# No Batteries Needed: Providing Physical Context with Energy-Harvesting Beacons

Nurani Saoda
University of Virginia
ns8nf@virginia.edu

Bradford Campbell
University of Virginia
bradjc@virginia.edu

## Abstract

Battery-powered digital beacons have played a significant role in shrinking the gap between physical and digital world. At the same time, ubiquitous sensing encourages tiny, unobtrusive, energy-harvesting devices to eliminate the limited lifetime of battery-powered devices. In this paper, we design a new fire-and-forget room number broadcaster beacon to investigate the feasibility and performance of such a design point. We study how several factors including different deployment spaces, the storage capacity of the harvester, indoor light intensity levels, and spatial position of the receiver impact the performance in three real-world deployments. We find that the 95th percentile of inter-packet reception time is 35 s or less in a lab space with exposure to sunlight and indoor lights, 29 s or less in an industrial plant with indoor lights, and 405 s or more in office rooms. With strategic beacon placement and a light intensity level of only 390 lx, performance can be improved by 61%. We believe that these results will help guide future energy-harvesting beacon deployments. We also outline possible improvements for future energy-harvesting beacon designs.

## CCS Concepts

• **Computer systems organization → Sensor networks**; **Embedded systems**.

## Keywords

Energy Harvesting Systems, Bluetooth Low Energy Beacons, Deployment

**ACM Reference Format:**
Nurani Saoda and Bradford Campbell. 2019. No Batteries Needed: Providing Physical Context with Energy-Harvesting Beacons. In *7th International Workshop on Energy Harvesting Energy-Neutral Sensing Systems (ENSsys'19), November 10, 2019, New York, NY, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3362053.3363489

## 1 Introduction

Energy-harvesting sensors and other wireless devices that scavenge their own energy are becoming more prevalent as avoiding battery replacement is seen as necessary to enable scaling up the number of Internet of Things devices. With the increased growth and interest in this area, a range of energy-harvesting system designs encompassing both hardware and software has emerged as potential solutions to the intermittency and unreliability inherent to using harvested energy sources. Researchers have proposed schemes with supercapacitor buffers [11, 14, 16], multi-tier energy backup [9, 11], software checkpointing [1, 8, 15], purely intermittent operation [3, 5], and energy-adaptive operation [4], among other techniques, and at least in the near future it is unlikely that a single platform will emerge as the de-facto energy-harvesting standard. Even with battery-powered devices designs span from fixed primary cells to rechargeable batteries [12, 13].

A variety of energy-harvesting approaches may be viable, and different architectures may be well suited for different applications. In this work we study a particular design point, namely fire-and-forget devices with a simple operating principle: harvest a usable amount of energy, perform some useful and complete operation, and then resume harvesting. This design point prioritizes simplicity to reduce both hardware and software costs. Our goal is not to advocate for this approach instead of alternatives, but to investigate the performance of this approach in different environments and contexts to understand its viability (or lack thereof) for different applications. This will help guide future system developers when considering energy-harvesting options for their application.

These simple energy-harvesting sensors that push towards nearly-invisible, stick-on or smart dust devices can be quite useful as digital beacons that periodically chirp sensor data, proximity information (as in iBeacon or Estimote), pointers linking the physical and digital worlds (as with Eddystone), or heartbeat packets. We extend this concept by introducing an energy-harvesting device designed for office settings that periodically sends the room number physically dialed in on the device itself. This provides a digital counterpart to the ubiquitous room number familiar to occupants, and can provide rough indoor location information to user-carried devices or other sensors in the environment.

While the room-number transmitter and other beacons can easily be, and often are, battery powered, will they work as small, indoor photovoltaic powered energy-harvesting devices? How long will a listener have to wait for a beacon? How many beacons should be deployed in a space, and what is the impact of the placement? How do office environments compare to more industrial environments? How do the ambient lighting level and on-device energy reservoir affect the performance? Understanding the answers to these and other questions will inform the utility of this approach and the results will set a benchmark for future deployments considering this approach.

Our study consists of deploying several of the fire-and-forget sensors in various environments and analyzing the resulting stream of beacons. We find that 95% of the time a listener has to wait 35


*ENSsys'19, November 10, 2019, New York, NY, USA*

ACM ISBN 978-1-4503-7010-3/19/11…$15.00
https://doi.org/10.1145/3362053.3363489

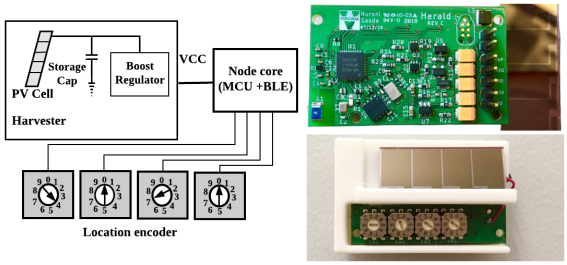


(a) System block diagram (b) Herald beacon (27mm×45mm)

**Figure 1: Design and implementation of Herald.**

seconds or fewer for a beacon in a lab space with windows, 405 seconds or more in interior office rooms with no windows, and 29 seconds or less in an industrial steam plant. While performance can be improved by 61% requiring fewer beacons with beacon placements close to light sources and at a light intensity of 390 lx, these placements do not always result in increased system availability. Our results also show that receiver position can affect the packet interval time by 56%, and reducing the energy buffer size can also increase performance.

Based on these results we also describe some future research directions that would help energy-harvesting beacons be viable in more application scenarios.

## 2 Related Work

We highlight existing works in low power beacon design, energy-harvesting system modeling and energy-harvesting system design.

**Energy harvesting beacon design.** A direction of work focuses on adopting light energy harvesting power supply for a variety of applications. Fraternali et al. propose ambient light energy harvesting sensor Pible based on supercapacitor storage but suffers from over four hours of outage due to longer recharge time [6]. Jean et al. design light energy harvesting Luxbeacon and show that the design is viable in low indoor light levels in their field test [10]. The study mostly focused on the design principles of solar energy harvesting under varying light intensity. In this paper, we design an intermittently powered light energy harvesting beacon and study a "harvest and use" design point for fire-and-forget devices and investigate the feasibility of such an approach in different application requirements.

**Energy-harvesting system modeling.** Another direction focuses on developing software tools to capture the dynamics in energy harvesting systems. Enspect [17] emulates the power output from different harvesting modalities closely matching the real power traces of the harvester. SunaPlayer [2] designs a non-linear analog device to emulate the characteristics of photovoltaic (PV) cells. In this work, we evaluate the performance of a specific energy-harvesting design in real-world harvesting conditions.

**Energy-harvesting system design.** Previously, UFoP architecture [7] has been proposed which partitions harvested energy in varying sizes of smaller storages. Also, Capybara [4] architecture adopts a hybrid hardware-software mechanism to reconfigure the storage size according to application demand. While all of these sophisticated design architectures have outperformed the traditional ones for a set of target sensing applications, the viability of such designs in different indoor light settings for digital beacon applications has not been explored.

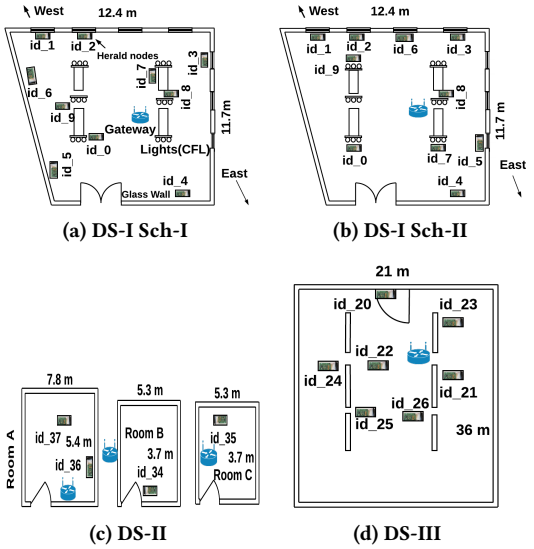


(a) DS-I Sch-I (b) DS-I Sch-II

(c) DS-II (d) DS-III

**Figure 2: High level schematic of different deployment scenarios with floor plans. (a) DS-I Sch-I in a lab space with beacon placement at a combination of window-based and wall-based positions at human heights. (b) DS-I Sch-II corresponds to the same space in DS-I Sch-II with an optimized placement. (c) DS-II is in 3 nearby conference rooms in a smart building without direct sunlight exposure. (d) DS-III corresponds to a steam plant mostly lit by indoor lights.**

## 3 Design and Implementation

To work towards a vision of location-aware perpetual IoT world, we design Herald, a photovoltaic energy-harvesting BLE beacon which opportunistically broadcasts location information. Herald transmits coarse location data and a URL for physical web interfaces in its chirp transmissions. Its form-factor is small and easy to install in indoor settings.

### 3.1 Design

The Herald design consists of three major components: an energy-harvesting power supply, location input interface, and a Bluetooth Low Energy (BLE) core. Its operating principle is simple: the node core activates when the harvested energy exceeds a threshold and performs an activity cycle. Enough energy is temporarily stored on small storage capacitors to send a radio packet. After activating, the node core cold boots, initializes the radio stacks, reads the location data from the peripheral, and creates a data packet to send as a BLE advertisement. The activity cycle concludes as the voltage on the storage capacitors falls below a usable threshold. The node core also maintains an activation counter and a packet sequence number, and includes them in each transmission. Figure 1(a) shows the block diagram of the Herald system design.

### 3.2 Implementation

We implement Herald with the following design goals: low power design, reconfigurable hardware, easy-to-deploy, and a small form factor for indoor smart spaces.

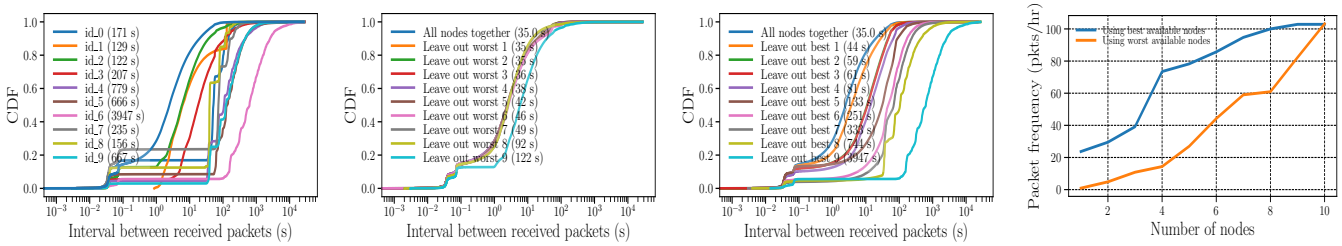


**(a) CDF of inter-packet reception time.** **(b) Leave out worst beacons (LOW)** **(c) Leave out best beacons (LOB)** **(d) Packet reception frequency vs number of beacons**

**Figure 3: Results from DS-I Sch-I. (a) CDFs of the interval between received packets for individual beacons with the 95th percentile value specified for each. (b) CDFs of the packet intervals starting with all beacons and removing the worst performing beacons at a time. (c) CDFs of the packet reception interval starting with all beacons and removing the best performing beacons each time. (d) Frequency of received packets per hour in two cases: using the best and worst available nodes.**

The energy harvesting power supply is implemented using prior work by Yerva, et al. [18]. It uses a 35 mm × 19.5 mm amorphous silicon PV cell suitable for indoor lighting ranging from 50 to 1000 lx. The harvested energy from the PV cell is stored on a bank of five 100 µF tantalum capacitors which provides a temporary energy buffer on the system. The Microchip MCP1640 boost regulator supplies the node core with a stable voltage.

We use the nRF51822, a BLE SoC by Nordic Semiconductor as the node core for its reasonably low TX current (10.3 mA) at 0 dbm TX power level. The MCU reads location data from four 7.3 mm × 7.1 mm rotary encoded switches that allow a user to encode a room number by specifying each digit on the device. The prototype enclosure as shown in Figure 1(b) allows the device to be reconfigured by the user if moved to a different location, and the solar panel can be positioned based on the available light source. Lastly, the MCU leverages a 4K FM25L04B FRAM to read and write the activation number and packet sequence number at each wake up cycle.

## 4 Real-world Deployment

We employ several real-world deployments to evaluate the system performance in multiple uncontrolled settings with a mixture of natural and artificial light sources.

### 4.1 Deployment Scenarios

Our deployment consists of three different scenarios: a hardware research lab, conference rooms and an industrial steam plant. The steam plant represents reasonably harsh (i.e. above room temperature, high signal attenuation) environment. Figure 2 shows the high level floor plans of different deployment schemes. We categorize the scenarios as follows:

- **Deployment Scenario I (DS-I)**: a hardware lab with exposure to sunlight and indoor LED lights, with two versions:
  - Unplanned deployment scheme (DS-I Sch-I): Here, we position beacons to maximize coverage of the space without considering irradiance levels. We install the beacons at human height level.
  - Planned deployment scheme (DS-I Sch-II): We optimize positions in the area with high irradiance levels, such as close to sunlight exposure and wall mounted lights.
- **Deployment Scenario II (DS-II)**: three conference rooms lit with indoor LED lights.
- **Deployment Scenario III (DS-III)**: an industrial steam plant with mostly indoor lights.

### 4.2 Setup

We installed a total of 38 Herald beacons in different deployment scenarios and collected data packets with a BLE receiver for at least one week for each scenario. DS-I Sch-I and DS-I Sch II each contain ten Herald beacons and one gateway receiver. DS-II contains four beacons with one in each of the first two rooms and the remaining two on different walls of a relatively larger third room, plus two gateway receivers. DS-III has seven Herald beacons installed with one gateway receiver.

## 5 Evaluation

In this section, our goal is to analyze how physical variables (i.e. different deployment schemes, positions of beacon placement, indoor light levels, and presence of multiple receivers) and design variables (i.e. the energy storage size) affect the system-level performance of fire-and-forget energy-harvesting devices.

### 5.1 Overall System Performance

We start by evaluating the overall performance of the system across different DSs (deployment scenarios). Specifically, we evaluate the Cumulative Distribution Function (CDF) of inter-packet reception time, number of beacons required, and system availability.

*5.1.1 DS-I Sch-I* Figure 3(a) shows the CDFs of the packet reception interval at the gateway. With 10 nodes, we achieve a 95th percentile of 35 s for inter-packet reception time which corresponds to the maximum time (with 95% probability) a receiver needs to wait until it receives a message from any of the 10 beacons in this particular DS. Different beacons see a wide range of transmission intervals, with the best node achieving a 95th percentile interval of 122 s (~2 minutes) and the worst node at 3947 s ( ~1 hour). Though these numbers are acceptable for certain non-time critical applications, they are relatively high for applications that need fast updates. These findings suggest a denser deployment of beacons or coordination among beacons is needed.

Next, we investigate how beacon count affects inter-packet arrival time. In this experiment, we start by removing one beacon (either the best or worst) from the dataset at a time and calculate the 95th percentile at each step Figure 3(b) and Figure 3(c) show the results from the leave-out-best and leave-out-worst cases. Removing the best performing node from the system degrades overall

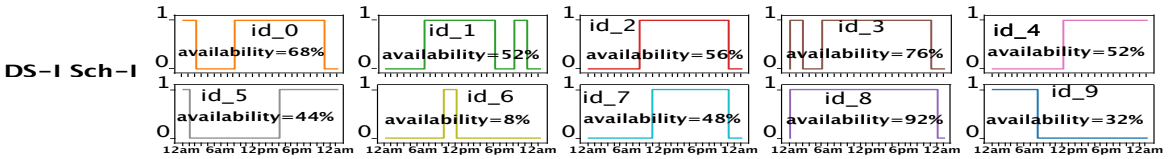


(a) Availability of individual beacons over the course of a day in DS-I Sch-I

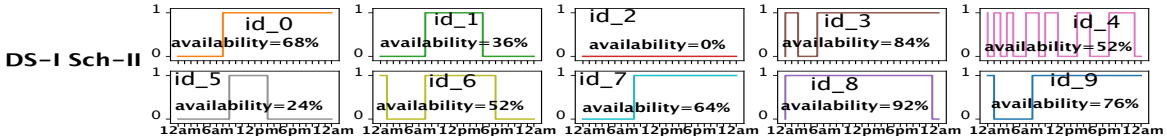


(b) Availability of individual beacons over the course of a day in DS-I Sch-II

**Figure 4: Availability of individual beacons in DS-I. In (a) 40% of the beacons are unavailable for more than 50% of the day. In (b) we change the beacon positions and achieve only an overall 2% point improvement of average system availability. Three beacons actually had worse availability, and this suggests the complexity of finding ideal beacon placement.**

performance by 31%, whereas removing the worst performing node has almost zero impact. Figure 3(d) shows the frequency of received packets per hour in optimistic and pessimistic cases, where the best nodes or the worst performing nodes are deployed, respectively. This shows that adding more beacons after a certain number does not improve the throughput much as the curve reaches a plateau, and shows the number of beacons required to meet a specific throughput requirement. For example, in DS-I four well-placed beacons are sufficient to receive at least one packet per minute.

With intermittent energy-harvesting, a lack of energy results in reduced availability of the system. In Figure 4(a), we plot the time over a certain sunny day when a beacon is "available", that is it successfully sent a packet to the receiver at least once within each hour. We find that over a day at least one beacon was available in each of the 24 one-hour periods, resulting in an aggregate availability of 100%. While no individual device is 100% available, the overall energy-harvesting system in an unplanned deployment is able to consistently operate through changing conditions, suggesting redundancy can support applications that require consistent sensing.

*5.1.2 DS-I Sch-II* In DS-I Sch-II, we changed the positions of beacon id_3, id_4, id_5, id_6, id_7, and id_9 to positions with higher perceived irradiance level such as windows and bright walls (Figure 2(b)) and shifted the positions of other nodes already in proximity to the light sources above normal human height. This will help understand if a planned and cautious deployment of these intermittently powered devices leads to less wait time and better system availability. Figure 5(b) shows an improvement of 31.43% over DS-I Sch-I in overall 95th percentile of inter packet reception time, and using the best four beacons matches the minimum packet interval of DS-I Sch-I. However, even with a planned and improved deployment, only 6 out of 10 beacons achieved a lower 95th percentile inter-packet interval.

To evaluate if the better deployment positions result in increased availability over the course of a sunny day, we plot the same availability graph in Figure 4(b). Though the overall average system availability is improved by 2% point, we find that three beacons incurred worse availability. While optimizing beacon placement can significantly improve packet interval times, optimizing for availability was more difficult. This suggests that more sophisticated placement strategies are needed.

*5.1.3 DS-II* DS-II tests the performance of the beacons in a space which is mostly lit by indoor lights in the presence of humans. Figure 6(a) shows the CDFs of interval between received packets on a week long dataset. The best performing node has a 95th percentile interval of 405 s (~7 minutes). This suggests that a denser deployment or improved PV cell design might be required for interior rooms. Figure 6(b) shows the availability of those four beacons over the course of a week. Interestingly, id_36 was not sending packets for 4 out of 7 days over a week, though the other co-located beacon

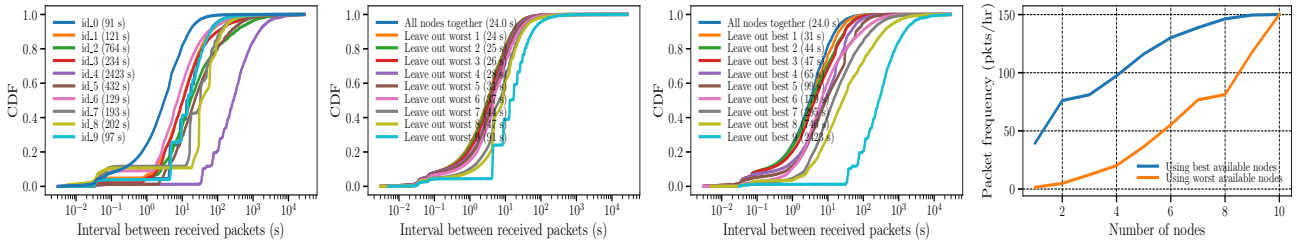


(a) CDF of inter packet reception time for individual beacons

(b) CDF of LOW

(c) CDF of LOB

(d) Received packet frequency vs number of nodes

**Figure 5: Results corresponding to DS-I Sch-II. (a) the CDFs of packet reception interval for individual beacons. Six out of ten beacons reduced in 95th percentile value. (b) and (c) show the LOW and LOB experiments. (d) shows the number of nodes vs the packet reception frequency where the maximum throughput is 40% higher than DS-I Sch-I and fewer nodes are required for a desired frequency.**

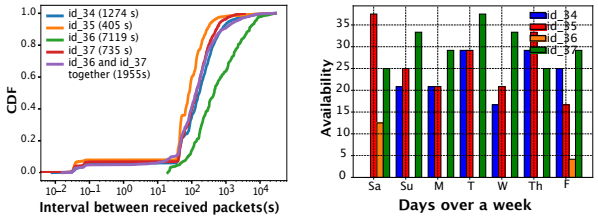


(a) (b)

**Figure 6: Performance in DS-II. (a) CDFs of interval between received packets in DS-II. (b) The percentage of hours each beacon was available in a day over the course of a week.**

id_37 which faces a bright glass wall with sunlight was mostly operating, highlighting the variation in available energy and the need for systems to adapt to the individual deployment.

*5.1.4 DS-III* This deployment scenario models industrial spaces that usually lack direct sunlight, have low indoor irradiance level, have temperature above room level, and are enriched with wireless signal attenuators. Figure 7 shows the 95th percentile interval between received packets for all nodes and the required number of sensors for a particular throughput. We find that although the 95th percentile corresponds to 29 s and a maximum throughput of 125 pkts/hr (on par with DS-I), some beacons have packet intervals in the range of several hours. Also, Figure 7(b) confirms that the system performance is highly dependent on the best performing beacon.

## 5.2 Impact of Reduced Energy Storage

To understand how energy storage capacity affects energy-harvesting beacons, we build an analytical model to capture the relation between storage size, packet transmission interval, and throughput, and also test experimentally.

*5.2.1 Re-configuring the storage size of the harvester* We start by identifying the minimum capacitance needed to perform the most energy-hungry atomic task, i.e. a radio transmission, for a particular radio. We then sweep the storage capacitance in a range of 0-500 µF and plot the analytical relation between capacitance and the interval between transmitted packets and packets transmitted per hour in Figure 8(a) and Figure 8(b) for harvesting levels of 20 µW and 200 µW, respectively. The minimum required energy for

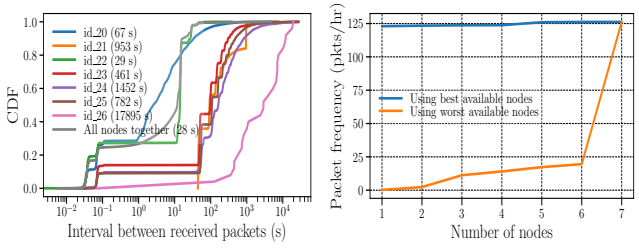


(a) (b)

**Figure 7: Performance in DS-III. (a) CDFs of interval between received packets with the best node having 95th percentile at 28 s. (b) Received packet frequency vs Number of nodes. The maximum throughput is comparable to DS-I Sch-I.**

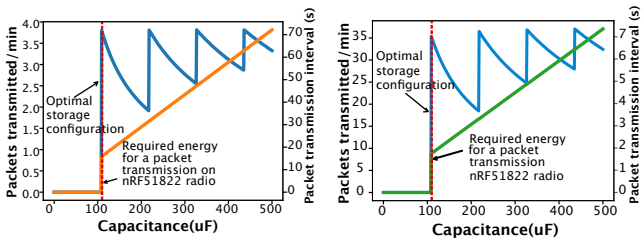


(a) $P_{in}$=20 µW (b) $P_{in}$=200 µW

**Figure 8: Modeling the relation between capacitance, packets transmitted per hour and transmission interval for two different harvesting conditions. The blue line corresponds to the packets transmitted per minute. The dotted red line with maximum throughput and minimum transmission interval denotes the optimal storage configuration.**

one atomic task (in our case, transmitting a 14 B payload packet) for the nRF51822 corresponds to 119 µF, below which the device never transmits. The optimal storage configuration with the highest throughput and lowest 95th interval is at the same capacitance. A larger storage capacitance eventually allows multiple transmissions per activation, improving throughput but not helping packet interval timing.

*5.2.2 Results from different deployment scenarios.* We conduct the same deployment study with the storage capacitor size set to 200 µF and re-run the experiments for at least a week of time. Figure 9 shows the 95th percentile values from the customized storage size over the three different deployment schemes. We can see that as we decrease the storage size the overall inter-packet reception time indeed decreased for DS-I Sch-II and DS-III which matches with our analytical model. But for DS-I with poor harvesting conditions, the overall inter packet reception time deteriorated which we believe is due to two beacons sending too infrequent packets over three days in a week.

## 5.3 Impact of Indoor Light Levels

To understand how light levels affect overall system performance, we measure the ambient light levels at the gateway receivers, and then calculate the packet interval times during different lighting conditions. Figure 10 shows the typical brightness level in DS-I over a day. From Figure 11(a), we observe that as the average light level increases packet interval times decrease, but increasing the light levels from the 50th percentile to the 95th percentile has little

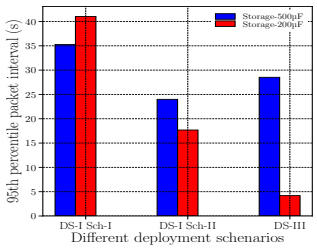


**Figure 9: 95th percentile inter-packet reception time for all beacons together in different DSs. In DS-I Sch-II and DS-III reducing capacitance lowers inter-packet time.**

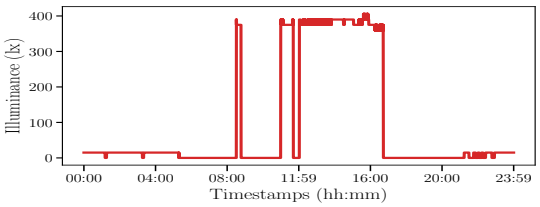


**Figure 10: Room brightness over one day.**

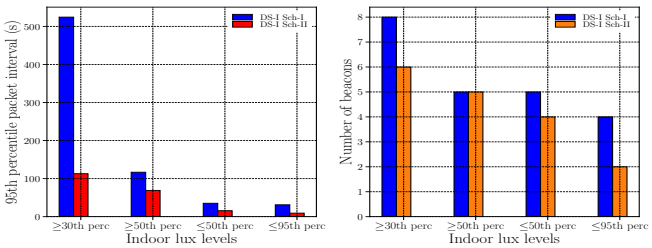


**Figure 11: Impact of indoor irradiance level on performance. (a) The 95th percentile inter-packet time improves with light levels. (b) shows the required number of beacons for a 90% of packet reception frequency. Only two beacons are required above 95th percentile lux value in DS-I Sch-II. Here, ≥ and ≤ indicates that we only consider the light levels above and below the percentile value respectively.**

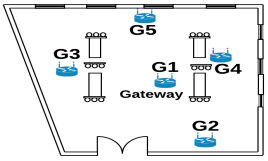


**(a) Positioning of 5 gateways in DS-I Sch-I with the beacon positions omitted.**

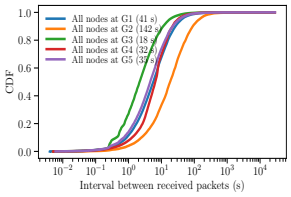


**(b) CDFs of inter-packet reception time over all the gateways.**

**Figure 12: Performance with multiple receivers at different positions in DS-I. (a) High level schematic with relative positioning between receivers. (b) CDFs of the inter-packet time for 10 beacons together at different receivers.**

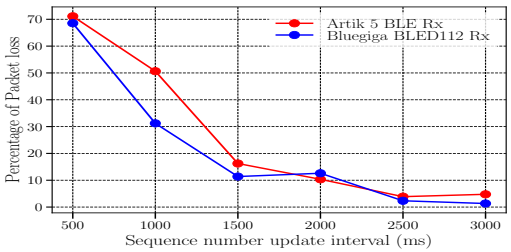


**Figure 13: Average percentage of packet loss vs Sequence number update interval on Artik 530 and BLED112 BLE receiver. Percentage of packet loss is above 65% for a sequence number update interval below 500ms.**

effect. With a lux level of only 390 lx, the 95th percentile value reduces by 61% in DS-I Sch-II from DS-I Sch-I. In the low lighting condition (30th percentile), two window based beacons in DS-I Sch-I and four beacons in DS-I Sch-II did not send any packets. Figure 11(b) shows the number of beacons required to achieve 90% of the maximum packet reception frequency at different light levels, and better lighting requires fewer beacons.

## 5.4 Impact of Receiver Positioning

To evaluate how the presence and position of multiple receivers in a real world environment affect the overall performance, we deployed five gateways in DS-I Sch-I with the beacon storage (Figure 12(a)) set to 200 µF. Figure 12(b) shows the 95th percentile inter-packet reception time for all beacons in DS-I Sch-I where G3 is the receiver with the smallest 95th percentile value. With an optimized receiver position, we achieve a 56% improvement over the central position as in DS-I. The mean and standard deviation of the 95th percentiles are 54 s and 49 s respectively. Also, we notice that the receiver G3 which is closest to the best performing node (in this case, id_1) outperforms the rest, instead of G1 positioned at a central point.

## 6 Directions for Future Systems

Based on our deployments, we highlight future research directions to design viable applications using energy-harvesting beacons.

- Typical BLE beacon applications leverage smartphones where scanning for beacons is a relatively energy-inexpensive operation. However, can low power devices leverage these beacon packets as well? Based on the results from the bright lighting conditions, a new "sense and scan" architecture where a low power receiver only listens in high-light conditions may be feasible.
- We find that Herald incurs a high percentage of packet loss using off-the-shelf BLE receivers (Figure 13). This suggests improved receivers with better BLE stack implementations are needed.
- Placing light-based energy-harvesting beacons to prioritize bright lighting positions does not actually guarantee increased system availability. New beacon placement strategies or runtime-optimization systems are needed that consider overall system availability and not per-node harvesting capability.

## 7 Conclusion

While battery-powered beacons have enabled many promising applications for Internet of Things devices, making them more sustainable requires energy-harvesting. In this paper, we investigate how different variables such as different deployment spaces, placement of beacons, harvester storage configuration, indoor light intensity levels, and positions of the receiver influence the performance of fire-and-forget light energy harvesting beacons. We design Herald, a room number beacon and conduct our experiments in real-world deployments with different lighting conditions. The work demonstrate a benchmark for future deployments of fire-and-forget energy-harvesting beacons with possible future research directions.

## 8 Acknowledgements

We thank the anonymous reviewers for their insights on improving this paper. This work is supported in part by the National Science Foundation under grant CNS-1823325, and the Strategic Investment Fund at the University of Virginia.

## References


[1] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi, D. Brunelli, and L. Benini. 2015. Hibernus: Sustaining Computation During Intermittent Supply for Energy-Harvesting Systems. *IEEE Embedded Systems Letters* 7, 1 (March 2015), 15–18. https://doi.org/10.1109/LES.2014.2371494

[2] Stanislav Bobovych, Nilanjan Banerjee, Ryan Robucci, James P. Parkerson, Jackson Schmandt, and Chintan Patel. 2015. SunaPlayer: High-accuracy Emulation of Solar Cells. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks (IPSN '15)*. ACM, New York, NY, USA, 59–70. https://doi.org/10.1145/2737095.2737110

[3] Bradford Campbell and Prabal Dutta. 2014. An Energy-harvesting Sensor Architecture and Toolkit for Building Monitoring and Event Detection. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings (BuildSys '14)*. ACM, New York, NY, USA, 100–109. https://doi.org/10.1145/2674061.2674083

[4] Alexei Colin, Emily Ruppel, and Brandon Lucia. 2018. A Reconfigurable Energy Storage Architecture for Energy-harvesting Devices. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '18)*. ACM, New York, NY, USA, 767–781. https://doi.org/10.1145/3173162.3173210

[5] Samuel DeBruin, Bradford Campbell, and Prabal Dutta. 2013. Monjolo: An Energy-harvesting Energy Meter Architecture. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys '13)*. ACM, New York, NY, USA, Article 18, 14 pages. https://doi.org/10.1145/2517351.2517363

[6] Francesco Fraternali, Bharathan Balaji, Yuvraj Agarwal, Luca Benini, and Rajesh Gupta. 2018. Pible: Battery-free Mote for Perpetual Indoor BLE Applications. In *Proceedings of the 5th Conference on Systems for Built Environments (BuildSys '18)*. ACM, New York, NY, USA, 168–171. https://doi.org/10.1145/3276774.3282822

[7] Josiah Hester, Lanny Sitanayah, and Jacob Sorber. 2015. Tragedy of the Coulombs: Federating Energy Storage for Tiny, Intermittently-Powered Sensors. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys '15)*. ACM, New York, NY, USA, 5–16. https://doi.org/10.1145/2809695.2809707

[8] 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 '17)*. ACM, New York, NY, USA, Article 17, 13 pages. https://doi.org/10.1145/3131672.3131673

[9] Neal Jackson, Joshua Adkins, and Prabal Dutta. 2019. Capacity over Capacitance for Reliable Energy Harvesting Sensors. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks (IPSN '19)*. ACM, New York, NY, USA, 193–204. https://doi.org/10.1145/3302506.3310400

[10] K. E. Jeon, J. She, J. Xue, S. Kim, and S. Park. 2019. luXbeaconâĂŤA Batteryless Beacon for Green IoT: Design, Modeling, and Field Tests. *IEEE Internet of Things Journal* 6, 3 (June 2019), 5001–5012. https://doi.org/10.1109/JIOT.2019.2894798

[11] Xiaofan Jiang, Joseph Polastre, and David Culler. 2005. Perpetual Environmentally Powered Sensor Networks. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN '05)*. IEEE Press, Piscataway, NJ, USA, Article 65. http://dl.acm.org/citation.cfm?id=1147685.1147765

[12] Robert Margolies, Maria Gorlatova, John Sarik, Gerald Stanje, Jianxun Zhu, Paul Miller, Marcin Szczodrak, Baradwaj Vigraham, Luca Carloni, Peter Kinget, Ioannis Kymissis, and Gil Zussman. 2015. Energy-Harvesting Active Networked Tags (EnHANTs): Prototyping and Experimentation. *ACM Trans. Sen. Netw.* 11, 4, Article 62 (Nov. 2015), 27 pages. https://doi.org/10.1145/2831236

[13] Paul Martin, Zainul Charbiwala, and Mani Srivastava. 2012. DoubleDip: Leveraging Thermoelectric Harvesting for Low Power Monitoring of Sporadic Water Use. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (SenSys '12)*. ACM, New York, NY, USA, 225–238. https://doi.org/10.1145/2426656.2426679

[14] C. Park and P. H. Chou. 2006. AmbiMax: Autonomous Energy Harvesting Platform for Multi-Supply Wireless Sensor Nodes. In *2006 3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks*, Vol. 1. 168–177. https://doi.org/10.1109/SAHCN.2006.288421

[15] Benjamin Ransford, Jacob Sorber, and Kevin Fu. 2011. Mementos: System Support for Long-running Computation on RFID-scale Devices. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XVI)*. ACM, New York, NY, USA, 159–170. https://doi.org/10.1145/1950365.1950386

[16] F. Simjee and P. H. Chou. 2006. Everlast: Long-life, Supercapacitor-operated Wireless Sensor Node. In *ISLPED'06 Proceedings of the 2006 International Symposium on Low Power Electronics and Design*. 197–202. https://doi.org/10.1145/1165573.1165619

[17] Nick F. Tinsley, Stuart T. Witts, Jacob M.R. Ansell, Emily Barnes, Simeon M. Jenkins, Dhanushan Raveendran, Geoff V. Merrett, and Alex S. Weddell. 2015. Enspect: A Complete Tool Using Modeling and Real Data to Assist the Design of Energy Harvesting Systems. In *Proceedings of the 3rd International Workshop on Energy Harvesting & Energy Neutral Sensing Systems (ENSsys '15)*. ACM, New York, NY, USA, 27–32. https://doi.org/10.1145/2820645.2820648

[18] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta. 2012. Grafting energy-harvesting leaves onto the sensornet tree. In *2012 ACM/IEEE 11th International Conference on Information Processing in Sensor Networks (IPSN)*. 197–207. https://doi.org/10.1109/IPSN.2012.6920957