

UCLA

Technical Reports

Title

SCALE: A tool for Simple Connectivity Assessment in Lossy Environments

Permalink

<https://escholarship.org/uc/item/2q49z78g>

Authors

Alberto Cerpa
Naim Busek
Deborah Estrin

Publication Date

2003

SCALE: A tool for Simple Connectivity Assessment in Lossy Environments

Alberto Cerpa, Naim Busek and Deborah Estrin

CENS Technical Report # 21

Center for Embedded Networked Sensing, University of California, Los Angeles (UCLA)

Los Angeles, CA 90095, USA

September 5, 2003

Abstract— Wireless sensor networks will allow fine-grained monitoring in a wide range of environment (indoor and outdoor). Many of these environments, present very harsh conditions for wireless communication using low-power radios, including multipath/fading effects, reflections from obstacles, and attenuation from foliage. In this paper, we introduce *SCALE*, a network wireless measurement tool that uses packet delivery as the basic application-level metric. *SCALE* facilitates the gathering of packet delivery statistics using the *same* hardware platform and in the *same* environment targeted for deployment. Using up to 55 nodes, we were able to measure and study the connectivity conditions of two hardware platforms, Mica 1 and 2 motes, in three different environments: an outdoor habitat reserve, an urban outdoor environment in a university campus, and an office building, under systematically varied conditions. Among other things, we found that there is no clear correlation between packet delivery and distance in an area of more than 50% of the communication range, temporal variations of packet delivery are correlated with mean reception rate of each link, and the percentage of asymmetric links varies from 5% to 30%. Data collected using *SCALE* have interesting implications in the design, evaluation, and parameter tuning of sensor network protocols and algorithms.

I. INTRODUCTION

The advent of wireless sensor networks will allow detailed spatial and temporal environmental monitoring in a wide range of environments, from urban to wilderness; indoor and outdoor. Wireless radio communication is an essential component of these systems and enables sensor nodes to perform significant local coordination, distributed signal processing, and network self-configuration to achieve scalable, robust and long-lived networks [1], [10], [11]. The quality of the wireless channel depends on multiple factors, such as the environment, the radio frequency, the modulation scheme, and even the RF

transceiver hardware in use.

These networks will be deployed in harsh environments from the communication perspective, with significant multi-path effects. In addition, the low power radios typically used in sensor networks do not have sufficient frequency diversity to be resilient to multi-path communication. Under these conditions, wireless communication is known to be unpredictable and has been shown to vary drastically with small spatial changes and on different time scales. Even though most sensor network algorithms are designed to be adaptive to the variations in the communication channel [16], [4], there are several parameters that need to be adjusted to the operating conditions in order to improve performance. Furthermore, the real communication channels are very difficult to model for the wide range of target environments and the different type of radios, frequencies, and modulation schemes in use [6], [27], [12]. Thus, it is difficult to extensively test the algorithms under development in simulations under realistic conditions. Given the variability of the communication channel, and the difficulty to model it accurately, it is essential to get *quantitative* data that may allow us to better understand the channel characteristics in the target deployment area.

In this paper, we present *SCALE*, a measurement tool to study wireless communication channels with low power radios in new environments. It facilitates the characterization of the most basic communication metric from the application point of view: packet delivery. The tool enables the collection of packet delivery statistics using the *same* specific hardware platform and in the *same* environment intended for deployment. The data gathered by *SCALE* may allow protocol developers and engineers to better estimate the appropriate density, system parameter tuning constants, and expected performance of protocols and algorithms (data capacity, convergence time, latency). Table I shows some examples of how the connectivity statistics

TABLE I

EXAMPLES OF THE USEFULNESS OF CONNECTIVITY STATISTICS IN PROTOCOL DEVELOPMENT AND PARAMETER TUNING

Design Parameter	Data Collected	Utility
Physical density	Delivery rate vs. distance	Expected mean topological density
	Delivery rate vs. environment type and distance	Expected standard deviation in topological density
Algorithmic selection		Expected performance of in-network processing, e.g. opportunistic (geographical) data aggregation
Protocol selection		Expected performance of spatial correlation, e.g. geographical and topological routing
Protocol parameters (time constants)	Link asymmetry vs. distance	Expected performance of routing mechanisms that assume bidirectional links
	Delivery rate vs. time	Find reasonable routing and application soft state refresh time; find neighbor discovery probe period as a function of the stddev.
Packet size selection	Link asymmetry vs. delivery rate	Find neighbor discovery period as a function of mean and stddev.
	Delivery rate vs. packet size	Find optimal packet size to maximize <i>efficiency</i> ^a

^a Metric defined in section V-D.

collected from a specific target environment can be useful in this regard.

SCALE is fully configurable. Several parameters are configurable, such as, the packet probe size, the inter-packet period time, the transmission power gain, among others. This flexibility permits performing experiments under multiple different varied conditions. More importantly, it allows to repeat the measurements while constraining all parameters other than the one being varied, allowing us to systematically probe the effects of that particular parameter. The tool can be run transparently in a centralized way with all the software running in a central PC and connected to the nodes via serial cables, or in a fully distributed way with the software running in different distributed nodes. *SCALE* also provides a visualization screen to help viewing the connectivity data in real-time and after each experiment completes. Using up to 55 nodes, we were able to measure and study the connectivity conditions of two hardware platforms, Mica 1 and 2 nodes [14], [7], in three different environments: an outdoor habitat reserve, an urban outdoor environment on a university campus, and an office building.

In our experiments, we distributed the nodes in an ad-hoc manner in each of the different environments. Once all the nodes were deployed, the system made each node a transmitter, going through all the nodes in a round-robin fashion, one node at a time. When a node was transmitting packet probes, the rest of nodes in the experiment were in receiver-only mode collecting packet delivery statistics from the sender. The results were centrally logged. In all our experiments, we studied the effect of the environment under different conditions in the absence of interfer-

ing transmissions.

The results of our measurements using *SCALE* revealed some interesting findings. By analyzing data from a rich set of links with different distances, directions, antennae elevations from the ground, with or without line of sight—conditions that we expect to find in sensor network deployments [3]—, we found that there is no clear correlation between packet delivery and distance in an area of more than 50% of the total communication range. In addition, we found that temporal variations of packet delivery are not correlated with distance from the transmitter or transmission power level, but to the mean reception rate of each particular link. We also found that the percentage of link asymmetries varies from 5% up to 30% in some cases, and there was no obvious correlation between link asymmetries and distance and/or transmission power levels. By using this tool, we provide significant quantitative evidence that supports the commonly held belief that link asymmetries are due to hardware calibration differences.

Before we proceed, we would like to highlight the primary contributions of our paper. These are:

- The development of a measurement and visualization tool based on an application level metric (e.g. packet delivery), which facilitates *qualitative* and *quantitative* characterizations of the wireless channel in a particular target environment and using the same hardware platform intended for the actual deployment.
- The report of an initial set of *qualitative* and *quantitative* results using *SCALE* that investigates previous measurements, supplies data to support previous hypotheses in the literature, and provides new data from experiments performed in three different type

of environments, and with two different type of radios, under systematically varied conditions.

The rest of the paper is structured as follows. In the next section we review the related work in the area. Section III provides a complete description of the measurement tool, including the hardware and software components. The methodology used for the data collection experiments is discussed in Section IV. In Section V, we present some initial experimental results using the measurement tool. Finally, we conclude in Section VI.

II. RELATED WORK

There is currently a dearth of wireless communication measurement data for low power devices. Most of the previous related work has focused on experimental measurements with ad-hoc programs specialized to fit a particular platform.

In Ganesan *et al.* [12] a testbed of 150 nodes (Mica 1 motes) was used to measure the effects of link, MAC, and application layers in data communication. The experiments were carried out in a single outdoor environment, with no obstacles in the vicinity and with all the nodes near the ground. This work provided some empirical data to prove that radio connectivity was not isotropic (exhibit directionality) and also provided some measurements of number of asymmetric links as a function of distance. This work also speculated that links asymmetries may be caused by small differences in the hardware (radios) and slight differences in the nodes energy levels. In our work we provide substantial evidence that the cause of link asymmetries is in fact due to differences in hardware calibration and provide a more in depth analysis of the different factors affecting wireless communications in more than one environment and with more than one radio.

Woo *et al.* [29] examined packet loss between pair of motes and constructed packet loss models used to evaluate link quality estimators. Building on that work, in a more recent study by Woo *et al.* [30] and using up to 100 nodes in an open tennis court, they constructed packet loss models based on the mean and standard deviation reception rate values. Using these models in simulation and with a network of 50 nodes in a building lobby, they provide an illuminating evaluation of link quality estimators, neighborhood management policy, and routing strategies under varied conditions. Our study is complementary to this work; while we do not conduct any algorithmic evaluations, we do study the characteristics of packet delivery in the absence of concurrent transmissions, and using more than one radio in multiple environments.

A recent study by Zhao *et al.* [31] using up to 60 nodes (Mica 1 motes) showed some of the effects of link and

MAC layers in wireless communication. Using a simple linear topology, with a single sender, the work studied the packet delivery performance in three different environments, power levels, and coding schemes. This study provided experimental data showing heavy variability of packet reception in almost one third of the communication range for some scenarios. Our work is complementary to this. In our study, we gathered connectivity data using more than one sender and non-linear topologies, and our results show even further variability of packet delivery in more than half of the communication range. Our work does not consider the impact of multiple coding schemes, but study how the packet delivery is affected by packet size and using different hardware platforms.

Near ground effects in the 800-1000 MHz band is studied by Sohrabi *et al.* [28]. This work uses a particular model for power loss, and finds the constants in the model for different type of environments. The study provided experimental validation of the power drop off with higher exponents at smaller distances than the same channels with higher antennae. Our work considers near ground effects as one of the multiple effects affecting radio propagation. Our measurements also include data gathered from the 400 MHz band and use an application level metric, mean packet loss, instead of path loss.

There has been several studies for the characterization of cellular networks [19]. In our study, we use different (low-power) radios, and different coding schemes (less complex due to resource constraints); thus, we cannot rely completely on previous results from cellular networks.

Our previous work with ASCENT [4] motivated us to build this measurement tool to help us gain a quantitative understanding of some of the radio channel features. In ASCENT, we showed that due to the spatial and temporal variability of the wireless channel, the use of adaptive algorithms that constantly adapt to the local connectivity conditions was a *sine qua non* prerequisite to build any real sensor network system. Nevertheless, when faced with the challenge of defining some of the algorithm constants (e.g. heartbeat period), we were forced to use ad-hoc values and intuition for the parameter tuning. We believe *SCALE* fills this gap. Our work has also been inspired by the large number of measurement tools [20], [23] developed to understand protocol performance issues in the Internet. These tools have had a significant role in the development of Internet protocols like TCP [24], multicast routing protocols [8], [2], and many more. The data collected by these tools allowed Internet researchers to detect flaws in the design, adjust the parameter tuning, and improve the general performance of these protocols. Similarly, we hope *SCALE* could become a useful tool for



Fig. 1. *SCALE* hardware. The portable array is composed of a laptop PC attached to a serial multiplexer. Several UTP cables run from the multiplexer to the deployment locations where a mote is attached at the end.

TABLE II
NODES CHARACTERISTICS

	Mica 1	Mica 2
CPU Processor	Amtel 128	Amtel 128
Prog. Memory (KB)	128	128
Data Memory (KB)	4	4
Serial RS232	needs adapter	needs adapter
Clock Speed (MHZ)	4	7.38
RF Manufacturer	RFM [21]	Chipcon [5]
RF Transceiver	TR1000	CC1000
Radio frequency (MHz)	916	433
Modulation	ASK	FSK
Throughput (kbps)	13.3	19.2
TX power [0dBm] (mW)	< 1	< 1
Hardware Encoding	none	Manchester
Antenna	Omni whip	Omni whip

researchers working with sensor networks in often harsh and lossy environments for wireless communication.

III. SYSTEM DESCRIPTION

A. Overview

The system is built using the EmStar programming model [9]. It consists of a number of sensor nodes (motes) attached using long serial cables to one or more serial multiplexers that are connected to a standard laptop PC. This PC centrally runs the different processes that perform the data collection as if they were run by individual nodes. A visualization tool is integrated to help visualize in real time the progress of the experiment and to analyze and display the final results.

B. Hardware and Firmware

In our experiments we use two versions of nodes based on Mica motes (Mica 1 and 2)[14], [7]. Test Figure 1(a) shows the Mica 1. Figures 1(a)(c) show the mote platforms. Table II shows the main features of the hardware platforms used.

The ceiling and portable arrays [9] used in the experiments are composed of one or more serial port multiplexers attached to a laptop PC. Figures 1(b) show an image of the portable array with one serial multiplexer. The only difference between the arrays is that the ceiling array is permanently deployed in the ceiling of our lab, and the portable array is a completely mobile system that can be deployed anywhere. We use UTP Cat 5 cables of different lengths (up to 30 meters) and attach on end of the cable to the multiplexer and the other end to a node. The nodes are wall powered in the ceiling array and battery powered in the portable array. The portable/ceiling array is used as a logging/control channel through which we interface to the software.

The Mica motes firmware comes with an event-driven operating system called TinyOS [15]. When using Mica 1, it provides a DC-balanced single-error correction and double bit error detection (SECDED) scheme to encode each byte transmitted by the RF transceiver (RFM). When using Mica 2, it relies on the hardware encoding. The system supports variable packet sizes, and uses a 16-bit CRC that is computed over the entire packet for error detections (for both Mica 1 and 2). A simple driver (*Transceiver*) was used to run on the motes in TinyOS. Its function is to send/receive packets to/from the radio and pass them from/to the PC using the host-mote protocol over the serial connection.

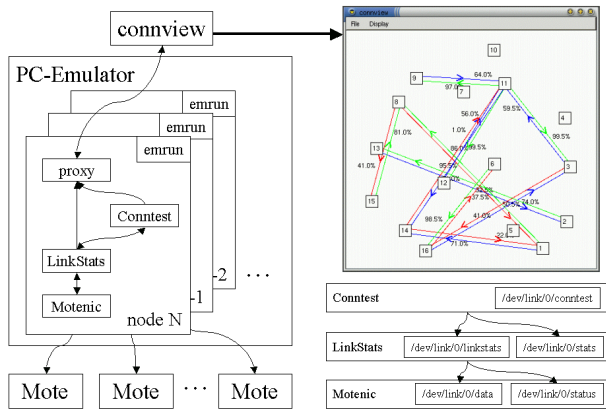


Fig. 2. *SCALE* software architecture. Multiple independent modules that export devices for IPC run in their own address spaces, all controlled by *emrun*. A user can interact with each module by simply using *cat* or *echo* Unix commands, or let the system proxy all the information to a central place. *Connview*, the visualization tool, allows checking the state of the experiments in real-time and performing post-processing analysis.

C. Software

SCALE has been designed to make full use of the EmStar programming model and software framework. Due to lack of space, we refer to [9] for further details on EmStar. Figure 2 shows a diagram of the software architecture. *SCALE* is completely modularized and all the modules have been written in C. Each node participating in the experiment runs a *software stack*, which consists of a series of modules interconnected in a certain way. Each module is represented by a process with its own address space. There are three modules for each node *software stack*: *Conntest*, in charge of sending and receiving probe packets, doing the control coordination among nodes (when to start/stop sending packet probes); *LinkStats*, responsible for maintaining the packet delivery statistics from all neighbors; and the low level channel driver, in charge of performing the communication with the radio. There are two channel drivers implemented: *MoteNic*, which implements the host-mote protocol to communicate to the radio over the serial port, and *Udpd*, which uses the UDP network interface as a communication driver. The collection of processes is managed by *emrun*, which starts each of the above modules in the correct dependency order based on the configuration file we provide (e.g. *Conntest* depends on *LinkStats*, and should only start once *LinkStats* is active). If a module terminates unexpectedly, *emrun* automatically restarts it and the other modules can reconnect to it without losing state. When using the system with the ceiling and portable arrays, all the processes are run in emulation mode in a central PC. Multiple copies of *emrun* are started—one for each node in the system—,

each of which forks a copy of the *software stack*. *SCALE* also provides a visualization tool, *Connview*, and its purpose is two-fold. First, it allows checking the status of the experiment in real time. Second, it permits the analysis and display of the final experimental results. Among some of its features, it includes the on/off display of any node or link, the coloring of links based on different percentages of packet delivery, display of asymmetric links, screen capture and file saving in graphical formats (jpeg and png), and many more.

We note that the *SCALE* could be used in a completely distributed fashion. For example, nodes could be connected to handheld-type battery power devices, like Compaq iPAQs [22] or Intel XScale [17], each of them being able to run a copy of the *software stack*. The coordination and data transfer for visualization could be done by an out-of-band channel, like an 802.11 network (in order to avoid interference with the radio channel we are measuring). One of the advantages of using the EmStar environment is that no software changes are required to run in a centralized or fully distributed way; the transition between the two modes is completely transparent. The advantage of the fully distributed mode is the elimination of the serial cables and the multiplexor to connect to the central PC. The main disadvantage is the increased total cost of the system and the limited battery lifetime of the handheld-devices. In our study, we opted for the centralized solution.

The basic data collection experiments work as follows. Each node transmits a certain number of packet probes in a round robin fashion (one transmitter at a time). Each probe packet contains the sender's node id and a sequence number. The rest of the nodes record the packets received from each neighbor and keep updated connectivity statistics, using the sequence numbers to detect packet losses. There are multiple variables that can be configured for each experiment. The number of round robin passes, the total number of packet probes to be sent (and the number of probes in each round), the packet probe size, the inter-packet period time, and the transmission output power are all fully configurable. If a user wants to evaluate the performance of an algorithm (e.g. routing algorithm) under different traffic workload and allowing multiple transmitters at a time, it simply deactivates the *Conntest* module in the configuration file. The measured packet delivery results will be the aggregate effect of the environment and the traffic workload in use (which may include collisions depending of the MAC layer used).

SCALE is also script-ready, and it is easy to configure an entire set of experiments varying one or more parameters at a time, leaving the system running with no human



(a) Outdoor Habitat, Will Rogers State Park

(b) Outdoor Urban, UCLA Boelter Hall Court Yard

(c) Indoor Office, UCLA CENS lab ceiling array

Fig. 3. Different environments used in our experiments using *SCALE*.

intervention. At the end of each experiment, all the data is automatically stored in log files with date and time of the experiment, the location, and the values of all the parameters used.

IV. METHODOLOGY

In this section we discuss the methodology used for our experiments.

The most important aspect of wireless communication for us is packet delivery performance, which is a metric that directly affects the performance perceived by the application. More precisely, our primary measure of performance is *packet loss* (the percentage of packets transmitted but not received), and its complement, *reception rate*.

The topology used for our experiments consisted of 16 nodes (portable array) distributed in an ad-hoc manner in different environments. We also used up to 55 nodes for our indoor experiments distributed in the ceiling of our lab (ceiling array). When using the portable array, nodes were placed in a variety of different positions, such as near the ground or elevated from the ground, with or without line of sight (LOS) between them, and with different levels of obstructions (furniture, walls, trees, etc.). The placement of the nodes also took into account the distance between them, in order to create a rich set of links at distances varying from 2 to 50 meters and in multiple different directions from any particular sender. In most of our experiments, each node sends up to 200 packets per round, transmitting 2 packets per second (unless otherwise noted). We verified that the transmission rate was low enough to guarantee no packet losses as a result of system issues (e.g. internal queue overflow).

Using this setup, we varied four factors in our experiments: the choice of environments, the radio type (and

frequency), the output transmit power settings, and the packet size settings.

The first factor we varied was the environment type. We selected three environments for our experimentation:

- *Indoor Office*. We chose our lab to perform some indoor connectivity experiments. It consists of a typical office type environment with an area of approximately 20m by 20m. It has partition panels, desks, chairs, cabinets, computers, monitors, etc. This environment is harsh for wireless communication due to multi-path reflections from walls and the possibility of interference from electronic devices. The choice of this environment is motivated by sensing applications in indoor environments [25].
- *Outdoor Urban*. We picked the UCLA Engineering courtyard as another environment for our experiments. It is an area of 70m by 35m surrounded by buildings and with some vegetation, trees, and an open area around the center. The vegetation and the walls from the buildings are expected to produce some signal attenuation and multi-path reflections as well. This environment is an intermediate measuring point between indoor places and outdoor natural habitats.
- *Outdoor Habitat*. We use a 200m by 150m section of the Will Rogers State Park, Pacific Palisades, California. The area consists of a small valley, surrounded by a 35 degree slope hill with very dense vegetation, including different type of plants, bushes and trees. Multi-path effects and signal attenuation due to the dense vegetation contribute to a harsh environment for wireless communication. There has been several efforts to monitor habitats in sensor networks [3], which motivate this environment.

The second factor we varied was the radio type. We

used two different type of radios with different transmission frequency and different modulation schemes. The Mica 1 transmits in the 916MHz band, and uses an amplitude shift keying (ASK) modulation scheme. The Mica 2 transmits in the 433MHz band, and uses a frequency shift keying (FSK) modulation scheme. The FSK modulation is more resilient to voltage supply variations since each symbol detection includes multiple zero-crossings. This is one of the reasons why the Mica 1 board needs an additional voltage regulator in place in order for the radio to be effective.

The third factor we varied was the output transmission power. The motes hardware allows discrete control of the output transmission power of the RF transceiver. This capability permits sensor network applications to control the power gain of the transceiver, allowing them to trade-off energy usage versus transmission range. The Mica 1 motes have a potentiometer circuit that allows controlling the amount of current delivered to the RFM radio [21]. The dynamic range of the output power selection with Mica 1 ranges from -10dBm to 0dBm. The Mica 2 Chipcon radio chip (CC1000) [5] has programmable output power from -20dBm to 10dBm controlled directly with the microcontroller. In our experiments we explored the -15dBm to +5dBm range of transmit power for the Mica 2 platform. Due to the differences in the dynamic ranges between the two platforms, we decided to qualify the power levels with respect to the dynamic range of *each* platform. For example, when using Mica 2 the -1dBm power level is considered *medium* power level (with respect to its own dynamic range), but when using Mica 1 the -1dBm power level is considered *high* power. In all our graphs we included the power level used in dBm units in order to facilitate the comparison. For the Mica 1 and outdoor experiments, we only explored the high-power settings (near 0 dBm) that were the only power values delivering enough signal strength to get meaningful connectivity results.

Finally, we varied the packet probe sizes in our experiments using two qualitatively different power settings (high and low power). The set of different packet sizes used was 25, 50, 100, 150 and 200 bytes. The payload of the packets was filled with random data up to the maximum size in use.

Nodes were localized manually. For each experiment, we built a local coordinate system and find the local coordinates of all the nodes in three dimensions. For the indoors experiments we localized the nodes using a measuring tape. The measuring error of the instrument is ± 0.1 cm. For the outdoors experiments we use a sonic ranger device (Zircon DM S50)[32]. The measuring error of the instrument is ± 1 cm. A conservative estimate of the lo-

calization error would be one order of magnitude larger than the instrument measuring error, so we estimate the localization error of each node to be ± 1 cm for indoors and ± 10 cm for outdoors. We note that the manual localization of the nodes is the *only* part of the entire procedure that requires human intervention.

Summary: We collected packet delivery data from more than 300,000 packet probes in experiments performed in 3 different environments, with 2 different type of radios, with 6 different power settings, and 5 different packet sizes. We used up to 16 nodes in our outdoor experiments and up to 55 nodes in our indoor experiments distributed in an ad-hoc manner, each node transmitting 200 packets. In each experiment, we measured the packet delivery performance of 240 links for the outdoor experiments and 2970 links for the indoors experiments.

V. EXPERIMENTAL RESULTS

In this section we present the results of using *SCALE* in different environments, and describe the different aspects of packet delivery performance. In all the results from our experiments shown in this section, we use confidence intervals with 95% degree of confidence based on large sample size ($n > 30$).

After some initial experimentation we have characterized the primary features discussed in the literature [12], [31] of our radio channels:

- **Asymmetrical links:** the connectivity of node *A* to node *B* ($A \rightarrow B$) might be significantly different than from node *B* to node *A* ($B \rightarrow A$).
- **Non-isotropic connectivity:** the connectivity is not necessarily the same in all the directions (same distance) from the source.
- **Non-monotonic distance decay:** nodes that are geographically far away from the source *may* get better connectivity than nodes that are geographically closer.

In the following sections we will take a closer look at the different aspects of packet delivery under systematically varied conditions using *SCALE*.

A. Spatial Characteristics

In this section we examine the qualitative and quantitative *spatial* characteristics of packet delivery in our experiments. We are interested in understanding how the reception rate varies with distance from the transmitter under different conditions and environments.

Figure 4 plots the raw packet delivery data in three example scenarios as a function of distance. The goal of these graphs is to show qualitatively the drastic variation

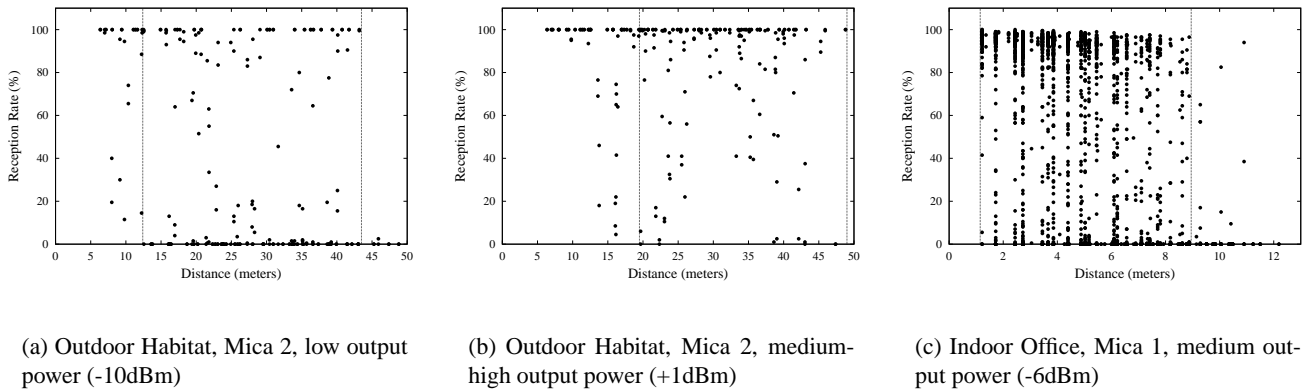


Fig. 4. Packet delivery percentage as a function of distance for different environments using different radios and power settings. In all cases, there is a region in which the reception rate varies dramatically, with delivery rates varying from near 100% to 0%. The width of the region where this phenomenon occurs is a significant portion (more than 50%, and up to 80% in some cases) of the communication range.




Fig. 5. Mean reception rate over distance for multiple environments, radios and transmission power levels. Each graph shows that the useful radio range tends to increase when the transmission output power increases. In addition, the graphs show that there is a great variability in some intermediate regions, as shown by the large values of the confidence intervals.

in reception rate for all the scenarios and platforms used in our experiments.

In Figure 4(a), we plot the raw connectivity data for the outdoor habitat experiment using Mica 2, and with low power settings. In this case, we observe that links with the same distance from the source can have reception rates that vary drastically from 100% to 0%, i.e., the area between the vertical lines. Figure 4(b) shows the same setup (environment and platform used), but using bigger transmission output power. When increasing the transmission power, we see the expected significant improvement in reception rate with respect to (a) for most of the links in our experiment. This can be seen by a larger density of data points near the 100% mark for almost all the distance range tested. We also see that links with reception rate lower than 50% appear at a larger minimum distance from the source (13 meters in the high power case b vs. 7 meters in the low power case a). Links with reception rates of 100% also appear at the limit of the maximum

range tested.¹

Figure 4(c) shows the raw connectivity data for our indoors experiments using our ceiling array. Note that the scale on the x axis (distance) is different from the previous graphs since the measurements are limited by the physical dimensions of our lab (the area is smaller than in the previous outdoors experiments). The bigger density of measuring points is due to the larger number of nodes available for our experiments (55 nodes). In this case, we also noticed great variation in reception rate for almost all the distance ranges tested in our experiments.

As expected, increasing the transmission output power produces an increase in the number of links with good reception rate at any given distance. However, the exist-

¹The 50 meters maximum range limit was due to the hardware availability, i.e. the total number of nodes available for our experiments to cover the entire distance range with minimum density and the number of serial multiplexors. There is no explicit maximum distance limit when using *SCALE*.

tence of bad links (links with small reception rate) is *not* completely eliminated when increasing the transmission output power and bad links tend to appear at almost *any* power setting used (although fewer when large power setting is used). We have verified this behavior even with maximum power settings using both Mica 1 and 2, and in the 2 outdoors environments we tested (Mica 1 and 2 at maximum power get high reception rates in our space-limited indoor lab). The graphs are omitted for brevity.

Next we analyze the mean behavior of the reception rate. In Figure 5, we plot the mean reception rate as a function of distance for different transmission power levels, environments, and radios. In these graphs, links were sorted based on distance from the source, and aggregated in 5 meter bins. Each measuring point represents the mean of all the links included in each 5 meter bin. There are more than 30 links in each bin.²

The large confidence intervals at some points show the high variability that could be visually observed in Figure 4. In all the cases shown in Figure 5 there is a general decrease in the reception rate as we increase the distance from the source. This is expected due to attenuation of the signal over distance for any transmission power level.

Discussion. The significant spatial variation in packet delivery using low power devices was first noted in previous work [4], which showed that nodes that are geographically further away from the source could, in practice, obtain better reception rate than nodes that are closer.

In [31], using nodes placed in a line, the area where the variability in packet reception was significant had a width of 20% to 30% of the communication range, and it was always located near the maximum radio range. In our experience, when using network topologies that extend in multiple directions from the source (not necessarily in a line) with different probability of obstruction depending on the node placement (as one would expect in real sensor network deployments [3]), we observed the width of the highly variable reception rate area to be in most cases larger than 50%, and up to 80% of radio range in some cases. In our experiments, this area starts well before the limit of the radio range. This result indicates that assumptions of packet delivery based exclusively on distance from the source can be erroneous in practice.

Multipath and fading effects can explain the level variability in packet delivery seen in our experiments. When the direct signal is strong and the reflected components are attenuated, the reception rates are high. When the di-

²The rightmost bin (largest distance) for the *outdoor* experiments has less than 30 links, so its confidence interval has less statistical significance (cannot assume a population normal distribution). The sample mean is still the best estimator of the population mean though.

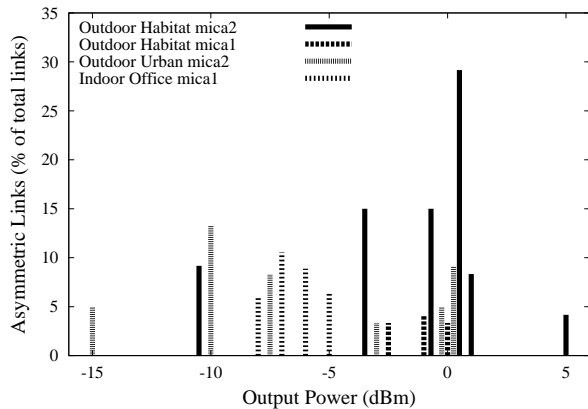


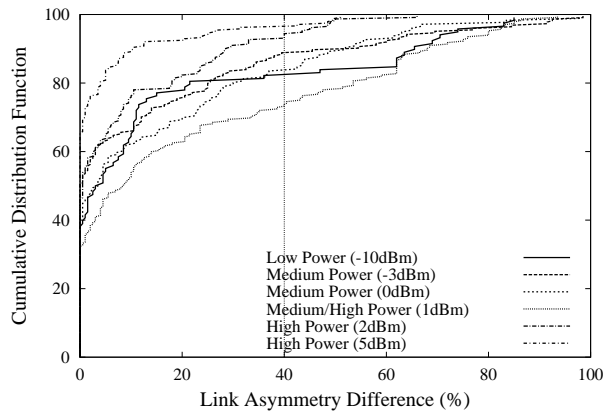
Fig. 7. Percentage of asymmetric links as a function of transmission output power for different environments, and radios. There is no clear correlation between the transmission output power and the total number of asymmetric links using a large range of environments, transmission output power and radios.

rect signal is attenuated, the reflected components might produce constructive or destructive interference of the final signal. Thus, small variations in the attenuation due to obstructions and node position can affect the reception rate. In our experiments, due to the harshness of the environments for low-power radio communication, nodes at the same distance from the source can have different levels of obstruction and attenuation (since the signal travels on different directions from the source toward the different receivers), experiencing significantly different packet delivery depending on the strength of the direct signal and the type of constructive or destructive interference.


We argue that the great variability in the reception rate over an extended area of the communication range is a common characteristic shared by a *family* of low-power radio devices commonly used in sensor network systems. This is sustained by the fact that we got the same *qualitative* results using two different radio platforms (widely accepted in the sensor research community). The lack of frequency diversity in these devices might be one of the reasons why these radios are more likely to suffer multipath effects (as opposed to more power-hungry spread spectrum radios).

B. Link Asymmetries


In the previous section, we discussed how packet delivery varies greatly over a large portion of the radio range. In this section, we focus on quantitative analysis of asymmetric links. Link asymmetries occur infrequently in 802.11 wireless networks, and are often filtered out by protocol levels [18], [26]. The study in [12] reported that asymmetric links were far more common when using low



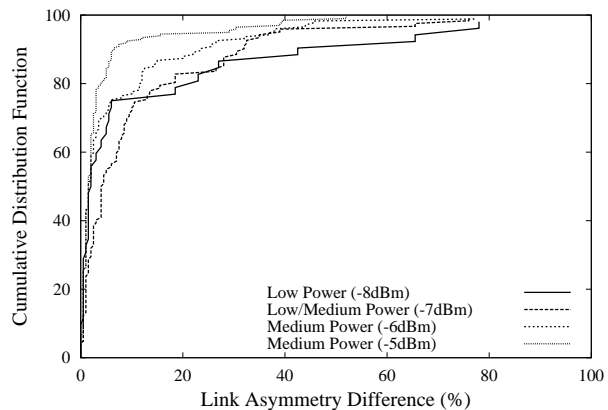
(a) Outdoor Habitat, Mica 2



(b) Outdoor Urban, Mica 2



(c) Outdoor Habitat, Mica 1



(d) Indoor Office, Mica 1

Fig. 6. Link asymmetry distribution for Mica 1 and 2 in three different environments. In all cases there is at least 5% of link pairs with a difference in reception rate larger than 40%, and in some cases the percentage of asymmetric links is as big as 30%


power radios, even when all the nodes were set to use the same transmission power level.

In this study, an asymmetric link is defined as one where the difference in the reception rate between the link in one direction and the other direction is larger than a certain threshold. We have chosen 40% as our threshold. We used two qualitatively different packet sizes (25 bytes and 200 bytes) in the experiments performed in these section, and we did not observe important variations based on packet size.


Figure 6 presents the cumulative probability distribution of link pair asymmetry for several environments and transmission power levels using both Mica 1 and 2. The vertical line on the 40% shows the threshold for asymmetric links used in this study. This graph shows how the percentage of asymmetric links would change if we had picked a different threshold value. More than 5% of the

link pairs have reception rate differences larger than 40%, and sometimes up to 30% of the link pairs have asymmetric properties. These asymmetric links are known for their impact on higher level protocols, such as routing [26].

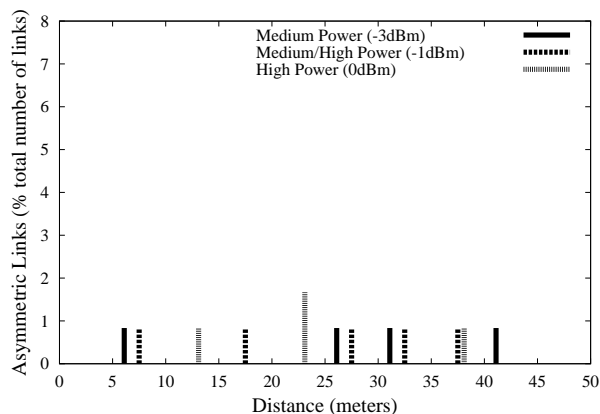
Figure 7 shows the total percentage of asymmetric links –with respect to the total number of links in each experiment– as a function of the transmission output power for three different environments using both Mica 1 and 2. Each bar represents an entire set of experiments performed at a particular transmission power level. Note that we did not systematically cover the entire dynamic range of transmission output power, but rather picked sample measuring points. In other words, the absence of a bar in certain power region is due to the absence of a measuring point, *not* the result of zero asymmetric links in that power level. For each radio platform, we covered almost the entire power range in different environments.




(a) Outdoor Habitat, Mica 2



(b) Outdoor Urban, Mica 2



(c) Outdoor Habitat, Mica 1



(d) Indoor Office, Mica 1

Fig. 8. Percentage of asymmetric links (with respect to the total number of links) as a function of distance for different environments, radios, and power levels. It is clear from the graphs that there is no obvious correlation between the asymmetric links and distance

Mica 1 was explored from -8 dBm to 0 dBm and Mica 2 was explored from -15 dBm to $+5$ dBm (in both cases near the entire dynamic range allowed by each RF transceiver hardware). Some of the bars have been offset in the x axis value (power) to improve readability, mainly around the cluttered 0 dBm region.

From the graph we can see that for each platform in each environment, there is no clear correlation between transmission power level and the percentage of asymmetric links. Furthermore, the percentage of asymmetric links seems to oscillate between 5% to 15% of the total number of links depending on the hardware platform and the environment, and in some cases being up to 30% of the total.

Figure 8 plots the percentage of asymmetric links as a function of distance for three different environments and two platforms for different transmission power levels.

Note that in this case we systematically explored the entire distance space for each environment, and the absence of a bar at a particular distance indicates the absence of asymmetric links at that distance from the source.

Figure 8(a) and (b) show the results of using Mica 2 in two different environments with different transmissions power levels. There is no clear correlation between the number of asymmetric links and the distance from the source. Asymmetric links tend to appear in a wide range of distances from the source, increasing and decreasing alternatively as we move further. In Figure 8(c) and (d) we show the results of using Mica 1 in two different environments with several different transmission power levels. Note that the scale for the x axis (distance) in Figure 8(d) is different from the outdoors experiments, since the indoors experiments were performed in a smaller area. When using Mica 1 we notice the same phenomena than

TABLE III
ASYMMETRIC LINK-PAIRS NODE SWAPPING RESULTS

Node Type	Location Type	Asymmetric link-pairs <i>before</i> swapping	Inverted link-pairs <i>after</i> swapping
Mica 2	Outdoor Urban	11	10
Mica 2	Indoor Office	10	9
Mica 1	Indoor Office	24	22

when using Mica 2; i.e. asymmetric links seem not to be correlated with distance from the source, and they appear in all the distance ranges tried in our experiments.³

Discussion. In [12] the spatial distribution of the asymmetric links was concentrated around the limit of the communication range for two different power settings tried. Our results show that there was no spatial correlation of asymmetric links; asymmetries were equally likely to happen well before the limit of the radio range. In that study they argued that at the limit of the communication range, small differences between nodes' transmit power and reception sensitivity may become significant and resulted in asymmetries. In other words, the link in one direction may have a direct signal that is strong enough (above a certain threshold) to get good reception rate while in the other direction the signal may be below the threshold and reflected signal components may affect the reception rate, causing link asymmetries. One interesting observation is that the experiments performed in [12] were done in a flat, open parking structure with no obstacles in the immediate vicinity. The difference between the environments where the experiments were conducted might explain the differences between results of the two studies. In [12], in the absence of obstacles, sufficient attenuation to produce link asymmetries was only existent in the limit of the radio range, while in our experiments with cluttered environments we experienced different level of attenuation at the same distance from the source, potentially producing the same effect at distances other than near the radio range.

One question that still remained unanswered was whether the cause of link asymmetries was *primarily* due to differences in hardware calibration. In both, Indoors Office and Outdoor Urban, we run experiments using different transmission power levels. Using the *SCALE* visualization tool (*Connview*), we quickly identified the pair of nodes that experienced asymmetric links. We emphasize that the online nature and ease of use of *SCALE* made this

³Mica 1s got systematically smaller percentages of asymmetric links than Mica 2s. We do not have an explanation for this behavior other than hardware differences between the two radios.

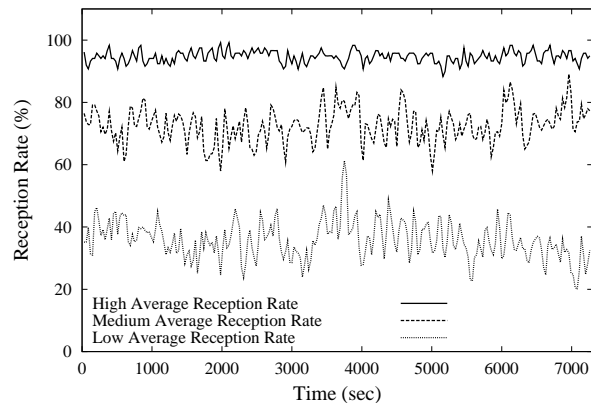



Fig. 9. Reception rate as a function of time for Mica 1 in the Indoor Office environment with medium power level (-5dBm). Links with higher mean reception rate tend to have less variability over time.

task very simple. If a node experienced link asymmetries with more than one node, we picked the pair with larger reception rate difference. Then we proceeded to carefully mark all the nodes physical placement (for the outdoor experiments we even took pictures of each node exact position/placement). We first verified the sensitivity of very small manual displacements by removing the nodes from the end of the serial cable and re-attaching them again in the same previously marked position. We re-ran the experiments and verified that the each pair of nodes had the same level of link asymmetry as before. In all cases the level of asymmetry in each pair remained the same. This result gave us confidence that minor manual displacements that happen when removing and re-attaching nodes in the same positions would not affect our final results.


Once this was verified, we proceeded to swap positions for each pair of nodes, being very careful to place the opposite node of each pair into exactly the same position of the original node. Table III shows the summary of our results. We tested 45 asymmetric link-pairs in both environments using both Mica 1 and 2. In most cases, when swapping the nodes' positions, the link asymmetries got inverted. This phenomenon happened 91.1% of the time with a confidence interval of $\pm 8.32\%$ and a degree of confidence of 95%. This result suggests that there is a strong indication that link asymmetries are primarily caused by small differences in hardware calibration and energy levels between nodes. We believe this is the first study that presents *quantitative* data supporting this hypothesis.

C. Temporal Characteristics

In this section, we examine how packet delivery varies with time, and what are the spatial characteristics of this



(a) Indoor Office, Mica 1, reception rate standard deviation as a function of distance



(b) Indoor Office, Mica 1, reception rate standard deviation as a function of the mean reception rate

Fig. 10. Figure 10(a) shows that there is no clear correlation between the variability of the reception rate (σ , standard deviation) and the distance from the transmitter. Figure 10(b) shows an interesting correlation. Links with very high reception rate over time ($> 90\%$) tend to be more stable (small σ), followed by links with very low reception rates (near 0%). The links with intermediate reception rate tend to be highly unstable, with very large variability over time (up to values of 50% for σ).

variation.

For this experiment we configured *SCALE* to run with just one sender (no round-robin) at a data rate of 2 packets/sec, with data packet size of 200 bytes. We configured *SCALE* to try multiple power levels, and let it run for more than 2 hours for each power level selected in the Indoor Office environment. The mean reception rate was computed every 30 seconds, and the window size for the reception rate calculation was set to 60 seconds (each packet sent affects two mean reception rate calculations). We present the results only using Mica 1, since Mica 2 traces present the same *qualitative* characteristics (not shown for brevity).

Figure 9 shows the mean reception rate variability over time for three different links with different mean reception rate over the entire time of the experiment. The figure illustrates that the variability for the link with a high mean reception rate (95% , the top curve) is quite small, and varies between 92% to 98% . On the other hand, the link with low mean reception rate ($\sim 40\%$, the bottom curve) has high variability reception rate, and varies between 20% to 60% over the entire time of the experiment.

Figure 10(a) shows the relation between the standard deviation of the reception rate and the distance from the transmitter. Each point in the graph represents the sample variance, which is the best estimator of the population variance. The errorbars show the confidence interval of the standard deviation estimation with a degree of confidence of 95% . The confidence intervals for each point

were obtained using the Chi-Square distribution⁴. The graph shows that, using two different power levels, there is no clear correlation between the variability of the reception rate in time (standard deviation σ) and the distance from the transmitter. High values of standard deviation appear in a wide range of distances from the source.

In Figure 10(b) we plot the relationship between the standard deviation and the mean reception rate. On the right side of the graph, we can see that links with high mean reception rate ($> 90\%$) show very little variation over time and tend to remain stable with good connectivity. Similarly, links with very low mean reception rate (near 0%) are also stable over time and tend to remain bad links over the time period tested. On the contrary, links with mean reception rates that range from 20% to 80% show great variability over time, and in some cases present standard deviation values in the order of 50% ! It is not uncommon for some of these links to go from 100% reception rate to 0% in the course of a two hour window.

The results presented in Section V-A showed that links with poor/medium reception rate were present in a wide range of distances from the transmitter. In addition, the correlation between high variability over time and poor/medium reception rates shown in Figure 10(b) can help explain the results we obtained in Figure 10(a); links with poor/medium reception rate appear across a wide

⁴The points with very large standard deviation have confidence intervals with less statistical robustness. In some of these points, we do not *always* have a normally distributed population necessary by the Chi-Square method to make robust statistical inferences.




Fig. 11. Mean reception rate over distance for Mica 1 and 2 using multiple transmission power levels, and packet sizes in the Outdoor Urban environment. There is no significant difference in packet delivery between large and small packet sizes, with only a small decrease in performance for larger packet sizes.

range of distances from the transmitter, and these are precisely the links with higher standard deviation values.

D. Transmission Efficiency

In this section we take a look at the packet delivery efficiency. For any given payload size, there is an associated efficiency overhead given by the size of the radio packet header and the packet preamble. In addition, one would expect that as we increase the packet size, the probability of successfully receiving an error-free packet would decrease since there is a larger probability of any part of the packet being corrupted (for any given channel BER).

In our experiments, we used different coding schemes depending on the RF transceiver used (software-based SECDED for Mica 1, and hardware-based Manchester for Mica 2). In this study, we did not consider the efficiency differences between the different coding schemes. The packet overhead is related to the packet header added to each transmitted packet (it contains addressing and CRC error-checking information), and the preamble/start-symbol overhead that is used by each radio to detect a new packet over the air and lock into the incoming signal to achieve bit-level synchronization. The packet header overhead is the same in both Mica 1 and 2 and consists of 7 bytes/packet (addressing + CRC). The preamble/start symbol overhead is different for each radio. The RFM transceiver (Mica 1) uses 12 bytes/packet, while the CC1000 transceiver (Mica 2) uses 20 bytes/packet.

In Figure 11 we show the reception rate as a function of distance for both Mica 1 and 2 at different power levels in the Outdoor Urban environment. We show only small (25 bytes) and big (150 bytes) payload sizes. We gathered data for several other payload sizes (25, 50, 100, 150 and 200 bytes), but we did not include them in the graphs to

improve readability, since all the curves were very similar. In general, for all three graphs and using different transmission power levels for Mica 2, we notice only a *slight* decrease in the reception rate as we increase the payload size.

Based on the results shown in Figure 11, we investigated whether the small decrease in reception rate as we increase payload size could be compensated by the decrease in the packet overhead. We defined a new metric to measure this relationship. For any given link, we want to measure the relation between the useful information received versus the information sent over a link using a metric called *efficiency*⁵, which is defined as:

$$\frac{UsefulBitRX}{TotalBitTX} = \frac{PayloadSize_i}{TotalPacketSize_i} \times P_i(s) \quad (1)$$

where $TotalPacketSize_i$ is the total size of the packet type i (preamble + header + payload); $PayloadSize_i$ is the payload size of packet type i ; and $P_i(s)$ is the probability of successfully receiving a packet of type i . The fraction $PayloadSize_i/TotalPacketSize_i$ determines the ideal efficiency under optimal reception rates. The different packet types i are given by the different payload sizes we used in our experiments.

Figure 12 plots the *efficiency* metric defined in (1) as a function of distance from the source. Each horizontal line determines the theoretical efficiency value under ideal reception rate (100%) for the different payload sizes. We can see from the graphs that larger efficiency is achieved by using larger payload sizes. For Mica 2, the maximum efficiency is achieved with the larger payload size tried in most of the radio range. For Mica 1, it seems that the

⁵A related metric was defined in [31] in the context of MAC layer retransmissions




Fig. 12. Efficiency rate as a function of distance for Mica 1 and 2 in the Outdoor Urban environment for multiple payload sizes. The vertical lines show the ideal efficiency for each packet payload size. Larger efficiency is achieved by using larger payload sizes than the default values set up in the mote software.

optimal packet size to maximize efficiency is around 150 bytes. Using values larger than that do not improve performance.

The default value established in the mote software [15] for the payload size is 29 bytes, and could probably be increased with a noticeable improvement in terms of efficiency. Nevertheless, larger packets may produce fairness problems for channel utilization when there are multiple senders in a region, and even increase the probability of collisions depending of the MAC scheme used in the system. The final choice of packet size should also consider the particular MAC layer scheme used and the expected traffic pattern and workload introduced by the application.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented *SCALE*, a network wireless measuring and visualization tool that enables the *qualitatively* and *quantitative* characterization of the wireless channel in a particular target environment and using the same hardware platform intended for deployment. Using *SCALE* with two different radio transceivers in three different environments, we found that there is no clear correlation between packet delivery and distance in an area of more than 50% of the communication range, temporal variations of packet delivery are not correlated with distance from the source but with the mean reception rate of each link, the percentage of asymmetric links varies from 5% to 30% and strong indications suggesting that the cause of links asymmetries is primarily due to hardware calibration differences. The data collected using *SCALE* provide some useful insights for protocol developers and engineers working in sensor networks.

In the near future, we plan to integrate *SCALE* with some self-localization systems under development [13] to

eliminate the extent of human intervention needed and make the system more autonomous. In addition, we plan to extend the low-level radio interface to collect some of the signal-to-noise information available in some of the RF transceivers used in sensor networks.

Acknowledgments

SCALE was made possible through the support of the USC/NSF SCOWR project (USC subcontract #059565), the GALORE project (USAF/Wright-Patterson Air Force Base contract #F33615-01-C-1906) and the National Science Foundation Cooperative Agreement #CCR-0120778 supporting the Center of Embedded Networked Sensing.

We would like to thank to all members of the UCLA LECS lab that made contributions to EmStar software developing environment. Many individuals made valuable comments that helped improve previous versions of this paper, including Pierpaolo Bergamo, Henri Dubois-Ferriere, Deepak Ganesan, Ramesh Govindan, John Heidemann, Joe Polastre, Mohammed Rahimi, Hanbiao Wang, Alec Woo, and Jerry Zhao.

REFERENCES

- [1] B.R. Badrinath, M. Srivastava, K. Mills, J. Scholtz, and K. Sollins. Special issue on smart spaces and environments. October 2000.
- [2] Tony Ballardie, Paul Francis, and Jon Crowcroft. Core based trees (cbt). In Deepinder P. Sidhu, editor, *SIGCOMM Symposium on Communications Architectures and Protocols*, pages 85–95, San Francisco, California, USA, September 1993.
- [3] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao. Habitat monitoring: Application driver for wireless communications technology. In *Proceedings of the SIGCOMM Workshop on Communications in Latin America and the Caribbean*, San Jose, Costa Rica, April 3–5 2001. ACM.

- [4] Alberto Cerpa and Deborah Estrin. ASCENT: Adaptive self-configuring sensor networks topologies. In *Proceedings of the Twenty First Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, NY, USA, June 23–27 2002. IEEE.
- [5] Chipcon. Cc1000 low power radio transceiver, http://www.chipcon.com/files/cc1000_data_sheet_2.1.pdf.
- [6] D. C. Cox, R. R. Murray, and A. W. Norris. Measurements of 800 mhz radio transmission into buildings with metallic walls. *The Bell System Technical Journal*, 9(62):2695–2717, November 1983.
- [7] Crossbow. Mica2 wireless measurement system datasheet, http://www.xbow.com/products/product_pdf_files/datasheets/wireless/6020-0042-03_a_mica2.pdf.
- [8] Stephen Deering, Deborah L. Estrin, Dino Farinacci, Van Jacobson, Ching-Gung Liu, and Liming Wei. The PIM architecture for wide-area multicast routing. *IEEE/ACM Transactions on Networking*, 4(2):153–162, 1996.
- [9] Jeremy Elson, Solomon Bien, Naim Busek, Vladimir Bychkovskiy, Alberto Cerpa, Deepak Ganesan, Lewis Girod, Ben Greenstein, Tom Schoellhammer, Thanos Stathopoulos, and Deborah Estrin. EmStar: An environment for developing wireless embedded systems software. Technical Report CENS Technical Report 0009, Center for Embedded Networked Sensing, University of California, Los Angeles, March 2003.
- [10] D. Estrin, R. Govindan, and J. Heidemann. Special issue on embedding the internet. 43(5), May 2000.
- [11] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom 1999)*. ACM, August 1999.
- [12] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report UCLA CSD-TR 02-0013, Center for Embedded Networked Sensing, University of California, Los Angeles and Intel Research Lab, University of California, Berkeley, February 2002.
- [13] Lewis Girod, Vladimir Bychkovskiy, Jeremy Elson, and Deborah Estrin. Locating tiny sensors in time and space: a case study. In *Proceedings of the International Conference on Computer Design (ICCD)*, Freiburg, Germany, September 2002. <http://lecs.cs.ucla.edu/Publications>.
- [14] Jason Hill and David Culler. Mica: A wireless platform for deeply embedded networks. 22(6):12–24, Nov/Dec 2002.
- [15] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for networked sensors. In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, pages 93–104, Cambridge, MA, USA, November 2000. ACM.
- [16] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, N.Y., August 2000. ACM.
- [17] Intel. Intel xscale technology, <http://www.intel.com/design/intelxscale>.
- [18] D. Kim, C. K. Toh, and Y. Choi. On supporting link asymmetry in mobile ad hoc networks. In *Proceedings of IEEE GLOBECOM*, pages 2798–2803, San Antonio, Texas, USA, November 25–29 2001. IEEE.
- [19] K. K. Leung, M.V. Clark, B. McNair, and Z. Kostic. Outdoor ieee 802.11 cellular networks: Radio link performance. In *Proceedings of IEEE ICC 2002*, New York, April 2002. <http://www.bell-labs.com/user/kin/papers/802.11.radio.pdf>.
- [20] S. McCanne. The bsd packet filter: A new architecture for user-level packet capture. In *Proc. of the Winter 1993 USENIX Conference*, pages 259–270, San Diego, California, USA, 1993.
- [21] RFM Monolithics. Tr1000 low power radio system, <http://www.rfm.com/products/data/tr1000.pdf>.
- [22] Hewlett Packard. ipaq h1910, <http://www.shopping.hp.com/shopping/pdf/311324.pdf>.
- [23] Vern Paxson. Automated packet trace analysis of TCP implementations. In *SIGCOMM*, pages 167–179, Cannes, France, September 14–18 1997.
- [24] J. Postel. RFC 793: Transmission Control Protocol, September 1981.
- [25] J. Rabaey, E. Arens, C. Federspiel, A. Gadgil, D. Messerschmitt, W. Nazaroff, K. Pister, and P. Varaiya S. Oren. Smart energy distribution and consumption: Information technology as an enabling force. Technical report, Center for Information Technology Research in the Interest of Society (CITRIS), 2001.
- [26] Venugopalan Ramasubramanian, Ranveer Chandra, and Daniel Mosse. Providing a bidirectional abstraction for unidirectional adhoc networks. In *Proceedings of the Twenty First Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, NY, USA, June 23–27 2002. IEEE.
- [27] S. Y. Seidel and T. S. Rappaport. 914 mhz path loss prediction models for indoor wireless communications in multifloored buildings. *IEEE Transactions on Antennas and Propagation*, 2(40):207–217, February 1992.
- [28] K. Sohrabi, B. Manriquez, and G. Pottie. Near ground wideband channel measurement in 800-1000 mhz. *IEEE 49th Vehicular Technology Conference*, 1:571–574, March 1999.
- [29] Alec Woo and David E. Culler. Evaluation of efficient link reliability estimators for low-power wireless networks. Technical Report UCB/CSD-03-1270, University of California, Berkeley, April 2003.
- [30] Alec Woo, Terence Tong, and David Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proceedings of ACM Sensys 2003 (to appear)*, Los Angeles, CA, USA, November 5–7 2003. ACM.
- [31] Yonggang Jerry Zhao and Ramesh Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of ACM Sensys 2003 (to appear)*, Los Angeles, CA, USA, November 5–7 2003. ACM.
- [32] Zircon. Dms50 sonic ranger datasheet, http://www.zircon.com/sellpages/measuring/dms50/dms50instructions/dms50_inst.pdf.