

Opportunistic Communication with Latency Guarantees for Intermittently-Powered Devices

Kacper Wardega*, Wenchao Li*, Hyoseung Kim[†], Yawen Wu[‡], Zhenge Jia[‡] and Jingtong Hu[‡]

*Department of Electrical and Computer Engineering, Boston University, USA

[†]Department of Electrical and Computer Engineering, University of California, Riverside, USA

[‡]Department of Electrical and Computer Engineering, University of Pittsburgh, USA

Abstract—Energy-harvesting wireless sensor nodes have found widespread adoption due to their low cost and small form factor. However, uncertainty in the available power supply introduces significant challenges in engineering communications between intermittently-powered nodes. We propose a constraint-based model for energy harvests that together with a hardware model can be used to enable consistent, opportunistic communication with worst-case latency guarantees. We show that greedy approaches that attempt communication whenever energy is available lead to prolonged latencies in real-world environments. Our approach offers bounded worst-case latency while providing a performance improvement over a conservative, offline approach planned around the worst-case energy harvest.

I. INTRODUCTION

Advances in energy-harvesting technologies combined with the maturation of ultra-low-power computing systems have led to the development of batteryless, intermittently-powered devices. These devices can operate only when sufficient energy has been harvested and banked. The removal of tethered power or battery coupled with an energy harvesting source gives rise to energy-harvesting wireless sensor networks (EH-WSNs). EH-WSNs enable deeper penetration of remote sensing and computing technologies in applications such as infrastructure monitoring [1], wearable devices [2], and industrial applications [3]. On the flip side, the unreliable power sources and the intermittent nature of computing give rise to many new challenges, such as the need for proper checkpointing to ensure data consistency, and how to synchronize a collection of such devices in a distributed setting [4], [5].

While there are plenty of works addressing the intermittent computation problem, the study of communication between intermittently powered devices still has open problems preventing their use in dependable systems. Prior work often leverages an Access Point, a strong centralized radio beacon with stable power supply, as a means to manage the functioning of the EH-WSNs. In the case of RF-powered EH-WSNs, the Access Point can also co-schedule energy delivery and data retrieval [6]. The authors of [7] propose a data delivery scheme where data must make several hops between intermittently-powered nodes, but it is still assumed that regular energy deliveries are coordinated by the Access Point. Additionally, non-RF EH-WSNs such as those powered by solar, wind, kinetic, or thermal sources, present unique challenges due to respective operating environments and considerable effort must be made to model the environment and determine the

approach that should be taken with respect to scheduling both computation and communication. The difficulty of modeling new environments has led to the proliferation of stochastic-modeling and model-free approaches [8]–[11].

As radio usage is the dominant consumer of energy in a typical WSN, energy overhead of communication coupled with uncertainty in energy availability impacts a node's ability to reliably communicate. Approaches that aim to limit radio usage either through synchronous duty cycling or on-demand wake up schemes via wake up radio hardware are therefore essential for reliable communications between energy-harvesting nodes [12], [13]. Hardware considerations and protocol choices that reduce energy overhead of communication aside, no guarantees on the worst-case performance can be obtained without taming the uncertainties in energy availability.

In this work, we address this gap by detailing how to check the worst-case behavior of general hardware and protocol choices within an energy-harvesting context. Our approach is to verify the worst-case communication behavior that a node can support given an energy-harvesting context offline, and then check online whether additional opportunistic communication can be performed without violating the offline guarantees. Our contributions are as follows.

- We propose a constraint-based approach to symbolically encode the types of energy-harvesting traces that can be generated by a given environment.
- We show that locally greedy approaches to communication in EH-WSNs are prone to unbounded latency even for just two nodes.
- We propose an offline worst-case approach that guarantees bounded communication latency between two nodes while still allowing for locally opportunistic online communication attempts in the event of better-than-worst-case energy harvests.
- Based on real-world energy harvesting traces, we observe that our approach to opportunistic communication that respects worst-case harvesting scenarios not only maintains the average performance of the locally greedy approach to opportunistic communication but also maintains the worst-case latency guarantees of a conservative duty-cycling approach, leading to a 44% decrease in worst-observed communication latency on average.

II. ENERGY-HARVESTING SENSOR NODE MODEL

Our primary goal in this work is to obtain worst-case guarantees on communication performance between energy-harvesting nodes and to determine when and how nodes can attempt to communicate opportunistically without violating the guarantees. Our performance metric is the *latency* between successful communications, defined (with more detail in II-C) as the number of elapsed rounds in our synchronous communication scheme between two consecutive successful communications. We adopt a model-driven approach to tame the uncertainties in the energy availability. The primary components of our modeling framework are the device model, the communication model, and the harvesting model. Throughout this work we make use of the standard notation for sets; denote by $\mathbb{N}_k$ the first k natural numbers, by $\mathbb{R}$ the reals and by $\mathbb{R}^+$ the positive reals. Denote by $\text{lcm } S$, $S \subseteq \mathbb{N}_k$, the least common multiple of the elements in S .

A. Device Model

A device $d = (A, \ell, \phi, \Gamma)$ models a physical compute node in the environment and consists of four elements:

- 1) A : a set of allowable actions
- 2) ℓ : the device's energy storage limit
- 3) ϕ : a function mapping actions to energy costs
- 4) Γ : an energy-harvester model describing the types of energy traces the environment may provide

In each round a device selects an action from its action set A , which incurs an *energy cost* determined by $\phi : A \rightarrow \mathbb{R}^+$. The energy-harvesting traces that d may experience in its environment are described symbolically by the harvesting model Γ , which we elaborate on in II-B. The limit on the amount of energy that d can store is given by ℓ . The actions in A can be tailored to suit the type of communication protocol and hardware being analyzed – it could be as simple as $A = \{\text{sl}, \text{tx}, \text{rx}\}$ which are equivalent to “radio off (sleep)”, “radio transmit,” and “radio receive” respectively. Device models for hardware with additional modes of communication (e.g. main radio, wake up radio) or further options specifying data rate, radio output power, and packet length, can of course include additional actions corresponding to the other modes and options. The only mathematically important action that A must contain is sl , or sleep. The sleep action has the important property that it should always be valid for the device to sleep regardless how much energy the device has. We discuss the hardware details of the representative energy-harvesting platform in more detail in IV-A. Note that while we consider devices with homogeneous A and ϕ , this is only to simplify the presentation but not a requirement.

B. Harvesting Model

Typically, significant effort must be made when designing harvesting-based systems to model the harvesting environment, i.e. reasoning about the availability of energy to a given compute node. To avoid the effort of modeling the environment and to provide a general approach that can be leveraged across environments, we propose a trace-based

approach for modeling power availability. As we will show later, this approach allows us to obtain guarantees about the system without significant modeling effort. At each device d , we assume as given a set of harvesting constraints Γ , where each element of Γ is of the form $[m, k]$ read as “ d is guaranteed to receive at least m units of power every k rounds.” Such specifications can be readily mined from real power traces. Note that a *set* of minimum harvesting guarantees, as opposed to a singleton, allows for refining the worst-case harvest provided by the environment at decreasing intervals, e.g. even if the environment only guarantees that d will harvest at least 10 units of power in every minute-long window, the environment may also guarantee 100 units of power every five minute-long window (more power on average). For $[m, k] \in \Gamma$, as m increases, m/k approaches the average worst-case harvest per round.

While we assume that a suitable Γ is known, we provide here a simple way to obtain such a Γ from energy harvesting traces. Suppose you have sampled the energy harvests $\hat{h}(t)$ for a device at discrete time instants $t \in \mathbb{N}_T$ with a constant sampling interval of $1/f$ (preferably, $1/f$ is a multiple of the round duration). Then compute from the trace $\forall k \in \mathbb{N}_T$,

$$m_k = \min_{1 \leq t \leq T-k+1} \sum_{s=t}^{t+k-1} \hat{h}(s) \quad (1)$$

Giving us a natural choice for Γ ,


$$\Gamma = \{[m_a, a/f] : a \in \mathbb{N}_T \wedge (\forall b \in \mathbb{N}_T)(m_b/b > m_a/a \Rightarrow b > a)\} \quad (2)$$

Given an energy-harvesting trace $h : \mathbb{N} \rightarrow \mathbb{R}^+$ and harvesting constraints Γ , we say that h models, or satisfies, the constraints Γ iff h satisfies $(\forall t > 0, [m, k] \in \Gamma) \sum_{s=t}^{t+k-1} h(s) \geq m$, or $h \models \Gamma$ for short.

C. Communication Model

The purpose of the communication model is to define when communication between two nodes is successful. Importantly, we assume that devices are time-synchronized. Time is partitioned into equal-duration rounds, and the devices agree on which round corresponds to $t = 0$. In III we will obtain guarantees on the worst-case communication behavior for each device. Therefore once the devices have successfully established a first communication, the synchronicity assumption can be checked against the computed bounds on worst-case latency and bounded clock drift given that the devices will be able to communicate periodically to exchange timing information.

In each round, the devices pick actions to execute from their respective action sets. Given that two devices d_1, d_2 with action sets A_1, A_2 respectively are within range in the physical space, consider the joint action $(a_1, a_2) \in A_1 \times A_2$. Which joint actions are successful again depends on the hardware and protocol – in the simplest case where $A_1 = \{\text{sl}, \text{tx}\}$ and $A_2 = \{\text{sl}, \text{rx}\}$, we would say that (a_1, a_2) is a successful joint action if $(a_1, a_2) = (\text{tx}, \text{rx})$, and unsuccessful otherwise. A visualization of the round-based scheme is shown in Fig. 1. For more complicated settings with additional radio



(b) greedy policy that is consistent w.r.t. a future communication at $t = 5$
 Fig. 1. An illustration showing the energy levels and communication attempts of sender and receiver nodes. At top, the sender node uses harvested energy greedily whenever it is available. At bottom, we see another possible run where the sender chooses not to communicate at $t = 1$, but rather conserve energy in order to communicate at a future time $t = 5$ when it knows that the receiver will have gathered enough energy to communicate. In this work, we aim to bound the worst-case latency observed by communicating nodes by (1) computing a valid offline policy that identifies times when successful communication can occur even under worst-case energy harvests and (2) compute if and when nodes can communicate opportunistically in a consistent manner that respects the actions planned by the offline policy.

modes and options, the joint action would depend on more than simply neither device sleeping; success would also be contingent on the devices agreeing on the correct radio mode and options. We define the latency at time t as the number of rounds elapsed since the last successful communication.

Given the device and communication models, we wish to guarantee the worst-case latency of communication between nodes. Coordination between distributed energy-harvesting nodes presents many challenges, however. At any instant in time, each node cannot be certain how much energy will be available to its neighbors. *Each node must carefully plan when it should communicate.* Otherwise, it risks attempting to communicate at a time when neighbors cannot spare the energy to listen, and wasting its own energy on the attempt.

III. ANALYSIS OF WORST-CASE LATENCY

In this section, we apply the modeling framework proposed in II to obtain a worst-case upper bound on communication latency for intermittently-powered communication nodes.

A. The Offline Scheduling Problem

We start by considering the simplest case of two devices attempting to perform periodic communication under intermittent power supply constraints. In other words, we have two synchronized devices d_1 and d_2 and we wish to obtain an upper bound on the communication latency between the devices. The first step is to obtain a worst-case lower bound on the amount of energy that has been harvested by a device by round t . Let $d = (A, \mathcal{E}, \phi, \Gamma)$ be any device, $h : \mathbb{N} \rightarrow \mathbb{R}^+$ any energy harvesting trace. A policy π for d takes the current state of the system x : a finite history of length $\mathcal{K}$ of the harvesting trace, the current energy level and the time, and returns an action in A . Then the state of the system and the energy on the device is given by the system of difference equations

$$\begin{cases} x_{d,h,\pi}(t) = (h(\max(0, t - \mathcal{K})), \dots, h(t-1), \\ E_{d,h,\pi}(t-1), t) \\ E_{d,h,\pi}(t) = \max\{0, \min\{\mathcal{E}, \\ E_{d,h,\pi}(t-1) + h(t) - \phi(\pi(x(t-1)))\}\} \end{cases} \quad (3)$$

for $t > 0$ and boundary condition $E_{d,h,\pi}(0) = E_0$. It is often true that harvesting efficiency and action costs are functions also of the current energy level, which we ignore here for simplicity of presentation as our approach readily generalizes to the case where h or ϕ have non-linear scaling dependent on $E_{d,h,\pi}$. Since our only offline knowledge about h is that it satisfies the harvesting specification, i.e. $h \models \Gamma$, our offline lower bound for the amount of energy harvested by d at time t is

$$E_{d,\pi}(t) = \min_{h \models \Gamma} E_{d,h,\pi}(t) \quad (4)$$

Definition III.1 (Valid policy). A policy $\pi : (\mathbb{R}^+)^{\mathcal{K}+1} \times \mathbb{N} \rightarrow A$ is considered valid for a device $d = (A, \mathcal{E}, \phi, \Gamma)$ iff $(\forall h \models \Gamma, t \in \mathbb{N})(\pi(x_{d,h,\pi}(t)) \neq \text{sl} \Rightarrow \phi(\pi(x_{d,h,\pi}(t))) \leq E_{d,h,\pi}(t))$, and invalid otherwise. In other words, a valid policy for d is one where d does not attempt anything other than sleep unless it has the energy for a non-sleep action.

Definition III.2 (Offline policy). A policy $\pi : (\mathbb{R}^+)^{\mathcal{K}+1} \times \mathbb{N} \rightarrow A$ is considered to be an offline policy for a device $d = (A, \mathcal{E}, \phi, \Gamma)$ iff $(\forall h_1, h_2 \models \Gamma, t \in \mathbb{N})(\pi(x_{d,h_1,\pi}(t)) = \pi(x_{d,h_2,\pi}(t)))$, and online otherwise. In other words, an offline policy for d is one that makes the same decisions regardless of which harvesting signal it encounters.

The important first question we'd like to answer is whether a given offline policy is also valid – two devices executing valid offline policies have a guaranteed worst-case communication behavior that is known offline. Consider the following template for a periodic policy π_α with period p that repeatedly executes the finite sequence of actions $\alpha = \alpha_0, \dots, \alpha_{p-1}$.

$$\pi_\alpha(h_1, \dots, h_{\mathcal{K}}, e, t) = \alpha_s \text{ s.t. } t \equiv s \pmod{p} \quad (5)$$

Clearly, this is an offline policy since it is agnostic to the harvesting history $h_1, \dots, h_{\mathcal{K}}$ and energy level e . We can check the validity of a policy by checking that $\neg \exists h, h \models \Gamma$ s.t. the validity property is violated. Although difficult to check for general π , provided that π is periodic of the form π_α , validity can be checked by proving non-existence of a falsifying h through the use off-the-shelf constraint solvers – our Satisfiability Modulo Theories (SMT) encoding is given in III-A1. The special case $\mathcal{E} = \infty$ can be solved efficiently without the use of constraint solvers; we discuss the approach in III-A2. Regardless how it is determined that a policy is offline and valid, we state the following result.

Theorem III.1. Let d_1 and d_2 be two devices in communication range. If π_α is a valid periodic offline policy for d_1 and π_β is a valid periodic offline policy for d_2 s.t. $|\alpha| = |\beta| = p$, then the worst-case latency for d_1, d_2 is given by the maximum number of rounds between successful joint actions in the sequence

$(\alpha_0, \beta_0), \dots, (\alpha_{p-1}, \beta_{p-1}), (\alpha_0, \beta_0), \dots, (\alpha_{p-1}, \beta_{p-1})$.

This result follows directly from the periodicity of α and β , Def. III.1, and the definition of latency in II-C.

1) *Finite energy capacity: $\theta < \infty$* : For a device $d = (A, \theta, \phi, \Gamma)$, $\theta \in \mathbb{R}^+$, and a periodic offline policy π , we can check validity of π via an encoding to Linear Real Arithmetic (or Linear Integer Arithmetic if all of the actions costs are integers) and solve using off-the-shelf SMT constraint solvers. As θ is finite and π is periodic, $E_{d,\pi}$ is also eventually periodic and therefore we only need to solve for a finite time horizon. As π is offline, the sequence of costs is a fixed sequence of values $\phi(t) \in \mathbb{R}^+$ independent of h (since the sequence of actions is a fixed $\pi(t) \in A$). We allow $E_{d,h,\pi}(t), h(t)$ to take symbolic values in $\mathbb{R}^+$. The energy on the device is constrained by the update equation:

$$\bigwedge_{t>0} E_{d,h,\pi}(t) = \max\{0, \min\{\theta, E_{d,h,\pi}(t-1) + h(t) - \phi(t)\}\} \quad (6)$$

while the harvesting trace is constrained by $h \models \Gamma$, i.e.

$$\bigwedge_{t>0, [m,k] \in \Gamma} \sum_{s=t}^{t+k-1} h(s) \geq m \quad (7)$$

Finally, we specify that the policy should be invalid, i.e.

$$\bigvee_{t>0} \pi(t) \neq s1 \wedge E_{d,h,\pi}(t) < \phi(t) \quad (8)$$

If the solver returns that the constraints are unsatisfiable, then we have confirmed that $\neg \exists h, h \models \Gamma$ that invalidates π . If however the solver returns a satisfying assignment for h , then we have a counterexample where π is invalidated by a harvesting trace h .

2) *Special case: $\theta = \infty$* : In this case, we can leverage the fact that harvested energy is never lost due to arriving at a time when the energy storage is full. Under $[m, k]$ -style constraints, we cannot be *certain* that the m units of power have been harvested until k rounds have been elapsed. Operating under this assumption, the worst thing that can happen is that all of the energy arrives at the beginning of the k window, and we lose as much as possible during our inaction. Putting this in terms of a lower bound on $E_{d,\pi}$, consider the harvesting trace h_{late} where the minimum amount of energy arrives as late as possible, defined recursively as

$$h_{\text{late}}(t) = \max\{m - \sum_{s=t-k+1}^{t-1} h(s) : [m, k] \in \Gamma \wedge t \geq k-1\} \cup \{0\} \quad (9)$$

with initial conditions $\forall t \leq 0, h(t) = 0$. It can be shown that h_{late} is eventually periodic with period $\text{lcm}\{k : [m, k] \in \Gamma\}$ and furthermore that as long as π is an offline policy,

$$\begin{aligned} E'_{d,h_{\text{late}},\pi}(t) &= E'_{d,h_{\text{late}},\pi}(t-1) \\ &\quad + h_{\text{late}}(t) - \phi(\pi(x_{d,h_{\text{late}},\pi}(t))) \\ &\leq E_{d,\pi}(t) \end{aligned} \quad (10)$$

This leads to the following result for $\theta = \infty$,

Theorem III.2. *Let $d = (A, \infty, \phi, \Gamma)$ be a device with infinite energy storage capacity and let π be an offline policy for d . If $\forall t$*

$$\pi(x_{d,h_{\text{late}},\pi}(t)) \neq s1 \Rightarrow \phi(\pi(x_{d,h_{\text{late}},\pi}(t))) \leq E'_{d,h_{\text{late}},\pi}(t)$$

holds, then π is a valid policy for d .

Briefly, Thm. III.2 is a result of the fact that Eq. 9 is the harvesting trace that minimizes $E'_{d,h,\pi}$. By the definition of max we have that $(\forall h) E'_{d,h,\pi} \leq E_{d,h,\pi}$ which in turn implies that $E'_{d,\pi}$ (defined analogously to $E_{d,\pi}$) is less than or equal to $E_{d,\pi}$. Since h_{late} minimizes $E'_{d,\pi}$ for all t , we have that $E'_{d,h_{\text{late}},\pi}$ is a lower bound for $E_{d,\pi}$.

Naturally, there are as of yet no real devices with an infinite capacity to store energy. However, the important distinction being made in the $\theta = \infty$ case is that no harvested energy is lost due to being discarded by the completely full energy storage. As long as θ is “large enough” that it never becomes full, then Thm. III.2 will still hold – in practice, this is a requirement that is often satisfied while executing online policies that are consistent w.r.t. a valid offline policy, discussed next.

B. Opportunistic Communication

Valid offline policies have the attractive property of guaranteed worst-case latency without requiring any additional cooperation between the nodes online. Unfortunately, valid offline policies tend to be (practically by definition) overly conservative in situations with better-than-worst-case energy harvests. To this end, we would like to allow nodes to opportunistically communicate in the event of windfall harvests while somehow maintaining the worst-case latency guarantee.

Definition III.3 (Consistent policy). *Let π be a valid policy for a device d . A policy π^* is considered to be consistent w.r.t. π iff π^* is a valid policy for d and $(\forall h \models \Gamma, t \in \mathbb{N})(\pi(x_{d,h,\pi}(t)) \neq s1 \Rightarrow \pi(x_{d,h,\pi}(t)) = \pi^*(x_{d,h,\pi^*}(t)))$. In other words, a policy is consistent w.r.t. to a base policy iff it is valid and “promises” to perform all of the non-sleep actions that the base policy performs, given the same harvesting trace.*

Consistent policies help us to achieve precisely what we want in the following way: Using the analysis in III-A, we first prove a base policy π to be a valid offline policy for a device d . Any policy π^* that is consistent w.r.t. π can be executed by d as a substitute for π with the same worst-case behavior as π . For example, a corollary of Thm. III.1 is that two devices executing consistent substitutes for valid offline periodic policies with worst-case latency L will now have worst-case latency less than or equal to L .

As the space of possible policies is quite large, synthesizing consistent policies directly would be difficult. We instead opt to apply similar reasoning as is used for the validity of offline policies to check in an online fashion whether, given the current state, a round of opportunistic communication can be performed consistently. Essentially, if the base policy performs $s1$ at time s , we can simulate taking an opportunistic action by replacing the cost of $s1$ at s by the cost of the opportunistic action, updating the boundary conditions for the

possible harvesting traces to match the observed trace, and then following the base policy for $t > s$. This gives us a new projected lower bound for the energy level at times $t > s$ computed by the same procedures as in III-A that we can check against $E_{d,\pi}$ – if the new projected lower bound is greater than $E_{d,\pi}$ for $t > s$, then the opportunistic action is consistent to perform. The computational cost of checking the consistency of a single opportunistic communication is therefore equivalent to the cost of checking the validity of the base offline policy. In other words, for the special case $\ell = \infty$ the consistency check is a fast constant-time operation. We leave more efficient methods for $\ell < \infty$ for future work.

IV. EXPERIMENTS


In this section, we give details about the representative EH-WSN architecture and instantiate a hardware model, evaluate the prediction accuracy of the constraint-based harvesting model on real-world traces, and compare the performance of a valid inconsistent online locally greedy policy (GREEDY) and a valid periodic offline policy (OFFL) to our proposed approach (CONS-GR) that opportunistically follows the locally greedy policy GREEDY but is consistent w.r.t the valid periodic offline policy OFFL.

A. Experimental Setup


Our hardware model is based on the energy-harvesting platform described in [11]. Each simulated node is controlled by a low-power MCU modeled after the MSP430FR5994 with active power 2.2mW and standby power 0.88 μ W. The main radio is a CC1101 operating at 433MHz, consuming 70.4mW and 35.2mW power on the sender and receiver respectively. We select a relatively low data rate of 4Kbits/s to ensure a low error rate at the receiver. In order to exploit asymmetric energy availability at the sender and receiver, the nodes are also equipped with a second, custom-made backscattering radio [14] operating at 915MHz with 1mW and 70mW power consumption at the sender and receiver respectively. We define one round to last one second and assume that the data to be sent has a payload of 1Kbits.

The sender device has action set $A_{tx} = \{sl, tx-br, tx-ba\}$ corresponding to “sleep,” “transmit with main radio,” and “transmit with backscattering radio” respectively whereas the receiver device has action set $A_{rx} = \{sl, rx-br, rx-ba\}$ corresponding to “sleep,” “receive with main radio,” and “receive with backscattering radio” respectively. Joint actions that match mode, i.e. $(tx-br, rx-br)$ and $(tx-ba, rx-ba)$, are considered successful and unsuccessful otherwise.

We generate harvester traces from real-world meteorological data gathered from [15], [16]. The dataset records observations of the global horizontal irradiance, wind speed, air pressure, and temperature at one-minute intervals. We suppose a solar panel of area 0.1m² that has power efficiency 6%. For the wind harvester, we suppose a small wind turbine with sweep area 0.1m² and power efficiency 15%.



(a) Four wind energy harvesting traces from [15] corresponding to the month of March for the years 2018-2021.



(b) Lower bound energy harvest versus window size computed by Eq. 1
Fig. 2. Energy harvesting traces with different behaviors can be captured with a single harvesting model. Here, historical wind harvesting data for the month of March across four years can be modeled with the candidate Γ -set shown in dashed black line in (b).

B. Specification Mining

We briefly justify the harvesting model proposed in II-B. Given a training trace from real-world wind and solar energy harvesters obtained via IV-A, we mine a Γ -set from one-month-long training traces according to Eq. 2 and check if the same month in the following year satisfies the mined set. Given that the mined trace is $\Gamma = \{[m_k, k]\}_k$, we find that the set $\Gamma_{60\%} = \{[0.6 * m_k, k]\}_k$ models the energy harvests for the same month in the following year $> 95\%$ of the time. An example of this straightforward method using historical data to come up with a candidate harvesting model is shown in Fig. 2.

C. Performance Evaluation

To evaluate the impact of planning around the worst-case energy harvests in two-node settings, we mine specification sets $\Gamma_1, \Gamma_2, \dots, \Gamma_D$ from D samples of real-world solar and wind harvesting traces $\hat{h}_1, \hat{h}_2, \dots, \hat{h}_D$. The length of the harvesting traces varies between one day in length to half a week in length. For each $i \in \mathbb{N}_D$, we instantiate device models $tx_i = (A_{tx}, \infty, \phi, \Gamma_i)$ and $rx_i = (A_{rx}, \infty, \phi, \Gamma_i)$ that model sender and receiver nodes respectively in the environment Γ_i . Note that implicitly we are assuming that $\forall i, \max\{m/\tau : [m, \tau] \in \Gamma_i\} > \phi(sl)$, i.e. that each environment provides at least enough energy on average for the device to accumulate enough power beyond what is needed to sleep so that it can eventually communicate in the worst case. For all ≈ 150 pairs $(i, j) \in \mathbb{N}_k^2$ where $\hat{h}_i, \hat{h}_j$ are of the same length, we compute a valid offline policy OFFL from Γ_i, Γ_j by simulating a run of h_{late} , as described in III-A2, on tx_i and on rx_j – OFFL instructs nodes to perform a valid successful joint action at every time t whenever it was possible in the simulation and instructs them to sleep at all other times. The online policy GREEDY, as the name suggests, performs

TABLE I
THE RELATIVE AGGREGATE PERFORMANCE OF
CONS-GR (C-G), GREEDY (GR), AND OFFL.

Worst-Case Latency	C-G vs Gr	Gr vs OFFL	C-G vs OFFL
Minimum	-96%	-92%	-93%
Maximum	0%	1289%	0%
Average	-44%	112%	-37%

a valid non-sleep action whenever energy is available. The policy CONS-GR opportunistically attempts actions suggested by GREEDY, but is consistent w.r.t. OFFL.

We then measure the worst-case communication latencies observed by $\tau_{x_i, \tau_{x_j}}$ on the traces $\hat{h}_i, \hat{h}_j$ running policies OFFL, GREEDY, and CONS-GR. In order to compare the performance of the different communication policies, we examine the percent difference in worst-case latency in aggregate over the scenarios. For example, given resultant worst-case latencies $l_{i,j}^{\text{GREEDY}}, l_{i,j}^{\text{OFFL}}$ in scenario (i, j) with communication policies GREEDY and OFFL respectively, the maximum relative performance of GREEDY vs OFFL is given by $\max_{i,j} (l_{i,j}^{\text{GREEDY}} - l_{i,j}^{\text{OFFL}}) / l_{i,j}^{\text{OFFL}}$. The aggregate percent difference over all i, j of GREEDY vs OFFL, of CONS-GR vs OFFL, and of CONS-GR vs GREEDY is presented in Table I. Per-experiment worst-observed latencies are shown in Fig. 3.

Since CONS-GR respects the worst-case behavior, we expect it to result in worst-case latency that is never larger than the upper bound. This is indeed what we observe; CONS-GR is at worst as large as OFFL. The comparison between GREEDY and OFFL is meant to reinforce the motivation of our work: if it was the case that the locally greedy policy always performed well, there would be no reason to seek worst-case guarantees on latency. Indeed, we see that at best GREEDY outperforms our worst-case guarantee by as much as 92% in certain scenarios. On average however, we observe that GREEDY leads to worst-observed latency 112% in excess of what our worst-case upper bound is able to guarantee. Furthermore, in some scenarios we observe “run away” latencies that are much larger (up to 1289%) than what our guarantees can provide. Interestingly, we observe that CONS-GR *never* performs worse than GREEDY, and on average results in worst-observed latency that is 44% smaller than the worst-observed latency under GREEDY. Alternatively, pathological inputs that drive the performance of CONS-GR to the worst-case upper bound but where GREEDY outperforms the worst-case upper bound are theoretically possible, but we observe no such inputs among the real-world traces.




Fig. 3. We compare the worst-observed latencies for devices (τ_{x_i}, τ_{x_j}) on the harvesting scenarios $\hat{h}_i, \hat{h}_j$. Given an identical scenario, we find that CONS-GR outperforms both OFFL and GREEDY in terms of worst-observed latency by up to $\sim 10x$ and $\sim 100x$ respectively.

V. CONCLUSION & FUTURE WORK

We examine a basic, but general, constraint-based approach for modeling the energy provided by the environment. We have demonstrated on real-world traces that locally greedy communication policies can and do lead to run-away latencies, and that planning against the worst-case harvesting scenario allows us to obtain latency guarantees and performance benefits. Our work provides a first step towards reliable communication of EH-WSNs in unpredictable environments without a battery-tethered node, and we envision networked EH-WSNs as the prime subject of future work.

Acknowledgements We gratefully acknowledge the support from NSF grants 1646497, 2007274 and ONR grant N00014-19-1-2496.

REFERENCES

- [1] G. Park, T. Rosing, M. D. Todd, C. R. Farrar, and W. Hodgkiss, “Energy harvesting for structural health monitoring sensor networks,” *Journal of Infrastructure Systems*, vol. 14, March 2008.
- [2] M. Magno and D. Boyle, “Wearable energy harvesting: From body to battery,” in *12th International Conference on Design Technology of Integrated Systems In Nanoscale Era (DTIS)*, April 2017, pp. 1–6.
- [3] M. Magno *et al.*, “SmarTEG: An Autonomous Wireless Sensor Node for High Accuracy Accelerometer-based Monitoring,” *Sensors (Switzerland)*, vol. 19, no. 12, 2019.
- [4] B. Lucia, V. Balaji, A. Colin, K. Maeng, and E. Ruppel, “Intermittent Computing: Challenges and Opportunities,” in *2nd Summit on Advances in Programming Languages (SNAPL 2017)*, vol. 71, 2017, pp. 8:1–8:14.
- [5] B. Lucia and B. Ransford, “A simpler, safer programming and execution model for intermittent systems,” in *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2015, pp. 575–585.
- [6] X. Lu, P. Wang, D. Niyato, D. I. Kim, and Z. Han, “Wireless Networks with RF Energy Harvesting: A Contemporary Survey,” *IEEE Communications Surveys and Tutorials*, vol. 17, no. 2, pp. 757–789, 2015.
- [7] W. K. Seah and J. P. Olds, “Data Delivery Scheme for Wireless Sensor Network Powered by RF Energy Harvesting,” *IEEE Wireless Communications and Networking Conference*, pp. 1498–1503, 2013.
- [8] Z. Zhou, C. Fu, C. J. Xue, and S. Han, “Transmit or discard: optimizing data freshness in networked embedded systems with energy harvesting sources,” in *56th ACM/IEEE Design Automation Conference (DAC)*, 2019, pp. 1–6.
- [9] P. Blasco, D. Gunduz, and M. Dohler, “A learning theoretic approach to energy harvesting communication system optimization,” *IEEE Transactions on Wireless Communications*, vol. 12, no. 4, pp. 1872–1882, 2013.
- [10] A. Ortiz, H. Al-Shatri, X. Li, T. Weber, and A. Klein, “Reinforcement learning for energy harvesting point-to-point communications,” in *IEEE International Conference on Communications (ICC)*, 2016, pp. 1–6.
- [11] Y. Wu, Z. Jia, F. Fang, and J. Hu, “Cooperative communication between two transiently powered sensor nodes by reinforcement learning,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2021.
- [12] R. C. Carrano, D. Passos, L. C. S. Magalhaes, and C. V. N. Albuquerque, “Survey and Taxonomy of Duty Cycling Mechanisms in Wireless Sensor Networks,” *IEEE Communications Surveys Tutorials*, vol. 16, no. 1, pp. 181–194, 2014.
- [13] S. L. Sampayo, J. Montavont, F. Prégaldiny, and T. Noël, “Is wake-up radio the ultimate solution to the latency-energy tradeoff in multi-hop wireless sensor networks?” in *14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2018, pp. 1–8.
- [14] P. Hu, P. Zhang, M. Rostami, and D. Ganesan, “Braidio: An integrated active-passive radio for mobile devices with asymmetric energy budgets,” in *ACM SIGCOMM*, 2016, pp. 384–397.
- [15] D. Jager and A. Andreas, “NREL National Wind Technology Center: M2 Tower,” <http://dx.doi.org/10.5439/1052222>.
- [16] A. Andreas and T. Stoffel, “NREL Solar Radiation Research Laboratory: Baseline Measurement System,” <http://dx.doi.org/10.5439/1052221>.