DUCHY: Double Cost Field Hybrid Link Estimation for Low-Power Wireless Sensor Networks

Daniele Puccinelli and Martin Haenggi
{dpuccine, mhaenggi}@nd.edu

Network Communication and Information Processing Laboratory
Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN 46556, USA

ABSTRACT

We present a novel link estimator for cost-based routing in low-power sensor networks. Our design is hybrid in two different ways: it leverages on both broadcast control traffic and unicast data traffic, and it exploits channel state information and delivery estimates. Broadcast control beacons and data traffic are used to set up a double cost field: an outer field bootstrapped by depth estimates, and an inner field of link estimates. Both sets of estimates are based on control traffic and corrected with feedback from the data plane. We benchmark our link estimation technique and evaluate its performance on Berkeley motes using a public testbed.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing Protocols*; C.2.1 [**Network Architecture and Design**]: Wireless Communication

General Terms

Algorithms, Design, Experimentation, Measurement, Performance

Keywords

Wireless Sensor Networks, Link Estimation, Routing

1. INTRODUCTION

1.1 Motivation and contribution

Lossy links are one of the main challenges of low-power wireless sensor networks (WSNs). Far from Boolean connectivity, low-power wireless links have a very complex behavior due to propagation effects, noise, and interference. The large-scale path loss, shadowing, and multipath fading weaken the signal power; if the attenuation brings the signal strength too close to the noise floor, it becomes impossible to correctly receive packets. Noise-induced losses can be controlled through link estimation, which enables the selection of high quality links. Lower-quality links can also be strengthened through the use of coding and Automatic Repeat Request (ARQ). Coding has not found much use in WSNs due to the computational complexity of decoding, which has been shown to offset the energy savings made possible by the reduction of the transmit power. ARQ schemes, on the contrary, are commonly used. The main problem with ARQ in WSNs is that sensing nodes are typically static: if the channel between two nodes is in a deep fade, it will stay the same until the fading patterns of the deployment area change, which may take an indefinite amount of time. This makes link estimation a critical component of any routing protocol for low-power WSNs: lossy links requiring many retransmissions should be used only if no other option is available. This paper presents the design and experimental evaluation of a novel link estimator for WSN cost-based routing protocols, DUCHY (DoUble Cost field HYbrid link estimator), informed with recent results and observations on the properties of wireless links.

1.2 Notation

We denote the set of all nodes in a given network with $\mathcal{N}$. We assume the presence of only one sink in the network, $s \in \mathcal{N}$. We denote a directed wireless link between nodes i and j as (i, j). This notation indicates that the link is *physical*: if i transmits packets using a given physical layer and a set transmit power, j receives at least one of the packets over a given time period T (or else, (i, j) is said not to exist within T). We define the packet delivery rate (PDR) over (i, j), $\pi_{i,j}$, as the ratio between the number of packets from i received by j and the number of packets sent by i to j over T. Node j is said to be i's neighbor if (i, j) or (j, i) exists.

HotEmNets'08, June 2–3, 2008, Charlottesville, Virginia, USA.

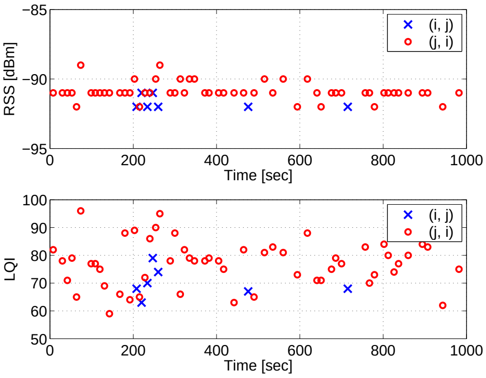

Figure 1: RSS and LQI of a direct and reverse link between TMoteSky motes.

The set of neighbors of s is called the *critical set*. Its members, the *critical nodes*, must relay all upstream traffic to s, and the extra workload drains their batteries at a faster rate; since they are responsible for one-hop data delivery to the sink, their lifetime upper-bounds the network lifetime.

In the context of routing, the parent of i is denoted as $p(i)$. We use the symbol z for an invalid node. We further define the concept of *generalized link*, which we indicate as $[i, j]$, to represent a route between i and j whose intermediate relays are unknown. We denote as $[i, s]_k$ the set of all the relays used by the routing protocol (at a given time) to get i's packets to s using k as the first hop. The network goodput is defined as the number of packets successfully delivered to the sink per time unit, averaged over the course of an experiment and measured in pkts/sec. Finally, we define the hop count of i, h_i, as the number of hops that separate i from s (averaged over the course of a given experiment).

2. LOW-POWER WIRELESS LINKS

The transitional region. Several studies have focused on the properties of low-power wireless links. In [1, 2, 3], it is shown that a *transitional reception region* separates a region with high PDR from a disconnected region. The extent of the transitional region varies dramatically depending on where the nodes are deployed; the transitional region is particularly wide indoors, where multipath fading is more severe. These observations have raised the awareness that standard routing solutions are not viable for low-power WSNs.

Link estimation. Link estimation in WSNs typically relies on either Channel State Information (CSI) from broadcast control traffic or delivery cost estimates from unicast traffic. The most common form of CSI, Received Signal Strength (RSS), used to be considered a poor predictor of link quality, mostly because of the limitations of early platforms [1]. The community has focused on the IEEE 802.15.4 stack, and, in particular, on motes built around the CC2420, which provides a much more reliable RSS indication. The RSS has recently been recognized as a good predictor of link quality; specifically, it has been shown that the RSS, if higher than about -87dBm, correlates very well with the PDR [4]. Another form of CSI, the Link Quality Indicator (LQI) field, is specific to 802.15.4. In the CC2420, the LQI is implemented as the sum of the first 8 correlation values after the Start of Frame Delimiter represented as a 7 bit unsigned integer value.

As for link estimation by means of delivery cost estimates, the ETX (Expected Number of Transmissions) link metric is proposed in [5]; the idea is to estimate the total number of transmissions needed to get a packet across a link, and use the route with the minimum ETX. ETX has been shown to be very robust, especially on top of an ARQ scheme [6]. We note that the traditional way of estimating the ETX of (i, j) as $E_{(i,j)} \approx 1/(\pi_{i,j}\pi_{j,i})$ relies on the assumption that the reception over (i, j) is independent from the reception over (j, i). While this may be reasonable in MANETs due to mobility, it is not true in a typical WSN deployment where nodes are static: losses on the direct and reverse channel are correlated (a similar observation is made in [7]). It is therefore all the more crucial to use 802.15.4's level 2 acknowledgments to measure the ETX, as done in [8], which proposes a hybrid link estimator that fuses the broadcast beacon delivery rate with the Measured Number of Transmissions (MTX) of data traffic, obtained by counting level 2 ACKs.

Link asymmetry. Given an RSS value within the transitional region, there is a very large variance in the corresponding packet delivery rate because a small fluctuation in the signal strength (due for instance to fading or shadowing) or in the noise floor may bring the RSS below the sensitivity threshold of the receiver. Indeed, asymmetric links are concentrated in the transitional region, as observed in [2] and [9]. In Fig. 1, we show how a small difference in the noise floor can make a transitional region link virtually unidirectional. If the noise floor is the same in both the transmitter and the receiver, then this asymmetry is only apparent and due to the fact that the channel is sampled at different times. Forward and reverse link estimates cannot be performed at the same time; depending on the coherence time of the channel, changes in the fading state may occur. Therefore, in such cases the packet delivery rate for the reverse link is different from its counterpart for the forward link as a consequence of the time-varying nature of the channel. Moreover, interference also represents a major reason for link asymmetry [10]. Asymmetric links have been shown to be particularly

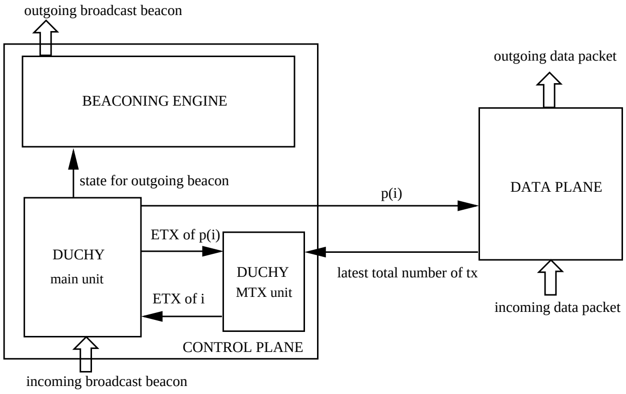

Figure 2: **Structure of DUCHY.**

detrimental to cost-based routing protocols [11], if not properly accounted for.

3. THE DUCHY LINK ESTIMATOR

We propose DUCHY, a novel link estimator designed for network layer integration in cost-based protocols. DUCHY is table-free: it does not maintain a routing table, but only keeps state for the current parent at any given time. We employ RSS to obtain soft information about good links, LQI to get soft information about bad links, and borrow from [8] by using feedback in the form of level 2 ACKs for unicast outgoing data packets. DUCHY is hybrid in two different ways: it is driven by both CSI and PDR estimates, and it leverages on both broadcast control traffic and data traffic (like [2] and [8]). Our design is specific to 802.15.4 (it uses LQI and link-level ACKs), but it can be adapted to non-802.15.4 stacks.

DUCHY includes a main unit and an MTX unit. The former is set up to receive control beacons and extract state information from them. As in all cost-based routing schemes for WSNs, broadcast control beacons are used to set up a distributed cost field. DUCHY builds a *double cost field*: an outer ETX field, bootstrapped with depth estimates and refined using per-link MTX counts as data packets are unicast, and an inner field based on CSI, also refined through MTX. Given the joint use of CSI and data-based information, both fields can be considered to be hybrid. For the inner field, both RSS and LQI are used, as they contain complementary information: RSS tells us how good a good link is, LQI tells us how bad a bad link is. Given link (i, j), we map LQI and RSS measured at j to $[0, 1]$ so that low dBm values of RSS and low raw values of LQI map to high values of RSS cost $r_{i,j}$ and LQI cost $l_{i,j}$. Rather than adding up link costs to obtain route cost, DUCHY employs LQI and RSS bottlenecks [12, 13]. Let us represent the ETX from i to s, $E_{[i,s]}$, as E_i, and the MTX from i to $p(i)$ as M_i. A control beacon from i must contain the following state information: i, $p(i)$, E_i, $B(l_i)$, and $B(r_i)$, where $B(y_i) = \max_{r\in[i,s]_{p(i)}} y_{r,p(r)}$. Upon reception of a beacon from its parent $p(i)$, node i measures $l_{p(i),i}$ and $r_{p(i),i}$ and extracts $B(r_{p(i)})$, and $B(l_{p(i)})$. For the outer field, node i needs to keep $E_i = E_{p(i)} + M_i$.

The MTX unit receives feedback from the data plane of the routing protocol in the form of the latest value of the total number of transmissions x, which it uses to compute M_i as a moving average of N_x samples of x in order to maintain an updated ETX estimate, E_i, which it passes to the main unit. The main unit is also responsible for providing the control plane of the routing protocol with the state information to be inserted in the outgoing beacons.

Suppose node i is parentless, which is the case at startup; with our notation, $p(i) = z$. As a beacon is received from a node, say c, the beacon sender is awarded parent status if it advertises $p(c) \neq z$. The cost of using c to get to s is computed using the inner cost field as

$$C_{i,c} = l_{c,i} + B(l_c) + r_{c,i} + B(r_c). \qquad (1)$$

As i sets $p(i) = c$, no MTX information is available ($M_i = 0$), so ETX is bootstrapped using depth, and node i sets $E_i = E_c + 1$. As i routes data traffic to c, the MTX is computed in the dedicated unit and used in the main unit to refine ETX as $E_i = E_c + M_i$. When a beacon from another node, say w, is received, w is considered as a potential parent. Eligibility is determined based on the outer ETX cost field: node w is eligible for parent status if $E_w < E_i$, *i.e.*, node w offers a more efficient route to s, in the sense that its route needs fewer transmissions. If node w meets this condition, it is tested again using the inner field, which is set up merging local link estimates with global state information diffused by the beacons. This inner hybrid field is used to finalize routing decisions: the cost of using w is computed replacing c with w in (1), and node w is chosen to replace $p(i) = c$ if $C_{i,w} < C_{i,p(i)} + M_i$. Adding MTX to the CSI-based cost helps with asymmetric links, which may show excellent CSI in the upstream direction (direction of control traffic) while being lossy in the downstream direction (direction of data traffic). The double cost field concept inherently promotes routes with a low hop count, which is shown in [14] to be extremely beneficial. In particular, reducing the hop count increases energy efficiency and end-to-end path reliability, improves load balancing, and reduces congestion. Neighbors leading to unnecessarily long routes are filtered out by the outer field, because they have a higher cost in terms of ETX. Using a double cost field also provides an enhanced immunity to routing loops and enhanced route stability: unnecessary route changes are avoided by narrowing down the set of neighbors that are eligible for parenthood.

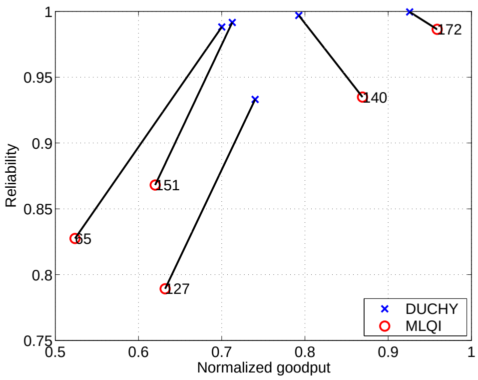

Figure 3: **Gooodput-reliability performance.**

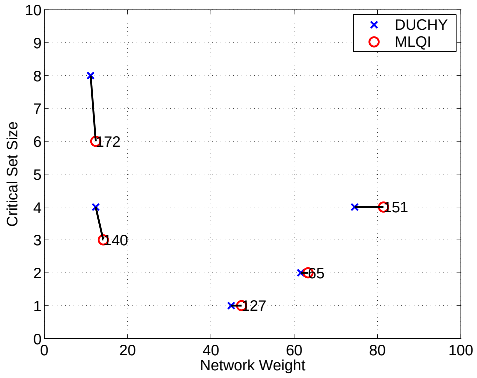

Figure 4: **Quantitative characterization of the topologies used in the experiments.**

4. EXPERIMENTAL RESULTS

The application assumed in this paper is many-to-one data collection: every node generates data packets at the rate $f_{\text{gen}} = 1$pkt/sec with the goal of having them delivered to the sink. We integrate DUCHY into a cost-based protocol, Arbutus, which we are currently developing (an early version is described in [13]). Arbutus implements infinite retransmissions [15]: all packets are assumed to be equally important, and there is no reason why a given packet should be dropped to favor other packets. To avoid retransmissions that are useless due to long coherence times [16], the interval between retransmissions increases linearly with the number of consecutive unacknowledged transmissions. The data plane suspends transmissions if $p(i) = z$, but keeps a FIFO buffer for incoming packets. Congestion is prevented with backpressure: the control plane broadcasts a beacon indicating $p(i) = z$ if the buffer occupancy exceeds a threshold.

As a benchmark, we choose a widely used table-free protocol based on MintRoute [2], MultiHopLQI [1], which only employs LQI for link estimation and implements a maximum of 5 retransmissions; we will refer to MultiHopLQI's link estimation scheme as MLQI. Note that, while our protocol is optimized for reliability (infinite retransmissions), MultiHopLQI is optimized for goodput (it performs 5 retransmissions before moving on). Using MoteLab [17], a public testbed of TMote Sky nodes, we ran 10 experiments of variable duration (between 5 and 15 minutes) at 5 different sink assignments; in this paper we show the average results for each sink assignment. For ease of reference, the sink assignments are 65, 127, 140, 151, and 172. The total number of functioning nodes at the time of the experiments (February 2008) was, on average, about 50 (in disjoint clusters, due to the particular connectivity conditions of the testbed at that time). We ran DUCHY and MLQI back-to-back, and only considered the outcome if the number of reachable nodes was the same for both runs (*i.e.*, no nodes ceased to work during either run).

Fig. 3 shows the reliability-goodput performance of DUCHY and MLQI. As expected, using DUCHY with infinite retransmissions results in a superior reliability performance. DUCHY may also outperform MLQI in terms of goodput; this happens for sink assignments 65, 127, 151. Since changing the sink assignment means dealing with a different network, we characterize a given topology using two quantities: the network weight, defined as $\sum_{i \in \mathcal{N}} h_i$, and the critical set size, defined as $|\mathcal{N}_{1.5}|$, $\mathcal{N}_{1.5}$ being the set of critical nodes i with $1 \leq h_i \leq 1.5$.

Fig. 4 shows where each sink assignment stands in terms of these newly defined topology indicators. We immediately note that sink assignments 65, 127, 151 correspond to a relatively large network weight and a small critical set; we call these *high-pressure topologies*, since the sink only has a few neighbors that relay the load from many other upstream nodes. In particular, sink assignment 127 only has one critical node. On the contrary, sink assignments 140 and 172 are *low-pressure topologies*: the sink has a large number of neighbors that relay the load from a few other nodes. It should be noted that DUCHY never yields a smaller critical set than MLQI; a larger critical set becomes less critical (pressure is lowered) and inherently promotes load balancing. Going back to Fig. 3, DUCHY completely outperforms MLQI (in terms of both reliability and goodput) in the three high-pressure cases, while in low-pressure topologies MLQI yields a slightly better goodput. DUCHY's outer cost field eliminates unnecessarily long routes, while its inner field selects the best links among those that made it through the outer field. For instance, if the choice is between a two-hop route over lossless links or a one-hop route over a link (i, j) with $M_i < 2$, DUCHY chooses the latter, as the former does not pass through the outer field. This way, DUCHY provides Arbutus with all the benefits of a lower hop count.

---

[1] Available at `http://www.tinyos.net/tinyos-1.x/tos/lib/MultiHopLQI`.

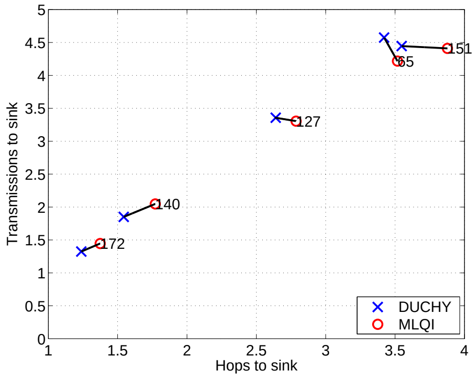

Figure 5: Average number of transmissions needed for delivery to s as a function of the average hop count.

It is, however, fair to argue whether it is DUCHY or Arbutus that outperforms MLQI. Arbutus is definitely instrumental to the reliability performance, and it does contribute to the goodput performance through its use of backpressure to prevent congestion. An extensive evaluation of Arbutus is currently in progress; an early assessment can be found in [13]. Nonetheless, DUCHY's effectiveness is evident in Fig. 5, which shows that, despite the infinite retransmissions implemented by Arbutus, DUCHY does not significantly increase the average number of transmissions per delivered packet in high-pressure topologies and reduces it in low-pressure ones. Moreover, DUCHY consistently yields shallower trees. Over all our experiments, DUCHY outperforms MLQI in terms of reliability in 96% of the nodes. The goodput performance is slightly worse in 20% of the nodes, and better in 60% of them.

5. CONCLUSIONS AND FUTURE WORK

DUCHY outperforms the benchmark, MLQI, in terms of reliability and offers a competitive goodput performance. DUCHY is integrated within Arbutus, our routing protocol optimized for reliability, which employs infinite retransmissions; DUCHY maintains a low number of transmissions per delivered packet (routing cost) by finding the best links, thus avoiding a degradation of the goodput performance.

Given these results, we conclude that WSN link estimation should adopt a hybrid approach; MLQI's sole reliance on one form of CSI is the main reason behind its inferior performance compared to DUCHY. The use of CSI from broadcast beacons is, however, crucial to bootstrapping link estimation; ETX-based feedback is post facto information that is only available once an initial routing choice has been made. The optimization of Arbutus for reliability and energy-efficiency, the impact of the network topology on the routing and load balancing performance, and the scalability of our solutions are the subject of our ongoing work.

ACKNOWLEDGMENTS

The authors would like to thank Geoff Werner-Allen and Matt Welsh for providing public access to their MoteLab testbed. The support of NSF (CNS 04-47869) and CRANE/DTRA (N00164-07-8510) is gratefully acknowledged.

6. REFERENCES

[1] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *1st International Conference on Embedded Networked Sensor Systems (SenSys'03)*, Los Angeles, CA, USA, November 2003.

[2] A. Woo, T. Tong, and D. Culler. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In *1st International Conference on Embedded Networked Sensor Systems (SenSys'03)*, Los Angeles, CA, USA, November 2003.

[3] M. Zuniga and B. Krishnamachari. An Analysis of Unreliability and Asymmetry in Low-Power Wireless Links. *ACM Transactions on Sensor Networks*, 3(2):1–30, 2007.

[4] K. Srinivasan and P. Levis. RSSI is Under Appreciated. In *Third Workshop on Embedded Networked Sensors (EmNets'06)*, Cambridge, MA, USA, May 2006.

[5] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. A High-Throughput Path Metric for Multi-Hop Wireless Routing. In *9th Annual International Conference on Mobile Computing and Networking (MobiCom'03)*, San Diego, CA, USA, 2003.

[6] O. Gnawali, M. Yarvis, J. Heidemann, and R. Govindan. Interaction of retransmission, blacklisting, and routing metrics for reliability in sensor network routing. In *First IEEE Conference on Sensor and Adhoc Communication and Networks*, pages 34–43, Santa Clara, CA, USA, October 2004.

[7] A. Cerpa, J. Wong, M. Potkonjak, and D. Estrin. Temporal Properties of Low Power Wireless Links: Modeling and Implications on Multi-Hop Routing. In *"ACM/IEEE Fourth International Conference on Information Processing in Sensor Networks (IPSN'05)"*, Los Angeles, CA, USA, April 2005.

[8] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis. Four-Bit Wireless Link Estimation. In *Sixth Workshop on Hot Topics in Networks (HotNets-VI)*, Atlanta, GA, USA, November 2007.

[9] A. Cerpa, J. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical Model of Lossy Links in Wireless Sensor Networks. In *" 6th ACM international Symposium on Mobile Ad Hoc Networking and Computing"*, Urbana-Champaign, IL, USA, May 2005.

[10] K. Woyach, D. Puccinelli, and M. Haenggi. Sensorless Sensing in Wireless Networks: Implementation and Measurements. In *Second International Workshop on Wireless Network Measurement (WinMee'06)*, Boston, MA, USA, April 2006.

[11] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic. Impact of Radio Irregularity on Wireless Sensor Networks. In *MobiSys'04*, Boston, MA, USA, June 2004.

[12] B. Chen, K. Muniswamy-Reddy, and M. Welsh. Ad-Hoc Multicast Routing on Resource-Limited Sensor Nodes. In *Second International Workshop on Multi-hop Ad Hoc Networks: from Theory to Reality*, Florence, Italy, May 2006.

[13] D. Puccinelli and M. Haenggi. Arbutus: Network-layer load balancing for wireless sensor networks. In *IEEE Wireless Communications and Networking Conference (WCNC'08)*, Las Vegas, NV, USA, March 2008.

[14] M. Haenggi and D. Puccinelli. Routing in Ad Hoc Networks: A Case for Long Hops. *IEEE Communications Magazine*, 43, October 2005.

[15] M. Wachs, J. Choi, J. Lee, K. Srinivasan, Z. Chen, M. Jain, and P. Levis. Visibility: A New Metric for Protocol Design. In *5th ACM Conference on Embedded Networked Sensor Systems (ACM SenSys'07)*, Sydney, Australia, November 2007.

[16] K. Srinivasan, M. Kazandjieva, S. Agarwal, and P. Levis. The Beta-factor: Improving Bimodal Wireless Networks. Technical Report SING-07-01, Stanford University, 2007.

[17] G. Werner-Allen, P. Swieskowski, and M. Welsh. MoteLab: a wireless sensor network testbed. In *4th international symposium on Information Processing in Sensor Networks (IPSN'05)*, Los Angeles, CA, USA, April 2005.