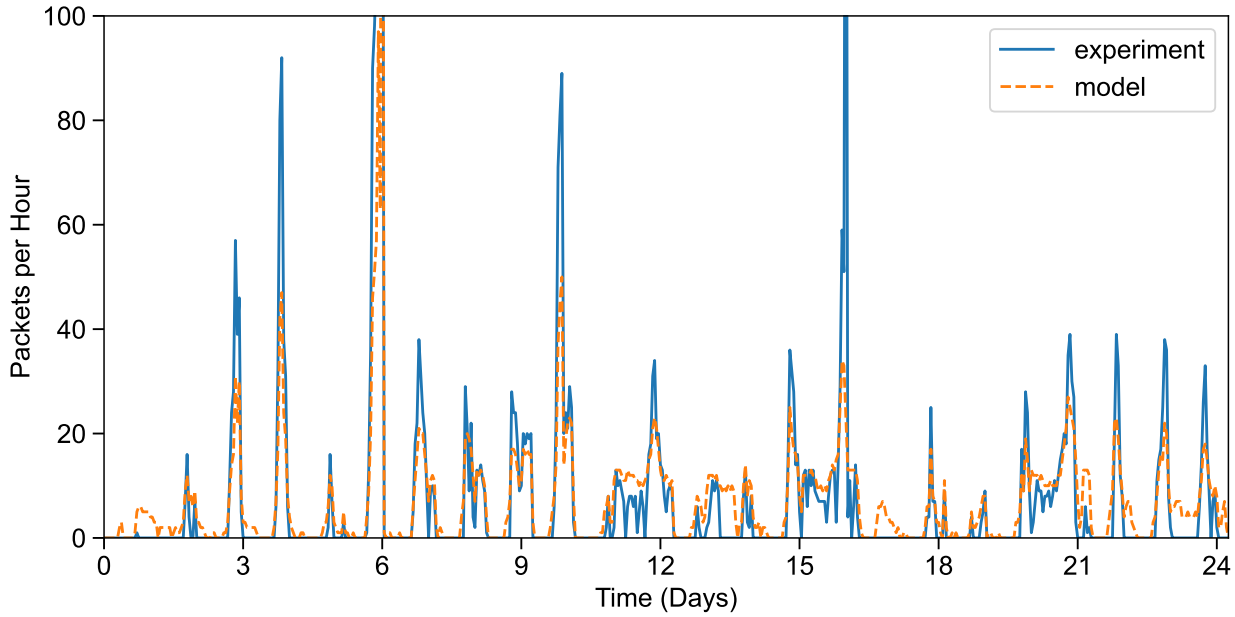

Figure 6.3: Performance comparison of model expectation versus real batteryless system. Data from a three month deployment of two systems is used to verify our model. We use three weeks of illuminance measurements to estimate irradiance and model the number of packets transmitted by an batteryless node. Average daily error is 15%, with a standard deviation of 17%.

Batteryless

We model the number of packets sent each hour by the batteryless system over a three week period, and compare against the results of an actual device in Section 6.1. Like the simulation of the primary cell system, the model also predicts a more conservative result for the batteryless system. The simulation predicts fewer successful packets sent compared to the actual batteryless system under test. The average daily error of the model versus our results is 15%, with a standard deviation of 17%. This error can attributed to two primary sources. Illuminance is measured close to, but not exactly at the solar panel of the test device. Occasional direct sunbeams, like that experienced on day 16, can illuminate the solar panel but not the sensor, or vice versa. This results in a substantial over or underestimate of available light. In addition to inaccurate light measurements, we introduce error through our estimation of irradiance. We measure illuminance instead of irradiance, and must resort to a piecewise linear estimation, when in reality the relationship is not well defined and non-linear when considering different light sources [168]. In the case of our estimation, results indicate that the model consistently underestimates high irradiance measurements.

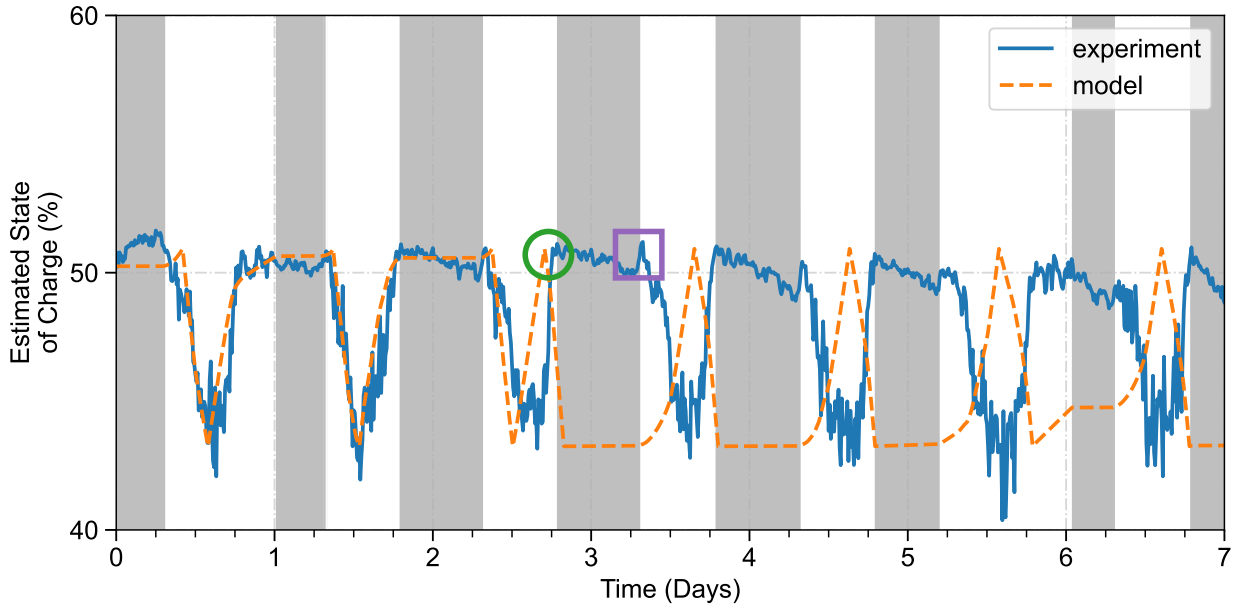

Figure 6.4: The estimated state of charge of a week of a workload compared to our model's estimation. We use three weeks of illuminance measurements to estimate irradiance and model the number of packets transmitted by an batteryless node. Average daily error is 15%, with a standard deviation of 17%. (b) We model and measure a Permamote's state of charge while running a "sense and send" workload with a 1 s period for a week, beginning at midnight on the first day. Charging hysteresis limits of the devices are set at 51% and 43%. Shaded regions represent periods of low harvestable potential ($< 15\,\mu W/cm^2$), i.e. nighttime. For the first two days, model predictions closely track the experimental measurements. Errors in hysteresis and irradiance estimation cause the model to reach its upper hysteresis sooner than the experiment does, annotated by the **green circle**. In actuality, the device exits charging hysteresis at the peak marked with the **purple square**. More importantly, the frequency and length of periods spent using harvested energy collected in the secondary-cell (downward slopes) are identical.

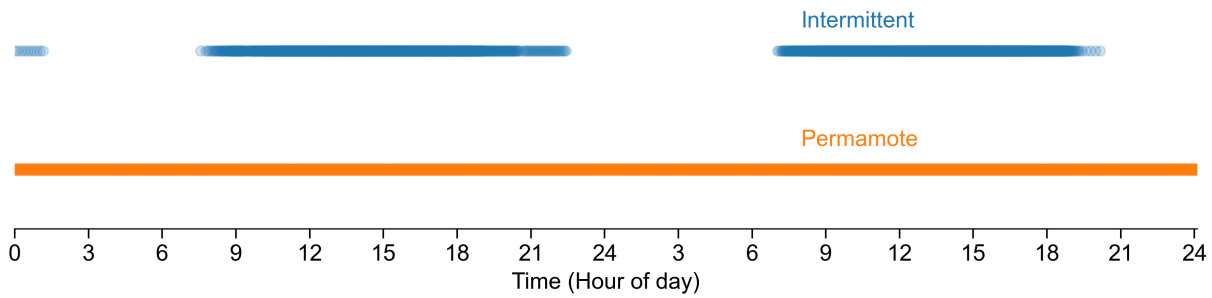

Figure 6.5: Packets received over two days. This figure compares the availability of an batteryless design and Permamote. Permamote sends a packet every second and does so without fail, while the batteryless system is only able to send when light is available.

Secondary and Primary-Cell

We compare our model's predicted state of charge to a deployed Permamote over a seven day period in Section 6.1. We estimate state of charge from the reported secondary-cell voltage, and irradiance from lux measurements. In this figure, the state of charge cycles between configured battery hysteresis limits, as the workload is too intense to be sustained by energy harvesting alone. Flat and upward slopes of the curve represent the device in hysteresis, using the primary battery to perform its workload. Upper slopes indicate the secondary cell is charging from harvested energy. Downward slopes indicate the device is out of hysteresis and is using harvested energy stored in its secondary battery to perform its workload. The shaded "nighttime" regions are not uniform, as the deployment environment is occupied by graduate students that occasionally work late hours or forget to turn off the lights. The model correctly predicts the cycling behavior of the deployed device for two days, but deviates during the third day. The model predicts that the device would charge above the upper hysteresis limit and begin supplying energy from the secondary-cell before the test device actually does. This inaccuracy, like that of the last of experiment, is partially due to our inexact estimation of irradiance. In addition, real device hysteresis limits are set using resistor networks. The resistors used have 1-5% tolerance, and are susceptible to temperature changes, which introduces dynamic errors that is not accounted for in our model. Even though the predicted state of charge deviates after two days, the length and frequency of periods in which harvested energy is stored and used are identical to our experimental measurements. The amount of energy that is charged and discharged from the secondary cell is also identical.

Permamote Performance and Lifetime

We also compare the performance of the deployed batteryless system and Permamote. In Figure 6.5, we show the number of packets sent per hour for two days. Permamote sends data

Platform	Successful Events (%)		Long-Running Time to Completion Ratio		Lifetime (yrs)
	Periodic	Reactive	Average	95th Percentile	
Telos [28]	100	100	1	1	8.55
Hamilton [103]	100	100	1	1	6.75
BLEES [97]	100	100	1	1	1.11
Gecko [2]	39.5	64.9	387	981	∞ [g]
Capybara [18] [a]	46.3	72.8	37.6	1	∞ [g]
Capybara [18] [b]	41.1	67.1	2730	8900	∞ [g]
Flicker [7]	39.3	64.2	1307	5670	∞ [g]
EnHANTs [52]	79.4	96.0	1	1	— [h]
DoubleDip [53]	77.9	66.5	1	1	— [h]
[169]	78.4	66.9	1	1	— [h]
Permamote [c]	81.2	98.3	1	1	— [i]
Permamote [d]	100	100	1	1	35.8
Permamote [e]	100	100	1	1	30.2
Permamote [f]	100	100	1	1	6.27

[a] With capacitors: 400 µF ceramic + 330 µF tantalum + 67.5 mF supercapacitor.
[b] With capacitors: 300 µF ceramic + 1100 µF tantalum + 7.5 mF supercapacitor.
[c] No primary-cell. [d] AA primary-cells like Telos. [e] CR123A primary-cell like Hamilton.
[f] CR2032 like BLEES. [g] Lifetimes are theoretically infinite for capacitor-based systems.
[h] Not enough information to predict cycling failure time for theses systems.
[i] Expect cycling degradation in 20-50 years, but do not attempt to estimate.

Table 6.1: Simulated performance of energy-harvesting systems performing the same workloads. For each platform considered, we model the performance of its energy storage architecture. Periodic workload and lifetime estimates are based on a 10 s period, and the reactive workload is scaled to generate a maximum of 2000 events per hour (3.4 s average daily period). Permamote is the only energy harvesting platform that can provide 100% availability, while also offering a lifetime of more than triple that of similar battery-only platforms.

every second, while the batteryless system sends as fast as possible. Permamote is able to collect and send its data continuously, while the batteryless system is limited to sending only during the day. This demonstrates the increased availability afforded by increasing secondary capacity and including a backup energy store.

We also use our model to explore the estimated performance of Permamote compared to the power supply architectures of historical systems for workloads including our illuminance sense and send application and the other applications defined in Section 4.1. To isolate the analysis to just power supply types and sizing, we assume each system uses the same low-power hardware and is performing the same workload as Permamote. The results of this modeling are shown in Table 6.1. Our model estimates that Permamote can expect several decades of 100% reliable lifetime when configured as it was deployed for this evaluation, albeit configured with the less intense 10 or 30 second periodic workload. For some workloads, Permamote can provide over double the availability of similar batteryless platforms and more than triple the lifetime of battery-based architectures.

6.2 Image-based Occupancy Detection

Besides the utilization of LED lighting and daylighting, another way to increase lighting efficiency in buildings is to automate lighting power states based on human occupancy. Human occupancy is also a useful metric for regulating building heating and cooling. The presence and number of people in a space directly effects the amount of heating, ventilation, and air conditioning (HVAC) a building must perform to maintain a temperature set point [170]. Compared to lighting, building HVAC is much more energy intensive, with cooling consisting of 10% of the total U.S. electricity consumption in 2021 [167]. Occupancy is measured in lighting control systems, however they often rely on a binary indication of occupancy based on simple ultrasonic or PIR motion detection [171]. This technique cannot quantify the occupancy of a space, and it is prone to false positives and negatives. Non-human objects that move through a space can trigger an ultrasonic sensor, while any rapid changes in surface temperatures can result in false positives for a PIR sensor. When a person is not moving sufficiently, both ultrasonic and PIR sensors may sense a false negative. Instead of, or in addition to these binary occupancy measurements, image sensing can be used to more accurately capture the occupancy, including person count, of a space. Image inference, including classification and object detection, has been one of the most active areas of modern computer science and machine learning research. However, due to the cost and difficulty of deploying wired cameras, long-lived applications based on continuous image sensing has traditionally been untenable. This is especially true for indoor building-centric applications, where camera density must be higher for sufficient coverage and lifetime must be sufficient to avoid frequent maintenance.

Related Work

Indoor wireless camera sensors have been heavily researched and commercialized over the past fifteen years, but due to the technology available at the time, as well as incompatibilities between design decisions and longevity, these platforms are typically limited to lifetimes of at most weeks to months [172–176]. Deployment of these platforms beyond small or temporary installations remain a challenge due to the cost of frequent battery replacement.

Modern image sensing platforms have taken advantage of technology improvements, as well as employed new techniques to reduce transmit power and increase (or abolish) lifetime. WISPCam [177], is a battery-free camera that utilizes RFID for power and backscatter communication. It can capture an image every 15 minutes when an RFID reader is 5 meters away. BackCam [176] is another camera platform that utilizes backscatter for communication, but over commodity WiFi. BackCam can stream video at 4 frames per second for a lifetime of 32 days. Similar to WISPCam, BackCam also requires a nearby wall-powered transmitter to generate excitation packets for backscatter communication. Camaroptera [56] is another batteryless platform, focused on wide-area image sensing. The platform uses a long-range LoRa radio and harvests energy from solar panels. It performs local inference to detect people within captured photos and can achieve an end-to-end latency of less than 20 seconds in well-lit outdoor environments. Commercial platforms like the Blink Indoor wireless security camera [174] and the Wyze outdoor camera [175] utilize low power motion detection to minimize energy usage. The Blink Indoor claims a 2 year lifetime consisting of 40000 seconds of recording 720p video on two AA batteries. Wyze estimates a lifetime of three to six months on its internal 5.2 Ah battery if capturing 10 to 20 video clips per day.

While these modern platforms push the envelop for video streaming or batteryless imaging, they have drawbacks that limit their deployability in indoor spaces. Backscatter-based systems require a nearby carrier transmitter to communicate. Even with a dedicated, nearby transmitter, WISPCam can only periodically capture images every 15 minutes. Camaroptera is designed for outdoor use, and it is unclear if it can capture and send images indoors, especially as its energy harvesting system requires 197 µW when idle. The Blink camera lasts two years if it is placed in locations infrequently occupied by people. Placing the camera in a kitchen that sees an average 26% occupancy on weekdays leads to a lifetime of only 1.78 days [176]. The Wyze camera faces a similar fate if placed in a frequently occupied area.

Image sensing platforms developed by industry and in research exhibit the same design disconnect that we noted in Chapter 2. Commercial image sensing systems largely utilize batteries and avoid energy harvesting, while modern research systems are largely energy harvesting and batteryless. These two different design result in a deceptive choice between reliable operation with a preallocated energy source but a limited lifetime, or an unlimited but unreliable lifetime with batteryless energy harvesting. Neither option is satisfactory for many applications, including our indoor human occupancy detection example application that simultaneously demands a long sensor lifetime and high availability.

The Performance of a Batteryless Image Sensor

In this dissertation, we have argued that a batteryless approach is not appropriate for applications with QoS requirements. Designers of batteryless systems rarely evaluate their platforms and frameworks in regards to availability and reliability. However, utilizing the heuristics and simulation tools we have developed in Chapters 3 and 4, we can estimate the performance of these systems and identify alternate design decisions that may result in a more performant image sensing system.

In this section, we consider the design and performance of Camaroptera, a batteryless image sensor [13, 56]. We capture the design details of Camaroptera and use them to simulate the system and examine its performance. Camaroptera is designed as a primarily outdoor system, and utilizes LoRa for communication and image backhaul. It performs local inference to detect people in captured frames. A positive detection results in transmitting an image, and if no person is detected, it discards the captured image to preserve energy. It harvests solar energy with four small high-efficiency (20%) monocrystaline solar panels (6.2 cm^2 total area). Camaroptera stores harvested energy in a 33 mF supercapacitor. Camaroptera, like many batteryless systems, operates opportunistically. Whenever its capacitor is charged sufficiently, it turns on and performs its workload. Camaroptera's workload consists of capturing an image, performing some computations and inference on this image, compressing the image, and finally transmitting the packetized image. Camaroptera's energy storage is sized to support the most energy intensive operation in its workload: transmitting a single LoRa packet. A single packet cannot fit an entire compressed image, so Camaroptera relies on multiple power cycles to send between seven and eight packets to transmit an entire image. All combined, a single image capture and transmission requires 781 mJ[3]. When idle, Camaroptera's power subsystem consumes 197 µW. With this information, along with operation timing provided by the authors [13], we are able to complete a simulation configuration like Table 4.3 for Camaroptera's hardware design and workload.

For energy income, we synthesize an outdoor irradiance trace based on the EnHANTs Setup D trace. The authors evaluated Camaroptera from 5–95 klx, which roughly corresponds to 4–77 mW/cm^2 for natural light[4]. We scale the mean of the synthesized outdoor trace to represent 10 and 50 mW/cm^2 for an estimate of outdoor irradiance on the same scale. The simulation results of Camaroptera's average packet distribution under these two harvesting conditions are summarized by Figure 6.6. With either income, Camaroptera is able to maintain high image capture and transmit rates when light is available. However, as expected of a batteryless system, when light is unavailable, Camaroptera is unable to capture and send many images. Since Camaroptera performs its workload in response to its energy storage state, it cannot maintain a periodic schedule or react and detect the presence of people.

This limitation is primarily due to the energy capacity available to the platform. The available energy in an outdoor setting is essentially limitless to a low power embedded system like Camaroptera. For the two simulations we consider, our simulated Camaroptera with its

[3] 3.06 (capture) + 0.253 (difference) + 66 (inference) + 40 (JPEG) + 96×7 (transmission) = 781 mJ [13]

[4] For natural sunlight: 1 W/m^2 ≈ 122 lx [168]

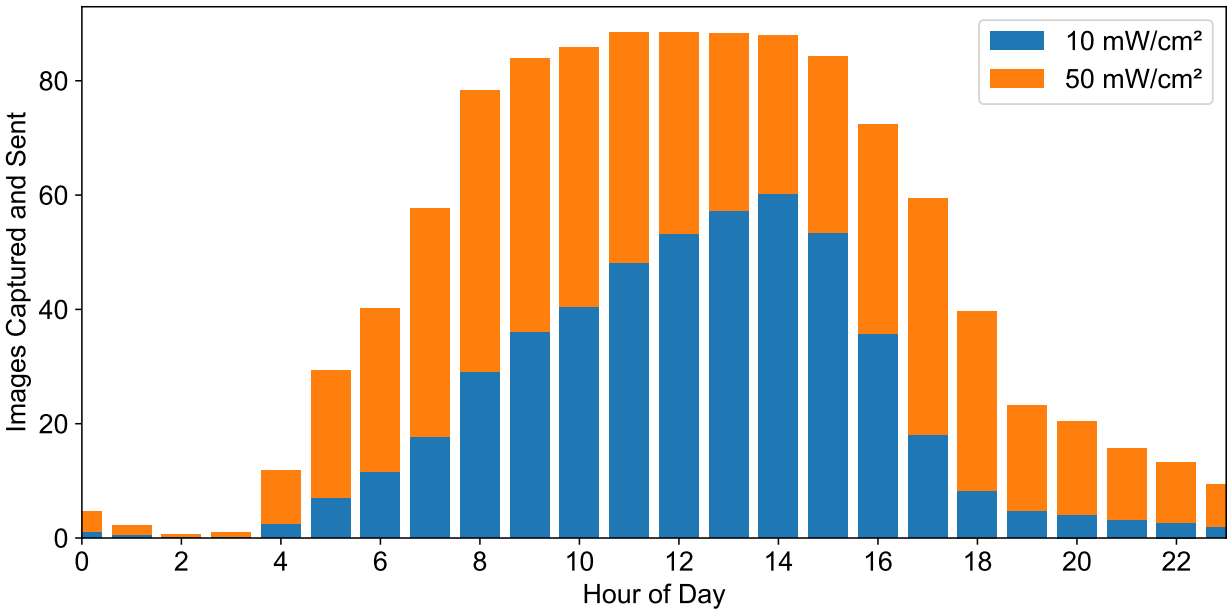

Figure 6.6: The distribution of simulated Camaroptera transmitted image packets per hour in a day. The distribution represents an average over the length of two synthesized outdoor traces (10 and 50 mW/cm^2) based on the EnHANTs Setup D trace. Camaroptera operation is limited to times when daylight is available, regardless of the scale of average input power. The average number of packets is significantly lower between 6PM and 6AM.

capacitor energy buffer is only able to capture 48% and 21% of the available energy for the 10 and 50 mW/cm^2 traces, respectively. Due to its designed energy capacity, Camaroptera is simultaneously limited in the amount of energy it can capture and the amount of instantaneous energy available to it at any given time. Given the design heuristics developed in Chapter 3, we can determine an estimate for Camaroptera's average income and average workload power if we assume a fixed sensing period instead of the opportunistic operation of the actual platform. From these averages and using the capacity sizing heuristics from Section 3.3, we can determine an estimate for the minimum sufficient capacity to sustain this workload given our synthesized incomes. As mentioned earlier, each Camaroptera activation requires 781 mJ over 40 seconds, assuming all images go through the entire inference pipeline. The platform's idle power is 197 µW. While the platform is capable of transmitting an image every 45 seconds to two minutes depending on lighting conditions, we assume a relaxed periodic rate of capturing and transmitting an image every five minutes. This results in a 2.8 mW average workload power[5].

From this example workload and Camaroptera's income, we can use the huristics developed in Chapter 3 to determine the proper capacity sizing for the platform. With Camaroptera's designed solar panel size, maximum power point voltage, and efficiency, it can expect between

[5] $\frac{781\text{mJ}+(197\,\mu\text{W}\times(300-40)\text{s})}{300\text{s}} = 2.77\text{mW}$

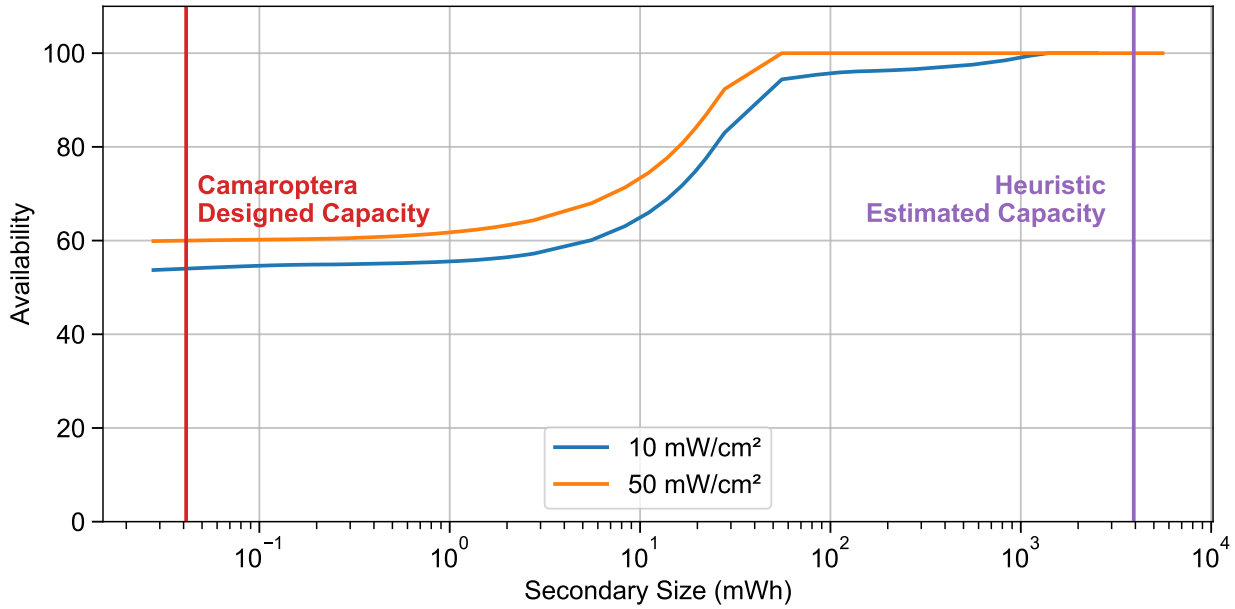

Figure 6.7: The availability of Camaroptera running a five minute sense and send workload, as energy capacity is increased from that offered by its original 33 mF supercapacitor to the minimum sufficient capacity estimated by the heuristics developed in Section 3.3. Capacity must be increased by five orders of magnitude in order to achieve near 100% availability.

12.3 and 61.6 mW average income power from our two synthesized traces. On average, this theoretical energy income should be more than sufficient to power this periodic sense and send workload, assuming the platform has enough capacity to capture it. The lower end of this income corresponds to an income margin of over 300%. Considering the sizing factor for income resembling the distribution of Setup D, determined in Section 3.3, we should expect a minimum sufficient capacity that is 1.4×10^3 times Camaroptera's average workload power. This corresponds to an energy capacity of 3920 mWh. The 33 mF supercapacitor on Camaroptera only provides 41 µWh. This capacity represents five orders of magnitude less than the minimum sufficient energy capacity predicted by our heuristic. To further explore the impact of capacity sizing for Camaroptera, we simulate the platform with a sweep of different amounts of rechargeable capacity. We change our simulated Camaroptera's workload behavior from an opportunistic strategy to the periodic schedule discussed above. We sweep capacity from the size of Camaroptera's original energy capacity to the amount predicted by our sizing heuristic. The results of the simulated platform performance versus capacity are presented in Figure 6.7. For each successive simulation that increases energy capacity, the periodic Camaroptera workload is able to achieve higher availability. With an energy capacity above 100 mW h, simulations with the 50 mW/cm^2 income are able to achieve near 100% availability, while it requires at least 1000 mW h for the 10 mW/cm^2 simulations to achieve similar levels of availability. Both of these simulation results suggest less energy capacity is

required than the previously estimated 3920 mWh. This is for the same reason as explained in Section 6.1: the simulation begins with a partially charged energy buffer to mimic how a real system would be deployed while the heuristic analysis assumes starting from empty. Thus, the heuristic provides a more conservative estimate that is more likely to result in energy neutral operation over longer periods of time.

Camaroptera was implemented with an energy capacity that was sized to support its most energy intensive task. This arbitrary design decision results in a minimally feasible design that is fails to fully capture and utilize the available harvestable energy. With significantly more energy capacity, Camaroptera would be able to capture more energy, persist through periods of no harvesting potential, and provide significantly higher availability. Based on the capacity requirements determined by this analysis, Camaroptera could be redesigned with a medium-sized rechargeable battery, on the order of 1000 Ah, to achieve the higher performance estimated in Figure 6.7.

The Design of an Indoor Wireless Image Sensor

Due to Camaroptera's relatively high active and idle power requirements, it is unsuitable for use in many indoor environments. Designing an image sensor for indoor use will require different design decisions to achieve lower power operation. An indoor sensor does not need to depend on higher power wide-area communications like LoRa. Instead, with modest infrastructure investment, an indoor platform can rely on low power personal- or local-area networks. Besides data transmission efficiency, power efficiency can be improved through the use of a more efficient and capable processor. Idle power requirements can be reduced through the elimination of Camaroptera's hysteresis circuitry, and instead utilizing simple power gating to reduce component quiescent power. Using the heuristics and simulation tools developed, we can also identify the appropriate size of our energy harvesting, buffer capacity, and non-rechargeable storage to maximize energy available to our application. The next few sections expand on these design changes for indoor image sensor we name Permacam.

Indoor Wireless

We design Permacam to be untethered from wired communication and power. There are several options we consider for Permacam's networking, including WiFi, and personal area networks like BLE and 802.15.4. WiFi is attractive as it supports high bandwidth, which would be appropriate for transmitting large data like images. However, many low power WiFi SoCs like the ESP32 have significant idle and startup power requirements [178]. If included on Permacam, a WiFi radio or SoC must act as an external, power-gated component to reduce idle current. This complicates system design, and it is unclear if WiFi start up cost is worth the bandwidth advantage. Instead, we focus on technologies like BLE and networks built on top of 802.15.4. These networks are specifically designed to cater to devices that spend the majority of their lifetimes in an ultra low power sleep state. We trade off bandwidth and the time to send images for a simpler design and lower overall power.

Processor Selection

Processing and manipulating images requires significant computing and memory resources when compared to sensor data with lesser dimensionality, like periodic illuminance and color sensor data that is collected by Permamote. Camaroptera's 16–bit, 16 MHz processor requires 25 seconds of constant computation to compress a 160×120 JPEG image with floating point emulation, and seven seconds to compress the same image with a fixed point algorithm. This represents a significant amount of time for a low power device to be active and continuously computing. The Camaroptera design is severely limited in digital signal processing by its processor selection. For Permacam, we consider more capable, faster, processors with a 32–bit path and floating point support.

Power Supply

Like Permamote, to support fully wireless operation for a long lifetime, Permacam optimizes energy harvesting and storage through the use of rechargeable and non-rechargeable batteries. As we have explored previously in our analysis of Camaroptera, vision applications require significant energy to capture images, process them, and transmit them. We size Permacam's battery in the same manner as Permamote, utilizing the heuristics and simulation tool from Chapters 3 and 4 to determine a sufficient capacity. To determine this, we take benchmarks of capturing and transmitting images over an 802.15.4 network, and utilize these benchmarks to estimate an average power. We utilize the same indoor EnHANTs irradiance traces that we considered previously [99].

Indoor energy harvesting is inconsistent and variable, and Permacam will operate close to the edge of harvestable power that is reasonably available in many indoor environments. To augment harvested energy, we include a backup nonrechargable battery on Permacam. A backup battery allows continuous operation regardless of energy harvesting conditions. It safeguards the system from the volatile nature of energy harvesting, at the cost of a finite, but very long, lifetime.

Other Considerations

Permacam includes additional sensors besides a camera to limit duty cycle and minimize idle power. A PIR sensor is used to sense motion, providing an ultra low power wake up mechanism for capturing images. Performing imaging on event detection instead of periodically can save considerable amounts of energy and extend lifetime. The platform also has a light sensor that can interrupt on large changes in light illuminance. This provides another wake up mechanism, allows Permacam to detect when it is too dark to capture a useful image, and enables illuminance calibration for captured images. Both PIR and illuminance sensors require many orders of magnitude less energy than an image sensor to sense motion or illuminance.

To address the long-term relevancy of the platform, we also support an over-the-air update mechanism such that Permacam can receive regular updates to improve energy management. Additionally, updates can be used to deploy new inference algorithms and models, as well as

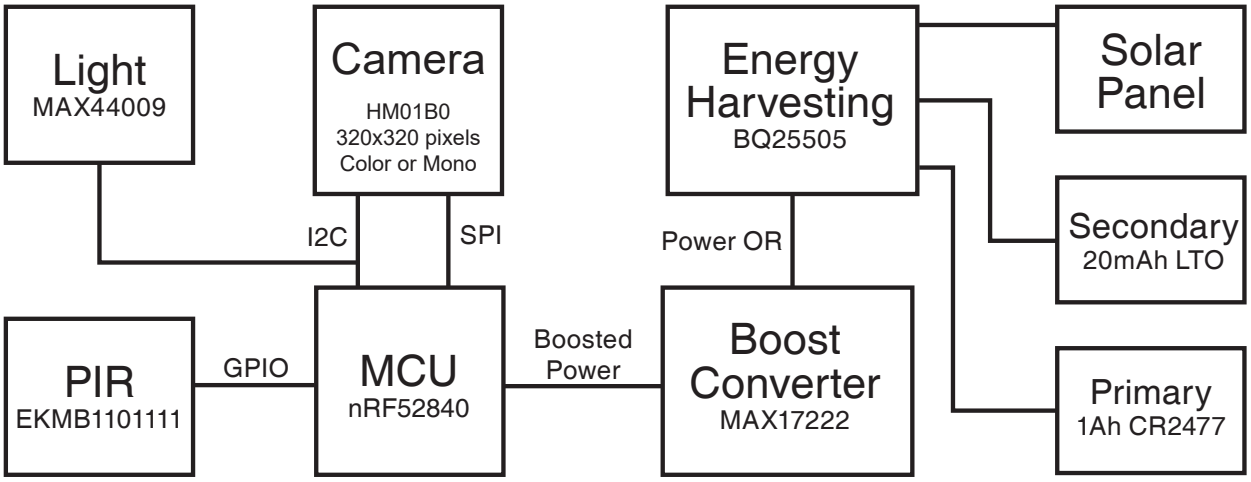

Figure 6.8: Permamote system block diagram. The system is based on the Himax HM01B0 camera and the Nordic NRF52840 MCU. We include a light and PIR sensor to provide a low power wake up mechanism to drive image capture. A hierarchical energy harvesting system with a rechargeable and non-rechargeable battery are utilized to provide a long, reliable lifetime to the system.

change image capture workload behavior. The platform can be configured to periodically take pictures, use the low power wake up mechanisms mentioned previously, or use some other software heuristic to only send interesting images.

Permacam Implementation

In this section, we explore the implementation of Permacam, including component selection, techniques to enable a low power camera interface, and architectural details about the energy harvesting and storage architecture

We implement the design of Permacam in an **open source** hardware[6] and software[7] implementation. The platform is built around the Himax HM01B0 ultra low power camera [179]. This camera has also enabled other platforms like BackCam [176] and Camaroptera [56]. The Himax camera is available in two different versions: monochrome and color. We use the color version of the sensor to ensure the ability to capture color information of images, but the two versions are pin compatible and can be interchanged. The other key component of the design is the SoC processor. For a processor, we select the Nordic nRF52840 Cortex-M4F SoC as it is one of the most power efficient processors available, includes a Floating Point Unit (FPU), and is relatively fast compared to other low power microcontrollers at 64 MHz. An FPU allows the platform to perform floating point dependent tasks like image compression.

[6] https://github.com/lab11/permamote/tree/master/hardware/permacam/rev_a
[7] https://github.com/lab11/permamote/tree/master/software/apps/permacam/camera_coap

Component	Function	Active Current	Idle Current
Himax HM01B0	Image sensor	2.0 mA	25.1 µA [a]
Nordic NRF52840	Processor	52 µA/MHz	3.16 µA [b]
	Radio	4.8 mA @ 0 dbm	— [b]
Ambiq AB1815-T3	Real time clock	55 nA	N/A [c]
Maxim MAX44009	Light sensor	650 nA	N/A [c]
Panasonic EKMB11011	PIR Occupancy	100 µA	1 uA

[a] Power gated when not in use.
[b] Sleep current for both processor and radio, full RAM retention, wake on low freq. timer.
[c] No shutdown or idle mode.

Table 6.2: The components used in Permacam, many of which are shared by Permamote. They represent some of the lowest power options currently available. Due to the extremely low idle power of all included components, Permacam is able to sleep at 4.4µA.

It has a multiprotocol BLE/802.15.4 radio that is well supported by OpenThread, an open source implementation of the Thread 6LoWPAN mesh network protocol [180]. It features comparatively large amounts of flash (1 MB) and SRAM (256 MB), which is enough to capture images and perform some local processing. We include an ultra-low power real time clock (RTC) to provide accurate image timestamps. The platform also features an ultra low power PIR sensor and illuminance sensor to provide a low power wakeup mechanism and illuminance calibration for captured images. A block diagram of the major system components is displayed in Figure 6.8. Characteristics of the sensors and the MCU are summarized in Table 6.2. In the next few sections we describe Permacam's implementation, including the MCU to camera interface, energy harvesting architecture, image compression methods, the end-to-end image transmission architecture, and an implementation of onboard person classification.

Camera Interface

The majority of image sensors do not support normal sensor interfaces like I2C or SPI to transfer data. Instead they rely on high frequency ($\geq$8 MHz) serial or parallel data busses. This is because images are many kilobytes or megabytes in size, and data transfer must be at a high frequency or parallel to transfer multiple frames per second. The majority of embedded processors, the nRF52840 included, are not designed to interface with parallel camera busses. Traditionally, platforms have used intermediate hardware like dedicated processors, CPLDs or FPGAs to interface with image sensors and buffer frames [172, 173]. The Himax HM01B0 has two interfaces, an I2C command interface and a serial, 4x, or 8x data interface. The Himax HM01B0 requires an input master clock (MCLK) of 3-36 MHz to drive internal sensor timings and outputs a pixel clock (PCLKO), frame valid (FVLD), and line valid (LVLD) signals [179]. It can be configured to capture full frame (320x320), QVGA (320x240), and if the mono version of the sensor, downscaled QQVGA (160x120) resolution images. In Permacam we capture full frame images and use the color version of the sensor. This version produces color filter array (CFA) images. A CFA is an alternating "mosaic" of green-blue-green and red-green-red rows of pixels [181]. Each pixel corresponds to a single

Figure 6.9: An image from Permacam, displayed as a mosaiced image (left) and demosaiced (right). Permacam is capable of capturing 320x320 resolution color images, which can be directly fed to machine learning inference tools for object, person, or in this case: dog detection.

color, and through post-processing known as "demosaicing", a three channel RGB image is generated from the single channel representation. An example of a mosaiced image and its demosaiced counterpart are displayed in Figure 6.9. Permacam's processor does not possess enough memory to locally demosaic an full image so any inference or transmission must be done with the mosaiced image. We power-gate the HM01B0 camera on Permacam, allowing the system to completely power the camera off and enter the lowest possible power state.

Other systems that utilize the Himax HM01B0 camera including BackCam, Camaroptera, and the Sparkfun Edge have developed a bit bang protocol to emulate a parallel interface [56, 176, 182]. This is possible because their processors support GPIO direct memory access (DMA). This allows the GPIO peripheral to directly write to memory, avoiding costly processor cycles to load GPIO state into registers and then write it to memory. This optimization allows a bit-bang protocol on these platforms to operate faster than the HM01B0's minimum operating frequency. However, a bit-bang protocol is undesirable because it requires the processor to be on during an image transfer, consuming additional energy. The nRF52840 does not support GPIO DMA. Because of this, developing a bit-bang protocol that adheres to the required camera timing is infeasible. Instead, we note that the camera serial protocol closely resembles SPI, and the nRF52840 implements an 8 MHz SPI peripheral with DMA support. The camera PCLKO signal is identical to SCLK, and can be gated by the LVLD signal. The LVLD signal from the camera indicates when a line of the image is being transmitted, and is analogous to a SPI chip select line. However, this line is active high, and most SPI

Operation	Latency (s)	Energy (mJ)
Image Capture	1.10	4.96
JPEG Compression (Q=90)	0.203	1.71
Image Transmission	6.67	73.8
Total	7.97	80.5

Table 6.3: Latency and energy measurements for key operations on Permacam, including image capture, compression, and image transmission. Measurements are averaged over 20 images.

implementations, including the hardware on the nRF52840, expect active low. Inserting a low latency inverter in between the LVLD output of the camera and the MCU allows the nRF52840 SPI peripheral to interface with the camera. This has the added benefit of reducing power required to read in images as the processor can sleep while the SPI peripheral directly writes images to memory.

Image Compression

To reduce energy and time required to transmit images, we employ JPEG compression on images prior to transmission. We use the Moodstocks JPEC encoder, a simple and portable monochrome JPEG encoder written in C [183]. JPEG compression, like most DSP algorithms, relies heavily on floating point arithmetic. Due to its integrated FPU, Permacam's processor is able to compress full resolution 320x320 images in 210 ms. This encoder is also used in Camaroptera [56], although this platform is based on the MSP430, which lacks an FPU. It requires 25 second to compress an downsampled 160x120 image using floating point emulation, and 7 second when JPEC is adapted to fixed point math. This fixed point adaptation also reduces image quality slightly. It is clear that a camera platform like Permacam benefits greatly from the inclusion of a faster, more capable processor. JPEG is not designed to compress mosaiced images. However, we find that monochrome compression with a high quality setting does not overly degrade image fidelity or color representation. We explore the effects of monochrome compression on Permacam's captured images further in Section 6.2.

Energy Harvesting and Storage

The energy storage architecture for Permacam is built around the TI BQ25505 [25] maximum power point tracking boost converter, like Permamote. The output of the BQ25505's power OR is boosted to the system voltage by a MAX17222 regulator to power all components under a single voltage domain. We use a 10.9 cm^2 amorpohous photovoltaic to charge a rechargeable LTO battery. We size this LTO battery according to the capacity heuristic from Section 3.3. The Permacam workload consists of an image capture, JPEG compression, and image transfer. These operations are benchmarked on the nRF52840 SoC and HM01B0 image sensor and summarized in Table 6.3. If we assume a periodic workload, capturing,

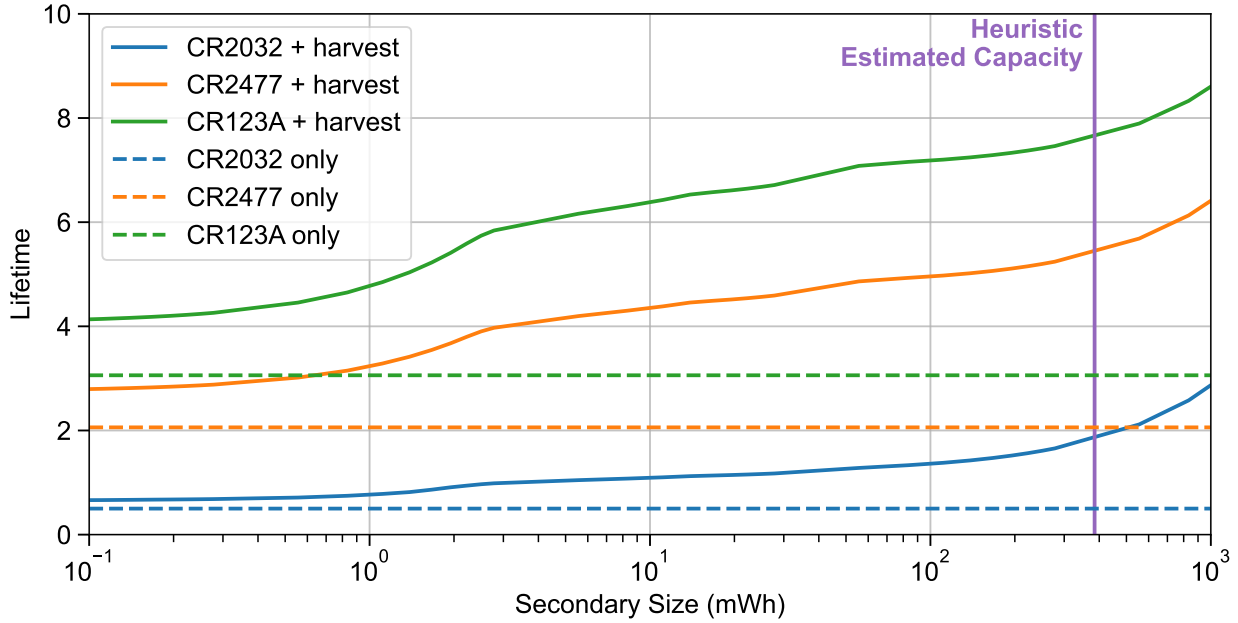

Figure 6.10: Lifetime estimates from Permacam simulations with various rechargeable and non-rechargeable energy capacity. The addition of energy harvesting results in capturing more than double the energy provided by a non-rechargeable cell alone. The addition of a CR2477 results in over five years of lifetime for Permacam.

compressing (at JPEG quality 90), and transmitting an image very ten minutes, Permacam will require 148 µW on average[8]. This workload for Permacam is pushing the bound on the energy that is harvestable within an indoor environment. If we consider there is no harvesting margin, this suggests a minimum sufficient capacity on the order of mWh. This capacity corresponds closely to a 125 mAh LTO battery with a nominal 2.4 V, which is the closest size available for purchase. We simulate Permacam's workload under the conditions of EnHANTs Setups A and D, while sweeping capacity and considering several different sizes of backup batteries. From the results, presented in Figure 6.10, we can determine an appropriate backup battery size to achieve an acceptable lifetime. We select a lithium 3 V, 1 Ah CR2477 coin cell battery. This amount of backup energy results in a design that can persist more than five years while capturing an image every ten minutes.

End-to-end Image Pipeline

To support transmitting full images, we implement a full application stack based on standard IP-based protocols. This end-to-end image architecture is depicted in Figure 6.11. We choose to use OpenThread, an open source implementation of the Thread network protocol for Permacam devices. Thread is a 6LoWPAN mesh network, and allows packets to be sent

8 $\frac{80.5\,\text{mJ}+14.5\,\mu\text{W}\times(600-7.97)}{600}$

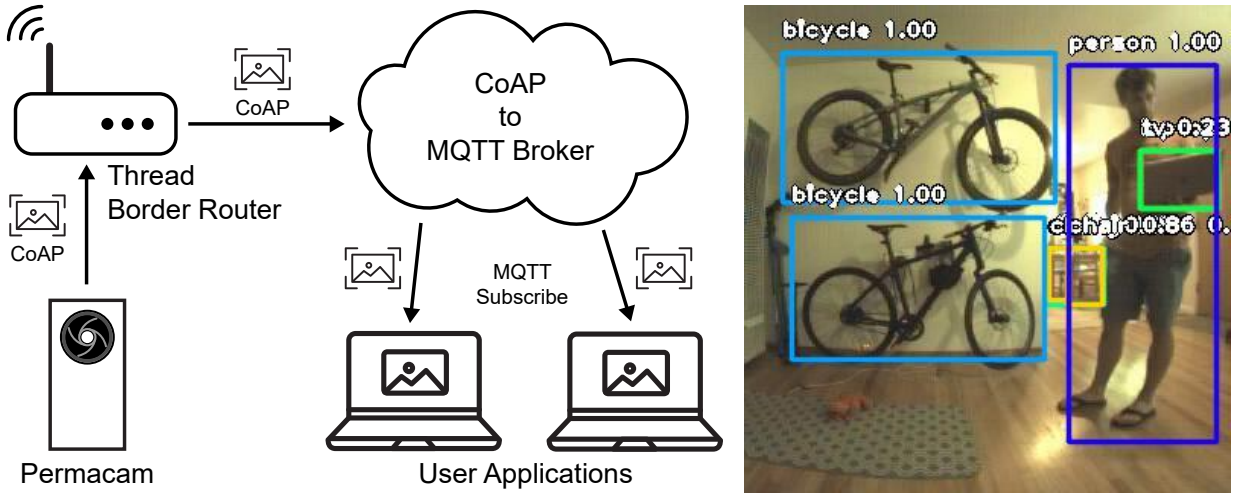

(a) End-to-end image transfer architecture

(b) YOLOv3 detection

Figure 6.11: The Permacam end-to-end image transfer architecture. Permacam uses OpenThread, a 6LoWPAN network. This allows it to transmit images over the CoAP block protocol directly to any IP endpoint. We implement a CoAP server to receive and reassemble image, demosaic them, and publish them over an MQTT stream. User applications can easily subscribe to incoming images and use them as inputs to object detection machine learning models, like YOLOv3.

end-to-end over a standard IP network. OpenThread also provides implementations of useful protocols like CoAP, SNTP, and DNS. CoAP is a lightweight restful protocol similar to HTTP, but over UDP [184]. On top of OpenThread's CoAP API, we implement the CoAP Block add-on feature. This allows us to fragment large data like an image across multiple CoAP payloads [185]. The choice to use CoAP block is a reluctant one however, and was based on the available implementations of reliable transfer protocols for low power networks. Permacam image transmission is a perfect application for TCP, although TCP is rarely implemented for low power networks. TCP has been shown to be provide a 40% higher throughput compared to CoAP block over an OpenThread network [186] over an OpenThread network. However, TCP was not implemented in the version of OpenThread used during the implementation of Permacam, limiting us to CoAP block.

Permacam captures an image, compresses it using JPEG, and transmits the image using a CoAP blockwise transfer. The CoAP block messages are sent to an endpoint that reassembles them, parses their contents, and publishes a JSON message over an MQTT broker. An application running on the endpoint listens for the transmitted mosaiced images, demosaics them, and republishes the color image for application use. High level applications wishing to interface with images captured by Permacam devices simply must subscribe to the relevant device topics to receive images. From there, the images can be easily processed in frameworks like OpenCV, Tensorflow, or Pytorch [187–189].

In addition to our data backhaul pipeline, we implement an independent update server. Permacam devices arbitrarily poll the update server every 24 hours for a new application image. If a newer image exists, it downloads it over a CoAP blockwise transfer and applies it. This allows users to deploy new applications quickly across an entire deployment with minimal intervention.

Local Image Classification

In addition to transmitting images, we also demonstrate the ability of the platform to perform local inference. While performing local object detection is not feasible with the memory and compute constraints of Permacam,image classification (with a limited number of classes) is not. To explore the capabilities of machine learning inference on Permacam, we modify and train our own MobileNets v1 network to perform person classification. We then deploy it to Permacam's processor using TensorFlow Lite for Microcontrollers.

The parameter flexibility of MobileNets v1 allows us to reduce the size of the model so it can fit on a microcontroller. We reduce the size of the network by reducing the depth of each convolutional layer by 75% ($\alpha = .25$). The model consists of a regular convolution with batch normalization followed by multiple depthwise separable convolutions [190]. There is insufficient memory to perform the forward pass on full sized 320x320 images, so images need to be downscaled. We configure this network for different input image dimensions including 48, 72, 96, and 120. We train this model for 80 epochs on the Visual Wake Word dataset [191] to achieve a validation accuracy of 78%, compared to a 90% accuracy with an unmodified MobileNet. The model weights require 230kB after post-training quantization and achieve an accuracy of 76% for input images of size 120x120 when using TensorFlow in Python. We deploy this model on Permacam and evaluate its performance in Section 6.2.

Permacam Evaluation

We evaluate Permacam on our goals of deployability and capability through a number of experiments. We begin with an analysis of the effects of JPEG compression on image size, time to send, and energy. From this analysis, we use these measured metrics to estimate platform lifetime using different representative workloads. We explore the capability of the platform by writing an object detection application on top of the Permacam end-to-end image transfer architecture. We evaluate the ability to detect objects in images captured by Permacam at varying qualities of compression and distance from the camera. We also implement local image inference in the form of person classification. We evaluate the performance of local classification, and compare it to image transmission.

Image Compression

A raw, full frame image from Permamote's camera is over 100 kB in size. Sending these raw images requires a significant amount of time and energy. We employ JPEG compression to

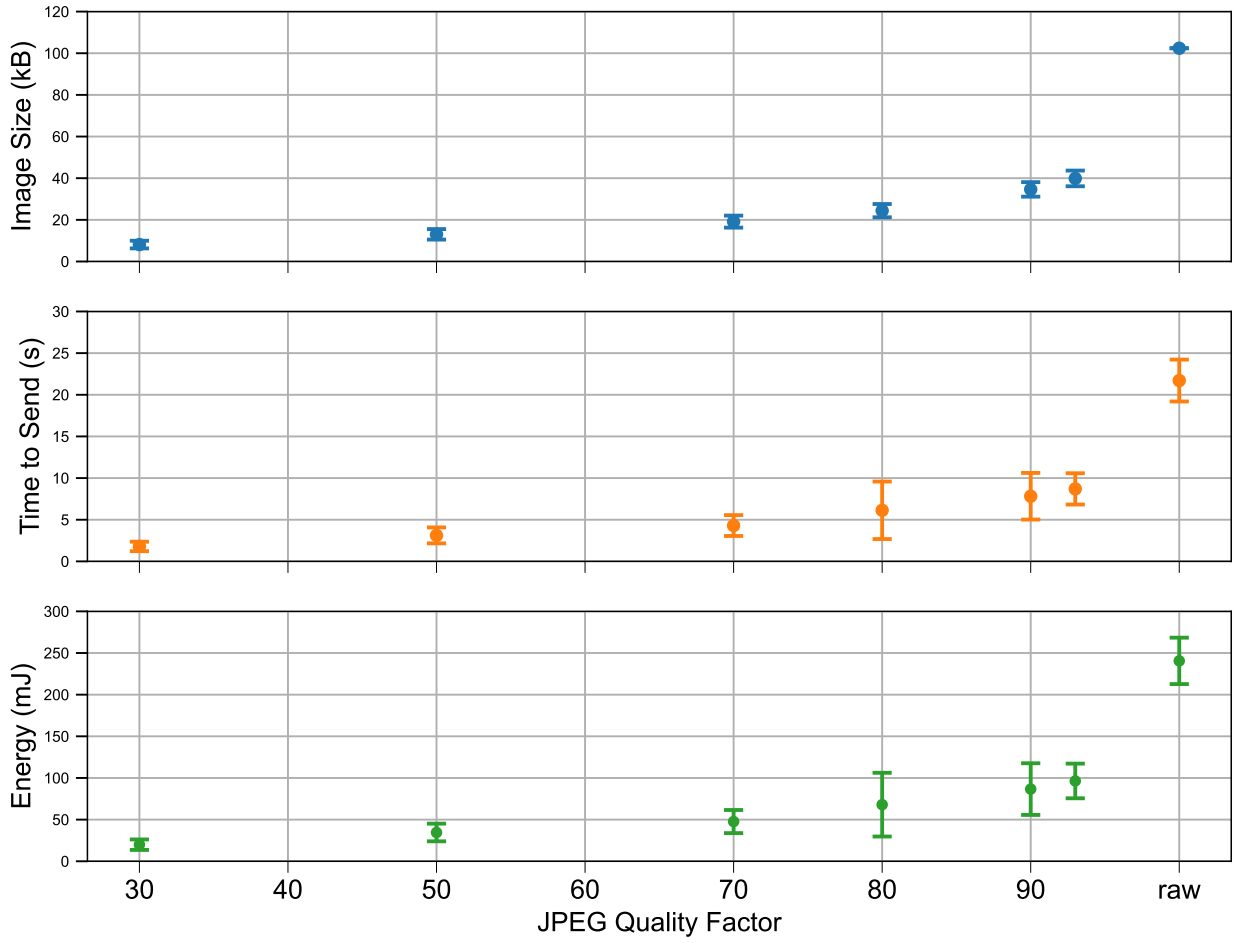

Figure 6.12: Effects of JPEG compression on image size, time to send, and energy to send. Compression provides exponential decrease in image size, which directly relates to decreases in the time and energy required to send images. Using a low quality factor results in images that are 8.2% the size of the original raw image. The amount of time and energy required to send images exhibit more variance than image size. This is the result of occasional packet loss and backoff during image transmission. Larger images require more packets and thus exhibit a higher probability of this occurring.

significantly reduce image size. This decrease in size comes at the cost of reduced image fidelity and color representation.

We use the Moodstocks JPEC encoder to compress captured images in JPEG format. JPEC is a monochrome JPEG encoder, and we have configured Permamote to use the color version of the HM01B0 image sensor. The images captured are a "mosaiced" color filter array, and performing monochrome JPEG compression on a mosaiced image reduces color representation in addition to image fidelity. JPEG can be configured with different quality factors from 1 to 100. A lower quality factor results in a less accurate representation and smaller compressed size. Image size relates linearly to the time and energy required to send images. The nRF52840's FPU is capable of performing floating point JPEG compression on 320x320 images in 210 ms, compared to 25 seconds with Camaroptera's MSP430.

To evaluate the effect of compression on image quality, latency, and energy, we configure Permamote to capture images, compress them with 6 different JPEG quality factors and transmit them. We collect 60 images, each with 6 compressed versions, and analyze the effects of compression. The results are displayed in Figure 6.12. JPEG image compression allows an exponential decrease in image size with respect to the quality factor. Compressing images at a high quality factor (JPEC's default 93) results in greater than a 2x reduction in image size, time to send, and energy required. There are diminishing gains after the knee of the curve at quality 90, but an image compressed with quality 30 is only 8% the size of a raw image, on average. Image compression on mosaiced images from Permamote also does not appear to degrade the perceived quality of the image. We display four compressed and raw versions of the same image in Figure 6.13. Qualitatively, and from a distance, these images appear nearly identical. However, closer inspection reveals artifacts and a significant loss in color fidelity with lower compression quality factors. Quality 93 and a raw image are nearly identical, while quality 30 is more obviously a lower quality compression. Even at low quality compression levels, important image details are preserved and still visible to the human eye.

To more quantitatively analyze the effects of compression on mosaiced images, we measure the Structural Similarity Index (SSIM) [192] of all 360 compressed images compared to their raw counterparts. SSIM is often used to measure the quality degradation of images due to compression or transmission. The results are summarized in Figure 6.14. A higher SSIM indicates that an image is a closer representation to the original raw image. Even with a low quality factor of 30 or 50, the SSIM for images average above 0.75. With an image compressed with a high quality factor, the perceived similarity is above 0.9 and, as seen in Figure 6.13, is almost indistinguishable from the raw image. These results suggest that using a quality factor of 90 or above on to compress images on Permamote is advantageous if we can afford the energy to send them.

Camera Density

While our numerical model can provide an accurate estimate of lifetime, it makes an important assumption. In our modelling, we assume that every image transfer requires the same amount of time and energy. The reality is that wireless environments can have interference, especially

Figure 6.13: An image compressed with different quality factors. In this scene, we set smart lights to bright, intense colors to determine the effects of compression on color representation. Compression is performed on the mosaiced version of the image, which after transmission is demosaiced into the color representations displayed. Due to this, an image compressed with a low quality factor loses significant color information compared to the raw image. Luckily, high quality factors produce a near-indistinguishable representation of the raw image. We explore a more quantitative view of image similarity in Figure 6.14.

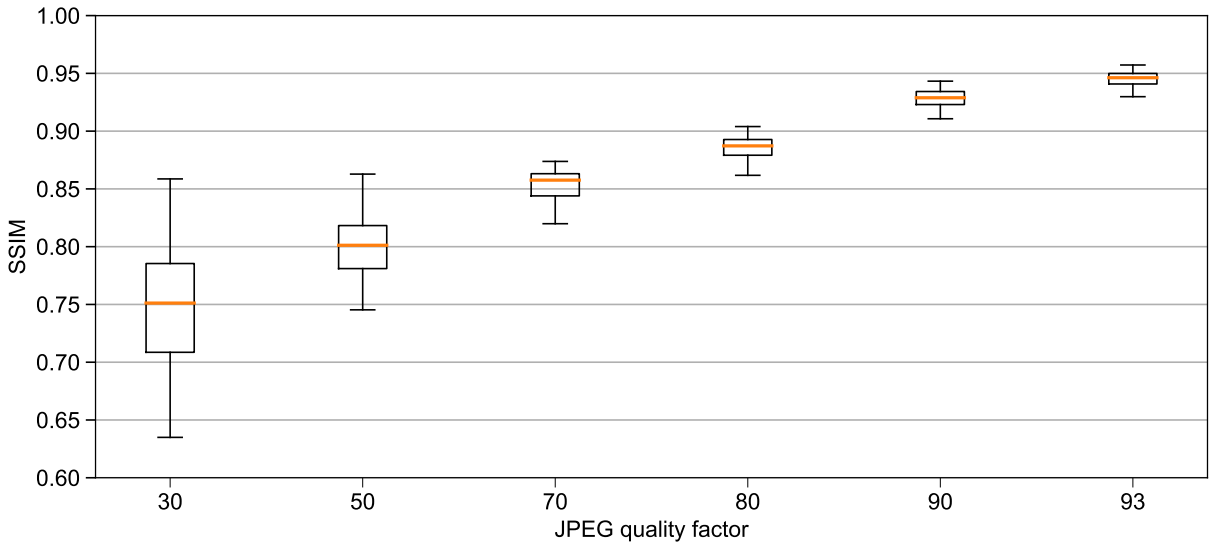

Figure 6.14: The image structural similarity index (SSIM) of images compressed at various JPEG quality factors. A higher SSIM indicates that an image is a closer representation to the original raw image. While a low quality factor results in smaller compressed images, it results in a significant loss in image structural similarity. Quality factors 90 and 93 provide a >90 SSIM, suggesting that they are near identical representations of the original image.

for low bandwidth networks like 802.15.4. The result of interference from other networks, devices, or other Permacams, means that packets will be dropped and retransmissions are necessary. This extends the length of image transmission and the increases required energy. We explore the scalability and effects of interference when multiple Permacams on the same network. We place all cameras within a meter radius of each other facing the same direction. The network consists of only one Thread router. Each Permamote is configured to simultaneously capture on a motion event to generate a worst-case collision scenario. Each camera is sending images compressed with quality factor 90. The Permacams are configured with a 2 minute backoff period after a motion event and camera capture. At the start of the experiment, a person walks into view of all the cameras, triggering them simultaneously. The person then remains in view of the cameras for half an hour. We vary the number of cameras active and measure the average time to send images over the duration of the experiment. The results are displayed in Figure 6.15.

We do not currently implement any application layer collision avoidance and we use the default 802.11.4 CSMA MAC protocol specified by Thread. Based on our naive implementation, having more than three Permamote sensors connected to a single router and on the same network leads to a dramatic increase in time to send images. Implementing an actual coordination protocol between Permacams would dramatically reduce collisions. Additionally, the use of TCP instead of CoAP block would significantly reduce transmission latency

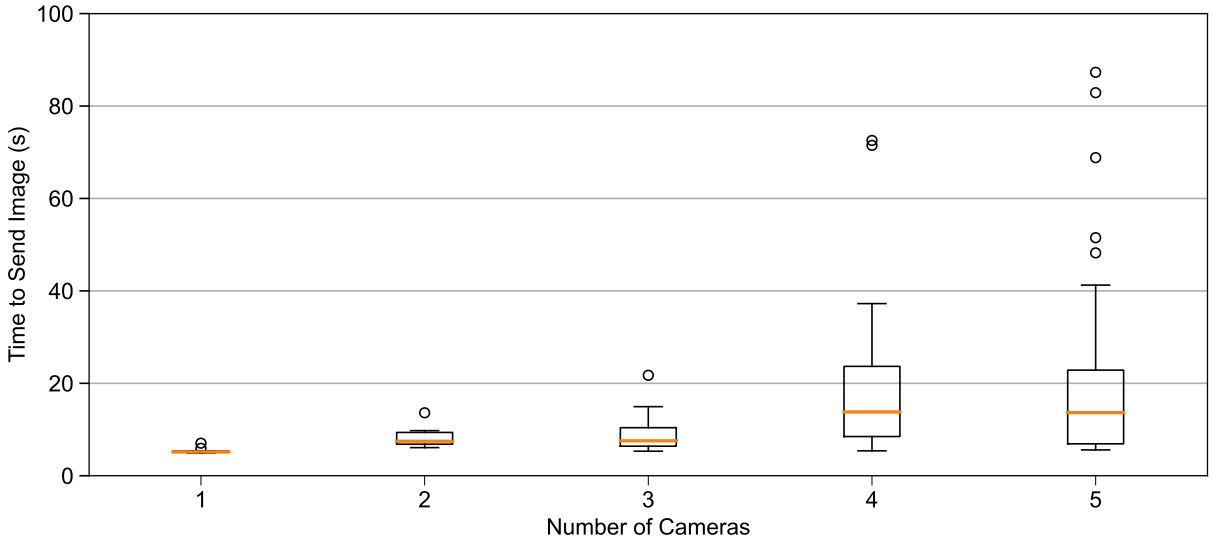

Figure 6.15: An image's time to send as number of Permacams on the network increases. All sensors are within one meter radius of each other, and configured to capture and transmit images on motion events to generate the worst case collision condition. Images are compressed with quality 90.

and possible overlap. Irregardless of these improvements, we expect most deployments of Permamote will be aware of the limitations and will only require one to two cameras per room.

Object Detection Performance

To illustrate the capability of the platform and end-to-end image transfer system, we evaluate the ability to perform object detection with Permamote. The images produced by Permamote are published over an MQTT stream, which feeds into a script that performs object detection using the pre-trained YOLOv3 network included in the Python ImageAI package [193]. An example of classification and bounding box detection results are illustrated by Figure 6.11b. We do not wish to measure the accuracy and performance of the ImageAI YOLOv3 model, but instead isolate the ability to perform detections on images captured by Permamote's camera and the effects of different compression quality factors on detection performance.

Compressed detection accuracy

While we measure the structural similarity of compressed images using SSIM, this does not directly relate to the performance of object detection on compressed images. To evaluate the effect of image compression on detection accuracy, we use the same 360 images used previously to measure the effects of JPEG compression. These images feature various representative

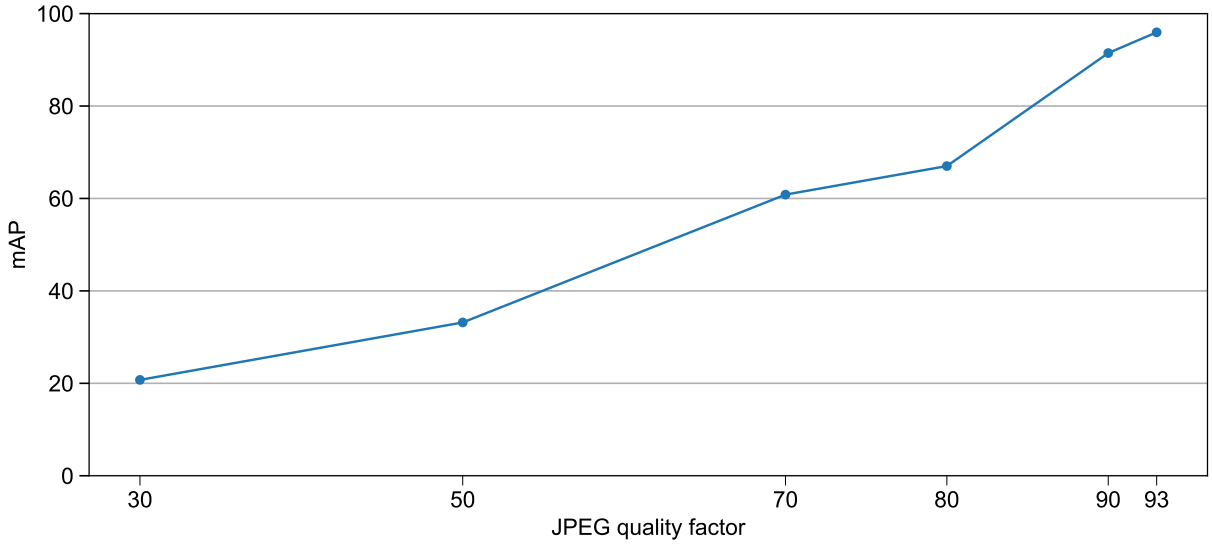

Figure 6.16: Mean Average Precision (mAP) of YOLOv3 person detection on compressed images compared to original raw versions. An IoU of 0.5 is used. Based on the qualitative and quantitative results of image similarity from Figure 6.13 and Figure 6.14 it is surprising that mAP degrades so quickly as the compression quality factor decreases. A quality factor of 90 or higher is necessary to achieve the near the same detection performance as a raw image

objects of the classes supported by the YOLOv3 model in ImageAI. To quantify the effects of compression, we calculate the mean average precision (mAP) using an intersection overlap union (IOU) of 50% of the detections compared to a raw image from Permamote. The metrics mAP and IoU=0.5 are both common metrics used to evaluate object detection algorithms. We refer readers unfamiliar with this metric to the Pascal VOC challenge paper [194].

In Figure 6.16, we display the mAP for different JPEG quality factors. Surprisingly, images compressed with a lower quality factor perform much worse than their SSIM suggests. Images compressed with quality factor 30 and 50 achieve less than 40% mAP, while factors above 90 achieve near identical detections to their raw counterparts.

Limitations of resolution

Permamote's camera is limited in resolution and color representation compared to most image sensors in modern cameras and cell phones. For example, the Google Pixel 3 features a 12.2 megapixel image sensor, which offers two orders of magnitude more resolution than the HM01B0's 0.1 megapixel. A lower resolution places a limit on the size of objects and the distance at which they can be detected. If an object is only represented by a few pixels, there is unlikely to be enough information to successfully detect it. To evaluate the capability to perform object detection with images captured by Permamote, we compare detection accuracy to images captured by a modern cell phone camera. We successively capture images

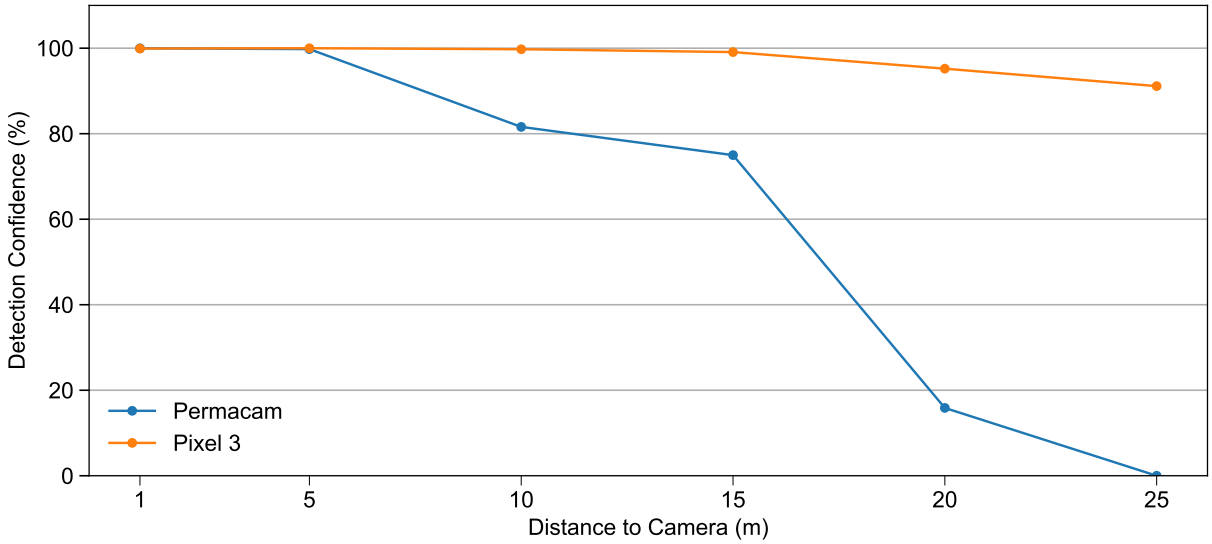

Figure 6.17: Detection confidence as distance from camera to person is increased. Compared to a modern smartphone camera, the camera on Permacam cannot compete due to limited resolution. However, images captured by Permacam still enable person detection at a distance of 15-20 meters. This distance is generally sufficient for most indoor spaces.

of a person at varying distances and measure the resultant detection accuracy, if any. The results are summarized in Figure 6.17.

A person is successfully detected at a distance of 20 meters with images from Permamote, though confidence is only 16%. Unsurprisingly, a person is easily detected at all distances tested with images taken by a Pixel 3. While images captured by Permamote cannot compete with the resolution of those captured by cell phones, they are sufficient for detection at reasonable distances. Many indoor spaces do not have sight lines longer than 20 meters, and if they do, additional camera coverage can help mitigate the lack of resolution of a single camera. The ability to clearly depict objects in the distance is partially dependent on the lens configuration of a camera. During this experiment, Permamote was configured with a 200°, wide angle lens. Better performance could be achieved with a lens with a narrower field of view and more magnification, at the cost of less coverage. We envision most applications will prioritize field of view to cover large areas instead of additional magnification.

Local Inference

In addition to end-to-end image transmission, Permamote is also capable of local image inference. We implement and train a modified MobileNets v1 network using TensorFlow in Python [187] and the Visual Wake Word dataset [191]. Using TensorFlow Lite for Microcontrollers, we quantize and deploy the model to Permamote's processor.

Dimension (pixels)	Latency (s)	MOPs	Energy (mJ)	Memory (kB)	Accuracy (%)
48	1.72	0.450	17.2	73.7	68.6
72	2.71	1.01	27.0	91.0	69.2
96	3.64	1.80	36.4	115.	72.4
120	5.10	2.81	50.9	146.	74.7

Table 6.4: Latency, millions of operations, energy, peak memory, and accuracy of local person classification. Images must be downscaled from full resolution, as inference on a 320x320 image requires too much runtime memory. The quantized version of model weights are used to measure accuracy of the validation set. The highest accuracy achieved is only 74.7%, and requires 5.1 seconds of continuous computation.

We upload 20 images from the validation set to Permamote, and measure inference latency and energy. We measure inference accuracy across our validation set of 8059 images, using quantized weights in TensorFlow. The results are summarized in Table 6.4. Inference on the largest image dimension (120) requires just over 5 seconds of continuous computation and has a peak memory usage of 146kB to achieve a 75.1% accuracy. This accuracy pales in comparison to what can be achieved with a full sized model running on more powerful hardware. Can local inference provide benefits to energy or latency to make up for its lackluster accuracy?

To answer this question, we also perform end-to-end image transmissions on Permamote to compare. Images are downsampled from 320x320 to 160x160 to compare with the downsampled images used for inference. They are 1.77x as large as the 120x120 images used for classification. Images are compressed using JPEG at varying qualities. We measure time and energy required to send images. The results of this experiment are summarized in Table 6.5. Surprisingly, the average energy required to transmit images of this size is generally less than performing inference on them. For example, we are able to transmit a 160x160 image compressed at quality 93 for under half the energy needed to perform inference on a smaller 120x120 image. Considering the effort to deploy machine learning inference to a microcontroller, the marginal accuracy of shrunken models, and the demanding energy requirements, a simpler and more beneficial solution is to just transmit images to more capable endpoints.

6.3 Summary

The designs of Permamote and Permacam are results of the design process developed in Chapter 3, Chapter 4, and Chapter 5. These designs are simultaneously long-lived and highly available solutions for illuminance sensing and image sensing. We avoid adhering to a single design strategy, such as a batteryless or energy preallocation-only design. Instead, we consider the requirements of each application and identify the type and size of energy harvesting, rechargeable buffers, and non rechargeable storage to achieve these goals. The end results are hybrid energy harvesting systems that utilizes a rechargeable LTO battery for

Compression Quality	Size (kB)	Latency (s)	Energy (mJ)
30	1.98	0.513	4.23
50	3.25	0.764	6.95
70	4.82	1.12	10.3
80	6.19	1.31	13.2
90	9.12	1.86	19.5
93	10.6	2.13	22.7
raw	25.6	4.91	54.7

Table 6.5: Time and energy required to send 160x160 images end-to-end. Images are compressed with varying JPEG qualities. Sending compressed images requires less time and energy than performing person classification on lower resolution images.

buffering, and either coin cell or cylindrical non-rechargeable batteries for backup. Through simulation, we show that these designs can achieve the intended application goals better than any other proposed battery-based or batteryless wireless sensor power supply in the literature. Permamote is estimated to provide decades of lifetime with high availability, which is more than triple the lifetime of similarly sized battery-only sensors, and more than double the availability of batteryless platforms in some cases. Permacam, inspired by other recent image sensor work like Camaroptera, is the first long-lived, energy harvesting indoor image sensor. Unlike Camaroptera, its design considers the effect of energy capacity and makes alternate design decisions to dramatically lower system power to enable indoor use.

We utilize Permamote and other battery-based and batteryless systems to evaluate our simulation. We find that our simulation produces lifetime and performance estimates that match reality. We simultaneously utilize our simulation to identify non-rechargeable backup battery sizing and estimate lifetime. Permamote is estimated to last for over a decade while maintaining a reliable and consistent periodic workload schedule.

Through many experiments, we evaluate Permacam's deployability and capability. For deployability, we demonstrate the effectiveness of image compression on image size and resulting time and energy to transmit images. We find that with high compression quality factors, images sent by Permacam are nearly indistinguishable from raw images, and are half the size. We also explore network congestion and the impact on the time to send images when multiple Permacam's are on the same network. We find that our 6LoWPAN network can easily tolerate 1-3 Permacams without any significant increase in time to send. We utlize our simulation tool to measure the performance and lifetime of Permacam under a periodic workload. Our simulation produces an estimated lifetime of over 5 years with compression using quality factor 90, and a period of 10 minutes. We also evaluate the effect of compression on the resulting usefulness of images for object detection. We find that high quality compressed images produce near identical detection results to raw images, and also find that Permacam is able to produce images that allow a person to be detected at 20 meters, despite the low resolution of its image sensor.

In addition to end-to-end image transmission, we also evaluate Permacam on its capability

of performing local image classification. We find that it is possible to deploy a neural network for person detection on Permacam. However, the resulting accuracy is only 74.7% and the required energy to perform this inference is more than double what is required to simply transmit an image. When considering energy and inference performance, it is preferable for Permacam to forward images to more capable servers that can perform more complex and accurate machine learning inference.

Chapter 7

Conclusion

In this dissertation, we have developed an application-focused framework for the design of energy harvesting wireless sensor system power supplies. We have identified common application requirements, and common power supply design archetypes, and identified which requirements are met by different archetypes. In particular, we find that batteryless energy harvesting systems fail to address many common and critical application requirements, including reliable and consistent operation. We examine the efficacy of energy harvesting in general when compared to non-rechargeable batteries, and determine that energy harvesting is a worthwhile investment *if* there is enough harvestable energy *and* the system has enough rechargeable capacity to capture the energy. We identify that a hybrid power supply architecture that combines energy harvesting and non-rechargeable batteries results in a system that can achieve consistent operation coupled with a long lifetime.

A main contribution of this work is the development of a reasoned approach to system-level rechargeable capacity sizing for energy harvesting sensor applications. Previous platforms in the literature have sized rechargeable capacity using ill-conceived heuristics or via arbitrary means. We develop a more reasoned heuristic for capacity sizing that is based on the relationship between minimum sufficient capacity, energy harvesting income, and average workload power. This novel heuristic allows system designers to select a rechargeable capacity based on their intended workload, expected income distribution, and a safety margin between the average income and workload power.

We expand upon this heuristic by developing an energy simulation of wireless sensor systems. We use this simulation tool in tandem with our heuristic to determine appropriate rechargeable capacity sizing to meet the requirements of a desired application. This simulation tool also aids in determining the appropriate non-rechargeable backup capacity to achieve high availability, reliability, and lifetime should an application require it.

Finally, we utilize this new heuristic and simulation tool to design and implement two new wireless sensor systems to address real indoor sensing applications that require high sensing availability and long sensor lifetimes. The result is two indoor energy harvesting sensors that utilize energy harvesting paired with non-rechargeable backup energy. These systems simultaneously provide high sensing availability and five to ten year lifetimes under

conservative energy harvesting conditions.

Essentially, the results of this work allows system designers to build energy harvesting sensors that can capture and utilize a greater amount of available energy. Beyond the design changes mentioned in this work, future sensors will require additional primitives, architectures, and more sophisticated embedded processing to capitalize on this additional energy.

7.1 Design Directions for Energy Harvesting Sensors

With more energy, future sensors will be capable of longer lifetimes, more reliable operation, and more capable local processing. To take full advantage of this energy and to maximize their lifetime, future sensors will need methods for dynamically reducing or increasing their quality of service to match that of available harvestable and stored energy. For applications suitable for batteryless designs, state retention can be greatly simplified with new SoC power domain architectures, resulting in more consistent and less wasteful operation. Future sensors will also require more processing and more memory to tackle increased computational complexity of new sensing modalities and inference methods.

Measuring Energy for Dynamic Adaptation

With greater capacity compared to what is usually allocated on batteryless sensors, embedded programmers do not need to optimize and develop intermittent software that reacts to energy state that changes on time scales of milliseconds or seconds. With greater capacity, programmers can instead develop software that reacts over days, weeks, months, or seasons [195]. These long-term adaptation strategies are accurate and effective, but would improve greatly with a new primitive: an accurate measure of the state of charge of a battery and the rate of energy entering or leaving it. Estimating the state of charge of a battery via its voltage potential alone is difficult. This is due to the relatively flat voltage charge and discharge curve of many battery technologies, compared to the linear curve of capacitors. Many commercially-available integrated circuits exist to estimate the state of charge of batteries via the battery's voltage and coulomb counting of the charge entering and exiting the battery. However, they are generally power prohibitive for many ultra low power applications. For example, the MAX17260, an "ultralow" power battery fuel gauge IC requires nearly 20 µW to estimate battery state of charge [196]. The inclusion of this fuel gauge would increase Permamote's idle power from 3 µW to 23 µW. While acceptable for consumer devices like wearables or smartwatches that are frequently recharged, this increase in idle power is inappropriate for long-lived energy harvesting sensors in adverse harvesting environments. These fuel gauges generally utilize periodic coulomb counting with a sense resistor and an ADC.

An alternate method for coulomb counting proposes counting the switching cycles of a DC-DC regulator [197] In the iCount approach, each cycle represents a fixed quanta of energy, and the switching frequency of regulators increases linearly with current supplied to a load. This measurement has the potential to be low power in comparison to other approaches. For

the initial implementation of iCount, a counter peripheral on an MCU was used and only required an extra 90 nA quiescent current to implement. This current draw could be improved further with dedicated hardware and the benefits of newer technology. The iCount technique has been included and improved upon in commercial products like the LTC3335, a buck-boost DC-DC regulator with an integrated coulomb counter. The LTC3335 only requires a little over 2 µW quiescent power for simultaneous regulation and coulomb counting. Integrated circuits like the LTC3335 provide a solution for measuring energy expended from a rechargeable battery, but there is currently no commercially available counterpart for measuring the energy being captured and stored in the battery. It is slightly more difficult to implement iCount on an energy harvesting DC-DC regulator, as the voltage of the harvester can vary significantly compared to that of a battery. Having an accurate measure of voltage is important for the iCount approach as the switching frequency of a DC-DC regulator is dependent on both the difference in voltage between the input and output as well as the current draw on the output. Implementing a iCount-based coulomb counter for a energy harvesting front end regulator would solve this problem and provide system designers the low power tools to accurately measure energy entering and leaving the system and write algorithms to adjust workload dynamically to detect and mitigate energy failure.

Alternate Power Domain Architectures

Beyond measuring the energy entering and leaving an energy harvesting sensor system, there also exists opportunities for exploring the utility of new power domain architectures for embedded processors. In particular, it would be useful to separate the power domains for an embedded processor's idle state retention and when it is active. With separate power domains for idle and active states, it would be easier to dedicate different energy storage options for either purpose. As we discussed in Chapter 2, one of the most challenging problems that batteryless researchers have attempted to address is maintaining forward progress over system power outages. They have developed state retention strategies that save volatile state to non-volatile flash or other memory. All of these strategies assume that the entire system will lose power at some point, and some state will be lost and must be recomputed the next time energy is available.

With a different power domain architecture that isolates active and idle operation, it would be possible to dedicate a small amount of non-rechargeable backup energy to greatly simplify and improve state retention for batteryless systems. This backup energy store would exist solely to retain state when the system runs out of harvested energy. For example, the Nordic nRF52840 utilizes 256 kB of volatile SRAM and requires 2.35 µW for full register and RAM retention [148]. A small CR2032 coin cell battery, with 240 mAh capacity, could preserve the state of registers and RAM for over a decade. New processors that utilize non-volatile memory like FRAM or MRAM would allow for even longer idle lifetimes, as they would only require energy to finish storing data on power outages or load data on restoration. Separate power domains would provide a simple and long-lasting method for state retention that is

incorporated into the processor power management architecture, making it easy for designers to use compared to current software-based approaches.

Sensing Dimensionality Requires Interfaces and Memory

In addition to new power architectures, low power embedded processors and sensors would also benefit from the inclusion of standard, high-speed, parallel interfaces for fast data transfer, and additional memory to hold and manipulate that data. Sensing modalities will continue to increase in power efficiency and the data they produce will increase in size and dimensionality.

This increase in data will require high speed and high bandwidth interfaces that are not commonly built into existing off-the-shelf embedded processors. For this reason, many recent research systems built to interface with image sensors have resorted to developing slow and inefficient bit-bang protocols, or in Permacam's case, shoehorning existing interfaces to interface with image sensors [13, 176]. The lack of appropriate interfaces for image sensors is not ideal, and future commercial embedded processors should consider the addition of a standard high speed data streaming interface.

Even with an appropriate interface, the increase in sensor data dimensionality requires significant memory to store it, manipulate it, or infer meaning from it. While the memory built into embedded processors has steadily increased in size, the ratio of memory size to processor performance is outpaced by other classes of computing. For example, a desktop personal computer may have 16 GB of RAM and a processor clocked at 3 GHz. This represents a ratio of 5.3 B/Hz. Likewise, an iPhone 14 has 6 GB and a processor speed of 3.24 GHz [198], with a ratio of 1.9 B/Hz. Conversely, the nRF52840 has 256 kB of SRAM and is clocked at 64 MHz. This represents a ratio of only 0.004 B/Hz. Compared to user-oriented devices, embedded SoCs provide significantly less memory when normalized for processor clock. While the multi-core CISC processors in desktop personal computers and smartphones are not necessarily comparable to single-core RISC embedded processors, the multiple orders of magnitude disparity in memory is notable. Partially, this is due to the workloads commonly performed on desktops and phones, which are usually graphics intensive and require fast reaction to user input. If embedded sensors are to collect and manipulate image data, or other sensor data with higher dimensionality, they will require more memory. Permacam had enough memory to capture a single raw image frame, but did not have enough to demosaic the image locally. The nRF52840 used on Permacam was capable of performing fast JPEG compression as well as other simple image manipulation like downsampling, but was limited by the amount of memory available to it. Similarly, the neural network models for image inference were limited by the memory on the system, resulting in models that sacrificed substantial accuracy just so they could feasibly fit and run within the memory constraints. Future embedded SoCs will require greatly increased memory to match the increased complexity and data dimensionality of new sensing modalities and the inference to make sense of them.

7.2 Implications for Future Sensing

This dissertation has developed and presented a novel design framework for building energy harvesting wireless sensors that can achieve application goals. This design framework provides a novel heuristic for sizing capacity to better capture harvestable energy for a given application, and identifies a hybrid architecture that combines harvesting with non-rechargeable backup energy as a design point that provides longevity and reliability.

At its core, this work enables wireless system designers to combine the reliability of traditional battery powered sensors with the longevity of energy harvesting technology. Designers are now able to determine correct capacity sizing, increasing their sensor energy budget. This increased budget not only enables longer-lived sensor deployments, but also permits the development of new and more complex sensing systems. Applications like image-based person detection and counting are now possible in indoor environments with off-the-shelf components. We believe that the tools developed and presented in this dissertation will allow future sensor designers to not only utilize more complex and power intensive sensors but also create systems that utilize sophisticated local inference and data processing than is currently practical. This will enable future sensors to better utilize improvements in machine learning to detect complex phenomena and to do so locally, preserving privacy and reducing costly network communication. Regardless of what direction future wireless sensors take, one thing has always proven true: sensor designers will always find a way to better utilize any available energy, and this work will enable them to capture as much as they need.

Bibliography

[1] Mark Weiser. "The Computer for the 21 st Century". In: *Scientific American* 265.3 (1991), pp. 94–105.

[2] Lohit Yerva et al. "Grafting Energy-harvesting Leaves Onto the Sensornet Tree". In: *Proceedings of the 11th International Conference on Information Processing in Sensor Networks*. IPSN'12. Beijing, China: ACM, 2012, pp. 197–208. ISBN: 978-1-4503-1227-1. DOI: 10.1145/2185677.2185733. URL: http://doi.acm.org/10.1145/2185677.2185733.

[3] Bradford Campbell, Branden Ghena, and Prabal Dutta. "Energy-Harvesting Thermoelectric Sensing for Unobtrusive Water and Appliance Metering". In: *Proceedings of the 2nd International Workshop on Energy Neutral Sensing Systems*. ENSsys'14. Memphis, TN: ACM, Nov. 2014.

[4] Bradford Campbell, Ye-sheng Kuo, and Prabal Dutta. "From Energy Audits to Monitoring Megawatt Loads: A Flexible and Deployable Power Metering System". In: *2018 IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)*. IEEE. 2018, pp. 189–200.

[5] Colleen Josephson, Neal Jackson, and Pat Pannuto. "Farming Electrons: Galvanic Versus Microbial Energy in Soil Batteries". In: *IEEE Sensors Letters* 4.12 (2020), pp. 1–4.

[6] Google. *Google Nest Temperature Sensor*. https://store.google.com/us/product/nest_temperature_sensor_specs. 2022.

[7] Josiah Hester and Jacob Sorber. "Flicker: Rapid Prototyping for the Batteryless Internet-of-Things". In: *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. SenSys '17. Delft, Netherlands: ACM, 2017, 19:1–19:13. ISBN: 978-1-4503-5459-2. DOI: 10.1145/3131672.3131674. URL: http://doi.acm.org/10.1145/3131672.3131674.

[8] Alan Mainwaring et al. "Wireless sensor networks for habitat monitoring". In: *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*. WSNA'02. Acm. 2002, pp. 88–97.

[9] Philo Juang et al. "Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet". In: *Proceedings of the 10th international conference on Architectural support for programming languages and operating systems*. 2002, pp. 96–107.

[10] Williot. *Williot Platform: Iot Pixel*. https://www.wiliot.com/product/iot-pixel. 2022.

[11] Mikhail Afanasov et al. "Battery-less zero-maintenance embedded sensing at the mithræum of circus maximus". In: *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 2020, pp. 368–381.

[12] Dhananjay Jagtap and Pat Pannuto. "Repurposing cathodic protection systems as reliable, in-situ, ambient batteries for sensor networks". In: *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2021)*. 2021, pp. 357–368.

[13] Harsh Desai et al. "Camaroptera: A Long-Range Image Sensor with Local Inference for Remote Sensing Applications". In: *ACM Transactions on Embedded Computing Systems (TECS)* (2022).

[14] Josiah Hester and Jacob Sorber. "The Future of Sensing is Batteryless, Intermittent, and Awesome". In: *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. SenSys '17. Delft, Netherlands: ACM, 2017, 21:1–21:6. ISBN: 978-1-4503-5459-2. DOI: 10.1145/3131672.3131699. URL: http://doi.acm.org/10.1145/3131672.3131699.

[15] Josiah Hester, Lanny Sitanayah, and Jacob Sorber. "Tragedy of the Coulombs: Federating Energy Storage for Tiny, Intermittently-Powered Sensors". In: *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. SenSys '15. Seoul, South Korea: ACM, 2015, pp. 5–16. ISBN: 978-1-4503-3631-4. DOI: 10.1145/2809695.2809707. URL: http://doi.acm.org/10.1145/2809695.2809707.

[16] Josiah Hester, Kevin Storer, and Jacob Sorber. "Timely Execution on Intermittently Powered Batteryless Sensors". In: *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. SenSys '17. Delft, Netherlands: ACM, 2017, 17:1–17:13. ISBN: 978-1-4503-5459-2. DOI: 10.1145/3131672.3131673. URL: http://doi.acm.org/10.1145/3131672.3131673.

[17] Josiah Hester and Jacob Sorber. "The future of sensing is batteryless, intermittent, and awesome". In: *Proceedings of the 15th ACM conference on embedded network sensor systems*. 2017, pp. 1–6.

[18] Alexei Colin, Emily Ruppel, and Brandon Lucia. "A Reconfigurable Energy Storage Architecture for Energy-harvesting Devices". In: *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS '18. Williamsburg, VA, USA: ACM, 2018, pp. 767–781. ISBN: 978-1-4503-4911-6. DOI: 10.1145/3173162.3173210. URL: http://doi.acm.org/10.1145/3173162.3173210.

[19] Brandon Lucia et al. "Intermittent Computing: Challenges and Opportunities". In: *SNAPL*. 2017.

[20] Amjad Yousef Majid, Patrick Schilder, and Koen Langendoen. "Continuous sensing on intermittent power". In: *2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE. 2020, pp. 181–192.

[21] Neal Jackson, Joshua Adkins, and Prabal Dutta. "Reconsidering Batteries in Energy Harvesting Sensing". In: *Proceedings of the Sixth ACM International Workshop on Energy Harvesting and Energy-Neutral Sensing Systems*. ENSsys'18. ACM. 2018.

[22] Neal Jackson, Joshua Adkins, and Prabal Dutta. "Capacity over capacitance for reliable energy harvesting sensors". In: *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*. 2019, pp. 193–204.

[23] Gilman Tolle et al. "A macroscope in the redwoods". In: *Proceedings of the 3rd international conference on Embedded networked sensor systems*. 2005, pp. 51–63.

[24] Analog Devices. *ADP5091 Datasheet*. http://www.analog.com/media/en/technical-documentation/data-sheets/ADP5091-5092.pdf. 2017.

[25] Texas Instruments. *BQ25505 Datasheet*. http://www.ti.com/lit/ds/symlink/bq25505.pdf. 2015.

[26] MATRIX Industries. *MATRIX Prometheus*. https://static1.squarespace.com/static/5fc7b8c2a4492a057e255bb3/t/6230d3034a3ac04c33b20acc/1647366915465/DS-Prometheus.20220218.D.U.pdf. 2022.

[27] Gabriel Marcano and Pat Pannuto. "Soil Power? Can Microbial Fuel Cells Power Non-Trivial Sensors?" In: *Proceedings of the 1st ACM Workshop on No Power and Low Power Internet-of-Things*. 2022, pp. 8–13.

[28] Joseph Polastre, Robert Szewczyk, and David Culler. "Telos: enabling ultra-low power wireless research". In: *Proceedings of the 4th international symposium on Information processing in sensor networks*. IPSN'05. IEEE Press. 2005, p. 48.

[29] Mobashir Mohammad and Mun Choon Chan. "Codecast: Supporting data driven in-network processing for low-power wireless sensor networks". In: *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE. 2018, pp. 72–83.

[30] Kai Li, Yousef Emami, and Eduardo Tovar. "Privacy-preserving control message dissemination for PVCPS". In: *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE. 2019, pp. 301–302.

[31] Michael P Andersen, Hyung-Sin Kim, and David E Culler. "Hamilton-A Cost-Effective, Low Power Networked Sensor for Indoor Environment Monitoring". In: (2017).

[32] Ecobee. *Ecobee SmartSensor*. https://www.ecobee.com/en-us/accessories/smart-temperature-occupancy-sensor. 2022.

[33] Honeywell Home. *T9 Smart Thermostat With Sensor*. https://www.honeywellhome.com/us/en/products/air/thermostats/wifi-thermostats/t9-smart-thermostat-with-sensor-rcht9610wfsw2003-u. 2022.

[34] Lutron. *Lutron Solutions Brochure*. https://www.lutron.com/technicaldocumentlibrary/367-1768.pdf. 2014.

[35] Phillips Hue. *Motion sensor*. https://www.philips-hue.com/en-us/p/hue-motion-sensor/046677570972. 2022.

[36] Emerson. *Rosemount™ 648 Wireless Temperature Transmitter*. https://www.emerson.com/documents/automation/product-data-sheet-rosemount-648-wireless-temperature-transmitter-rosemount-x-well-technology-en-us-192638.pdf. 2022.

[37] GE Measurement & Control. *Essential Insight.mesh Wireless System Product Datasheet*. https://perpetuapower.com/wp-content/uploads/2015/12/185301k1-GE-Essential-Insight.mesh-Product-Datasheet.pdf. 2021.

[38] Honeywell. *Honeywell OneWireless*. https://prod-edam.honeywell.com/content/dam/honeywell-edam/pmt/hps/products/pas/wireless/OneWireless-Network-PIN.pdf. 2020.

[39] Xiaofan Jiang, Joseph Polastre, and David Culler. "Perpetual environmentally powered sensor networks". In: *Proceedings of the 4th international symposium on Information processing in sensor networks*. IPSN'05. IEEE Press. 2005, p. 65.

[40] Aman Kansal et al. "Power management in energy harvesting sensor networks". In: *ACM Transactions on Embedded Computing Systems (TECS)* 6.4 (2007), p. 32.

[41] Peter Corke et al. "Long-duration solar-powered wireless sensor networks". In: *Proceedings of the 4th workshop on Embedded networked sensors*. SenSys'07. ACM. 2007, pp. 33–37.

[42] Kris Lin et al. "Heliomote: enabling long-lived sensor networks through solar energy harvesting". In: *Proceedings of the 3rd international conference on Embedded networked sensor systems*. SenSys'05. ACM. 2005, pp. 309–309.

[43] Jay Taneja, Jaein Jeong, and David Culler. "Design, modeling, and capacity planning for micro-solar power sensor networks". In: *2008 International conference on information processing in sensor networks (ipsn 2008)*. IEEE. 2008, pp. 407–418.

[44] Joshua Adkins et al. "The Signpost Platform for City-Scale Sensing". In: *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks*. IPSN'18. New York, NY, USA: ACM, Apr. 2018.

[45] Davis. *Vantage Pro2*. https://www.davisinstruments.com/pages/vantage-pro2. 2022.

[46] Spypoint. *Link-Micro-S-LTE Solar Cellular Trail Camera*. https://www.spypoint.com/en/products/solar-cellular-trail-camera/product-link-micro-s.html. 2022.

[47] WANCO. *Traffic Solar Surveillance System*. https://www.wanco.com/product/traffic-solar-surveillance-system/. 2022.

[48] Perpetua. *Perpetua® Power Puck® Solutions*. https://www.emerson.com/documents/automation/product-data-sheet-perpetua-power-puck-solutions-en-79966.pdf. 2019.

[49] ReVibe. *ReVibe Anura*. https://revibeenergy.com/revibe-anura/. 2022.

[50] Kinergizer. *Kinergizer Motion Energy Harvesting*. https://kinergizer.com/. 2022.

[51] Maria Gorlatova et al. "Challenge: ultra-low-power energy-harvesting active networked tags (EnHANTs)". In: *Proceedings of the 15th annual international conference on Mobile computing and networking*. 2009, pp. 253–260.

[52] Robert Margolies et al. "Energy-harvesting active networked tags (EnHANTs): Prototyping and experimentation". In: *ACM Transactions on Sensor Networks*. TOSN 11.4 (2015), p. 62.

[53] Paul Martin, Zainul Charbiwala, and Mani Srivastava. "DoubleDip: Leveraging thermoelectric harvesting for low power monitoring of sporadic water use". In: *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*. SenSys'12. ACM. 2012, pp. 225–238.

[54] Alanson P Sample et al. "Design of an RFID-based battery-free programmable sensing platform". In: *IEEE transactions on instrumentation and measurement* 57.11 (2008), pp. 2608–2615.

[55] Samuel DeBruin, Bradford Campbell, and Prabal Dutta. "Monjolo: An energy-harvesting energy meter architecture". In: *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. SenSys'13. ACM. 2013, p. 18.

[56] Matteo Nardello et al. "Camaroptera: A batteryless long-range remote visual sensing system". In: *Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*. 2019, pp. 8–14.

[57] Francesco Fraternali et al. "Pible: battery-free mote for perpetual indoor BLE applications". In: *Proceedings of the 5th Conference on Systems for Built Environments*. 2018, pp. 168–171.

[58] Hoang Truong et al. "Capband: Battery-free successive capacitance sensing wristband for hand gesture recognition". In: *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. 2018, pp. 54–67.

[59] Rishi Shukla et al. "SkinnyPower: enabling batteryless wearable sensors via intra-body power transfer". In: *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 2019, pp. 68–82.

[60] Kemet. *Not All Lifetime Calculations are Created Equal*. https://ec.kemet.com/wp-content/uploads/sites/4/2020/05/2020-05-27-Lifetime-calculation.pdf. 2020.

[61] Bradford Campbell, Joshua Adkins, and Prabal Dutta. "Cinamin: A Perpetual and Nearly Invisible BLE Beacon." In: *EWSN*. 2016, pp. 331–332.

[62] Nurani Saoda and Bradford Campbell. "No batteries needed: Providing physical context with energy-harvesting beacons". In: *Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*. 2019, pp. 15–21.

[63] Bradford Campbell and Prabal Dutta. "An Energy-harvesting Sensor Architecture and Toolkit for Building Monitoring and Event Detection". In: *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. BuildSys'14. Memphis, Tennessee: ACM, 2014, pp. 100–109. ISBN: 978-1-4503-3144-9. DOI: 10.1145/2674061.2674083. URL: http://doi.acm.org/10.1145/2674061.2674083.

[64] Benjamin Ransford, Jacob Sorber, and Kevin Fu. "Mementos: System support for long-running computation on RFID-scale devices". In: *Acm Sigplan Notices* 47.4 (2012), pp. 159–170.

[65] Domenico Balsamo et al. "Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems". In: *IEEE Embedded Systems Letters* 7.1 (2014), pp. 15–18.

[66] Josiah Hester, Timothy Scott, and Jacob Sorber. "Ekho: Realistic and repeatable experimentation for tiny energy-harvesting sensors". In: *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*. 2014, pp. 330–331.

[67] Alexei Colin et al. "An energy-interference-free hardware-software debugger for intermittent energy-harvesting systems". In: *ACM SIGOPS Operating Systems Review* 50.2 (2016), pp. 577–589.

[68] Kiwan Maeng, Alexei Colin, and Brandon Lucia. "Alpaca: Intermittent execution without checkpoints". In: *Proceedings of the ACM on Programming Languages* 1.OOPSLA (2017), pp. 1–30.

[69] Alexei Colin and Brandon Lucia. "Chain: tasks and channels for reliable intermittent programs". In: *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. 2016, pp. 514–530.

[70] Kasım Sinan Yıldırım et al. "Ink: Reactive kernel for tiny batteryless sensors". In: *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. 2018, pp. 41–53.

[71] Alexander Curtiss et al. "FaceBit: Smart Face Masks Platform". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5.4 (2021), pp. 1–44.

[72] Pressac. *Pressac CO2 Sensor Datasheet*. https://www.pressac.com/co2-temperature-and-humidity-sensors/. 2022.

[73] EnOcean. *EnOcean Multisensor*. https://www.enocean.com/en/products/self-powered-sensors/multisensor/. 2022.

[74] Everactive. *Batteryless Systems to Solve the IoT's Scaling Crisis*. https://everactive.com/batteryless-technology/. 2022.

[75] Yoonmyung Lee et al. "\"A modular 1 mm ^{3} die-stacked sensing platform with low power I ^{2} C inter-die communication and multi-modal energy harvesting\"". In: *IEEE Journal of Solid-State Circuits* 48.1 (2012), pp. 229–243.

[76] Jason Hill and David Culler. *A wireless embedded sensor architecture for system-level optimization*. Tech. rep. Citeseer, 2002.

[77] Jeeva. *Jeeva Wireless*. https://www.jeevawireless.com/. 2022.

[78] G Tuna and VC Gungor. "Energy harvesting and battery technologies for powering wireless sensor networks". In: *Industrial Wireless Sensor Networks*. Elsevier, 2016, pp. 25–38.

[79] Nordic Semiconductor. *Online Power Profiler for Bluetooth LE*. https://devzone.nordicsemi.com/power/w/opp/2/online-power-profiler-for-bluetooth-le. 2022.

[80] CAL Valk et al. "Designing for technology acceptance of wearable and mobile technologies for senior citizen users". In: *Academic Design Management Conference Proceedings (ADMC 2018)*. 2018, pp. 1361–1373.

[81] Vizi Metering. *Vizi Metering*. https://www.vizi.net/. 2022.

[82] Libelium. *Waspmote Technical Guide*. https://development.libelium.com/waspmote-technical-guide/. 2021.

[83] Digi. *XBee® 868LP Radio Frequency (RF) Modules*. https://www.digi.com/resources/documentation/digidocs/pdfs/90002126.pdf. 2018.

[84] Noshin Omar et al. "Lithium iron phosphate based battery—Assessment of the aging parameters and development of cycle life model". In: *Applied Energy* 113 (2014), pp. 1575–1585. ISSN: 0306-2619. DOI: https://doi.org/10.1016/j.apenergy.2013.09.003. URL: http://www.sciencedirect.com/science/article/pii/S0306261913007393.

[85] John Wang et al. "Cycle-life model for graphite-LiFePO4 cells". In: *Journal of Power Sources* 196.8 (2011), pp. 3942–3948. ISSN: 0378-7753. DOI: https://doi.org/10.1016/j.jpowsour.2010.11.134. URL: http://www.sciencedirect.com/science/article/pii/S0378775310021269.

[86] Brett Warneke et al. "Smart dust: Communicating with a cubic-millimeter computer". In: *Computer* 34.1 (2001), pp. 44–51.

[87] Mohammad Hassan Ghaed et al. "Circuits for a cubic-millimeter energy-autonomous wireless intraocular pressure monitor". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 60.12 (2013), pp. 3152–3162.

[88] Yen-Po Chen et al. "An injectable 64 nW ECG mixed-signal SoC in 65 nm for arrhythmia monitoring". In: *IEEE Journal of Solid-State Circuits* 50.1 (2014), pp. 375–390.

[89] Meisam Nazari, M Sencan, et al. "A novel needle-injectable millimeter scale wireless electrochemical glucose sensing platform for artificial pancreas applications". In: *Scientific reports* 9.1 (2019), pp. 1–11.

[90] Xiao Wu et al. "A 0.04 mm 3 16nW wireless and batteryless sensor system with integrated Cortex-M0+ processor and optical communication for cellular temperature measurement". In: *2018 IEEE Symposium on VLSI Circuits*. IEEE. 2018, pp. 191–192.

[91] Gyouho Kim et al. "A millimeter-scale wireless imaging system with continuous motion detection and energy harvesting". In: *2014 Symposium on VLSI Circuits Digest of Technical Papers*. IEEE. 2014, pp. 1–2.

[92] Bo Zhai et al. "Energy-efficient subthreshold processor design". In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 17.8 (2009), pp. 1127–1137.

[93] Ronald G Dreslinski et al. "Near-threshold computing: Reclaiming moore's law through energy efficient integrated circuits". In: *Proceedings of the IEEE* 98.2 (2010), pp. 253–266.

[94] Nordic Semiconductor. *NRF5340 Reference Manual*. https://infocenter.nordicsemi.com/pdf/nRF5340_PS_v1.1.pdf. 2021.

[95] *Hamilton H7C (PIR)*. 2018. URL: https://hamiltoniot.com/products/hamilton-h7c-pir-20-sensors-pack.

[96] Timothy W Hnat et al. "The hitchhiker's guide to successful residential sensing deployments". In: *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*. 2011, pp. 232–245.

[97] Joshua Adkins et al. "Demo: Michigan's IoT Toolkit". In: *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. SenSys'15. ACM. 2015, pp. 485–486.

[98] Bernhard Großiwindhager et al. "SnapLoc: An ultra-fast UWB-based indoor localization system for an unlimited number of tags". In: *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE. 2019, pp. 61–72.

[99] Maria Gorlatova, Aya Wallwater, and Gil Zussman. "Networking low-power energy harvesting devices: Measurements and algorithms". In: *IEEE Transactions on Mobile Computing* 12.9 (2013), pp. 1853–1865.

[100] Panasonic. *CR2032*. https://industrial.panasonic.com/cdbs/www-data/pdf2/AAA4000/AAA4000C321.pdf. 2006.

[101] Panasonic. *CR123A*. https://industrial.panasonic.com/cdbs/www-data/pdf2/AAA4000/AAA4000C285.pdf. 2022.

[102] Energizer. *Lithium Coin Handbook and Application Manual*. https://data.energizer.com/pdfs/lithiumcoin_appman.pdf. 2018.

[103] Hyung-Sin Kim et al. "System architecture directions for post-soc/32-bit networked sensors". In: *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. 2018, pp. 264–277.

[104] Philip Levis et al. "Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks". In: *Proc. of the 1st USENIX/ACM Symp. on Networked Systems Design and Implementation*. Vol. 25. NSDI '04. 2004.

[105] AVX. *TPS Series Capacitor Datasheet*. https://datasheets.kyocera-avx.com/TPS.pdf. 2018.

[106] Murata. *DMF Series EDLCs*. https://www.murata.com/en-us/products/productdetail?partno=DMT334R2S474M3DTA0. 2016.

[107] A.J. Bernheim Brush et al. "Home Automation in the Wild: Challenges and Opportunities". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 2115–2124. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979249. URL: http://doi.acm.org/10.1145/1978942.1979249.

[108] W. Keith Edwards and Rebecca E. Grinter. "At Home with Ubiquitous Computing: Seven Challenges". In: *Proceedings of the 3rd International Conference on Ubiquitous Computing*. UbiComp '01. Atlanta, Georgia, USA: Springer-Verlag, 2001, pp. 256–272. ISBN: 3-540-42614-0. URL: http://dl.acm.org/citation.cfm?id=647987.741327.

[109] Erika Shehan and W. Keith Edwards. "Home Networking and HCI: What Hath God Wrought?" In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '07. San Jose, California, USA: ACM, 2007, pp. 547–556. ISBN: 978-1-59593-593-9. DOI: 10.1145/1240624.1240712. URL: http://doi.acm.org/10.1145/1240624.1240712.

[110] Kai Geissdoerfer, Mikołaj Chwalisz, and Marco Zimmerling. "Shepherd: A portable testbed for the batteryless IoT". In: *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 2019, pp. 83–95.

[111] Francesco Fraternali et al. "Ember: energy management of batteryless event detection sensors with deep reinforcement learning". In: *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 2020, pp. 503–516.

[112] Samsung Electro-Mechanics. *Multi layer Ceramic Capacitor*. http://product.samsungsem.com/mlcc/CL31A476MQHNNN.do. 2013.

[113] Samsung Electro-Mechanics. *Multi layer Ceramic Capacitor*. http://product.samsungsem.com/mlcc/CL32A107MPVNNN.do. 2022.

[114] Seiko Instruments. *Micro Battery Catalog*. https://www.sii.co.jp/en/me/datasheets/chip-capacitor/cpx3225a752d/. 2017.

[115] Kyocera-AVX. *BestCap Ultra-low ESR High Power Pulse Supercapacitors*. https://catalogs.kyocera-avx.com/BestCap.pdf. 2020.

[116] Kemet Electronic Components. *FY Series Datasheet*. https://content.kemet.com/datasheets/KEM_S6015_FY.pdf. 2017.

[117] Murata. *Murata Supercapacitor Technical Note*. https://www.mouser.mx/pdfdocs/muRataEDLCTechnicalnote.PDF. 2013.

[118] Millibat. *Nimbus Product Line*. https://www.millibatt.com/product. 2022.

[119] A. H. Zimmerman. "Self-discharge losses in lithium-ion cells". In: *IEEE Aerospace and Electronic Systems Magazine* 19.2 (Feb. 2004), pp. 19–24. ISSN: 0885-8985. DOI: 10.1109/MAES.2004.1269687.

[120] T. Guena and P. Leblanc. "How Depth of Discharge Affects the Cycle Life of Lithium-Metal-Polymer Batteries". In: *Twenty-Eighth International Telecommunications Energy Conference*. INTELEC'06. Sept. 2006, pp. 1–8. DOI: 10.1109/INTLEC.2006.251641.

[121] A. Millner. "Modeling Lithium Ion battery degradation in electric vehicles". In: *2010 IEEE Conference on Innovative Technologies for an Efficient and Reliable Electricity Supply*. Sept. 2010, pp. 349–356. DOI: 10.1109/CITRES.2010.5619782.

[122] AA Battery Corp. *Rechargeable Li-ion Button Battery*. http://www.batteryspace.com/productimages/aa/20060224/LIR2032_new1.pdf. 2006.

[123] Data Power Technology. *Polymer Li-ion Rechargeable Battery Specification*. https://cdn.sparkfun.com/datasheets/Prototyping/SPE-00-301120-40mah-en-1.0ver.pdf. 2015.

[124] HuaHui Energy. *HTC-Lithium Titanate Series*. http://www.huahuienergy.com/?post_type=products&page_id=15707. 2022.

[125] Florian Hall et al. "Experimental investigation of the thermal and cycling behavior of a lithium titanate-based lithium-ion pouch cell". In: *Journal of Energy Storage* 17 (2018), pp. 109–117. ISSN: 2352-152X. DOI: https://doi.org/10.1016/j.est.2018.02.012. URL: https://www.sciencedirect.com/science/article/pii/S2352152X17302475.

[126] HuaHui Energy. *LTO Battery Specification*. http://www.batteryspace.com/prodspecs/7455.pdf. 2013.

[127] HuaHui Energy. *HFC-Lithium Iron Phosphate Series*. http://www.huahuienergy.com/?post_type=products&page_id=15704. 2022.

[128] M. Swierczynski et al. "Investigation on the Self-discharge of the LiFePO4/C nanophosphate battery chemistry at different conditions". In: *2014 IEEE Conference and Expo Transportation Electrification Asia-Pacific (ITEC Asia-Pacific)*. ITEC-AP'14. Aug. 2014, pp. 1–6. DOI: 10.1109/ITEC-AP.2014.6940762.

[129] E. Sarasketa-Zabala et al. "Cycle ageing analysis of a LiFePO4/graphite cell with dynamic model validations: Towards realistic lifetime predictions". In: *Journal of Power Sources* 275 (2015), pp. 573–587. ISSN: 0378-7753. DOI: https://doi.org/10.1016/j.jpowsour.2014.10.153. URL: http://www.sciencedirect.com/science/article/pii/S0378775314017728.

[130] STMicroelectronics. *EnFilm™ - rechargeable solid state lithium thin film battery*. https://www.st.com/content/ccc/resource/technical/document/datasheet/cd/ac/89/0b/b4/8e/43/0b/CD00270103.pdf/files/CD00270103.pdf/jcr:content/translations/en.CD00270103.pdf. 2014.

[131] TDK. *CeraCharge - Rechargeable Multilayer Ceramic Battery*. https://product.tdk.com/system/files/dam/doc/product/solid-state-batt/solid-state-batt/smd-assbs/ceracharge_bct1812m101ag.pdf. 2020.

[132] Ming-Jen Pan and Clive A Randall. "A brief introduction to ceramic capacitors". In: *IEEE electrical insulation magazine* 26.3 (2010), pp. 44–50.

[133] John Gill. "Basic tantalum capacitor technology". In: (1994).

[134] Kemet Electronic Components. "Technical Update - Comparison of Ceramic and Tantalum Capacitors". In: (2008).

[135] Jiří Libich et al. "Supercapacitors: Properties and applications". In: *Journal of Energy Storage* 17 (2018), pp. 224–227.

[136] Manisha Vangari, Tonya Pryor, and Li Jiang. "Supercapacitors: review of materials and fabrication methods". In: *Journal of energy engineering* 139.2 (2013), pp. 72–79.

[137] Naoki Nitta et al. "Li-ion battery materials: present and future". In: *Materials today* 18.5 (2015), pp. 252–264.

[138] Daniel H Doughty and E Peter Roth. "A general discussion of Li ion battery safety". In: *The Electrochemical Society Interface* 21.2 (2012), p. 37.

[139] Yuliya Preger et al. "Degradation of commercial lithium-ion cells as a function of chemistry and cycling conditions". In: *Journal of The Electrochemical Society* 167.12 (2020), p. 120532.

[140] Ilias Belharouak, Gary M. Koenig, and K. Amine. "Electrochemistry and safety of Li4Ti5O12 and graphite anodes paired with LiMn2O4 for hybrid electric vehicle Li-ion battery applications". In: *Journal of Power Sources* 196.23 (2011), pp. 10344–10350. ISSN: 0378-7753. DOI: https://doi.org/10.1016/j.jpowsour.2011.08.079. URL: http://www.sciencedirect.com/science/article/pii/S0378775311016259.

[141] CP Sandhya, Bibin John, and C Gouri. "Lithium titanate as anode material for lithium-ion cells: a review". In: *Ionics* 20.5 (2014), pp. 601–620.

[142] Joo Gon Kim et al. "A review of lithium and non-lithium based solid state batteries". In: *Journal of Power Sources* 282 (2015), pp. 299–322.

[143] Maxim Integrated. *MAX17222 Datasheet*. https://datasheets.maximintegrated.com/en/ds/MAX17220-MAX17225.pdf. 2017.

[144] Illinois Capacitor. *DCN Super Capacitors*. https://www.illinoiscapacitor.com/pdf/seriesDocuments/DCN%20series.pdf. 2017.

[145] Data Power Technology Limited. *DTP301120 Datasheet*. http://cdn.sparkfun.com/datasheets/Prototyping/SPE-00-301120-40mah-en-1.0ver.pdf. 2015.

[146] Murata. *What are impedance/ ESR frequency characteristics in capacitors?* https://article.murata.com/en-us/article/impedance-esr-frequency-characteristics-in-capacitors. 2013.

[147] *Fixed electric double-layer capacitors for use in electronic equipment*. Standard. Geneva, CH: International Electrotechnical Commission, 2006.

[148] Nordic Semiconductor. *NRF52840 Reference Manual*. https://infocenter.nordicsemi.com/pdf/nRF52840_PS_v1.2.pdf. 2021.

[149] Branden Ghena et al. "Challenge: Unlicensed lpwans are not yet the path to ubiquitous connectivity". In: *The 25th Annual International Conference on Mobile Computing and Networking*. 2019, pp. 1–12.

[150] Analog Devices. *Li-Ion/Polymer Shunt Battery Charger System with Low Battery Disconnect*. https://www.analog.com/media/en/technical-documentation/data-sheets/LTC4071.pdf. 2020.

[151] Ken Kundert. "Modeling dielectric absorption in capacitors". In: *The designer's guide community* 18 (2008).

[152] Julia Kowal et al. "Detailed analysis of the self-discharge of supercapacitors". In: *Journal of Power Sources* 196.1 (2011), pp. 573–579.

[153] Sanyo. *Amorphous Silicon Solar Cells / Amorphous Photosensors*. https://media.digikey.com/pdf/Data%20Sheets/Sanyo%20Energy/Amorphous_Br.pdf. 2007.

[154] Silicon Labs. *SI7021 Datasheet*. https://www.silabs.com/documents/public/data-sheets/Si7021-A20.pdf. 2016.

[155] Paul Kreczanik et al. "Study of supercapacitor aging and lifetime estimation according to voltage, temperature, and RMS current". In: *IEEE Transactions on Industrial Electronics* 61.9 (2013), pp. 4895–4902.

[156] Frank Richter et al. "Measurements of ageing and thermal conductivity in a secondary NMC-hard carbon Li-ion battery and the impact on internal temperature profiles". In: *Electrochimica Acta* 250 (2017), pp. 228–237.

[157] Ana-Irina Stroe et al. "Accelerated lifetime testing of high power lithium titanate oxide batteries". In: *2018 IEEE Energy Conversion Congress and Exposition (ECCE)*. IEEE. 2018, pp. 3857–3863.

[158] Xuebing Han et al. "Cycle life of commercial lithium-ion batteries with lithium titanium oxide anodes in electric vehicles". In: *Energies* 7.8 (2014), pp. 4895–4909.

[159] Michael Brunell et al. "Effect of Zero Volt Storage on Commercial Lithium Titanate Cells". In: *Meeting Abstracts*. 6. The Electrochemical Society. 2016, pp. 896–896.

[160] TXC. *SMD Seam Sealing Crystals Datasheet*. http://www.txccorp.com/download/products/quartz_crystals/2015TXC_7M_17.pdf. 2010.

[161] Joris Jaguemont et al. "Lithium-ion battery aging experiments at subzero temperatures and model development for capacity fade estimation". In: *IEEE Transactions on Vehicular Technology* 65.6 (2015), pp. 4328–4343.

[162] Feng Leng, Cher Ming Tan, and Michael Pecht. "Effect of temperature on the aging rate of Li ion battery operating above room temperature". In: *Scientific reports* 5.1 (2015), pp. 1–12.

[163] Kebin Chen et al. "Evaluation of the low temperature performance of lithium manganese oxide/lithium titanate lithium-ion batteries for start/stop applications". In: *Journal of Power Sources* 278 (2015), pp. 411–419. ISSN: 0378-7753. DOI: https://doi.org/10.1016/j.jpowsour.2014.12.051. URL: http://www.sciencedirect.com/science/article/pii/S0378775314020795.

[164] NASA. *Mars 2020 Perseverance Launch Press Kit*. https://www.jpl.nasa.gov/news/press_kits/mars_2020/download/mars_2020_launch_press_kit.pdf. 2020.

[165] GOMspace. *High capacity lithium-ion battery pack with a heater for CubeSats*. https://gomspace.com/shop/subsystems/power/nanopower-bpx.aspx. 2020.

[166] Fredrik Larsson and Bengt-Erik Mellander. "Abuse by External Heating, Overcharge and Short Circuiting of Commercial Lithium-Ion Battery Cells". In: *Journal of The Electrochemical Society* 161.10 (2014), A1611–A1617. URL: http://jes.ecsdl.org/content/161/10/A1611.abstract.

[167] "Annual energy outlook 2022 (AEO2022)". In: *Energy Information Agency* (2022).

[168] Peter R Michael, Danvers E Johnston, and Wilfrido Moreno. "A conversion guide: solar irradiance and lux illuminance". In: *Journal of Measurements in Engineering* 8.4 (2020), pp. 153–166.

[169] Hynek Raisigel et al. "Autonomous wireless sensor node for building climate conditioning application". In: *Sensor Technologies and Applications (SENSORCOMM), 2010 Fourth International Conference on*. SENSORCOMM'10. IEEE. 2010, pp. 68–73.

[170] Peter Wei et al. "A deep-reinforcement-learning-Based recommender system for occupant-driven energy optimization in commercial buildings". In: *IEEE Internet of Things Journal* 7.7 (2020), pp. 6402–6413.

[171] Leviton. *MULTI-TECHNOLOGY PIR/ULTRASONIC WALL SWITCH AND OCCUPANCY SENSOR*. https://store.leviton.com/products/multi-technology-pir-ultrasonic-wall-switch-and-occupancy-sensor-black-ossmt-gde?variant=18216120131. 2022.

[172] Anthony Rowe, Dhiraj Goel, and Raj Rajkumar. "Firefly mosaic: A vision-enabled wireless sensor networking system". In: *28th IEEE International Real-Time Systems Symposium (RTSS 2007)*. IEEE. 2007, pp. 459–468.

[173] Mohammad Rahimi et al. "Cyclops: in situ image sensing and interpretation in wireless sensor networks". In: *Proceedings of the 3rd international conference on Embedded networked sensor systems*. 2005, pp. 192–204.

[174] Amazon. *Blink Indoor*. URL: https://www.amazon.com/Blink-Home-Security-Detection-Included/dp/B018HC6PVA.

[175] Wyze. *Wyze Outdoor*. URL: https://wyze.com/wyze-cam-outdoor.html.

[176] Colleen Josephson et al. "Wireless computer vision using commodity radios". In: *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE. 2019, pp. 229–240.

[177] Saman Naderiparizi et al. "WISPCam: A battery-free RFID camera". In: *2015 IEEE International Conference on RFID (RFID)*. IEEE. 2015, pp. 166–173.

[178] *ESP32 Series Datasheet*. ESP32. Ver. 3.4. Espressif. 2020. URL: https://www.espressif.com/sites/default/files/documentation/esp32_datasheet_en.pdf.

[179] *HM01B0 Ultra Low Power Image Sensor*. HM01B0. Ver. 08. Himax. July 2018.

[180] Google. *OpenThread*. 2022. URL: https://openthread.io/.

[181] Bryce E Bayer. *Color imaging array*. US Patent 3,971,065. July 1976.

[182] Sparkfun. *Sparkfun Edge*. 2020. URL: https://www.sparkfun.com/products/15170.

[183] Maxime Brénon and Cédric Deltheil. *Moodstocks JPEC JPEG Encoder*. Mar. 2018. URL: https://github.com/Moodstocks/jpec.

[184] Zach Shelby, Klaus Hartke, and Carsten Bormann. "The constrained application protocol (CoAP)". In: (2014).

[185] Carsten Bormann and Zach Shelby. "Block-wise transfers in the constrained application protocol (coap)". In: *Internet proposed standard RFC* 7959 (2016).

[186] Sam Kumar et al. "Performant {TCP} for Low-Power Wireless Networks". In: *17th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 20)*. 2020, pp. 911–932.

[187] Martín Abadi, Ashish Agarwal, and Paul Barham. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: http://tensorflow.org/.

[188] Adam Paszke et al. "PyTorch: An Imperative Style, High-Performance Deep Learning Library". In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach et al. Curran Associates, Inc., 2019, pp. 8024–8035. URL: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

[189] Itseez. *Open Source Computer Vision Library*. https://github.com/itseez/opencv. 2015.

[190] Andrew G Howard et al. "Mobilenets: Efficient convolutional neural networks for mobile vision applications". In: *arXiv preprint arXiv:1704.04861* (2017).

[191] Aakanksha Chowdhery et al. "Visual wake words dataset". In: *arXiv preprint arXiv:1906.05721* (2019).

[192] Zhou Wang et al. "Image quality assessment: from error visibility to structural similarity". In: *IEEE transactions on image processing* 13.4 (2004), pp. 600–612.

[193] Moses Olafenwa and John Olafenwa. *ImageAI, an open source python library built to empower developers to build applications and systems with self-contained Computer Vision capabilities*. Mar. 2018. URL: https://github.com/OlafenwaMoses/ImageAI.

[194] Mark Everingham et al. "The pascal visual object classes (voc) challenge". In: *International journal of computer vision* 88.2 (2010), pp. 303–338.

[195] Rehan Ahmed et al. "Optimal power management with guaranteed minimum energy utilization for solar energy harvesting systems". In: *ACM Transactions on Embedded Computing Systems (TECS)* 18.4 (2019), pp. 1–26.

[196] *5.1*µA *1-Cell Fuel Gauge with ModelGauge m5 EZ and Optional High-Side Current Sensing*. MAX17260. Rev. 1. Maxim Integrated. June 2018.

[197] P. Dutta et al. "Energy Metering for Free: Augmenting Switching Regulators for Real-Time Monitoring". In: *2008 International Conference on Information Processing in Sensor Networks*. ISPN'08. Apr. 2008, pp. 283–294. DOI: 10.1109/IPSN.2008.58.

[198] Apple. *iPhone 14 Tech Specs*. 2022. URL: https://www.apple.com/iphone-14/specs/.