

Continuous sensing on intermittent power

Majid, Amjad Yousef; Schilder, Patrick; Langendoen, Koen

DOI

[10.1109/IPSN48710.2020.00-36](https://doi.org/10.1109/IPSN48710.2020.00-36)

Publication date

2020

Document Version

Final published version

Published in

2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)

Citation (APA)

Majid, A. Y., Schilder, P., & Langendoen, K. (2020). Continuous sensing on intermittent power. In L. O'Conner (Ed.), *2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)* (pp. 181-192). Article 9111002 IEEE. <https://doi.org/10.1109/IPSN48710.2020.00-36>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Continuous Sensing on Intermittent Power

Amjad Yousef Majid
a.y.majid@tudelft.nl
Delft University of Technology

Patrick Schilder
p.t.schilder@student.tudelft.nl
Delft University of Technology

Koen Langendoen
k.g.langendoen@tudelft.nl
Delft University of Technology

ABSTRACT

The main obstacles to achieve truly ubiquitous sensing are (i) the limitations of battery technology - batteries are short-lived, hazardous, bulky, and costly - and (ii) the unpredictability of ambient power. The latter causes sensors to operate intermittently, violating the availability requirements of many real-world applications. In this paper, we present the *Coalesced Intermittent Sensor* (CIS), an intermittently-powered “sensor” that senses continuously! Although a single node will frequently be off charging, a group of nodes can –in principle– sense 24/7 provided that their awake times are spread apart. As communication is too expensive, we rely on inherent component variations that induce small differences in power cycles. This basic assumption has been verified through measurements of different nodes and power sources. However, desynchronizing nodes is not enough. An important finding is that a CIS designed for certain (minimal) energy conditions will become synchronized when the available energy exceeds the design point. Nodes employing a sleep mode (to extend their availability) do wake up collectively at some event, process it, and return to charging as the remaining energy is typically too low to handle another event. This results in multiple responses (bad) and missing subsequent events (worse) due to the synchronized charging. To counter this undesired behavior we designed an algorithm to estimate the number of active neighbors and respond proportionally to an event. We show that when intermittent nodes randomize their responses to events, in favorable energy conditions, the CIS reduces the duplicated captured events by 50% and increases the percentage of capturing entire bursts above 85%.

CCS CONCEPTS

• **Human-centered computing** → *Empirical studies in ubiquitous and mobile computing.*

KEYWORDS

Embedded systems, Energy harvesting, Ubiquitous computing, Intermittent Systems

1 INTRODUCTION

Batteries may compromise the viability of sensor nodes in various ways. Batteries are bulky, short-lived, hazardous, and expensive. To ameliorate the battery problem, researchers have been investigating different alternatives to extend lifetime and reduce costs and form factor. The reduction in power consumption of recent microcontrollers (MCUs) and the advances in energy-harvesting (EH) circuitry have enabled the emergence of battery-free EH sensors. These sensors elide the constraints of batteries and extract power from ambient energy sources such as sunlight and RF emissions.

Ambient energy sources provide perpetual power. However, ambient power is usually too weak to directly power a sensor node [1]. Therefore, an EH node first buffers the harvested energy until a usable amount has been accumulated; then it operates, for a short period of time, until the buffered energy has been exhausted [2]. Consequently, battery-less EH sensors operate intermittently (Figure 1).

Intermittent power introduces a set of new challenges that are under ongoing investigation. For example, [2–7] studied the intermittent computation problem, which is concerned with the preservation of application progress and data consistency under frequent power failures; the authors of [8] investigated the timely operation challenge, which is concerned with data freshness after a power interrupt; and [9] introduced event-driven execution for the intermittent domain, which deals with input and output operations under arbitrarily-timed power loss.

Despite these notable advances, intermittently-powered sensors suffer from a new fundamental shortcoming: *the intermittent availability of the system*. Being frequently off charging compromises the value of these devices. For example, a sensor that has a low probability (e.g., 10% [10]) to be available (on) when an event of interest occurs has no value. Overcoming the intermittent availability challenge without changing the size of the device or re-including batteries requires a novel approach that explores new design dimensions.

1.1 Vision and Application

Miniaturized sensors are non-intrusive devices. Therefore, they can be embedded in locations that are not suited for the others (enabling new applications). Miniaturizing sensors, however, introduces the significant challenge of powering them. On the one hand, batteries make these sensors *continuously* available -for sensing opportunities-, but with an environmental footprint and only for a short period of time. On the other hand, removing batteries and relying on ambient energy make them available for a long period




Figure 1: Harvested-energy profile. Ambient power is weak; therefore, it is usually buffered. The buffered energy is then consumed to operate the device. The operation period is often short as power consumption is much higher than the energy harvesting rate.




Figure 2: A Coalesced Intermittent Sensor (CIS) is a group of intermittently-powered nodes that sense continuously despite the intermittent power supply. CIS exploits the inherent randomization of energy harvesting systems, if available, and introduces artificial randomization, when needed, to preserve continuous sensing.

of time, but intermittently. Our vision¹ is that by combining *multiple* battery-less EH sensors we can create a new *virtual* sensor that operates permanently (no batteries) and reliably (continuously available): we call this sensor the *Coalesced Intermittent Sensor* (CIS).

Sensors with such characteristics would allow us to add a cheap and maintenance-free sensing layer to many objects, making them smart and interactive. For example, one can imagine developing smart wallpaper that users can interact with. Smart wallpaper with embedded microphones can enable direct in-building human-to-object communication (Figure 2). Such a permanently operating sensor can be deployed, for example, in kids’ playgrounds to monitor their occupancy. These battery-less sensors can enable interactive and safe-to-dispose sports rugs (that count how many times a person has jumped on them) or play rugs for kids. In short, we would like to develop small sensors with permanent and continuous sensing capabilities.

1.2 Research Challenges

Many sensing applications require the sensor to be available when there is a change in the monitored environment. EH battery-less sensors can provide cheap and maintenance-free sensing, but they do not meet the availability requirements of many real-world applications.

C1-Approach continuous availability on intermittent power: An EH battery-less sensor is frequently off, spending most of the time charging. One way to increase the availability of the system is by using multiple nodes. However, coordinating the nodes’ awake times using communication may introduce prohibitive overhead as a scattering algorithm must be regularly executed, and messages for synchronizing nodes’ clocks and reserving time slots need to be repeatedly exchanged. Thus, the challenge is *can we exploit some of the inherent characteristics of EH battery-less sensors to distribute nodes’ awake times without the need for communication?*

C2-Continuous sensing on intermittently powered sensors: Even when the collective availability of intermittent sensors approaches 100%, the emerging overall sensing behavior may still be

intermittent. Event-trigger sensors sleep in low-power mode waiting for an event to wake them up. When ambient energy rises, the EH rates of these sensors may equal (or approximate) their sleeping mode power consumption. Under such energy conditions, these sensors become available for an extended period of time. Therefore, when an external event arrives, nodes respond collectively, which exhausts their energy buffers, making them unavailable for the next set of events. This is, particularly, a significant problem when events arrive in bursts, like a command of a few words (e.g., “light on”). Thus, the challenge is, *how to prevent EH battery-less sensors from synchronizing their power cycles on some of the incoming events?*

C3-Efficient sensing on intermittent sensors: One of the main factors that determine the intermittency pattern of an EH battery-less sensor is the richness of ambient energy. For example, at mid-noon under direct sunlight, even a small solar panel can power a sensor node continuously. In such conditions (favorable energy conditions), using plenty of intermittent sensors would only result in duplicated work that leads to duplicated messages when the data is being communicated to a sink node: a continuously-powered node acts as a gateway for such sensors to communicate with other layers of the Internet of Things. These messages will collide as they will be generated at approximately the same time, and if some of them are received by the sink, then they waste energy as they carry the same information. Thus, the challenge is, *how to reduce the number of duplicated event detections?*

1.3 Contributions

In this paper, we tackle the paradox of continuous sensing on intermittently-powered sensors. We studied the inter-relationship between the power cycles of EH battery-less devices, the emerging collective behavior, and the effect of the change in ambient energy on this behavior. In particular, this paper makes the following key contributions:

- We show how to approach continuous sensing using multiple intermittently-powered sensors. For that, we **modeled** the collective effective availability—the system availability that leads to successful sensing—of a group of intermittent sensors and **validated** our models using simulation and on real hardware against different ambient energy sources.
- We introduce a new type of virtual sensor, showing its capabilities and limitations. This **Coalesced Intermittent Sensor (CIS)** is the abstraction of a group of intermittently-powered sensors that achieves maximum statistical availability by exploiting (inherent) randomization to spread nodes’ awake times uniformly.
- Contrary to common sense, we show how favorable energy conditions can deteriorate the performance of a CIS. We, therefore, equipped the CIS with **an new algorithm** that makes it ambient-energy aware. This algorithm enables the nodes to determine their own duty cycles (without requiring additional hardware), and the average number of alive nodes (without requiring communication). This information can effectively be used by the nodes to decide when to back off to avoid duplicated event detection and availability interruptions (implicit synchronization in favorable harvesting conditions).
- We prototype, evaluate, and demonstrate the feasibility of the Coalesced Intermittent Sensor concept in the form of a voice-control

¹An alternative approach is to combine EH with a (small) rechargeable battery [11, 12].

commands recognizer, the **Coalesced Intermittent Command Recognizer (CICR)**. We chose to develop a command recognizer as voice is a natural way for the human to interact with small devices. Moreover, words allow us to easily experiment with individual event arrivals and events that arrive in bursts. However, the goal of this paper is *not* to present a novel word recognition technique. Instead, we adapt a classical word recognition algorithm to make it power-failure immune. Yet, our CICR prototype is the first intermittent command recognizer, shedding light on the potential of intermittent systems.

2 RELATED WORK

Recent advances in ultra-low-power microcontrollers along with the development of energy harvesters have enabled the creation of stand-alone battery-free sensors. These sensors operate intermittently because the power that they harvest is weak and volatile.

2.1 Energy-harvesting systems

Energy harvesters have the potential to power devices indefinitely as they collect energy from perpetual energy sources. Sunlight, vibration, and radio frequency (RF) waves are examples of such energy sources. The power harvested from these sources vary wildly, for example, RF harvestable power ranges from nW-scale when harvested from ambient signals to μ W-scale when collected from a dedicated RF signal emitter, and solar power varies from tens of μ W to tens of mW when it is harvested by a solar panel of a few cm^2 illumination surface [2, 13].

Many battery-less EH platforms have been proposed. Some of them rely on dedicated external energy sources such as WISP - and its variants-, a general wireless sensing and identification platform [14–16]; WISPCam, an RF-powered camera [17] and, the battery-free cellphone [18]. Others, harvest from ambient sources such as the ambient backscatter tag [1], and the solar-powered tag [19]. Platforms that facilitate the development of battery-less EH systems have also been proposed. For instance, Flicker [20], a prototyping platform for battery-less devices; EDB [21] an energy-interference-free debugger for intermittent devices; Capybara [22], a re-configurable energy storage architecture for EH devices; and Stork [23] is protocol for over-the-air programming for batteryless EH devices.

However, *there is no EH platform that considers the abstraction of many intermittent sensors (or nodes) and exploits the statistical energy harvesting differences between them to provide reliable sensing.*

2.2 Intermittent execution

Intermittent execution models enable applications to progress despite frequent power failures [5, 24–28]. To this end, they decompose an application into several small pieces and save the state of the computation on the transitions between these code segments. Therefore, intermittent applications do not return to the beginning of the program (i.e., `main()`) after each power failure. Instead, they resume execution from the last successfully saved progress state.

Mementos [3] proposed a volatile memory *checkpoint-based* approach to enable long-running applications on intermittently powered devices. DINO [4] enables safe non-volatile memory access despite power failures. Chain [5] minimizes the amount of data

needed to be protected by introducing the concepts of *atomic tasks and data-channels*. Hibernus [6, 29] measures the voltage level in the energy buffer to reduce the number of checkpoints per power cycle. Ratchet [24] uses compiler analysis to eliminate the need for programmer intervention or hardware support. HarvOS [26] uses both compiler and hardware support to optimize checkpoint placement and energy consumption. Mayfly [8] enables time-aware intermittent computing. InK [9] introduces event-driven intermittent execution. *For our prototype implementation we adopt a power failure protection approach similar to that of DINO [4], see Section 4.2.*

2.3 Explicit duty-cycle desynchronization

Explicit duty-cycle desynchronization has been proposed in the wireless sensor networks literature [30–32]. These (biologically-inspired) algorithms, however, cannot be applied to desynchronize intermittently-powered nodes as they assume that nodes (i) are able to listen to other nodes, and (ii) can maintain a notion of global time (slots). Listening is expensive, and keeping track of time is difficult at best when nodes can power down at random moments. We therefore adopt a best-effort approach.

2.4 Speech recognition

The speech recognition problem has been tackled from many angles and has experienced many breakthroughs. For example, the dynamic time warping (DTW) algorithm enables matching voice signals with different speeds (or time) [33]. Approaches based on Hidden Markov Models showed much better performance than DTW-based ones [34]. Hence, they became the standard techniques for general-purpose speech recognition until artificial intelligent algorithms [35] outperform them.

From a recognition complexity standpoint, we can classify the speech into *spontaneous speech*, *continuous speech*, *connected word*, and *isolated word* [36]. The *continuous* and *spontaneous speech* are the closest to natural speech, but they are the most difficult to recognize because they need special methods to detect words boundaries [36]. This is less the case for the *connected word* type, where a minimum pause between the words is required. The type with the least complexity is the *isolated word*, as it requires a period of silence on both sides of a spoken word.

Speech recognition on resources—memory, computation power, and energy—limited platforms is challenging, to say the least. Therefore, *our command recognizer targets isolated-word type of speech.*

3 COALESCED INTERMITTENT SENSOR

The Coalesced Intermittent Sensor (CIS) is the abstraction of a group of EH battery-less sensor nodes seeking to approximate the continuous sensing availability characteristic of a battery-powered sensor. The design of a CIS needs to consider four main aspects: (i) how the nodes' awake time is distributed; (ii) the consequence of emulating continuous sensing availability by chaining multiple short on-times; (iii) the effect of the environment on the CIS's availability; and (iv) the spacial coverage of the event of interest, which determines the diameter of the CIS.

However, let us first characterize the power cycle of an EH battery-less device. An EH intermittent node frequently switches




Figure 3: A Coalesced Intermittent Sensor's availability is the emerging collective on-time of its intermittent nodes' on-times. The difference between the power cycles leads to a constant relative shift between the nodes' duty cycles. This, in turn, causes their on-times to be uniformly distributed on the overall power cycle. The red bars indicate a minimum CIS time span—CIS nodes are overlapping—whereas the green bars show the maximum time span of the CIS.

between off and on, charging energy and operating. We can characterize, from a time perspective, this charge-discharge (or power) cycle using the following notation, (t_{on}, t_p) , where t_{on} is the node's uptime interval, and $t_p := t_{\text{on}} + t_{\text{off}}$, where t_{off} is the node's charging time interval.

3.1 Sensing

The ability of a CIS to sense depends on the availability of its intermittent nodes and on the characteristics of the event of interest.

3.1.1 Coalesced availability. The CIS's availability is the projection of its underlying intermittent nodes' on-times on the time axis. To determine the expected availability of a CIS, the strategy being employed to distribute its nodes' on-times must be first specified.

Explicit on-time division strategy. A CIS can build on top of the recent advancements in passive (light or RF) communication [1] and ultra-low-power timers [8] to apply a time-division multiplexing strategy, minimizing overlapping on-times. For example, a node calculates its average on-time \bar{t}_{on} and off-time \bar{t}_{off} for a certain number of power cycles. Then, it encodes the information $(\bar{t}_{\text{off}}, \bar{t}_{\text{on}})$ in a message and broadcasts it to others at the beginning of its next power cycle. When a node receives this message, it can then adjust its power cycle, relative to the transmitting node's cycle, by increasing (or decreasing) its power consumption to shorten (or lengthen) its on-time and subsequently shift its power cycle to a different time slot.

With such explicit on-times control strategy, a CIS of N nodes with on-time of \bar{t}_{on} and off-time of \bar{t}_{off} will have an availability $= \min\left(N \times \frac{\bar{t}_{\text{on}}}{t_p}, 100\%\right)$. However, we expect such an approach to introduce significant overhead as a scattering algorithm (e.g., [31]) must be frequently executed, messages need to be exchanged, and clocks should be synchronized. Therefore, we propose a different on-times spreading strategy.

Implicit on-time division strategy. With no information being exchanged between intermittent nodes, the best CIS can do is to




Figure 4: Coalesced Intermittent Sensor availability for a different number of nodes and different duty cycles. The nodes are uniformly distributed and the CIS on-time evolves, when adding new nodes, according to equation 1.

uniformly distribute its node's on-times and maintaining this distribution over time. The key observation to approach uniform distribution is to ensure that the lengths of the node's power cycles are randomized, avoiding nodes being in lockstep indefinitely.

Let us start by assuming that we have a CIS of two nodes with idealized power cycles and these nodes have the same initial conditions. The availability of this CIS equals t_{on} as the nodes are in perfect synchronization (the two nodes wake up and power down together). To extend the availability of this CIS, one of the nodes should shift its on-time away from the other. If one of the nodes sleeps for t units of time, then the on-time of this power cycle will be $t_{\text{on}} + \Delta t$. Consequently, the length of this power cycle will be $t_p + \Delta t$, delaying the next awake time by Δt . If the node sleeps only once, then availability of the CIS will equal $\min(2 \times t_{\text{on}} | t_{\text{on}} + \Delta t)$

However, if the initial conditions are unknown, then shifting a node's on-time a constant number of times may cause the initially desynchronized nodes to become synchronized, collapsing the CIS's availability instead of extending it. Therefore, a safer option is to *constantly* shift the awake time of the node. In this case, the on-time will shift over the entire power cycle of the other node, spending $\frac{t_{\text{off}}}{t_p}$

and $\frac{t_{\text{on}}}{t_p}$ of the time overlapping with the other node's off-time and on-time, respectively. This behavior is illustrated in Figure 3, where node 1 and node 2 have power cycles of (2,6) and (1,5). Following the time axis from the left to the right, we can observe that the position of the on-time of node 2 is shifted by -1 unit of time relative to the on-time of node 1 after each power cycle of node 1. This implies that the on-times of the two nodes are $\frac{1}{3}$ of the time cluster together and $\frac{2}{3}$ of the time they are apart (from an external event standpoint, the on-times are uniformly distributed over the longest power cycle, as they have the same probability to be anywhere when the event arrives). To model the availability of a CIS of N nodes, we first model the nodes' on-times and power cycles. If we represent the on-time of a node with a random variable R_n and find its expected value $E(R_n)$ then we can approximate *any* CIS node's on-time with the mean of the expected values of the nodes' on-times, i.e., $t_{\text{on}} = \frac{1}{N} \times \sum_{i=1}^N E(R_n)^i$ (intuitively, since we assume CIS nodes have the same energy buffer, their expected on-times should approach the same value; we relax this assumption in Section 5.2.4). Using a similar analogy, we can define the mean of the expected




Figure 5: Nodes' power cycles length for different ambient energy sources, and different energy buffer sizes.

values of the power cycles lengths as $t_p = \frac{1}{N} \times \sum_{i=1}^N E(R_p)^i$. Now, we can model the availability of a CIS of N nodes as:

$$A_v(N) = A_v(N-1) + (1 - A_v(N-1)) \times \frac{t_{on}}{t_p}, \quad (1)$$

for the initial case where $N = 1$ we define $A_v(0) := 0$. Figure 4 shows the availability of CIS when $N \in \{1, 2, \dots, 20\}$ and nodes' duty cycles $\frac{t_{on}}{t_p} \in \{10\%, 20\%, \dots, 50\%\}$. We can conclude from the above discussion that to approach uniform distribution of nodes' on-times, the lengths of the power cycles need to be randomized².

The power cycles of EH battery-less devices are inherently randomized and different because the power source (ambient energy) is volatile and the harvesters are not perfect devices (notice that, even battery-powered wireless sensor nodes require a synchronization protocol to correct for the drift in their local clocks). Our own measurements using different EH devices and different energy sources, i.e., solar and RF, also confirm that the power cycles of intermittent nodes are different and randomized (Figure 5). Therefore, we expect their on-times to be uniformly distributed (we will challenge our expectation in Section 5).

3.1.2 Events classification. The availability of a CIS is not a single stretched interval: it is a chain of short intervals. Therefore, it is important to classify from a CIS perspective which types of events the CIS is best suited for.

- *Short events:* are events that can be captured using a single intermittent node. For example, a spoken word can be seen as a short event if the energy needed to record it is less than what the energy buffer, i.e., the capacitor, can store.
- *Long events:* are events that need more energy to be completely captured than what the energy buffer can store. Long events can be subdivided into three categories:
 - *Simple:* is a long event that can be captured using a single intermittent node—capturing part of it is sufficient to obtain all the information of interest—such as the sound produced by the friction between two moving parts of an engine.
 - *Burst:* is a group of short events that requires multiple intermittent nodes to be captured such as a command of a few words (e.g., “room temperature up”).
 - *Complex:* is a long event that must be fully captured to be recognized. For example, sampling a gyroscope attached to a moving device (e.g., a toothbrush).

Based on the above classification, we can argue that designing a CIS for long events is not like designing it for capturing short ones. For

²Note that, having power cycles of lengths that are multiples of each other is a very unlikely as nodes' energy buffers are assumed to be of the same size.




Figure 6: Simulating the availability, the effective availability, and successfully captured events of a CIS of 10 nodes with a node duty cycle $\in \{10\%, 20\%, \dots, 50\%\}$.

example, while capturing a short event may require continuous CIS availability, capturing a long simple event that is longer than the power cycle t_p does not require extending the availability of a single intermittent node. Furthermore, capturing a long complex event may require data fusion and processing that require the CISs' nodes to communicate the raw data to a more powerful node, which may lead to significant overhead. However, this paper focuses on short and long burst events as they cover a wide range of applications (e.g., voice-controlled human-object interface).

3.1.3 Effective Availability. Approaching continuous availability does not mean that a CIS can successfully capture all events. It can happen that an event is being only partially captured by one or more nodes, which may lead to unsuccessful event detection. Therefore, it is important to specify the effective availability of a CIS that leads to a successful event capturing (which we assume leads to successful sensing).

Polling-based Sensing. Let us assume that we have a CIS of a single intermittent node monitoring a short event of length t_e . For capturing the entire event, the event has to arrive within the interval, $t_{on} - t_e$, which we call, the effective on-time of an intermittent node. Therefore, the effective availability of a CIS of N nodes is the joined effective on-times of the underlying intermittent nodes, which can be modeled as,

$$A_v(N) = A_v(N-1) + (1 - A_v(N-1)) \times \frac{t_{on} - t_e}{t_p}, \quad (2)$$

Event-driven Sensing. An intermittent sensor has a limited energy budget per power cycle. When it is tasked with a polling-based sensing activity, its energy consumption, generally, switches between two levels: zero when charging and maximum when sensing. However, in event-based sensing, a node puts its MCU into low-power mode and waits (or listens) for an external event to wake up the MCU. For example, in our prototype, a voice-controlled command recognizer, we exploit the microphone's wake-on-sound feature to send an interrupt to the MCU, which will then start recording the sound samples from the microphone. This wake-on-event style of operation is important as the minimal energy consumption during sleep significantly prolongs the period during which an event can be handled (for example, our prototype consumes 7 times less energy during sleep compared to being active). To model the effective CIS availability when it is tasked with event-based sensing, the

change in energy consumption between the sleep and active mode must be taken into account. Since the event itself times when the node changes its energy consumption, we can model the effective availability as,

$$A_v(N) = A_v(N-1) + (1 - A_v(N-1)) \times \frac{t_s - \left(t_e \times \frac{p_a}{p_s}\right)}{t_{sp}}, \quad (3)$$

where t_s is the expected sleep time of the CIS nodes, $t_{sp} := t_s + t_{off}$, and p_a and p_s are the power consumption in active and sleeping mode, respectively. Notice that, there is a subtle point about equation 3 as when an event arrives the node wakes up, consuming more energy. Therefore, its uptime shrinks. We, for simplicity, modeled this effect by extending the event time with the same factor. This is sufficient to say if that the event will be fully captured or not (effective availability).

3.1.4 Simulation. As a first sanity check on our models, we simulated 10^5 power cycles of a CIS of 10 nodes (Figure 6). The duty cycles of the nodes range from 10% to 50%, while the event length is fixed at 3% of the power cycle length, t_p . The on-times and event arrivals were uniformly distributed over the power cycles. The results clearly confirm our models and support our argument about the distinction between CIS's availability and effective availability (notice that the percentage of captured events matches the effective availability). The importance of this distinction—availability versus effective availability—is a function of the value $\frac{t_e}{t_{on}}$; observe the difference between availability and effective availability when nodes' duty cycle is 10% (large effect) and 50% (negligible effect).

3.2 Environment

Ambient energy controls the availability of CIS nodes. Consequently, it also controls their collective response to external events. When it rises, it extends nodes' on-times that may lead node's power cycles to be synchronized on the arrival of some external events, compromising the CIS's overall availability. To overcome this problem the CIS nodes must be power-state aware and able to estimate the number of active nodes in the CIS.

3.2.1 Power States. A CIS can experience a wide range of ambient power intensities. For example, a solar-powered CIS may harvest no energy at night, modest energy from artificial light, and abundant energy from direct sunlight. Generally, we can identify four different CIS powering states:

- **Targeted power state**—These are the powering conditions that a CIS is designed for. In these conditions, the CIS should work intermittently and have sufficiently randomized power cycles to uniformly distribute its intermittent nodes on-times and meet the desired availability (Figure 4). In general, the targeted powering conditions should be near worst energy harvesting conditions to ensure that the system is properly functioning for the majority of the time.
- **Under-targeted power state**—Ultimately, the ambient energy is an uncontrollable power source, and it is not hard to imagine scenarios where a CIS will be under-powered or even comes to complete and long power down (for example, a solar CIS will come to a perpetual power down in darkness). In general,




Figure 7: Capturing events may lead to synchronizing the power cycles of nodes that were in low-power mode. In particular, if some of the nodes power down while capturing an event, then these nodes should have slept (prior to the event) longer than the nodes that capture the event. This means that the nodes that have died spent their buffered energy slower (they stayed longer in sleep mode); therefore, their overall uptime is longer than the uptime of the nodes that captured the entire event. In other words, the nodes that woke up earlier have stayed longer on, and therefore, the next power cycles are more synchronized, see the Δt before and after the event.

for under-targeted energy conditions, the CIS behavior can be considered as undefined.

- **Hibernating power state**—In event-driven sensing, nodes sleep in low-power mode waiting for an event to wake them up. This mode extends CIS nodes' on-times and makes them overlap significantly. Moreover, the on-times overlap even more, when ambient energy level rises (favorable energy conditions). If an event arrives in such conditions, it will wake up many nodes, causing them to consume their buffered energy much faster. Some of these nodes will power down before capturing the entire event, while others survive. This difference in how much energy is spent in active and sleep mode causes these nodes to tend to synchronize their power cycles after the event. To understand why let us analyze the example presented in Figure 7. The figure shows the power traces of two nodes. the nodes consume 7 times more power in active than in sleep mode (our prototype has a similar power consumption ratio, Table 3). Further, it shows that a node sustains the low-power mode for 90 units of time, $t_s = 90$; therefore, the maximum buffered energy can be calculated as

$$E_{buf} = t_s \times p_s,$$

where p_s is a node's power consumption in sleep mode. If we focus on the power cycle with the first event, then we see that node R powers up at t_0 , and it remains in low-power mode for 50 units of time, whereas node G spends only 20 units of time before the event arrives. Now, we can calculate when these two nodes will power down and compare the difference, Δt , before and after the event. A node will turn off when the buffered energy is depleted. This can be expressed as follows,

$$E_{buf} = t_{se} \times p_s + t_{on} \times p_a,$$

where t_{se} is the sleep time of the power cycle that the event arrives in, t_{on} is a node's on-time and p_a is the power consumption in active mode, which can be expressed as $p_a = \delta \times p_s$. t_{se} and δ are given and p_s can be eliminated; thus to find when the nodes will




Figure 8: The discharge time (on-time) of an EH node when it is allowed to harvest energy while executing and not.

power down we need to find t_{on} for both nodes. By substituting the given values we find t_{on} to be 10 and 5.7 units of time for the G and R node, respectively. Therefore, the difference between the end points of the power cycles, Δt , becomes 4.3 while it was 30 before the event (notice, $\frac{30}{7} \approx 4.3$). In summary, a node that dies while capturing the event should have started its power cycle *before* a node that captures the entire event. Further, the uptime of the node that is died while capturing is *longer* than the node that captures the event because it spends less time in active mode. Therefore, the difference, Δt , between the power cycle of the died-while-capturing node and the power cycle of the node that successfully captures the event becomes smaller. When the events arrive in burst this becomes a significant problem, as a CIS will capture multiple copies of the first event, while missing the subsequent ones.

- *Continuous power state*—Under direct mid-noon sun a tiny solar panel may provide sufficient power to run a sensor node continuously. In such conditions, a CIS node will be available and able to sense continuously. Therefore, the job of a single node will be repeated N times, and instead of sending a single message to a sink—to push the data to the Internet— N identical messages will be sent. These messages will collide as they are sent at about the same time, causing the information to be lost; if they arrive, however, they -except the first one- will waste energy of the sink as they carry the same information.

The inefficiencies highlighted in the hibernating and continuous power states can be mitigated by enforcing randomization on the response of intermittent nodes: when a node is woken up by an external event it responds to that event with a certain probability. However, if the randomized response is enforced all the time, then the CIS will have a lower probability of catching events during the targeted energy conditions state. Therefore, the CIS has to distinguish between the targeted and above-targeted energy conditions and randomize its response only during the hibernating and continuous power states.

Furthermore, responding with a constant probability during the above-targeted energy conditions is inefficient, as the number of active nodes is a function of the total number of intermittent nodes and the power intensity at that time. Therefore, efficient randomization requires intermittent nodes to estimate the number of active nodes and respond proportionally. Our proposed algorithm for estimating the number of active nodes depends on the nodes' ability to measure their on-times and off-times.

Algorithm 1 off-time estimation

```

1:  $R_{ctr}++$                                 ▶ reboot counter
2:  $E_{buf}$                                   ▶ Size of energy buffer
3:  $t_a$                                      ▶ time of discharging  $E_{buf}$  at load  $a$ , no harvesting
4:  $X_{cy}$                                    ▶  $X$  power cycles
   ▶ Code executed on each  $X$  power cycles
5: if ( $R_{ctr} == X_{cy}$ ) then
6:    $f_{LOAD}(a)$                             ▶ set node load to  $a$ 
7:    $t_{on} \leftarrow TIME()$                 ▶ measure time until power down
8: end if
   ▶ Code executed on each  $X + 1$  power cycles
9: if ( $R_{ctr} == X_{cy} + 1$ ) then
10:   $\Delta t = t_{on} - t_a$                     ▶ time difference due to charging
11:   $E_{har} \leftarrow E_{buf} \times \frac{t_a}{\Delta t}$   ▶ harvested energy
12:   $P_{in} \leftarrow E_{har} \div t_{on}$          ▶ incoming power
13:   $t_{off} \leftarrow E_{buf} \div P_{in}$ 
14:   $R_{ctr} = 0$ 
15: end if

```

3.2.2 Intermittent Timing. Timing is a key building block of sensing systems. It is, however, missing on intermittent nodes unless an additional dedicated (RC-based) timer is included [8]. Here we propose an alternative that does not require additional hardware. This alternative does not only enable time estimation but also ambient energy richness, which is very important for estimating the number of a node's active neighbors. But, *how a node can time its own on/off cycle?*

Intermittent nodes fail abruptly; therefore, a persistent timer is needed to measure node's on-time. A simple way to emulate persistent timer is by using a persistent counter, or sampling the volatile built-in timers of the MCU and save the obtained values in the non-volatile memory. To estimate the off-time, t_{off} in Figure 8, a node needs to determine the incoming power (harvesting rate). The average harvesting rate can be induced from the on-time as follows. The node measures its on-time while harvesting, see t_{on} in Figure 8, and compares it to the time required to drain the energy buffer *without* charging, see t_a in Figure 8. The additional on-time, Δt , is the result of the energy accumulated while executing. If t_{on} and t_a are measured on the same load—thus, they have the same power consumption—then the amount of the energy harvested while the device is on can be calculated as in Line 11, Algorithm 1. And, the average input power can be found as in Line 12 that, in turn, enables the node to estimate its own t_{off} (Line 13). Since calculating the off-time requires constant load, the sensor cannot run arbitrary code during time measurement. Therefore, the sensor needs to sacrifice a certain percentage of its power cycles for measuring time (Line 1-7). Once the on-time and off-time are found the node's power cycle for load a is determined.

Notice that, when the harvested power is very low the accuracy of inferring the charging time from the discharging degrades. However, for the Coalesced Intermittent Sensor this is not a serious problem as the intermittent nodes need to randomize their response to events only in favorable energy conditions.

3.2.3 Alive nodes estimation. To estimate the number of active nodes, a CIS node needs to determine the following information: (i) the total number of nodes in its CIS, which is a typically constant




Figure 9: The average duty cycles of 8 solar-powered and 6 RF-powered intermittent nodes for different ambient energy intensities. In general, the average duty cycle of a node is a good indicator of the average duty cycle of the other CIS nodes.

value that can be loaded to the device memory; (ii) the on-times distribution, which is uniform in our case; and (iii) its own average $\overline{t_{on}}$ and $\overline{t_{off}}$.

Since, we assume that a CIS nodes have the same energy buffers and are in the vicinity of each other (thus, they are exposed to the same energy conditions) then their duty cycles should approach the same value. Figure 9 shows the average duty cycles of the nodes of a solar- and RF-powered CISs. In general, we can conclude that a node’s average duty cycle is a good estimator of other CIS nodes’ duty cycles. Now, a node can estimate the maximum time span, t_{max} , of its CIS, which is the total duration of the nodes’ on-times when they are aligned next to each other, as follows

$$t_{max} = N \times \overline{t_{on}}. \quad (4)$$

Then, from equation 1, the node calculates the CIS availability, $A_v(N)$. As we argued in Section 3.1.1, nodes on-times are uniformly distributed; therefore, the overlapping on-time is also uniformly distributed. As such, a node can calculate the average number of active intermittent nodes, N_{active} , using the following formula,

$$N_{active} = \frac{t_{max}}{t_p \times A_v(N)}. \quad (5)$$

3.2.4 Response randomization factor. Once a node has estimated the number of active neighbors, N_{active} , it can use the following formula to determine the response probability,

$$P_{resp} = \begin{cases} \frac{N_{resp}}{N_{active}} & \text{if } \frac{N_{resp}}{N_{active}} < 1 \\ 1 & \text{otherwise,} \end{cases} \quad (6)$$

where N_{resp} is a system parameter that reflects the desired redundancy factor required by an application.

Table 1 shows the average number of active nodes of an 8-nodes CIS for different light intensities. These measurements provide a sanity check on equation 5. For example, at 1200 lux an individual node of our CIS has a duty cycle of $\approx 62\%$, i.e., it is on average $0.62 t_p$ operating. If we multiply that by the number of nodes (equation 4) we get about $5 t_p$. Figure 4 indicates that a CIS with eight nodes of duty cycles above 50% has near 100% availability. From equation 5, we find that the expected number of clustered nodes is 5, confirmed by the measurements presented in Table 1.

Table 1: Measuring intermittent nodes overlapping of a CIS of 8 intermittent nodes for different light intensities.

Light (lux)	On/Off cycle (%)	N_{active}	Std
300	8	1.01	0.85
500	17	1.63	0.98
800	31	2.88	1.50
1200	62	5.05	1.08

4 PROTOTYPE: COALESCED INTERMITTENT COMMAND RECOGNIZER

The coalesced intermittent command recognizer (CICR) is a prototype of the Coalesced Intermittent Sensor. The CICR consists of eight³ battery-less intermittent nodes. Each node is capable of performing isolated word recognition.

4.1 Hardware

A CICR node consists of three main parts: a microphone, a MCU, and a harvester. MSP430RF5994 [37], an ultra-low-power MCU, is used for data acquisition and processing. This MCU has a 16-bit RISC processor running on 1 MHz, 8 KB of SRAM (volatile), 256 KB of FRAM (non-volatile), and a 12-bit analog to digital converter (ADC). It also features a Low Energy Accelerator (LEA), which offloads the main CPU for specific operations, such as FFT. For recording we use the PMM-3738-VM1010-R piezoelectric MEMS microphone, which features Wake on Sound and ZeroPower listening technologies [38], allowing both the MCU and the microphone to sleep in a low-power mode until a sound wave is detected. The MCU and microphone are powered by a BQ25570 solar power harvester [39] connected to an IXYS SLMD121H04L solar cell [40] and a super-capacitor of 470 μF . For debugging we used the Saleae logic analyzer [41].

4.2 Software

The CICR runs power interrupts immune command recognizer. The recognizer is capable of recognizing isolated-word type of speech. The main parts of the recognizer are illustrated in Figure 10 and explained below:

Data acquisition. The *Wake-on-Sound* feature of the microphone triggers the data acquisition process once the energy level in the sound signal crosses a certain level. The ADC, then, samples the output of the microphone at 8 kHz. This sampling rate is sufficient to cover most of the frequency range of the human voice. The recording length was set to 285 ms, which suffices to get all the acoustic features needed to recognize the words.

Feature Extraction. For word recognition, we adopted the method presented in [42]. Here, we briefly describe the algorithm for convenience. The CICR starts by dividing the signal into frames of 256 samples (≈ 33 milliseconds). Then, it computes a 256-point Fast Fourier Transform for each frame. The resulting feature vectors are normalized (by computing the binary logarithm of each entry of that vector) to reduce detection errors that result from differences in the amplitude of the speech input. This feature vectors are the basis for the word-identification process.

³The number of nodes is bounded by the hardware available to us.




Figure 10: The Coalesced Intermittent Command Recognizer (CICR) features a power-failure-immune word recognizer. First, a word is recorded. Then, its spectral features are extracted. The resulting feature vector is compared against previously-stored words for recognition.

Feature Matching. Feature matching is achieved by computing the squared Euclidean distance between the normalized feature vectors of the recorded word and the feature vectors of the words stored during the training phase (templates, see Table 2). Once the recorded word has been compared to all template words, the template with the smallest distance to the recorded word is considered the correct word. However, if the smallest distance is bigger than a confidence threshold, then the CICR will return “undefined word”.

We have experimented with two feature matching algorithms: the Linear Distance Matching (LDM) and Dynamic Time Warping (DTW) algorithm. While LDM compares the feature vectors of two words successively, DTW looks for the minimum distance between the two vectors. In our implementation, the DTW was about 10 times slower than LDM, whereas the detection accuracy was comparable; therefore, we default our implementation to LDM.

Power-Failure Protection. In order to preserve the progress state and to protect CICR data against randomly timed power failures, we split the recognition program into 19 atomic regions. We ensured that each of these regions requires less energy than what the energy buffer can provide with a single charge. The program state is checkpointed in non-volatile memory on the transition between these regions. This prevents the program from falling back to its starting point (`main()`) after each power failure. Data in non-volatile memory with Write-After-Read dependency is double-buffered to ensure data integrity when the power supply is interrupted.

Code profiling. The entire command recognition software was written in C. The total program consists of 973 lines of code, excluding the FFT function, which is imported from the Texas Instrument DSP library. The memory footprint on the MCU is 20,064 B of FRAM and 1,134 B of SRAM.

The power usage of a node differs according to its activity. When a node is waiting for a voice event, it is in low-power mode. Recording a voice event activates the microphone, ADC and MCU (maximum power consumption). Processing the recorded data requires only the MCU to be on. Table 3 lists a node’s power consumption for each of these states (sleeping, recording, and processing), as measured with a Monsoon power monitor [43].

5 EVALUATION

To evaluate the performance (availability) of the Coalesced Intermittent Sensor, we conducted several experiments in different energy conditions and with different event arrivals patterns.

Table 2: Test set

on	off
stop	clear
load	go
pause	resume
edit	cancel

Table 3: Power usage

State	Current (μA)	Time (ms)
<i>Sleeping</i>	64 ± 20	—
<i>Recording</i>	423 ± 20	285
<i>Processing</i>	282 ± 20	600

5.1 Availability


Irrespective of the energy source (RF or light) we showed in Figure 5 that the power cycles of a CIS’s nodes are different, which leads to a uniform distribution of their on-times, as we argued in Section 3.1.1. We captured the expected joined availability of these nodes in Model 1. Here, we validate model by comparing the modeled availability of a CIS against data captured with different hardware (solar- and RF-powered nodes) in different scenarios. Figure 11 shows the availability of three CISs when they are powered by sunlight, artificial light, and RF and for a different number of intermittent nodes. The results clearly confirm our expectation: when the power cycles are slightly different, the on-times are uniformly distributed. The results also validate our model; the dashed lines represent the modeled availability when the nodes duty cycle is 15%.

5.1.1 Availability on a Fine Scale. Since the nodes’ on-times are in a constant shift relative to each other (Section 3.1.1), the collective availability of the CIS fluctuates when it is observed in a short time window. Figure 12 captures CIS availability on a time window of 5 seconds for three different ambient energy conditions. In these experiments, the average power cycles of the CIS nodes are (3,18), (3,9,12,3), and (4,3,11,5) seconds when ambient light intensity are 500, 800, and 900 lux, respectively. If we focus on the line graphs associated with 500 and 800 lux and compare the system availability within the interval [20, 50] seconds and the rest, we can observe that the CIS gradually alternates between low and high collective availability; nodes’ on-times gradually transition from maximum to minimum separation and vice versa (Section 3.1.1). Notice that, when ambient light intensity was 800 lux the CIS collective availability transitioned from low to high, while this pattern happened to be reversed when light intensity was 500 lux. For the 900 lux the 8-node CIS achieved near-continuous 100% availability.


5.2 Sensing

5.2.1 Experiment setup. After validating our observation on different energy sources, we designed a testbed with controllable light intensity for clarity and reproducibility. To this end, we blocked uncontrollable light sources with a box of $60 \times 40 \times 40$ cm. On the box ceiling, we attached a light strip of 2.5 m with 150 LEDs that can produce 15 different light intensities. On the bottom a Coalesced Intermittent Command Recognizer of 8 intermittent nodes is placed (see Section 4.1 for hardware description).


The events in our experiments are spoken words (Table 2). Short events (see events classification in Section 3.1.2) are represented with individual words, while burst events are represented with



(a) The CIS is powered by uncontrollable light sources—artificial light (night) and sunlight (day).



(b) The CIS is powered by a controllable array of LEDs.



(c) The CIS is powered by an RF reader [44] located 30-70 cm away from the RF tags (WISPs [14]).

Figure 11: The Coalesced Intermittent Sensor’s availability for differed energy sources and number of nodes. The modeled availability (dashed red lines with nodes’ duty cycles of 15%) approximates the measured availability with high accuracy.




Figure 12: CIS availability smoothed with a 5-second time window.

phrases of a few words. We recorded different patterns of inter-event and inter-burst arriving time. We used a Bluetooth speaker [45] to replay a certain recording.

5.2.2 *Events detection rate.* Here we experiment with the behavior of a CIS when events arrive individually or in bursts *without* enabling randomized response in favorable energy conditions.

Individual events. Figure 13 shows the percentages of capturing duplicate and unique events when light intensity varies from 300 lux to 1400 lux and the inter-event arrival time ranges from 1 second to 6 seconds. For each experimental trial 20 words were played, resulting in a total of 240 playbacks.

Figure 13 shows a positive correlation between light intensity and the number of detected events. In particular, the number of duplicate detections rises dramatically when light intensity increases, *demonstrating the overpowering problem* (Section 3.2.1). Moreover, increasing the inter-event arrival time also surges the number of duplicated events. The reason for this is that when the time between events increases, the intermittent nodes get the chance to sleep longer in low-power mode, consuming less energy. Thus, nodes’ on-times expand, reducing their inherent randomization, which leads them to be in *hibernating power state* (Section 3.2.1).

Bursty events. Figure 14a shows the capturing behavior of a CIS when the events arrive in bursts. A burst of four events with one second between the individual events was fired every 20 seconds. Each burst was repeated 10 times and under four different light





Figure 13: Duplicate and unique events captured by a coalesced intermittent command recognizer of eight solar-powered nodes. In general, the number of captured events increases in two case: when light intensity rises and when inter-event arrival time increases. Red numbers indicate events arrival interval in seconds.



(a) When capturing a burst of 4 events without randomizing the response, the majority of the nodes reacts to the first event capturing rates. It also reduces in the burst and powers down shortly after, missing the rest of the burst. Red numbers indicate events in the burst. (b) Response randomization enables a CIS to capture the entire burst of events with high capturing rates. It also reduces the number of duplicated events. Red numbers indicate events in a burst.

Figure 14: Capturing bursts of events when CIS nodes’ response is immediate (a) or randomized (b).

intensities. The nodes sleep in a low-power mode when they finish processing an event, waiting for the next one.

In general, we observe that in favorable energy conditions (above 500 lux) intermittent nodes react to the first event of a burst and power down shortly after, missing the rest of the burst. These results confirm our argument about the side effect of the *hibernating power state* of a CIS (Section 3.2.1). These results also demonstrate

that the hibernating power problem happens on a wide range of power intensities, showing its significance. Next, we will show how randomizing the response can mitigate the problems generated when ambient energy exceeds the design point.

5.2.3 Events detection rate with randomization. Here, we examine the effect of enabling artificial randomization on the CIS’s response.

Individual events. Table 4 compares the number of detected events when the CICR’s response is randomized and when not. When randomization is enabled, nodes respond to events with a probability of 65% for the scenario of (800 lux, 6 seconds) and (1400 lux, 4 seconds), and for the highest energy level and the longest inter-event arrival time the responding probability was set to 30%.

Table 4 shows that randomizing the response reduces duplicated events by an average of $\approx 50\%$, while only marginally lowers the number of the uniquely detected events (7% on average).

Bursty events. To enable a CIS to capture events that arrive in bursts, the response probability for each event in a burst should be different. The CIS should respond with a minimum probability to the first event in a burst and gradually increase the response probability for the subsequent events in the burst (we assume that between the bursts the CIS resumes to its expected collective availability). This gradual increment to the responding probability is motivated by the observation that when a node captures an event it becomes unavailable for the subsequent ones in the burst. In this experiment, the nodes reacted with a probability of 40%, 50%, 70% and, 100% on the first, second, third and fourth event, respectively. Since the event distribution is known these probabilities were fixed during the development stage.

Figure 14b shows how randomizing the CIS response spreads the nodes’ awake times—as compared to Figure 14a—and enables the CIS to capture the entire burst with a high probability, i.e., above 85%. However, capturing the signal of a word does not ensure correct word recognition as the recognition algorithm is not robust to ambient noise. As a consequence, our word recognizer achieved an acceptable recognition accuracy of 76.6% on average.


5.2.4 Different energy harvesting rates. A valid concern is that neighboring nodes may harvest energy at different rates (e.g., being in direct sunlight vs. in the shade) leading to a different view on the effective power cycle, in turn, yielding different response probabilities. Figure 15a shows (in green) that the length of a power cycle can differ by as much as a factor of two between low and high intensity lighting. Note that this data is from nodes that employ the wake-on-event strategy and—in good lighting conditions at the right— spend most of their on time in low-power mode (in blue), allowing them to charge while sleeping. That also explains their very short off times (in orange).

To see what the effect of the different energy harvesting rates (and corresponding response probabilities) is we ran a simulation experiment with 10 nodes where we varied the number of nodes in the shade (400 lux) and sunlight (1200 lux) running at 16% and 50% duty cycle, respectively. The response probabilities were set to achieve a 2.5 factor of redundancy following Equation 6. Figure 15b shows that the CIS availability (i.e., the number of uniquely captured events in green) is not dramatically affected by the distribution of nodes over sun/shade. This is in part due to the redundancy (cf. the

Table 4: The number of *unique/total* detected events. Randomizing the response reduces the number of duplicated events by 50% while losing only 7% of the unique events.

	(800 lux, 6 seconds)	(1400 lux, 4 seconds)	(1400 lux, 6 seconds)
<i>randomization*</i>	205/432	236/675	223/493
<i>no randomization</i>	240/831	240/938	240/1802

* A node’s response probability is 65% in the first two scenarios, (800,6) and (1400,4), and 30% in the third.



(a) In wake-on-event style of operation, nodes go into sleep mode after charging. The low power-consumption rate of this mode makes the on-time of a node an important factor for its duty cycle and power-cycle length. **(b) Simulated results showing how the number of captured events is affected when shadow is covering a CIS. Detected events are #times events trigger nodes. Captured events are #times nodes processed events.**

Figure 15: The effect of non-uniform energy distribution.

number of captured events in orange), but mainly caused by the adequate response (sleep on it) of the “sunny” nodes.

6 CONCLUSION AND FUTURE WORK

This paper addresses the availability problem of intermittent sensors that fail to capture (and process) events while charging their energy buffer. As the power to drive a node is much higher than what can be harvested from ambient sources, the chance of capturing an event can be as low as just 8% (light) and 4% (RF) (cf. the duty cycles reported in Figure 9). To address this problem of missing most events we presented the Coalesced Intermittent Sensor (CIS), which is the abstraction of a group of intermittently-powered sensors, whose collective duty cycle (on-time) can approach the desired 100% availability. The inherent differences in the powering subsystem of intermittent sensors result in (slight) differences in the sensor nodes’ power cycles causing the nodes’ on-times to be uniformly distributed. This implies that simply selecting the right number of nodes is all that is required. To this end we have modeled the (effective) availability of a CIS and validated its accuracy against data collected on real hardware.

Experimentation with an 8-node prototype CIS, a basic voice-control application recognizing up to 4-word commands, showed that the inherent randomization in the power cycles can easily be disrupted. In case the ambient power exceeds the (worst-case) design point and nodes employ an efficient wake-on-event sleep mode, all nodes wake-up on the same (rare) event. If the energy buffer is small then they all enter the charging state at approximately the same time (unwanted synchronization) and subsequent events (words) will be missed (compromising availability). To counter this

unwanted behavior, we proposed to use a probabilistic approach in which the number of active neighbors is determined and nodes respond proportionally to an event. This approach was shown to be effective for our prototype, capturing burst events with above 85% detection accuracy.

In the future we plan to dive deeper in studying partially captured data. Intermittent sensors may partially capture events. Classical recognition algorithms face difficulties dealing with partially captured data. Therefore, we want to investigate *how much machine learning algorithms can improve the sensing quality of intermittent sensing?*

7 ACKNOWLEDGEMENTS

We like to thank Dr. Stephan Wong for his contributions in the early stages of this work, and the anonymous referees for their comments helping us to polish the paper to perfection.

REFERENCES

- [1] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith, "Ambient backscatter: wireless communication out of thin air," in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4. ACM, 2013, pp. 39–50.
- [2] B. Lucia, V. Balaji, A. Colin, K. Maeng, and E. Ruppel, "Intermittent computing: Challenges and opportunities," in *LIPICs-Leibniz International Proceedings in Informatics*, vol. 71. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [3] B. Ransford, J. Sorber, and K. Fu, "Mementos: system support for long-running computation on rfid-scale devices," in *ACM SIGARCH Computer Architecture News*, vol. 39, no. 1. ACM, 2011, pp. 159–170.
- [4] B. Ransford and B. Lucia, "Nonvolatile memory is a broken time machine," in *Proc. MSPC*. Edinburgh, United Kingdom: ACM, Jun. 13, 2014.
- [5] A. Colin and B. Lucia, "Chain: tasks and channels for reliable intermittent programs," *ACM SIGPLAN Notices*, vol. 51, no. 10, pp. 514–530, 2016.
- [6] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi, D. Brunelli, and L. Benini, "Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems," *IEEE Embedded Systems Letters*, vol. 7, no. 1, pp. 15–18, 2014.
- [7] A. Rodriguez Arreola, D. Balsamo, G. V. Merrett, and A. S. Weddell, "Restop: Retaining external peripheral state in intermittently-powered sensor systems," *Sensors*, vol. 18, no. 1, p. 172, 2018.
- [8] J. Hester, K. Storer, and J. Sorber, "Timely execution on intermittently powered batteryless sensors," in *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. ACM, 2017, p. 17.
- [9] K. S. Yildirim, A. Y. Majid, D. Patoukas, K. Schaper, P. Pawelczak, and J. Hester, "Ink: Reactive kernel for tiny batteryless sensors," in *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2018, pp. 41–53.
- [10] A. Y. Majid, D. Carlo, K. Maeng, A. Ytiddtm, Colin, K. Sinan, P. Pawetczak, and B. Lucia, "Dynamic task-based intermittent execution for energy-harvesting devices," *ACM Transactions on Sensor Network*, 2019.
- [11] N. Jackson, J. Adkins, and P. Dutta, "Reconsidering batteries in energy harvesting sensing," in *Proceedings of the 6th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*, 2018, pp. 14–18.
- [12] N. Jackson, J. Adkins, and P. Dutta, "Capacity over capacitance for reliable energy harvesting sensors," in *Proc. IPSN*. Montreal, QC, Canada: ACM/IEEE, Apr. 16–18, 2019.
- [13] V. Rao, "Ambient-energy powered multi-hop internet of things," 2017.
- [14] J. R. Smith, A. P. Sample, P. S. Powlledge, S. Roy, and A. Mamishev, "A wirelessly-powered platform for sensing and computation," in *International Conference on Ubiquitous Computing*. Springer, 2006, pp. 495–506.
- [15] Y. Zhao, J. R. Smith, and A. Sample, "Nfc-wisp: A sensing and computationally enhanced near-field rfid platform," in *RFID (RFID), 2015 IEEE International Conference on*. IEEE, 2015, pp. 174–181.
- [16] H. Zhang, J. Gummesson, B. Ransford, and K. Fu, "Moo: A batteryless computational rfid and sensing platform," *University of Massachusetts Computer Science Technical Report UM-CS-2011-020*, 2011.
- [17] S. Naderiparizi, A. N. Parks, Z. Kapetanovic, B. Ransford, and J. R. Smith, "Wisp-cam: A battery-free rfid camera," in *2015 IEEE International Conference on RFID (RFID)*. IEEE, 2015, pp. 166–173.
- [18] V. Talla, B. Kellogg, S. Gollakota, and J. R. Smith, "Battery-free cellphone," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 2, p. 25, 2017.
- [19] A. Y. Majid, M. Jansen, G. O. Delgado, K. S. Ytiddtm, and P. Pawetczak, "Multi-hop backscatter tag-to-tag networks," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 721–729.
- [20] J. Hester and J. Sorber, "Flicker: Rapid prototyping for the batteryless internet-of-things," in *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. ACM, 2017, p. 19.
- [21] A. Colin, G. Harvey, B. Lucia, and A. P. Sample, "An energy-interference-free hardware-software debugger for intermittent energy-harvesting systems," *ACM SIGPLAN Notices*, vol. 51, no. 4, pp. 577–589, 2016.
- [22] A. Colin, E. Ruppel, and B. Lucia, "A reconfigurable energy storage architecture for energy-harvesting devices," in *ACM SIGPLAN Notices*, vol. 53, no. 2. ACM, 2018, pp. 767–781.
- [23] H. Aantjes, A. Y. Majid, P. Pawelczak, J. Tan, A. Parks, and J. R. Smith, "Fast downstream to many (computational) rfids," in *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*. IEEE, 2017, pp. 1–9.
- [24] J. Van Der Woude and M. Hicks, "Intermittent computation without hardware or programmer intervention," in *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, 2016, pp. 17–32.
- [25] B. Lucia and B. Ransford, "A simpler, safer programming and execution model for intermittent systems," *ACM SIGPLAN Notices*, vol. 50, no. 6, pp. 575–585, 2015.
- [26] N. A. Bhatti and L. Mottola, "Harvos: Efficient code instrumentation for transiently-powered embedded sensing," in *Information Processing in Sensor Networks (IPSN), 2017 16th ACM/IEEE International Conference on*. IEEE, 2017, pp. 209–220.
- [27] G. Gobieski, B. Lucia, and N. Beckmann, "Intelligence beyond the edge: Inference on intermittent embedded systems," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2019, pp. 199–213.
- [28] D. Patoukas, K. S. Yildirim, A. Y. Majid, J. Hester, and P. Pawelczak, "Feasibility of multi-tenancy on intermittent power," in *Proceedings of the 6th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*, 2018, pp. 26–31.
- [29] D. Balsamo, A. S. Weddell, A. Das, A. R. Arreola, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini, "Hibernus++: a self-calibrating and adaptive system for transiently-powered embedded devices," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 35, no. 12, pp. 1968–1980, 2016.
- [30] J. Degeysys, I. Rose, A. Patel, and R. Nagpal, "Desync: self-organizing desynchronization and tdma on wireless sensor networks," in *Proceedings of the 6th international conference on Information processing in sensor networks*. ACM, 2007, pp. 11–20.
- [31] A. Giusti, A. L. Murphy, and G. P. Picco, "Decentralized scattering of wake-up times in wireless sensor networks," in *European Conference on Wireless Sensor Networks*. Springer, 2007, pp. 245–260.
- [32] C. Zheng and D. C. Sicker, "A survey on biologically inspired algorithms for computer networking," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1160–1191, 2013.
- [33] T. K. Vintsyuk, "Speech discrimination by dynamic programming," *Cybernetics and Systems Analysis*, vol. 4, no. 1, pp. 52–57, 1968.
- [34] F. Jelinek, *Statistical methods for speech recognition*. MIT press, 1997.
- [35] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath et al., "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [36] S. K. Gaikwad, B. W. Gawali, and P. Yannawar, "A review on speech recognition technique," *International Journal of Computer Applications*, vol. 10, no. 3, pp. 16–24, 2010.
- [37] Texas Instruments. (2019) MSP430FR5994 16 MHz ultra-low-power microcontroller product page. [Online]. Available: <http://www.ti.com/product/MSP430FR5994>
- [38] pui audio, vesper. (2019) Pmm-3738-vm1010-r: A zeropower listening piezoelectric mems microphone. [Online]. Available: <http://www.puiaudio.com/pdf/PMM-3738-VM1010-R.pdf>
- [39] Texas Instruments. (2018) Ultra low power management IC, boost charger nanoperovided buck converter evaluation module. [Online]. Available: <http://www.ti.com/tool/BQ25570EVM-206>
- [40] IXYS Corporation. (2018) IXOLAR™ high efficiency SLMD121H04L solar module. [Online]. Available: http://ixapps.ixys.com/DataSheet/SLMD121H04L_Nov16.pdf
- [41] Saleae. (2017) Logic 16 analyzer. <http://www.saleae.com>. Last accessed: Jul. 28, 2017.
- [42] G. Hopper and R. Adhami, "An fft-based speech recognition system," *Journal of the Franklin Institute*, vol. 329, no. 3, pp. 555–562, 1992.
- [43] Monsoon Solutions Inc., "High voltage power monitor," 2018. [Online]. Available: <https://www.monsoon.com/hvpm-product-documentation>
- [44] Impinj Inc., "Impinj speedway R420 RFID reader product information," <https://www.impinj.com/platform/connectivity/speedway-r420/>, 2018, last accessed: Apr. 8, 2019.
- [45] JBL. (2019) JBL Go+ bluetooth speaker. [Online]. Available: <https://www.jbl.com/>