

Understanding Packet Delivery Performance In Dense Wireless Sensor Networks*

Jerry Zhao
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781
zhaoy@usc.edu

Ramesh Govindan
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781
ramesh@usc.edu

ABSTRACT

Wireless sensor networks promise fine-grain monitoring in a wide variety of environments. Many of these environments (*e.g.*, indoor environments or habitats) can be harsh for wireless communication. From a networking perspective, the most basic aspect of wireless communication is the packet delivery performance: the spatio-temporal characteristics of packet loss, and its environmental dependence. These factors will deeply impact the performance of data acquisition from these networks.

In this paper, we report on a systematic medium-scale (up to sixty nodes) measurement of packet delivery in three different environments: an indoor office building, a habitat with moderate foliage, and an open parking lot. Our findings have interesting implications for the design and evaluation of routing and medium-access protocols for sensor networks.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; C.4 [Performance of Systems]: Performance attributes, Measurement techniques

General Terms

Measurement, Experimentation

Keywords

Low power radio, Packet loss, Performance measurement

1. INTRODUCTION

Wireless communication has the reputation of being notoriously unpredictable. The quality of wireless communication depends on the environment, the part of the frequency

*This work is supported in part by NSF grant CCR-0121778 for the Center for Embedded Systems.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SenSys'03, November 5–7, 2003, Los Angeles, California, USA.
Copyright 2003 ACM 1-58113-707-9/03/0011 ...\$5.00.

spectrum under use, the particular modulation schemes under use, and possibly on the communicating devices themselves. Communication quality can vary dramatically over time, and has been reputed to change with slight spatial displacements. All of these are true to a greater degree for ad-hoc (or infrastructure-less) communication than for wireless communication to a base station. Given this, and the paucity of large-scale deployments, it is perhaps not surprising that there have been no medium to large-scale *measurements* of ad-hoc wireless systems; one expects measurement studies to reveal high variability in performance, and one suspects that such studies will be non-representative.

Wireless sensor networks [5, 7] are predicted on ad-hoc wireless communications. Perhaps more than other ad-hoc wireless systems, these networks can expect highly variable wireless communication. They will be deployed in harsh, inaccessible, environments which, almost by definition will exhibit significant multi-path communication. Many of the current sensor platforms use low-power radios which do not have enough frequency diversity to reject multi-path propagation. Finally, these networks will be fairly densely deployed (on the order of tens of nodes within communication range). Given the potential impact of these networks, and despite the anecdotal evidence of variability in wireless communication, we argue that it is imperative that we get a *quantitative* understanding of wireless communication in sensor networks, however imperfect.

Our paper is a first attempt at this. Using up to 60 Mica motes, we systematically evaluate the most basic aspect of wireless communication in a sensor network: packet delivery.

Particularly for energy-constrained networks, packet delivery performance is important, since that translates to network lifetime. Sensor networks are predicated using low-power RF transceivers in a multi-hop fashion. Multiple short hops can be more energy-efficient than one single hop over a long range link. Poor cumulative packet delivery performance across multiple hops may degrade performance of data transport and expend significant energy. Depending on the kind of application, it might significantly undermine application-level performance. Finally, understanding the dynamic range of packet delivery performance (and the extent, and time-varying nature of this performance) is important for *evaluating* almost all sensor network communication protocols.

We study packet delivery performance at two layers of the communication stack (Section 3). At the physical-layer and in the absence of interfering transmissions, packet de-

livery performance is largely a function of the environment, the particular physical layer coding scheme, and perhaps individual receiver characteristics. We place a simple linear topology, with a single sender, in three different environments: an office building, a local habitat, and an open parking lot. For these three environments, we study the efficacy of packet delivery under different transmit powers and physical layer codings.

At the medium access layer, interfering transmissions contribute to poor packet delivery performance. Many MAC layers contain mechanisms, such as carrier sense and link-layer retransmissions, to counteract these effects. We study the efficacy of such mechanisms in our three environments discussed above.

Our measurements (Sections 4 and 5) uncover a variety of interesting phenomena. There are heavy tails in the distributions of packet loss, both at the physical layer and at the MAC layer. In our indoor experiments at the physical layer, for example, fully half of the links experienced more than 10% packet loss, and a third more than 30%. At the physical layer, this variability can be characterized by the existence of a *gray area* within the communication range of a node: receivers in this gray area are likely to experience choppy packet reception, and in some environments, this gray area is almost a third of the communication range. The gray area is also distinguished by significant variability in packet reception over time. Relatively sophisticated physical layer coding schemes are able to mask some of the variability, but with a loss in bandwidth efficiency. At the MAC layer, link-layer retransmissions are unable to reduce the variability; packet losses at the MAC layer also exhibit heavy tails. Moreover, the efficiency of the MAC layer is low: 50% to 80% of communication energy is wasted in overcoming packet collisions and environmental effects. Finally, in our harsher environments, nearly 10% of the links exhibit asymmetric packet loss.

Taken together, this appears to paint a somewhat pessimistic picture of wireless communication for sensor networks. However, we contend that there might be a simple set of mechanisms that can greatly improve packet delivery in the environments that sensor networks are targeted for. Such *topology control*¹ mechanisms would carefully (*i.e.*, through measurement of actual performance) discard poorly performing neighbors or neighbors to whom asymmetric links exist. This represents a departure from traditional lower-layer design, where decisions are made at packet granularity (collision avoidance using RTS/CTS or link-layer retransmissions). At least for static (non-mobile) sensor networks, because pathological loss performance depends upon spatial positioning (*cf.* our gray area), it is meaningful to make decisions at the granularity of links to neighbors.

2. RELATED WORK

There is very little work that has extensively evaluated packet delivery performance on dense ad hoc wireless sensors. Woo *et al.* [23] examine a packet loss trace between

¹Our use of this term is slightly different from its use in the literature. Topology control in the ad-hoc context has meant the adaptation of transmit powers to enable higher spatial reuse [17], and some sensor networks work has used this term to denote mechanisms that selectively turn off nodes to reduce density and/or increase lifetime [24, 3, 2].

a pair of motes to construct packet loss models to evaluate link quality estimators. Zhao *et al.* [26] describe results from a measurement on a testbed of 26 motes and show the existence of links with high packet loss and link asymmetry. Most related to our work is that of Ganesan *et al.* [6], where packet loss is studied on a large-scale (approximately 180 motes) testbed grid on an unobstructed parking lot. That research also focuses on the loss and asymmetry of packet delivery at both the link layer and the MAC layer. In this paper, our study of packet delivery performance has more control of the topology that allows us to more carefully examine spatial and temporal characteristics. Moreover, our study examines packet delivery performance in harsher environments (indoor and habitat). We also examine different physical-layer encoding schemes and a wider variety of performance characteristics.

Measurements of infrastructure based wireless networks have been studied in [13, 21]. However, those studies focus more on the patterns of user mobility and their impact on traffic. Maltz *et al.* [15] describes a full scale testbed constructed for studying ad-hoc routing protocols. More recently, De Couto *et al.* [4] finds high variability in link quality, both on a wireless local network and a roof-top radio frequency network. They argue the given such variability, the widely accepted shortest path routing criterion is not enough. This class of measurement work is clearly complementary to ours since it focuses on a different kind of radio environment and different deployment densities than ours.

Finally, signal strength measurement has been used to understand different aspects of radio propagation properties such as modeling path loss for in-door environment [20]. The SpotON system [8] measures signal strength from low power radio transceivers to improve precision in localization systems. Our measurements of signal strengths are complementary, designed to examine the efficacy of signal strength estimation as an indication of link quality.

3. OVERVIEW, METRICS AND METHODOLOGY

In this paper, we take a first step towards understanding the performance of wireless communication in environments and at the densities that we expect sensor networks to be deployed. The primary aspect of wireless communication performance of interest to us is packet delivery performance. More precisely, our primary measure of performance is *packet loss rate* (the fraction of packets that were transmitted within a time window, but not received) or its complement, *the reception rate*.

There are many, many factors that govern the packet delivery performance in a wireless communication system: the environment, the network topology, the traffic patterns and, by extension, the actual physical phenomena that trigger node communication activity. It is difficult to isolate these phenomena in order to study the impact of different factors on packet delivery performance. Rather, we take a somewhat mechanistic view in this paper, and look at the packet delivery performance at two different *layers* in the networking stack: the physical layer and the medium-access layer.

We do this in a systematic fashion, in the sense that we exert some control over network topology, traffic generation, and the timing and duration of our experiments. Our experiments are not entirely *controlled*, however, since our mea-




Figure 1: Experiments in an indoor environment \mathcal{I}



Figure 2: Experiments in a habitat environment \mathcal{H}




Figure 3: Illustration of node placement in multi-hop experiments \mathcal{I}

measurements are subject to external factors, such as vagaries in the environment. This is deliberate, since (at least in part), we wish to understand how environmental factors affect communication.

Given our goal of understanding packet delivery in sensor networks, we employ a commonly-used sensor network platform: the Mica *mote* [10], its RF Monolithics radio [19], and the networking stack as implemented in TinyOS [9]. Clearly, this is a moving target; as the platform continues to evolve, the radio and the various protocols will continue to change. We address this by not making fine distinctions that could be invalidated by incremental improvements in the existing platform, and by pointing out which of our observations are likely to be affected by changes in technology.

In the following subsections, we discuss these experiments in a bit more detail.

3.1 Packet Delivery at the Physical Layer

The physical layer of most wireless networking stacks has two simple functions: framing and bit error detection or correction. These two functions are affected by many different factors. First, environmental characteristics can cause multi-path signal reception, or signal attenuation. Second, the spatial separation between sender and receiver can determine the received signal strength. Finally, minor variations in receiver and sender circuitry or in battery levels can adversely affect these functions of the physical layer.

To measure packet delivery at the physical layer, we use the following general setup. We place approximately sixty nodes in a chain topology. The precise pattern of node separation in this chain topology is discussed later. There is a single sender: the node at the head of the chain sends out a message periodically, and all other nodes receive. This simple setup measures the impact of the environment and the spatial separation between sender and receiver. It does not measure individual receiver or sender diversity; in fact, we are interested in the collective behavior or distributions of performances. We show that these distributions are not qualitatively affected by sender or receiver variations (we do this by permuting the physical setup).

We then place this setup in, and take measurements from, three different environments: an office building, a natural habitat, and an empty parking lot. The first two have been proposed as target environments for sensor networks, and the last represents a relatively benign environment that provides some calibration and context for our results.

With this setup, we can study several interesting ques-

tions: How does packet loss vary across environments? What is the spatial dependence on packet loss behavior? How are environmental effects and spatial dependence masked by different physical coding (error correction and detection) schemes? Are there spatial correlations in packet delivery? What are the temporal characteristics in packet delivery?

Note that even with such a carefully defined methodology, we will have only obtained a few data points on packet delivery performance. We do not claim that our experiments or our environments or the particular conditions under which we conducted our experiments are “typical” in any way. But we are fairly confident that our experimental conditions were not pathological either; we repeated some of our experiments at different times and did not observe any qualitative differences in our results. If anything, the actual behavior of these environments is likely to be worse than that reported in the paper, since we were careful to choose quiet times (*e.g.*, late nights in the indoor environment) for experimentation.

3.2 Packet Delivery at the MAC Layer

The medium-access layer has two functions² that impact packet delivery performance: arbitrating access to the channel, and (optionally) some simple form of error detection. In addition to factors that impact the physical layer, and hence the performance of medium-access, two factors affect the medium-access layer. First, the application workload (and, in the case of sensor networks, the sensed environment) determines the traffic generated by nodes and hence the efficacy of channel access. Second, the topology (or, equivalently, the spatial relationship between nodes) affects how many nodes might potentially contend for the channel at a given point in time.

To understand packet delivery performance as observed at the MAC layer, we use the following general setup. We place sixty nodes in a somewhat ad-hoc fashion, but at densities that we expect of sensor network deployments. Each node periodically generates a message destined to one of its neighbors; the periodicity of this message generation defines an artificial workload. We then place this setup in three environments as before, and measure several aspects of packet delivery performance.

The particular medium-access layer we choose is the default MAC that is implemented in TinyOS (henceforth called

²Other functions, such as node addressing, are orthogonal to the *performance* of packet delivery. Many traditional medium-access layers are also interested in fairness, a subject we do not evaluate in this paper.

the TinyOS MAC). It incorporates a simple collision avoidance scheme, and has a link-layer acknowledgment scheme to which we added a retransmission mechanism that enables us to study the efficacy of link-layer error recovery. Of this MAC, we then ask the following questions: What is the overall packet delivery performance observed by the applications upon MAC layer? Given such a density, what is capability of the MAC layer deal with interference introduced by simultaneous transmission? What is the *efficiency* of the MAC layer?

More than our physical layer experiments, there are many caveats to be aware of in our medium-access layer experiments. First, the TinyOS MAC is quite simplistic in that it does not include virtual carrier sense mechanisms like RTS and CTS for hidden-terminal mitigation. Our conclusions are somewhat limited by this; we intend to address this shortcoming by evaluating against S-MAC [25] when a stable implementation becomes available. We note, however, that many deployments of sensor networks using the TinyOS MAC are already under way; our performance measurements can give some understanding of behavior observed in the field in these and other deployments planned for the near future. As we discuss later, our results also give some insight into the design of future MAC layers for sensor networks. Second, we investigate one topology (or one node density) that we expect to be somewhat typical. We have no experimental data to justify this, but in our indoor office deployment, our network size corresponds to roughly one node per office.

Despite these caveats, we believe that there are many lessons to be learned from our experiments, as we discuss in Sections 4 and 5.

3.3 Instrumentation

Before we discuss the experiments, we discuss our experimental platform and the experimental instrumentation. This, together with a description of the actual experiments in later sections, should help convince the reader both of the logistical difficulty of conducting any kind of systematic study in these networks, as well as the care we have taken (to the extent possible).

We use Mica motes [10] in this study as the experimental platform. It is widely available and has been used in wireless sensor network research. Each Mica mote has a 4MHz Atmel processor (128K EEPROM and 4KB RAM), 512KB flash memory, and an ASK (amplitude shift keying) low power 433 Mhz radio [19]. We installed an omni-directional whip antenna to replace the built-in trace antenna on motes. In our experiments, the radio has a nominal throughput of 20Kbps. The low-level radio interface also supports the measurement of received signal strength, in a manner we describe later. Finally, the Micas come with an event-driven operating system called TinyOS [9]. TinyOS's networking stack includes a default physical layer that supports single-error correction and double bit error detection (SECDED) capabilities. On top of this, its default MAC layer implements a simple CSMA/CA scheme, together with link-layer acknowledgments.

To simplify experimental control and data collection, we used or wrote several pieces of instrumentation and experimentation software. The first such software module is a simple traffic generator. Driven by a clock which has an accuracy of one millisecond, the traffic generator repeatedly

sends out packets tagged with a sequence number. The exact periodicity depends on the experiment. A second module allows us to upload experimental parameters (such as packet sending rate, experiment duration) wirelessly to all motes within the radio range. We do this using a laptop connected to a mote's interface board. To store information about received data, we use the logger component built into TinyOS. At the end of an experimental run, we collect the motes and download the data from the logger to a central database.

In order to study the impact of more physical layers than the default one that comes with TinyOS, we implemented two schemes: a simple 4-bit/6-bit (or 4b6b) coding and a Manchester coding. We describe these schemes in greater detail in a later section. In order to study the correlation between packet loss and signal strength, we implemented a careful signal strength measurement module, whose details we reveal later. The TinyOS MAC delivers an acknowledgment; we added an optional re-transmission mechanism to understand the efficacy of link-layer ARQ (Automatic Repeat Request) to the MAC layer. Finally, in some cases, we used randomized intervals for packet generation based upon a precomputed set of random numbers. This allowed us to use different distributions (*e.g.*, exponential) for packet generation times than that allowed by the uniform random number generator available on the motes.

One of practical challenges in our experiments was to efficiently reprogram so many motes and download log content from them. We reduced the frequency of programming motes by parameterizing the experiments as much as possible. In addition, we implemented a simple negotiation protocol such that reprogramming and downloading is as simple as "plug-and-play" on multiple programming boards on multiple PCs. Together with audible feedback of the download progress, this reduces human intervention as much as possible and expedites the process of conducting experiments.

4. PACKET DELIVERY AT THE PHYSICAL LAYER

Our first set of experiments analyzes packet delivery performance at the physical layer. In this section, we discuss the methodology we use for our experiments, and then describe various aspects of packet delivery performance at the physical layer.

4.1 Detailed Methodology

The topology for this set of experiments consisted of approximately 60 motes, most of which were placed in a line at 0.5m apart. Guided by results from preliminary experiments, we intentionally removed some nodes from near the transmitter and placed more nodes at a finer granularity (0.25m apart) close to the edge of the communication range, giving us finer resolution in that region. Our node placement was therefore slightly non-uniform, and we are careful to account for this in our analysis. Finally, because we conducted experiments over several days, we were careful to mark node positions so that nodes could be precisely placed.

The traffic pattern for this experiment consisted of the node at one end of the line transmitting one packet per second, each with a monotonically increasing sequence number. All the other nodes merely received packets and recorded received packets in local storage. In order to purely measure




Figure 4: Packet Loss with 4b6b coding, high Tx power

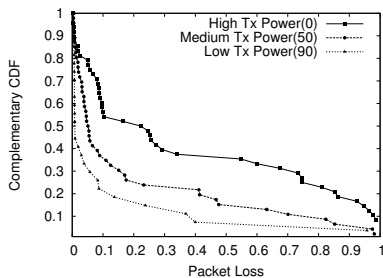


Figure 5: Packet loss v.s. Tx power in \mathcal{I} , 4b6b coding




Figure 6: Packet loss v.s. coding schemes in \mathcal{I} , high Tx Power

packet delivery at the physical layer, we had to disable the TinyOS MAC layer. Because we have a single transmitter, the MAC layer’s carrier-sense and collision avoidance strategy is effectively non-operational. However, we had to modify TinyOS’s MAC so that its acknowledgment mechanism could be optionally disabled.

Using this basic setup, we varied three factors in our experiments: the choice of environments, the physical layer coding schemes, and the transmit power settings.

We chose three environments for experimentation:

- \mathcal{I} is an office building. The choice of this environment is motivated by in-building sensing applications [16]. In this office building, we placed our setup in a long hallway (2 meter by 40 meter) (Figure 1). This hallway poses a particularly harsh wireless environment, because of significant likelihood of multi-path reflections from the walls. This particular placement does not result in signal attenuation through walls or other obstacles, but may suffer from interference with other electronic devices (this is a somewhat remote possibility; we were operating the radios in the 433MHz band, which is allocated for amateur radio use in the US).
- \mathcal{H} is a 150m by 150m segment of a local state park (Figure 2). The choice of this environment is motivated by several recent efforts that seek to monitor habitats [1]. To conduct our experiment, we chose a downhill slope with foliage and rocks. As with \mathcal{I} , multi-path due to scattering from foliage and rock would contribute to a fairly harsh wireless communication in this environment as well.
- Finally, \mathcal{O} is a large, spacious parking lot (150m by 150m). Compared to the our two environments, it is relatively benign (no obstacles, and multi-path only due to ground reflections). \mathcal{O} provides some context for interpreting our other environments; it is hard to envision a sensor net in an open parking lot since there would be no interesting phenomena to sense.

Many of our experiments were conducted on different days. In conducting our experiments, we tried to keep the environment’s gross characteristics as consistent as possible (in addition to making sure we were able to replicate placement exactly, using markers). For example, in \mathcal{I} , we kept all the doors along the hallway closed, and conducted our experiments at late night hours, to minimize (but of course, not completely eliminate) interference from human activity. In this sense, our measurements from \mathcal{I} and \mathcal{H} report their “quiescent” state.

The second factor we varied in our experiments was the physical layer coding scheme. The default TinyOS SECDED coding encodes each byte into 24 bits. SECDED can detect 2 bit errors and correct one bit error. By contrast, the 4-bit/6-bit (or 4b6b) scheme encodes one 8-bit byte into 12 bits, with the capability of detecting 1 bit error out of 6 bits. The well-known Manchester coding scheme encodes each byte into 16 bits, with capability of detecting erroneous bit out of 2 bits. All of these coding schemes are DC-balanced. Of these schemes, 4b6b is the least error tolerant, followed by Manchester and SECDED. However, it is the most bandwidth efficient, using the fewest extra encoding bits.

Finally, the motes have hardware that allows discrete control of transmit powers. Specifically, the motes have a potentiometer that regulates the voltage delivered to the transmitter. Rather than explore the entire range of transmit power settings, we chose three qualitatively different settings: *high* (potentiometer 0), *medium* (potentiometer 50), and *low* (potentiometer 90).

Each experiment takes approximately 8 hours, though in the following sections, the analysis is based on the data in a window from hour 2 to hour 4. This allowed us to have an analyzable data set; we also examined other time windows and found the results to be in qualitative agreement.

4.2 Aggregate Packet Delivery Performance

Our basic metric for packet delivery performance is *packet loss*: the fraction of packets not *successfully* received (*i.e.*, passed CRC check) within some time window, where the time window will be clear from the context.

Sometimes, we measure its complement, the *packet reception rate*. We measure *packet loss* by analyzing the sequence numbers received at each receiver.

We first discuss a very gross measure of overall packet delivery performance to summarize our findings. For each experiment, we plot the *distribution* of packet loss within a two hour frame (*i.e.*, 7200 transmitted packets) across all the receivers. Such a metric can bring out the variability (or conversely the uniformity) of packet loss radially from a node.

Figures 4 through 6 illustrate the aggregate packet delivery performance for different environments, coding schemes and transmission power settings. Several interesting observations emerge from these graphs. These observations provide fodder for a more detailed analysis of packet loss. (The actual distributions plotted in all of these graphs is likely to be slightly different than if we had had more sample points




Figure 7: Spatial profile of packet delivery for \mathcal{I} , 4B6B, High Tx Power

Figure 8: Spatial profile for \mathcal{O} , 4B6B, High Tx Power

Figure 9: Spatial profile for \mathcal{H} , 4B6B, High Tx Power

closer to the source, as described in Section 4.1. However, the observations that we make below will hold in a qualitative sense.)

Almost all of these graphs show the existence of a significant tail in the distribution. In all graphs at least 20% of the nodes had at least 10% packet loss, and at least 10% of the nodes had greater than 30% packet loss. In later sections, we try to understand the possible causes for this distribution.

Figure 4 shows that different environments can be significantly different. Notice, first, that in plotting these graphs, we have used the least error tolerant coding scheme (4b6b). Our intent in doing this was to expose as close to worst case packet delivery performance as we could obtain. This gives us some measure of the harshness of the communication environment. \mathcal{I} is harshest for packet delivery; of our 60 nodes, almost 30 experienced a greater than a 10% loss rate and 20 over 30% over a 2-hour time window. \mathcal{H} isn't that much better, with about 20 nodes experiencing over 10% loss. Even in the "benign" \mathcal{O} , about eight nodes experience more than 10% loss.

Our next set of plots (Figure 5) depict the packet loss distribution for different transmit power settings. At lower transmit powers, the packet delivery performance improves significantly. At the lowest transmit power, the packet distribution of the in-door environment starts to resemble that of \mathcal{O} at the highest transmit power. We are not exactly sure why this is, but note that lower transmission power implies a shorter communication range. The reduced spatial extent of communication may reduce the likelihood of multi-path, contributing to better performance.

Finally, in our harshest environment and at highest transmit power, the differences between coding schemes (Figure 6) is quite evident. SECCDED coding has a noticeably lower incidence of high packet loss. This is not surprising, given its resilience to packet losses, and its error correction capability. What is surprising is that the other two schemes seem to statistically indistinguishable, even though Manchester coding is more tolerant to bit corruption.

While these distributions indicate significant packet loss pathologies, particularly indoors and in a habitat, they are coarse measures at best. In the next few sections, we look at more detailed measures of packet delivery performance.

4.3 Spatial Characteristics of Packet Delivery

In the previous section, we have observed some pathological packet loss behavior in our experiments. We now examine the *spatial* characteristics of loss in our experiments.

Specifically, for our linear topology experiments, we ask the question: *How does reception rate vary with distance from the transmitter?*

We visualize packet delivery as a function of distance by plotting the packet reception rate at each receiver over a 2 hour window, as before. In Figures 7 through 9), the reception rate is plotted as solid vertical bars. In addition, we plot, using gray boxes, those packets that were received completely, but did not pass the CRC (*i.e.*, were corrupted). The remaining packets were those for which it was not possible to detect the start symbol of the physical layer frame. Because the placement of node is not strictly uniform, we mark the existence of each node by a dot on the top region of each graph.

Figure 7 plots the spatial reception profile for \mathcal{I} and 4b6b coding at maximum transmit power. For our harshest environment and our least capable physical layer, one notices a very interesting phenomenon. There are two distinct regimes of reception rate: up to a certain distance from the sender, even in \mathcal{I} , packet reception rates are uniformly high. Beyond this, however, there exists a *gray area* in which reception rate varies dramatically; some nodes see near 90% successful reception, while neighboring nodes sometimes see less than 50% reception rate.

More interesting than the existence of the gray area is the *extent* of the gray area. In Figure 7, the width of the gray area is almost *one third* the total communication range, while for the habitat (Figure 9) it is one-fifth! This is surprising, because while it has been known (at least anecdotally) that nodes at the "edge" of the communication range often see erratic packet reception, we are not aware of any literature that has suggested that this "edge" is significantly thick in some environments. Of our three environments, only in \mathcal{O} is the gray area relatively thin (about 10% of the communication radius). However, as we have argued, \mathcal{O} is uninteresting from a sensor network perspective, since it is relatively featureless.

Going back to our discussion of Section 4.2, we now understand that the links that see pathological packet loss are those that are in the gray area. Gray area symptoms also persist at lower transmit power settings. We have omitted these graphs for brevity. In a later section, we study how well different physical layer codings are able to mask the extent of the gray area. As we shall also see later, the temporal variation of packet reception is significant for nodes in the gray area.

Multi-path signal delivery constitutes a plausible explanation for the gray area phenomenon. Close to the trans-




Figure 10: Signal strength vs. distance (\mathcal{I} High Tx power)

	In Door	Out Door	Habitat
s_0	773.649	734.829	717.114
n	2.81	2.29	2.24

Figure 11: Path loss model parameter, high transmission

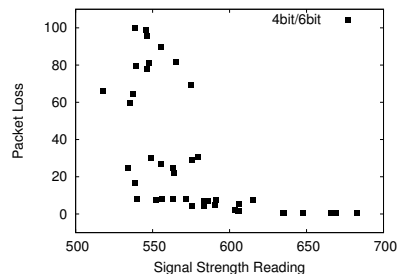


Figure 12: Signal strength vs. packet loss (high Tx power, \mathcal{I})

mitter, the direct signal is strong enough, and the reflected or scattered signals attenuated enough that reception rates are consistently high. Further away from the transmitter, the direct signal is weaker, and the reception rates depend upon the exact placement of nodes; at some nodes the signals destructively combine to ensure poor reception rates, while at others they combine constructively. It also makes intuitive sense that, in more severe multi-path environments such as \mathcal{I} , the physical extent of this phenomenon would be larger. It is difficult to precisely establish this hypothesis, but we did additional experiments to gain confidence in this assertion. In these experiments, we first empirically found the gray area corresponding to one transmitter; then, we placed a single receiver at different locations a few inches from each other. Sure enough, we found significant variability in packet reception rate across these locations.

An interesting question to ask is: How much of what we observe is an artifact of the particular radio that we use? We believe that our observations are not so much an artifact of the particular radio we use, but of the *class* of radios into which the RFM radio falls. Low-power radios without frequency diversity, we conjecture, are all likely to have similar multi-path rejection capabilities, although they may have widely different ranges if they use different modulation schemes (the Chipcon radio on the newer notes uses frequency shift keying and is reported to have a longer reach). Thus, unless sensor network nodes start using spread spectrum radios (which may use higher power transceivers and which are difficult to architect in an ad-hoc setting because they require fairly complex lower layer self-organization mechanisms that manage frequency or code acquisition), we believe the existence of the gray area will continue to hold qualitatively in these networks.

The width of the gray area has fairly deep implications for sensor networks. Sensor networks are designed to provide fine-grain monitoring of physical phenomena, which implies dense deployment in possibly harsh environments. The existence of a gray area implies that the likelihood of links falling into the gray area is high. For example, consider \mathcal{H} where the width of the gray area is $1/5$ th of the communication range. If we assume uniform deployment of sensor nodes around a given node, it follows that about $9/25$ th (or nearly a third) of a node's neighbors are likely to be in the gray area!

The existence of links with heavy packet loss can impact protocol performance in sensor networks. Consider the example of routing a message by selecting the next hop from neighboring nodes. Certainly selecting a shortest path

merely according to geographical distance or hop count is not sufficient. Such routing decisions tend to include long range links with high packet loss, which could result in poor end-to-end packet delivery³. In addition, including neighbors with high packet loss can also frequently trigger resets in soft-state or keep-alive mechanisms [11]. This fluctuation of system state not only leads to less consistent services, but also incurs heavier overhead and energy dissipation.

All of this suggests that, more than for other kinds of networks, nodes need to carefully select neighbors based on measured packet delivery performance. This kind of *topology control* is crucial in sensor networks, and is different from the kinds of topology control studied in the literature (power control [17], or density adaptation [24, 3, 2]).

4.4 Signal Strength and Packet Delivery

In the previous section, we saw how nodes in the gray area had unpredictable, and sometimes pathological reception rates. This suggested that sensor network protocols should detect (through measurement), and avoid using, poorly performing neighbors. Given this, a natural question to ask is: What kind of measurement strategy works for estimating poor quality links? This section discusses whether *signal strength* measurements can be used to estimate link quality. Such a technique is attractive because it can provide a simple technique for measuring link quality. (Of course, such a measurement would be made at the receiver and would somehow have to be communicated to the sender, which would then decide to not use the link.)

In other wireless environments, received signal strength provides a coarse indication of connectivity. For example, the received signal strength (and thus the estimated SNR) is provided by many 802.11-compatible devices as one primary metric to select preferred access points. It is not immediately obvious from this that signal strength can be used to determine whether a link will observe pathological reception rates.

Mica notes have a very simple functionality to provide received signal strength information: The RFM radio receiver provides the baseband output which rides on a DC level of approximately 1.1V. A decay in received signal causes the baseband modulation peak-to-peak amplitude to drop from a maximum of 685mV by approximately 10mV/dBm. The baseband output from the radio is directly routed to a 10bit A/D converter on the microprocessor. Note there is no inte-

³This observation has also been made for 802.11 links in [4], but that work does not remark on the existence of a gray area.


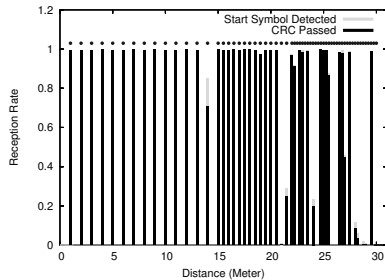



Figure 13: Spatial reception rate profile (\mathcal{I} high Tx power, 4B6B coding)

Figure 14: Spatial reception rate profile (\mathcal{I} high Tx power, SECCED coding)

Figure 15: Spatial reception rate profile (\mathcal{I} high Tx power, Manchester coding)

grator nor sample-and-hold circuit on the mote as suggested in [19, 8]. Thus, the reading from the A/D converter is the *instantaneous* voltage level of baseband signal.

We slightly modified the TinyOS radio communication stack, so that after transmitting the last byte of every packet, the transmitter sends out an 8-bit stream of alternating 0s and 1s. This last byte is not subject to any physical layer coding. Upon the receiving of the last byte of data packet, the receiver starts the A/D converter to continuously collect 10 samples of the baseband level every 15 microsecond, which covers approximately 3 bits of transmission at 20Kbps. This guarantees that at least one 0 and one 1 bit are covered. The maximum value of the samples is taken as the indication of signal strength. (We tried several other metrics, such as average and min, but found that the max value obtained most consistently matched our expectation of a signal strength measurement).

We now validate that what we have measured plausibly captures signal strength. The signal strength decay along distance can be summarized with a well known path loss model [20], where s_0 is the reading of the signal strength at the unit distance, n is the path loss exponent, $\Delta ADC = 1.707$ is the reading change in A/D converter when signal strength changes $1dBm$.

$$s(d) = s_0 - \frac{\Delta ADC}{dBm} * n * 10 * \log d \quad (1)$$

We fit this model to the observed data, and notice a good fit (Figure 10 shows the fit for \mathcal{I} ; values for other environments and other transmit power settings also yield qualitatively similar results and we omit these for space reasons), validating that indeed what we measure plausibly indicates signal strength. The exponents in Figure 11 are consistent with typical path loss environments in the range of 2 to 3.

We now return to the question of how well signal strength can be used to predict pathological loss behavior. Figure 12 depicts the correlation between signal strength and packet loss for the in-door environment at highest transmit power (other environments and transmit powers show similar behavior). It is generally true that high signal strength corresponds to low packet loss. For example, in our figure, links with a packet loss of less than 5% all have a signal strength greater than 550. However, the converse is not true; not all links with a signal strength greater than 550 correspond to low loss; some very pathological links have signal strengths higher than this threshold.

Thus, while it is seductive to assume that one can assess packet loss purely from signal strength alone, this may not

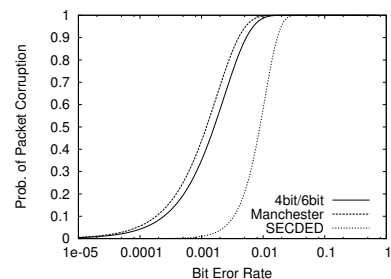


Figure 16: Packet loss rate as a function of bit error rate




Figure 17: Effective bandwidth for different coding schemes

work very well. At low signal strength, we suspect that individual receiver characteristics and local noise dominates, leading to significant variability in packet delivery performance.

4.5 Coding Schemes

Physical layer coding schemes provides resilience to bit error in packet transmission. Given our finding that there exists a noticeable gray area in packet reception (Section 4.3), we now ask the question: can sophisticated physical layer coding schemes *mask* this gray area? We have already seen that in an aggregate sense (Figure 6), SECCED performs noticeably better than 4b6b and Manchester coding given its error correction capabilities. The performance difference between these schemes is also evident from Figure 16, which plots the theoretical packet loss rate for a 36 byte packet as a function of bit error rate, for our various schemes.

Figures 13-15 show the spatial profiles of packet loss for our three coding schemes. Clearly, SECCED alleviates some




Figure 18: Correlation of packet loss for \mathcal{I} **Figure 19: Correlation of packet loss for \mathcal{O}** **Figure 20: Correlation of packet loss for \mathcal{H}**

of the variability in the gray area although it does not eliminate the gray area completely. However, this comes at a price. SECDDED reduces the effective bandwidth significantly, coding one bit into three. In the sensor network context, this translates to a relatively low energy-efficiency. Figure 17 shows the theoretical bandwidth (computed using a simple model of the three schemes) provided by these three schemes. By comparing Figure 16 with Figure 17, we notice that for packet loss rates of up to 50% (or 5% in SECDDED coding), 4b6b is more bandwidth efficient than SECDDED. Contrary to intuition, SECDDED does *not* increase the communication range. Packet reception is initiated if and only if a start symbol is detected. Thus the communication range is decided by the robustness of start symbol which is the same in those coding schemes.

This discussion suggest an interesting possibility. Given that the links of pathological loss are in the gray area, we can be more bandwidth efficient by using topology control (that prevents the use of pathological links), together with a more bandwidth efficient coding scheme. Of course, this possibility needs to be examined through actual implementation and experimentation, and depends on the efficacy of the topology control schemes.

4.6 Spatial Correlation

In this section, we examine whether there is any *spatial correlation* of packet loss among individual receivers: are two receivers in our linear topology likely to see similar loss patterns? If losses were correlated, that might be important from a modeling/simulation perspective (most simulations assume independent losses), or from a protocol design perspective (for example, if losses were independent, nodes could attempt to retrieve lost packets from nearby receivers rather than the source).

We formally define the packet delivery correlation coefficient between between two receivers i and j as:

$$R_{i,j} = \frac{\sum_{k=1}^n x_{ik}x_{jk} - n\bar{x}_i\bar{x}_j}{\left[\sum_{k=1}^n x_{ik}^2 - n\bar{x}_i^2\right]^{1/2}\left[\sum_{k=1}^n x_{jk}^2 - n\bar{x}_j^2\right]^{1/2}} \quad (2)$$

where $x_{ik} = 1$ if the k th packet is successfully received by node i , otherwise $x_{ik} = 0$. \bar{x}_i is the reception rate of n packets. This metric reflects the correlation in packet delivery both from packet loss and successful reception. Figure 18-20 plot the correlation coefficient between each pair of distinct nodes, as a function of the distance of the transmitter to each of the nodes (a 3-dimensional plot). For each environment, the graph is plotted based on 7200 packets (in 2 hours) transmitted with high power with 4b6b coding.

The plot shows significantly different correlation characteristics for different environments. Interestingly, \mathcal{I} and \mathcal{O} show noticeably higher correlated packet loss than \mathcal{H} . However, these two environments have different patterns of loss correlation. In \mathcal{O} , the correlations between nearby nodes are relatively strong almost everywhere except those near the transmitter. \mathcal{I} , on the other hand, shows an noticeable peak region in the middle of communication range.

We do not precisely understand the reasons for these differences (and in particular for the low correlation in the habitat), but we have some hypotheses. A more careful examination of \mathcal{I} reveals three regimes of correlation. Close to the source, the correlation between nearby receivers is minimal. In this region, we expect the direct signal is strong, and any packet loss is primarily due to strong local interference at receivers. Further away packet loss in the middle of the gray area shows relatively stronger correlation. When the signal is weaker, such fluctuation of environmental noise is more likely to cause correlated packet loss at nearby nodes. At the edge of the communication range, the signal strength is very weak due to direct path loss and can be different at nearby nodes due to severe multi-path cancellation. The environmental noise dominates so that a packet can only be successful received when the signal instantaneously becomes stronger. We also conjecture that the in-building has possible intermittent stronger interference from other devices which can be regional, which may also contribute to the difference between in-door and out-door environments.

However, overall, the highest correlation coefficient value is less than 0.7, indicating very moderate correlations, especially in the gray area. To a first order approximation, we conclude that *at the physical layer*, independent losses are a reasonable assumption. Of course, at the MAC layer, losses due to colliding transmissions can induce correlations.

4.7 Temporal Characteristics of Packet Delivery

In this section, we discuss one last aspect of packet delivery at the physical layer: How does packet loss vary with time, and what are the spatial characteristics of this variation? This question attempts to quantify how much variability there is within the environment, over our measurement window, and where the effect of this variability is felt.

If we compute average packet loss over a 40 second interval, Figure 21 shows that in a 2 hour window, packet delivery performance can be quite different at different times. A receiver near the edge of the communication range can have a packet reception rate that varies between 20% to 60%! Con-

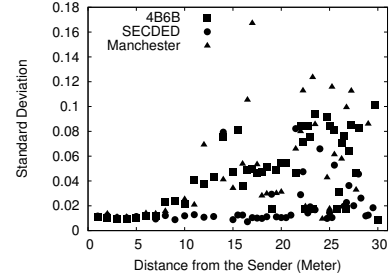




Figure 21: Packet reception rate over time (Window size=40 seconds) **Figure 22: Standard Deviation in Reception Rate for different environments (4b6b, high Tx Power)** **Figure 23: Standard Deviation in Reception Rate for different coding schemes (\mathcal{I} , high Tx Power)**

versely, over this averaging interval, a receiver close to the transmitter has little or no variability in reception rate.

Figure 22 describes the computed standard deviation of average reception rate over a 40 sec time window, computed from data taken over 2 hours for each receiver at the specified distance. (We have also plotted—but do not include—the standard deviation when packet loss is averaged over larger windows of 10s and 80s; those plots are similar to the ones shown here). These plots all have an interesting pattern. Radiating outward from the sender, the reception rate variance is low for all receivers up to some distance. Beyond a certain distance, the variance increases suddenly, and successive nodes see distinctly different packet reception variance. This again is a very graphic illustration of the gray area phenomenon we identified in Section 4.3. It indicates not only that some nodes in the gray area can see pathological loss, but that those nodes also see time varying packet loss. This has an important implication for schemes that perform topology control by excluding low quality wireless links. It is important for such schemes to continuously measure link quality, since reception rates can vary significantly over larger time-scales.

Finally, Figure 23 shows that the SECDED coding can mask packet delivery variability at most of the receivers. As such its gray area is smaller, but not non-existent.

4.8 Sensitivity of our Results

Transceiver Characteristics

Other work has observed that different transceivers can have significantly different characteristics [8, 22]. We ask: Are our results sensitive to the particular set of receiver placements we chose? To verify this, we conducted the same experiments using a different permutation of the transceivers: We obtain qualitatively the same results; for example, Figures 7 and 13 are from two different permutations but show the very similar behavior in “gray area”. While transceiver characteristics may well affect individual behavior (and indeed may affect RF ranging accuracy), they don’t seem to impact the ensemble results we present here.

Impact of the radio

It might be argued that our results are an artifact of the particular radio on the motes. We are not able to directly verify with the generation of motes with a frequency shift key radio. However, based on the result in a sparse measurement [12], we conjecture that our results will continue

to hold for low-power baseband radios—that unless the sensor nodes transition to wide-band radios, the multi-path rejection characteristics are likely to remain the same. Also note that our radios were in the 433 MHz ISM band. Using a 916 MHz ISM band, we might see slightly different attenuation characteristics, but our qualitative conclusions should be unchanged.

5. PACKET DELIVERY AT THE MEDIUM ACCESS LAYER

Thus far, our measurements have been with a single transmitter to study the packet delivery performance on the physical layer with a null MAC. In this section, we pop up a level in the stack and examine the packet delivery observed with the TinyOS MAC layer, under different traffic loads and environments. The intent here is to examine packet delivery performance that sensor network applications will see in the relatively dense deployments.

5.1 Detailed Methodology

For the medium-access layer experiments, many aspects of our experimental methodology are similar to that of our physical layer experiments. We conducted experiments in the same three environments (\mathcal{I} , \mathcal{H} and \mathcal{O} as described in Section 4.1), and our instrumentation infrastructure was very similar. Since our interest was in measuring the MAC layer performance we did not vary the physical layer coding schemes (except for one experiment designed to test the sensitivity of our results) More substantively, our MAC layer experiments differed in two other aspects: the topology, and the traffic pattern.

Unlike the regular placement of nodes in the physical layer experiments the Section 4, we placed nodes in a somewhat ad-hoc fashion but at densities we expect of sensor network deployments.

For the in-building environment \mathcal{I} , we placed 62 motes in the hallways as illustrated in Figure 3. The placement is ad-hoc but at the granularity of one mote per office. To us, this represents a realistic deployment in the sense that any sensor network in that building has to have a density of at least one sensor per office. With careful measurement of the deployment configuration, the experiments were also repeated on approximately the same topology on the outdoor parking lot environment \mathcal{O} . However, it was difficult to recreate the same topology for the habitat environment \mathcal{H} , where the terrain is irregular. Furthermore, the communication range in the habitat environment is considerably




Figure 24: Packet loss distribution in \mathcal{I}




Figure 25: Packet loss distribution in \mathcal{O}




Figure 26: Packet loss distribution in \mathcal{H}

shorter than the other two. For the habitat, therefore, we placed nodes on a 4 by 12 grid with an approximate 0.75m spacing between neighboring nodes. For each environment, we carefully marked the position for each node so the deployment can be recreated for experiments on different dates.

Despite the fact that the topologies are not identical in three environments, our goal was to obtain a reasonably *dense* deployment. Based on some preliminary experiments, we set the transmit power to medium (a potentiometer setting of 50) which resulted in node degrees between 15 to 18 for \mathcal{I} , 17 to 20 for \mathcal{O} , and between 6 to 8 for \mathcal{H} , and the diameter of the topology in hop count (3 to 4 hops in each case).

The traffic pattern is very simple: each node sends roughly k packets per second with an exponentially distributed inter-packet interval (to avoid synchronization). The traffic load on the network can be adjusted by changing the average load k . This traffic pattern is not intended to model application traffic; sensor networks are still in their infancy, and it is unclear what traffic models will be meaningful. Rather, our goal is design a simple tunable workload using which we can examine the both how simultaneous transmissions (at high loads) and environmental characteristics (at low loads) affect MAC layer mechanisms.

Each node unicasts its packets to its neighbors in a round-robin fashion. Periodically, each node broadcasts a packet so that all nodes can construct their neighbor lists. Nodes log received packets in the mote’s flash memory. Each experiment is executed long enough such as that at least of 200 packets are transmitted to each neighbor. Most experiments last at least two hours.

The TinyOS MAC protocol uses a CSMA/CA protocol. It has no virtual carrier sense (RTS/CTS), but will backoff for a random time whenever it detects a concurrent transmission in the carrier. The TinyOS MAC provides a link layer acknowledgment; upon successful reception, the receiver sends a 4-byte acknowledgment packet to the sender. The default MAC layer does not take advantage of this information. We added a layer to retransmit the packet immediately when no acknowledgment is received, up to a tunable limit which we set to 3. We did this after noticing the variability in packet reception in our physical layer experiments, to see how well a link-layer loss recovery scheme⁴ could mask the vagaries in wireless communication. Finally,

⁴Although, for ease of exposition, we present our results as a link-layer loss recovery scheme, what we are really measuring is the efficacy of simple ARQ schemes (with bounded number of retransmissions) at any layer to overcome the packet loss rates seen in our various environments.

the physical layer we chose uses the 4b6b coding, given its high bandwidth efficiency. We later consider whether a different choice of physical layer coding (such as SECDED) would have produced a different result.

In the following sections, we describe several different metrics of packet delivery performance at the MAC layer.

5.2 Aggregate Packet Delivery Performance

Our first metric is the distribution of packet loss across all links, shown in Figures 24 through 26. We say a packet is lost at the MAC layer only if our link-layer recovery scheme fails to deliver the packet. Packet losses could be due to corruption at the physical layer or due to collisions.

There are two distinct regimes in \mathcal{I} (Figure 24): *low-load* (up to approximately 1 pps) and *high-load* (above 1 pps). Note that 3 pps is close to the nominal capacity of the 20Kbps radio channel for 36-byte packets encoded with 4b6b scheme transmitted on a network with average degree of 15. Even at a low load of 0.5 pps, we note that there is a significant heavy-tail in the distribution: around 35% of the links with packet loss worse than 50%. This indicates that even a reasonable link layer loss recovery is unable to mask the high packet losses one sees in the harsher environments. Predictably this behavior worsens with load: more than 50% of the links observe 50% packet loss under traffic load of 2 pps. Also, the habitat results are qualitatively similar, even though the densities were much lower.

It would be very interesting to distinguish packet losses due to collision of simultaneous transmissions from those due to environmental noise. However, we have found it difficult to design experiments to do this. We have attempted to correlate packet loss to simultaneous transmission. One way to do this, perhaps, is to time-stamp each transmission and then infer that a lost packet at a receiver might have been caused by simultaneous transmissions. This needs fine-grain clock synchronization, which is not available in current TinyOS. Even so, we cannot be absolutely certain that the transmissions actually collided at the receiver (for example, at the instant those packets were sent, the propagation characteristics may have changed and the receiver may not have been within range of one of the senders).

Thus, the range of losses that one sees in these environments translates to pessimistic packet delivery performance. We argue that the incapability of the TinyOS MAC with retransmission results from the nature of *dense* deployment together with the relatively high occurrence of pathological connectivity. Certainly, a more capable MAC-layer with a virtual carrier sense (*e.g.*, S-MAC) can eliminate many hidden-terminal effects that occur in dense deployments.




Figure 27: Delivery efficiency in \mathcal{I}




Figure 28: Average Efficiency

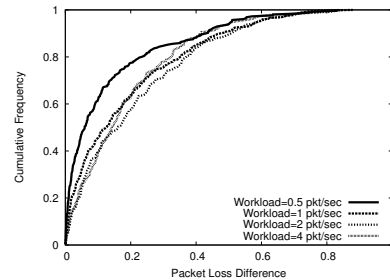


Figure 29: Packet Loss difference distribution (\mathcal{I})

Unfortunately, we could not validate this since no stable implementation of such a MAC exists for the motes. In addition, we believe that topology control mechanisms which reject poorly performing links can greatly improve MAC-layer performance.

5.3 Packet Delivery Efficiency

Packet loss distributions tell only part of the story. Recall that our MAC has link-layer error recovery. In this section, we try to measure the *useful* work done by the system in the presence of such an error recovery scheme. For a given link, we measure the useful work done over that link using a metric we call *efficiency*, which is defined as the ratio of the *distinct* packets received and the packets transmitted including retransmission.

We intend to capture the efficiency of link layer retransmission, so our definition does not count the overhead from coding schemes or preamble for packets. Note that the efficiency metric does *not* measure channel utilization. Rather, because it measures the useful work done as a fraction of total work done, it gives us some indication of the energy wasted by the system in overcoming packet losses.

Like the packet loss distributions, distributions of efficiency for different environments (for example Figure 27 for \mathcal{I}) show heavy tails. The performance is fairly pessimistic. In Figure 27 at light loads nearly 50% of the links have an efficiency of 70% or higher, but at heavy loads, nearly 50% of the links have an efficiency of less than 20%. The habitat environment is a little more benign with higher efficiency. This is evident in the average efficiency curves (Figure 28 as well). With increasing load, the average efficiency drops from 50% down to 20%. It also shows that coding with SECDED scheme in \mathcal{I} does improve the efficiency, however the advantage is reduced at higher workload. In addition, coding overhead is doubled in SECDED scheme thus the actual goodput (*i.e.*, effective bandwidth times efficiency) is actually less than with 4b6b coding.

Thus, depending on the load, *anywhere between half and 80% of the communication energy is wasted on repairing lost transmissions*. Even under lightly loaded conditions, the prevalence of pathological links dramatically reduces the efficiency of the system. This, to us, is a colossal expenditure of energy in these systems and warrants an investment of effort in the development of a good MAC layer for sensor networks.

5.4 Asymmetry in Packet Delivery

The final aspect of MAC layer performance that we explore is asymmetry in packet delivery. Asymmetry occurs

when a node can transmit to another node but not vice versa. The existence of asymmetry in wireless communication is well-known [4, 6, 26]. However its *extent* is less well understood, particularly in densely deployed wireless networks. In this section, we examine the asymmetry in packet delivery using a *packet loss difference* metric for a link pair between i and j , defined as follows:

$$D_{asym} = |P_{i \leftarrow j} - P_{j \leftarrow i}| \quad (3)$$

Notice that we are measuring the asymmetry observed at the MAC layer, which is complicated by possible packet collision in addition to environmental factors. However, on the other hand, the measurement is more “realistic” in a sense that it reflects what application experiences in reality.

Figure 29 shows distribution of packet delivery asymmetry in \mathcal{I} . Asymmetric links are quite common. More than 10% of link pairs have packet loss difference $> 50\%$, even for light loads where one expects fewer collisions contributing to packet loss. The results for the habitat (not shown) are similar.

A possible explanation for asymmetry is the difference in transceiver calibration (slightly different transmit powers, or differences in receiver circuitry). We have experimentally observed that for a given transmitter, different receivers exhibit slightly different reception rates at the same spatial separation. The reverse is also true; with a fixed receiver, different transmitters result in different reception rates at the same spatial separation. However, these differences are not enough to quantitatively explain our observed asymmetry. More extensive experimentation is needed to establish the cause of asymmetry.

Such asymmetric links are well-known for their impact on routing [18] and network aggregation [11, 14, 26]. The fraction of asymmetric links is high enough that topology control mechanisms should, we argue, carefully target such links, in addition to rejecting links exhibiting pathologically performing links.

6. CONCLUSIONS

In this paper, we have described results from a collection of measurement experiments designed to understand the packet delivery performance in dense sensor network deployments under realistic environments. Our findings quantify the prevalence of “gray areas” within the communication range of sensor radios, and indicate significant asymmetry in realistic environments. We have not yet been able to devise experiments that indisputably establish *causes* for these findings (although we have plausible conjectures, such

as multi-path, and have ruled out other causes, such as transceiver calibration); we leave this for future work. While our measurements indicate that the performance in these environments is fairly pessimistic, we believe simple topology control mechanisms will go a long way towards improving performance.

Our work has several other contributions. It provides some insights for modeling packet loss. As well, our measurement data can be useful as trace-driven simulation input to sensor network simulations. Finally, our experimental methodology itself is a first step towards a set of systematic techniques to study the performance of sensor networks in various environments.

7. ACKNOWLEDGMENTS

We thank Deborah Estrin for some stimulating discussions at various stages of the work. Wei Ye (USC/ISI) and Jie Nie from RFM provided their valuable help and feedback in analyzing our results, and John Heidemann (USC/ISI) provided logistical support and encouragement. Members of USC ENL Laboratory, the anonymous reviewers and our shepherd Joe Hellerstein provided insightful comments that significantly improved this paper.

8. REFERENCES

- [1] A. Cerpa, J. Elson, D. Estrin, L. Girod, M. Hamilton, and J. Zhao. Habitat Monitoring: Application Driver For Wireless Communications Technology. In *Proceedings of the ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, San Jose, Costa Rica, April 2001. ACM.
- [2] A. Cerpa and D. Estrin. ASCENT: Adaptive Self-Configuring sEnsr Networks Topologies. In *Proceedings of the IEEE Infocom*, volume 3, pages 1278–1287. IEEE Computer Society Press, June 2002.
- [3] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An Energy-efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks. In *Proceedings of the IEEE Infocom*, pages 85–96. IEEE Computer Society Press, 2001.
- [4] D. De Couto, D. Aguayo, B. Chambers, and R. Morris. Performance of Multihop Wireless Networks: Shortest Path is Not Enough. In *Proceedings of the First Workshop on Hot Topics in Networks (HotNets-I)*, New Jersey, USA, October 2002.
- [5] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next Century Challenges: Scalable Coordination in Sensor Networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, Seattle, WA, USA, August 1999. ACM.
- [6] 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. In *Technical Report UCLA/CSD-TR 02-0013, Computer Science Department, UCLA*, July 2002.
- [7] W. R. Heinzelman, J. Kulik, and H. Balakrishnan. Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 174–185, Seattle, WA, USA, August 1999.
- [8] J. Hightower, C. Vakili, G. Borriello, and R. Want. Design and Calibration of the SpotON Ad-Hoc Location Sensing System, 2001.
- [9] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System Architecture Directions for Network Sensors. In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 93–104, Cambridge, MA, USA, November 2000. ACM.
- [10] M. Horton, D. Culler, K. Pister, J. Hill, R. Szewczyk, and A. Woo. MICA, The Commercialization of Microsensor Motes. In *Sensors Magazine*, pages 40–48, April 2002.
- [11] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 56–67, Boston, MA, USA, August 2000. ACM.
- [12] J. Jeong and S. Kim. CS262A Class Project: DOT3 Radio Stack. Technical Report Unpublished, University of California, Berkeley, January 2003.
- [13] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proceedings of the Eighth Annual International Conference on Mobile Computing and Networking*, pages 107–118, September 2002.
- [14] S. Madden, M. Franklin, J. Hellerstein, and W. Hong. TAG: A Tiny Aggregation Service for Ad hoc Sensor Networks. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation*, 2002.
- [15] D. Maltz, J. Broch, and D. Johnson. Quantitative Lessons from a Full-Scale Multi-Hop Wireless Ad Hoc Network Testbed. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, September 2000.
- [16] J. Rabaey, E. Arens, C. Federspiel, A. Gadgil, D. Messerschmitt, W. Nazarov, K. Pister, S. Oren, and P. Varaiya. Smart Energy Distribution and Consumption: Information Technology as an Enabling Force, 2002.
- [17] R. Ramanathan and R. Hain. Topology Control of Multihop Wireless Networks Using Transmit Power Adjustment. In *Proceedings of the IEEE Infocom*, pages 404–413. IEEE Computer Society Press, 2000.
- [18] V. Ramasubramanian, R. Chandra, and D. Mosse. Providing A Bidirectional Abstraction for Unidirectional Ad-Hoc Networks. In *Proceedings of the IEEE Infocom*, June 2002.
- [19] RF Monolithics Inc., <http://www.rfm.com/>. *ASH Transceiver TR1000 Application Notes*.
- [20] S. Y. Seidel and T. S. Rappoport. 914 MHz Path Loss Prediction Model for Indoor Wireless Communication in Multifloored Buildings. In *IEEE Transactions on Antennas and Propagation*, volume 40(2), pages 207–217, February 1992.
- [21] D. Tang and M. Baker. Analysis of a local-area wireless network. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 1–10, N. Y., August 6–11 2000. ACM Press.
- [22] K. Whitehouse and D. Culler. Calibration as a Parameter Estimation Problem in Sensor Network. In *Proc. ACM Workshop on Sensor Networks and Applications*, Atlanta, GA, 2002.
- [23] A. Woo and D. Culler. Evaluation of Efficient Link Reliability Estimators for Low-Power Wireless Networks. Technical Report number to be assigned, University of California, Berkeley, April 2003.
- [24] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed Energy Conservation for Ad Hoc Routing. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 70–84, Rome, Italy, July 2001. ACM.
- [25] W. Ye, J. Heidemann, and D. Estrin. An Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *Proceedings of the IEEE Infocom*, June 2002.
- [26] J. Zhao, R. Govindan, and D. Estrin. Computing Aggregates for Monitoring Wireless Sensor Networks. In *Proceedings of the IEEE ICC Workshop on Sensor Network Protocols and Applications*, Anchorage, AK, May 2003.