

Empirical Models for Predicting Radio Link Quality in Outdoor Deployment Environments

Sally K. Wahba
 School of Computing
 Clemson University
 Clemson, SC 29634-0974
 Email: sallyw@cs.clemson.edu

Jason O. Hallstrom
 School of Computing
 Clemson University
 Clemson, SC 29634-0974
 Email: jasonoh@cs.clemson.edu

Abstract—In this paper, we present two environment-specific models for predicting radio link quality in embedded wireless systems as a function of radio transmission power and inter-node distance. The models are empirically-based and developed using regression analysis. The underlying data was collected from over 1400 experiments conducted in open grass fields and dense forest environments using *Tmote Sky* nodes. We focus on this hardware platform due to its popularity in the domain of wireless sensing. Our models predict radio link quality in typical outdoor deployment environments, and achieve a goodness of fit of over 0.83.

Keywords—Radio link quality modeling; wireless sensor networks; embedded network systems; radio link quality prediction.

I. INTRODUCTION

Large-scale embedded network systems have moved from imagination to reality. Applications of these systems vary from social networking to saving lives in the battlefield [3, 7, 15], and the domain is still evolving.

Although the community is growing, and interest in the field is increasing, the domain is still in its infancy. Developing embedded applications that behave as expected is a challenge. The main issue is the lack of tools available at the design stage of the application life cycle to help developers implement sound applications. In particular, there are few tools to assist in predicting radio link quality within an embedded network. This in turn leads to designers implementing applications that suffer from unpredictable and undesirable wireless performance [12].

We present two empirical models of radio link quality based on experiments conducted in common deployment environments (*i.e.*, an open grass field and a dense forest). We use *packet reception rate* (PRR) as the link quality metric, defined as the ratio of the number of messages received to the number of message sent. Using regression analysis, the resulting models predict the radio link quality within an embedded network based on transmission distance and radio power level. The models are based on data collected using *Tmote Sky* motes, a widely adopted hardware platform in the domain of wireless sensing.

Our models are different from existing empirical models in two significant ways. First, existing models predict link quality based on data collected from indoor testbeds [6].

Link quality in indoor deployments often varies significantly from link quality in outdoor deployments [9]. Our models rely on empirical data collected from open field and forest environments to predict link quality in such environments. These environments represent a large number of wireless sensor deployments. Second, existing models require users to perform their own empirical analyses (*e.g.*, measuring signal decay, noise fading), which limits ease-of-use [5, 6, 17]. Our models are more straightforward. Users select the deployment environment (*i.e.*, open grass fields and dense forests), radio power level, and inter-node distance. The models are then used to predict link quality based on these parameters.

The remainder of the paper is organized as follows. Section II discusses elements of related work. Section III presents the process for collecting data for our models. Section IV describes the process of filtering the data. Section V presents the link quality models. Finally, Section VI concludes and discusses future work.

II. RELATED WORK

Seada *et al.* [10] study various energy-efficient forwarding strategies for routing in lossy wireless sensor networks. The authors conclude that PRR is a good metric for evaluating channel conditions. Accordingly, we use PRR as the link quality metric in our models.

Liu and Cerpa [6] use a three-phase model for predicting the estimated link quality between nodes. The authors collect PRR data from two indoor testbeds. The data is used as input for training a prediction model using machine learning. This model is then used by the network to adjust message routing. This approach suffers from two main drawbacks. First, the authors rely on indoor data to train their model, which may render the model inaccurate for outdoor deployments. Second, the authors require users to provide their own models for various network conditions.

Kamthe *et al.* [4] describe a statistical approach to modeling link quality variation over time. The authors collect PRR data from a testbed in an indoor lab. The data collected is used as a training set for a learning algorithm used to predict packet reception over time. The authors use two statistical models simultaneously, one for modeling short-term link quality dynamics, and another for long-term dynamics.

This work is different from ours in that the authors provide a statistical model to describe the *correlation* between successive packet failures and receptions. Our work provides a model for predicting radio link quality (*i.e.*, the probability that a packet sent on a link will be received).

Xu and Lee [16] present a regression-based algorithm for on-line estimation of link quality using spatial correlation of nodes. While the model can be used to dynamically adjust routing protocols, it relies on the network being deployed before estimating link quality. The drawback of this approach is that a better network layout can be achieved given a priori knowledge of link quality.

Cerpa *et al.* [1] provide a probability density function that characterizes the relationship between distance and link quality. The authors use the absolute physical location of nodes on a grid, as well as the relative physical proximity of nodes and their neighbors to represent distance. One drawback of the model is that the authors do not account for physical obstructions, although the data was collected from experiments run in environments that include physical obstructions.

Finally, Wahba *et al.* [14] provide a model for predicting radio link quality as a function of inter-node distance and radio transmission power. The model is limited to open grass fields that are obstruction- and interference-free. With the introduction of TinyOS 2.x, a new radio MAC protocol was adopted, which renders the TinyOS 1.x-based model obsolete. Consistent with the findings of Cerpa *et al.* [2], the authors observe that link quality generally falls into three regions: low, mid, and high. Low- and high-quality regions tend to be stable over time; mid-quality regions tend to be unstable. In other words, when PRR is in the high- or low-quality region, it tends to be immune to temporal changes and minor equipment adjustments (*e.g.*, distance and orientation). However, when PRR is in the mid-quality region, it exhibits significant variability, by as much as 100%.

III. DATA COLLECTION

Our work began with preliminary studies to control for factors that are not included as independent variables in our models. We conducted the first study to understand the noise floor in the environments where the experiments were conducted. This allowed us to select the radio channel with the least noise variability. In the second study, we investigated the effect of height differences between communicating devices (because we place the motes on risers). In the third study, we investigated the effect of device orientation on communication quality. The last study involved investigating the transmission rate used in our experiments to ensure that the chosen rate did not lead to network saturation and packet loss. The results of these studies appear in [13]. We use these results in the data collection process to ensure data quality.

To collect the empirical data for our models, we developed two nesC applications designed to run on three motes. The first, DC_1, is designed to run on a sender, S, and a receiver, R. The second application, DC_2, is designed to run on a noise floor data collector, N, and samples the RSSI register

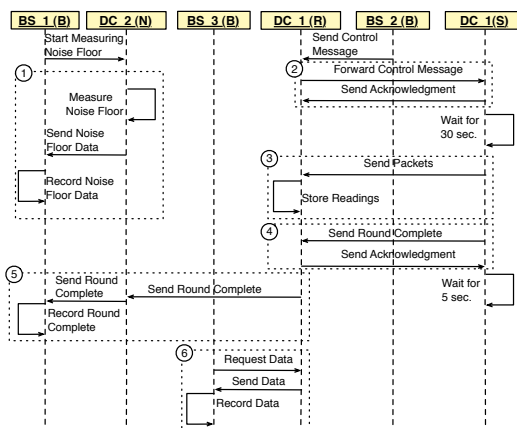


Fig. 1: Data Collection Sequence Diagram

during data collection. We measured noise floor during our experiments to ensure that there was no significant interference affecting the quality of our data.

We also developed three Java applications designed to run on a basestation, B. The first, BS_1, controls noise floor data collection. The second, BS_2, controls the main transmission experiment. The third, BS_3, controls data collection after the transmission experiment is complete.

To run an experiment, we first install DC_1 on S and R, and DC_2 on N. S, R, and N wait for a control message from B to start execution. (N and R are connected to the basestation via serial.) Figure 1 shows a sequence diagram representing the data collection process. The numbers in the figure highlight operations that can be repeated, as explained later. (It is worth mentioning that operations on the timeline are not to scale.) We first run the BS_1 application, which sends a serial message to N to start measuring the noise floor. Upon receiving the message, N begins to continuously measure the noise floor on the channel used by S and R, and sends the information to B. Operation (1) repeats for the duration of the experiment. N also marks the end of transmission for each experimental configuration for S and R. (This is discussed later.)

BS_2 and BS_1 begin at the same time. BS_2 sends a serial message to R with the parameters for the experiment, and then terminates. This message includes the following information: (i) initial radio power level, (ii) final radio power level, (iii) transmission rate, and (iv) period duration (for each power level). R, in turn, sends this control message via radio to S. Operation (2) repeats until R receives an acknowledgment. S then waits for 30 seconds before it starts sending the desired messages for the first radio power level. After the 30 seconds are over, operation (3) repeats. In this operation, S sends radio messages at the specified transmission rate for the specified duration. Throughout, R counts the number of messages received. When S finishes sending its messages for the first radio power level, it sends a message to R indicating that it has finished that power level. Operation (4) repeats until S receives an acknowledgment from R. R then saves the information associated with that power level in a corresponding array location and clears its counters. R

also sends a message to N indicating that the messages for the first power level are complete. N, in turn, sends a serial message to BS_1, which records an end of round marker. S then waits for 5 seconds before it starts sending the next sequence of messages using the next radio power level. These steps repeat until the final power level is reached. At this time, N continues to measure the noise floor for another 30 seconds before it terminates. After noise floor data collection is complete, BS_1 terminates. Accordingly, operations 1, 3, 4, and 5 are repeated for each experimental setup (*i.e.*, distance). Operation 2 executes at the beginning of each trial.

To communicate the information stored on R to B, BS_3 executes. This program sends a serial message to R to instruct it to send the information for each power level. R, in turn, sends this information; the basestation saves it to a file and terminates. Operation (6) represents the steps performed at the end of each trial to collect all data points stored on R.

It is worth noting that each data point represents the PRR for a certain radio power level and inter-node distance value (within a specific deployment environment). For each data point (PRR), S sends 30 messages per second for 30 seconds.

A. Experimental Parameters

Here, we summarize the experimental parameters considered in our data collection experiments.

Radio Power Level. We ran the experiments using all available power levels (1 – 31), in increments of 1 unit. These values correspond to non-linear changes in dBm, ranging from less than -37 to 0 dBm, respectively [11].

Deployment Environment. We ran the experiments in an open grass field at Clemson University, and in a forest environment in the Clemson Experimental Forest. Figure 2 shows a typical deployment in the open grass field. Figure 3 shows a typical deployment in the forest. Note from the figures that S, R, and N were placed on 4-foot wooden stakes to avoid the effects of physical obstructions from the grass. Further, N and R were placed 6 inches apart. All nodes had the same orientation. For the experiments conducted in the forest, we ensured that the nodes were placed within a clear line of sight. We chose these two environments as they are typical outdoor deployment environments. The open grass field represents an outdoor environment that is free of physical obstructions, while the forest represents an outdoor environment containing physical obstructions (*i.e.*, trees).

Inter-node Distance. The distance between nodes was varied between 1 – 416 feet. Note that according to [11], the radio range is 410 feet. The distance increment applied in each step was based on the radio power level. For power levels 1 – 3, the distance was varied in increments of 2 feet. For power levels 4 – 16, the distance was varied in increments of 8 feet. For power levels 17 – 31, the distance was varied in increments of 16 feet. Without the increment variation, the number of experiments would have been prohibitively large.

Figures 4 and 5 illustrate the distances covered for each power level in the field and the forest, respectively. The distances covered in the forest are relatively sparse. From




Fig. 2: Field Deployment for Data Collection



Fig. 3: Forest Deployment for Data Collection

the experiments, no messages were received beyond 256 feet, as opposed to 416 feet in the open grass field. Accordingly, the furthest distance covered was 280 feet. In the dense forest, experiments were conducted over 20% of the distances covered in the grass field.

IV. DATA PROCESSING

After all of the experimental data was collected, we applied a filtering process to ensure data validity. The first step was to eliminate outliers; the second was to eliminate data points collected during noise spikes.

We explain the process of eliminating outliers with example data shown in Figure 6. If a data point was different from the two “surrounding” data points by more than 20%, the experiment was repeated. For example, in Figure 6a, the PRR data point at a distance of 4 feet differs from the preceding distance (2 feet) and the succeeding distance (6 feet) by more than 20% (surrounding distances were within 20% of one another). Hence, the experiment at distance 4 feet was




Fig. 4: Distances Covered in the Field




Fig. 5: Distances Covered in the Forest

repeated. If the new data was consistent with the surrounding data, the old data was replaced with the new data, as shown in Figure 6b. However, if the new data was still inconsistent with the surrounding data points, two more experiments mid-way between the outlier data point and the surrounding distances were run (*i.e.*, at distances of 3 and 5 feet), respectively. If the data points were within 20% of distances 2 and 6, we discard the data at distance 4, but keep the data from distances 3 and 5. This is the case in Figure 6c. However, if the data was still inconsistent, all three data points were kept (*e.g.*, distances 3, 4, 5 feet), as these data points were not outliers. This is the case in Figure 6d. This approach resulted in 67 reruns and 21 outliers in the field, and 17 reruns and 2 outliers in the forest.

When examining the noise floor data, we noticed that some experiments were associated with noise spikes, suggesting external interference in the environment. When examining some of the corresponding PRR data, we noticed that they were inconsistent with the PRR data collected from the surrounding distances. At these instances, the noise spikes were more than twice the standard deviation of the noise samples.

From the noise floor data and the associated PRR data collected, it was not feasible to determine whether the spikes in the noise floor occurred at the same time a transmission was sent/received, hence affecting the collected PRR data. In other words, we could not adjust PRR to account for noise spikes. Accordingly, we discarded the PRR data collected when spikes

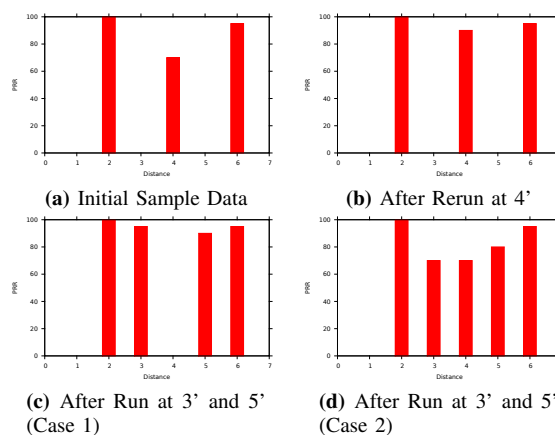


Fig. 6: Filtering Data based on Experiment Reruns




Fig. 7: Example of the Noise Filtering Process

in the noise floor occurred.

To eliminate the data associated with spikes in the noise floor, we implemented a Java application that processed all noise samples collected during our experiments. If we found a noise spike above a certain threshold (*i.e.*, twice the standard deviation above the series), we discarded the PRR associated with the experiment during which the noise spike occurred.

To determine a noise spike, we consider all the noise floor measurements collected during a single experimental setup (*i.e.*, one inter-node distance and radio power level). We describe the process of determining a noise spike through an example. Figure 7 represents an example of 100 noise floor samples. The X-axis represents the sample number. The Y-axis represents the noise floor measurement in dBm. The data is divided into segments of 10 noise floor samples each, resulting in a total of 10 segments. For the sake of exposition, we refer to these segments as segments 1 to 10. The cumulative average and standard deviation are maintained, from the start of the series through the end of the last processed segment. For example, when processing the noise floor samples in segment 4, the average and standard deviation for all data points from segments 1 – 3 are maintained. After reading the noise floor samples associated with segment 4, the average of the noise floor samples in that segment is calculated (*i.e.*, the average across noise floor samples 31 – 40). If the average of the noise floor samples in segment 4 is greater than the cumulative




Fig. 8: Data Collected from an Open Grass Field




Fig. 9: Data Collected from the Experimental Forest

average plus twice the cumulative standard deviation, a noise spike is identified. When processing segment 4 in Figure 7, the cumulative average (*i.e.*, the average for segments 1 – 3) is -98.33 dBm, and the cumulative standard deviation is 0.479. (The minimum value the RSSI register can store is ≈ -100 dBm [11].) The average of the noise floor samples in segment 4 is -97.4 dBm, which is greater than the cumulative average plus twice the cumulative standard deviation. Accordingly, a noise spike is detected. This is consistent with the noise floor samples in the figure, since the noise floor at sample 35 is -90 dBm, which is indeed a spike in the noise floor. At this point, when a noise spike is identified, the PRR data associated with the experimental setup is discarded. This process resulted in 56 data points being discarded from the field experiments and 4 data points being discarded from the forest experiments.

After filtering the data using both procedures, we were left with a total of 1,211 and 213 data points for the field and forest, respectively.

V. LINK QUALITY MODELS

We used the processed data to design two environment-specific empirical models that predict link quality as a function of radio power level and inter-node distance. Figures 8 and 9 summarize the processed data collected from the field and forest, respectively. Consistent with our findings from [14], the link quality falls into three regions – high-, mid-, and low-quality. We characterize the high-, mid-, and low-quality regions as the regions where PRR is between 90% – 100% inclusive, 90% – 10% exclusive, and 10% – 0% inclusive, respectively. Figure 10 shows a comparison of the data collected from the field versus the data collected from the forest. Notice that the data is consistent in the high- and low-quality regions. However, in the mid-quality region, some data points in the forest have a higher PRR than their counterparts




Fig. 10: Field vs. Forest Data

in the field and vice-versa. We suspect the mid-quality region as the cause.

Using one algebraic model to predict link quality is not feasible given the variation among regions. Accordingly, for each deployment environment, we provide a model for the low-quality region, and another for the high-quality region. (Designers usually try to avoid the mid-quality region due to its increased variability [2].) To determine the data points needed for each model, the data was processed as follows. For each distance, we identified the value *a*, corresponding to the highest power level that resulted in a low PRR. Similarly, we identified the value *b*, corresponding to the lowest power level that resulted in a high PRR. For the field data, this resulted in 60 and 63 samples for *a* and *b*, respectively. The number of *a* and *b* data points was different because at shorter distances, some *b* values did not have corresponding *a* values. In some cases, PRR was in the high-quality region from the lowest radio power level. For the forest data, this resulted in 11 samples for both *a* and *b*.

For each deployment environment, we used the *a* values and linear regression analysis to determine the appropriate formula for the low-quality region. The process resulted in the following formulae:

$$Field : power = 2.213 + 0.0289 * distance \quad (1)$$

$$Forest : power = -0.1674 + 0.0514 * distance \quad (2)$$

These formulae predict the highest radio power level resulting in a PRR belonging to the low-quality region for a given distance. The R^2 values were 0.85 and 0.892 for the field and forest, respectively.

Similarly, we used the *b* values and linear regression analysis to determine the appropriate formula for the high-quality region. This process resulted in the following formulae:

$$Field : power = 3.307 + 0.0341 * distance \quad (3)$$

$$Forest : power = 1.4278 + 0.06155 * distance \quad (4)$$

These formulae predict the lowest radio power level resulting in a PRR belonging to the high-quality region for a given distance. The R^2 values were 0.848 and 0.762 for the field and forest, respectively.

In these models, radio power level is measured in discrete units, from 1 – 31, and distance is measured in feet. From the R^2 values, the models achieve a good fit to the actual




Fig. 11: Link Quality Model Compared to Field Data

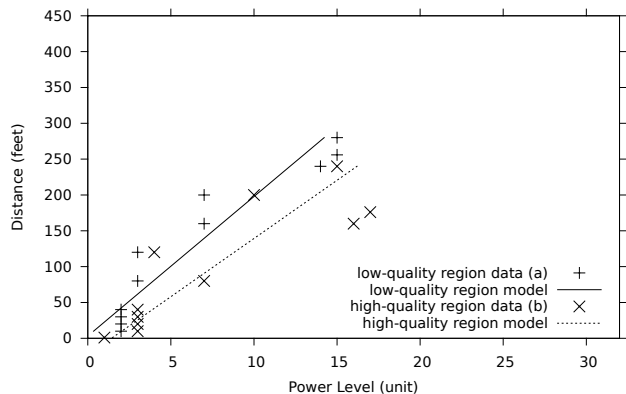


Fig. 12: Link Quality Model Compared to Forest Data

data. Figure 11 shows the plots for a and b values, along with the corresponding models for the field experiments. Similarly, Figure 12 shows the plots for a and b along with the corresponding models for the forest experiments.

Limitations. Three limitations of the current models should be noted. First, the results are specific to a single hardware platform. We do not expect the models to provide high accuracy for network deployments that employ different hardware platforms, such as those operating outside the Zigbee band (*i.e.*, 2.4Ghz). Second, our models are limited to interference-free environments. Given the effect of interference on network link quality [2, 8], these models are not likely to provide high accuracy in the presence of significant interference. Third, we note that the models assume only a single transmitting process. In scheduled transmission networks (e.g., using TDMA or FDMA) and networks with few concurrent transmitters, the models are directly applicable. However, in the presence of many concurrent transmitters, the accuracy of the models is expected to degrade.

VI. CONCLUSION AND FUTURE WORK

The behavior of embedded network systems depends largely on radio link quality. Without a priori knowledge of link quality, reliable system behavior is difficult to achieve. As a result, designers often develop applications characterized by unpredictable wireless behavior.

We have developed environment-specific empirical models for predicting radio link quality as a function of inter-node

distance and radio transmission power. The models help developers understand the behavior of radio links in open grass fields and dense forests, which correspond to a large number of network deployments. Hence, designers should be better able to develop applications that yield predictable performance. For example, designers will be able to use the models to determine which node layout and radio transmission power yield high performance for a given deployment site. Additionally, by adjusting radio transmission power, designers will be able to prolong application lifetimes by saving energy without sacrificing performance.

We have two elements of future work to extend our models. First, we plan to measure the degradation in the accuracy of the models when used in the presence of interference. Second, we plan to measure the accuracy of the models in the presence of common occurrences of noise spikes. In other words, we plan to measure the sensitivity of the models to the frequency of noise spikes as opposed to noise threshold – our current approach.

Acknowledgments. This work is funded by the National Science Foundation (awards CNS-1126344 and CNS-0745846). The authors wish to thank Alex Propst and Kevin Poole for their help in collecting the empirical data, as well as Dr. Wayne Goddard and Dr. Pradip Srimani for their assistance with the link quality models.

REFERENCES

- [1] A. Cerpa, J. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical model of lossy links in wireless sensor networks. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks*, pages 81–88, Los Alamitos CA, USA, April 2005. IEEE Computer Society.
- [2] A. Cerpa, J. Wong, M. Potkonjak, and D. Estrin. Temporal properties of low power wireless links: Modeling and implications on multi-hop routing. In *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 414–425, New York NY, USA, May 2005. ACM Press.
- [3] R. Jafari, A. Encarnacao, A. Zahoor, F. Dabiri, H. Noshadi, and M. Sarrafzadeh. Wireless sensor networks for health monitoring. pages 479–481, Washington DC, USA, July 2005. IEEE Computer Society.
- [4] A. Kamthe, M. Carreira-Perpi, and A. Cerpa. M&M: multi-level Markov model for wireless link simulations. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, pages 57–70, New York NY, USA, November 2009. ACM.
- [5] J. Leskovec, P. Sarkar, and C. Guestrin. Modeling link qualities in a sensor network. *Informatica (Slovenia)*, 29(4):445–452, 2005.
- [6] T. Liu and A. Cerpa. Foresee (4C): Wireless link prediction using link features. In *Proceedings of the 10th International Conference on Information Processing in Sensor Networks*, pages 294–305, Washington DC, USA, April 2011. IEEE.
- [7] E. Miluzzo, N. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. Eisenman, X. Zheng, and A. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*, pages 337–350, New York NY, USA, November 2008. ACM.
- [8] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, pages 95–107, New York NY, USA, November 2004. ACM Press.
- [9] N. Reijers, G. Halkes, and K. Langendoen. Link layer measurements in sensor networks. In *Proceedings of the IEEE International Conference on Mobile Ad-hoc and Sensor Systems*, pages 224–234. IEEE, Oct 2004.
- [10] K. Seada, M. Zuniga, A. Helmy, and B. Krishnamachari. Energy-efficient forwarding strategies for geographic routing in lossy wireless sensor networks. In *Proceedings of the 2nd International Conference on*

- Embedded Networked Sensor Systems*, pages 108–121, New York NY, USA, November 2004. ACM Press.
- [11] Texas Instruments. CC2420 2.4 GHZ IEEE 802.15.4 zigbee-ready RF transceiver data sheet (rev. 1.3). <http://www-s.ti.com/sc/ds/cc2420.pdf>, March 2012. (*last access*).
 - [12] G. Tolle, J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong. A macroscope in the redwoods. In *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems*, pages 51–63. ACM, Nov 2005.
 - [13] S. Wahba and J. Hallstrom. An empirical analysis of communication links in embedded wireless networks. In *Proceeding of the 49th ACM Southeast Conference*, pages 185–190, New York NY, USA, 2011. ACM.
 - [14] S. Wahba, K. LaForce, J. Fisher, and J. Hallstrom. An empirical evaluation of embedded link quality. In *Proceedings of the 2007 International Conference on Sensor Technologies and Applications*, pages 430–435, Washington DC, USA, Oct. 2007. IEEE Computer Society.
 - [15] G. Werner-Allen, K. Lorincz, M. Welsh, O. Marcillo, J. Johnson, M. Ruiz, and J. Lees. Deploying a wireless sensor network on an active volcano. *IEEE Int. Comp.*, 10(2):18–25, 2006.
 - [16] Y. Xu and W. Lee. Exploring spatial correlation for link quality estimation in wireless sensor networks. In *Proceedings of the 4th Annual IEEE International Conference on Pervasive Computing and Communications*, pages 200–211, Washington DC, USA, March 2006. IEEE Computer Society.
 - [17] M. Zuniga and B. Krishnamachari. Analyzing the transitional region in low power wireless links. In *Proceedings of the 1st IEEE Conference on Sensor and Ad Hoc Communications and Networks*, pages 517–526. IEEE, Oct 2004.