

Demo Abstract: Power-Aware Neighbor Discovery for Energy Harvesting Things

Tingjun Chen¹, Gregory Chen², Saahil Jain², Robert Margolies¹, Guy Grebla¹,
Dan Rubenstein², Gil Zussman¹

¹Electrical Engineering and ²Computer Science, Columbia University, New York, NY 10027, USA
{tingjun@ee., glc2121@, sj2675@, robm@ee., guy@ee., danr@cs., gil@ee.}columbia.edu

ABSTRACT

Object tracking applications are gaining popularity and will soon utilize energy harvesting low-power wireless nodes where power is mostly consumed for *neighbor discovery*. Such applications require the design and experimentation with low-power neighbor discovery protocols. We demonstrate the Panda protocol [4, 5] implementation using commercial off-the-shelf energy harvesting devices, based on the TI eZ430-RF2500-SEH prototype. The prototypes harvest indoor light energy to perform *power-aware* neighbor discovery, while maintaining a power budget. A custom-designed online monitoring system interactively demonstrates the wireless network dynamics, including the energy levels of the devices, the neighbor discovery events, and aggregate discovery statistics.

Keywords

Neighbor Discovery; Energy Harvesting; Low-Power Communications; Energy Adaptive Networking.

1. INTRODUCTION

Object tracking and monitoring applications are gaining popularity within the realm of the Internet-of-Things (IoT) [3]. Emerging low-power wireless nodes that can be attached to physical objects are enablers for such applications. Often, these nodes are meant to interact with a reader, but architectures are emerging that handle scenarios where no reader is present, or where the number of nodes overwhelms the readers' availability. These scenarios can be supported by Energy Harvesting (EH) tags (e.g., [6]) that are able to communicate peer-to-peer and are powered by an ambient energy source (e.g., indoor light).

Such EH nodes will enable applications in healthcare, smart buildings, assisted living, manufacturing, supply chain management, and intelligent transportation [2]. An example application is identifying misplaced objects: often when an item is misplaced, its ID is significantly different from the IDs of its neighbors. In such a case, the misplaced node can, for instance, flash a low-power LED to indicate that it is lost. One enabler of such applications is Neighbor Discovery (ND), in which nodes identify neighbors that are within wireless communication range.




Figure 1: The prototype EH node stores harvested energy in a capacitor and powers the microcontroller and transceiver.




Figure 2: Panda demonstration setup: 5 EH nodes harvest energy through the solar cells with neighbor discovery events monitored by a listening node.

To maintain perpetual tracking of the mobile objects, ND must be run continuously with the EH nodes operating in a low-power mode, consuming power at the rate of power harvested. We proposed, analyzed, and experimented with the Panda protocol in [4, 5], which is configured to *maximize the rate at which nodes discover their neighbors, given a constrained power budget at each EH node*.

In this demonstration, we present the Panda protocol¹ implementation using Commercial Off-The-Shelf (COTS) EH nodes, based on the TI eZ430-RF2500-SEH prototype [1] shown in Fig. 1. The demonstration testbed (shown in Fig. 2) is composed of EH nodes harvesting light from a custom-designed software-based light control system. The EH nodes contain a low-power microcontroller and a wireless transceiver, and are configured to discover their neighbors using the Panda protocol. A custom-designed online monitoring system interactively visualizes both the wireless network dynamics and the neighbor discovery events and statistics.

¹We focus on the implementation of the dynamic version of the protocol, Panda-Dynamic, in [4, 5]. For simplicity, we refer to it as Panda throughout this abstract.




Figure 3: The Panda online monitoring system which illustrates: (a) network dynamic and discovery events, (b) energy storage level (capacitor voltage) of each node, (c) the neighbor discovery rate at each node, and (d) the discovery rate and discovery latency statistics of both individual nodes and the whole network.

2. PROTOTYPE DESCRIPTION

To maximize the discovery rate, Panda follows a probabilistic approach in which nodes transition between sleep, receive, and transmit states (for more details, see [4]). The Panda protocol is implemented in the EH prototype based on the TI eZ430-RF2500-SEH. However, to fully experiment and demonstrate with the COTS devices, we made some modifications to the hardware, as summarized in [5]. The prototype consists of the following components:

Power Harvesting Source: The prototype harvests light from a Sanyo AM 1815 amorphous solar cell, which has a fixed harvesting voltage of 1.2 V.

Energy Storage: The energy harvested by the solar cell is stored in a 10 mF capacitor. To ensure stable operation of the microcontroller, the capacitor is prevented from reaching a low voltage level by imposing a software cutoff at 3.6 V. If a node reaches such a low voltage, it enters an ultra-low-power sleep state until the voltage level reaches a safe operating point.

Low-Power Microcontroller: A TI-MSP430 microcontroller is used to provide computational capabilities. These include (i) sampling the capacitor voltage using an analog to digital converter (ADC), (ii) operating a low-power 12 kHz clock with an idle power draw of 1.6 μ W to instruct the node to enter and exit an ultra-low-power deep sleep state, and (iii) sending and receiving messages to and from the radio layer.

Low-Power Transceiver: The prototype utilizes a 2.4 GHz CC2500 wireless transceiver, which is designed to provide low-power wireless communications. The transceiver consumes 64.85 mW in receive state, and the transmission power level is set in software to -10 dBm, with a resulting power consumption of 59.23 mW. At this transmission power level, nodes within the same room typically have little or no packet loss.

3. DEMONSTRATION

The demonstration testbed (see Fig. 2) includes: (i) five EH prototypes running the Panda protocol, (ii) a custom-designed software-based light control system, and (iii) a custom-designed online monitoring system. In the light control system, a Java-based script and an Arduino-based light control module precisely controls the irradiance (light energy intensity) generated by light-emitting diodes (LEDs). Dark box enclosures and 3D printed mounting fixtures ensure full control over the light conditions at the solar cells.

Demo participants can change the energy harvesting rates at the nodes by increasing or reducing the amount of light incident on the solar cells. From the monitoring system, demo participants can observe power-aware neighbor discovery: nodes with a higher en-

ergy harvesting rate are more active, and listen and transmit more frequently than nodes with a lower energy harvesting rate. In addition, demo participants can also observe the following features via a set of graphs shown in Fig. 3.

Node States and Discovery Events: The network dynamics are demonstrated by a topology graph (see Fig. 3(a)) in which each numbered node represents an EH prototype. A node in the graph “blinks” when it transmits a discovery message. If a discovery event happens, a directed edge will occur, pointing from the transmitting node to the receiving node (e.g., in Fig. 3(a), node 2 receives a discovery message from node 1).

Energy Storage Level: Fig. 3(b) shows the capacitor voltage of each node, which corresponds to its energy storage level. Because of the stochastic nature of the Panda protocol as well as the varying energy harvesting rates, each node’s capacitor voltage varies with time. In addition, energy neutrality is demonstrated by the voltage oscillating within the limits of the capacitor storage.

ND Statistics: Fig. 3(c) and Fig. 3(d) show the ND statistics, from each individual node to the overall network. The ND rate of one node is the average number of discovery messages per unit time it receives from all other nodes, and the ND rate of the network is the sum of ND rate of individual nodes. Fig. 3(d) also contains the discovery latency, which is the time between consecutive discoveries received by one node from some other node.

Acknowledgements

This research was supported by NSF grants CCF-09-64497, CNS-10-54856, and CNS-16-50685, the People Programme (Marie Curie Actions) of the European Union’s Seventh Framework Programme (FP7/2007-2013) under REA grant agreement no. [PIIF-GA-2013-629740].11.

4. REFERENCES

- [1] Texas Instruments eZ430-RF2500-SEH solar energy harvesting development tool user’s guide, 2013.
- [2] O. Abari, D. Vasisht, D. Katabi, and A. Chandrakasan. Caraoke: An e-toll transponder network for smart cities. In *Proc. ACM SIGCOMM’15*, 2015.
- [3] L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer networks*, 54(15):2787–2805, 2010.
- [4] R. Margolies, G. Grebla, T. Chen, D. Rubenstein, and G. Zussman. Panda: Neighbor discovery on a power harvesting budget. In *Proc. IEEE INFOCOM’16*, 2016.
- [5] R. Margolies, G. Grebla, T. Chen, D. Rubenstein, and G. Zussman. Panda: Neighbor discovery on a power harvesting budget. *IEEE J. Sel. Areas Commun. (to appear)*, 2016.
- [6] S. Ulukus, A. Yener, E. Erkip, O. Simeone, M. Zorzi, P. Grover, and K. Huang. Energy harvesting wireless communications: A review of recent advances. *IEEE J. Sel. Areas Commun.*, 33(3):360–381, 2015.