

An Energy-Harvesting Sensor Architecture and Toolkit for Building Monitoring and Event Detection

Bradford Campbell and Prabal Dutta
Electrical Engineering and Computer Science Department
University of Michigan
Ann Arbor, MI 48109
{bradjc,prabal}@umich.edu

Abstract

Understanding building usage patterns and resource consumption, particularly for existing buildings, requires a sensing infrastructure for the building. Often, deploying these sensors and obtaining real-time information is hindered by installation and maintenance difficulties resulting from scaling down and powering these devices. Devices that rely on batteries are limited by the scale of the batteries and the maintenance cost of replacing them while AC mains powered sensors incur high upfront installation costs. To mitigate these burdens, we present a new architecture for designing building-monitoring focused energy-harvesting sensors. The key to this architecture is masking the inevitable intermittency provided by energy-harvesting with a trigger abstraction that activates the device only when there is useful work to be done. In this paper, we describe our architecture and demonstrate how it supports existing energy-harvesting sensor designs. Further, we realize three additional design points within the architecture and demonstrate how the sensors are effective at building monitoring and event detection. The sensors, however, are classically disruptive: they improve ease of installation and maintenance, but to do so, they sacrifice some fidelity and reliability. Whether this trade-off is acceptable remains to be explored, but the technology needed to do so is now here.

Categories and Subject Descriptors

B.4.m [HARDWARE]: Input/Output and Data Communications—*Miscellaneous*

General Terms

Architecture, Design, Experimentation, Performance

Keywords

Building monitoring, Energy-harvesting, Sensing

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.

Copyright is held by the authors.

BuildSys'14, November 5–6, 2014, Memphis, TN, USA.
ACM 978-1-4503-3144-9/14/11
<http://dx.doi.org/10.1145/2674061.2674083>

1 Introduction

Improving occupant comfort while reducing energy usage is a key goal of recent buildings research. The motivation is strong: people spend approximately 90% of their time indoors [30] and buildings consume 41% of the energy produced in the United States [8]. Recent work proposes building control algorithms [9, 26], dataset analyses [13, 15, 19], feedback mechanisms [28, 29], and personalized HVAC systems [14, 23] in order to increase comfort or decrease resource utilization. A common thread of these strategies is that they require sensing devices to feed information about the state of the building and its occupants to models or algorithms. Often, these sensors are tailored to the particular application and not designed for scalability or installability. The U.S. Department of Energy has responded by announcing a new program in 2014 [1] that seeks “innovative [sensor] approaches that reduce the cost and power consumption for data collection of common building operation variables” and “open-source sensor packages that allow for data acquisition and transmission with increased lifespan between manual calibrations.” Without these sensors, building management strategies become challenging to implement in practice.

The absence of such sensors make both quick and efficient benchmarking of buildings difficult, and their long-term monitoring very costly. To address this issue and to aid building-scale control systems in the future, we propose an architecture for developing sensors that directly addresses the issues of scalability and installability in indoor spaces while supporting sensing applications that are common to building monitoring tasks. Two main factors limit how well sensors scale in buildings: size and power. This is in stark contrast to outdoor sensor networks, where large sensors with batteries are acceptable and the dominating limitation is network connectivity. Indoors, connectivity with Ethernet, Wi-Fi, Bluetooth, and potentially 802.15.4 abounds, but large sensors are unacceptable due to their unappealing aesthetics, visual intrusiveness, and installation challenges. As sensors shrink to make them closer to invisible, so too does their energy storage capability [32], making long-term battery powered operation impossible without large maintenance costs. Powering sensors with AC mains is also infeasible as it dramatically increases installation costs and may be impossible to retrofit into existing buildings.

Our proposed architecture enables sensors with power budgets that allow them to scale in size to easily deployable dimensions by leveraging energy-harvesting as the primary power supply. By accruing energy at run-time and only storing enough for immediate or short-term operation, the volume of the sensor node reserved for the power supply can be dramatically reduced in comparison to battery or mains powered devices. While this addresses the scaling issue, energy-harvesting, and the intermittency that often accompanies harvesting from unreliable sources, makes designing and building sensors capable of periodic sampling, event detection, networking, and other common sensor features that can accomplish typical building monitoring tasks, such as occupancy detection, temperature sensing, and air quality monitoring, more challenging than when using reliable power supplies.

Our architecture addresses this by noting common patterns in building monitoring applications and decomposing sensors into four main subsystems: an energy-harvesting power supply, an activation trigger, a sensing device, and a data communication module. The energy-harvesting power supply accrues and stores energy to power the rest of the sensor node, often only intermittently. The activation trigger is responsible for cold-booting the node, initiating computation, and sampling a sensor. The optional external sensing device captures information about some phenomena. Lastly, the data communication module is responsible for immediately transmitting the sensor's data to an interested receiver or logging the data locally and transferring later. Each of these subsystems can be adapted for different applications.

The hypothesis underlying this architecture is that a system composed of these layers allows for and simplifies the design of energy-harvesting based sensor nodes that can perform many sensing applications. The activation trigger is key to enabling this. By providing a trigger based abstraction for the device's computation, sensing, and communication components, the intermittency of the underlying energy-harvesting power supply is masked by requiring that a node need only be powered and active in response to a trigger. Triggering at the circuit level provides a familiar event-based programming abstraction at the software level and allows the device to stay in a dormant state, with or without available energy, when it is not triggered.

To illustrate this architecture, we overview the design space of sensor designs that are supported within this architecture. We explore viable choices for the four subsystems and explore which combinations are viable for a range of indoor monitoring and event detection applications. To demonstrate the effectiveness of this architecture and its encompassed sensor designs, we survey three recent energy-harvesting designs that are suitable for building monitoring applications and show how their designs and characteristics fit well with our proposed architecture. Additionally, we realize three new indoor energy-harvesting sensors designed to detect vibration, airflow, and light/occupancy. We evaluate the capabilities, performance characteristics, and suitability of each sensor for three building monitoring applications. Finally, we discuss limitations of the energy-harvesting approach, possible solutions, and directions for future research.

2 Related Work and Background

As tools and methods for improving building energy use often begin with or rely on understanding the state of buildings, much previous work has focused on building sensing technology. In this section we describe a selection of sensing systems that are designed for buildings.

2.1 Ambient Sensing

Ambient sensors monitor general conditions, like temperature [5], humidity, light, and air quality [2, 17] inside of a building. These sensors are often battery powered or built into the building and powered by AC mains. Typical sampling rates are on the order of once per minute and data are typically fed to a server for processing and presentation.

2.2 Energy Consumption

When trying to improve the energy efficiency of buildings a common first step is to submeter loads to determine where energy is being consumed. Many plug-load level metering [3, 16, 18, 21] systems exist for measuring the power draw of individual loads and appliances. By their very nature they have access to AC mains power and avoid the sensor node powering problem. However, size is still important and efforts have been made to scale down the AC-DC power supply to enable power-strip friendly power meters [12].

To meter loads that do not plug into an outlet or to reduce the number of required meters, AC power can also be measured at the circuit level. Panel level meters [20, 24] provide reasonable sensing insight while only having to be installed in panel boxes and not at every load. However, meters that use a direct connection to a spare breaker to power themselves typically require an electrician to install, raising installation costs.

2.3 Occupancy Detection

Tailoring HVAC systems for efficiency and occupant comfort requires knowing where people are in a building to optimally direct resources. Also, many responsive applications at a building scale require localizing occupants in the building. Systems for detecting occupancy are often based on thermopile arrays [10], which can detect people using heat signatures, or PIR motion sensors [22], which rely on detecting movement to infer occupancy. These sensors are often bulky and power inefficient or require direct installation into ceilings or walls, failing the size, maintenance, and installability requirements of building monitoring sensors.

Many of the building monitoring sensors in prior work are tailored to a specific application and designed to sense a specific phenomena. Our work is designed to not only implement three sensors that are useful for building monitoring, but also to provide a framework for developing other energy-harvesting sensors by defining hardware abstractions and providing reusable components. We also address the issues of size and maintenance by adapting an energy-harvesting power supply to eliminate the need for large batteries.

We describe related work in building monitoring sensors that incorporate energy-harvesting power supplies in [Section 4](#) to better show how the projects fit well into our proposed architecture.




Figure 1: Energy-harvesting node architecture.

3 Sensor Node Architecture

The overall node architecture for the series of energy-harvesting building monitoring sensors is depicted in Figure 1. Each sensing device is comprised of four main subsystems: the energy-harvesting power supply, a node triggering mechanism, the primary sensing apparatus, and a data communication module. Each sensor accrues energy at run time with a suitable energy-harvesting frontend. The architecture specifies minimal constraints on the power supply for successful system operation, however, the harvesting supply determines the sensing capabilities and performance of the overall device. Power supplies that range from ones that can supply the device continuously to anemic ones that offer very limited, intermittent runtimes are supported. Application goals define the power supply requirements.

The trigger mechanism is responsible for activating the computational, sensing, or communication resources on the sensor, depending on application needs. The trigger may fire as soon as sufficient energy is available or in response to an external event, such as a door opening.

An on-board sensor may be used to monitor some building condition. The sample rate is bound by the power supply harvesting rate and the range of usable sensors is governed by the energy store. Low power sensors that can obtain a reading quickly are best suited for this architecture.

Finally, the data communication module conveys the information collected by the device to a central building control system or other data management service. This can be handled in real-time, batched, or logged locally and transferred in bulk at a later time.

3.1 Design Space

Each of these subsystems has a spectrum of design options when instantiating a physical sensor. Different environmental factors, monitoring or detecting goals, and data rate requirements will guide which options are best for a particular application. Table 1 highlights the main options for each subsystem and a series of potential building-based sensor applications. A mark in a column for a particular sensor application indicates that that option is viable for that application. The remainder of this section provides a more detailed breakdown of the range of options.

3.1.1 Power Supply

The power supply is responsible for harvesting energy for the node's operation. Depending on the harvester's capabilities and node power requirements, the power supply may be continuous, i.e. its power output is equal to or greater than the average power draw of the node, or it may be intermittent, i.e. it is only capable of powering the node for short periods of time between charging cycles. The architecture supports both paradigms, however, particular applications may require one or the other.

A common harvesting source for indoor applications is small ($\sim 5 \text{ cm}^2$) photovoltaic solar cells because when a building is occupied it is typically lit. These photovoltaics can often be sized to fit node dimension constraints or average power requirements. Larger solar cells in sufficiently lit environments are capable of supporting continuous operation, while smaller solar cells can only provide intermittent power [32]. Other indoor harvesting sources have similar tradeoffs, including airflow based [31] and water pipe based [25].

Optionally, the power supply system may include backup energy storage to buffer variations in harvesting capability, much as the power grid does with renewable sources. A rechargeable battery or supercapacitor could be charged with excess harvested energy when the harvesting source is plentiful and discharged when it is needed, such as at night in the case of the solar cell. The stored energy may also be useful to continuously power a real-time clock or other timebase or to provide a burst of energy to power a wireless radio for bulk data transfer.

One important property of the power supply that enables the proposed architecture is a trigger, latch, and shutdown mechanism. Typically, energy-harvesting power supplies supply power on their output rails immediately when power is available. In the case of continuous supplies, power is always available. For some applications this may be useful, but others may only want to consume the limited energy reserves after a period of time has elapsed or an event has occurred. To support this, the power supply must have the capability to store energy without activating the remainder of the sensor until a trigger pin is asserted. Upon assertion, the power supply latches and keeps the power rail activated until energy is exhausted or the device requests that it be disconnected by asserting the shutdown pin. The power supply should also be capable of disabling this feature and supplying the power rail immediately after energy is available. This key property enables successful node operation and event detection with energy budgets on the order of millijoules.

The power supply is a very important subsystem as its characteristics and capabilities directly affect the operation of the entire sensor. The sample rate, supported sensors, and measurement capability of the device are directly based on the output and duty cycle of the node's power supply. In this way, the power supply's characteristics define many of the constraints on the rest of the sensor, such as the runtime duration, maximum instantaneous power draw, per-event energy consumption, and other operational limitations. Therefore, the power supply and application must be well paired to ensure a successful device.

Sensor Application	Power		Trigger			Sensing			Communication		
	Sustained	Intermittent	Periodic	Opportunistic	Event	External	Event	Energy	Immediate	Batch	Log
Door	✓	✓			✓		✓		✓		
Occupancy	✓	✓	✓	✓		✓		✓	✓	✓	✓
Temperature	✓	✓	✓	✓		✓		✓	✓	✓	✓
Light	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Motion	✓	✓			✓	✓	✓		✓		
Air Quality	✓	✓	✓	✓		✓			✓	✓	✓
Water Usage	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓
Solar Radiation	✓	✓	✓	✓		✓			✓	✓	✓
Chair Occupancy	✓	✓			✓		✓		✓		

Table 1: Energy-harvesting node architecture properties for a variety of building monitoring applications. This table highlights a number of potential building monitoring applications and the system properties that are likely sufficient for implementing an energy-harvesting sensor for the particular application. Each subsystem—power, trigger, sensing, and communication—needs one property to enable a functioning sensor. Some applications are best suited by event detection, some just need to detect harvestable energy, others are best served with periodic sampling, but all support intermittent power supplies making them candidates for energy-harvesting sensors.

3.1.2 Trigger

The trigger subsystem determines when the sensor should take a sample, perform some computation, or make note of an event. Using the trigger abstraction and an event based model to control the operation of the device allows the architecture to natively support intermittent power supplies as the intermittency can behave as a trigger. Different applications will use different types of triggers to best accomplish their sensing goals. Trigger types fall into three main categories: periodic, opportunistic, and event based.

The **periodic trigger** is a trigger designed to fire at a regular interval. If the harvesting source is predictable and the sensor has some timebase, it may be feasible to have the trigger activate periodically in order to take regular samples. This periodicity matches how many battery-powered sensors operate when monitoring an environment and provides a fixed sample rate. If the harvesting potential changes there may not be sufficient stored energy to sample and the sample will be lost, just as a sample can be lost due to faulty packet transmission in a battery powered network.

The **opportunistic trigger** is based on an intermittent power supply. This trigger is activated by the power supply detecting that the stored energy has surpassed a threshold and supplying power to the rest of the sensor. This allows the device to sample at the maximum rate allowed by the intermittent power supply and not be limited to a precomputed rate that optimizes for average power. If harvestable energy is abundant, the sensor will potentially be able to sample faster than the rate budgeted for a battery-powered node.

The **event based trigger** activates the sensor in response to an external event. This allows the device to store up energy from an intermittent power supply until the event occurs. Detecting the event will likely require dedicated hardware that is capable of asserting the trigger pin on the power supply in response to the desired event occurring. Also, the rate at which the node can detect events is bounded by the rate the power supply can charge as the trigger is only active when there is sufficient energy to otherwise activate.

3.1.3 Sensing

Certain building monitoring applications may need one or more dedicated sensors to fulfill the application’s goals. Others may just rely on the trigger mechanism for sensing: the mere occurrence of the trigger provides the desired event detection. Yet other applications may simply rely on the ability to harvest energy or not.

Applications that employ dedicated sensors must ensure the sensors fit into the energy budget provided by the power supply. Most sensors will likely not be active continuously, but instead for very short intervals when power is available and sensing is warranted. Coupling sensing with event detection may be useful, and the trigger mechanism can provide a useful rate limiting for sensing: when the trigger fires, sample the sensor. This motivates using sensors that require little time between power-on and valid measurements. Once the sensor has been sampled it can be power-gated again allowing for a quicker recharge.

Using the trigger as a sensor is useful for applications that must detect events, such as doors or windows opening and closing, lights turning on, airvents activating, and chairs being sat on. These applications can rely solely on a trigger mechanism activating the output rail of the power supply to sense a particular phenomena. It is possible this sensing style may be combined with a dedicated sensor to both detect an event and measure the conditions at the time of the event. For instance, a temperature reading may be taken immediately after a window open event to gain insight into the motivation for opening the window.

In some cases the ability of the power supply to harvest at all may be sufficient for the application’s sensing needs. Similar to the Monjolo [11] project, the mere presence of energy from which to harvest can often be a sufficient sensor. For example, a solar cell powered device placed close to a light can tell when the light is on simply based on when it is able to harvest. Further, the relative rate of harvesting may provide additional information such as the brightness of the light. In this way, an energy-harvesting sensor can be constructed with no true sensor at all.

3.1.4 Data Communication

The final operating phase is to communicate, store, or otherwise handle the result from the sensing subsystem. To best enable real-time data processing and visualization, one option is to transmit the data immediately after sampling. This would likely require a wireless receiver to be within range and always listening. We propose this is feasible with a combination of gateways designed to bridge low power networks and the Internet, and other, mains-powered sensors, such as power meters, distributed throughout the building that can forward the data packets.

If delaying data transfer from the sensing device to the building control or data management service is tolerable for a certain application, batching multiple samples and transmitting the group may be a superior solution. This could potentially allow the harvesting device to save power by transmitting a single packet with multiple readings and incurring less payload overhead, and therefore transmit time and power, rather than sending multiple data packets each with their own headers and other overhead. Buffering packets may also allow the device the flexibility to probe for online receivers at each activation and, if possible, join a wireless network temporarily or pair with a device capable of relaying its data.

If a network with which to transmit the data is unavailable, the device has the option of logging the samples locally into some nonvolatile storage, such as flash or FRAM. This has the advantage of being a stand-alone sensor, but the drawbacks of requiring a data retrieval system and being limited by the storage space available for data. Likely this would require physically retrieving the sensor, downloading the stored data, and reinstalling the device. Other schemes may be feasible, such as a targeted wireless retrieval mechanism where a data mule [27] individually activates each device and requests its data. The device could store the energy it would otherwise use for periodic transmissions to enable the bulk wireless transfer.

4 Previous Harvesting Designs

Some previous work has also identified the need for eliminating batteries from indoor sensing devices by moving to energy-harvesting power supplies. While these projects presented an energy-harvesting approach to a particular application, they can be well described as a part of our node architecture. We highlight three such projects here—Monjolo, DoubleDip, and Trinity—and describe how their designs help support our architecture as a viable framework for building energy-harvesting sensors.

The first approach is the Monjolo architecture [11]. Monjolo leverages an energy-harvesting power supply to sense energy-consuming loads. That is, the rate at which the harvester can recharge is proportional to the energy consumption of the attached load. In this way, any load that produces a side-channel of harvestable energy can be metered by attaching a suitable energy-harvesting power supply. The Monjolo device is designed around an intermittent power supply, uses an opportunistic trigger and an implicit sensor (sufficient energy has accrued), and transmits immediately. By embracing intermittency, this work provides a great example of a particular design point within our architecture.




Figure 2: Solar cell based energy-harvesting power supply.

DoubleDip [25] is a water flow monitoring system attached to pipes that uses an energy-harvesting thermoelectric generator (TEG) as both a power supply and water event detector. A sudden water flow event causes a temperature gradient to be present on the pipe and a resulting spike in the output voltage of the TEG. DoubleDip detects this spike and records a water event in the pipe. DoubleDip blurs the lines between a sustained power supply and an intermittent one by including a backup battery. A DoubleDip node that has limited opportunity to harvest can still respond to water events by powering itself from the battery, however, the system design caters to intermittent operation. It uses an event based trigger (a water flow event), an accelerometer as an external sensor once it has detected water movement, and transmits data immediately or batches and waits for sufficient power to be available. Again, intermittency is masked by waiting for an event-based wakeup before performing any sensing or computation.

The third design is another indoor energy-harvesting sensor named Trinity [31], a system for detecting and measuring airflow. Trinity uses one piezoelectric device to harvest and another, simpler piezoelectric device as a sensor to measure airflow. Trinity’s runtime power draw and harvesting capability are closely aligned to provide operation similar to a battery powered node, but at a lower duty cycle. While this aids inter-node communication, it forces Trinity to be relatively large to provide sufficient power. Because of Trinity’s harvesting capabilities we consider its power supply to be sustained. It uses a periodic sensor to initiate sensing, a dedicated vibration sensor for measuring airflow, and is able to transmit after sampling. Trinity is the most similar to a battery powered node and demonstrates that our architecture supports applications where the node design is less limited by intermittency.

5 Implementation

To further verify and evaluate the energy-harvesting building monitoring architecture, we implement three new sensors that represent three points in the design space: a vibration detector, an airflow meter, and a light/occupancy sensor. Each sensor is designed to monitor a particular phenomena common to buildings that would likely be of interest to measure. Because the power supply for the three sensors is shared we describe it independently.

5.1 Solar Harvesting Power Supply

Common to all three sensors is an energy-harvesting power supply based on an amorphous silicon solar cell. The circuit design is from prior work presented by Yerva, et al. [32]. Fundamentally, it charges a 500 μF bank of capacitors using the solar cell and then waits for a trigger signal. Upon triggering, the supply enables an output regulator that



Figure 3: Vibration sensor.

maintains a 3 V power rail. A latch circuit keeps the regulator enabled until the capacitors discharge or shutdown is requested. The power supply supports the power supply interface described in Section 3.1.1 by allowing for switching between opportunistic and external triggering when the sensor is assembled. It also has a shutdown pin that allows the sensor to disable power to itself. Further, the power supply interface provides access to the voltage on the storage capacitors which may be useful when the output regulator is not enabled. The power supply board is 3.81 cm by 1.52 cm with one side nearly completely covered by the solar cell and the other by the power supply electronics, as shown in Figure 2.

Using a photovoltaic based energy-harvester is compelling due to the relative abundance of indoor light. Spaces occupied by people are typically lit, and light propagates to most areas of a room. This allows us to put the solar cell on one side of the sensor and face it towards the light when it is mounted. The primary drawback to solar based harvesting is it is ineffective in dark rooms. We discuss possible solutions for this problem in Section 7.1.1.

While we present a power supply that only harvests from one source (indoor photovoltaics), other sources may be useful for building indoor sensors. Our intent is to demonstrate the effectiveness of our proposed architecture and one harvester is sufficient. Demonstrating the design and evaluation of alternative harvesters is left as future work.

5.2 Buzz: Vibration Detector

Our first sensor, named Buzz and seen in Figure 3, operates as a vibration sensor. It uses a piezo-film vibration sensor [7] that, when vibrated, generates enough current to activate the trigger on the power supply. Once activated, the sensor uses its onboard radio to transmit a packet announcing that the vibration event occurred. This sensor is intermittently powered, event triggered, senses only that the event occurred, and transmits immediately.

One application of this sensor is door monitoring. Buzz is attached to the door and when the door opens and closes it vibrates the sensor enough for the door open event to be logged. Similarly, the sensor can be attached to chairs to detect when the chair becomes occupied.

5.3 Breeze: Airflow Meter

Breeze (Figure 4) uses an analog MEMS airflow sensor [6] and opportunistic triggering to detect when an air vent is blowing air. The airflow sensor is placed near the vent to be monitored and the solar cell and power supply are placed in a nearby location conducive to photovoltaic harvesting. After each recharge, the power supply triggers automatically. Upon booting, the device powers the airflow sensor and waits for the output line to stabilize before beginning to sample




Figure 4: Airflow sensor.




Figure 5: Light based occupancy sensor.

the sensor at 1 kHz. After each sample the measurement is stored locally. This continues until Breeze depletes its stored energy reserves. On the next recharge and activation, Breeze averages the measurements it sampled on the last activation and wirelessly transmits the result. For airflow sensing this split operation is acceptable if subsequent activations occur quickly enough that the samples do not become out of date before they are transmitted. In our deployment we expect the power supply to be near an abundant source of energy, allowing for activations separated by no more than 5 s. If this was not the case, we could enlarge the energy store on the power supply to allow for sampling and transmission on the same activation, although at the expense of longer intervals between sampling.

Because Breeze does not need to keep time and its operation is simple enough to operate after a cold boot, Breeze uses the power supply's intermittent mode. It triggers opportunistically, uses an external sensor, and batches data before transmitting.

5.4 BigBen: Light/Occupancy Sensor

Our third sensor is a lights on/off sensor named BigBen (Figure 5). BigBen is designed to be installed very close to a light source allowing BigBen to do light state detection directly with the power supply: if the power supply is able to harvest then the light must be on. BigBen does not measure light intensity, but rather the binary state of the light. Specifically, it detects when the light is on based on when activations occur and when the light is off based on the lack of the trigger firing. When the light is on, many triggers will occur that signify that the light has simply remained on. Rather than store this redundant information, BigBen only makes note of when the light goes from off to on or vice-versa.

Due to the intermittent nature of BigBen, a state change of the light manifests as a change in the time between activations. If BigBen is activating at a certain rate and that rate suddenly drops, the light turned off, and if the rate suddenly increases, the light turned on. Detecting these rate change events presents an interesting challenge: how does BigBen know if the activations rate changes? Each activation appears identical and by default there is no time keeping. To address this, we supplement BigBen with a low-power Micro Crystal RV-3049-C3 [4] real-time clock (RTC) powered by a small (17 mAh) battery. This provides a reliable time source that BigBen can use to timestamp each activation. Now calculating activation deltas, and therefore light change events, is straightforward.

Using the RTC provides an additional advantage: change events can be timestamped against a global clock. This is useful because BigBen logs light change events locally rather than transmitting them. By having access to the RTC, BigBen can provide information about when the lights turned on and off, not just how many times. BigBen logs rather than transmits because it was designed to monitor rooms with poor connectivity over a short period of time (~3 months). This makes retrieving data from all of the sensors feasible. If the trial run demonstrates that light monitoring is worthwhile, the sensor can be redeployed and configured to transmit rather than log.

When lighting state can be considered a proxy for occupancy, BigBen can serve as a reasonable occupancy detection sensor. This is based on the observation that many rooms in buildings and offices contain motion sensitive lights. That is, when motion, and therefore people, are detected in a room the lights stay illuminated, and when no motion is detected the lights automatically turn off. BigBen can detect these light changes and use them to roughly infer when the room was occupied. This is still an estimation, but in certain cases it may be a suitable method for deploying an energy-harvesting occupancy sensor.

BigBen fits the architecture as a intermittent, opportunistic, energy-sensing, local-logging sensor.

6 Evaluation

To gauge the performance of our energy-harvesting sensor architecture we evaluate our three instantiations in the applications for which they were designed.

6.1 Recharge Rate

Because all three of our sensors rely on a intermittent solar cell based power supply, all three are limited to the recharge rate provided by the solar cell for any measurements or event detection. That is, the maximum sample rate of the sensor is capped by the maximum recharge rate. Figure 6 shows recharge rates as measured in various room types under different lighting conditions. Unsurprisingly, the office and atrium, which have incoming natural light, support much higher recharge rates than the lab. Directly attached to and facing out a window the power supply can recharge its capacitor every 1.9 s. On a wall in a bright, open atrium the capacitors recharge in 6.2 s. Near the middle of office the recharge period is 11 s. The recharge periods when the power supply is placed on the door in a lab that only has artificial light are




Figure 6: Recharge rate in varying lighting conditions. We measure the time the solar cell based energy-harvester takes to recharge in opportunistic trigger mode under different lighting conditions. As expected, the brighter the room the faster the recharge rate. Rooms with natural light (atrium and office) can support relatively fast recharge rates (in the 10s of seconds). Rooms with only artificial light (lab) cause the sensor to recharge more slowly, but can still support a sample every two minutes.




Figure 7: CDF of the interval between door open events. Plotted on a log scale x-axis is the CDF of time intervals between subsequent door opening events of a door over a month period. Also shown are the recharge times for the solar based energy-harvesting power supply in different lighting conditions. Sensors in rooms with natural light would be able to detect most door open events, and even in moderately lit rooms at least 65% of door open events would be detected.

44 s when the lights are at maximum brightness and 116 s in more moderate lighting. Different times of day, lighting types, the angle of light on the solar cell and many other factors affect the recharge rate of the power supply. The numbers presented here provide an estimate of maximum sensor sample rate in different building locations and not a guarantee of performance.

6.2 Door Open Detector

We configured Buzz, the vibration sensor, to detect door open events. This application lends itself naturally to the event trigger mechanism. The primary limiting factor of this sensor is the recharge rate of the power supply: if a door opens before the supply has recharged from the last open event, the sensor will be unable to detect the second opening. To estimate how the recharge rate enforced sample rate would affect the sensor's ability to detect when a door is opened, we analyze the time between subsequent door open events for a door over a one month period. Figure 7 shows the CDF of the interval between subsequent door openings plotted with the x axis in a log scale. Also marked in the fig-




Figure 8: Door open events across a 24 hour period. Buzz was deployed on the door of a lab with moderate lighting and no exterior windows. Each vertical line represents a door open event as detected by a ground truth sensor and the energy-harvesting Buzz sensor. The red dashed lines in the ground truth row are door open events that Buzz did not detect. Buzz successfully detected 66 out of the total 100 door open events. This aligns well with our expected detection rate from Figure 7.




Figure 9: Airflow metered over time. A Breeze sensor is placed near an air vent as the vent cycles on and off and changes its flow volume. The airflow sensor is uncalibrated and therefore we don't include y-axis units. While noisy, the data from Breeze are sufficient to determine when air is blowing from the vent and to determine when the flow changes.

ure are the approximate sensor recharge rates for the doors in the lab and office. In the office with natural light the sensor would miss only 2% of door open events. In the lab with normal lighting the sensor would have missed 35% of door open events over the month period. The missed door events would all have been within two minutes of each other, however, and it is feasible that finer temporal resolution may not be required for high-level building management algorithms.

To test Buzz in practice we attach the sensor to a door over a 24 hour period. We also use a mains powered door opening sensor for ground truth. Detected door openings versus actual door openings are displayed in Figure 8. Buzz is largely effective at detecting door open events, catching 66% of the openings. Openings that occurred after a long period of no activity were missed, likely due to the power supply being unable to charge when room was unoccupied and the lights were off. Also, events that occurred in rapid succession did not allow Buzz enough recharge time, causing those events to be missed as well. We discuss possible solutions as future work in Sections 7.1.1 and 7.1.2.




Figure 10: Light state over time. BigBen monitors the on/off state of a light fixture for 24 hours. It detects change events well and shows only a 40 s offset error from ground truth, likely due in part to jitter errors when initializing the RTC.

6.3 Airflow Meter

To demonstrate external sensing on an energy-harvesting node we developed Breeze, a sensor that periodically measures airflow. To evaluate Breeze, we placed its airflow sensor near an air conditioning vent and its power supply very close to an overhead light. Again, to collect ground truth we placed a mains powered sensor on the same vent. Figure 9 shows the results from the two sensors. Because Breeze is uncalibrated we omit the absolute magnitude from its samples. The on-to-off transition at the 42 minute mark and the corresponding drop in Breeze's readings shows that Breeze is, at a minimum, able to detect the on/off state of the vent. When the volume of airflow increases, at 2:06, Breeze's readings drop. This occurs because Breeze measures air speed, and the volume rate limiting is occurring at the vent. Since the HVAC unit continues to supply the same amount of air, the increased opening at the vent causes the air to move slower. Breeze was attached very close to an overhead light for this experiment and averaged a sample period of 7 s. When light is available, Breeze is a suitable sensor for airflow detection. When the lights are off or no daylight is present, Breeze will be unable to sample. Using the moving air to harvest, as in the Tinity [31] design, may be a viable solution for extending Breeze's sampling times.

6.4 Light/Occupancy Sensor

We deploy BigBen, our timestamping sensor that senses only based on its ability to recharge and eschews connectivity for local logging, on a light that is illuminated automatically only when there are occupants in the room. At the end of the measurement period the data is dumped from the node's storage and compared against the control system's ground truth commands as shown in Figure 10. In a 24 hour evaluation period, BigBen tracks ground truth state changes very well, with only a constant 40 s timestamp error.

Because BigBen does not transmit immediately after detecting a lighting change, it must timestamp the events itself, and this requires an RTC powered by a small rechargeable battery. To estimate how long BigBen will be able to successfully timestamp light change events we compare the battery and the RTC. We use a 17 mAh battery and the RTC [4] draws 800 nA, which means the RTC will run for approximately 2.4 years. In practice, however, we expect BigBen to serve as a temporary sensor which will be replaced with a networked device before the battery completely discharges.

7 Discussion

In this section we discuss possible issues raised by building, deploying, and managing energy-harvesting sensors.

7.1 Limitations

Driving down the size and power of building monitoring sensors comes with several limitations and drawbacks. We describe a few here and suggest possible solutions or workarounds.

7.1.1 *Darkness*

The key drawback to powering sensors with photovoltaics is that their energy source runs dry in dark conditions. Buildings are well lit when occupied or during the daytime, but after people leave or when the sun sets solar cell based energy-harvesters will cease charging and therefore stop sensing. For event based sensors this may be fine as the lack of people may imply the lack of events to detect. For other sensors, a rechargeable backing energy store, such as a battery, may be useful for continuing to sense while the power supply is unable to harvest. Take, for instance, a 3 V sensor node with a one minute sampling period that requires 1 mJ per sample. To continue sampling for a 63 hour weekend, the node would only need to have a 0.35 mAh battery on reserve. If the node's power supply harvests a surplus during the day, it is feasible the node could store the required energy to sample throughout the weekend.

7.1.2 *Recharge Rate too Slow*

When using an intermittent power supply the frequency of system operation is limited by the recharge rate of the power supply. If the recharge rate is slower than the required sample rate, the sensor fails to operate correctly. A possible solution may again be using a rechargeable battery as a buffer to store energy when harvesting a surplus and to fill in when not.

7.1.3 *Always On Receiver*

To facilitate real-time sensing each sensor must transmit shortly after collecting any data. Due to the relatively unpredictable run times provided by energy-harvesting power supplies, sensors likely cannot participate in any complex MAC layer wireless protocols. Therefore, these sensors likely need an always-on receiver in range for receiving transmissions. We propose three methods to overcome this issue. First, using long-range, sub-1 GHz radios may reduce the number of receivers needed to cover a building. Second, other, non-energy-harvesting devices may be able to receive and route packets for the energy-harvesting sensors. Plug-load or panel level power meters, for instance, typically have access to mains power and could easily handle the additional packets. Third, using Bluetooth and occupant's mobile phones as gateways may be a feasible option for sending data to the cloud.

7.2 Power Supply Implementation

A critical feature of the architecture is the trigger and shutdown interface between the power supply and the rest of the sensor node. The trigger is critical for supporting event detection and shutdown is necessary for maximizing sample rate by conserving stored energy. However, this functionality is not prevalent in many energy-harvesting power supply designs and ICs. Typically, these designs try to power the node immediately and for as long as possible until energy drops

beneath a threshold. This is reasonable for nodes that support a sustained power supply, but as size and energy budgets scale down the lack of functionality is limiting. Hopefully new circuit designs and energy-harvesting ICs emerge that enable these features.

7.3 Diagnosing Faulty Sensors

Identifying when a sensor is faulty is an important task in any sensor network deployment. When using energy-harvesting nodes, however, this becomes more challenging as a lack of communication with a node could simply be the result of the node being unable to harvest instead of being faulty. Discriminating the cause of a lack of packets from a sensor may require coordinated node placement to gain a consensus on the harvesting potential in a particular part of a building. If several sensors are in a similar area and only one is not communicating that might suggest the node is faulty.

7.4 Querying Sensors

Sensor nodes in our proposed architecture that are to report immediately or in batched transmissions need only be able to transmit. However, certain applications may wish to be able to query the state of sensor rather than wait for them to transmit. Because the time in which the node will be active is variable and difficult to predict, sending packets to energy-harvesting nodes may require a scheme similar to receiver initiated MAC layers. First, the intermittent node must announce its presence and that it is able to receive for a short period of time. Any always-on routing node in range that has a packet destined for the energy-harvesting sensor can then transmit it to the node. As long as the energy storage on the node is sized correctly, the sensor should be able to respond to the query as well, if needed.

8 Conclusion

Models of building conditions, occupant preference and activity, and energy consumption continue to advance, along with building control algorithms designed to improve energy efficiency and occupant comfort in buildings. Typically these models rely on input data from installed sensors to monitor the real-time activity inside the building. Sensors for this application are often one-off designs to validate a particular algorithm or are difficult to scale to an entire building for long-term operation due to size and power issues. To address this, we propose an architecture for designing energy-harvesting sensors targeted toward building monitoring applications. Leveraging energy-harvesting allows for small, deployable sensors that do not require AC mains power or battery replacement, and defining the architecture allows for a range of different sensors to be constructed from a set of common building blocks and interfaces.

To validate our architecture, we identify examples of application-specific energy-harvesting sensors from the recent literature and show how they decompose well into our framework. Further, we instantiate three new energy-harvesting sensors for common building applications and evaluate their feasibility and performance. We find that using a trigger abstraction is key for managing intermittent energy and that energy-harvesting sensors within our framework can solve the size and power problem for many common indoor activity recognition and building monitoring tasks.

9 Acknowledgments

This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STAR-net phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA. This material is based upon work partially supported by the National Science Foundation under grants CNS-1111541 and CNS-1350967, and generous gifts from Intel and Texas Instruments.

10 References

- [1] Building Energy Efficiency Frontiers and Incubator Technologies (Benefit) - 2014 Funding Opportunity Announcement. http://www1.eere.energy.gov/financing/solicitations_detail.html?sol_id=740.
- [2] Cube Sensors. <https://cubesensors.com/>.
- [3] Kill-A-Watt Wireless. <http://www.p3international.com/products/consumer/p4220.html>.
- [4] Micro Crystal RV-3049-C3 RTC. <http://www.microcrystal.com/index.php/products/real-time-clocks>.
- [5] Ninja Blocks Temperature Sensor. <http://shop.ninjablocks.com/products/temperature-and-humidity-sensor>.
- [6] Omron D6F-V03A1. [http://components.omron.com/components/web/pdfib.nsf/0/343C62397126E69185257201007DD6C6/\\$file/D6F-V_1010.pdf](http://components.omron.com/components/web/pdfib.nsf/0/343C62397126E69185257201007DD6C6/$file/D6F-V_1010.pdf).
- [7] Piezo Sensor - LDT Series. http://meas-spec.com/product/t_product.aspx?id=2484&terms=LDT0-028K*.
- [8] Department of Energy (DOE) Buildings Energy Data Book. <http://buildingsdatabook.eren.doe.gov/>, 2012.
- [9] B. Balaji, H. Teraoka, R. Gupta, and Y. Agarwal. Zonepac: Zonal power estimation and control via HVAC metering and occupant feedback. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 18:1–18:8, New York, NY, USA, 2013. ACM.
- [10] A. Beltran, V. L. Erickson, and A. E. Cerpa. Thermosense: Occupancy thermal based sensing for hvac control. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 11:1–11:8, New York, NY, USA, 2013. ACM.
- [11] S. DeBruin, B. Campbell, and P. Dutta. Monjolo: An energy-harvesting energy meter architecture. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, SenSys '13, 2013.
- [12] S. DeBruin, J. Grunnagle, and P. Dutta. Scaling the wireless AC power meter. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks*, IPSN '12, pages 153–154, New York, NY, USA, 2012. ACM.
- [13] A. Frye, M. Goraczko, J. Liu, A. Prodhan, and K. Whitehouse. Circulo: Saving energy with just-in-time hot water recirculation. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 16:1–16:8, New York, NY, USA, 2013. ACM.
- [14] L. A. Hang-yat and D. Wang. Carrying my environment with me: A participatory-sensing approach to enhance thermal comfort. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 21:1–21:8, New York, NY, USA, 2013. ACM.
- [15] S. R. Iyer, V. Sarangan, A. Vasani, and A. Sivasubramaniam. Watts in the basket?: Energy analysis of a retail chain. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 4:1–4:8, New York, NY, USA, 2013. ACM.
- [16] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity AC metering network. In *IPSN '09: Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, pages 253–264, Apr. 2009.
- [17] Y. Jiang, K. Li, L. Tian, R. Piedrahita, X. Yun, O. Mansata, Q. Lv, R. P. Dick, M. Hannigan, and L. Shang. MAQS: A mobile sensing system for indoor air quality. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11, pages 493–494, New York, NY, USA, 2011. ACM.
- [18] D. Jung and A. Savvides. Estimating building consumption breakdowns using on/off state sensing and incremental sub-meter deployment. In *SenSys'10: In Proceedings of the Eighth ACM Conference on Embedded Networked Sensor Systems*, Nov. 2010.
- [19] J. Kadengal, S. Thirunavukkarasu, A. Vasani, V. Sarangan, and A. Sivasubramaniam. The energy-water nexus in campuses. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 15:1–15:8, New York, NY, USA, 2013. ACM.
- [20] N. Klingensmith, D. Willis, and S. Banerjee. A distributed energy monitoring and analytics platform and its use cases. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 5:1–5:8, New York, NY, USA, 2013. ACM.
- [21] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso. A platform for ubiquitous sensor deployment in occupational and domestic environments. In *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*, Cambridge, Massachusetts, Apr. 2007.
- [22] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: Using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 211–224, New York, NY, USA, 2010. ACM.
- [23] P. Mansourifard, F. Jazizadeh, B. Krishnamachari, and B. Becker-Gerber. Online learning for personalized room-level thermal control: A multi-armed bandit framework. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 20:1–20:8, New York, NY, USA, 2013. ACM.
- [24] A. Marchiori and Q. Han. Using circuit-level power measurements in household energy management systems. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '09, pages 7–12, New York, NY, USA, 2009. ACM.
- [25] P. Martin, Z. Charbiwala, and M. Srivastava. DoubleDip: Leveraging thermoelectric harvesting for low power monitoring of sporadic water use. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, SenSys '12, pages 225–238, New York, NY, USA, 2012. ACM.
- [26] A. Parisio, D. Varagnolo, D. Risberg, G. Pattarello, M. Molinari, and K. H. Johansson. Randomized model predictive control for HVAC systems. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 19:1–19:8, New York, NY, USA, 2013. ACM.
- [27] U. Park and J. Heidemann. Data muling with mobile phones for sensor networks. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, SenSys '11, pages 162–175, New York, NY, USA, 2011. ACM.
- [28] J. Ploennigs, B. Chen, A. Schumann, and N. Brady. Exploiting generalized additive models for diagnosing abnormal energy use in buildings. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 17:1–17:8, New York, NY, USA, 2013. ACM.
- [29] A. Rogers, S. Ghosh, R. Wilcock, and N. R. Jennings. A scalable low-cost solution to provide personalised home heating advice to households. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 1:1–1:8, New York, NY, USA, 2013. ACM.
- [30] U.S. EPA. The inside story: A guide to indoor air quality, Sept. 1993.
- [31] T. Xiang, Z. Chi, F. Li, J. Luo, L. Tang, L. Zhao, and Y. Yang. Powering indoor sensing with airflows: A trinity of energy harvesting, synchronous duty-cycling, and sensing. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, SenSys '13, pages 16:1–16:14, New York, NY, USA, 2013. ACM.
- [32] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta. Grafting energy-harvesting leaves onto the sensor tree. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks*, IPSN '12, pages 197–208, New York, NY, USA, 2012. ACM.