

Poster Abstract: Network-Wide Energy Profiling of CTP

Marcelo Martins Rodrigo Fonseca
Department of Computer Science
Brown University
{martins, rfonseca}@cs.brown.edu

Thomas Schmid Prabal Dutta
Computer Science & Engineering Division
University of Michigan
{thschmid, prabal}@umich.edu

Abstract

We present our experiences evaluating the power-performance tradeoffs of a sensor network protocol on a power-aware testbed. We characterize the power draw of the entire network while running the Collection Tree Protocol (CTP), as a function of low-power listening interval. We find that message transmission counts are poor predictors for energy consumption on the CC2420 radio, that CTP routinely creates energy hotspots in the routing tree, and that conclusions based on protocol evaluation performed without low-power listening enabled provide little insight about the same protocol performance using low-power listening.

1 Introduction

Wireless sensor networks, perhaps more than any other area of systems and networking research, has been driven by low-power design goals. Yet, with few exceptions, publications do not report empirical energy consumption results even though hardware, software, and simulators for energy metering and estimation exist. It is remarkable that after a decade of research in sensor networks, so many efforts still continue to focus on improving energy efficiency¹, yet so few publications report simulated, modeled, or empirical results.

In this paper, we make a point on the importance of evaluating the energy efficiency claims in the literature using *empirical metrics* and *realistic workloads*. Much like experimental methods supplanted simulation results in wireless networking and exposed many false assumptions, we show that empirical measurements should supplant proxy metrics in order to truly understand energy complexity.

To show the dangers of empirical negligence, we consider the case of energy consumption in a network protocol. A great amount of simulation-based evaluation of protocols for wireless sensor networks (WSNs) rely on simple assumptions for energy consumption, such as the number of packets sent and received, while disregarding other parameters,

¹Energy-efficiency pervades the sensor network research agenda: in 2009, over 70% of SPOT papers, 60% of Sensys papers, and 40% of IPSN papers were directly or indirectly related to energy, where directly means that energy is part of the title, abstract, or central argument, while indirectly means that energy is part of the motivation, evaluation, or a key tradeoff.

such as the node's microcontroller working when processing packets, or idle listening to the channel. In the case of the CC2420 chipset used in many sensor nodes, the latter accounts for a significant portion of the energy budget.

In reality, understanding the power profile of a network protocol deployment involves many more factors, such as the radio transmit power, the imposed workload, the retransmission policy, the network topology, external interference, and the MAC layer in use and its parameters. Of course, these factors also influence the delivery ratio and the efficiency of the network, which should be taken into account.


To explore some of these factors, we evaluate the Collection Tree Protocol (CTP). CTP serves as an anycast protocol that provides a best-effort, multihop delivery of packets to the root(s) of a tree. We chose CTP since it is the *de facto* standard collection layer in TinyOS. CTP has many mechanisms that can affect energy consumption in non-trivial ways, such as adaptive beaconing, the use of data packets to perform datapath validation and loop detection, and an aggressive policy of retransmissions with no backoff. It also optimizes network paths according to the ETX metric, and measures cost as the average number of transmissions in the network required to deliver each packet to the base station.

2 Evaluation and Results

For our evaluation, we used a subset of 33 testbed nodes deployed on a university open floor lab space, running the default implementation of CTP in TinyOS 2.x (CTP Noe [3]). We instrumented each device with Quanto [2] to track the radio's power states over time, as well as the platform's overall energy consumption given by iCount [1], and obtained detailed network-wide runtime power profiles.

In our experiments, we ran the `TestCollection` application, in which all nodes send packets to the root at the same average rate, properly jittered to avoid synchronization. We set up a single root in one of the corners of the network, and ran each test for 1 hour. We varied the transmission power of the radio and Low-Power Listening (LPL) settings in TinyOS 2.x. When necessary, we also varied the aggregate load to avoid strong congestion in the medium.

Our first experiment has the radio always on, *i.e.*, without Low Power Listening. As we increase transmit power, average path lengths decrease, as longer links become available. We saw little variation in the number of transmissions across nodes and power settings, and even less variation in the total



(a) Measured average power draw per node. (b) Time per radio power state. The sum of the two states is the radio duty cycle. (c) Breakdown of total transmissions per node, a poor predictor of energy usage.

Figure 1. Per-node measurements for CTP with LPL @ 1000ms. The radio duty cycle per node (b) is a very good predictor for measured power. On the other hand, number of transmissions (c) is not.

energy used by the network. In this setting, the radio is on for a very short period for each transmission (on the order of tens of ms), and remains on receive/listen mode for the rest of the time, explaining where most of the energy goes.

Things gets more interesting when we switch to low power listening, as nodes try as much as possible to turn their radios off. The MAC layer is controlled by one parameter, the receive check interval, which dictates how often nodes have to turn their radios on to check for activity in the channel, before going to back to sleep. Because of LPL, at the 500ms setting, the average power draw is significantly reduced to around 25% of the non-LPL case.

However, we also see an increase in the variability of power draw across nodes, with some nodes getting close to the non-LPL power draw! We find that CTP is prone to create these power outliers for different reasons. When the transmit power is low (-25, -15, and -10 dBm), link qualities are poor, and nodes engage in multiple retransmissions. In some cases a node just cannot find a parent, and others nodes become overloaded and queues overflow. When the transmit power increases, overall link quality increases, and we see a dramatic decrease in the total number of transmissions. Interestingly, though, *we do not observe a similar drop in the total energy usage*. In this setting we see outliers appear for another reason: congestion created because of larger radii of interference. A particular node has none of its packets sent for transmit powers above -1.5 dBm. By examining the logs, we noticed the node can hear all of its neighbors, but all neighbors fail to hear it. Thus, most of its receive checks succeed, forcing it to be awake, and all of its transmissions fail, which means they are all maximum length. The node also used the maximum number of retransmissions (30 in CTP) for all of its messages.

To examine this behavior in more detail, we use Quanto to measure how much time the radio spends in each of its power states. The dominant power states are listen/receive and transmit. We set the LPL interval to 1000 ms, to give more weight to transmissions, and decrease the aggregate traffic rate to 1 message/second over the entire network. We found a very strong correlation between the sum of the times the radio spent in listen and transmit modes and the power draw (Figures 1(a) and (b)), as indeed the radio is the largest

component of the power draw in this application. On other hand, the number of transmissions is a poor predictor of the radio duty cycle or the node’s power draw (Figure 1(c)).

3 Conclusion

The strong correlation between radio duty cycle and energy in our last experiment might suggest that the actual measurements are unnecessary. However, this may be an early conclusion. Our application is quite simple in terms of its energy sinks, since only the radio is active. More realistic applications would have other energy sinks like sensors and flash, and would require more complex proxies for energy consumption. More importantly, we see that proxies like “messages sent”, although popular in some circles, cannot be used by itself as a measure of energy efficiency.

Our results suggest that CTP wastes energy in cases when a node loses connectivity. In these cases, all of its transmissions reach the maximum number of retransmissions, which is 30 for CTP. It would be interesting to see the impact of a different policy, such as an exponential backoff for probing whether the node regains connectivity.

We plan to extend our evaluations to other protocols and scenarios, as we believe actual network-wide energy measurements can confirm and/or challenge current assumptions about protocols. We also plan on using our measurements to build models of power consumption based on time profiling of activities, and to assess their accuracy for situations in which it is not possible to have hardware-assisted measurement in all deployed nodes.

4 References

- [1] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler. Energy metering for free: Augmenting switching regulators for real-time monitoring. In *Proc. of IPSN 2008*, pages 283–294, 2008.
- [2] R. Fonseca, P. Dutta, P. Levis, and I. Stoica. Quanto: Tracking energy in networked embedded systems. In *OSDI’08*, pages 323–338, 2008.
- [3] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *Proc. of the 7th ACM Conf. on Embedded Networked Sensor Systems (SenSys)*, Berkeley, CA, 2009.