Connectivity and Energy Usage in Low-Power Wireless: an Experimental Study

Daniele Puccinelli and Silvia Giordano Networking Laboratory University of Applied Sciences of Southern Switzerland CH-6928 Manno, Switzerland firstname.lastname@supsi.ch

Abstract—We present an experimental study of the impact of network connectivity on energy usage in low-power wireless. We focus on the ends of the connectivity spectrum and investigate the energy footprint of excessive connectivity as well as poor connectivity. We observe that high connectivity boosts the energy footprint of broadcast traffic, because it is conducive to contention and overhearing. With respect to poor connectivity, we observe that the presence of connectivity outliers may also result in a significant increase of the energy consumption. To enable fair comparisons between experimental runs, we augment our results with quantitative data regarding the network topology during each run.

I. INTRODUCTION

We present an experimental study of the interplay between low-power wireless connectivity and energy consumption, whose minimization is the top concern of wireless sensor network designers and users. We focus on the key primitive on data collection, which to date remains the primary use of wireless sensor networks. Specifically, we consider the traditional many-to-one converge-cast collection scenario wherein the goal of each node is to get its data delivered to the sink, directly over one hop if possible, or else by way of a multi-hop path across the network. In this paper, connectivity is to be interpreted in a network-wide context and refers to the ability of a given node to deliver data to the sink. If a node has poor connectivity, it means that there exist no stable paths between the node and the sink (and therefore, no stable links to its neighbors). If a node has good connectivity, the opposite is true: there exists at least one stable path between the node and the sink. A stable path is a collection of stable links, and an individual link is considered stable if it has a Required Number of Packets (RNP) [3] below 2 (on average over the course of a whole experiment, less than 2 transmissions are required to get a packet across the link). We use the concept of connectivity because it is more general than other closely related concepts such as node density and node degree. A node can have good connectivity even if the network is not dense, and even if it has a low node degree. Another reason is that it is easy to quantify connectivity in terms of link stability.

We run a tree collection protocol over a low-power link layer and perform an experimental study of the interplay of low-power wireless connectivity and energy consumption. Duty-cycling is essential to low-power wireless: without it, network lifetime will be measured in days as opposed to months or years [4]. At the same time, duty-cycling increases the energy cost of broadcast traffic with respect to unicast traffic [4], [5], and it may prevent nodes from exploiting tidbits of fleeting connectivity, thus making it harder to use lossy links. If nodes have poor connectivity as a baseline, sleeping may cause them to miss out on connectivity opportunities. These are certainly minor matters compared to the enormous benefits of duty cycling, but their impact is unclear, and our goal is to quantify it. As it has been shown in [6], the radio duty cycle correlates well with the energy consumption; therefore, we estimate the energy consumption by measuring the duty cycle of the nodes through online software-based estimation [7].

We use the Arbutus collection protocol [8] because it is optimized for reliability and is immune to network-layer packet loss. After investigating its performance in dutycycled networks in Section II, we employ the protocol as a tool to study the impact of poor connectivity on the energy footprint of unicast traffic in Section IV and the impact of high connectivity on the energy footprint of broadcast traffic in Section III.

II. COLLECTION PERFORMANCE IN DUTY-CYCLED NETWORKS

To illustrate how collection behaves on top of a dutycycled link layer, we begin by showing the results of an extensive set of experiments carried out on MoteLab [9] (offline since Fall 2012 but still available at the time of the experiments) and Twist [10] using Arbutus [8], which we run on top of a Low-Power-Listening (LPL) [11] link layer called BoX-MAC [12]. With BoX-MAC, which is informed with the basic principles of B-MAC [11] and X-MAC [13], all nodes are duty-cycled; the minimum on time is dictated by the minimum time required by the radio to sample the medium, while the theoretical sleep time is defined by the programmer. In the absence of traffic, all nodes transition between a radio sleep state, in which they remain for the duration of the prescribed sleep time, and a radio on state, in which they remain for the minimum on time. Whenever a node wishes to perform a unicast transmission, it unicasts a train of copies of the outgoing packet (in principle, a long strobed preamble) over a time window set to slightly exceed the duration of the sleep time. The train of copies is cut short as soon as a link layer acknowledgment is received. In the case of broadcast traffic, the duration of the train of copies of the outgoing packet must necessarily exceed the duration of the sleep time to ensure that all nodes within radio range have a chance to receive the outgoing packet. For this reason, broadcasts take longer than unicasts and are inherently more energy-costly.

Arbutus [8] is a cost-based collection routing protocol that employs various mechanisms to boost reliability under tough network conditions. Most notably, it uses a hybrid form of link estimation [14] that merges Channel State Information (CSI) with data-driven feedback and sets up a cost field to make routing decisions that account for the vagaries of real-world low-power wireless propagation [1]. CSI is derived from control beacons that are periodically broadcast by each node, while data-driven feedback is obtained in the form of RNP. We assume a standard collection application whereby each node injects data packets at a constant Inter-Packet Interval (IPI). Because our focus is on low data rate operation, we use an IPI of 5 minutes. For BoX-MAC, we choose a sleep time of 1 second. We choose this operating point to replicate one of the low data rate setups in [15]. Arbutus uses a fixed Inter-Beaconing Interval (IBI) with on-demand extra control traffic on top: the interbeacon interval always stays the same, and the extra traffic is injected only when needed. A fixed beaconing interval is, in general, not as efficient as an adaptive beaconing interval [16], but it is useful for our purposes because we can span the control overhead space by modifying the beaconing interval. Specifically, we focus on a high overhead point at IBI=IPI and a low overhead point at IBI=20 IPI (100 minutes), chosen (arbitrarily) to exceed the upper limit (one hour) of the Trickle adaptive beaconing scheme used by CTP. Most of our low overhead runs are from Twist because of its looser constraints on the usage time windows. In Twist we also span the transmit power space by using the highest transmit power setting of the CC2420 (0 dBm) and its lowest setting (-25 dBm). On MoteLab, all experiments are run at 0 dBm because of the relatively low density of the testbed (lower transmit power settings are conducive to network partitioning).

For each experiment, we capture the network connectivity by employing the methodology in [17], which consists in measuring the connectivity matrix of the network *in vivo*, *i.e.*, as it is being used by a protocol under test. Passive network measurements are taken based solely on the protocol's own traffic in order to estimate the expected path delivery from each node to the sink (based on optimal routing choices that maximize the overall delivery to the sink). We employ the principal metric from [17], the Expected Network Delivery (END), to distil the network topology conditions estimated during each run as the average expected path delivery. The END is computed as the average expected path delivery (over all nodes). In essence, the END quantifies the connectivity over the achievable paths to the sink focusing on the key links that keep the network connected and neglecting the redundant links that might as well not exist. The rationale behind the END is that not all links were created equal: a few are essential to keep the network connected, while most of them are redundant. Networks whose key links are stable result in high END topologies, while networks whose key links are volatile and unstable result in low END topologies. Note that the study in [17] only considers an always-on link layer, while in this paper we focus on a duty-cycled link layer. The END is a normalized value within [0, 1]; a network whose key links are challenged will have a low END value, while a network with good key link connectivity will have a high END value. The END makes it possible to tease out the impact of the topology from the impact of the protocol. If the END is low, then we know that the topology is hostile, and we cannot draw any conclusions regarding the quality of the protocol under test. If the END is high, then we know that the topology is benign; if the protocol does not perform as expected, there must be something wrong in the protocol itself.

Twist has a regular grid-like topology and a much higher node density than MoteLab. According to our measurements, the average number of neighbors in Twist is more than four times the number in MoteLab, and the poorest sink assignment in Twist typically has 30% more neighbors with respect to the most well-connected MoteLab sink. Due to Twist's dense connectivity, the sink placement in Twist does not affect the network topology to a significant extent, and the END is therefore generally insensitive to the sink placement: all the sink placements in the Twist experiments resulted in high-end END ranges. On the other hand, because of its complex layout and sparser connectivity, MoteLab displays the same dichotomy between high-performing (high END) and poorly-performing (low END) topologies observed in the always-on experiments in [17].

BoX-MAC sends out packet trains whose baseline length is matched to the sleep time. The packet trains, however, are interrupted upon receipt of an ACK, as in X-MAC [13]. On average, with a low-power link layer, an individual transmission has a much longer air-time than with an alwayson link-layer. For this reason, low-power operation implies greater contention for reduced channel access. Even for topologies with an END above 0.9, low power operation increases the link layer failure rate and takes a toll on reliability. The packet loss averaged over all our runs is 2%, which compares favorably with the 4.9% reported in Table 4 of [15] that corresponds to the same IPI/sleep operating point on MoteLab (CTP's packet loss is higher due to its finite number of retransmissions and the presence of connectivity outliers).

Testbed	END Range	Packet Loss		Goodput	Hops		Cost		Delay		Duty Cycle	
				[pkts/sec]					[s]		[%]	
		mean	σ	mean	mean	σ	mean	σ	mean	σ	mean	σ
MoteLab	[0, 0.7)	0.22	0.13	1.7e-3	3.0	0.7	4.7	1.0	37.1	18.8	6.1	2.8
MoteLab	[0.7, 1]	0.02	0.01	3.1e-3	2.8	0.5	3.7	0.6	4.6	3.9	4.3	1.1
Twist	[0.8, 0.9)	6.7e-3	3.6e-3	3.0e-3	2.3	0.8	2.6	0.9	3.4	3.1	2.5	0.5
Twist	[0.9, 1]	7.7e-3	6.0e-3	3.1e-3	1.8	0.6	2.1	0.7	1.2	2.6	6.5	2.9

 Table I

 PERFORMANCE OF ARBUTUS WITH BOX-MAC (LPL) AT IPI=300s ON MOTELAB AND TWIST AVERAGED OVER DIFFERENT END INTERVALS.



Figure 1. END vs. duty-cycle in the MoteLab and Twist experiments with BoX-MAC.



Figure 2. A handful of connectivity outliers drastically increase the through traffic (and the packet loss) in MoteLab.

Table I reports a detailed breakdown of our results for specific END ranges. We note a large discrepancy between the duty cycles on MoteLab and Twist, also evident in Figure 1, which shows the END vs. the duty cycle. While the mean duty cycle in MoteLab is 4.3% in the high END regime and 6.1% in low END regime, the mean duty cycle in the Twist runs (all at high END regime) is as high as 5.8%. In MoteLab, a higher END means a lower duty cycle; the duty cycles in high-END Twist, however, are much closer to MoteLab's low-END duty cycles than they are to MoteLab's high-END duty cycles. This is due to the interplay between the topology and the overhearing effect, a major duty cycle driver at ultra-light offered load points (such as our operating point, IPI=300s).



Figure 3. The empirical CDF of the duty cycle for the two runs in Figure 2 confirms the impact of the outliers in MoteLab.



Figure 4. The empirical CDF of the duty cycle for these two Twist runs shows that even a single outlier can have a huge energy impact.



Figure 5. A poorly connected leaf can operate the equivalent of a denialof-sleep attack to the victim, which may be a well-connected hub. The crosses indicate the low-PRR links.

III. HIGH CONNECTIVITY: BROADCAST IS COSTLIER

To study the effect of dense connectivity on the energy footprint of the broadcast traffic, we vary the transmit power: lowering the transmit power reduces the number of available links in the network. In Twist, our measurements indicate that, when the transmit power is reduced from 0 dBm to -25 dBm, the average number of stable neighbors drops by 75%. Denser or sparser does not necessarily mean higher END or lower END, because the END captures the availability of the key links: a sparse network whose key links are solid may very well have the same END as a very dense network. A lower transmit power typically means longer paths and higher costs (more transmissions per successful delivery), but the performance can be expected to be stable if the END remains high in spite of the transmit power reduction. On the other hand, the extra energy required by the longer paths may very well be compensated for by the fact that the nodes have fewer neighbors and are less exposed to overhearing, which has a direct impact on the LPL duty cycle. For this reason, although the reduction of the transmit power does not yield proportional energy savings on the transmitter side [18], it may be conducive to significant energy savings on the receiver side. Given the IPI, the overhearing is mostly driven by the control overhead, also because the broadcast nature of control traffic makes it more energy-costly as a baseline [4].

Table II provides a breakdown of the parameter space that we explored in our experiments in the two dimensions of transmit power (high at 0 dBm and low at -25 dBm) and control overhead (high at IBI=IPI and low at IBI=20 IPI); the MoteLab results are only broken down in the dimension of control overhead. The Twist experiments were performed using two different transmit power settings (high at 0 dBm and low at -25 dBm) and beaconing rates (high at IBI=IPI and low at IBI=20 IPI). The MoteLab experiments were all performed using one power setting (high at 0 dBm) and only explore different beaconing rates (the same two rates as in the Twist runs). All MoteLab experiments are run at 0 dBm because of the relatively low density of the testbed (lower transmit power settings are conducive to network partitioning). We put the MoteLab results into the proper topological context by leveraging the methodology in [17]: in order to ensure a fair comparison between the MoteLab and the high-END Twist runs, we separate out the MoteLab runs based on the END in Table II.

The results in Table II provide a clear indication of the interplay of overhearing, contention, and connectivity. A higher beaconing rate means higher overhead, more overhearing, and more contention, while a higher transmit power means more connectivity but also more overhearing and more contention.

In MoteLab, at high transmit power, the impact of the control overhead (beaconing rate) is significant: the duty

cycle ranges from just 1.4% at low overhead to over three times as much at high overhead and almost five times as much at high overhead and low END. Aside from the duty cycle, the packet loss, the transmissions per successful reception (cost), and the delay also deteriorate at high overhead. In particular, the packet loss is four times higher at high overhead than it is at low overhead, and the reason for this is that the extra broadcast traffic competes with data traffic for channel access (a scarce resource in low-power operation) and makes link-layer failures more likely. In Twist, at high transmit power the impact of the control overhead is also considerable: the duty cycle ranges from just below 3% at low overhead to as much as 9% at high overhead. At high overhead, the packet loss doubles but the other performance dimensions do not suffer considerably (the mean delay does, but the median delay remains constant). At low transmit power, the performance is not sensitive to the overhead in the dimensions of packet loss, cost, and delay; the impact of the overhead on the duty cycle is relatively small compared to its impact at high transmit power.

The joint effect of the transmit power and the overhead is the main driver of energy consumption through overhearing at high END. This result suggests that more connectivity is not necessarily better. A lower transmit power typically reduces the connectivity and increases the path length, and therefore the number of transmissions per delivery. However, it also reduces contention and overhearing, resulting in a lower duty-cycle because overhearing dominates the energy consumption by the radio on a node, even with a low-powerlistening link layer. Control overhead over stable routes is not only unnecessary, but also detrimental in energy and reliability terms.

The side effects of Twist's denser connectivity, exacerbated by the control overhead, are the reason behind the high duty-cycles recorded at high END values in Figure 1. Indeed, Table II shows that the duty cycle is highest in the runs at high transmit power and high control overhead, while relatively normal (below 5%) in all other runs.

The energy footprint of high connectivity could be drastically reduced with the use of a broadcast-free collection protocol. A practical solution for broadcast-free collection has been recently proposed and evaluated by the authors in [5].

IV. POOR CONNECTIVITY: THE ROLE OF THE OUTLIERS

Nodes with no stable links to the rest of the network represent connectivity outliers. Their poor connectivity has a significant energy footprint, particularly in the presence of unicast traffic. When dealing with unicast traffic, such nodes may require numerous retransmissions to communicate with any of their radio neighbors. To study the impact of connectivity outliers on the energy footprint of unicast traffic, we performed a set of runs on Motelab with a high END sink assignment (node 114). We singled out the connectivity outliers based on their average RNP: nodes that required, on average, at least 10 transmissions per successful delivery were labeled as outliers. Figure 2 shows the overall amount of traffic per time unit (packets/sec) that is sent over the air in the MoteLab network with and without the connectivity outliers. Figure 2 also displays the lost traffic, *i.e.*, the injected packets that are dropped due to a link layer failure. As shown in [8], Arbutus eliminates all causes of packet loss [19] other than link layer failures. The difference in the total traffic through the network caused by the 16 outliers is impressive: with the outliers there are 12.8 injected packets for every successful delivery, while without them there are only 3. Without the outliers the total network traffic decreases drastically, and the packet loss also decreases by almost 50% (from 1.73% to 0.92%): the outliers make link layer failures much more likely.

The energy impact of the outliers is also very large, as the duty-cycle distribution for the two runs shows (Figure 3). The outliers are responsible for an average increase of the duty-cycle of more than 30%. The average duty-cycle of the outliers is 6.6%, while the average duty-cycle of the rest of the nodes in the presence of the outliers is 4.17%, which is an extra 24% with respect to their average duty cycle in the absence of the outliers. The implication is that the outliers, in spite of their apparently harmless leaf status, drain a quarter of the energy reserve of their peers.

Log inspection showed that leaf outliers are harmful mostly through outbound loss: the outliers flood their upstream peers with duplicates. The poorly connected leaves keep their one-hop neighbors listening, thus creating the equivalent of a denial-of-sleep attack. This is particularly detrimental in the likely event that one of the neighbors is a well-connected hub or a bottleneck, as is the case in the example in Figure 5.

We also found a leaf outlier in Twist and turned it on and off to isolate out its impact. Figure 4 shows the empirical distribution of a Twist run with the outlier on and one without it (-1), and we find that a single outlier boosts the duty cycle by over 5%.

Because outliers with predominantly outbound loss typically have very high duty cycles, a node can easily selfdiagnose its outlier status. Including a drastic measure such as a duty cycle-based auto-shutdown mechanism in a networking protocol for sensor networks may be an advantageous tradeoff: the outlier is sacrificed to extend the lifetime of its neighbors.

V. RELATED WORK

In recent years, low-power wireless connectivity has been the object of extensive investigations. Transitional connectivity has been studied in [2], and its protocol-level implications have been investigated in [20] and [21]. The methodology from [17] that we employ to compute the END value that characterizes the network topology during our experiments is part of the body of work for the characterization of the topology of low-power wireless. The work in [17] is significantly different from other studies that also pursue a low-power *wireless lexicon* [22], such as [23] and [24], because it addresses the *in vivo* characterization of the network topology as the network is in use. Collection protocols have been the object of several investigations [20] [15] [8]. The behavior of collection on top of a duty-cycled link layer has been studied in [25], [15], [26], [27], and [28], and the energy footprint of collection is the specific focus of [6]. The energy footprint of broadcast in the presence of duty-cycling is illustrated in [4], and the relative energy footprints of broadcast and unicast traffic are discussed in [29]. Recently, the operation of a broadcast-free collection protocol and the related energy benefits have been illustrated in [5].

VI. CONCLUSION

Poor connectivity has a significant energy footprint. Not only do connectivity outliers drain their own resources, but they also drain the resources of their neighbors. The impact of poor connectivity may only be mitigated with drastic measures, such as auto-shutdown mechanisms for nodes that self-diagnose their connectivity outlier status. Overly redundant connectivity also has a significant energy footprint, because it increases the cost of broadcast traffic (control traffic). The impact of excessive connectivity may be mitigated with transmission power control and the aggressive reduction of unnecessary control traffic, as shown in [5].

ACKNOWLEDGEMENTS

This work was partially supported by the European Commission under the SCAMPI Project (ICT grant agreement 258414).

References

- K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. An Empirical Study of Low-Power Wireless. ACM Transactions on Sensor Networks (To appear), 2010.
- [2] 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.
- [3] 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.
- [4] A. Dunkels, L. Mottola, N. Tsiftes, F. Österlind, J. Eriksson, and N. Finne. The Announcement Layer: Beacon Coordination for the Sensornet Stack. In *European Conference* on Wireless Sensor Networks (EWSN'11), Bonn, Germany, February 2011.
- [5] D. Puccinelli, M. Zuniga, S. Giordano, and P. Marrón. Broadcast-Free Collection Protocol. In 10th ACM Conference on Embedded Networked Sensor Systems (SenSys '12), Toronto, Canada, November 2012.
- [6] M. Martins, R. Fonseca, T. Schmid, and P. Dutta. Poster Abstract: Network-Wide Energy Profiling of CTP. In 8th ACM Conference on Embedded Networked Sensor Systems (SenSys'10), Zurich, Switzerland, November 2010.

Testbed	Power	Overhead	Packet Loss	Hops	Cost	Delay		Duty Cycle
						μ [s]	med.[s]	[%]
MoteLab (high END)	high	low	5e-3	2.4	2.9	9.7	0.2	1.4
MoteLab (high END)	high	high	0.02	2.8	3.7	4.3	0.9	4.4
Twist	high	low	2.4e-3	1.5	1.6	0.14	0.03	2.9
Twist	high	high	5.6e-3	1.4	1.7	0.93	0.03	9.0
Twist	low	low	9.1e-3	2.7	3.0	3.5	0.3	2.4
Twist	low	high	0.017	2.7	3.2	3.1	0.3	3.9

Table II

PERFORMANCE BREAKDOWN: TRANSMIT POWER (HIGH CORRESPONDS TO 0 DBM AND LOW CORRESPONDS TO -25 DBM) AND OVERHEAD (HIGH CORRESPONDS TO IBI=IPI AND LOW CORRESPONDS TO IBI=20 IPI) FOR RUNS IN THE END RANGE [0.8, 1].

- [7] A. Dunkels, F. Osterlind, N. Tsiftes, and Z. He. Softwarebased on-line energy estimation for sensor nodes. In Proceedings of the Fourth Workshop on Embedded Networked Sensors (Emnets IV), Cork, Ireland, June 2007.
- [8] D. Puccinelli and M. Haenggi. Reliable Data Delivery in Large-Scale Low-Power Sensor Networks. ACM Transactions on Sensor Networks, Jul. 2010.
- [9] 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, April 2005.
- [10] V. Handziski, A. Koepke, A. Willig, and A. Wolisz. TWIST: A Scalable and Reconfigurable Testbed for Wireless Indoor Experiments with Sensor Networks. In 2nd International Workshop on Multi-hop Ad Hoc Networks: from Theory to Reality (REALMAN'06), Florence, Italy, 2006.
- [11] J. Polastre, J. Hill, and D. Culler. Versatile Low Power Media Access for Wireless Sensor Networks. In 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys'04), Baltimore, MD, November 2004.
- [12] D. Moss and P. Levis. BoX-MACs: Exploiting Physical and Link Layer. Technical Report 08-00, Stanford University, 2008.
- [13] M. Buettner, G. Yee, E. Anderson, and R. Han. X-MAC: a short preamble MAC protocol for duty-cycled wireless sensor networks. In 4th ACM Conference on Embedded Networked Sensor Systems (SenSys'06), Boulder, CO, USA, November 2006.
- [14] D. Puccinelli and M. Haenggi. DUCHY: Double Cost Field Hybrid Link Estimation for Low-Power Wireless Sensor Networks. In *Fifth Workshop on Embedded Networked Sensors* (*HotEmNets* '08), Charlottesville, VA, USA, 2008.
- [15] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection Tree Protocol. In 7th ACM Conference on Embedded Networked Sensor Systems (SenSys'09), Berkeley, CA, November 2009.
- [16] P. Levis, N. Patel, D. Culler, and S. Shenker. Trickle: A Self-Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks. In *First USENIX/ACM Symposium on Networked Systems Design and Implementation* (NSDI'04), San Francisco, CA, March 2004.
- [17] D. Puccinelli, O. Gnawali, S. Yoon, S. Santini, U. Colesanti, and L. Guibas. The Impact of Network Topology on Collection Performance. In 8th European Conference on Wireless Sensor Networks (EWSN'11), Bonn, Germany, February 2011.
- [18] M. Haenggi. The Impact of Power Amplifier Characteristics on Routing in Random Wireless Networks. In *IEEE 2003 Global Communications Conference (GLOBECOM'03)*, San Francisco, CA, December 2003.
- [19] 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 (SenSys'07), Sydney, Australia, November 2007.

- [20] A. Woo, T. Tong, and D. Culler. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In 1st ACM Conference on Embedded Networked Sensor Systems (SenSys'03), Los Angeles, CA, November 2003.
- [21] 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 Ad Hoc Communication and Networks (SECON'04)*, pages 34–43, Santa Clara, CA, October 2004.
- [22] K. Srinivasan. Towards a Wireless Lexicon. PhD thesis, Stanford University, September 2010.
- [23] K. Srinivasan, M. Kazandjieva, S. Agarwal, and P. Levis. The Beta-Factor: Improving Bimodal Wireless Networks. In 6th ACM Conference on Embedded Networked Sensor Systems (SenSys'07), Raleigh, NC, November 2008.
- [24] K. Srinivasan, M. Jain, J. Choi, T. Azim, and E. Kim. The Kappa Factor: Inferring Protocol Performance Using Interlink Reception Correlation. Technical Report 09-02, Stanford University, 2009.
- [25] L. Filipponi, S. Santini, and A. Vitaletti. Data Collection in Wireless Sensor Networks for Noise Pollution Monitoring. In International Conference on Distributed Computing in Sensor Systems (DCOSS'08), Santorini, Greece, June 2008.
- [26] L. Mottola, G. P. Picco, M. Ceriotti, S. Guna, and A. Murphy. Not All Wireless Sensor Networks Are Created Equal: A Comparative Study On Tunnels. ACM Transactions on Sensor Networks, 7, August 2010.
- [27] J. Vanhie-Van Gerwen, E. De Poorter, B. Latré, I. Moerman, and P. Demeester. Real-Life Performance of Protocol Combinations for Wireless Sensor Networks. In *IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC'10)*, Newport Beach, CA, USA, June 2010.
- [28] D. Puccinelli and S. Giordano. Duty-Cycle Distribution in Low-Power Collection. In European Conference on Wireless Sensor Networks (EWSN'10), Coimbra, Portugal, February 2010.
- [29] G. Halkes and K. Langendoen. Practical considerations for wireless sensor network algorithms. *Wireless Sensor Network*, 2(6):441–446, June 2010.