

Capacity over Capacitance for Reliable Energy Harvesting Sensors

Neal Jackson
University of California, Berkeley
neal.jackson@berkeley.edu

Joshua Adkins
University of California, Berkeley
adkins@berkeley.edu

Prabal Dutta
University of California, Berkeley
prabal@berkeley.edu

ABSTRACT

Today, most sensors that harvest energy from indoor solar, ambient RF, or thermal gradients buffer small amounts of energy in capacitors as they intermittently work through a sensing task. While the utilization of capacitors for energy storage affords these systems indefinite lifetimes, their low energy capacity necessitates complex intermittent programming models for state retention and energy management. However, recent advances in battery technology lead us to reevaluate the impact that increased energy storage capacity may have on the necessity of these programming models and the reliability of energy harvesting sensors.

In this paper, we propose a capacity-based framework to help structure energy harvesting sensor design, analyze the impact of capacity on key reliability metrics using a data-driven simulation, and consider how backup energy storage alters the design space. We find that for many designs that utilize solar energy harvesting, increasing energy storage capacity to 1-10 mWh can obviate the need for intermittent programming techniques, augment the total harvested energy by 1.4-2.3x, and improve the availability of a sensor by 1.3-2.6x. We also show that a hybrid design using energy harvesting with a secondary-cell battery and a backup primary-cell battery can achieve 2-4x the lifetime of primary-cell only designs while eliminating the failure modes present in energy harvesting systems. Finally, we implement an indoor, solar energy harvesting sensor based on our analysis and find that its behavior aligns with our simulation's predictions.

CCS CONCEPTS

• **Computer systems organization** → **Sensor networks**; • **Hardware** → **Sensors and actuators**; **Batteries**; **Reusable energy storage**;

KEYWORDS

Energy Harvesting, Sensor Network, Battery, Capacity

ACM Reference Format:

Neal Jackson, Joshua Adkins, and Prabal Dutta. 2019. Capacity over Capacitance for Reliable Energy Harvesting Sensors. In *The 18th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2019) (IPSN '19)*, April 16–18, 2019, Montreal, QC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3302506.3310400>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IPSN '19, April 16–18, 2019, Montreal, QC, Canada
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6284-9/19/04.
<https://doi.org/10.1145/3302506.3310400>

1 INTRODUCTION

Non-rechargeable (primary-cell) batteries have been the preferred method of powering sensors for both academic experimentation and commercial applications. They enable sensors that are easy to design, simple to program, and reliable to operate until their batteries are exhausted. However, as we strive towards ubiquitous sensor deployments aimed at supporting applications such as building automation and industrial monitoring, the human cost of frequent battery replacement may become untenable.

With the goals of alleviating battery replacement costs and deploying sensors in difficult-to-access environments, researchers have explored sensor designs that can rely solely on harvested energy. Due to declining active and idle power of core system components, subsisting on the small amounts of energy available from indoor solar, ambient RF, and thermal energy sources has become possible in the last decade. Initially, harvested energy was stored in capacitors because contemporaneous rechargeable lithium ion batteries required complex charging circuitry and offered low cycle lifetimes—undermining the goal of long-lifetime sensors.

Unfortunately, capacitors offer small energy capacity relative to the sensing, computing, communication, and storage tasks performed by sensors, and systems that employ them must cope with this limitation. Specifically, capacitor-based systems cannot perform atomic operations that require more energy than can be stored in the capacitor. They must intermittently work through non-atomic tasks over several iterations of starting up, performing some work, depleting energy reserves, and recharging. Still, significant progress has been made in making these systems more reliable and programmable. Progress latching and checkpointing techniques [30, 40] enable forward progress through reboots, special debugging tools [9] can emulate and replay energy state, and finely-tuned or reconfigurable power supplies [10, 16] increase sensor availability under varying workloads. Even with these techniques, capacitor-based energy harvesting sensors are less reliable and more difficult to use than their battery-powered counterparts.

Recent trends in technology, including new energy harvesting battery management ICs and advances in rechargeable battery (secondary-cell) chemistries, suggest that we no longer may be limited to using capacitors in low power, long-lifetime designs [21]. New energy harvesting ICs are low leakage and offer high efficiency max power point tracking even at low harvesting voltages and currents [3, 43]. New battery chemistries like lithium titanate (LTO) [19] and lithium iron phosphate (LiFePO₄) [42] can withstand 4,000-10,000 cycles before cell degradation [36, 44] and come in small (100-700 mm³), and inexpensive (\$1 USD) packages [19, 20]. This work seeks to analyze and characterize the impact of energy storage capacity on energy harvesting sensor systems. We define a design space that is characterized by system capacity, harvesting

potential, and workload, identifying where different intermittent computing techniques are necessary and helpful. Further, we explore the benefits of capacity on energy utilization and reliability. Finally, we show the potential gains offered by including a pre-charged, backup energy store, such as a primary-cell battery, in a sensor's power supply and examine the impact on lifetime.

In our initial analysis, we consider human-occupied indoor environments and solar energy harvesting sensors. We find that in these scenarios, increasing energy storage capacity to an amount much greater than that required by a sensor's workload eliminates the need for checkpointing and significantly reduces the impact of finely-tuned and reconfigurable power supplies. Eliminating the necessity of these techniques further increases operating efficiency because the techniques themselves draw considerable power.

To analyze the impact of capacity under variable harvesting conditions or variable workload, both of which are common in real-world environments, we develop a numerical model that utilizes energy traces and several workloads that are representative of real hardware. We use this model in the context of solar energy harvesting, and find that sensors with higher capacity, on the order of that of a battery, capture and utilize 1.3-2.6x the energy of systems with capacitor-sized energy stores. This energy allows systems to perform near 100% of scheduled tasks on time compared to only 20-80% for systems with insufficient energy capacity to continue operating through energy droughts, or periods in which there is not sufficient harvestable energy to sustain operation.

We find that the inclusion of a backup energy store like a primary-cell in energy harvesting systems can eliminate nearly all drawbacks of energy harvesting designs. Backup energy can be used to retain state across reboots, cold-start harvester front ends to increase their efficiency, and boost the availability of a sensor when harvested energy is depleted. While a backup energy store does come with a finite lifetime, our model predicts that hybrid designs performing all of these operations achieve 2-4x the lifetime of primary-cell only designs. A backup energy source that is only used for state retention and harvester cold-start would offer longer lifetimes.

Based on our design exploration, we implement a solar energy harvesting platform, Permamate, to enable autonomous lighting control for indoor spaces. It has a 20 mAh secondary-cell for more energy capture and a backup energy store to ensure state retention and high reliability even during long energy harvesting droughts. We present the new and low-power components selected for its design, demonstrating the viability of the design point and power measurements of these components, which are used as the basis for the standard workloads that we define and model. To validate the model, we deploy Permamate along with battery-powered and intermittent systems and compare their performance to that predicted by the model. We show that lifetime estimates, the frequency and timing of charge-discharge cycles, and expected number of transmissions are similar to that of the deployed systems.

When we began the design of Permamate, we initially intended to use capacitors as the rechargeable energy store. At the time, the assertion that batteries were expensive, bulky, and had extremely limited lifetimes was prevalent [10, 15–18, 29]. Upon further examination, we find that these claims are no longer true for many applications due to technology improvements, and the gains provided by increased energy storage capacity are numerous. Greater

rechargeable capacity allows a system to obviate the need for intermittent techniques and harvest more energy. More collected energy translates to higher reliability and capability. The addition of a primary-cell allows an otherwise unreliable system to operate without interruption for years or decades, or to support state retention and eliminate expensive cold start. With these advances, we envision future energy harvesting sensors which are not limited by the visions of immortality but are instead long-lived, reliable, and capable of enabling interesting and useful applications.

2 RELATED WORK

Prior work regarding energy harvesting sensor systems can be broadly divided into two categories: those which make use of intermittent computing techniques and those which do not. Intermittent systems often exist in a regime of unreliable and ultra low harvester power, where operation and uptime are not guaranteed. As such, they often lose power and reboot while intermittently working through a sensing task. A wealth of work has been devoted to making these systems usable and reliable. Other energy harvesting systems, especially those deployed outside, have access to significantly more harvestable energy and are able to store more of this energy for later use, so they do not use intermittent computing techniques to complete their workloads.

2.1 Intermittent Sensors

Energy harvesting systems that rely exclusively on repeatedly buffering small amounts of energy to operate are commonly referred to as intermittent systems. Many choose to employ capacitors as an energy buffer due to their theoretically infinite lifetime, but are limited to small energy capacities, and are only as reliable and lively as their source of harvested energy. In situations of energy drought, these platforms often deplete their small energy stores, and they power off and lose state, potentially in the middle of an important operation or for an extended period of time.

One-shot Intermittency. The Gecko and Monjolo platforms ignore the difficulties associated with completing longer workloads and instead allocate just enough capacitance to turn on and perform a simple task. Sometimes, the rate of harvesting is the sensor itself [8, 12, 45]. However, this approach can require tedious and non-standard optimization of the cold start process and is severely limited in its simplicity. Performing any sensing or computing outside of the hardware's intended use case is often not possible, and it is difficult to distinguish sensor failure from a lack of energy.

Checkpointing. Other work in this space attempts to cope with intermittency by developing tools and programming language primitives that allow complex and energy intensive tasks to execute despite limited energy storage. Intermittent-aware programming models and compilers were developed to enable checkpointing and progress latching over workloads that may require more energy than can be stored or harvested in a reasonable time [18, 30, 40]. New debugging tools spanning both the hardware and software domains measure the energy required for specific code operations and restore energy state during code breakpoints [9].

Hardware Hysteresis Management. In addition to intermittent software techniques, hardware platforms have been developed to

increase availability and responsiveness through the fine-grained management of capacitor charging hysteresis. For these systems, it is often assumed that the upper hysteresis threshold, the point at which a charging device turns on, is the voltage at which a capacitor is full, and the lower threshold, the point at which the system turns off or sleeps, is the minimum operating voltage of components in a system. Smaller capacitors can charge to an upper threshold and turn on faster, but store less energy. Hysteresis management techniques attempt to combine different sized capacitors to minimize charging time while also maximizing available energy.

The Flicker platform employs federated energy storage in which each peripheral has its own storage tuned to the task it is expected to perform [15, 16]. This has the effect of allowing various components to charge their storage faster, as well as isolating power failure to independent components. The Capybara platform is similar, but provides even more flexibility by dynamically resizing its banked capacitor store to match the energy required by a task [10]. This leads to the lowest possible cold start and capacitor recharge times to support a given operation.

We believe the assumption that operation should be coupled to full-swing hysteresis is not valid for many systems. Capybara does explore the possibility of setting an upper threshold to less than the maximum voltage and using an adjustable bottom threshold instead of dynamically resizing their capacitor store. However, they disregard these options due to high voltage comparator power and long cold start times, respectively. While these decisions make sense in this context, the importance they place on cold start optimization is specific to their design. Capybara's power supply has a significant "power overhead of the power system" that limits the effectiveness of sleeping. Due to this, they opt to fully discharge their storage on every operation and optimize cold start. In practice, if a system has the ability to enter a low power sleep mode or power off, it can avoid cold start and control its energy usage by willfully entering these states. With this operating principle, the benefits of hardware hysteresis management are limited to reducing cold start time and are workload independent.

These complex software and hardware solutions, while increasing usability and reliability, do not address the singular problem for capacitor-based energy harvesting systems: in the face of plentiful harvestable energy, they are not able to store the energy for later use (in times of energy drought). As a result, they must micro-optimize the little energy they have. In many applications, if these systems had sufficient capacity, they could instead adjust sensing rate and workloads over periods of days or weeks. We show that the energy captured by these systems and their subsequent availability could be substantially improved by using larger energy buffers.

2.2 Non-intermittent Sensors

Non-intermittent energy harvesting sensors have largely existed in environments with plentiful harvestable energy and have been designed with sufficient capacity to capture this energy. Some devices have also embraced backup primary-cells to further ensure reliable operation.

Rechargeable Batteries. Most examples of such devices are deployed outdoors. For these systems, the obvious choice is to use secondary-cells, as they can better capture a significant portion of

copious solar energy [2, 11, 22, 23, 28]. Most notable of this group is Prometheus, which utilizes a supercapacitor as a short term energy store, and when full, charges a backup rechargeable lithium battery [22]. At the time of its design, lithium cells offered highly limited recharge cycles, and by utilizing a supercapacitor, much of this charge-discharge volatility was masked from the secondary-cell, extending its lifetime. Rechargeable batteries have also been applied to indoor sensing. The EnHANTs sensor uses an intentionally oversized NiMH battery, with plans to eventually use a thin-film battery [32]. While the choice to use batteries allows for more energy capacity, NiMH and thin film chemistries offer poor energy density and lower cycle life compared to lithium based chemistries. DoubleDip and other sensors [33, 39] use a lithium-manganese battery. DoubleDip notes that supercapacitors offer lower energy density and higher leakage when compared to batteries, but admits that the lithium-manganese chemistry suffers from low maximum output currents and a limited number of charge-discharge cycles. While the limitations of past batteries have slowed their adoption in low energy harvesting scenarios, we claim that recent developments in battery technology will enable higher capacity energy storage without these trade offs.

Backup Energy Store. Regardless of which energy store is used, energy harvesting systems will experience some degree of intermittency. We advocate that a non-rechargeable backup energy store can be utilized to mask this intermittency, cold start electrical components, and provide consistent, reliable, and lively operation.

The only system we find that employs a non-rechargeable backup is the Pressac line of capacitor-based energy harvesting sensors which use battery backup to obtain an estimated 10 years of continuous operation [38]. This work suggests that these sensors could significantly increase their lifetime by using a larger secondary energy store. There has been little exploration on the benefits of this hybrid design and the use of primary-cells to avoid intermittency, cold start harvesting circuits, and provide baseline reliability.

3 AN ENERGY HARVESTING FRAMEWORK

We seek to illustrate the design space for energy harvesting sensors in two ways. The first defines an energy harvesting sensor framework to examine when designs are feasible and when they require intermittent techniques. The second examines dynamic income energy and device behavior through numerical modeling and simulation. The framework is based on three key metrics: harvested energy income, workload, and capacity.

3.1 Sensor Regimes

The framework splits the design space into four main regimes: always on, infeasible, checkpointing required, and no intermittent techniques required. These regimes and their constraints are illustrated in Figure 1, and explained in more detail below.

Always On. If the energy harvester supplies a sensor with more power than the max power it will ever draw, then the device needs no energy buffer capacity to remain operational. If this is not the case, then a sensor must have some ability to buffer energy to use when its operating power exceeds than the harvester input power.

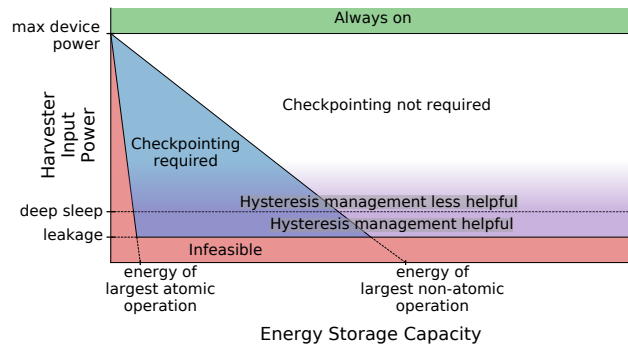


Figure 1: Design space for energy harvesting sensors based on their energy income (which we assume is constant for this analysis), energy storage capacity, and workload. Workload is represented by the largest atomic/non-atomic operations supported by a design, as well as the deep sleep and leakage power. The plot breaks into four regions: 1) always on or effectively powered, 2) Infeasible due to lack of energy storage or leakage higher than harvesting rate 3) Feasible but requires checkpointing to make forward progress, and 4) Enough energy storage to not require or benefit from checkpointing. 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.

Infeasible. If the energy harvester supplies less power than the system leakage, the energy buffer will never charge. If the energy buffer capacity is less than the energy required to perform a workload’s largest atomic operation, with energy harvested during the operation itself, then that operation will not have enough energy to complete. Neither of these designs will make forward progress and are therefore infeasible. Common atomic operations on energy harvesting sensors include sampling a sensor, sending a radio packet, booting the processor, and performing a checkpoint.

Checkpointing Required. If the energy buffer can hold enough energy to perform atomic operations, but not enough to complete workloads composed of multiple, chained atomic operations (such as sampling a sensor and then sending a radio packet), then a mechanism for saving state and continuing progress on the next reboot must be employed.

No Intermittent Techniques. A sensor that has enough harvester potential and energy capacity to complete a workload’s longest non-atomic operation can operate without checkpointing. Such systems also benefit as energy devoted to checkpointing can be used for a workload instead.

Hysteresis Management. Finally, if a sensor’s deep sleep power draw is a substantial fraction of the harvester power, as is the case with Capybara [10], then it is beneficial to continue operating until the energy buffer is depleted, power off, and recharge quickly rather than stop early and recharge slowly. Under this scenario, hysteresis management techniques, such as reconfigurable capacity and federated energy can increase sensor performance as discussed in Section 2. The utility of hysteresis management is diminished

when the ratio of harvester power to deep sleep power increases. For sensors that can willfully power off or sleep, operating thresholds can be controlled, disentangling capacity and charging hysteresis. Their deep sleep power is equal to leakage, and such techniques will not improve recharge times.

For all systems, these techniques can decrease cold start time by reducing the capacity that must be charged to achieve cold start. This is more beneficial for systems that cold start frequently and have a higher energy capacity, however, a higher-capacity energy store also has a lower probability of needing to cold start. Therefore the benefits of hysteresis management for cold start with respect to storage capacity are in conflict with their necessity, and we do not attempt to quantify these subtle nuances in Figure 1.

3.2 Framework Limitations

This framework makes a couple simplifying assumptions that prevent it from fully capturing the richness of the design space.

Backup Energy Store. This framework 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 checkpointing 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 checkpointing 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 power supplies, or to increase the efficiency of the energy harvesting front-end at low voltages. Finally, in periods of long energy drought, the backup energy store could increase sensor availability.

While the use of a backup energy store does constrain the sensor to a finite lifetime, energy harvesting can substantially extend these lifetimes under certain harvesting conditions.

Constant Harvester Power. The framework assumes an energy harvester will supply a constant energy income, when in reality income is often highly variable. In practice, a sensor platform both defines the regions of the plot, and occupies a vertical line which represents the energy storage capacity of the sensor combined with the range of harvester input powers it might experience. We expect this line will span multiple regions for most sensors.

However, by ignoring variability, the plot also fails to illustrate key benefits of capacity under varying energy incomes and workloads. Intuitively, higher energy buffer capacity can store energy in times of excess and supply that energy in times of drought. This balancing out of energy income effectively raises the minimum power supplied by the energy harvester. Because the extent of this impact is completely dependent on the variability of the energy income and workload of the sensor, we also develop a numerical simulation to quantify the impact of capacity on key metrics including energy utilization, availability and reactivity.

4 MODELING THE COMMON CASE

The framework we have presented does not consider harvesting and workload variability. To explore the dynamic effects that energy capacity and backup storage have on sensor performance in the face of variability, we develop a simple numerical model. We use representative environmental conditions, measurements of real hardware, and synthesized workloads to determine the common case for our model and simulate the behavior of energy harvesting sensors. From these data, our model produces estimates of energy utilization, availability, responsiveness, and lifetime.

Environmental Conditions. We assume an indoor environment that is built for and used by people. 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. Even the example applications of intermittent, energy harvesting systems are nearly all centered around monitoring indoor and human-centric phenomena [8, 10, 16, 18]. We expect our environment to be lit, and that it may occasionally get direct or indirect sunlight. We use indoor photovoltaic energy harvesting because under these conditions, it offers an order of magnitude more energy than other methods. This does not mean that the conclusions we draw are not applicable to other environments and harvesting methods, but that the sizing and lifetime conclusions may be different.

To model the harvestable light in an occupied indoor environment, we use the EnHANTS indoor irradiance dataset [14]. We find that it is the most complete and extensive dataset for indoor light irradiance traces, capturing over a year of data in several situations. The traces we use for our model are summarized in Table 1.

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 solar energy harvesting sensor platform that utilizes available, state-of-the-art, commercial components. 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 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 1. The physical size of the solar panel used by this sensor is assumed to be 10.9 cm² and the volume of the sensor node is similar to prior work like the Hamilton sensor [24]. We implement this design and describe it further in Section 8.

Representative Workloads. We find that sensing workloads generally fall into three categories: (i) periodic sense-and-send [31], (ii) reactive event detection [8], and (iii) infrequent, long-running, high-power events [27]. We choose a representative workload for each of these categories to use in our model. We characterize our “periodic sense-and-send” workload as periodically sampling a light and color sensor and sending a BLE advertisement containing the data. Our “reactive workload” is represented by sending a BLE advertisement upon motion detection of the main entrance of a university building, and we linearly scale the frequency of these

Irradiance Trace	Total Days	Average Power ($\mu\text{W}/\text{cm}^2$)	90 th Percentile Daily Power ($\mu\text{W}/\text{cm}^2$)	10 th Percentile Daily Power ($\mu\text{W}/\text{cm}^2$)
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 8.

^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 1: Representative harvesting conditions and workloads. To evaluate different energy storage architectures, we define a set of energy harvesting conditions and workloads that are representative of common sensing applications. We choose two real, 1 Hz, irradiance traces with different magnitudes of available energy. 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 8.

events to represent varying amounts of usage. We treat these workloads as atomic. 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 conservatively assume these long running tasks can be interrupted and resumed at any point during execution and that any checkpointing is free.

A Predictive Model. 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 source¹ numerical model that allows parameterization of various system characteristics, including regulator efficiency, solar harvester size and efficiency, energy storage capacity, leakage, ESR, and charge-discharge efficiency. These parameters are summarized in Table 2.

The simulation of our model operates as a second-by-second calculation of the amount of energy entering and exiting a device. At every step, the simulation calculates the net energy gain or loss of the system based on its current state 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 workload) or from a random distribution (reactive event detection or a high-power event). For our modeling, workload schedules are generated from values listed in Table 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 events’

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 2: Simulation configuration parameters. A representative set of available configuration options for our simulation of a sensor with secondary storage and 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`.

time to completion are collected, and if applicable, the primary-cell lifetime is estimated from a trace of its state of charge.

During simulation, modeled devices can be online or offline and idle or performing work. These states are shown in Figure 2. 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 2. 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 the assumption that the duration of atomic events are less than the one second simulation step. We 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 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.

5 CAPACITY INCREASES CAPABILITY

In addition to relaxing the need for intermittent techniques, we find that an increase in sensor energy capacity achieves a significant

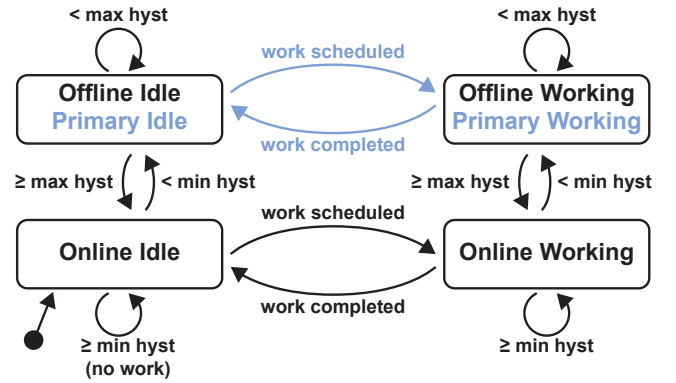


Figure 2: Model state machine. A modeled device can be in one of four states: Offline Idle, Online Idle, Online Working, and Offline Working. When a device is Offline Idle, it has run out of energy and is off. If a device is Online Idle, it is on and in deep sleep, ready to perform work if triggered. 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.

return on sensor reliability and capability in the face of variable energy income. Higher energy capacity allows higher energy capture during periods of abundant harvestable energy.

5.1 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 store 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, which 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 [8]. For an example of the second scenario, consider 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 [10, 16]. Another example of this are sensors in which the harvesting rate is proportional to the sensed phenomena [12]. While compelling for their simplicity, these sensors use energy that could be saved and used later for more useful and relevant tasks.

To explore ambient energy utilization as a function of storage size, we model the charge-discharge patterns of idealized energy stores under the harvesting conditions and workloads described in Section 4. 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




Figure 3: 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 Section 4 and Table 1. The x-axis is split by energy capacities possible with capacitors, supercapacitors, and batteries. The upper extents of capacity for (super)capacitors represent ten 100 μF tantalum capacitors [4], and one large 220 mF supercapacitor [35]. Larger capacitors exist, but are not appropriate for use on a small sensor node. As energy storage increases, the harvestable energy used in the application also increases, implying increased application performance. Some scenarios, such as the periodic 30 s, 15.1 $\mu\text{W}/\text{cm}^2$ case, reach 100% utilization at high secondary capacities indicating that available energy is not sufficient 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.

attempt to maximize the success rates for the specified sensing tasks. The results of this modeling are shown in Figure 3. No matter the amount of available energy, utilization increases as storage capacity increases. For scenarios with low harvesting potential and high power workloads, a sensor with at least 1 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. Small increases in utilization have significant impact on reliability and system lifetime.

5.2 Application Reliability

We also model the ability of varying energy stores to meet our defined sensing tasks and intensities. Our results are shown in Figure 4. Similar to the results of Section 5.1, we see marked increases in performance when energy stores reach 1-10 mWh of capacity. Simulations experience 100% reliability for all but the most energy constrained scenarios with high power workloads and ≥ 2 mWh of energy storage. We also see that even when low capacity sensors have large energy harvesting potentials and infrequent workloads, they experience low reliability. This is because they do not have enough storage to keep sensing throughout the night.

Finally, we analyze the ability of different storage configurations to perform a random, contiguous, higher energy task. We use a 100 mJ (30 μWh) event that is representative of a 50 KB over-the-air code update. While this is larger than a code update would be for simple programs, we see in Figure 4d that nearly all of the configurations with 0.28 mWh of energy storage complete the task in the minimum time (5 s). In comparison, even reducing the amount of energy storage to match the amount of energy required for the code updates causes significant latency. It is clear that 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 solar power.

6 RELIABILITY REQUIRES BACKUP

Increasing secondary capacity greatly improves reliability. However, some environmental conditions and workloads do not reach 100% reliability regardless of the size of the secondary store. Some results in Figure 4 appear to achieve perfect reliability, but still miss 0.1-2% of events. Others achieve well below 100% reliability, and simply require much more energy than is harvestable. A backup energy store can increase reliability of a system to 100%, at the cost of a finite lifetime.

6.1 Reliability Required

We argue that 100% reliability is a significant improvement over even low failure rates with respect to reliability and simplicity due to the lack of intermittency. Many applications, especially human facing ones, must be reliable to function, and research shows that unreliability leads to frustration and unwillingness to adopt automated solutions [7, 13, 41]. To use energy harvesting sensors for control or feedback of systems with potential safety issues, inherent unreliability is intolerable.

Worse, intermittent system failures are difficult to detect because there is no method for distinguishing 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 operation such as Flicker [16], Gecko [45], and some configurations of Cappybara [10].

Finally, it is more difficult to program intermittent systems because programmers or the underlying programming model must monitor and adapt to available energy with fine granularity, both of which are non-trivial tasks. These systems have little ability to correct for failures even when they are detected.

6.2 Lifetime of a Backup Energy Store

To achieve 100% reliability and avoid intermittency, designs can utilize a backup energy store. In instances where the rechargeable source is depleted, the system can operate from the backup, masking the effects of variable energy income. When the backup energy store

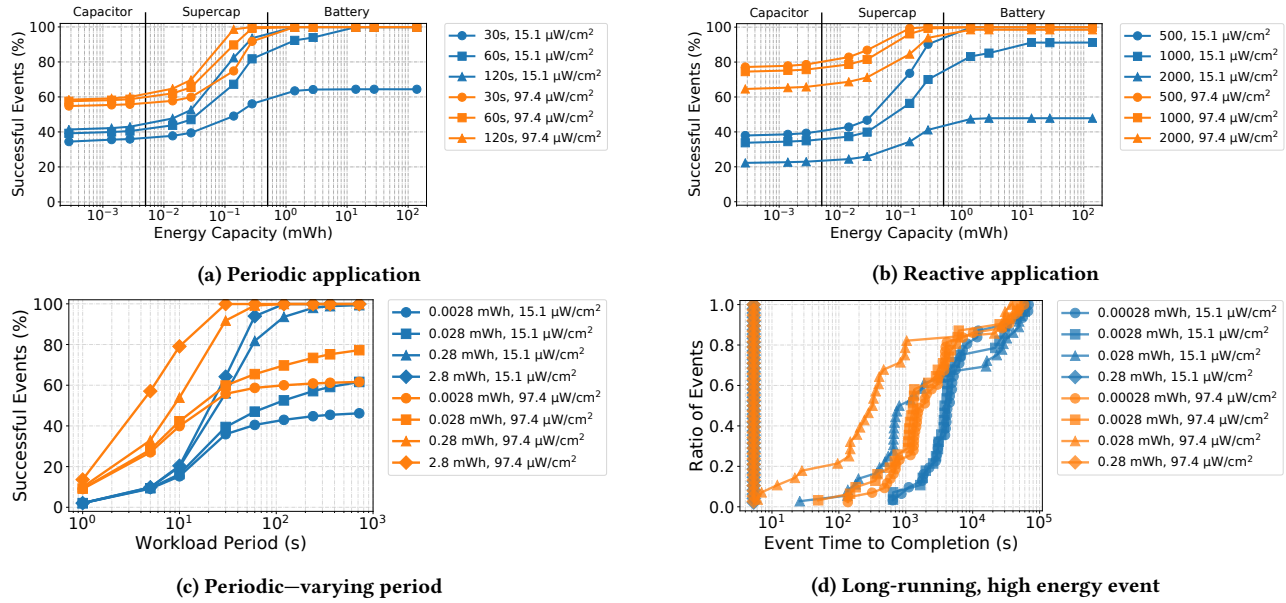


Figure 4: Workload reliability for different harvesting scenarios, workloads, and idealized secondary storage sizes. We define reliability as the percentage of successfully completed events. As expected, workload reliability follows the same trend as energy utilization, improving with increased secondary energy storage. For both periodic and reactive workloads, from the smallest to largest capacity simulated, we see a 1.4-2.7x improvement in availability. In (c) we investigate the period at which different secondary storage sizes meet a specific reliability, showing that even with infrequent periodic workloads, small amounts of secondary storage have low reliability while larger secondary stores reach near 100% reliability. (d) shows a CDF of time to completion for events in the long-running workload. 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.

is depleted, we consider the node’s lifetime to be complete, although it could continue operating intermittently and with lower reliability. This energy store should be a primary-cell, as they offer very low self-discharge, long shelf life, and substantial energy density.

An analysis of the reliable lifetime of a node with both energy harvesting and a backup energy store is shown in Figure 5. 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 reliability in uncharacteristically adverse conditions. Even for conditions with limited energy availability we still observe significant lifetime improvements due to energy harvesting.

7 OPTIONS FOR CAPACITY

Capacitors and supercapacitors have been the preferred option for storing energy in energy harvesting systems due to their purported indefinite lifetime. Batteries have been largely abandoned by energy harvesting researchers even though they offer performance and lifetime benefits. Many intermittent systems papers have dismissed batteries as expensive [15–18], temperature-sensitive [10, 15–18,

29], less efficient [15–18], bulky [15–18, 45], dangerous [15–18], and short-lived [10, 15–18, 29, 45].

Due to recent developments in both battery technology and management techniques, we believe many of these claims may no longer be true. Recent work has highlighted new battery chemistries such as LTO and LiFePo₄ that do not possess these shortcomings to the same extent as traditional lithium-based chemistries [21].

This work finds that 2-40 mAh LTO cells are available in small quantities for \$1 USD, are of similar or smaller size, and offer orders of magnitude more capacity than the capacitor configurations used in recent intermittent system designs. These batteries are somewhat temperature sensitive, but have similar temperature ratings to some supercapacitors and better temperature performance than traditional lithium-based chemistries. Specifically, they are sufficient for common temperatures expected for indoor and non-extreme outdoor spaces. While these battery chemistries possess more self discharge and ESR than ceramic and tantalum capacitors, they exhibit an order of magnitude lower self discharge than supercapacitors. At worse, the self discharge of LTO and LiFePo₄ chemistries are lower than the leakage of other sensor components [21].

More importantly, LTO and LiFePo₄ chemistries offer 3000-4000 full cycle lifetimes [19, 36, 44], and this increases exponentially with decreasing depth-of-discharge. By reducing depth-of-discharge to 10-20%, LiFePo₄ cells can achieve greater than 10,000 cycles before substantial cell degradation [36, 44]. We expect LTO cells will have




Figure 5: 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. 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. We see a 2-4x increase in lifetime estimates from the smallest to largest capacity simulated, if we only consider bounded results. 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 degradation that would be experienced over decades of use.

similar gains in lifetime. In addition, the capacity of LTO cells do not degrade when undervolted to zero volts, indicating that long term storage or absences in charging are not as destructive as they are with other lithium chemistries [6]. Finally, these new chemistries are safer than older lithium chemistries. They exhibit significantly lower thermal runaway under electrical, mechanical, and thermal stress [5, 26], and LTO cells have been shown to not release toxic gasses under temperature abuse like conventional lithium cells [5].

Finally, all batteries have the advantage of producing a stable voltage if they have any stored energy. This helps significantly with cold start issues caused by insufficient voltage, and prevents the waste of energy stored in the lower voltages of a capacitor or supercapacitor based system.

8 PERMAMOTE IMPLEMENTATION

We implement the design principals discussed in Sections 5 and 6 in a new sensor called Permamate. The Permamate sensing platform integrates a processor, BLE/802.15.4 radio, and various environmental, lighting, and occupancy sensors. A picture and system diagram of Permamate is shown in Figure 6. All hardware and software for the platform is open source².

Energy Harvesting and Storage. Permamate is powered by an energy harvesting front end that realizes the benefits of using batteries. It uses the TI BQ25505 energy harvesting IC, which harvests energy while monitoring both rechargeable and backup energy stores, switching between them at user-configurable voltages [43]. A 20 mAh (48 mWh) LTO battery is charged by an 10.9 cm² amorphous silicon solar panel [19, 20]. We limit the apparent capacity of this battery to ensure longer cycle lifetime as described in Section 7, but still have 24 mWh of energy storage, more than the capacity required to achieve the reliability and energy utilization improvements described in Section 5. For the backup energy store, Permamate uses lithium primary-cells which can be configured to 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 [34].

Processor, Radio and Sensor Selection. In designing Permamate, we search for the newest and lowest power components. To benefit other platform builders, we document our component selections along with their key performance metrics. A summary of these components can be found in Table 3.

We note our choice of the Nordic NRF52840 MCU over the more commonly used MSP430FR series because of its higher efficiency in active mode while offering comparable sleep currents. Specifically, it only draws 56 μ A/MHz compared to over 100 μ A/MHz for the MSP430. Unlike intermittent systems, we do not rely on the 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 an SoC-based design.

Energy Benchmarks. The data presented in Table 3 are benchmarks taken on the Permamate platform. We find that a BLE advertisement at 0 dbm consumes 86 μ J and that sampling both light and color sensors and transmitting them in a BLE advertisement consumes 586 μ J. Additionally, the entire system, including the energy harvesting front end, consumes only 5.0 μ W in deep sleep with RAM retained and all sensors powered off. We use the energy numbers from Permamate as a basis for our workloads to fairly compare against prior energy storage architectures.

9 EVALUATION

To evaluate the model, we perform a three-month-long deployment in a partially sunlit room using i) a primary-cell only system, ii) an intermittent, capacitor-only system, and iii) Permamate, our system that features both a secondary and primary-cell. We model these systems over the same period and compare the availability of Permamate to the intermittent system.

²<https://github.com/lab11/permamate/tree/master/hardware/permamate>




Figure 6: The Permamate power supply architecture is informed by the findings in Sections 5 and 6. 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.

9.1 Model Analysis

We analyze the deployment of these systems and compare their behavior to our model’s predictions: i) ten CR2032 primary-cell only devices, ii) an intermittent system configured with just 500 μ F of capacitance (about 0.36 μ Wh at 2.2 V), and iii) Permamate, 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 intermittent system sends a beacon as soon as its capacitor bank is full. When its energy is depleted, it powers off and charges again. Permamate is running the “sense and send” workload that we described in Section 4, and sends illuminance measurements every second. This workload stresses the model and requires more charge and discharge cycles. We use Permamate illuminance readings to estimate irradiance using the same scaling factors used by Yerva et al. [45], and use these traces as model input.

Primary-Cell Only. We model the workload of the primary-cell system and produce estimates for lifetime. Our model predicts the platform lifetime to be 58 days. We find that the average lifetime of the 10 devices is 61 days.

Intermittent. We model the number of packets sent each hour by the intermittent system over a three week period, and compare against the results of an actual device in Figure 8a. The average daily error of the model versus our results is 15%, with a standard deviation of 17%. This error can be 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. In the case of our estimation, results indicate that the model consistently underestimates high irradiance measurements.

Secondary and Primary-Cell. We compare our model’s predicted state of charge to a deployed Permamate over a seven day period in Figure 8b. 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

Component	Function	Active Power	Idle Power
Nordic NRF52840	Processor	56 μ A/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 μ A @ 12.5 Hz	50 nA
Maxim MAX44009	Light sensor	650 nA	N/A ^b
Intersil ISL29125	Color sensor	56 μ A	500 nA
Silicon Labs SI7021	Humidity sensor	1.5 μ A @ 1 Hz	60 nA
TE Connectivity MS5637	Pressure sensor	0.6 - 5 μ A @ 1 Hz	10 nA
Panasonic EKMB11011	PIR Occupancy	100 μ A	1 μ A

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

Table 3: The components used in Permamate. These components are among the lowest power options available, and are even 2-4x lower power than those used on relatively recent systems such as BLEES, Flicker, Capybara, and Hamilton.




Figure 7: Packets received over two days. This figure compares the reliability of an intermittent design and Permamate. Permamate sends a packet every second and does so without fail, while the intermittent system is only able to send when light is available.

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.

9.2 Permamate Performance

We also compare the performance of the deployed intermittent system and Permamate. In Figure 7, we show the number of packets sent per hour for two days. Permamate sends data every second, while the intermittent system sends as fast as possible. Permamate is able to collect and send its data continuously, while the intermittent 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.

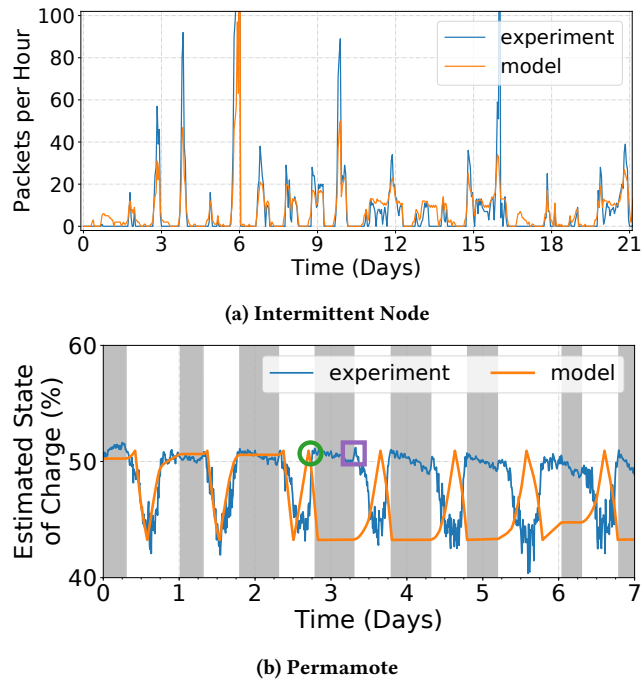


Figure 8: Model comparison to deployed hardware. Data from a three month deployment of two systems is used to verify our model. (a) We use three weeks of lux measurements to estimate irradiance and model the number of packets transmitted by an intermittent node. Average daily error is 15%, with a standard 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. Secondary charging hysteresis limits of the devices are set at 51% and 43%. Shaded regions represent periods of low harvestable potential ($< 15 \mu\text{W}/\text{cm}^2$), 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.

We also use our model to explore the estimated performance of Permamote compared to prior systems. 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. The results of this modeling are shown in [Table 4](#). 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 a less intense workload.

10 DISCUSSION

While we have verified that our model generally matches reality, many assumptions were made that impact its predictions. Specifically, lifetime estimates produced by the model can be described as

Platform	Successful Events (%)		Long-Running Time to Completion Ratio		Lifetime (yrs)
	Periodic	Reactive	Average	95th Percentile	
Telos [37]	100	100	1	1	8.55
Hamilton [24]	100	100	1	1	6.75
BLEES [1]	100	100	1	1	1.11
Gecko [45]	39.5	64.9	387	981	∞ ^g
Capybara [10] ^a	46.3	72.8	37.6	1	∞ ^g
Capybara [10] ^b	41.1	67.1	2730	8900	∞ ^g
Flicker [16]	39.3	64.2	1307	5670	∞ ^g
EnHANTs [32]	79.4	96.0	1	1	— ^h
DoubleDip [33]	77.9	66.5	1	1	— ^h
[39]	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 these systems.

ⁱ Expect cycling failure in 20-50 years, but do not attempt to estimate.

Table 4: Modeled performance of energy harvesting systems. 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). Generally, intermittent systems have significantly worse availability and responsiveness compared to battery-only systems and systems that use a secondary-cell. Battery-only systems achieve perfect operation, but have finite, sub-decade lifetimes.

optimistic. We are skeptical of the reality of these lifetimes when considering the myriad of factors not considered by our model.

Model Limitations. Many assumptions were made when designing the model and considering different parameters and behaviors. In particular, many dynamic parameters are modeled as static. For example, leakage is dependent on state of charge and age, but we assume a static leakage rate. Similarly, many components like solar panels and boost regulators have dynamic efficiencies, that we consider to be static. Nominal voltages for solar panels and batteries are used to calculate energy output and capacity, but they are really dynamic. In addition, we do not consider the effects of temperature and ignore cycle capacity degradation for supercapacitors and secondary-cells. We also assume that every platform can wake up and perform a task at any point in its voltage curve. In reality, many intermittent systems must recharge to a threshold voltage. We attempt to ensure that these assumptions favor intermittent systems over larger capacity designs so that our conclusions are not distorted by inaccuracies in the model.

Theoretically Infinite vs. Reality. We discuss capacitor-based systems as perpetual throughout the paper and even predict that Permamote will have 30-50 years lifetimes under some configurations and harvesting conditions. While in theory this is true, it is highly unlikely that any sensor nodes we build and deploy today will be functional, let alone relevant in 50 years. We expect physical degradation of not only energy storage components, but also the silicon itself. The MCU architectures in use on these systems have only been commercially available for 1-2 decades. This makes it

difficult to state that failure rates will remain low for over five. Components like MEMS sensors and oscillators will age and lose calibration well before this time. Outside of physical degradation, we must also consider the potential decline of common protocols and standards that dictate wireless and security functions, and special care will need to be taken to ensure long term operation [25]. Before advances in device and sensor technologies are realized, we expect that the increases in usable energy shown in this work may translate into sensor nodes that are smaller in size or more capable, rather than sensor nodes with multi-decade lifetimes.

11 CONCLUSIONS

Our results, along with recent advances in battery technology, suggest that we are now able to build energy harvesting sensors with greater energy storage capacity. These sensors will not need to micro-manage their energy state and can instead adapt their lifetime and energy usage over the course of weeks, months, or not at all. This push toward reliable and capable sensors will enable more standard programming models, energy management techniques, and ultimately useful applications to be built on top of dense and ubiquitous energy harvesting sensor deployments.

12 ACKNOWLEDGMENTS

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. 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. We would like to thank our anonymous reviewers for their insightful feedback, our shepherd, Koen Langendoen, for his extensive feedback and guidance on making the best possible representation of our work, and Will Huang and the rest of Lab11 for their thoughtful edits.

REFERENCES

- [1] J. Adkins, B. Campbelle, S. DeBruin, B. Ghena, B. Kempke, N. Klugman, Y.-s. Kuo, D. Natarajan, P. Pannuto, T. Zachariah, and others 2015. Demo: Michigan's IoT Toolkit (*SenSys '15*).
- [2] J. Adkins, B. Ghena, N. Jackson, P. Pannuto, S. Rohrer, B. Campbell, and P. Dutta 2018. The Signpost Platform for City-Scale Sensing (*IPSN'18*).
- [3] Analog Devices. ADP5091 Datasheet. <http://www.analog.com/media/en/technical-documentation/data-sheets/ADP5091-5092.pdf>. (2017).
- [4] AVX. TPS Series Capacitor Datasheet. <http://datasheets.avx.com/TPS.pdf>. (2018).
- [5] I. Belharouk, G. M. Koenig, and K. Amine, Electrochemistry and safety of Li4Ti5O12 and graphite anodes paired with LiMn2O4 for hybrid electric vehicle Li-ion battery applications. *Journal of Power Sources* 196, 23 (2011).
- [6] M. Brunell, B. Hanauer, M. Loveridge, R. Dashwood, and R. Bhagat 2016. Effect of Zero Volt Storage on Commercial Lithium Titanate Cells. In *Meeting Abstracts*.
- [7] A. B. Brush, B. Lee, R. Mahajan, S. Agarwal, S. Saroiu, and C. Dixon 2011. Home Automation in the Wild: Challenges and Opportunities (*CHI '11*).
- [8] B. Campbell and P. Dutta 2014. An Energy-harvesting Sensor Architecture and Toolkit for Building Monitoring and Event Detection (*BuildSys'14*).
- [9] A. Colin, G. Harvey, B. Lucia, and A. P. Sample, An energy-interference-free hardware-software debugger for intermittent energy-harvesting systems. *ACM SIGOPS Operating Systems Review* 50, 2 (2016).
- [10] A. Colin, E. Ruppel, and B. Lucia 2018. A Reconfigurable Energy Storage Architecture for Energy-harvesting Devices (*ASPLOS '18*).
- [11] P. Corke, P. Valencia, P. Si, T. Wark, and L. Overs 2007. Long-duration solar-powered wireless sensor networks (*SenSys'07*).
- [12] S. DeBruin, B. Campbell, and P. Dutta 2013. Monjolo: An energy-harvesting energy meter architecture (*SenSys'13*).
- [13] W. K. Edwards and R. E. Grinter 2001. At Home with Ubiquitous Computing: Seven Challenges (*UbiComp '01*).
- [14] M. Gorlatova, A. Wallwater, and G. Zussman, Networking low-power energy harvesting devices: Measurements and algorithms. *IEEE Transactions on Mobile Computing* 12, 9 (2013).
- [15] J. Hester, L. Sitanayah, and J. Sorber 2015. Tragedy of the Coulombs: Federating Energy Storage for Tiny, Intermittently-Powered Sensors (*SenSys '15*).
- [16] J. Hester and J. Sorber 2017. Flicker: Rapid Prototyping for the Batteryless Internet-of-Things (*SenSys '17*).
- [17] J. Hester and J. Sorber 2017. The Future of Sensing is Batteryless, Intermittent, and Awesome (*SenSys '17*).
- [18] J. Hester, K. Storer, and J. Sorber 2017. Timely Execution on Intermittently Powered Batteryless Sensors (*SenSys '17*).
- [19] HuaHui New Energy. LTO Battery Specification. <http://www.batteryspace.com/prod-specs/7455.pdf>. (2013).
- [20] HuaHui New Energy. LTO Battery Catalog. Self hosted. Removed for Anonymity.. (2018).
- [21] N. Jackson, J. Adkins, and P. Dutta 2018. Reconsidering Batteries in Energy Harvesting Sensing (*ENSys'18*).
- [22] X. Jiang, J. Polastre, and D. Culler 2005. Perpetual environmentally powered sensor networks (*IPSN'05*).
- [23] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems (TECS)* 6, 4 (2007).
- [24] H.-S. Kim, M. P. Andersen, K. Chen, S. Kumar, W. J. Zhao, K. Ma, and D. E. Culler 2018. System Architecture Directions for Post-SoC/32-bit Networked Sensors. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems. ACM*.
- [25] K. Kinningham, M. Horowitz, P. Levis, and D. Boneh 2016. CESEL: Securing a Mote for 20 Years (*EWSN '16*).
- [26] F. Larsson and B.-E. Mellander, Abuse by External Heating, Overcharge and Short Circuiting of Commercial Lithium-Ion Battery Cells. *Journal of The Electrochemical Society* 161, 10 (2014).
- [27] P. Levis, N. Patel, D. Culler, and S. Shenker 2004. Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks (*NSDI '04*).
- [28] K. Lin, J. Yu, J. Hsu, S. Zahedi, D. Lee, J. Friedman, A. Kansal, V. Raghunathan, and M. Srivastava 2005. HelioMote: enabling long-lived sensor networks through solar energy harvesting (*SenSys'05*).
- [29] B. Lucia, V. Balaji, A. Colin, K. Maeng, and E. Ruppel 2017. Intermittent Computing: Challenges and Opportunities. In *SNAPL*.
- [30] B. Lucia and B. Ransford, A simpler, safer programming and execution model for intermittent systems. *ACM SIGPLAN Notices* 50, 6 (2015).
- [31] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson 2002. Wireless sensor networks for habitat monitoring (*WSNA'02*).
- [32] R. Margolies, M. Gorlatova, J. Sarik, G. Stanje, J. Zhu, P. Miller, M. Szczodrak, B. Vignham, L. Carloni, P. Kinget, and others, Energy-harvesting active networked tags (EnHANTs): Prototyping and experimentation (*TOSN*).
- [33] P. Martin, Z. Charbiwala, and M. Srivastava 2012. DoubleDip: Leveraging thermoelectric harvesting for low power monitoring of sporadic water use (*SenSys'12*).
- [34] Maxim Integrated. MAX1722 Datasheet. <https://datasheets.maximintegrated.com/en/ds/MAX17220-MAX17225.pdf>. (2017).
- [35] Murata. DMF Series EDLCs. <https://www.murata.com/en-us/products/productdata/8796857270302/MFCDSF1E.pdf>. (2016).
- [36] N. Omar, M. A. Monem, Y. Firouz, J. Salminen, J. Smekens, O. Hegazy, H. Gaulous, G. Mulder, P. V. d. Bossche, T. Coosemans, and J. V. Mierlo, Lithium iron phosphate based battery—Assessment of the aging parameters and development of cycle life model. *Applied Energy* 113 (2014).
- [37] J. Polastre, R. Szewczyk, and D. Culler 2005. Telos: enabling ultra-low power wireless research (*IPSN'05*).
- [38] Pressac. Pressac CO2 Sensor Datasheet. <http://www.pressac.com/help/CO2TemperatureAndHumiditySensor.html>. (2017).
- [39] H. Raisigel, G. Chabanis, I. Ressejac, and M. Trouillon 2010. Autonomous wireless sensor node for building climate conditioning application (*SENSORCOMM'10*).
- [40] B. Ransford, J. Sorber, and K. Fu, Mementos: System support for long-running computation on RFID-scale devices. *Acm Sigplan Notices* 47, 4 (2012).
- [41] E. Shehan and W. K. Edwards 2007. Home Networking and HCI: What Hath God Wrought? (*CHI '07*).
- [42] Shenzhen Hibatt Technology. Mini LiFePO4 Battery. https://www.alibaba.com/product-detail/Mini-LiFePO4-battery-10130-30mAh-3_6071757578.html. (2018).
- [43] Texas Instruments. BQ25505 Datasheet. <http://www.ti.com/lit/ds/symlink/bq25505.pdf>. (2015).
- [44] J. Wang, P. Liu, J. Hicks-Garner, E. Sherman, S. Soukiazian, M. Verbrugge, H. Tatara, J. Musser, and P. Finamore, Cycle-life model for graphite-LiFePO4 cells. *Journal of Power Sources* 196, 8 (2011).
- [45] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta 2012. Grafting Energy-harvesting Leaves Onto the Sensor Tree (*IPSN'12*).