


Harvesting-Aware Optimal Communication Scheme for Infrastructure-Less Sensing

Journal Article**Author(s):**

Sigrist, Lukas ; Ahmed, Rehan; Gomez, Andres; Thiele, Lothar

Publication date:

2020-06

Permanent link:

<https://doi.org/10.3929/ethz-b-000412614>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

ACM Transactions on Internet of Things 1(4), <https://doi.org/10.1145/3395928>

Funding acknowledgement:

157048 - Transient Computing Systems (SNF)

Harvesting-Aware Optimal Communication Scheme for Infrastructure-Less Sensing

LUKAS SIGRIST, ETH Zurich, Switzerland

REHAN AHMED, Information Technology University, Pakistan

ANDRES GOMEZ, ETH Zurich, Switzerland

LOTHAR THIELE, ETH Zurich, Switzerland

Sensing systems for long-term monitoring constitute an important part of the emerging Internet of Things. In this domain, energy harvesting and infrastructure-less communication enable truly autonomous and maintenance-free operation of sensor nodes gathering long-term environmental data. Due to the infrastructure-less nature of the communication receivers are not always available. The variable energy provided by the environment and the receiver's mobility lead to non-deterministic node availability. In this work, we study infrastructure-less data transmission schemes to optimize communication when both senders and receivers exhibit intermittent behavior. We rely on the notion of data utility, describing the importance of sensed data to the receiver, to determine an optimal communication scheme. Deriving the communication policy that maximizes the utility of the received data is shown to be a convex optimization problem. The resulting scheme is implemented and validated on a batteryless Bluetooth Low Energy sensor node that communicates to commodity smartphones. Our evaluation demonstrates that the model accurately captures the application scenario with a maximum Root-Mean-Square Error (RMSE) of less than 0.016 in data reception probability. The communication scheme's adaptiveness to variable harvesting conditions is experimentally demonstrated under varying harvesting conditions and is shown to significantly increase the data utility.

CCS Concepts: • **Information systems** → **Mobile information processing systems**; • **Computer systems organization** → **Sensor networks**; *Sensors and actuators*; • **Hardware** → **Sensor applications and deployments**.

Additional Key Words and Phrases: Wireless Sensing, Optimization, Energy Harvesting, Batteryless, Infrastructure-Less, Internet of Things

ACM Reference Format:

Lukas Sigrist, Rehan Ahmed, Andres Gomez, and Lothar Thiele. 2020. Harvesting-Aware Optimal Communication Scheme for Infrastructure-Less Sensing. *ACM Trans. Internet Things* 1, 1, Article 1 (January 2020), 26 pages. <https://doi.org/10.1145/3395928>

1 INTRODUCTION

Networked sensing systems are widely used for long-term monitoring of our natural or built environment, personal health, assets, mobility, and many more. While specializing on diverse sensing methodologies and applications, they all accumulate time-dependent sensor information

Authors' addresses: Lukas Sigrist, ETH Zurich, Computer Engineering and Networks Laboratory, Zurich, Switzerland, sigristl@ethz.ch; Rehan Ahmed, Information Technology University, Department of Computer Engineering, Lahore, Pakistan, rehan.ahmed@itu.edu.pk; Andres Gomez, ETH Zurich, Computer Engineering and Networks Laboratory, Zurich, Switzerland, gomez@ethz.ch; Lothar Thiele, ETH Zurich, Computer Engineering and Networks Laboratory, Zurich, Switzerland, thiele@ethz.ch.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2577-6207/2020/1-ART1 \$15.00

<https://doi.org/10.1145/3395928>

and extract information from received data. Whereas in some scenarios, processing and transmitting the most recent data point is sufficient, the availability of longer time series of data is usually requested and necessary for the end user.

We consider long-term monitoring with sensor nodes that target a persistent and maintenance free deployment. These sensors record and locally accumulate environment data, such as temperature, humidity, soil moisture, etc. The sensed data is transmitted to mobile receivers in the vicinity, prioritizing the data most relevant/important to the receivers. To this end, the importance of data is modeled and quantified as a function of its age. The objective is sending a selection of the sensed data that maximizes the usefulness of the data accumulated on the receiver end. Targeting a persistent and maintenance free operation (deploy-and-forget scenario), a suitable communication scheme achieving this goal has to be found. At the same time such a scheme has to consider the specific requirements of the scenario.

An *infrastructure-less* communication scheme is required to not rely on deploying and managing any stationary communication or energy supply infrastructure. This facilitates the deployment, operation and long-term maintenance of the sensor network considerably. For *long-term operation* the system needs to operate autonomously for extended periods of up to several decades. *Scalability* of the communication scheme is demanded to allow deployment of many sensors in parallel and to guarantee reception of relevant information with a large number of receivers. These requirements heavily constrain our design and lead to the design choices presented in the following. The mobility of the receiver is a direct consequence of the *infrastructure-less* requirement, since any static data collection device is considered infrastructure.

Energy harvesting is seen as a key enabler for long-term, energy neutral operation [3]. However, the typical combination of harvesting with batteries is limited due to their finite recharge cycles [44]. To guarantee *long-term operation*, we design our sensor nodes in line with the new class of batteryless devices, which use an energy buffer in form of a capacitor only as in [2, 14, 29]. Though these systems promise virtually unlimited lifetimes, their operation is highly dependent on the environment. The resulting variability in both timing and energy requires our solution to tolerate and adapt to non-deterministic sensor node execution rates.

To satisfy the *scalability* requirement, we rely on broadcast-based communication, where the sensor nodes advertise data to nearby receivers without employing any request or acknowledgment based scheme. The reason for this design choice is two fold. Firstly, maintaining a network with batteryless devices is infeasible, as this would require some minimum service level, which cannot be guaranteed in these systems. Consequently, this rules out bidirectional communication schemes in the considered scenario. Secondly, requests and/or acknowledgments do not scale to many receivers with potentially conflicting data requirements. For these reasons, sensor nodes, henceforth referred to as *sensor beacons*, only broadcast a selection of sensed data. The unidirectional communication scheme preserves the privacy of the receivers. Neither a sensor beacon, nor other devices can detect the existence of a receiver.

To find an optimal data selection strategy, we introduce the notion of *data utility* that combines the concept of aging of information [16, 43] and the total amount of accumulated data. The challenge in finding the selection strategy that maximizes this *data utility* stems from the uncertainty in two different fronts. Firstly, the selection has to adapt to the non-deterministic activation intervals of batteryless devices that affect the rate at which sensing and transmitting is performed, thus altering the selection of data values providing maximal data utility to receivers. Secondly, mobile receivers are only intermittently within broadcasting range. To maximize the data utility they receive, it is therefore important to incorporate their behavior into the optimization by modeling their listening characteristics. This work is the first to address both of these problems simultaneously and derive the optimal communication scheme. More specifically, we summarize our contributions as follows:

- We propose a receiver-centric communication model for infrastructure-less monitoring in non-deterministic environments with batteryless sensors and stochastic receivers.
- We maximize the receivers' data utility by solving the corresponding convex optimization problem determining the optimized sender-side data selection strategy. This represents the first work to optimize the communication policy for infrastructure-less monitoring with batteryless systems.
- Based on a pool of optimized communication policies, we derive a run-time mechanism that adapts to dynamic environmental conditions with negligible overhead.
- We design a harvesting aware batteryless sensor node, implement the proposed communication scheme on top of it, and provide an extensive evaluation that confirms the suitability of the presented models and methods.

The real-world experiments demonstrate the feasibility of the proposed method and confirm that the run-time mechanisms are implemented with small overhead. With a maximal Root-Mean-Square Error (RMSE) of 0.0157 under controlled power conditions, the experimentally observed data reception probabilities accurately match the theoretical values. The optimized communication scheme is demonstrated to significantly outperform the baseline schemes in terms of data utility. A final deployment under real-world conditions proves that our scheme successfully adapts the communication policy to the harvesting conditions at run time.

The remainder is organized as follows: Section 2 covers related works. Section 3 introduces the overall architecture and the communication model. In Section 4 the optimization problem is formulated and solved, and Section 5 generalizes the model. The batteryless sensor design and implementation of the communication scheme is discussed in Section 6. We experimentally evaluate and validate the model in Section 7, and conclude in Section 8.

2 RELATED RESEARCH

Infrastructure-less sensing has been applied in different scenarios. This includes battery powered wireless beacons for localization [8] and harvesting powered ambient sensor nodes [27]. Recent batteryless applications include RF sensor tags that deployed cameras [29] or performed ambient sensing [23]. All of these systems must address issues related to data management, in order to decide what data to communicate. In addition, an appropriate communication scheme must be selected in accordance to the underlying system architecture support. In the following, we address works in these related fields.

Data Utility To abstract and quantify the usefulness of sensed and transmitted data, several researchers have used the notion of data utility. The authors of [31] used data utility to optimize adaptive sensing in wireless sensor networks. Kim et al. [22] applied the notion to data that becomes obsolete after a fixed timeout, and made scheduling decisions to optimize this utility. Hester et al. [16] introduced a specification language in order to formalize the age dependent utility of data and to perform task scheduling decisions in batteryless sensing. An overview on the use of data utility in information theoretical works [40] focused on offline and online schemes for optimizing the throughput in energy harvesting systems with constrained energy buffers. A more recent work in this domain [13] investigated optimal sampling and delay minimal communication strategies for energy harvesting system in single user and broadcast scenarios.

Following the same idea of age of information, we rely on the notion of data utility to derive an optimization formulation and determine the optimal communication strategy for infrastructure-less sensing scenarios.

Data Compression Data utility can be maximized using compression, as the sensed data is typically correlated in time and space. The compression allows for significant communication




Fig. 1. The considered scenario consists of a batteryless sensor node monitoring the environment and broadcasting (RF) data packets of N sensor values to smartphones with listening intervals described by distribution \tilde{R} . The application specification \tilde{Q}^{spec} describes the relative importance of individual data units.

energy savings with reasonable processing overhead [35]. The methods used in wireless sensor network can be broadly grouped as follows [34]: entropy coding based compression [11], prediction based compression [5], transformation based compression [26] or applying compressed sensing [9].

The communication scheme presented in the following leverages the low complexity Haar wavelet transformation for lossy compression of long-term historical values.

Infrastructure-Less Communication Specialized communication schemes are required to disseminate data without relying on any infrastructure like wires or base stations. In the area of delay tolerant networks, researchers addressed the mobility on the sender and/or receiver side [24, 42]. The various data dissemination schemes proposed to handle the intermittent connectivity are based on flooding or routing and make use of known or learned mobility patterns and neighbor discovery [32, 39, 41]. Similarly, energy harvesting sensor networks employed flooding [20, 27] or routing based schemes [25], relying on large batteries to mitigate the non-determinism in energy availability. Typically, these systems adapted the communication and sensing rate jointly to the harvesting conditions to sustain system operation [15, 38].

These techniques are not directly applicable to batteryless systems, as they rely on periodic discovery and updates with neighboring nodes for dynamic forwarding and routing. We therefore employ a broadcasting communication scheme.

Batteryless System Support The emerging class of batteryless systems avoids bulky batteries as energy supply and solely relies on energy harvesting and small energy buffers such as supercapacitors. Numerous proposals focusing on node architecture challenges have been published. They addressed the non-deterministic system operation by introducing state retention with periodic [33] or event-driven state saving mechanisms [19]. Others used an energy-driven approach by providing the energy in small, but guaranteed bursts and extending the hardware for more flexible energy management such as [10, 29] and the concepts introduced in [14]. Furthermore, the problem of persistent timekeeping has been addressed [17] and higher level abstraction for operating system integration and data management have recently been proposed [4, 16].

The batteryless sensor developed as part of the presented work builds on the energy burst principle introduced in [14]. Furthermore, a backup power domain is incorporated so that sensor nodes do not lose the notion of time over long periods of energy unavailability, for example during nighttime.

3 SYSTEM OVERVIEW

First, we introduce the infrastructure-less communication scenario more formally in this section. We state the models required for formulating the optimal communication problem, followed by an architectural overview in which we introduce the different components of the proposed communication scheme.

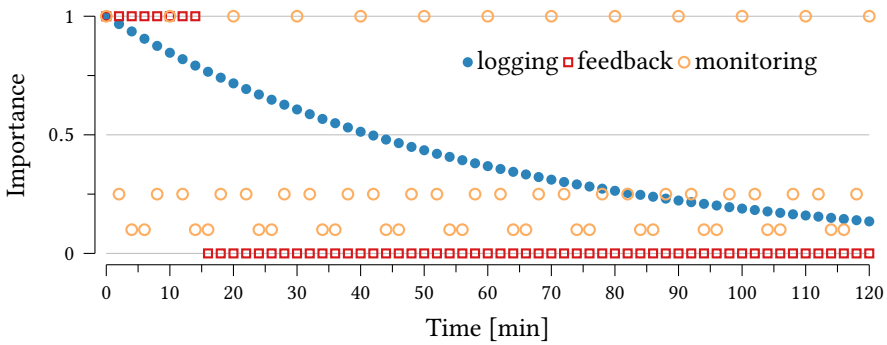


Fig. 2. Examples of the relative importance specification \tilde{Q}^{spec} for a few reference scenarios with difference requirements.

3.1 Communication Scenario

We focus on a long-term infrastructure-less ambient monitoring scenario, as illustrated in Fig. 1: a sensor node accumulates ambient sensor data and buffers them as timestamped *data units* in non-volatile memory. The sensor node broadcasts a selection of N data units per data packet, which can be received by devices within its range. There may be several receivers active simultaneously collecting this data. We consider smartphones as receiver devices, as they suit the infrastructure-less sensing scenario very well.

The scenario supposes that a receiver values the availability of data depending on their *age*. To this end, a specification vector \tilde{Q}^{spec} is introduced that enables an application designer to specify the relative importance of a data unit as a function of its age. Considering air quality monitoring as an example, this could be reflected by a specification with an importance that decreases with age, as represented by the *logging* specification shown in Fig. 2. For applications that provide feedback or control based on the received data, typically only values up to a maximum age are considered relevant, as is reflected by the step function of the *feedback* specification. For long-term monitoring where the availability of a long history is of higher importance than detailed coverage, a specification as presented by *monitoring* in Fig. 2 could be used: data units with a large interval have higher importance, while the importance decreases for intermediate, finer granular data.

Similarly, the mobility of receivers is specified by a listening interval distribution \tilde{R} . It characterizes the duration a specific receiver is within the range of the sensor and able to receive data. Examples for this specification are shown in Fig. 3. An *office* scenario reflecting the mobility of an office worker, a *meeting room* scenario in which people are present for one or two hour meetings, or a *kitchen* scenario in which people are either briefly dropping to get a coffee, or having an extended lunch break.

The system specification $(\tilde{Q}^{\text{spec}}, \tilde{R}, N)$ allows for a quantitative description of a scenario and enables us to derive an optimized communication strategy for the sensor node. The exact objective of the corresponding optimization problem, i.e., the exact notion of *data utility* used in this work, is defined in Section 4.1.

3.2 Notation and Communication Model

We introduce a model that abstracts the considered infrastructure-less monitoring scenario. The model consists of two parts: the sensor node, also referred to as *sender*, that transmits a selection of

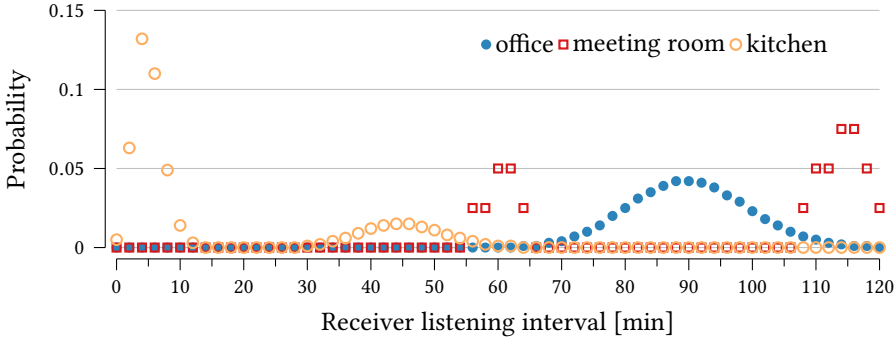


Fig. 3. Examples of receiver characteristics \tilde{R} for three different scenarios.

the sensed data, and a (potentially large set of) device(s) referred to as *receiver(s)* that collect the broadcast data.

Notation Throughout this paper, bold symbols are used to denote vectors and matrices and subscripts denote individual elements of the vector or matrix. For instance, if \mathbf{X} is a vector, X_i denotes its i^{th} element. In case \mathbf{X} is a matrix, \mathbf{X}_i denotes its i^{th} row, and $X_{i,j}$ denotes the element in its i^{th} row and j^{th} column. For a discrete random variable \mathcal{Y} , \mathbf{Y} denotes the vector of probabilities, whose i^{th} element is the probability that \mathcal{Y} is equal to i ; i.e. $Y_i = \Pr(\mathcal{Y} = i)$. Y^* denotes the largest index i for which $Y_i \neq 0$. All time dependent model components are specified in terms of *unit time* δ_0 that defines the time granularity.

Sender The sender performs sense-process-transmit operations based on the principle of energy bursts. Its operation is characterized by the following parameters:

- δ Integer specifying the time between sender activations, defined as a multiple of the base time unit δ_0 . For the presentation of the model and the optimization we assume that $\delta = 1$, i.e., the activation interval corresponds to the base time unit δ_0 . This can be generalized to arbitrary δ as discussed in Section 5.1.
- N Size of a data packet in number of data units.
- $E^{\text{activation}}$ Energy required per sensor node activation. This is the sum of the energy to sense a new data unit E^{sense} , buffer management and data selection for transmit E^{process} , and sending one data packet of N data units E^{send} .
- pwr^{in} Input power available to the sensor node. This is the harvesting power actually available after accounting for transducer, conversion, energy management, and storage inefficiencies.

Given these parameters, the following condition on the average input power $\widehat{\text{pwr}}^{\text{in}}$ must be satisfied for sustainable sensor node operation:

$$\widehat{\text{pwr}}^{\text{in}} \geq \frac{E^{\text{activation}}}{\delta \cdot \delta_0} \quad (1)$$

The interval δ has to be adapted accordingly by a run-time mechanism to satisfy the above condition.

The selection of the N data units is determined by the probability vector \mathbf{P} . The individual P_i specify the probability of including data that have been sensed at time $i \cdot \delta$ before the current time.

Receiver The receivers arrive at an unknown point in time and are active for a limited time only. As introduced above, the individual probability of occurrence of the listing intervals is described by vector $\tilde{\mathbf{R}}$. Furthermore, the specification \tilde{Q}^{spec} states the relative importance of data units for




Fig. 4. The overall architecture consists of an *online* harvesting aware data transmission scheme. The key component is a pool of communication policies that are optimized *offline* based on a model for the infrastructure-less communication scenario.

a receiver. Depending on the activation interval described by δ , the specifications \tilde{Q}^{spec} and \tilde{R} are rescaled to Q^{spec} and R . To understand why the rescaling is necessary, consider the following example: if a sensor is activated every minute, the entries defined for $[1, 2, \dots, 59]$ seconds of a \tilde{Q}^{spec} with second granularity are obsolete, since the corresponding data is never sensed. Therefore, we need to adapt \tilde{Q}^{spec} and \tilde{R} based on activation interval δ . The details of this rescaling step are explained in Section 5.1. For the assumption of $\delta = 1$, the probability vectors R and \tilde{R} , and Q^{spec} and \tilde{Q}^{spec} are identical. The receiver side reception probability is encapsulated in Q , with its elements Q_i representing the probability of receiving a data unit that is $i \cdot \delta$ time units old. They depend on both, the receiver characteristics R and selection probabilities P .

In the following, the sensor node architecture and its operation is detailed, followed by an overview of the communication scheme optimization that is discussed in depth in Section 4.

3.3 Sensor Node Architecture

The detailed operations of the sensor node performed as part of an activation are illustrated in the *online* part shown on the right in Fig. 4. The node starts with reading the ambient *sensors* and *timer* and then stores these values as one data unit in a non-volatile *buffer*. In addition, aggregate values like a compressed history are calculated in the *update aggregates* component. At the end of an activation, the node broadcasts (RF) a selection of N data units. The data units to transmit are selected in the *packet selector* unit according to the optimal selection policy ρ .

The optimal policy is dependent on the harvesting condition pwr^{in} that dictates the time interval between activations (1). Therefore, a *policy pool* is stored on the sensor node, containing optimal policies for different activation intervals δ . At run time, a *duty-cycle adaptation* calculates the activation interval that allows sustainable operation of the sensor and activates the corresponding optimal policy ρ . The derivation of the optimal policy for individual intervals is done offline based on the procedure outlined below.

3.4 Determination of the Policy Pool

The optimal communication policy depends on the system specifications $(\tilde{Q}^{\text{spec}}, \tilde{R}, N)$ and the current input power pwr^{in} . An energy model of the sensor node is used to calculate the activation interval δ that guarantees long-term sustainable operation for a constant input power pwr^{in} .

The scaled specifications Q^{spec} and R are passed together with the data packet size N to the model based *optimization*. This block solves the optimization formulation detailed in Section 4 and returns the resulting vector P specifying the optimal probabilities of selecting data units of a certain age. The probabilities are bin-packed to get individual probability mass functions for each of the N data slots of a packet. The resulting probability mass functions are summarized as the optimal selection policy ρ and used at run time to select the N data units to be transmitted.

The optimization flow of determining δ , rescaling \tilde{Q}^{spec} and \tilde{R} , finding the optimal probabilities P , and determining ρ by bin-packing P , is repeated for several discrete power levels. The corresponding set of optimal communication policies for different activation intervals δ define the policy pool $\{(\delta, \rho)\}$ that is deployed on the sensor node.

4 OPTIMAL COMMUNICATION POLICY

Based on the specifications and the communication model introduced before, we derive the optimal data selection policy. First, we formalize the notion of *data utility* that is serving as optimization objective. We then analytically derive the probability Q_i of receiving data with an age of i time units. Subsequently, we present the convex optimization problem to determine the optimal sender side data selection probability. Finally, the online data selection algorithm for implementing the optimal communication policy is introduced.

4.1 Data Utility

The data utility accumulated for given values of Q and Q^{spec} is given by the following equation:

$$U(Q, Q^{\text{spec}}) = \min_{i \in \{0, \dots, Q^{\text{spec}*}\}} \{Q_i / Q_i^{\text{spec}}\} \quad (2)$$

This utility formulation takes into account the relative importance Q^{spec} of data of a given age. Higher values of Q_i / Q_i^{spec} mean that important data is received with higher probability. Note that maximizing the above utility $U(Q, Q^{\text{spec}})$ for a given Q^{spec} is identical to the following formulation:

$$\text{maximize } \alpha \quad \text{s.t. } Q_i \geq \alpha \cdot Q_i^{\text{spec}} \quad \forall i \in \{0, 1, \dots, Q^{\text{spec}*}\} \quad (3)$$

i.e. determining a Q such that all its elements are no less than a maximally scaled version of Q^{spec} .

4.2 Probability of Data Reception

To derive the data reception probability Q_i we consider an introductory example. Let us examine a receiver that is within the range of a sensor node for 3 time units, i.e., at times 0, 1, 2. After time 2 the receiver moves out of range and cannot receive further data packets. Furthermore, we assume the sensor transmits one data packet every interval. To find the probability of having received data of age 5 at least once when moving out of range, there are three cases to consider:

- a) Data of age 3 is sent at time 0. In this case, the data ages for 2 time units on the receiver side and therefore has an age of 5 when the receiver moves out of range.
- b) Data of age 4 is sent at time 1 and ages for 1 time unit on the receiver side.
- c) Data of age 5 is sent at time 2.

Each of these cases occurs with a specific probability that depends on the selection probabilities P . The data reception probability Q_i of i time units old data on the receiver side is influenced by both, the receivers listening interval distribution R and the data selection probabilities P . For the analytical derivation of Q_i presented here, the communication link is assumed lossless for simplicity reasons, even though the stochastic model allows integrating a packet reception probability.

To derive Q_i we look at the $i + 1$ scenarios in which an i time units old data unit is not received. These scenarios are represented in Fig. 5: the enumerated gray nodes refer to the active time of the receivers, while the double circled red nodes indicate the individual scenarios in which the i time units old data is not received. While active, receivers transitions horizontally from the circled initial node to any of the gray nodes enumerated with the receivers' active time. The double circled red nodes are reached by vertical transitions referring to cases in which a receiver becomes inactive after the number of intervals indicated by the label of its source node. The probability of a scenario s_i is calculated by multiplying the probabilities indicated on the edges leading to that scenario. In




Fig. 5. Enumeration of the scenarios where i time unit old data is not received (s_0, \dots, s_i).

the first scenario s_0 , a receiver is active for exactly one time unit and the sender has not selected the i^{th} data for transmission. The probability of this scenario is $(1 - P_i) \cdot R_0$. Scenario s_1 encapsulates the situation where a receiver is active for exactly two time units and the sender does not transmit the i^{th} data unit at time 0 and the $(i - 1)^{\text{th}}$ data unit at time 1. The last scenario s_i refers to a receiver active for *at least* i time units and the sender does not send i^{th} data at time 0, $(i - 1)^{\text{th}}$ data at time 1, and so on. Summarizing all scenarios of not receiving i time units old data, the probability Q_i is calculated as:

$$Q_i = 1 - \underbrace{\sum_{j=i}^{j \leq R^*} R_j \cdot \prod_{j=0}^{j \leq i} (1 - P_j)}_{\text{scenario } s_i} - \underbrace{\sum_{m=1}^{m \leq i} \left(R_{i-m} \cdot \prod_{j=m}^{j \leq i} (1 - P_j) \right)}_{\text{scenarios } s_0, \dots, s_{i-1}} \quad (4)$$

The first sum in the formula computes the probability of scenario s_i , while the second sum covers the remaining i scenarios (s_0, \dots, s_{i-1}) by summing up the respective probabilities of occurrence for each scenario. The probabilities of all scenarios can be accumulated because all scenarios are disjoint. We also use the following equivalent formulation for conciseness:

$$Q_i = 1 - \sum_{j=0}^{j \leq R^*} \left(R_j \cdot \prod_{k \in S(i, j)} (1 - P_k) \right), \quad (5)$$

where the sets of indexes $S(i, j)$ are defined as:

$$S(i, j) = \begin{cases} \{i - j, \dots, i\} & j < i \\ \{0, \dots, i\} & \text{otherwise} \end{cases} \quad (6)$$

4.3 Optimization of Data Reception

The optimal data selection probability vector \mathbf{P} is derived using an optimization formulation based on the specified reception probability Q^{spec} , receiver listening interval distribution \mathbf{R} , and packet size N . For simplicity, it is still assumed that $\delta = 1$. A generalization to arbitrary activation intervals δ is discussed in Section 5.1.

ALGORITHM 1: First Fit Decreasing

```

1 Function FirstFit( $P, N$ ):
2    $\rho_{i,j} \leftarrow 0 \quad \forall i \in \{1, \dots, N\}, j \in \{0, \dots, P^*\}$ 
3   for  $(i, P_i) \mid P_i \in P$ ; in descending order of  $P_i$  do
4     for  $j = 1$  to  $N$  do
5       if  $\sum \rho_j \leq 1 - P_i$  then
6          $\rho_{j,i} \leftarrow P_i$ ; mark  $P_i$  as assigned
7         break
8     if  $P_i$  is not marked as assigned then
9       return failure
10  return  $\rho$ 

```

The optimization problem is stated as follows:

$$\text{maximize: } \min_{0 \leq i \leq Q^{\text{spec}}} \{Q_i / Q_i^{\text{spec}}\} \quad (7)$$

$$\text{subject to: } 0 \leq P_i \leq 1 \quad \forall i \in \{0, \dots, P^*\} \quad (8)$$

$$\sum_{0 \leq i \leq P^*} P_i \leq N - \epsilon \quad (9)$$

The objective function (7) maximizes the minimum ratio between the elements of \mathbf{Q} and the corresponding element of \mathbf{Q}^{spec} . Therefore, the objective maximizes the data utility $U(\mathbf{Q}, \mathbf{Q}^{\text{spec}})$ specified in (2). The constraints specified in (8) ensure that all elements of \mathbf{P} are valid probability values. Constraint (9) ensures that the number of data units selected does not exceed the data packet size N . The parameter ϵ is used to strengthen the requirement of at most N data units per data packet. This additional slack is used in the bin packing stage that is described in the next section, see also Fig. 4.

The optimization formulation specified in equations (7) to (9) is proven to be convex (see proof in Appendix A.1). Given the convexity of the problem, it can be solved efficiently and optimally by any of the available convex optimization solvers [6]. A feasible solution is returned as long as $N - \epsilon \geq 0$.

4.4 Optimal Data Selection Policy

The optimal selection probability \mathbf{P} of data units cannot directly be used to decide whether a data unit is transmitted in a given interval, as this would lead to a variable data packet size. Therefore, a final partitioning step is added to distribute these probabilities to the N slots of a fixed size data packet. To guarantee that no duplicates are selected for transmission, First Fit bin packing (Algorithm 1) is used to determine individual probability mass functions for each of the N data slots of a data packet. The parameter ϵ in (9) is used to account for partitioning loss, so that no more than the available N data slots are used. The probability mass functions for each of the N slots of a data packet are summarized in a data selection policy matrix ρ , where rows ρ_i refer to the probability mass function for selecting the data unit of the i^{th} data slot.

The exact procedure for obtaining an optimal data selection policy ρ is described in Algorithm 2. The slack variable ϵ is initialized to zero (Line 1). This is followed by finding the optimal selection probability vector \mathbf{P} using the optimization formulation (Line 2). The optimizer returns \mathbf{P} such that $\sum_i P_i \leq N - \epsilon$. Then the algorithm attempts partitioning \mathbf{P} into N unit sized bins using First Fit Decreasing bin packing (Algorithm 1). In case the partitioning succeeds, a matrix ρ with its

ALGORITHM 2: Optimal Data Selection Policy

Input: Q^{spec}, R, N
Output: Matrix ρ with N rows of probability mass functions

```

1 for  $\epsilon = \{0, 0.01, \dots, N\}$  do
2    $P \leftarrow$  Solution of the optimization problem given in (7) to (9)
3    $X \leftarrow \text{FirstFit}(P, N)$ 
4   if  $X \neq \text{failure}$  then
5      $\rho \leftarrow X$ 
6      $\rho_{i,j} \leftarrow \frac{\rho_{i,j}}{\sum \rho_i} \quad \forall i \in \{1, \dots, N\}, j \in \{0, \dots, P^*\}$ 
7     return  $\rho$ 

```

N rows being the required N partitions is obtained. The algorithm normalizes each row of ρ to receive valid probability mass functions (Line 6) and returns the resulting ρ . If bin packing is not successful, i.e., more than N bins are required, then ϵ is incremented by an arbitrary, predefined step of 0.01 and the procedure is repeated.

At run time, the sender makes the i^{th} selection decision by randomly sampling from the resulting ρ_i (see also Fig. 4). Since, each element of P exists in exactly one probability mass function, it is guaranteed that each data unit is selected at most once per data packet. Therefore, the calculation of Q_i (the probability of receiving data of age $i \cdot \delta$) in (5) holds when packet selection is done using ρ as explained here.

After discussing a number of generalizations in the next section, the implementation of the run-time mechanisms on a real-world batteryless sensor node is discussed in Section 6.

5 MODEL GENERALIZATIONS

In the following we generalize the model to allow for different time units δ , add support for aggregated sensor values, introduce run-time harvesting awareness, and to incorporate communication over a lossy channel.

5.1 Relaxing Activation Interval

The parameter δ characterizes the interval between consecutive activations of the sender. While $\delta = 1$ was assumed in the presentation of the model, in reality it needs to be adapted to the harvested input power pwr^{in} to guarantee sustainable operation. For this reason, a mechanism for scaling the specification of the receiver behavior \tilde{R} and the requirement Q^{spec} to arbitrary δ is needed.

Remember that \tilde{Q}^{spec} and \tilde{R} are defined for the lowest time granularity δ_0 . As the interval δ is defined to be an integer multiple of δ_0 , the scaled model inputs Q^{spec} and R can be received by down-sampling the specifications.

Several aggregation options like mean, min or max could be used for \tilde{Q}^{spec} . In this work, the probability of receiving a packet at any point in time within the down-sampling interval is used, corresponding to the following equation:

$$Q_i^{\text{spec}} = 1 - \prod_{j=i \cdot \delta}^{j < (i+1) \cdot \delta} (1 - \tilde{Q}_j^{\text{spec}}) \quad (10)$$

This avoids relaxing of the specification through aggregation and smooths hard transitions for large intervals.

To receive a valid \mathbf{R} that represents a listening interval probability mass function, the intermediate probabilities of $\tilde{\mathbf{R}}$ are accumulated for each aggregation interval:

$$R_i = \sum_{j=i \cdot \delta}^{j < (i+1) \cdot \delta} \tilde{R}_j \quad (11)$$

Therefore, when specifying \mathbf{Q}^{spec} or \mathbf{R} in the model, it is implicitly assumed that these values are determined using (10) and (11) for the contextual value of the activation interval δ .

5.2 Aggregated Sensor Values

While the introduction of the model focused on transmitting aging sensor values, the model can be extended to incorporate aggregate values. More specifically, we consider lossy compression of the long-term history using the Haar wavelet transformation [26]. Among the resulting coefficients the most significant ones are transmitted, with a decreasing probability specification.

To include P^{agg} aggregate values, the selection probability vector \mathbf{P}^{agg} is introduced, where its individual components P_k^{agg} define the probability of including the k^{th} aggregate value in a data packet. The components of the corresponding reception probabilities \mathbf{Q}^{agg} are calculated analogously to the probability of receiving aging values in (5). In contrast to aging values, the calculation is time-invariant:

$$Q_k^{\text{agg}} = 1 - \sum_{j=0}^{j \leq R^*} R_j \cdot \left(1 - P_k^{\text{agg}}\right)^{j+1} \quad (12)$$

By augmenting the vectors \mathbf{P} , \mathbf{Q} and \mathbf{Q}^{spec} with the corresponding variables and specifications for the aggregated values \mathbf{P}^{agg} , \mathbf{Q}^{agg} , and $\mathbf{Q}^{\text{spec,agg}}$, the original optimization formulation given in equations (7) to (9) is extended to support aggregate values. Note that in the calculation of Q_i for aging values (5), only elements of \mathbf{P} which are not in \mathbf{P}^{agg} are used. This extension does not impact the convexity of the optimization formulation (see proof in Appendix A.2). Optimal policy generation using Algorithm 2 is therefore guaranteed. Neither the algorithms, nor the run-time selection mechanisms need to be updated to support aggregate values. The only addition needed is the actual calculation of the aggregate values at run time.

5.3 Energy Harvesting Awareness

Inherent changes in the environment lead to a time varying input power pwr^{in} . Changes in pwr^{in} demand an update of δ to satisfy condition (1) for sustainable operation and consequently a model change. To support energy harvesting aware data selection, the optimal selection policy calculation is repeated for a set input power levels in the offline optimization: using the energy model of the sensor node, the minimum δ for long-term sustainable operation under the assumption of a constant power level is computed based on (1). With the specifications $\tilde{\mathbf{Q}}^{\text{spec}}$ and $\tilde{\mathbf{R}}$ scaled accordingly, the schedule is optimized using Algorithm 2. The generated optimal selection policy $\boldsymbol{\rho}$ and the corresponding δ for the considered power levels are summarized in a policy pool and deployed on the sensor node.

For selecting the optimal policy at run time, the sensor node deploys a duty-cycle adaptation mechanism that updates δ . The adaptation and energy estimation procedures used for this are beyond the scope of this work, we refer the reader to existing run-time adaptation methods [1, 7, 21, 28]. To adapt the communication policy accordingly, the optimal selection policy $\boldsymbol{\rho}$ with the activation interval closest to the new duty-cycle is loaded from the policy pool and forwarded to the run-time packet selection.




Fig. 6. The transient Bluetooth Low Energy (BLE) sensor node is based on the CC2650 system on a chip (SoC). The SoC interfaces a real-time clock (RTC) with separate backup power domain and a ferroelectric RAM (FRAM) to provide timekeeping capabilities and data retention across power failures. For batteryless operation, the node relies on the Energy Management Unit (EMU) and energy burst principle introduced in [14] and a solar panel as energy source.

6 SENSOR NODE IMPLEMENTATION

For testing the proposed communication scheme in a real-world sensing scenario, we designed a custom batteryless Bluetooth Low Energy (BLE) sensor node. Any BLE enabled Android smartphone can receive and decode the sent data using our receiver application.

The sensor node hardware design, firmware and Android application sources are released to the public as part of this publication¹.

6.1 Transient Sensor Node

Being a transient sensor node, the energy buffer is dimensioned to only support the minimum atomic operation that guarantees progress, therefore belonging to the most restricted share of batteryless systems. The execution is fully dictated by the environment, as the device accumulates ambient energy in a small buffer until enough energy is available for the next step of the application.

Hardware Platform The transient BLE sensor node is based on the Texas Instruments CC2650 system on a chip (SoC). Its architecture is detailed in the diagram in Fig. 6. For batteryless operation, the energy management circuitry presented in [14] is integrated with a small energy buffer of $C_{\text{buf}} = 200 \mu\text{F}$ and configured with a constant energy burst size. The sensor node is deployed with an AM-5412 solar panel of $50 \text{ mm} \times 33 \text{ mm}$ size covering the back of the node to harvest energy from ambient light. A backup power domain for the external AM0815 real-time clock (RTC) enables persistent timekeeping across several hours of energy unavailability using a backup capacitor of $C_{\text{bak}} = 320 \mu\text{F}$. The buffer of this backup domain is recharged as part of an activation of the sensor node. An external FM25V10 ferroelectric RAM (FRAM) memory is included for energy efficient storing of the sensor data history, aggregate values, and system state. For sensing the ambient temperature and humidity an SHT31 sensor is integrated on the node platform. The transient sensor node was implemented as a custom printed circuit board (PCB) that is shown in Fig. 7.

Software The application logic is implemented on top of the Contiki operating system [12] that comes with the required radio drivers to send standardized BLE data packets. Driver support for interfacing the Energy Management Unit (EMU), sensor, timer and FRAM was added to the operating system and the start-up sequence was optimized for fast and energy efficient start-up. The following sequence of operations are performed by the sensor node during one activation (see also Fig. 4):

- (1) Restore the system and data buffer states from the FRAM,

¹available at: https://gitlab.ethz.ch/tcc/public/employees/sigristl/transient_ble_node




Fig. 7. The custom designed transient sensor node used to implement and evaluate the communication scheme has a small size of only $35\text{ mm} \times 53\text{ mm}$, about the size of the match-box shown in the background. It is powered from the solar panel mounted on the back of the node.

- (2) Increase the system voltage to 3.3 V and start recharging the timer backup domain's capacitor C_{bak} ,
- (3) Readout the current time and sensor data, and store them as a new data unit in the FRAM data history,
- (4) Once a day, compress the accumulated 6 min aggregates using the Haar wavelet transform and store the 10 most relevant coefficients as aggregate values,
- (5) Estimate the time-based power level to load the corresponding optimal policy ρ from the policy pool,
- (6) Sample probabilistically according to ρ to get the selection of data units to include in the next data packet,
- (7) Load the selected data units from the FRAM, assemble them in a BLE packet, and broadcast it on all BLE advertisement channels (CH37-CH39),
- (8) Back up data buffer and system state and enter deep sleep at the minimal system voltage of 2.3 V .

In cases of extended periods of energy unavailability, the RTC backup domain eventually runs out of energy. This means that the timer is reset as soon as its supply is restored as part of the next successful sensor node activation. The loss of the time reference irrevocably invalidates the timestamps of the locally buffered sensor data, therefore demanding to clear buffered data units under these circumstances.

6.2 Smartphone Receivers

Commodity smartphones, Samsung Galaxy A3 (2016), were used as receiver devices for data collection. Our application scans for BLE advertisement packets in the background and logs timestamped packet data to internal memory. The collected traces are later post-processed for the performance analysis detailed in the following section.

7 EXPERIMENTAL EVALUATION

The experimental evaluation demonstrates the feasibility of implementing our optimized communication scheme for infrastructure-less sensing using batteryless sensor nodes. The scheme is realizable in an application with a small energy footprint (Section 7.2) and scalable for dense deployment of a high number of independent sensor beacons and receivers (Section 7.3). By comparison of the simulation and experimental results in Section 7.4, it is confirmed that the model abstracts




Fig. 8. The considered problem specifications consist of \tilde{Q}^{spec} specifying the relative importance of sensor data, and a stochastic model \tilde{R} of the receivers' listening interval.

the characteristics of the considered scenario with high accuracy. The optimized communication scheme leads to a significant improvement in terms of data utility as is shown in Section 7.5. Finally, the real-world deployments presented in Section 7.6 demonstrate the long-term autonomous operation of sensor beacons and the harvesting-aware adaptation of the communication policy in two different scenarios.

7.1 Evaluation Setups

We begin with the introduction of the setups used for all experiments and simulations. Each experiment or simulation has two different aspects: Firstly, a problem specification defining the required reception probability \tilde{Q}^{spec} and the receiver model \tilde{R} . Secondly, a setup which specifies how a given evaluation is conducted.

In the evaluation we focus on the problem specification $(\tilde{Q}^{\text{spec}}, \tilde{R})$ consisting of the \tilde{Q}^{spec} specification for monitoring and the \tilde{R} characteristics of an office deployment, both shown in Fig. 8, and a data packet size of $N = 4$ matching the sensor node implementation. It should however be noted that the proposed scheme is very flexible and supports arbitrary specifications of \tilde{Q}^{spec} , \tilde{R} , and N . To demonstrate this flexibility, an additional problem specification consisting of \tilde{Q}^{spec} for feedback controlled applications and a receiver characteristic \tilde{R} abstracting a kitchen scenario is evaluated under real-world conditions. In the considered problem specifications, \tilde{Q}^{spec} specifies the importance of the 200 considered data values as a function of their age. The additional 10 aggregates represent the 10 most significant wavelet coefficients which are used for reconstruction of long-term temperature/humidity time series. In the office scenario, the receiver active times \tilde{R} follow a Poisson distribution with arrival rate $\lambda = 30$ s. For the kitchen scenario, a superposition of two Poisson distributions with arrival rates of $\lambda = 5$ s and $\lambda = 30$ s are considered. Based on this problem specification the optimized communication policies are generated for 6 power levels that correspond to one sensor node activation every $\{1, 2, 3, 4, 5, 6\}$ s.

In addition to the optimal policy, we consider three baseline scenarios: a) a deterministic baseline transmitting the most recent values, i.e., not considering \tilde{Q}^{spec} nor \tilde{R} , b) a policy optimization only considering \tilde{R} , without distinguishing the importance of different sensor values, and c) a policy optimization in which only the relative importance \tilde{Q}^{spec} is considered. These four schemes are compared to investigate the importance of the model inputs required for policy optimization.

The above problem specifications are evaluated in the following setups:

Controlled Power Level Emulation In the first set of experiments the sensor node is supplied continuously from a power supply, bypassing the integrated Energy Management Unit (EMU). The node is triggered periodically from a signal generator emulating an constant harvesting power. Upon triggering, the sensor node performs all actions part of a single activation (sense values, process and store values, transmit values, see the detailed procedure in Section 6.1). The activation period is set in accordance with the stated power levels. This experiment lasts for 4 h at each power level.

Energy Harvesting-Driven Office Scenario In this experiment the sensor node is deployed with the EMU and a solar panel to operate fully autonomously. The real-world experiment is carried out in an office environment with the sensor node placed on a table, 0.5 m away from a window. The solar panel of the transient sensor node faces towards the ceiling and is exposed to a mixture of artificial office lighting and indirect natural light. A RocketLogger measurement device [37] instruments the sensor node to trace the sensor node activations by observing the EMU trigger signal and to monitor the ambient illuminance level next to the solar panel. Illuminance levels varying from 300 lx to 800 lx at daytime and a few shorter intervals of several minutes with illuminance levels up to 1400 lx were observed during the 45 h long experiment.

Energy Harvesting-Driven Kitchen Scenario In a second real-world harvesting experiment the sensor node is deployed at the wall in a kitchen, 0.3 m away from a window. Again, the sensor node is instrumented to trace its run-time behavior and a smartphone receiver is placed in a kitchen drawer 3 m away. The illuminance levels observed during the 9 day long experiment varied between 150 lx to 1 000 lx at daytime, and include 15 min of direct exposure to sunlight of up to 20 klx in the late afternoon.

Model Based Scenario Simulation The performance in terms of data utility $U(Q, Q^{\text{spec}})$ is evaluated in different scenarios using the model based simulation. It consists of the comparison of the optimized communication policy to three baseline scenarios, and a model sensitivity analysis with respect to the receiver characteristics \tilde{R} .

In the experimental setups, an always-on smartphone is placed near the sensors and continuously logs the received data packets. The intermittent receiver behavior characterized by \tilde{R} is introduced in the post-processing, by only considering the packets received within the time window where the specific intermittent receivers are active. This analysis was performed independently for each power level, and we randomly generate 4000 intermittent receivers for each level. The time at which they start listening is chosen from a uniform random distribution and their listening interval is sampled from \tilde{R} .

In the two *energy harvesting-driven* experiments, the power level is dictated by the environment. To perform the analysis in these setups, we identify time windows in which the average power remains relatively constant. Specifically, when performing the analysis for a power level with an activation interval δ , we consider 60 s periods where the average activation interval of the sensor node δ_{avg} is within $\delta \pm 0.1$. The intermittent receivers are activated such that their listening intervals are contained in the aforementioned time windows of relatively constant power.

7.2 Sensor Node Characterization

To calibrate the energy model, the energy required for one task activation, $E^{\text{activation}}$, is needed. The implemented sensor node was characterized using the *controlled power level emulation* setup. The resulting energy model is then used in the offline policy pool generation, as well as at run time for input power estimation. The RocketLogger measurement device [37] was used to characterize the energy consumption of a single activation (at 3.3 V), as well as the system's deep sleep power (at

Table 1. Transient Bluetooth Low Energy (BLE) sensor node operation characteristics in terms of energy and time: mean of 12 819 activations, including the 95 % percentile of the absolute deviation from the mean.

| | Energy/Power | Duration |
|-----------------------------------|------------------|-----------------|
| <i>Activation Total</i> (3.3 V) | 162.90 ± 1.01 μJ | 12.62 ± 0.14 ms |
| <i>Deep Sleep</i> (2.3 V) | 3.91 μW | - |
| <i>Run-Time Mechanism</i> (3.3 V) | 8.08 ± 0.12 μJ | 0.73 ± 0.01 ms |

2.3 V). To collect a significant number of samples the external triggering was kept constant at short periods. The results for more than 12 000 activations are summarized in Table 1 and demonstrate a very small energy requirement of only $E^{\text{activation}} \approx 163 \mu\text{J}$ per activation. Only a small amount thereof, i.e., $8 \mu\text{J}$, is needed for the communication policy run-time mechanism, demonstrating the low overhead of the presented communication scheme.

7.3 Model-Based Scalability Analysis

For the scalability analysis of many transmitting sensor beacons we consider M devices within communication range of each other. The wireless channel accesses are assumed random, as each sensor beacon independently harvests energy from the non-deterministic and spatially variable environment. With a time on air of 0.376 ms for a full BLE advertisement packet [36] and a sensor activation period of 10 s, the collision probability of the considered pure ALOAH communication scheme with M sensor beacons operating in parallel is calculated as:

$$\text{Pr}_{\text{coll}} = 1 - (1 - q)^{(M-1)} = 1 - \left(1 - 2 \cdot \frac{0.376 \text{ ms}}{10 \text{ s}}\right)^{(M-1)} \quad (13)$$

The beacons broadcast a data packet on all BLE advertisement channels to increase the chance of reception for scanning receiver devices. The collision probability Pr_{coll} is still calculated using (13), as the deterministic advertisement channel hopping sequence guarantees no collision on any channel if there is no collision on the first advertisement channel. For a dense deployment of $M = 100$ beacons within the restricted communication range, equation (13) evaluates to a collision probability of less than 0.8 %. Collisions with other sensor beacons can therefore be considered minimal when compared to the higher probability of interference from other sources like WiFi or Bluetooth communication [18]. The scalability of BLE advertisement based communication has also been confirmed experimentally in [30]. Reliable communication was demonstrated in a dense deployment with more than 200 sensor nodes within the range of a single receiver.

The fully passive receiver behavior guarantees by design a scalability to an unbounded number of receiver devices.

7.4 Model Accuracy

To validate the model accuracy, the *controlled power level emulation* setup described in Section 7.1 was used. The reception probabilities Q for the power level with $\delta = 1$ are shown in Fig. 9. This figure compares the ideal (Q lossless), the measured (Q experiment), and the packet loss compensated reception probability (Q lossy) with the scaled specification $\alpha \cdot Q^{\text{spec}}$. The scaling factor α relates to the maximum utility that resulted from the offline schedule optimization as specified in (3). The first Q refers to the theoretical value obtained by solving the optimization formulation introduced in Section 4. Comparing it to $\alpha \cdot Q^{\text{spec}}$ shows that the optimization formulation in (7) to (9) finds a correct packet selection policy ρ that results in a data reception probability Q that maximizes




Fig. 9. The model accuracy experiment results with $\delta = 1$ demonstrate a high model accuracy.

Table 2. The model error analysis for the model accuracy experiment shows very small errors for all power levels pwr^{in} .

| δ | Power pwr^{in} | α | Max. Error | RMSE |
|----------|--------------------------------|----------|------------|--------|
| 1 | 163 μW | 0.77 | 0.0481 | 0.0132 |
| 2 | 81.5 μW | 0.58 | 0.0420 | 0.0113 |
| 3 | 54.33 μW | 0.53 | 0.0419 | 0.0132 |
| 4 | 40.75 μW | 0.5 | 0.0388 | 0.0157 |
| 5 | 32.6 μW | 0.47 | 0.0341 | 0.0121 |

the data utility as specified in (3). The experimental results for Q are slightly lower than the model, because the model assumes a lossless communication channel. In the actual deployment the sensor nodes were placed in an office environment, where the frequency channel is heavily occupied. Augmenting the model to account for a time-invariant packet reception rate PRR and using the empirically observed value (91.8% for the shown power level), results in Q lossy. With this compensation factor the theoretical and the experimentally observed reception probabilities match. The model error analysis for all power levels is summarized in Table 2. The table states the maximum absolute error and Root-Mean-Square Error (RMSE) between the measured and packet loss compensated Q for all power levels. The consistently low errors across all power levels demonstrate the correctness of our model introduced in Section 3.2, and the correct implementation of the communication scheme.

7.5 Data Utility Analysis

We perform a sensitivity analysis of the communication policy with respect to the receiver's listening characteristic \bar{R} . Furthermore, we compare the proposed communication policy to three baseline policies. Both evaluations use the *model based scenario simulation* setup and do not consider aggregate values, since not all baseline policies support sending aggregate values.

To perform the sensitivity analysis, we examine the impact on data utility when the actual receiver behavior differs from the model used during optimization. We generate the optimal communication policies for 200 additional receiver models and calculate their data utility. In these models the receiver active times follow a Poisson distribution with arrival rates $\lambda \in \{0, 1, \dots, 200\}$.

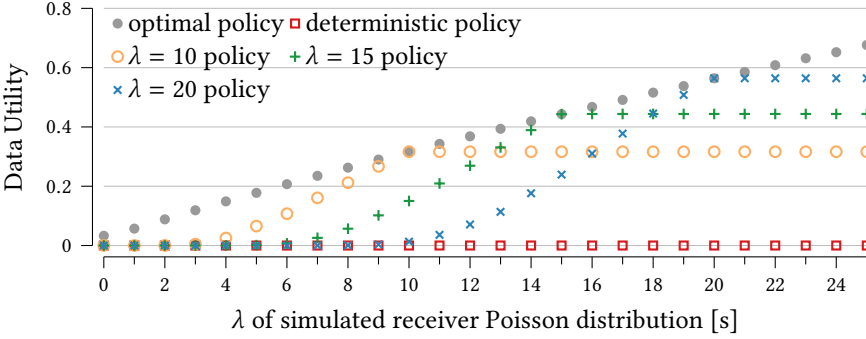


Fig. 10. The data utility provided by policies optimized for specific values of λ and the deterministic baseline scheme in comparison to the optimally achievable data utility.

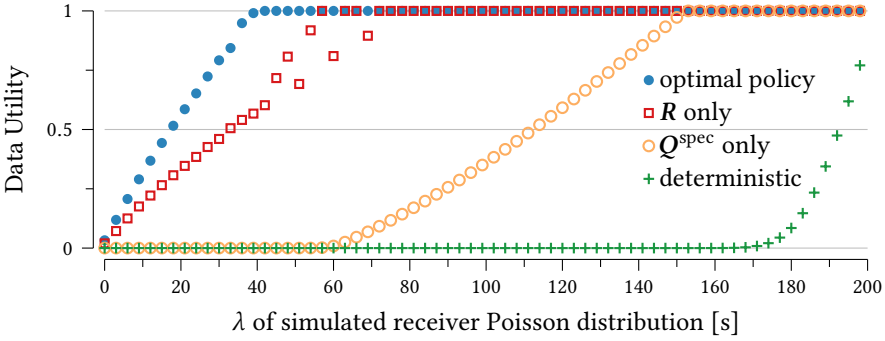


Fig. 11. Comparison of the data utility of the optimal scheme with three baseline policies that consider only a subset of the problem specification.

The results of the sensitivity analysis are given in Fig. 10. The *optimal policy* shows the utility when the actual receiver behavior (the specific value of λ) is identical to the value for which the optimal communication policy is generated. This represents the maximum achievable utility, which requires precise knowledge of the receiver behavior. The *policy for $\lambda = 15$* , henceforth called *policy 15* shows the data utility for different λ of the simulated receiver when the optimal communication policy is generated for $\lambda = 15$. The data utility of policy 15 is significantly lower than the optimum for $\lambda < 15$. This indicates that the data utility is fairly sensitive to the receiver characteristics. Therefore, it is important to optimize the policy for the specific receiver behavior. For $\lambda = 15$, policy 15 provides the optimal utility, since there is no mismatch between the actual receiver behavior and the one used for generating the optimal communication policy. For $\lambda > 15$, we see that the data utility remains unchanged. Even though the receiver is listening for longer, additional gains in utility cannot be observed. This is explained by the fact that the value of the data utility is dictated by the most recent sensor values (see (2)). The reception probability of these values do not improve if the receiver is active for longer. Figure 10 also shows the data utility for communication policies optimized with $\lambda = 10, 20$. The data utility of these policies follow a similar trend as policy 15, while reaching the optimal data utility at their respective optimization points.

For comparison we perform the data utility analysis with three baseline policies that consider Q^{spec} only, R only, and none of them, i.e. transmitting the most recent values only. The resulting data

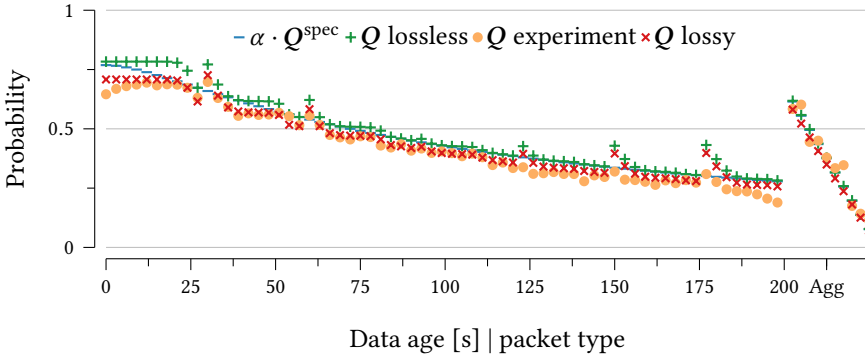


Fig. 12. Real-world harvesting office use-case results with $\delta = 1$: high model accuracy despite additional sources of error.

Table 3. The model errors analysis for real-world harvesting use-case experiment in an office deployment with illuminance levels varying from 300 lx to 800 lx demonstrates high model accuracy. The statistically not significant results for the $\delta = 2$ were omitted.

| δ | Power p_{wr}^{in} | Illuminance | α | Max. Error | RMSE | Duration |
|----------|---------------------|-------------|----------|------------|--------|----------|
| 1 | 163 μ W | 775 lx | 0.77 | 0.1099 | 0.0306 | 2028 s |
| 2 | 81.5 μ W | NaN lx | NaN | NaN | NaN | 35 s |
| 3 | 54.33 μ W | 459 lx | 0.53 | 0.0875 | 0.0334 | 1926 s |
| 4 | 40.75 μ W | 412 lx | 0.5 | 0.0970 | 0.0372 | 4539 s |
| 5 | 32.6 μ W | 391 lx | 0.47 | 0.1099 | 0.0481 | 1896 s |
| 6 | 27.17 μ W | 378 lx | 0.48 | 0.1457 | 0.0610 | 1087 s |

utility values for all four policies are shown in Fig. 11, again as a function of the λ of the simulated receiver. For low values of λ the data utility of the deterministic and Q^{spec} only baselines remains at 0. However, as λ increases the data utility eventually increases and reaches the optimal utility: for the deterministic baseline the utility increases as the receivers' listening interval approaches the considered history length of 200 s. For the Q^{spec} only baseline an increasing data utility is observed for considerably shorter listening intervals. This demonstrates the necessity of specifying the relative importance of the desired sensor values. Furthermore, the results for the R only baseline demonstrate that optimizing for the specific receiver behavior results in a considerable data utility increase already at low values of λ . However, not considering the relative importance Q^{spec} of sensor values can lead to significant variability in data utility, as is visible for values of λ between 40 and 70. This is because the optimization objective (7) relies on full knowledge of the Q^{spec} parameter. For the R only baseline Q^{spec} was assumed uniform, i.e. all sensor values were considered equally important. Considering both specifications, Q^{spec} and R , allows deriving the optimal policy presented in this work. It is important to note that the utility achieved by the optimal communication policy is no less than the one provided by the considered baseline policies for any value of λ . This substantiates the optimality of the proposed communication scheme.

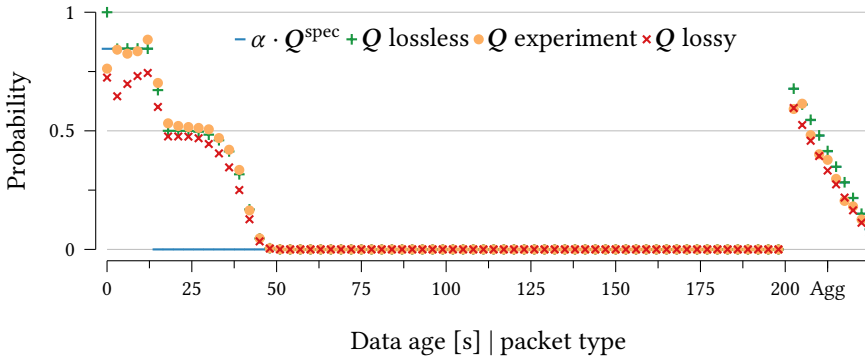


Fig. 13. Real-world harvesting results for the kitchen scenario with a feedback specification with $\delta = 3$: high model accuracy despite additional sources of error.

Table 4. The model errors analysis for the kitchen use-case experiment in the kitchen scenario with illuminance levels varying from 150 lx to 1 000 lx.

| δ | Power pwr^{in} | Illuminance | α | Max. Error | RMSE | Duration |
|----------|--------------------------------|-------------|----------|------------|--------|----------|
| 1 | 163 μW | 681 lx | 0.81 | 0.3463 | 0.0893 | 46522 s |
| 2 | 81.5 μW | 489 lx | 0.85 | 0.1097 | 0.0265 | 3375 s |
| 3 | 54.33 μW | 421 lx | 0.85 | 0.1966 | 0.0426 | 1250 s |
| 4 | 40.75 μW | 277 lx | 0.81 | 0.1672 | 0.0564 | 1165 s |
| 5 | 32.6 μW | 268 lx | 0.85 | 0.1467 | 0.0514 | 918 s |
| 6 | 27.17 μW | 250 lx | 0.83 | 0.1769 | 0.0600 | 890 s |

7.6 Real-World Harvesting Use-Case

To evaluate the model accuracy under real-world harvesting conditions we used the two *energy harvesting-driven* setups in an office and a kitchen environment, as detailed in Section 7.1.

Office Scenario with Monitoring Specification The resulting reception probabilities Q are presented in Fig. 12 for the power level with $\delta = 1$. The error observed here is slightly higher when compared to the model accuracy experiment in Section 7.4. This is expected, because the activation interval of the sensor node is not controlled, but variable and dictated by the environment. Small deviations in sensor node activation intervals are the primary cause of the higher errors in the data reception probabilities. However, the results still show a high model accuracy.

Table 3 summarizes the model error analysis for the packet loss compensated model at different power levels. The results for $\delta = 2$ are omitted, as the environment conditions in the experiment did not allow the sensor node to operate at this power level for long enough time. Therefore, no statistically significant results could be extracted using the analysis method described in Section 7.1. The table shows low errors for all power levels with a maximum RMSE of 0.061 for the $\delta = 6$ power level. The high model accuracy for all power levels demonstrates that the sensor node is able to dynamically adapt the communication policy to changing harvesting conditions.

Kitchen Scenario with Feedback Specification Similar to the above scenario, Table 4 summarizes errors of the packet loss compensated model for all power levels, and Fig. 13 shows

a representative example of the reception probabilities \mathbf{Q} for the power level with $\delta = 3$. With RMSE errors comparable to those of the office scenario, the results confirm the flexible applicability of the proposed communication scheme to various application scenarios. The higher error numbers observed for the $\delta = 1$ power level are most likely attributed to the abrupt change in \tilde{Q}^{spec} (step from 1 to 0 at $\delta = 15$, as shown in Fig. 8) in combination with the minimal timestamp resolution of 1 s of the sensor data. This can result in increased model errors for sensor values with an age around the sharp drop of \tilde{Q}^{spec} at $\delta = 15$.

Overall, these results confirm the accuracy and flexibility of the model and the adaptability of the presented communication scheme under varying harvesting conditions.

8 CONCLUSION

In this work, we motivate the use of infrastructure-less sensor beacons for autonomous long-term sensing. In the considered scenario, receivers are mobile and therefore intermittently within the communication range of sensor beacons. To model the receiver mobility we introduce a stochastic model. Based on this model and a specification of the relative importance of data, we derive an optimized communication scheme. The scheme is implemented and evaluated on a custom-designed batteryless BLE sensor node powered by ambient light. The results show that the sensor beacon architecture is realizable and can sense the ambient conditions at a frequency of up to 1 Hz under office harvesting conditions at 775 lx. In addition, we show that the proposed communication scheme can be implemented in resource constrained embedded platforms. Results from the real-world deployment match the theoretically expected values in all considered setups. A multi-day harvesting experiment under indoor office light conditions demonstrates that the proposed communication scheme adapts well to the changing harvesting conditions and closely matches the specification.

ACKNOWLEDGMENT

The authors would like to thank Alfonso Blanco for the sensor node PCB design, Oliver Brunecker for sensor node prototyping, and the anonymous reviewers for their valuable feedback.

This research was funded by the Swiss National Science Foundation under grant 157048: Transient Computing Systems.

REFERENCES

- [1] Rehan Ahmed, Bernhard Buchli, Stefan Draskovic, Lukas Sigrist, Pratyush Kumar, and Lothar Thiele. 2019. Optimal Power Management with Guaranteed Minimum Energy Utilization for Solar Energy Harvesting Systems. *ACM Transactions on Embedded Computing Systems* 18, 4 (June 2019), 1–26. <https://doi.org/10.1145/3317679>
- [2] D. Balsamo, A. S. Weddell, A. Das, A. R. Arreola, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. 2016. Hibernus++: A Self-Calibrating and Adaptive System for Transiently-Powered Embedded Devices. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 35, 12 (2016), 1968–1980. <https://doi.org/10.1109/TCAD.2016.2547919>
- [3] Naveed Anwar Bhatti, Muhammad Hamad Alizai, Affan A Syed, and Luca Mottola. 2016. Energy Harvesting and Wireless Transfer in Sensor Network Applications: Concepts and Experiences. *ACM Transactions on Sensor Networks* 12, 3 (Aug. 2016), 24:1–24:40. <https://doi.org/10.1145/2915918>
- [4] Naveed Anwar Bhatti and Luca Mottola. 2017. HarvOS: Efficient Code Instrumentation for Transiently-Powered Embedded Sensing. In *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 209–219. <https://doi.org/10.1145/3055031.3055082>
- [5] Davis Blalock, Samuel Madden, and John Guttag. 2018. Sprintz: Time Series Compression for the Internet of Things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (Sept. 2018), 1–23. <https://doi.org/10.1145/3264903>
- [6] Stephen Boyd and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9780511804441>

- [7] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. 2014. Towards Enabling Uninterrupted Long-Term Operation of Solar Energy Harvesting Embedded Systems. In *Proceedings of the 11th European Conference on Wireless Sensor Networks (EWSN 2014)*. Springer, 66–83. https://doi.org/10.1007/978-3-319-04651-8_5
- [8] N. Bulusu, J. Heidemann, and D. Estrin. Oct./2000. GPS-Less Low-Cost Outdoor Localization for Very Small Devices. *IEEE Personal Communications* 7, 5 (Oct./2000), 28–34. <https://doi.org/10.1109/98.878533>
- [9] E.J. Candes and M.B. Wakin. 2008. An Introduction To Compressive Sampling. *IEEE Signal Processing Magazine* 25, 2 (March 2008), 21–30. <https://doi.org/10.1109/MSP.2007.914731>
- [10] Alexei Colin, Emily Ruppel, and Brandon Lucia. 2018. A Reconfigurable Energy Storage Architecture for Energy-Harvesting Devices. In *Proceedings of the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) (ASPLOS '18)*. ACM, 767–781. <https://doi.org/10.1145/3173162.3173210>
- [11] Y. Collet and M. Kucherawy. 2018. *Zstandard Compression and the Application/Zstd Media Type*. Technical Report RFC8478. RFC Editor. <https://doi.org/10.17487/RFC8478>
- [12] A Dunkels, B Gronvall, and T Voigt. 2004. Contiki - a Lightweight and Flexible Operating System for Tiny Networked Sensors. In *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks*. 455–462. <https://doi.org/10.1109/LCN.2004.38>
- [13] Songtao Feng and Jing Yang. 2018. Optimal Status Updating for an Energy Harvesting Sensor with a Noisy Channel. In *Proceedings of the 2018 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, Honolulu, HI, 348–353. <https://doi.org/10.1109/INFOCOMW.2018.8406974>
- [14] Andres Gomez, Lukas Sigrist, Thomas Schalch, Luca Benini, and Lothar Thiele. 2017. Efficient, Long-Term Logging of Rich Data Sensors Using Transient Sensor Nodes. *ACM Transactions on Embedded Computing Systems* 17, 1 (Sept. 2017), 4:1–4:23. <https://doi.org/10.1145/3047499>
- [15] Maria Gorlatova, Aya Wallwater, and Gil Zussman. 2013. Networking Low-Power Energy Harvesting Devices: Measurements and Algorithms. *IEEE Transactions on Mobile Computing* 12, 9 (Sept. 2013), 1853–1865. <https://doi.org/10.1109/TMC.2012.154>
- [16] Josiah Hester, Kevin Storer, and Jacob Sorber. 2017. Timely Execution on Intermittently Powered Batteryless Sensors. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems (SenSys)*. ACM, 1–13. <https://doi.org/10.1145/3131672.3131673>
- [17] Josiah Hester, Nicole Tobias, Amir Rahmati, Lanny Sitanayah, Daniel Holcomb, Kevin Fu, Wayne P Bursleson, and Jacob Sorber. 2016. Persistent Clocks for Batteryless Sensing Devices. *ACM Transactions on Embedded Computing Systems* 15, 4 (Aug. 2016), 77:1–77:28. <https://doi.org/10.1145/2903140>
- [18] Anwar Hithnawi, Hossein Shafagh, and Simon Duquennoy. 2015. TIIM: Technology-Independent Interference Mitigation for Low-Power Wireless Networks. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks (IPSN)*. ACM, 1–12. <https://doi.org/10.1145/2737095.2737104>
- [19] Hrishikesh Jayakumar, Arnab Raha, Woo Suk Lee, and Vijay Raghunathan. 2015. QuickRecall: A HW/SW Approach for Computing Across Power Cycles in Transiently Powered Computers. *ACM Journal on Emerging Technologies in Computing Systems* 12, 1 (Aug. 2015), 1–19. <https://doi.org/10.1145/2700249>
- [20] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li Shiuian Peh, and Daniel Rubenstein. 2002. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. *ACM SIGARCH Computer Architecture News* 30, 5 (Dec. 2002), 96–107. <https://doi.org/10.1145/635506.605408>
- [21] Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B. Srivastava. 2007. Power Management in Energy Harvesting Sensor Networks. *ACM Transactions on Embedded Computing Systems* 6, 4 (Sept. 2007), 32–es. <https://doi.org/10.1145/1274858.1274870>
- [22] J E Kim, T Abdelzaher, L Sha, A Bar-Noy, and R Hobbs. 2016. Sporadic Decision-Centric Data Scheduling with Normally-Off Sensors. In *Proceedings of the 2016 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 135–145. <https://doi.org/10.1109/RTSS.2016.022>
- [23] Giang Truong Le, Thang Viet Tran, Hyeon-Sock Lee, and Wan-Young Chung. 2016. Long-Range Batteryless RF Sensor for Monitoring the Freshness of Packaged Vegetables. *Sensors and Actuators A: Physical* 237 (Jan. 2016), 20–28. <https://doi.org/10.1016/j.sna.2015.11.013>
- [24] Mengjuan Liu, Yan Yang, and Zhiguang Qin. 2011. A Survey of Routing Protocols and Simulations in Delay-Tolerant Networks. In *Wireless Algorithms, Systems, and Applications*, Yu Cheng, Do Young Eun, Zhiguang Qin, Min Song, and Kai Xing (Eds.). Vol. 6843. Springer, Berlin, Heidelberg, 243–253. https://doi.org/10.1007/978-3-642-23490-3_22
- [25] Robert Margolies, Peter Kinget, Ioannis Kymissis, Gil Zussman, Maria Gorlatova, John Sarik, Gerald Stanje, Jianxun Zhu, Paul Miller, Marcin Szczodrak, Baradwaj Vigraham, and Luca Carloni. 2015. Energy-Harvesting Active Networked Tags (EnHANTs): Prototyping and Experimentation. *ACM Transactions on Sensor Networks* 11, 4 (Nov. 2015), 1–27. <https://doi.org/10.1145/2831236>
- [26] R J E Merry. 2005. *Wavelet Theory and Applications - A Literature Study*. Technical Report DCT 2005.53. Eindhoven University of Technology, Eindhoven. 49 pages.

- [27] Masateru Minami, Takashi Morito, Hiroyuki Morikawa, and Tomonori Aoyama. 2005. Solar Biscuit: A Battery-Less Wireless Sensor Network System for Environmental Monitoring Applications. In *Proceedings of the 2nd International Workshop on Networked Sensing Systems (INSS)*. Citeseer, 2007.
- [28] C Moser, L Thiele, D Brunelli, and L Benini. 2010. Adaptive Power Management for Environmentally Powered Systems. *IEEE Trans. Comput.* 59, 4 (April 2010), 478–491. <https://doi.org/10.1109/TC.2009.158>
- [29] Saman Naderiparizi, Aaron N Parks, Zerina Kapetanovic, Benjamin Ransford, and Joshua R Smith. 2015. WISPCam: A Battery-Free RFID Camera. In *Proceedings of the 2015 IEEE International Conference on RFID (RFID)*. IEEE, 166–173. <https://doi.org/10.1109/RFID.2015.7113088>
- [30] Maciej Nikodem and Marek Bawiec. 2019. Experimental Evaluation of Advertisement-Based Bluetooth Low Energy Communication. *Sensors* 20, 1 (Dec. 2019), 107. <https://doi.org/10.3390/s20010107>
- [31] Paritosh Padhy, Rajdeep K. Dash, Kirk Martinez, and Nicholas R. Jennings. 2010. A Utility-Based Adaptive Sensing and Multihop Communication Protocol for Wireless Sensor Networks. *ACM Transactions on Sensor Networks* 6, 3 (June 2010), 1–39. <https://doi.org/10.1145/1754414.1754423>
- [32] Majharul Islam Rajib and Asis Nasipuri. 2017. Predictive Retransmissions for Intermittently Connected Sensor Networks with Transmission Diversity. *ACM Transactions on Embedded Computing Systems* 17, 1 (Sept. 2017), 1–25. <https://doi.org/10.1145/3092947>
- [33] Benjamin Ransford, Jacob Sorber, and Kevin Fu. 2011. Mementos: System Support for Long-Running Computation on RFID-Scale Devices. In *Proceedings of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, 159–170. <https://doi.org/10.1145/1950365.1950386>
- [34] M. A. Razzaque, Chris Bleakley, and Simon Dobson. 2013. Compression in Wireless Sensor Networks: A Survey and Comparative Evaluation. *ACM Transactions on Sensor Networks* 10, 1 (Nov. 2013), 5:1–5:44. <https://doi.org/10.1145/2528948>
- [35] Christopher M. Sadler and Margaret Martonosi. 2006. Data Compression Algorithms for Energy-Constrained Devices in Delay Tolerant Networks. In *Proceedings of the 4th ACM Conference on Embedded Network Sensor Systems (SenSys)*. ACM, 265–278. <https://doi.org/10.1145/1182807.1182834>
- [36] Gaoyang Shan and Byeong-Hee Roh. 2018. Advertisement Interval to Minimize Discovery Time of Whole BLE Advertisers. *IEEE Access* 6 (2018), 17817–17825. <https://doi.org/10.1109/ACCESS.2018.2817343>
- [37] Lukas Sigrist, Andres Gomez, Roman Lim, Stefan Lippuner, Matthias Leubin, and Lothar Thiele. 2017. Measurement and Validation of Energy Harvesting IoT Devices. In *Proceedings of the 2017 Design, Automation Test in Europe Conference and Exhibition (DATE)*. IEEE, Lausanne, Switzerland, 1159–1164. <https://doi.org/10.23919/DATE.2017.7927164>
- [38] Jacob Sorber, Aruna Balasubramanian, Mark D. Corner, Joshua R. Ennen, and Carl Qualls. 2013. Tula: Balancing Energy for Sensing and Communication in a Perpetual Mobile System. *IEEE Transactions on Mobile Computing* 12, 4 (April 2013), 804–816. <https://doi.org/10.1109/TMC.2012.52>
- [39] Thrasyvoulos Spyropoulos, Konstantinos Psounis, and Cauligi S. Raghavendra. 2005. Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks. In *Proceedings of the 2005 ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN)*. ACM, 252–259. <https://doi.org/10.1145/1080139.1080143>
- [40] Sennur Ulukus, Aylin Yener, Elza Erkip, Osvaldo Simeone, Michele Zorzi, Pulkit Grover, and Kaibin Huang. 2015. Energy Harvesting Wireless Communications: A Review of Recent Advances. *IEEE Journal on Selected Areas in Communications* 33, 3 (March 2015), 360–381. <https://doi.org/10.1109/JSA.2015.2391531>
- [41] Ricklef Wohlers, Niki Trigoni, Rui Zhang, and Stephen Ellwood. 2009. TwinRoute: Energy-Efficient Data Collection in Fixed Sensor Networks with Mobile Sinks. In *Proceedings of the 10th International Conference on Mobile Data Management: Systems, Services and Middleware*. IEEE, 192–201. <https://doi.org/10.1109/MDM.2009.30>
- [42] Sheng Yu, Baoxian Zhang, Cheng Li, and Hussein Mouftah. 2014. Routing Protocols for Wireless Sensor Networks with Mobile Sinks: A Survey. *IEEE Communications Magazine* 52, 7 (July 2014), 150–157. <https://doi.org/10.1109/MCOM.2014.6852097>
- [43] Jianhui Zhang, Zhi Li, and Shaojie Tang. 2016. Value of Information Aware Opportunistic Duty Cycling in Solar Harvesting Sensor Networks. *IEEE Transactions on Industrial Informatics* 12, 1 (Feb. 2016), 348–360. <https://doi.org/10.1109/TII.2015.2508745>
- [44] Pengyu Zhang, Deepak Ganesan, and Boyan Lu. 2013. QuarkOS: Pushing the Operating Limits of Micro-Powered Sensors. In *Proceedings of the 14th USENIX Conference on Hot Topics in Operating Systems*. USENIX.

A OPTIMIZATION FORMULATION CONVEXITY

Here, we prove that the optimization formulation proposed in (7) to (9) and its extension with aggregate values in Section 5.2 is convex in P . We will refer to Q_i as $Q_i(P, R)$ to explicitly show that it is a function of P and R .

A.1 Convexity of Optimization Formulation

The optimization formulation specified in equations (7) to (9) is convex.

LEMMA A.1. *The function $d(i, \mathbf{P}) = -Q_i(\mathbf{P}, \mathbf{R})/Q_i^{\text{spec}}$ is convex $\forall i \in \{0, 1, \dots, Q^{\text{spec}*}\}$.*

PROOF. To prove convexity, we need to prove that $d(i, \mathbf{P})$ is twice differentiable and that its Hessian is positive-semi-definite. Differentiating $d(i, \mathbf{P})$ w.r.t P_u and P_v , we get the following expression:

$$d_{P_u P_v}(i, \mathbf{P}) = \frac{1}{Q_i^{\text{spec}}} \cdot \begin{cases} 0 & \text{if } v = u \\ & \forall u > i \\ & \forall v > i \\ \sum_{j=i-w}^{j \leq P^*} R_j \cdot \prod_{k \in S(i,j)/\{u,v\}} (1 - P_k) & \text{otherwise} \end{cases} \quad (14)$$

where $w = \min\{u, v\}$. Since \mathbf{R} has no negative element and $0 \leq P_i \leq 1 \quad \forall i \in \{0, \dots, P^*\}$, $d_{P_u P_v}(\mathbf{P})$ is non-negative $\forall i, u, v \in \{0, \dots, P^*\}$. Therefore, the Hessian of $d(i, \mathbf{P})$ is positive-semi-definite in the specified domain of \mathbf{P} . Having a semi-positive-definite Hessian is a sufficient condition for convexity [6] and $d(i, \mathbf{P})$ is convex in \mathbf{P} . \square

LEMMA A.2. *The objective function (7) is convex with respect to \mathbf{P} .*

PROOF. From Lemma A.1, it can be seen that $-Q_i/Q_i^{\text{spec}}$ is convex with respect to \mathbf{P} . The objective in (7) can be transformed to the following dual form:

$$\text{minimize} \quad \max_{0 \leq i \leq Q^{\text{spec}*}} \{-Q_i/Q_i^{\text{spec}}\} \quad (15)$$

In (15), the max of n convex functions is computed. As the max of convex functions is also convex, the objective requires minimizing a convex function. \square

THEOREM A.3. *The optimization formulation specified in equations (7) to (9) is convex.*

PROOF. The objective (7) requires maximizing a concave function or minimizing a convex function (Lemma A.2). Constraints (8) and (9) are affine. Therefore, the optimization problem is convex. \square

A.2 Convexity of Aggregates Extension

Augmenting the optimization formulation in (7) to (9), retains its convexity.

LEMMA A.4. *The function $d^{\text{agg}}(i, \mathbf{P}) = -Q_i^{\text{agg}}/Q_i^{\text{spec,agg}}$ is convex $\forall i \in \{0, 1, \dots, Q^{\text{spec,agg}*}\}$.*

PROOF. To prove convexity, we need to prove that $d^{\text{agg}}(i, \mathbf{P})$ is twice differentiable and that its Hessian is positive-semi-definite.

$$d^{\text{agg}}_{P_u P_v}(i, \mathbf{P}) = \frac{1}{Q_i^{\text{spec,agg}}} \cdot \begin{cases} \sum_{j=1}^{j \leq R^*} j(j+1)R_j \left(1 - P_k^{\text{agg}}\right)^{(j-1)} & \text{if } u = v = i \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

Using arguments similar to the ones used in Lemma A.1, it can be seen that the Hessian of $d^{\text{agg}}(i, \mathbf{P})$ is positive-semi-definite, proving its convexity. \square

THEOREM A.5. *The optimization formulation specified in equations (7) to (9) remains convex after adding aggregate values to the formulation.*

PROOF. We have proven that the individual components of the objective for aging and aggregate values are convex (Lemmas A.1 and A.4). The objective (7) requires maximizing a concave function or minimizing a convex function (Lemma A.2). Constraints (8) and (9) are affine. Therefore, the optimization problem is convex. \square