On the Limits of Effective Hybrid Micro-Energy Harvesting on Mobile CRFID Sensors

Jeremy Gummeson, Shane S. Clark, Kevin Fu, Deepak Ganesan
University of Massachusetts Amherst, Dept. of Computer Science
{gummeson,ssclark,kevinfu,dganesan}@cs.umass.edu

ABSTRACT

Mobile sensing is difficult without power. Emerging Computational RFIDs (CRFIDs) provide both sensing and general-purpose computation without batteries—instead relying on small capacitors charged by energy harvesting. CRFIDs have small form factors and consume less energy than traditional sensor motes. However, CRFIDs have yet to see widespread use because of limited autonomy and the propensity for frequent power loss as a result of the necessarily small capacitors that serve as a microcontroller's power supply. Our results show that hybrid harvesting CRFIDs, which use an ambient energy micro-harvester, can complete a variety of useful workloads—even in an environment with little ambient energy available.

Our contributions include (1) benchmarks demonstrating that micro-harvesting from ambient energy sources enables greater range and read rate, as well as autonomous operation by hybrid CRFIDs, (2) a measurement study that stresses the limits of effective ambient energy harvesting for diverse workloads, (3) application studies that demonstrate the benefits of hybrid CRFIDs, and (4) a trace-driven simulator to model and evaluate the expected behavior of a CRFID with different capacitor sizes and operating under varying conditions of mobility and solar energy harvesting. Our results show that ambient harvesting can triple the effective communication range of a CRFID, quadruple the read rate, and achieve 95% uptime in RAM retention mode despite long periods of low light.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Measurement Techniques

General Terms

Design, Experiment, Measurement, Performance

Keywords

RFID, Solar, Energy Harvesting, Sensing

MobiSys'10, June 15–18, 2010, San Francisco, California, USA.

1. INTRODUCTION

An emerging model of RFID requires tags that go beyond mere identification, also carrying out sensing, computation, and storage duties. Such Computational RFIDs (CRFIDs) [14] come equipped with ultra low-power microcontrollers and a suite of low-power sensors such as those found in the Intel Wireless Sensing Platform (WISP) [18]. A CRFID resembles a sensor mote stripped of its battery and active radio, but augmented with an RFID front end for RF energy harvesting and backscatter communication.

CRFIDs offer several advantages over battery-powered embedded sensor platforms [3, 14]. Because CRFIDs are expected to be continually charged by a remote energy source, they can rely on small capacitors rather than large batteries for energy storage. This difference in the size and weight of the energy store makes CRFIDs more amenable to SoC integration and low-cost manufacturing. For instance, the lightweight SoCWISP has been implanted in the wing muscles of a hawkmoth [12]. In addition, capacitors are cheaper and more environmentally friendly than batteries—making them better suited for large-scale deployments. Finally, capacitors can endure a large number of recharge cycles—important for the frequent, bursty charge cycles common to energy harvesting.

CRFIDs also offer extremely low power profiles, consuming roughly an order of magnitude less energy than motes for many operations. Communication, in particular, is very energy efficient in CRFIDs, as they leverage backscatter communication rather than using active radio transmitters [1]. In addition, CRFIDs are optimized to exploit limited harvested energy rather than continuous long-term operation; hence they use far fewer hardware components and consequently consume less power than mote-class devices.

Limitations. Despite the advantages of CRFIDs, several challenges remain before they are suitable for widespread use in ubiquitous computing and sensing applications. A key limitation of CRFIDs is the reliance on a dense deployment of RFID readers. CRFIDs must be placed such that each device harvests enough energy to compute and sense. A CRFID cannot sustain operation for long outside of the effective harvesting range of a reader; this lack of autonomy necessitates carefully planned deployments of RFID readers relative to tags, making such networks expensive to deploy and maintain. A CRFID's lack of autonomy is exacerbated by the observation that its performance varies with distance from a reader. For example, our experiments show that an Intel WISP [18] within a few feet of an RFID reader receives sufficient energy to sample and transmit hundreds of times

per *second*, whereas one that is near its maximum reliable distance (a few meters) may be able to sample and transmit only a few times per *minute*. This rapid performance degradation limits the efficacy of CRFIDs for applications requiring complex processing or frequent sensing.

Approach. Our work exploits supplemental energy harvested using a miniature solar panel attached to the Intel WISP (Figure 2) to make the WISP more autonomous and therefore more useful for mobile sensing and computation. Even a tiny amount of hybrid harvesting enables a CRFID to better tolerate power interruptions. Our measurements show that the SolarWISP prototype is able to retain memory state in all but the most difficult lighting conditions that we encountered. Modest modifications to the WISP enable ambient energy harvesting and provide the platform with new capabilities, but there are many issues associated with batteryless harvesting. The specific challenges raised are (1) how a small amount of ambient energy can be used most effectively to enable autonomous operation and (2) how to select appropriate workloads. To determine appropriate component sizes and applications, we performed an application study and built a trace-driven software simulator that predicts the behavior of a CRFID's capacitor under an assortment of sensing and computation workloads.

1.1 Design Considerations

There are several design considerations for CRFIDs that leverage hybrid energy harvesting. In this section, we present some of these considerations, as well as the key questions that motivate the approaches in this paper.

Using limited energy. A central challenge in designing hybrid harvesting CRFIDs is dealing with the limited amount of harvested energy. CRFIDs have smaller footprints than motes, restricting harvesting units to only a few cubic centimeters in size. Micro harvesters of this size typically generate very little power. For example, our experiments show that solar harvesting provides power ranging from less than a μW to a few mW depending on the panel size and lighting conditions. Thus, the first question that we ask is:
Question 1: *What are the lower limits of usable energy for current CRFIDs?*

Handling harvesting variability. Ambient energy harvesting exhibits spatial and temporal variability, sometimes over very short time scales. Harvesting output can be relatively predictable for a static outdoor deployment, for example, but unpredictable and bursty when CRFIDs are attached to mobile objects or persons constantly moving past occlusions and altering the panel's orientation. Because CRFIDs use capacitors with potentially low storage capacity, even a few seconds without any incoming energy can have disastrous consequences. The second question that we address is then:
Question 2: *What impact do dynamics in the ambient energy source have on hybrid harvesting CRFIDs?*

Making component choices. Capacitor size plays an important role in the design of a CRFID system. Larger capacitors are desirable because they can more effectively buffer excess harvested energy, allowing a CRFID to survive periods when there is no power from a light source or reader. Smaller capacitors have the advantage of short charge times. This leads to our third question:
Question 3: *What capacitor size maximizes CRFID performance for a given application?*

Choosing suitable workloads. While CRFIDs have unique advantages for some applications, they have no clear benefit for others. Past work has focused almost entirely on applications requiring low throughput communication or near-constant reader interactions [4, 25]. No work has yet been done with applications specifically suited to hybrid harvesting CRFIDs. The final question is then:
Question 4: *For what applications are hybrid harvesting CRFIDs most suitable?*

Motivated by the questions above, we aim to quantify the performance of hybrid harvesting CRFIDs and determine an appropriate set of environmental conditions and use cases where these devices prosper while passive CRFIDs (such as WISPs) and active battery-powered sensors (such as motes) may be less appropriate.

1.2 Goals & Contributions

Our research seeks to discover how to design hybrid energy harvesting CRFID platforms that match hardware choices with application workloads. Thus, our work aims to:

- Determine the boundary between useful and unworkable energy for hybrid CRFIDs.
- Explore the design tradeoffs in hardware components that influence CRFID performance.
- Understand the effectiveness of CRFIDs for an appropriate set of applications.

Our contributions that address these goals include:

Tools. We created two tools to assist in the evaluation of CRFIDs. Our software contribution is the trace-driven CRFID Crash Test Simulator (CCTS) that models the expected behavior of a CRFID's capacitor under varying workloads with solar harvesting. Our hardware contributions are the SolarWISP and FrankenWISP. The SolarWISP allows us to test the effectiveness of miniature solar panels for CRFID energy harvesting and the FrankenWISP passively measures the capacitor voltage of a CRFID during application execution.

Measurements. A comprehensive measurement study is important in determining how miniature solar panels can be used to improve the performance of passive CRFIDs. We use our software and hardware tools to evaluate how solar harvesting improves CRFID performance by: (a) demonstrating that energy harvested from ambient lighting is sufficient for essential device operations, (b) quantifying how solar harvesting improves the communication range of a conventional WISP, (c) comparing SolarWISP performance against a micro-harvesting, mote-class platform under low-light conditions, (d) determining how capacitor size impacts platform responsiveness and survivability, and (e) showing how design decisions may be quickly evaluated using CCTS for a variety of gathered light traces.

Applications. We present an end-to-end study of two applications that are enabled by ambient solar harvesting. *Path Reconstruction* is an application that needs to communicate frequently with readers while harvesting minimal amounts energy from dynamic, indoor lighting sources. *Greenhouse Monitoring* is a sensing application that relies on the

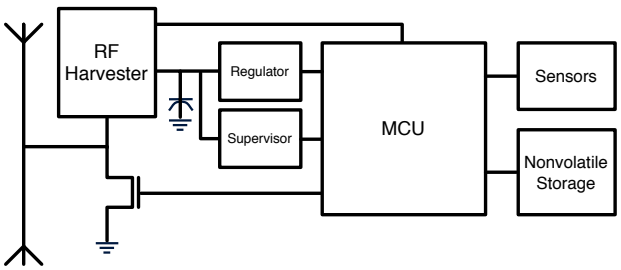

Figure 1: A CRFID strives to combine the versatility of a microcomputer with the energy efficiency of an RFID. With sensing capabilities, CRFIDs enable a host of new ubiquitous applications that go beyond mere identification.

predictable nature of static, outdoor deployments to buffer large amounts of energy in order to survive long harvesting outages. These applications are designed to stress two extremes along the spectrum of possible lighting variations. We use a combination of empirical and simulation results to show that CCTS can be used as a guide to size components for hybrid-powered CRFIDs given a particular application and expected lighting environment.

2. A CRFID PRIMER

Traditional passive RFIDs use small amounts of harvested RF energy from reader infrastructure to report static identifiers. Computational RFIDs, or CRFIDs, use the same basic operating principles, but provide general-purpose computation. The core components of CRFIDs are as follows (see Figure 1 for a block-level illustration):

- **RF harvester:** CRFIDs harvest small amounts of RF energy provided by an RFID reader. The energy is rectified to produce DC voltage and boosted to an appropriate level by a charge pump The RF harvester also includes an analog comparator circuit to decode reader transmissions.
- **Backscatter circuit:** To transmit data, a CRFID toggles the state of a transistor to detune its antenna and reflect a different signal back to the reader. This method of transmission requires little energy from the tag when compared to active communication circuits.
- **Storage capacitor:** Energy harvested from a reader needs to be accumulated before any computation or sensing can begin. A CRFID uses a small storage capacitor for this purpose.
- **MCU:** A microcontroller is used to participate in the RFID protocol, as well as to perform arbitrary computation. Ultra low-power MCUs are limited in capability but provide fine-grained duty-cycling options with low power state transition costs.
- **Supervisor:** The microcontroller does not know the state of its capacitor, thus additional hardware is required for energy awareness. A supervisor circuit generates an interrupt or releases the MCU from a reset state after climbing above a preset voltage threshold.

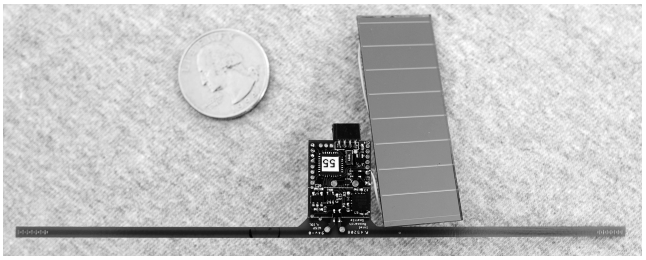

Figure 2: The WISP 4.1 with attached solar panel.

- **Regulator:** The voltage stored on a CRFID's capacitor is highly variable. A regulation circuit is used to provide a stable supply voltage for the MCU.
- **Sensors:** Sensors give a CRFID awareness of the world around it, unlike a traditional RFID. Sensors that are considered low power for other devices can dominate the power budget of a CRFID.
- **Nonvolatile storage:** Frequent power failure is common for CRFIDs, making nonvolatile storage a necessity. MCUs typically contain a small amount of onboard storage, but this may be augmented with offchip storage.

3. TOOLS AND TRACES

It is challenging to debug CRFID applications and evaluate their performance. Traditional methods of logging platform behavior are ineffective, as current CRFIDs possess no serial port for debugging and insufficient nonvolatile storage to create local logs of any significant length. Any attempt to actively log performance metrics will of course affect the energy consumption as well, which is unacceptable given CRFIDs' typically tiny energy store. We have developed two key tools that allow us to evaluate the effectiveness of ambient energy harvesting for CRFIDs. First, we present the FrankenWISP, a tool used to measure energy consumption and production of CRFIDs in several different hardware and software configurations. We also present CCTS, a trace-driven simulation tool that allows a system designer to quickly test a variety of hardware and software configurations against real harvesting traces. Finally, we present the illuminance traces gathered to carry out application studies.

3.1 The FrankenWISP

To overcome the introspective deficiencies of CRFIDs, we constructed the FrankenWISP (see Figure 3), which consists of a WISP observed by a TelosB mote. The TelosB, with the addition of a simple high-resistance voltage divider circuit, allows measurement of a WISP's capacitor voltage as a target application executes. The TelosB's 1MB of flash memory allows time-stamped voltage readings to be gathered at a frequency of 20 Hz for a duration of 2.5 hours, or at a lower rate for longer intervals. Depending on the size of the capacitor being measured, an appropriate sampling rate can be determined. The ability to gather runtime voltage traces during the deployment of an application is invaluable — these traces allow us to evaluate applications using differently sized capacitors and determine the length and number

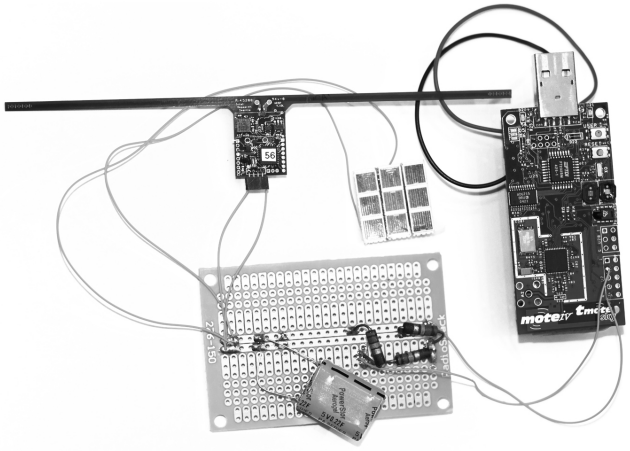

Figure 3: The FrankenWISP measures the supply voltage of a WISP during application execution without affecting the outcome or excessively restricting device mobility and deployability.

of power outages experienced by the WISP without affecting the outcome. We chose to construct the FrankenWISP in order to have a platform small enough for mobile test deployments. Oscilloscopes and Data Acquisition Units are too large to allow for realistic mobile testing.

3.2 Trace-Driven Simulator

We created the CRFID Crash Test Simulator (CCTS) in order to more fully explore the space of possible parameters and applications for CRFIDs — particularly the points of failure. It is difficult to profile a platform with power production and consumption values as small as those of CRFIDs. The difference of a microwatt is significant in many cases, leaving little room for error. Additionally, solar-powered devices that use capacitors for energy storage require the modeling of non-linear solar panel and capacitor outputs. Some platforms also use linear regulators, yet another non-linear component to model. With these modeling difficulties in mind, the goal for CCTS is to track the charge and discharge of the storage capacitor *closely enough to inform platform provisioning decisions.*

CCTS requires a number of parameters to adequately describe the platform and application to evaluate:

- A recorded or generated illuminance trace
- Maximum and minimum VI curve approximations
- Capacitor size and initial voltage
- Application power consumption characteristics

With the exception of a recorded illuminance trace, each of these inputs can be found in a datasheet or measured empirically with little more than a multimeter.

Internally, CCTS models the platform and mobility in the following logical order:

Illuminance trace. CCTS illuminance traces are time-stamped lists of lux readings used to estimate solar panel output. The simulator looks up the most recent illuminance measurement as time progresses and provides this as input to the simulated solar panel.

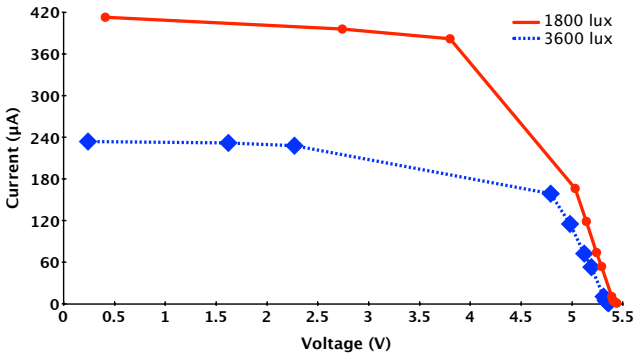

Figure 4: Example solar panel VI curves. Note that both curves are non-linear and that the relationship between the two varies greatly with voltage output.

A potential limitation with the use of illuminance traces as input is that illuminance is weighted toward portions of the spectrum perceptible by humans using the luminosity function. Irradiance, measured in W/m^2, would be the ideal metric for light intensity but we did not have ready access to any compact device capable of measuring irradiance. In order to keep conversions from illuminance to power output from being overly optimistic, we make the assumption that there is very little light available outside of humans' perceptible range. This assumption makes conversions from sources such as fluorescent bulbs the most accurate, while underestimating the power available from sources with greater bandwidth, such as incandescent bulbs or the sun. In practice, this inaccuracy may not have much of an impact because the efficiency of a given solar cell also varies with wavelength and it is impractical to model this efficiency curve for a large subset of those available.

Harvesting unit. The harvester is characterized using maximum and minimum functions to approximate the solar panel's family of VI curves. Given an illuminance value and voltage, CCTS interpolates between the two curves and estimates a current output value. We chose this approach because of the relationship between a solar panel's voltage and current output. As the load resistance of the platform increases, the solar panel's voltage output approaches its maximum and the current output approaches its minimum. A different VI curve exists for each light level that the panel could possibly observe and the curves do not simply scale with incident radiation. Their shapes change as well. See Figure 4 for an example of two curves from the same family. This behavior makes it clear that simply scaling a single curve is insufficient, but it is also infeasible to measure an infinite family of curves. To balance accuracy and usability, CCTS makes the simplifying assumption that a solar panel's output will scale linearly between that predicted by a minimum and maximum curve.

Storage capacitor. The energy output by the solar panel is next applied to the storage capacitor. The capacitor is the most difficult component to model accurately because of its dynamism and complex behavior. Changes in voltage on a capacitor depend on present voltage, effective resistance of the load, incoming energy, capacitor size, and internal leakage. None of these factors are modeled precisely in our

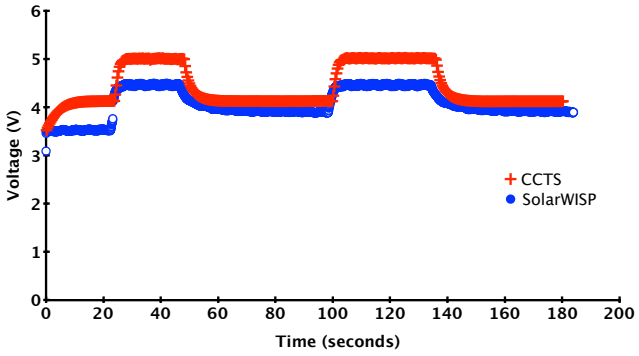

Figure 5: The two series in this plot are of a SolarWISP (as measured by the FrankenWISP) and a voltage trace produced by CCTS using a concurrently gathered illuminance trace. Notice that CCTS consistently overestimates the voltage, but the shape of the curve is similar. We believe that CCTS represents an upper-bound on empirical performance.

discrete-time simulator. In order to approximate continuous behavior, the simulator re-calculates the voltage on the capacitor at a rate of 10 kHz using the standard equation for voltage at time i given a load R (which scales with present voltage) applied for t seconds:

$$V(i) \quad = \quad V_f + Ae^{-t/RC}$$

t is equal to Δi so that all changes in load are taken into account. The effective resistance of the load, R, is calculated using the voltage at time $i - \Delta i$ and the difference between incoming and outgoing current. The constant A is similarly recalculated based on conditions at time $i - \Delta i$. Internal capacitor leakage is difficult to even approximate, as it depends on factors including: temperature, present voltage, time stable at present voltage, storage capacity, and capacitor type (ceramic, electrolytic, etc.). CCTS does not yet model internal leakage.

Platform power states. Finally, CCTS removes energy from the storage capacitor based on the platform's consumption. Platform power state information is a time-stamped list of operations used in the application and their associated consumption data. At runtime, CCTS applies the appropriate current drain for the given operation. Modeling the WISP, for example, required that we measure EEPROM operations, computation, and two low-power modes. Unfortunately the WISP uses a linear voltage regulator, which changes current consumption dependent upon voltage. Rather than using a single consumption number, we chose to measure current consumption at a variety of voltages and to use a regression function to scale the simulated consumption with voltage.

3.2.1 Simulator Validation

In order to gauge the accuracy of CCTS, we compared simulated versus actual SolarWISP performance. As previously noted, CCTS does not attempt to achieve perfect accuracy. The goal for CCTS is to produce results accurate enough to inform platform provisioning decisions. The

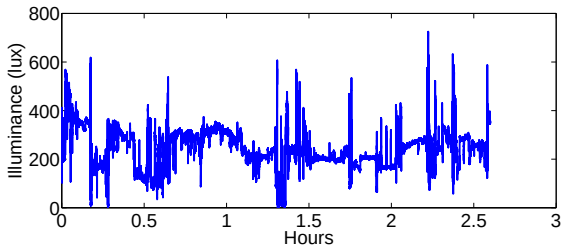

Figure 6: Mobile indoor environments are challenging for hybrid harvesting CRFIDs. Illuminance levels vary rapidly because light sources are localized and the mean illuminance is much lower than that typically observed outdoors (see Figure 7).

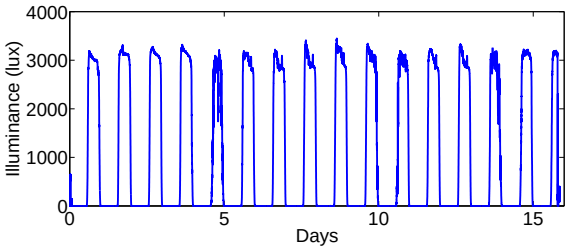

Figure 7: Outdoor environments exhibit diurnal and weather-induced variations in light level, but are generally much more amenable to energy harvesting than indoor environments. The main challenge for outdoor deployments is surviving the night, which is ∼11.5 hours in length in this trace.

comparison was performed using a FrankenWISP to monitor a SolarWISP continuously while it sat stationary and the lights were turned on and off to vary incoming energy. While this trace was captured, a TelosB mote was positioned as close as possible without shading the WISP's solar panel and set to sample illuminance continuously. Despite being placed close together, it is difficult to measure the illuminance at exactly the same angle and position as the solar panel. Other potential sources of error include the accuracy of the TelosB's illuminance sensor. We chose to use a 100 μF capacitor for this experiment.

As Figure 5 shows, CCTS consistently overestimates the voltage on the capacitor but stays within ∼0.5 V of the empirical result. Because CCTS consistently overestimates the voltage, we believe that it is an upper-bound for an empirical deployment.

3.3 Real-World Illuminance Traces

To ensure that our platform benchmarks and simulation results translate into real-world performance, we gathered a set of illuminance traces from diverse environments, including mobile indoor and outdoor deployments from all times of day. These traces allow us to calculate the amount of solar energy available to a CRFID in each scenario. Each trace was captured using a TelosB's photodiode to take periodic illuminance measurements. For mobile traces, the TelosB was carried on the shoulder. To accurately capture mobility-induced dynamics, all mobile traces were collected at a sampling rate of 20 Hz, providing a maximum trace

Environment	Trace length	Mean / Std. dev.
office1	2.5 hours	241.3 / 78.82 lux
office2	2.5 hours	34.71 / 24.25 lux
residential1	2.5 hours	49 / 62.91 lux
residential2	2.5 hours	124 / 161 lux
campus	15 minutes	3200 / 97.62 lux
yard	15 days	1076 / 1407 lux

Table 1: The length, mean, and standard deviation for each illuminance trace. The standard deviation varies widely based on mobility and diurnal cycles. See Section 3.3 for more detailed discussion.

length of 2.5 hours given TelosB flash memory constraints. The static outdoor trace was collected using a sampling period of 30 seconds; this is suitable because of slower dynamics and allowed the capture of a trace 15 days in length. See Table 1 for a summary of the traces, with mean and standard deviation values for each illuminance trace. The traces fall into a few rough qualitative categories.

Indoor office traces. The first set of traces, *office1* and *office2*, represents two different office buildings. In each case, the TelosB was carried during daylight hours. These mobile scenarios illustrate how changes in panel orientation and location affect harvestable light. Figure 6 shows a timeseries plot of the trace *office1*. The dataset shows short-lived illuminance peaks as high as 700 lux, with long-lived smoother variations between 200–300 lux. There are also several outage periods during which no significant amount of light is available. On average, this environment provides 241 lux. Dynamics for the *office2* trace were similar, although the overall illuminance was much less — on average, only 35 lux was available with peaks as high as 294 lux. The *office2* trace also has a much lower standard deviation (see Table 1), making it a much more consistent environment, though there is little available light.

Indoor residential. The second set of traces, *residential1* and *residential2*, gauges available light in a residential environment. The same experimental setup was used as in the *office* traces. This environment is of interest because there is little light available from incandescent bulbs but ample light available from windows. The *residential1* trace shows illuminance peaks of 1091 lux when the FrankenWISP approached windows, but a mean of only 49 lux over the duration of the deployment. Periods of insignificant available light were also commonplace, indicating a challenging harvesting environment. The *residential2* trace has roughly double the mean illuminance.

Outdoor mobile. There is one mobile trace gathered outdoors — labeled as *campus*. The trace chronicles a 15-minute walk, which includes periods when the sensor was obstructed by occlusions such as trees and buildings. Of particular interest is the availability of much more intense light available outdoors. The photodiode observed illuminance peaks of up to 3361 lux and a mean of 3200 lux — an order of magnitude greater than any of the indoor deployments. The second point of interest is the fact that mobility dynamics are smoothed out by large amounts of available diffuse light. The standard deviation is only 97.62 lux, among the lowest in any trace.

Outdoor static. The final trace, *yard*, uses a static sensor with a low sampling rate placed outdoors. It is intended to capture dynamics primarily induced by diurnal and weather variations. To assist in correlating light variations with weather patterns, we took note of conditions during the deployment. Our data represent days ranging from cloudy to sunny as well as several rainstorms. Nights, as defined by periods with no measurable illuminance, average $\sim$ 11.5 hours in this trace. In this variety of lighting conditions, Figure 7 shows an average of 1076 lux with peaks as high as 3445 lux. The standard deviation in this trace is by far the highest at 1407 lux.

4. EVALUATION

We now evaluate our hybrid CRFID prototype using platform benchmarks, real-world illuminance traces, and two application studies supplemented with CCTS simulation results. The first part of our evaluation explores several new capabilities enabled by ambient energy harvesting and quantifies them with platform benchmarks. Our measurements show that solar harvesting extends a WISP's effective communication range from 2 m to more than 7 m. We also show that the SolarWISP prototype can achieve perpetual with local timekeeping at an illuminance of 35 lux, while a solar harvesting mote prototype requires 200 lux to achieve perpetual operation. Next, we look at the implications of harvesting dynamics on platform performance with a simulation study. Our trace-driven simulation results show that the SolarWISP, even with the smallest capacitor size, can achieve nearly 95% uptime in RAM retention mode during a period of low light, while a standard WISP would fail completely in fewer than 15 seconds in the absence of an RFID reader. After examining the impact of harvesting dynamics, we provide a table to help platform designers predict response time and read rate. Finally, we present two case studies to validate our recommendations.

4.1 Is Micro-Harvesting Sufficient for CRFIDs?

Harvesting energy from an ambient source, such as indoor lighting, is fundamentally different than intentionally delivering RF energy to a conventional RFID tag. Most important to consider when assessing the viability of solar harvesting for a CRFID are scenarios where a CRFID is disconnected from reader infrastructure.

Energy harvesting CRFIDs, such as the SolarWISP, must provide a reasonable level of performance when operating autonomously and must simultaneously be energy efficient enough to survive short-lived interruptions to harvestable energy. To assess whether the SolarWISP is viable for typical sensing and computational workloads, we present microbenchmarks that quantify the performance of critical system components. Additionally, we show that solar harvesting increases the effective communication rate and range beyond that of a conventional tag, thus improving performance for basic identification applications. Finally, we present a comparison demonstrating that the SolarWISP can achieve higher uptime and communication rates than a mote prototype that uses a battery for energy storage.

4.1.1 Micro-Energy Harvester Benchmarks

To evaluate the effectiveness of solar harvesting for CRFIDs, we first benchmarked the SolarWISP's small photovoltaic cell (see Table 2). The particular panel used for

Conditions	Illuminance	Harvested power
Full shading	28 lux	6.6 μW
Partial shading	85 lux	35.9 μW
Diffuse	340 lux	62.5 μW
Direct	1300 lux	192.0 μW

Table 2: The amount of actual harvested power depends greatly on illuminance. A fully-shaded 11.4 cm^2 solar panel produces 29 times less power than the same panel under bright indoor lighting conditions.

Power state	Power draw	Energy consumption
Active	467.1 μW	-
LPM3	4.5 μW	-
ADC read	-	0.244 μJ
EEPROM read	-	0.216 μJ
EEPROM write	-	0.125 μJ

Table 3: WISP power consumption benchmarks. LPM3 is the lowest power mode that allows the WISP to maintain a persistent clock. The WISP can achieve a long lifetime despite limited energy reserves by leveraging low-power states.

this measurement study is optimized for artificial, indoor light sources. The power harvested varies widely with illuminance. In an indoor setting, the panel produces 6.6 μW of power while under full shading (28 lux) and 192.0 μW of power while under bright indoor light (1300 lux).

4.1.2 Computation, Sensing and Storage Benchmarks

Micro-power harvesting and capacitor-based energy storage together provide small amounts of buffered energy. This harvested energy is primarily used for three CRFID tasks: computation, sensing and storage. To evaluate the costs of these tasks, we benchmarked computation cost in terms of platform power consumption while in different power states and evaluated sensing and storage in terms of the energy required for an individual operation. A summary of these benchmarks appears in Table 3. The WISP consumes 4.5 μW in the lowest power mode that allows internal clocks (LPM3). This power state is sufficiently low to achieve perpetual operation while harvesting at a light level of 28 lux. The platform consumes 467 μW of power while in a fully active state. By combining low-power states with periods of active operation, the platform can achieve duty-cycles of 0.5% to 40.5% for illuminance values ranging from 28 to 1300 lux. Sensor readings require the WISP to sample the MCU's 12-bit ADC, with each read costing 0.244 μJ. Storage is accomplished by writing data over the MCU's I^2C bus to an off-chip EEPROM, with reads and writes costing 0.216 and 0.125 μJ respectively. These measurements agree with our claim that micro-power harvesting is sufficient for performing several computation, sensing, and storage operations on a CRFID.

4.1.3 Exploiting Hybrid-Harvesting for Improved Tag-Reader Interactions

CRFIDs must be able to communicate data to reader infrastructure in order to function effectively as autonomous

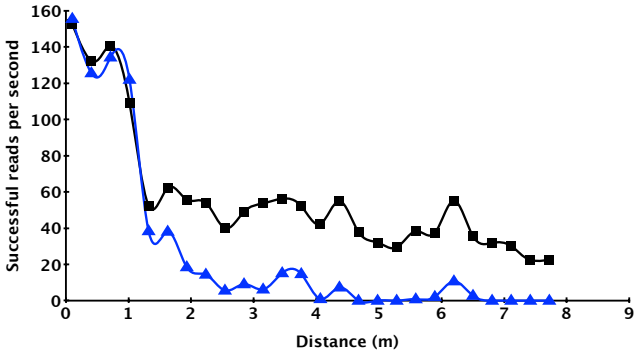

Figure 8: The number of successful tagID reads/sec indoors at a variety of distances with and without a 11.4 cm^2 solar panel. Note that the SolarWISP's energy harvesting gives it a consistent advantage at most ranges whereas the non-solar WISP encounters read rates of nearly zero beyond two meters.

sensors. We now consider computation and sensing in conjunction with communication and determine whether hybrid harvesting can improve performance while a CRFID interacts with a reader.

A CRFID generates responses to reader queries by modulating and reflecting a carrier waveform. Because the tags themselves do not generate the RF signal, wireless communication can occur at extremely low energy cost relative to an active radio circuit. Interestingly, the range at which a passive tag's RF circuit can correctly decode messages is longer than the range at which a tag harvests sufficient energy to generate a reply [5]. This imbalance suggests that a passive tag harvesting relatively continuous solar power has the potential for significantly improved communication ranges.

To investigate this performance enhancement, we measured the communication range of an RF harvesting WISP running the default firmware provided by Intel, as well as an identical SolarWISP. To find the maximum reliable communication range, we recorded the read rate for each device at a series of distances progressively farther from the reader. Figure 8 shows that read rates for a conventional WISP start at ~150 reads/second, but quickly reduce to 0 at a distance of 2 m, with constructive multipath interference allowing intermittent communication at greater ranges. This is not the case for the SolarWISP, which sustains 23 reads/second at a distance of 7.1 m while observing an illuminance of ~300 lux. According to our measurements, the SolarWISP is able to sustain more than four times the read rate of a standard WISP at any range greater than 2.5 m.

The SolarWISP has more than *triple the effective communication range of a standard WISP*, providing the following benefits for potential applications:

1. Fewer readers required to cover a given area.
2. More communication opportunities for mobile CRFIDs.

4.1.4 Hybrid CRFIDs vs. Motes

As an emerging research platform, it is important to explain how a SolarWISP-like approach to sensing differs from a similar mote-class platform. We now present several performance benchmarks that compare the SolarWISP with

	SolarWISP	EZ430
Minimal illuminance	35 lux	200 lux
Maximal read rate	4.2 reads/sec	0.1 reads/sec

Table 4: The SolarWISP in timekeeping mode can survive extremely poor lighting conditions compared to the EZ430-RF2500-SEH. The SolarWISP also achieves a higher read rate under identical lighting conditions because of the EZ430's limited current output.

the EZ430-RF500-SEH [20] development tool in a controlled lighting environment.

The EZ430 is similar to the SolarWISP in that it harvests solar energy with an indoor-optimized solar panel and uses an MSP430 microcontroller. The primary differences between the two platforms are the use of an active radio and battery on the EZ430. These design decisions have a major impact on how the devices are best used.

The EZ430 uses a thin film battery optimized for devices with a SolarWISP-like power budget. While more efficient than conventional batteries for ultra low-power operation thin film batteries still suffer from the same fundamental limitation as other batteries: a fixed number of available charge cycles. The sensitivity of batteries to charge and discharge power conditions limits the ways in which they can be used.

Stable charge voltage. Batteries store energy chemically and have charge voltage requirements that must be followed to ensure longevity. The EZ430 meets these requirements by using a large 32.5 cm^2, 2-cell solar panel in conjunction with a boost-converter to provide a stable battery input voltage. This circuit causes the EZ430 to incur overhead associated with voltage conversion and thus increases the minimal amount of energy required for perpetual operation. To find the minimum stable operating point of the EZ430, we measured battery charge voltage at varying light intensity levels in a well-controlled lighting environment, using a halogen bulb as the light source. We found that the minimum illuminance level required to activate the charging circuit and provide a stable battery voltage is approximately 200 lux (refer to Table 4), which agrees with the EZ430 datasheet. In the same lighting environment, the SolarWISP only requires 35 lux to operate perpetually while time-keeping. It is important for solar harvesting platforms to achieve energy neutrality. The device must use no more power than it receives via harvesting, lest the energy buffer exhaust. In Section 3 we presented several light traces from indoor environments. Trace *office2* illustrates the importance of energy neutral operation in low-light environments; 99% of the light readings are less than 200 lux, while only 45% of the readings are below 35 lux. The gap is similar for the other traces, indicating that a CRFID like the SolarWISP may be much better suited for mobile applications in low-light, indoor environments.

Limited discharge rate. Another drawback of batteries is the limited rate at which stored energy can be extracted and used. While thin film batteries offer vast amounts of energy storage compared to similarly sized capacitors, the rate at which energy can be used is much lower than that required for an active radio. The maximum possible rate of packet transmission on the EZ430 is thus dictated by battery discharge rate, not the amount of buffered energy. In Table 4, we show that the CC2500 is only capable of sending packets at a rate of 0.1 per second. The SolarWISP uses less energy per transmission and is not limited by battery discharge rates. Under the same lighting conditions as the EZ430 in the previous experiment (minimum light level for EZ430), we found that the SolarWISP is capable of achieving a read rate of 4.2 reads/second when harvesting 200 lux of light at a distance of 5 m from the reader. At this distance the SolarWISP must rely on harvested solar energy because it receives negligible energy from the reader (see Figure 8), but still maintains a communication rate an order of magnitude greater than that of the EZ430.

4.2 Capacitor Sizing for CRFIDs

Capacitor size is an essential design parameter for CRFIDs and capacitors have two key characteristics that must be taken into account during selection.

1. **Exponential charge curve.** Capacitors reach a useful voltage level slowly. As discussed in Section 3.2, the charge curve is described by an exponential function. This is in contrast to batteries, which reach nearly maximal voltage early in their charge cycles. Hardware solutions, such as boost converters, could potentially reduce this problem but have the potential to significantly increase quiescent current consumption. A low-input boost converter from TI, for example, consumes 5 μA of current while boosting voltages as low as 0.9 V [21]. This would more than double the timekeeping current consumption of a WISP.

2. **Low energy densities.** Capacitors, including supercapacitors, have lower energy densities than batteries. Due to size and weight limits, we only consider capacitors which will provide sufficient energy for at most *hours* of operation without incoming energy.

These characteristics create a tradeoff between responsiveness and survivability. Responsiveness is how quickly a capacitor can reach a suitable voltage to respond to reader contact events. Survivability is how long a capacitor can sustain operation without incoming energy. To demonstrate how capacitor size impacts responsiveness, we look at two benchmarks. First, we look at single capacitors' charge and discharge times. Next, we determine the performance of the SolarWISP in terms of the number of reads at different distances from a reader and for differently sized energy buffers.

Impact on tag-reader interactions. To demonstrate the effect that capacitor size has on aggregate read rate, we benchmarked tag performance in an environment similar to that which produced Figure 8. The SolarWISP was programmed with an application that ignores any reader queries when its voltage is below a threshold of 2 V and responds to the reader otherwise. Intuitively, a SolarWISP should generate a burst of reads when voltage surpassed this threshold, followed by a period of time spent recovering when it crossed below.

Figure 9 shows the average number of reads per second for several different capacitors at varying distances. The smallest capacitor size (10 μF) and the highest capacitor size (9400 μF) have the two worst read rates among all capacitor

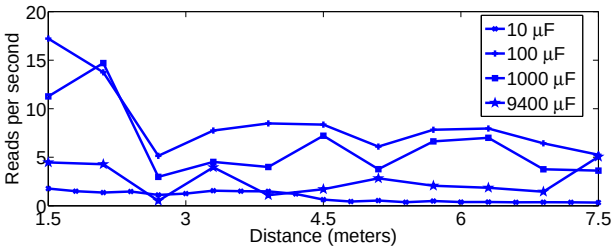

Figure 9: At distances greater than 3 meters, read rate becomes more dependent on harvestable solar energy than harvestable RF energy. A 100 μF capacitor provides the maximal read rate in this experiment.

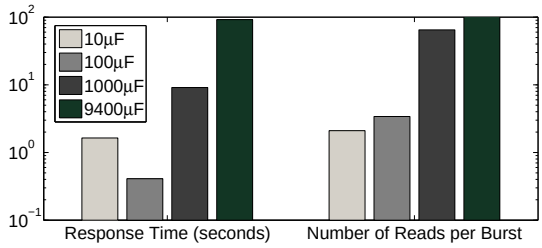

Figure 10: Progressively larger capacitors allow for increasingly large bursts of reads at the cost of responsiveness. We attribute the poor response time for the 10 μF capacitor to a pathology where the WISP attempts to wake prematurely, thus starving itself of energy.

sizes. The intermediate sizes perform better with 100 μF performing the best of the lot.

This behavior results from the combination of two factors. At one end of the spectrum, larger capacitors take more time to charge and therefore spend long periods of time with voltage below 2 V, thereby missing many read opportunities. At the other end of the spectrum, a small capacitor size triggers pathological behavior wherein a SolarWISP attempts to wake prematurely, starving itself of energy. An intermediate capacitor size (100 μF) works best since it balances the two factors.

Figure 10 provides further insight into our results. The figure shows the average burst size and recovery time for all four capacitance values at a distance of 5 meters. The burst size monotonically increases with capacitor size because of the increasing amount of energy available to the platform during one burst interval. Response time generally increases with capacitance, as it takes more energy to reach the same voltage.

Impact on survivability. While larger capacitors are less responsive to rapid changes in energy, they provide greater survivability, The larger the capacitor, the longer a SolarWISP can survive. As shown in Table 3, the SolarWISP consumes only a few μW for basic timekeeping operation, hence even the energy in a reasonably sized supercapacitor can last a long time. In our experiments, we find that the smallest capacitor (10 μF) lasts only seconds whereas the large capacitor (9400μF) lasts ~11.5 hours. Thus, for the SolarWISP, a modestly sized supercapacitor is sufficient as a power source for hours or even tens of hours while leveraging low-power states.

4.3 Coping with Harvesting Dynamics

While we have addressed the feasibility of autonomous operation, we have yet to account for the important issue of energy harvesting dynamics. Mobile devices that rely on harvested power for continuous operation are at the mercy of harvesting dynamics while trying to achieve robustness. Energy harvested by the SolarWISP varies greatly with time and movement. Indoors, this can be attributed to lighting distribution and short-term changes in panel orientation. Outdoors, variations are due mainly to temporary occlusions and diurnal variations.

To illustrate the impact of harvesting dynamics on the operation of a CRFID, we performed a simulation study using a 30-minute portion of the *office2* trace with little available

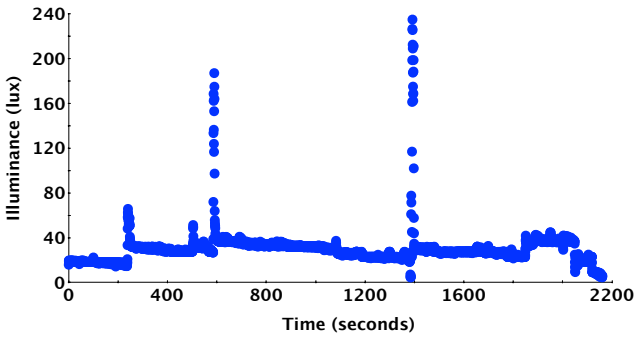

Figure 11: Time series plot of the *office2* trace excerpt used to illustrate dynamism tradeoffs. The mean illuminance is only 24 lux, but the maximum illuminance is more than 200 lux and the minimum is zero.

light and significant dynamics (see Figure 11). The trace excerpt stresses the ability of the SolarWISP to cope with dynamics and limited harvesting rates. The simulation also uses the WISP's standard 10 μF storage capacitor, which can only buffer enough energy for a several seconds of operation. As a result, harvesting dynamics are a major factor.

We considered several workloads for the SolarWISP. The lowest power state is RAM retention mode, which is only sufficient for refreshing SRAM between reader contact. A slightly higher power mode is RAM retention with periodic EEPROM writes every 5 minutes. Next, timekeeping requires waking up every ten seconds to bump a counter in order to maintain a persistent clock. We also simulated different duty-cycles for the SolarWISP, where the device periodically transitions from the RAM retention state to an active state. For a 1% duty-cycle, for example, the SolarWISP is up for 100 ms out of every 10 s. A summary of the CCTS results appears in Table 5.

The results show that the SolarWISP achieves an uptime of more than 90% in all cases, whereas a standard WISP would fail completely in fewer than 15 seconds away from an RFID reader. In this particular trace, writing data to EEPROM memory every five minutes has no impact on the WISP's time of death and saves most of the data recorded. This behavior is due to the 10 μF capacitor's high responsiveness. It typically achieves maximal voltage within the

Workload	# Failures	% Uptime
RAM retention	1	99.19
RAM retention + record	1	99.19
Timekeeping	1	98.51
1% Duty-cycle	1	98.14
2% Duty-cycle	174	95.14
5% Duty-cycle	169	91.69

Table 5: The total number of outages predicted by CCTS for the *office2* excerpt. "Record" is the process of writing one byte to EEPROM. The number of outages in these CCTS results highlight how little excess energy the SolarWISP can store in a 10 μF capacitor.

five-minute duration between EEPROM writes. The WISP is then able to perform the energy-expensive EEPROM writes without failing. The performance remains similar up to a duty-cycle of 2%, at which point performance degrades significantly. In terms of uptime, performance reduces only by about 3%, but the SolarWISP fails more far more frequently (174 failures in 30 minutes). Finally, a duty-cycle of 5% results in slightly fewer failures than 2% because of longer individual failures, but the overall uptime decreases as expected.

The simulation study has important implications for designing systems using SolarWISPs. Despite the fact that the trace represents bad lighting conditions with heavy dynamics, the SolarWISP achieves uptimes of close to 100% across a range of workloads. This result suggests that it is feasible to design systems that can exploit the SolarWISPs ability to do useful work between reader contact events. A designer may even be able to assume that the device rarely loses power or state between reader contacts. However, the results also show that it is difficult to guarantee that a SolarWISP will stay awake. For example, it is not possible for the device to survive the full 30 minutes — even if it remains in the lowest power mode at all times. Thus, it is critical to design under the assumption that, even with ambient harvesting, a small number of failures are inevitable and will need to be somehow mitigated.

While our study focused on a small capacitor size, dynamics are less of a concern with larger capacitor sizes. As mentioned in Section 4.2, larger capacitors can sustain a CRFID for more than 10 hours, but there is a price to be paid in responsiveness. For some applications a large capacitor will allow a CRFID to ignore harvesting dynamics, but for others, the rate of change in energy availability largely determines the minimal appropriate capacitor.

4.4 Application Studies

We now present case studies that validate our recommendations for balancing responsiveness and survivability. Each application stresses a different aspect of the *responsiveness* vs. *survivability* tradeoff. The first mobile application interacts with readers in frequent bursts of activity, making *responsiveness* critical. For the second application, a SolarWISP is statically deployed in an outdoor environment. The long harvesting outages caused by night make *survivability* the critical metric for success. We evaluate each application using empirical measurements of a deployment. For one application, we supplement the results with CCTS predictions to quantify the expected performance with a minimal capacitor size during a long deployment. For example, CCTS predicts that a SolarWISP can achieve 84% uptime in the worst case scenario that occurs in our traces.

4.4.1 Path Reconstruction

A mobile RFID tag or CRFID that periodically visits a set of networked readers generates data as a collection of individual read events. Unfortunately, these individual read events can be quite noisy. For example, Welbourne et al. report less than a 40% median success rate in detecting contact events using carefully placed readers and a variety of tag positions [23]. Heavy post processing was required in order to achieve reasonable inference accuracy for path reconstruction applications.

An attractive alternative to inferring tag behavior by examining reader logs is to directly report the desired application level measurement. Because CRFIDs can execute arbitrary instructions and store arbitrary state in memory, they are capable of reporting a result directly to a reader. In the case of path reconstruction, a single CRFID can report a timestamped list of all encountered readers. Providing mobile CRFIDs the ability to directly report pathing information is attractive because this removes the burden from the backend to infer paths based on large databases of tag visitations. Additionally, security may be added by reporting a cryptographically secure hash of path information instead of reporting a list of visitations in plain text.

To benchmark this application we implemented local path reconstruction on the SolarWISP. For tags to distinguish between readers and correctly store pathing information in memory, we embed a unique ID in the session field of the EPC Class 1 Generation 2 *query* command. Unfortunately, this protocol field contains 2 bits of information, thus placing a cap of 4 readers on our deployment experiments. A more flexible solution would be to make use of the write command, allowing a reader to report an arbitrary ID in memory, but the existing WISP firmware only supports the *query* and *read* commands.

The reader testbed consisted of 4 Impinj Speedway readers deployed in offices and labs around a research building. Each reader was programmed to use a unique session value for identification purposes.

Tag–reader interactions. To accurately reconstruct a path, at least one successful read must be recorded per reader contact event. The additional range and higher read rates enabled by ambient energy harvesting give the SolarWISP a higher chance of being read when passing by a reader.

Table 6 shows a comparison between the standard and SolarWISP in terms of successful reads per contact event. The experimental setup consisted of a mobile WISP moving past each deployed reader 30 times. Note that the SolarWISP can achieve up to four times the number of reads as compared to a standard WISP. Also of significance is that the SolarWISP performs *consistently* better at each reader location, in spite of different multipath environments, due to its greater energy availability. Contact duration is quantified in Table 7. Again, the SolarWISP consistently achieves longer contact durations than a conventional WISP because of its longer read range, typically by a factor of four or greater.

Responsiveness. The responsiveness of a tag to reader queries is key to the performance of path reconstruction.

Reader	Reads (WISP)	Reads (SolarWISP)	Improvement
0	25.9	31.8	1.23x
1	7.2	18.3	2.54x
2	7.7	31.5	4.09x
3	6.9	26.6	3.86x

Table 6: The average number of times a mobile tag can be read increases by as much as a factor of four when harvesting energy from ambient indoor lighting.

Reader	Contact duration (WISP)	Contact duration (Solar WISP)	Improvement
0	0.863 seconds	4.711 seconds	5.96x
1	0.539 seconds	2.978 seconds	5.53x
2	1.724 seconds	4.585 seconds	2.66x
3	1.124 seconds	5.211 seconds	4.64x

Table 7: The reader contact duration of a mobile WISP increases when ambient energy is available because of the greater maximum range that hybrid harvesting provides.

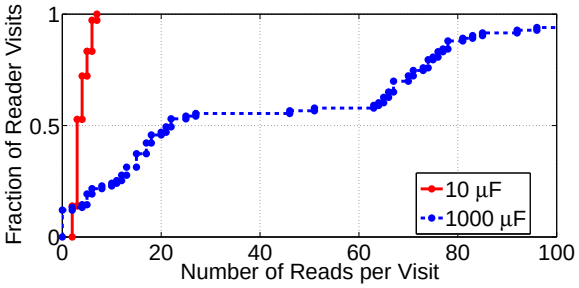

Figure 12: Large capacitors allow many reads per visit, but reduce responsiveness. A SolarWISP equipped with a 1000 μF capacitor misses 16% of reader visitations in this experiment.

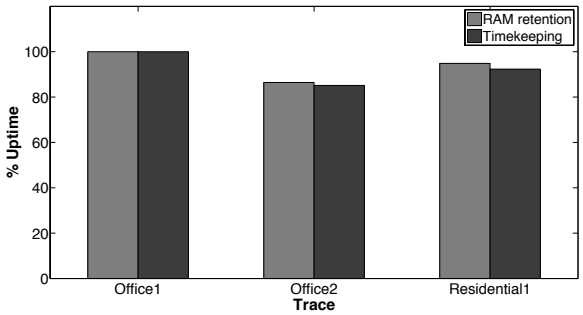

Figure 13: The percent uptime for a SolarWISP with two different workloads as predicted by CCTS. There are no outages predicted for the *office1* trace, but there are for the others. The percent difference in performance between the RAM retention and timekeeping workloads is ~2%.

Choosing an excessively large capacitor will result in large recovery times after responding to a burst of queries. If these recovery times become too large, it is possible that reader visitations could be missed because of insufficient voltage on the SolarWISP's capacitor.

Figure 14 illustrates the substantial effect that reader contact has on capacitor voltage. The time series depicts the voltage over time stored by a mobile SolarWISP with a 10 μF capacitor. The capacitor is pre-charged to a stable voltage by indoor lighting before logging begins. The test begins with the WISP in range of a reader—the large voltage swings represent individual reader transactions. Next, the SolarWISP moves out of reader range, resulting in the smooth variations observed in capacitor voltage. These variations are caused by movement under individual lighting units and slight variations in panel orientation during transit.

To test responsiveness, we moved a SolarWISP through the field of a reader ~50 times with two different sized capacitors. Figure 12 shows how the number of reads per visitation contributes to the fraction of total reader visitation events. A SolarWISP equipped with the stock 10 μF capacitor was read 2–7 times per visitation. This small capacitor size results in few reads per visit, but no visitations are missed completely. The SolarWISP augmented with a 1 mF capacitor is read 0–558 times per visit. This capacitor size, however, takes longer to charge and causes the SolarWISP to miss 16% of reader visitations.

Based on this experiment, it is clear that dense reader deployments or high rates of mobility require careful consideration of capacitor size. An excessively large capacitor will cause a CRFID to miss reader visitations. A capacitor that is too small will result in few reads per visitation, but may respond more consistently.

Survivability. Survivability is also desirable for path reconstruction, but is not essential since readers could supply timestamps while reader IDs may be potentially stored in EEPROM. Whether a CRFID keeps a local clock or relies on reader infrastructure, the power consumption of this application is minimal in between reader contact events. In the worst case, the application only needs to increment a timestamp once every 10 seconds and retain the state of SRAM, which contains past pathing data. These application

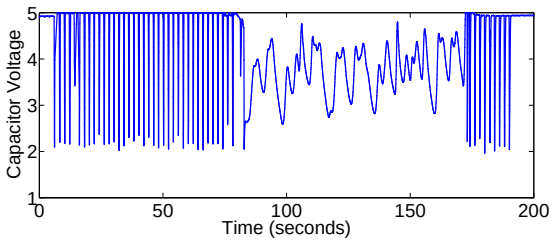

Figure 14: The capacitor is pre-charged by the attached solar panel; the large variations in voltage are caused by initial communication to a reader, while the smooth variations indicate lighting dynamics caused by mobility while disconnected.

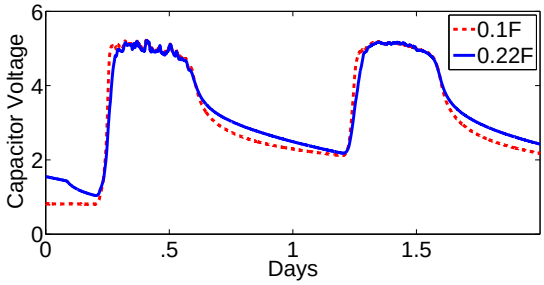

Figure 15: Supercapacitors of varying sizes store sufficient energy during daytime periods to prevent outages during night intervals. On an overcast snowy day, the SolarWISP still manages to store enough energy to survive.

characteristics suggest that a very small capacitor should be sufficient in many cases.

To gauge the performance of a large variety of capacitor sizes for disconnected operation in the path reconstruction application, we implemented the application in CCTS and tested it against the indoor illuminance traces discussed in Section 3.3. To test the application fully, we implemented both the version we measured empirically which records time locally, and the simple RAM retention version which assumes timestamps from the reader. Figure 13 shows the most interesting results from these simulation runs.

The only capacitor sizes to experience any outages were 10 μF and 100 μF, with the 10 μF performance being worse in all cases. Because these runs represent the worst performance, we present the percent uptimes for a 10 μF SolarWISP in Figure 13. For the *office1* simulation, there are no predicted outages for either workload. For the *residential1* simulation, the predicted percent uptime is between 92% and 95%. Finally, the predicted performance for the *office2* trace is significantly worse—with only 85% and 84% uptime for the timekeeping and RAM retention workloads respectively.

4.4.2 Greenhouse Monitoring

The temperature, relative humidity, and incident radiation observed in a greenhouse over time are of great interest to scientists studying biological responses to these factors. The sensing rate used for the experiment is 0.1 Hz. This sampling rate was suggested by a biologist who monitors greenhouse conditions.

From a sensing perspective, greenhouses are an environment of interest for solar-assisted CRFIDs because they can harvest light for long durations during the day and little to no artificial light at night. The sampling rates are also low enough that a mote-class device is greatly over-provisioned for the task. At the same time, greenhouse monitoring is an application for which a perpetual deployment is desirable, making batteryless computation a concrete maintenance advantage.

The SolarWISP contains an onboard temperature sensor, allowing the platform to achieve some of the functionality required by biologists. We implemented an application that records a 1-byte ADC sample once every 8 seconds and stores the result to a 1kB EEPROM. This limited amount of nonvolatile storage only provides sufficient space for 2.8 hours of temperature readings at this sampling rate, but this does not change the accuracy of the application power profile. A next generation platform could easily select a larger EEPROM while maintaining a similar form factor.

To survive the long power outages caused by night, we augment the SolarWISP with a supercapacitor. We also leverage the FrankenWISP to record the supercapacitor voltage once every 30 seconds to test survivability. Figure 15 shows voltage traces for a two-day application deployment using two different supercapacitor sizes. While both the 100 mF and 220 mF capacitor survive this deployment, they do decline to ∼2 V at night.

5. DISCUSSION AND FUTURE WORK

Micro-harvesters and small energy buffers. The key message of our work is that, despite the use of micro-harvesters and tiny energy buffers, ambient energy harvesting has the potential to greatly improve the performance and usability of CRFIDs. Unexpectedly, the benefits are not limited to well-lit environments. The SolarWISP was able to continue sensing or at least retain its RAM contents in all but the darkest conditions that we tested. While this demonstrates compelling potential, fully utilizing hybrid harvesting CRFIDs requires new systems that can adapt to short-term energy dynamics while achieving application objectives. Existing solutions do not fill this gap — most harvesting-based systems achieve deployment lifetime objectives by adapting to long term variations in energy availability (e.g. hours, days) using relatively large buffers that shield applications from unpredictable, short-term dynamics. Our work makes a compelling case for a new energy management paradigm, in which embedded devices are designed to use micro-harvesters and small energy buffers, and to handle short-term energy dynamics.

Alternative hardware designs. Our evaluation was performed using specific CRFID hardware (WISP), so one might ask if an alternative hardware design might change some of our conclusions. For example, we showed that there is a tradeoff in terms of responsiveness and survivability due to the size of the energy buffer. One might ask whether we can achieve both objectives simultaneously by pairing a large and small buffer on a CRFID. While this is a possibility, our comparison between CRFIDs and motes shows that it is important to keep the hardware minimal in order to fully exploit tiny amounts of harvested energy. Additional

hardware complexity typically results in inefficiency and energy wastage — in addition to introducing new constraints on the operating voltage and current.

Applications. While we have examined two potential CRFID applications, a broader question remains as to what class of applications is suitable for hybrid harvesting CRFIDs. Hybrid harvesting CRFIDs retain many of the benefits of RFIDs such as form factor and energy efficiency, but offer considerably more flexibility in application possibilities. However, CRFIDs may not be suitable for applications where individual points of failure are unacceptable, such as the detection of rare events. While these tasks cannot be supported reliably, CRFIDs can maintain a reasonable level of performance according to our measurements and are suitable for applications in which occasional failures are tolerable.

Simulation. We believe that CCTS is sufficiently accurate to aid in CRFID provisioning decisions, but there is room for improvement. The current characterization of solar panels, in particular, is both coarse-grained and labor-intensive. Rather than measuring a number of points along maximal and minimal VI curves, thus approximating the *output* of the panel, it may be possible to accurately model the panel itself in simulation. Other improvements to CCTS are possible and we intend to continue development with the goal of providing CCTS to the research community.

6. RELATED WORK

RFIDs and RFID sensors. As the first example of a CRFID, Intel's WISP has defined the class of devices [18]. Emerging variants on the WISP platform include Duke's Blue Devil WISP [15] and the UW SoCWISP [12]. Yeager et al. propose the use of supercapacitors to extend the fleeting lifetime of the WISP given a full charge [25]. Buettner et al. demonstrate activity inference enabled by the WISP's sensors [4]. This body of work is mostly concerned with designing RFID sensors, and does not consider how they can be augmented with ambient energy for autonomy.

Energy harvesting. There has been significant work on energy harvesting in sensor networks. Recent work [2, 9, 10, 17, 19, 22, 7] has explored scenarios in which nodes can harvest energy from their environment (e.g., from the sun) and use it to recharge their batteries. In the absence of such energy, nodes can then subsist on their replenished battery supply. There are also a growing number of solar-powered sensor network deployments, such as James reserve in Irvine, CA [8], Berkeley Angelo reserve, CA [19], CSIRO's Fleck sensor network in Australia [6], and the LUSTER system in Virginia [17]).

These systems are predominantly of the *macro-harvesting* type, wherein solar panel and energy buffer sizes are chosen to smooth out the short time scale variations in incoming energy. This greatly simplifies the design of a harvesting-aware sensor network, and enables a priori provisioning of resources. In many of these deployments, prior measurement studies of incident solar energy at the deployment location are used to select appropriate solar panel sizes for the sensor devices and to set system parameters such as the duty-cycling rate [24]. In contrast, our work tackles *micro-harvesting*, where the device, panels, and buffer are small. As a consequence, small-scale variations in energy conditions across seconds, minutes, or hours have a significant impact on the design of our system.

More broadly, the viability of various energy harvesting sources for computation and sensing has been considered many times in the past [13], but we focus exclusively on RFID harvesters that produce power at the μW scale. Sample et al. have demonstrated a WISP retrofitted with a directional TV antenna capable of harvesting energy from a TV transmitter over two miles away when positioned carefully [16]. The energy, however, was used to power a static load (small thermometer with LCD) rather than a WISP or other computation device.

Energy storage. At the core of our work is an understanding of the tradeoffs presented by capacitors for hybrid harvesting CRFIDs. The use of supercapacitors in sensor platforms is relatively common. For example, Prometheus is a harvesting-based sensor platform that integrates a conventional LiOn rechargeable battery and supercapacitors [9]. Capacitors are used in RFID systems as well. However, to the best of our knowledge, the tradeoff between survivability and responsiveness due to the use of capacitors has not been studied in any of the prior work.

Energy scheduling. Energy management for harvesting-based sensor networks has been studied in the past. Moser et al. [11] present optimal scheduling algorithms for harvesting networks that must meet deadlines; Vigorito et al. [22] present algorithms for adaptive duty-cycling based on harvested energy. In addition to adaptive duty-cycling, Kansal et al. [10] present a methodology for sizing energy buffers. While our work does not directly relate to such scheduling schemes, we believe that our work can inform such approaches. Many existing scheduling techniques are designed with implicit assumptions about the energy buffer and harvesting rate in mind. For example, [11] uses predictions of harvesting rates for its scheduling. Our work shows that for micro-harvesting, one needs to consider extremely small windows of time, and cannot assume that harvesting is smoothed by the energy buffer.

7. CONCLUSION

A small amount of ambient energy can grant autonomy to CRFID sensors. The SolarWISP augments the traditional WISP with a 11.2 cm^2 solar panel to supplement the energy harvested by the RF charge pump. This small change increases effective communication range threefold and quadruples read rate. The SolarWISP requires only 35 lux to remain in perpetual timekeeping mode, whereas the EZ430 solar mote requires 200 lux. Our trace-driven simulation results show that the SolarWISP, even with only a 10 μF capacitor, can achieve nearly 95% uptime in RAM retention mode during a period of low light, while a standard WISP would fail completely in less than 15 seconds in the absence of an RFID reader.

Our tools include a trace-driven simulator and an energy monitoring subsystem. The CRFID Crash Test Simulator estimates the survivability of a parameterized platform under a variety of lighting conditions. The FrankenWISP allows real-time monitoring of extremely resource-limited devices without greatly disturbing mobility or deployability. We hope that our tools and traces will help developers more easily evaluate the design space for future hybrid micro-energy harvesting CRFIDs.

8. ACKNOWLEDGMENTS

This work was funded by a Sloan Research Fellowship and NSF grants CNS-0923313, CNS-0845874, CNS-0615075, CNS-0627529, CNS-0546177, CNS-0855128, and CNS-0916577. We thank Intel Research Seattle for providing several WISPs; Priyanka Rajendra Iyer for help in performing experiments; Benjamin Ransford for help with simulation and measurement; and Negin Salajegheh and Quinn Stewart for feedback on drafts.

The free CRFID Crash Test Simulator software and energy traces are available from `http://www.cs.umass.edu/~ssclark/crfid/`.

9. REFERENCES

[1] TelosB Datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/TelosB_Datasheet.pdf`.

[2] D. Brunelli, L. Benini, C. Moser, and L. Thiele. An Efficient Solar Energy Harvester for Wireless Sensor Nodes. In *Proceedings of the Conference on Design, Automation and Test in Europe (DATE)*, 2008.

[3] M. Buettner, B. Greenstein, A. Sample, J. R. Smith, and D. Wetherall. Revisiting Smart Dust with RFID Sensor Networks. In *Proc. 7th ACM HotNets Workshop*, October 2008.

[4] M. Buettner, R. Prasad, M. Philipose, and D. Wetherall. Recognizing Daily Activities with RFID-based Sensors. In *Proc. 11th International Conference on Ubiquitous Computing (UbiComp)*, Orlando, Florida, USA, October 2009.

[5] S. S. Clark, J. Gummeson, K. Fu, and D. Ganesan. Towards Autonomously-Powered CRFIDs. In *ACM Workshop on Power Aware Computing and Systems*, October 2009.

[6] P. Corke, P. Valencia, P. Sikka, T. Wark, and L. Overs. Long-duration Solar-powered Wireless Sensor Networks. In *Proceedings of the 4th Workshop on Embedded Networked Sensors (EmNets)*, pages 33–37, 2007.

[7] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman. Challenge: Ultra-low-power Energy-harvesting Active Networked Tags (EnHANTs). In *Proceedings of the 15th annual international conference on Mobile computing and networking (Mobicom)*, 2009.

[8] J. Hicks, J. Paek, S. Coe, R. Govindan, and D. Estrin. An Easily Deployable Wireless Imaging System. In *Workshop on Applications, Systems, and Algorithms for Image Sensing (ImageSense)*, 2008.

[9] X. Jiang, J. Polastre, and D. Culler. Perpetual Environmentally Powered Sensor Networks. In *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks: Special track on Platform Tools and Design Methods for Network Embedded Sensors (IPSN/SPOTS)*, April 2005.

[10] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. Power Management in Energy Harvesting Sensor Networks. *ACM Transactions Embedded Computing Systems*, 6(4):32, 2007.

[11] C. Moser, D. Brunelli, L. Thiele, and L. Benini. Real-time Scheduling for Energy Harvesting Sensor Nodes. *Real-Time Systems*, 37(3):233–260, 2007.

[12] B. Otis and D. Yeager. SoCWISP: Ultra-low Power Wireless Sensing RFID Chip. WISP Summit Workshop, 2009. Presentation.

[13] J. Paradiso and T. Starner. Energy Scavenging for Mobile and Wireless Electronics. *IEEE Pervasive Computing*, 4(1):18–27, 2005.

[14] B. Ransford, S. Clark, M. Salajegheh, and K. Fu. Getting Things Done on Computational RFIDs with Energy-Aware Checkpointing and Voltage-Aware Scheduling. In *Proceedings of USENIX HotPower Workshop*, December 2008.

[15] M. Reynolds and S. Thomas. The Blue Devil WISP: Expanding the Frontiers of the Passive RFID Physical Layer. WISP Summit Workshop, 2009. Presentation.

[16] A. Sample and J. R. Smith. Experimental Results with two Wireless Power Transfer Systems. In *IEEE Radio and Wireless Symposium*, 2009.

[17] L. Selavo, A. Wood, Q. Cao, T. Sookoor, H. Liu, A. Srinivasan, Y. Wu, W. Kang, J. Stankovic, D. Young, and J. Porter. LUSTER: Wireless Sensor Network for Environmental Research. In *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems (SenSys)*, 2007.

[18] J. R. Smith, A. P. Sample, P. S. Powledge, S. Roy, and A. Mamishev. A Wirelessly-Powered Platform for Sensing and Computation. In *Proc. 8th International Conference on Ubiquitous Computing (UbiComp)*, 2006.

[19] J. Taneja, J. Jeong, and D. Culler. Design, Modeling, and Capacity Planning for Micro-solar Power Sensor Networks. In *Proceedings of the 7th International Conference on Information Processing in Sensor Networks (IPSN)*, 2008.

[20] Texas Instruments EZ430-RF2500-SEH Datasheet. `http://www.ti.com/litv/pdf/slau273b`.

[21] Texas Instruments TPS61097-33 Boost Converter Datasheet. `http://focus.ti.com/docs/prod/folders/print/tps61097-33.html`.

[22] C. Vigorito, D. Ganesan, and A. Barto. Adaptive Control for Duty-cycling in Energy Harvesting-based Wireless Sensor Networks. In *Fourth IEEE International Conference on Sensor and Ad hoc Communications and Networks (SECON)*, June 2007.

[23] E. Welbourne, K. Koscher, E. Soroush, M. Balazinska, and G. Borriello. Longitudinal Study of a Building-scale RFID Ecosystem. In *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 69–82, New York, NY, USA, 2009. ACM.

[24] Y. Yang, L. Wang, D. K. Noh, H. K. Le, and T. F. Abdelzaher. SolarStore: Enhancing Data Reliability in Solar-powered Storage-centric Sensor Networks. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 333–346, 2009.

[25] D. J. Yeager, P. S. Powledge, R. Prasad, D. Wetherall, and J. R. Smith. Wirelessly-Charged UHF Tags for Sensor Data Collection. In *IEEE RFID*, 2008.