Multiple Receiver Strategies for Minimizing Packet Loss in Dense Sensor Networks

Bernhard Firner, Chenren Xu, Richard Howard, and Yanyong Zhang
Dept. of Electrical and Computer Engineering, Rutgers, the State University of New Jersey
New Brunswick, New Jersey, United States of America
{bfirner,lendlice,reh,yyzhang}@winlab.rutgers.edu

ABSTRACT

A typical wireless sensor network consists of many small sensors that collect instrument data around their locations and forward it to a central location for data processing. These networks can be deployed to monitor livestock and agricultural assets, products in a store, patients in a hospital, and so on. In many cases sensors have to be densely deployed, and collisions or overhead due to collision avoidance will considerably degrade the system performance below an application's required levels. With the decreasing cost of radio devices the obvious solution to this problem is the use of multiple receivers on different radio channels. However, we show that if receivers can be placed in different locations then increasing the number of receivers on a single channel will increase the rate of the capture effect and decrease collision losses, while also increasing the fairness of the transmitters' radio links. Not only can this single channel approach be more effective than using multiple channels, it is also required for some techniques, such as localization, where each receiver must be able to detect a transmission from any transmitter. We also show that the optimal choice between these two solutions is influenced by the radio attenuation rate and the number of receivers in the system.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]:
Network Architecture and Design—Wireless communication

General Terms: Experimentation, Measurement, Performance

Keywords: Scalability, Capture Effect

1. INTRODUCTION

Wireless sensor networks are becoming increasingly useful in applications such as remote health, habitat, and infrastructure monitoring. However, some applications, such as localization, tomographic imaging, critical system monitoring, or surveillance require constant high rates of traffic and will not scale well as the number of transmitters increases. Collision avoidance and packet scheduling techniques are often used to increase the scalability and throughput of systems like these. However, application throughput and latency requirements might not be satisfied even after using such techniques, making lower level solutions necessary. Consider a system with 100 sensors attempting to deliver a 10 millisecond packet once every second - this is possible if there is a zero overhead TDMA system with perfect time synchronization between nodes, but this is not generally possible.

Many of these high traffic systems, such as tracking and localization, passive mobility detection, intrusion detection, and radio interferometry can operate as single hop networks, or as networks where a small, low power node transmits data to a larger, more powerful node. Each of these single-hop networks will likely be deployed very densely, meaning that there will be many transmitters in a small area. At some point the load is greater than what the network's MAC protocol support, setting a limit on the system's scalability. A typical solution might be to divide network traffic across multiple channels to increase the available bandwidth.

The situation in multi-hop networks used for applications such as surveillance or critical system monitoring is similar. In a multi-hop network, power control can be used to limit interference from a transmitter to a local cluster of nodes with cluster heads acting as bridges between clusters [8]. However, power control cannot reduce the traffic density within a cluster since densely deployed sensor nodes may remain single-hop neighbors even at their lowest power levels. As a result, the cluster cannot be further subdivided and a single cluster head would be overwhelmed with the offered load of the cluster. In this case multiple cluster heads might be used to collect packets within the cluster, again using multiple channels.

An alternative solution to increasing scalability with the amount of traffic is to store and compress data before transmission. However, this can be unrealistic if the sensors are constrained by any combination of energy (battery life), processing power, or storage. Also, if the application has any data latency requirements, such as the minimum time delay of tracking, then such a solution is not viable. Finally, data compression is not always desirable as many applications require detailed observations at a fine time/spatial granularity.

The obvious physical solution to this problem, as mentioned earlier, is to use multiple channels to decrease the possibility of collisions. An early MAC protocol that operates a wireless sensor network across multiple frequencies was proposed in [1], but it required multiple transceivers

MobiHoc'10, September 20–24, 2010, Chicago, Illinois, USA.

for each sensor and operated them simultaneously, leading to high energy consumption. Later work, such as [3] and [12], have the sensor switch channels dynamically and thus need only a single transceiver. In [3] clusters are formed as in Leach[10] and cluster heads assign different channels to source and destination pairs within the cluster so that multiple transmissions can occur at the same time on different channels. Leach can balance the load placed upon cluster heads (for the rest of this paper called receivers for simplicity) but the network is still overwhelmed when all traffic is destined to a single data sink, as in a convergecast traffic pattern.

In [11] a multi-channel protocol based upon LMAC is proposed that has nodes dynamically switch to new channels as current channels are filled. If a channel is filled but a node wishes to send messages to a node within that channel it finds a bridge node to relay messages between channels. Although extra packets may be sent when messages must be bridged from one channel to another the number of collisions is reduced and thus energy consumption per packet is reduced. However, this overhead would raise packet latency in convergecast systems and is not as efficient as possible.

Two other MAC protocols specifically designed to deal with dense wireless sensor networks are Y-MAC[12], and crankshaft[9]. Y-MAC is a multi-channel MAC whereas crankshaft is a single channel MAC but both were specifically tested in one-hop, convergecast scenarios. Although they have good performance compared to other protocols tested, convergecast does not take advantage of multiple receivers and Y-MAC is not optimized for convergecast traffic.

In a multi-hop system with convergecast traffic multiple receivers could be chosen as cluster heads simultaneously in order to increase packet throughput, with responsibility for being a cluster head rotating from node to node[10]. However, with the decreasing costs of sensors and the versatility of modern single chip radios it is also possible to simply use multiple specialized basestations as cluster heads, (or as backbone nodes, as in [2]), especially in close proximity to the data sink where the amount of traffic will be the greatest. In a single-hop system this is an even more feasible scenario since it would be easier to supply power to devices acting as cluster heads.

In either case, with multiple receivers operating at the same time it may be possible to take advantage of spatial diversity on a single channel either as an alternative or as a supplement to spectral diversity offered by multiple frequencies. The focus of this paper will be on the most effective way to use additional receivers in a dense sensor system. With multiple receivers the rate of the capture effect, the process by which a radio can receive the stronger of two or more simultaneous signals, will increase. We will show through calculation and experimentation the trade-offs between allocating new channels with each additional receiver, using multiple receivers on the same channel spatially deployed to maximize the rate of the capture effect, and a combination of the two approaches. We will also show that the optimal solution changes with the system's topology but that a dynamically controlled system can easily adapt to its current situation and choose the optimal behavior.

In Section 2 we will describe the systems that our work applies to. Then, in Section 3, we will discuss some related work that either serves as a foundation for our work or that offers alternative approaches. After establishing our goals in the previous section we will explore the theoretical foundation of our approach in Section 4 and will validate that theory in Section 5. Finally, we will discuss the results of our experiments in Section 6.

2. SYSTEM OVERVIEW

The assumptions we make in this paper reflect our focus on several specific applications. These systems have densely deployed sensors that send packets at frequent, regular intervals. The densities we are concerned with are too high for time scheduling and collision avoidance alone to maintain a desired packet delivery rate. High sensor density and power constraints make multi-hop networking within a cluster of nodes impractical so one or more sinks or cluster heads act as receivers and connect each cluster with the rest of the network.

One emerging application in this category is agricultural monitoring with sensor networks, as described in [18]. For instance, the state of plants and soil quality can be determined by moisture sensors in the ground and cameras that observe plant height and greenness. Sensors on livestock can identify behavior patterns, such as the time spend sleeping or ruminating, which can aid in livestock management. High transmitter duty cycles and low packet loss rates are important for the real-time components of these mobile systems, such as in some behavior monitoring systems.

One application of an agricultural monitoring system is to prevent fighting between livestock, which is costly when injuries occur. In [19], bull velocity and proximity were monitored and when sensor readings hinted that a fight might occur, the bull received a mild electrical as a deterrent. In order to gather enough data to make accurate predictions of livestock behavior the transmitter duty cycle of this system was one transmission every half second. In the evaluation trial only 5 bulls in a paddock were equipped with transmitters, but a full scale deployment could have sensors attached to every animal in the herd.

A different kind of application that has the same network requirements is inventory control. An item level tracking system in a warehouse or department store will have a small sensor attached to every item. The sensors will periodically send radio beacons to prevent theft, take inventory, track the location of items, or monitor the temperature or humidity of sensitive merchandise. High duty cycles and a high packet delivery rate are necessary to provide low-latency alarms for theft detection and environmental monitoring.

One deployed sensor network whose scale was large enough to cause density concerns is Project ExScal[2]. ExScal consisted of approximately 1000 sensors called *extreme scale motes* (XSM), 200 backbone communication nodes called *extreme scale stargates* (XSS), and a single master operator node. Each XSS was responsible for 20-50 XSMs and the area covered was 1.3km by 300m. The main goal of the network was intruder detection, which results in bursty convergecast packet transmissions from the XSMs to the XSSs.

Initially a packet delivery rate of only 33.7% was achieved, but with Logical Grid Routing[14] and Reliable Bursty Convergecast[21] packet delivery rate rose to 99%. However, this was with only 20 to 50 nodes generating messages. Since the cost of the XSMs is much less than the cost of the more powerful and battery powered XSS a better system would operate with more sensor nodes per backbone node. Also, there are applications for such a network that will generate even more traffic than intrusion detection - passive mobility detection and radio tomographic imaging[20] for instance. Packet delivery methods that give high packet delivery rates even in densely deployed sensor networks with bursty, convergecast traffic are required for future sensor network applications.

The main challenge of the considered systems is to guarantee a high packet reception rate given the density and packet delivery rate, which we refer to as *scalability* in this paper. In order to provide good scalability, the basic strategy is to have multiple receivers operating at the same time. These receivers can either operate at multiple channels to take advantage of spectral diversity (such as in [12]) or stay in the same channel, carefully placed spatially to take advantage of spatial diversity (such as capture effect [15]). These two strategies have been looked at before individually, but in this paper, we try to explore the utility of the capture effect and will compare these two strategies to find an optimal receiver strategy.

3. RELATED WORK

In this paper we seek to determine the best way to use multiple receivers in a single-hop network cluster and compare the advantages of receivers on multiple channels with the increase in capture effect when multiple receivers are on the same channel. We will consider the case where a network is so densely deployed that it cannot be subdivided because the small size and high density of the cluster prevents division by power scaling or because subdivision would increase latency to an unsatisfactory level. The power-scaling techniques from [8] can be used to reduce a network to this state so the reader can assume that power scaling has already been used, if possible, and the network has been reduced to a single-hop cluster whose density is still too high for good performance.

In [12] and [11] energy efficient multi-channel MAC protocols are presented. These protocols are intended for multi-hop networks with arbitrary traffic patterns. While Y-MAC has good performance in a dense single hop scenario, it was not optimized for convergecast traffic and does not consider the trade-offs between multiple channels and increasing the capture effect. Thus Y-MAC addresses a different need than our work.

To evaluate the possibility of capture-aware systems that rely upon the capture effect to increase system scalability and throughput, we will explore the capture effect in greater depth than in previous work. It has been evaluated experimentally in both sensor networks[16] and in 802.11 networks[13]. In [15] a link-layer protocol, called Shuffle, is proposed to stagger packet transmissions to take advantage of the message-in-message capabilities of some wireless cards. Although the radio chips used in sensor networks do not have this capability, we can use two radio chips working in tandem to duplicate the benefits of message-in-message packet reception, as we will describe in Section 4.1.1. There is also work showing that, if multiple receivers are present, they can combine corrupted versions of packets in a form of error correction[5].

All of this existing work shows that having multiple receivers operating on the same frequency within the range of the transmitters can increase the successful packet reception rate. If transmitters are spatially separated and there are multiple receivers, also spatially separated, then two beneficial effects could occur. First, the capture effect might prevent some packet losses when packet collisions are expected. Second, corrupted packets received at multiple receivers might have enough information to be successfully corrected and decoded. This paper focuses on just the capture effect.

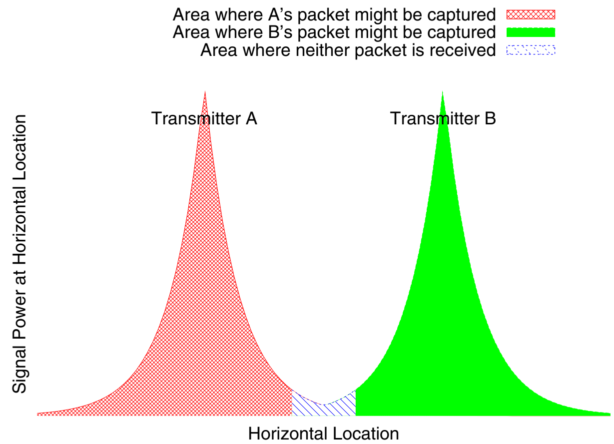

Figure 1: The signal strength and reception outcomes of signals simultaneously sent by two transmitters. If two transmitters send a packet at the same time then the receiver has a chance to receive the packet with the stronger signal strength. If the receiver is about equidistant from both transmitters then the packet will be lost, as depicted in the area between the two transmitters.

4. THEORETICAL BACKGROUND

Before designing any experiments to compare different multichannel and capture-aware transmission strategies, we need the ability to predict possible capture gains in different network environments. Work in [16] provides an initial look at the capture effect with the CC1100, a common radio transceiver in sensor networks. We will begin by expanding upon that work with analysis that includes details needed to create a general model to predict capture gains. We will also describe a modification to the CC1100 that increases packet reception and will investigate its effect upon the capture effect.

4.1 Capture Effect

A packet collision occurs when two or more radio signals overlap in time at a receiver. The capture effect occurs when the strongest one of those radio signals causes the other signals to be treated as noise and filtered out by the receiver. Thus a packet is received even though a collision occurred. This effect is illustrated in Figure 1.

Collision Region	% of Collisions
1	$100 \times \frac{\delta_{data}}{2\delta} = 10$
2	$100 \times \frac{\delta_{sync}}{2\delta} = 20$
3	$100 \times \frac{\delta_{preamble}+\delta_{data}}{2\delta} = 30$
4	$100 \times 1\frac{\delta_{sync}}{2\delta} = 20$
5	$100 \times \frac{\delta_{preamble}}{2\delta} = 20$

Table 1: The percent of collisions that fall into each collision region for this system.

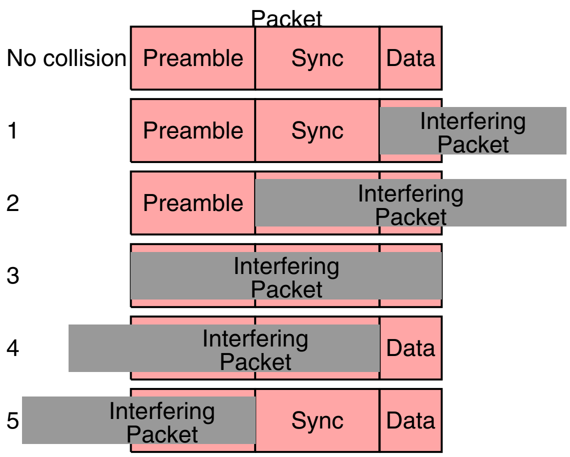

Figure 2: The five possible regions of packet interference with a three part packet. Packet regions with interference will have increased error rates.

A packet from the CC1100 radio has three main sections: a preamble, a sync word, and a data segment. More detailed information about how the CC1100 processes radio packets is available in [17]. Depending upon which parts of a packet suffer collision, the packet success rate will differ. There are five different regions of collision with this packet format, shown visually in Figure 2. The probability of each kind of collision depends upon the durations of the three sections, $\delta_{preamble}$, δ_{sync}, and δ_{data}, shown in Table 1.

4.1.1 Radio Chip Limitations and Redundant Receiving

Predicting exactly how often the capture effect will occur is difficult because it requires knowledge of signal strengths in a possibly changing environment and also requires detailed knowledge of the physical radio's behavior. An ideal radio would behave in an entirely predictable manner and would always receive the packet with the stronger signal strength during a collision. Unfortunately, the low cost transceivers used in many sensor networks are not ideal, as we will discuss.

There are two non-ideal behaviors of the CC1100. First, there is a delay after receiving a packet that causes packet losses when one packet immediately follows another. Second, once the radio commits to receiving a packet (after the sync word) it will not switch to receiving a different packet even if the new packet has greater signal strength and will cause decoding errors in the ongoing packet. Both of these problems can be partially fixed with what we call *redundant receiving*.

These problems occur because the CC1100 radio, and other similar transceivers, can only receive one packet at a time. Once the sync word has been received by the radio it commits to receiving that packet, even if a stronger packet arrives later. Thus the radio will attempt to decode a packet with a high error rate when it could successfully receive that stronger packet that began transmitting slightly later. The ability to receive the stronger, later packet in this situation is called *message-in-message* receiving and is discussed in detail in [15].

Although the CC1100 and other transceivers used in sensor networks do not have message-in-message capability, two radios can be used in tandem to achieve the same effect. When the first radio begins receiving a packet the second radio turns on and "covers" for the first one. Without a sync word the ongoing packet is just noise to the second radio but if a packet with a higher signal strength arrives it can be received successfully. We will refer to this method of achieving results similar to message-in-message receiving as *redundant receiving*. Redundant receiving does not completely fix this situation because there is a delay between when the first radio notifies the second radio to turn on and when the second radio is actually on. Even though it is still not ideal, redundant receiving does improve packet reception rates by a significant percentage, as will be shown in Section 6.

4.1.2 Evaluation Collision Behavior

To demonstrate the differences between collisions during different packet phases, we synchronized two radio transmitters together with a physical wire connection and had two transmitters create packet collisions in the different collision regions. Packets were sent using MSK modulation at 902.1 MHz with a 32 bit preamble, 32 bit sync word, and 16 bits of data with data whitening enabled[17]. Packet filtering based upon the quality of the preamble, an option of the CC1100, was not used as it would decrease packet reception rates. Differences in received signal strength (RSS) were measured by having each transmitter transmit a single packet without collision before each collision. One transmitter varied its transmission power over time to fill out all of the points on the curves.

Figure 3 shows packet reception rates under different collision conditions. Figure 3(a) shows the packet reception rates when interference starts after the packet begins and Figure 3(b) shows reception rates when interference was present when the packet began transmitting.

In Figure 3(a) the curve showing packet losses during collision region 1, when just the data segment suffers from interference, shows the effect of bit errors during decoding. The curve is exponential in appearance because it follows the probability of having no bit errors out of N data bits, $1 - N^{BER}$. Since we know the correct values of each bit in the packet we can calculate the bit error rate as a function of relative signal strengths. The result of this calculation appear in Figure 4, along with a best fit line. When interference is very strong, bit decoding succeeds as often as random guessing, or half of the time. As the packet being decoded becomes stronger than the interference, the BER falls quickly falls close to 0. As expected, once the packets are no longer overlapping or when the packet being decoded is much stronger than the interference, packet reception rates are near 100%.

The curve showing packet loss rates during collision region 2, with the sync word and data overlapping, shows that packet reception begins even before the packet has a stronger signal than the interfering packet. This is possibly because the radio locks on to the first signal it detects and thus misses the stronger packet. The weaker packet will have a large bit error rate however, and will thus often be received incorrectly.

When packets are transmitted at exactly the same time, as in collision region 3, successful packet decoding begins to

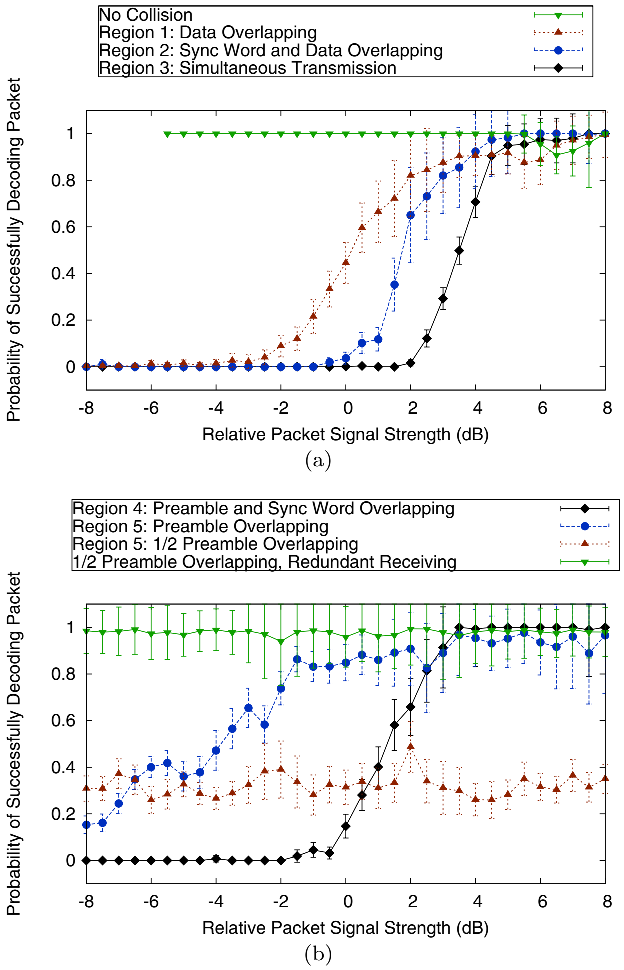

Figure 3: The packet reception rates for packets in collision regions 1-3 (a) and packets in collision regions 4-5 (b). The error bars are the estimated 95% confidence intervals.

occur when the packet is at about the same signal strength as the packet it collides with. As observed in [16] the probability of decoding a packet rapidly rises as one packet's signal strength becomes stronger, with a gray area in between a 0% probability of reception and 100% probability of reception. When there is a collision and a packet is 5 db stronger it will almost always be captured and successfully received.

Figure 3(b) shows packet collisions where the packet begins transmission in collision and ends transmission on a clear channel. In order for packet reception to occur, the sync word must be correctly detected. The curve where the preamble and sync word overlap, collision region 4, goes above a 0% chance of reception before collision region 3 because, for this packet size, this packet can draw out the sync word of the previous packet and force the radio to receive this one instead. If the data segment of the previous packet was longer by more than 60 $\mu seconds$ then redundant receiving would be necessary to receive these packets and the results of collision region 4 would be closer to the results of collision region 3.

With just the preamble in collision, as in collision region 5, a few different packet reception rates will be observed.

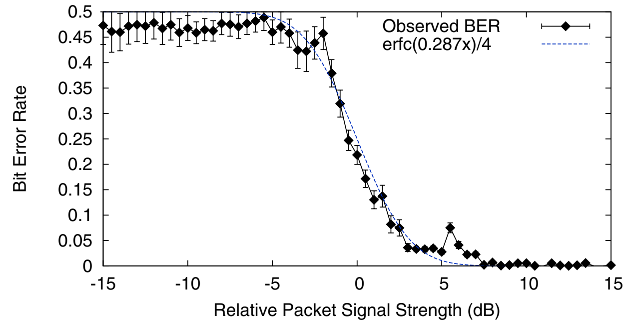

Figure 4: Observed bit error rates when the data segment of a packet has interference with the relative power shown in the x-axis. The error bars show the 95% confidence intervals. A line generated from regression analysis is also shown.

When the entire preamble is interfered with, packet reception depends upon the preamble of the packet drowning out the sync word of the previous packet and is thus correlated to the packet's received signal strength. The second receiver provided by redundant receiving will not turn on until the first radio begins receiving a packet so redundant receiving does not help in this case. If the data segment of these packets were longer, then we would see a curve similar to the simultaneous transmission case (region 3).

When just half of the preamble experiences interference, we see a flat line that does not vary with relative signal strength. These packet losses are due to a small, nondeterministic radio receiver delay in the CC1100 and packet processing after packet reception. As shown in the figure, these losses disappear with redundant receiving. The preamble is not decoded so it cannot be corrupted - it merely serves as a frame for the beginning of a packet - so interference during most of preamble has little adverse affects upon the packet. The CC1100 does allow packet filtering based upon the quality of the preamble, but this would actually decrease packet reception and was not used in our experiment.

4.2 Predicting Collision Losses and Capture Gains

We will begin our analysis by considering a transmit only protocol that periodically sends a fixed-length packet. This makes the analysis simple and the protocol is realistic for some energy constrained long lifetime systems[6], such as tracking and monitoring applications. For instance, if we wish to consider a MAC protocol that is split into a broadcast phase and an acknowledgement phase. During the broadcast, transmitters send packets to the sinks and during the acknowledgement phase the sink sends a single large packet with a bit field indicating the transmitter IDs of the packets that were received. The acknowledgement packet itself can be used to synchronize the transmitters and will lead to very low overhead.

Let us begin by calling the packet duration δ and duty cycle τ. Thus, every τ a sensor will transmit a packet of duration δ. To avoid successive collisions between sensors the value of τ could either be slightly different at each sensor or

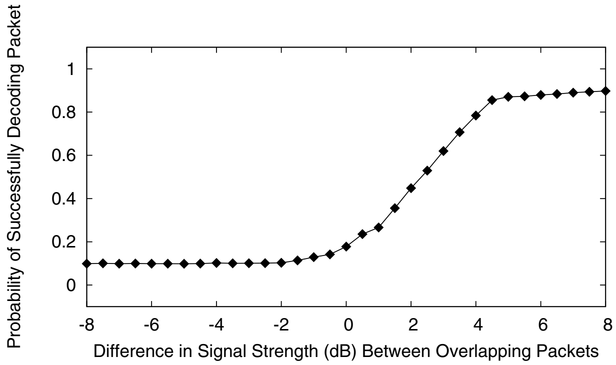

Figure 5: The cumulative probability of successfully receiving a packet despite a collision occurring. With redundant receiving, collisions during the beginning of the preamble (collision region 5) do not cause packet losses.

could change slightly from packet to packet. A collision will not always result in data loss however, because of the capture effect. We will call the probability of successful capture $P_{capture}$ and will provide more details shortly. Since transmission times are random and uncorrelated the probability that one sensor will have a packet collision with another is the probability that their packets will overlap with one another in time. This is the probability that one transmitter will begin transmission while the second is transmitting plus the probability that the second transmitter will begin transmitting when the first is already transmitting. Since the packet duration is δ the collision probabilities are

$$P_{2-waycollision} = P_{capture}\frac{2 \times \delta}{\tau} \quad (1)$$

$$\begin{aligned} &P(\text{collision}|\text{N transmitters}) \\ &= 1 - P(\text{no collision})^{N-1} \\ &= 1 - (1 - P_{capture}\frac{2 \times \delta}{\tau})^{N-1} \quad (2) \end{aligned}$$

The rate at which the capture effect occurs depends upon the environment, network topology, and specific radio in use, but the general equations for collision losses can still be formed.

If we take into account the duration of each of the 3 phases of the packet, we can construct the probability of capturing packets based upon relative signal strength. For each of the five phases in Figure 2 the probability of a collision falling into that region is shown in Table 1. From the values in this table and the results from Figure 3 we can construct the probability of packet loss given the relative signal strength of a packet to its interference. Those probabilities are shown in Figure 5.

To determine the probability of capturing a packet when the packet duration is different from the one used in these experiments, one need only recall that the error probability during data decoding for a packet with x bits of data is $(1 - BER)^x$. The bit error rates for the CC1100 appear in Figure 4 and can be used to adjust the capture rate in Figure 5 for different packet packets.

We can calculate $P_{capture}$ from the curve in Figure 5 once we find a function to predict probabilities of different relative signal strengths between two signals. Power loss during

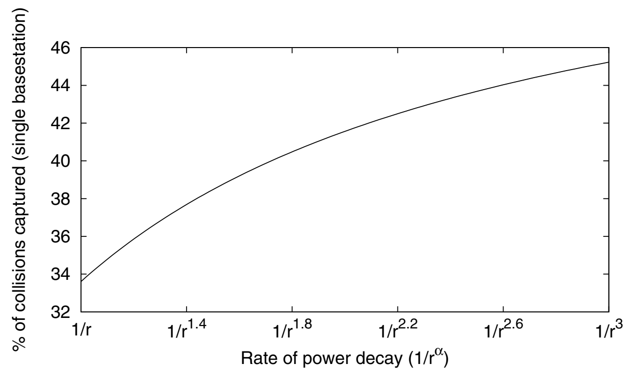

Figure 6: Theoretical capture effect gains when sensors are uniformly distributed about a receiver. Intuitively, rapid signal decay causes very different RSS levels at different receiver locations and increases the capture effect.

signal propagation is usually approximated by the formula P/r^{α}, where P is the power of the signal at the transmitter, r is the distance from the transmitter that the signal has travelled, and α is the attenuation factor. In free space $\alpha = 2$, but in our testing environment it was measured to be $\alpha = 2.69$. This measurement was done by taking four transmitters and four receivers and measuring the change in RSS from a distance of one foot to a distance of 40 feet in one foot increments. The attenuation factors for all 16 permutations of transmitters and receivers were calculated and then averaged to find $\alpha = 2.69$.

If we assume that the sensors are uniformly distributed then we can integrate over the uniform distribution to find the probability for each relative signal strength. We will convert a distance, Δ, to a relative dB amount for two transmitters at distances l_1 and l_2 and call the conversion factor C.

$$\begin{aligned} \frac{1}{l_1^{\alpha}} &\geq \frac{10^{\Delta/10}}{l_2^{\alpha}} \\ l_1 &\leq l_2 10^{-\Delta/10\alpha} \\ l_1 \leq l_2 C, \text{where C} &= 10^{-\Delta/10\alpha} \quad (3) \end{aligned}$$

We can now integrate over the probability density function of the uniform distribution to find the probability of a relative signal strength being greater than or equal to some value, Δ, by finding the probability that one transmitter is farther away than another transmitter by a factor of C.

$$\begin{aligned} &\int_a^b \frac{1}{b-a} \int_a^{cx} \frac{1}{b-a}\, dy\, dx \\ &= \frac{1}{(b-a)^2}\left(\frac{cb^2}{2} - ab - \frac{ca^2}{2} + a^2\right) \\ &= \frac{c}{2} \text{ if a} = 0. \quad (4) \end{aligned}$$

So if the receiver is placed in the center of a uniformly distributed set of sensors, or if the receiver is placed at the edge of a field of uniformly distributed sensors, then the probability of the relative signal strength values being greater than or equal to some value, Δ, is $10^{-\Delta/10\alpha}/2$, from equations 3 and 4. This relationship will be approximately correct for any receiver being used with uniformly deployed nodes. We

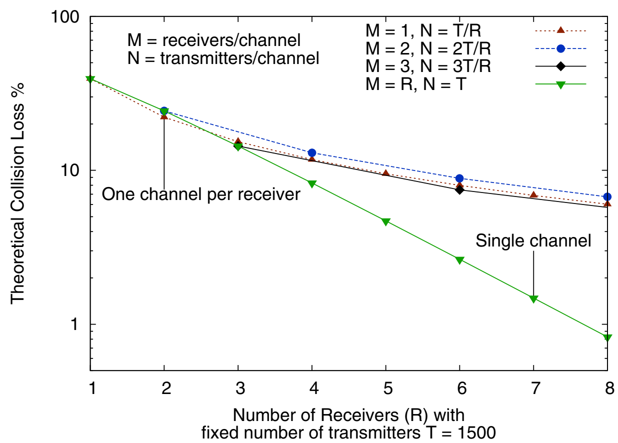

Figure 7: Theoretical packet losses with $\alpha = 2.69$, one 300 μ *second* transmission per second, and 1500 transmitters, using Equation 2 and $P_{capture}$ from Figure 5 and Equation 3. Once of the number of receivers passes a threshold a single channel is better than using multiple channels.

can now predict the rate of the capture effect for our packets with attenuation $\alpha = 2.69$. This is shown in Figure 6.

$P_{capture}$ from Equation 2 can now be determined from Figure 5 and Equation 4. Figure 5 shows the cumulative probabilites of packet success, which are broken down into probabilities for each relative signal value. These are multiplied by the probability of that signal value occuring, obtained from Equation 4 with $\alpha = 2.69$ to find the probability of a capture event at that relative signal strength level. These are then summed to find $P_{capture}$.

Using Equation 2 and inserting $P_{capture}$ we can create the theoretical packet loss curves shown in Figure 7. The value of α for the curves shown matches our test environment so these curves would be slightly different in other environments. These curves show that the most effective strategy for utilizing new receivers changes based upon the total number of receivers. When the number of receivers is small, using a new channel with each new receiver and distributing the transmitters evenly among the channels will decrease collisions most effectively. After the number of receivers passes a threshold though, it is better to use all of the receivers on the same channel and rely upon the capture effect to reduce collision losses. No mixture of the two approaches is worth pursuing.

We can also predict the theoretical packet loss rates for packets with larger data segments by using the bit error rate as a function of relative signal strength that was recorded previously. Figure 8 shows the theoretical packet loss rates for packets with the same sync and preamble length but with different data sizes. As the packet grows larger, the break-even number of receivers where it becomes better to operate on a single channel rather than on multiple channels shifts to higher numbers. Again, the attenuation factor used was $\alpha = 2.69$.

Our first experiment will verify the theoretical results in Figure 7, although for practical reasons instead of 1500 transmitters with a duty cycle of 1 second we will use 100 transmitters with a duty cycle of 0.1 seconds.

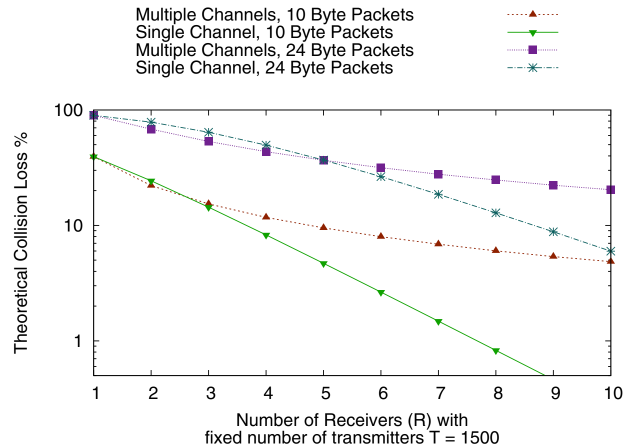

Figure 8: Theoretical packet loss with packets with different payload sizes.

5. EXPERIMENTAL SETUP

Our experiments will confirm that the magnitude of the capture effect has been predicted successfully by our analysis. Our experimental setup consists of 1 host PC serving as system controller, up to eight pairs of receivers (the receivers are paired to implement redundant receiving described earlier in subsection 4.1), and 100 transmitters sending a 10-byte ($300\mu second$) packet ten times per second, creating an offered load of 30%.

5.1 Hardware Description

The radio devices used in our experiments contain a Chipcon CC1100 radio transceiver and a 16-bit Silicon Laboratories C8051F321 microprocessor and are powered by a 20 mm diameter lithium coin cell battery, the CR2032. The receivers have integrated USB support for loss-free data collection but are otherwise identical to the transmitters. More complete information can be found in [7].

The radio link operates at 902.1 MHz. Transmitters use MSK modulation, a 250kbps data rate, and a programmed output power of 0dBm. Each packet contains 32 bits of preamble, 32 bits of sync word, and 16 bits of whitened data.

5.2 Test System Behavior

In our system, each transmitter will periodically send a 10-byte packet (8 bytes of sync and preamble and 2 bytes of payload) once every 0.1 seconds. The receivers will forward received packets to the host PC for analysis over a USB connection. The 10-byte packets being used in our system have an over-the-air duration of $300\mu seconds$.

We will perform tests to validate the predictions made in Figure 7. To validate the single channel, multiple receiver curve we will operate all of the transmitters and receivers on a single channel. We will then calculate the collision loss percent using data from a single receiver, from two receivers, and so on up to the results from all of the receivers.

To validate our collision loss predictions for multiple channel systems we will scale down the number of transmitters in a channel by the number of channels in use. For instance, we will generate results for a two channel system with two receivers by operating half of the transmitters with one of the receivers. Since there would be the same number of transmitters and receivers on channels 1 and 2 we can generate the overall collision loss percentage of both channels by just recording the collision losses on a single channel with half of the total transmitters. Likewise we can divide the total number of transmitters by four to simulate four channels, and so on. In our experiments we will generate results for a single channel, 2 channels, 4 channels, and 8 channels.

5.3 Test Topology

We will test a dense, short range topology in a 7 meter square area. Transmitters will be placed following a uniformly random distribution.

Receiver placement will be determined by using the landmark positioning work and the *maxL-minE* algorithm introduced in [4]. The maxL-minE (maximum lambda, minimum error) algorithm takes an optimal geometric pattern for the number of receivers and finds a deployment pattern by iteratively moving the receivers towards positions that achieve a local maximum in the deployment environment based upon a desired metric. In [4] the maxL-minE algorithm was used in a localization system so the optimization criteria we used was simpler - we started from a known optimal receiver pattern from [4] and maximized the distance between receivers in order to maximize the capture effect. Thus we expect our experimental results for the capture effect to be fairly close to the ideal analytical predictions. However, the analytical predictions assume receivers see each collision from a completely random vantage point. In reality transmitters at the edge of the deployment area are likely to be farther away to a receiver than other transmitters and will have slightly worse capture rates. Also, transmitters that are in the same location will always have a very low capture probability when their packets collide, no matter how many receivers are present.

A map of the experimental topology appears in Figure 9. Some of the marked transmitter locations actually have multiple transmitters, as can be seen in the photograph. The testing location was in the middle of a large open area to minimize differences in the attenuation factor within the area so that theoretical and experimental results could be fairly compared. The receivers are numbered in the order they were used. Results for a single receiver only used the receiver labelled "1", results for two receivers used "1" and "2", and so on.

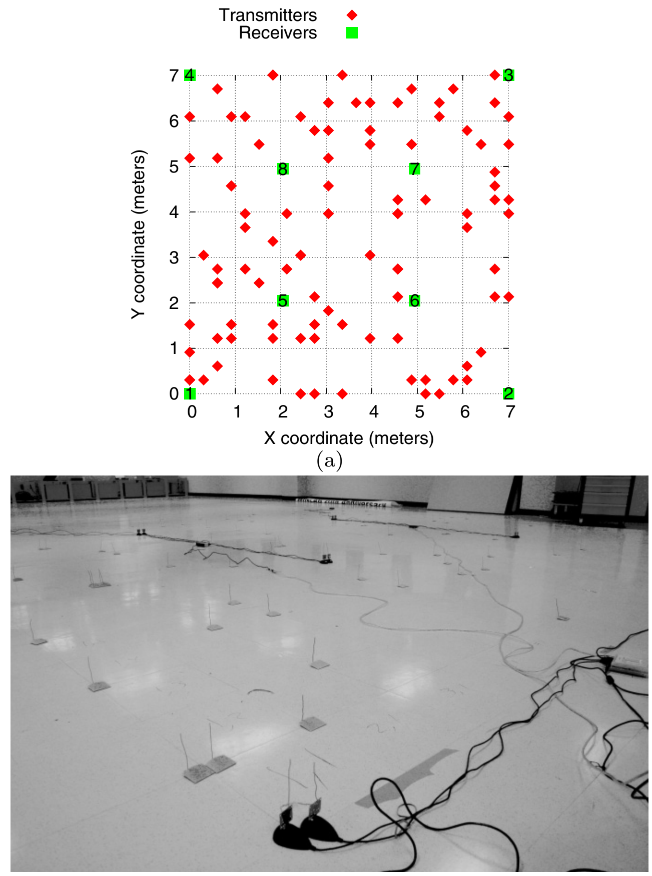

(a)

(b)

Figure 9: A gridded map of the experimental topology (a) and a photograph (b).

6. RESULTS

Experiment results are shown in Figure 10. These results confirm our theoretical model of the rate of the capture effect and also confirm that using multiple receivers on a single channel can result in fewer collision losses than using the same number of receivers spread over several channels. This gain is true whether or not the hardware supports message-in-message - Figure 10 shows that even without message-in-message the single channel approach outperforms the multichannel approach. This approach replaces increased spectrum usage with increased numbers of receivers in the same spectrum. The system offered 1000 packets per second and each packet had a duration of $300\mu seconds$ so the offered load was 30%. With this load 98.2% of packets were successfully received. The theoretical results from Figure 8 show that throughput should continue to rise even under higher offered loads.

The theoretical single channel results in Figure 10 are better than the experimental results. A slightly better experimental result might be achievable with careful receiver placement but worse results than the theoretical ideal should be expected in real deployment areas because after a certain number of receivers are deployed the differences between the ideal assumption of an infinite plane of transmitters and the reality of a finite space become more clear. Transmitters at the edge of deployment areas are more likely to have weaker signals than other transmitters since no transmitters are further away then they are, so their capture probabilities remain low. Larger deployment areas would probably achieve capture rates closer to the predicted values since they have smaller surface area to volume ratios.

Figure 10 also shows that the redundant receiving method was effective at decreasing packet losses by about 7.1% of what would be achievable with a single CC1100 radio chip. When multiple receivers are used on the same channel to increase the capture effect this increase is cumulative so the percent reduction in collision losses with 8 receivers on the same channel is 28%. The exact amount of this increase depends upon the packet size but these results show that using a second CC1100 transceiver as a redundant receiver can increase packet reception similar to the message-in-message capabilities of some 802.11 cards.

In addition to reducing packet collisions, the single channel approach also increases fairness in terms of relative packet loss between the transmitters. With a single receiver per

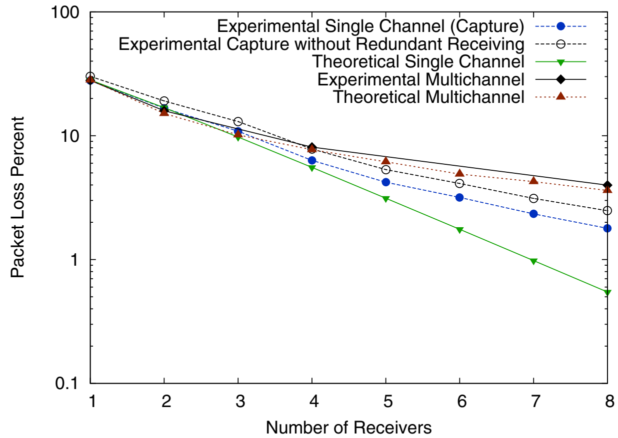

Figure 10: Experimental and theoretical packet losses with 100 transmitters sending ten packets per second. Results for single channel capture are shown with and without redundant receiving. As predicted, once of the number of receivers passes a threshold a single channel is better than using multiple channels, with or without redundant receiving.

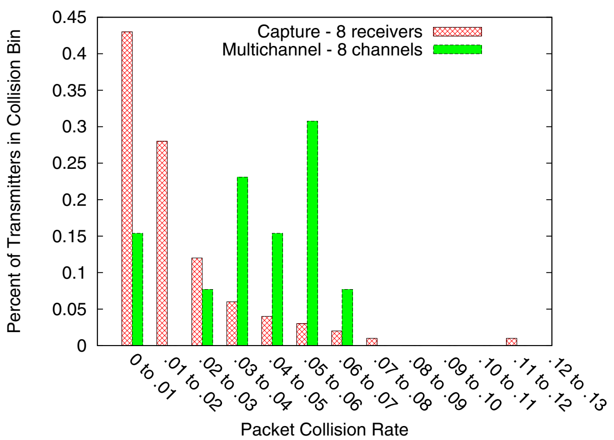

Figure 11: The frequency of different packet loss percentages for transmitters in the 8 receiver, single channel capture system and in the 8 receiver, 8 channel multichannel system.

channel the transmitters closest to the receiver will achieve much better packet reception rates than transmitters that are further away because of both the capture effect and increased bit error rates due to attenuation. With multiple receivers in the same channel at different locations though, each transmitter is more likely to have a nearby receiver so this non-uniformity is reduced and overall performance is improved. Figure 11 shows a histogram of packet loss percentages for the 8 receiver, single channel capture system with 13 transmitters and a single receiver per channel and the 8 channel, 8 receiver multichannel system with 100 transmitters. The packet loss rates of the multichannel system are scattered around several bins with varied performance. The packet loss rates in the capture system are mostly concentrated in the two bins representing the lowest collision rate and rapidly fall off. The average loss rate of the single channel system is only 1.8% compared to 4.0% in the multichannel system. Thus the loss rates of the transmitters in the capture system are more consistent from transmitter to transmitter. There is no way to achieve this level of fairness with a single receiver per channel without reducing the transmission rate of the transmitters closest to the receivers, which is also unfair in terms of transmission rate.

6.1 Implications for Capture-Aware MAC Protocols

With our results we can now make some initial recommendations for capture-aware MAC protocols in sensor networks. The transmit-only, convergecast system used for our experiment can be used as-is in tracking and localization, passive mobility detection, intrusion detection, and radio interferometry systems since they do not require a 100% packet success rate but do require frequent, regular transmissions. Transmitter battery lifetime is more important than guaranteed packet delivery in those systems so spending energy for synchronization, channel sensing, or retransmissions for a small increase in packet success rates would be detrimental.

These results may also be useful in networks that cannot use a transmit-only protocol. A capture-aware time-division MAC protocol for sensor networks would be similar to the proposal for 802.11 networks found in [15]. Transmitters using a collision avoidance capture-aware MAC protocol would still use channel sensing but might choose to transmit in the face of interference if the probability of a receiver being close enough to correctly receive its packet is high. A direction for future work is to study how a transmitter can estimate this capture probability.

It is possible to use the protocol that was tested in this paper as one phase of a MAC protocol with another phase for ACK messages. If there is little mobility of the nodes then it might be possible to compress the ACK messages to minimize the time spent in the ACK phase. The best way to schedule ACK messages in a capture aware system is another direction for future work.

7. CONCLUSIONS AND FUTURE WORK

We have shown that by taking advantage of antenna diversity by using multiple receivers on the same channel in different locations we can reduce packet collisions dramatically. Above a threshold in receiver number, we can obtain better performance, in terms of packet losses and fairness, than a multichannel system that uses the same number of receivers across multiple channels, whether or not the receivers have message-in-message capabilities. We have provided a theoretical model to estimate the expected gains from such a single-channel capture system. Using this model we also demonstrated that it is never worthwhile to use a hybrid approach - either all receivers should be on different channels to reduce temporal packet collisions or they should all be on the same channel to increase the capture effect. These results should serve as a starting point for the creation of capture-aware MAC protocols for sensor networks.

Using a single channel also means that the transmitters in such a system can use simpler protocols since no channel switching is required. This will lead to lower energy consumption and increased sensor lifetimes, especially in a single-channel system such as mobility detection or single-hop sensing systems where a transmit-only protocol is possible. An increase in fairness is also achieved, when compared to a multi-channel approach, because the increase in receivers on the same channel reduces the average distance between transmitters and receivers and reduces the disparity between the transmitters with the best radio links and the transmitters with the worst radio links.

This result is also important for research in localization and mobility detection systems that use signal strength measurements from multiple receivers. Although packet losses will increase as the number of transmitters increases, the receivers closest to a transmitter will still receive packets from that transmitter with high probabilities. Since these close receivers have the most important information about transmitter location and mobility, the information lost when distant receivers miss a transmitter's packet will have a smaller impact than might be expected.

We have also demonstrated a technique to achieve a message-in-message packet reception with the low-cost radio chips commonly used in sensor network research by using two transceiver chips on a single receiver. This can result in a significant reduction in packet collision losses and the required time between packet transmissions in a time scheduled system, as described in [15]. The power consumption of a receiver using two chips simultaneously would increase, but this technique is beneficial in systems with backbone nodes that are powered since the cost of deploying and maintaining additional nodes is much higher than the cost of an additional transceiver chip.

One advantage of a single channel system that was not explored in this paper is packet combining [5], where multiple corrupted versions of a packet that were received by different receivers can be combined to repair the packet and correct its data. This should yield further gains for the single channel system that relies upon the capture effect to reduce collisions because multiple receivers will listening to each packet.

8. REFERENCES

[1] Junaid Ansari, Xi Zhang, and Petri Mähönen. Multi-radio medium access control protocol for wireless sensor networks. In *SenSys '07: Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 403–404, New York, NY, USA, 2007. ACM.

[2] Anish Arora, Rajiv Ramnath, Emre Ertin, Prasun Sinha, Sandip Bapat, Vinayak Naik, Vinod Kulathumani, Hongwei Zhang, Hui Cao, Mukundan Sridharan, Santosh Kumar, Nick Seddon, Chris Anderson, Ted Herman, Nishank Trivedi, Chen Zhang, Mikhail Nesterenko, Romil Shah, Sandeep Kulkarni, Mahesh Aramugam, Limin Wang, Mohamed Gouda, Young-ri Choi, David Culler, Prabal Dutta, Cory Sharp, Gilman Tolle, Mike Grimmer, Bill Ferriera, and Ken Parker. Exscal: Elements of an extreme scale wireless sensor network. In *RTCSA '05: Proceedings of the 11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, pages 102–108, Washington, DC, USA, 2005. IEEE Computer Society.

[3] Xun Chen, Peng Han, Qiu-Sheng He, Shi-liang Tu, and Zhang-Long Chen. A multi-channel mac protocol for wireless sensor networks. In *CIT '06: Proceedings of the Sixth IEEE International Conference on Computer and Information Technology*, page 224, Washington, DC, USA, 2006. IEEE Computer Society.

[4] Yingying Chen, J.-A. Francisco, W. Trappe, and R.P. Martin. A practical approach to landmark deployment for indoor localization. In *Sensor and Ad Hoc Communications and Networks, 2006. SECON '06. 2006 3rd Annual IEEE Communications Society on*, volume 1, pages 365–373, Sept. 2006.

[5] Henri Dubois-Ferrière, Deborah Estrin, and Martin Vetterli. Packet combining in sensor networks. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 102–115, New York, NY, USA, 2005. ACM.

[6] Bernhard Firner, Prashant Jadhav, Yanyong Zhang, Richard Howard, Wade Trappe, and Eitan Fenson. Towards continuous asset tracking: Low-power communication and fail-safe presence assurance. In *Sensor, Mesh and Ad Hoc Communications and Networks, 2009. SECON '09. 6th Annual IEEE Communications Society Conference on*, pages 1–9, June 2009.

[7] Bernhard Firner, Shweta Medhekar, Yanyong Zhang, Richard Howard, Wade Trappe, Peter Wolniansky, and Eitan Fenson. PIP Tags: Hardware Design and Power Optimization. In *Proceedings of the Fifth Workshop on Embedded Networked Sensors (HotEmNets)*, June 2008.

[8] Gregory Hackmann, Octav Chipara, and Chenyang Lu. Robust topology control for indoor wireless sensor networks. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 57–70, New York, NY, USA, 2008. ACM.

[9] G. P. Halkes1 and K. G. Langendoen. Crankshaft: An energy-efficient mac-protocol for dense wireless sensor networks. In *EWSN 2007*, pages 228–244, January 2007.

[10] Wendi Rabiner Heinzelman, Anantha Ch, and Hari Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. pages 3005–3014, 2000.

[11] Ozlem Durmaz Incel, Stefan Dulman, and Pierre Jansen. Multi-channel support for dense wireless sensor networking. In *Smart Sensing and Context, EuroSSC 2006*, pages 1–14, October 2006.

[12] Youngmin Kim, Hyojeong Shin, and Hojung Cha. Y-mac: An energy-efficient multi-channel mac protocol for dense wireless sensor networks. In *IPSN '08: Proceedings of the 7th international conference on Information processing in sensor networks*, pages 53–63, Washington, DC, USA, 2008. IEEE Computer Society.

[13] Jeongkeun Lee, Wonho Kim, Sung-Ju Lee, Daehyung Jo, Jiho Ryu, Taekyoung Kwon, and Yanghee Choi. An experimental study on the capture effect in 802.11a networks. In *WinTECH '07: Proceedings of the second ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, pages 19–26, New York, NY, USA, 2007. ACM.

[14] Young ri Choi, Mohamed G. Gouda, Hongwei Zhang, and Anish Arora. Routing on a logical grid in sensor networks. Technical report, 2004.

[15] N. K. Santhapuri, J. Manweiler, S. Sen, R. R. Choudhury, S. Nelakuditi, and K. Munagala. Message in message mim: A case for reordering transmissions in wireless networks. In *HotNets VII*, 2008.

[16] Dongjin Son, Bhaskar Krishnamachari, and John Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 237–250, New York, NY, USA, 2006. ACM.

[17] Texas Instruments Norway AS. *CC1100 Single Chip Low Cost Low Power RF Transceiver*, June 2006. Data Sheet (Rev.1.1).

[18] T. Wark, P. Corke, P. Sikka, L. Klingbeil, Ying Guo, C. Crossman, P. Valencia, D. Swain, and G. Bishop-Hurley. Transforming agriculture through pervasive wireless sensor networks. *Pervasive Computing, IEEE*, 6(2):50–57, April-June 2007.

[19] Tim Wark, Chris Crossman, Wen Hu, Ying Guo, Philip Valencia, Pavan Sikka, Peter Corke, Caroline Lee, John Henshall, Kishore Prayaga, Julian O'Grady, Matt Reed, and Andrew Fisher. The design and evaluation of a mobile sensor/actuator network for autonomous animal control. In *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*, pages 206–215, New York, NY, USA, 2007. ACM.

[20] Joey Wilson and Neal Patwari. Radio tomographic imaging with wireless networks. *IEEE Transactions on Mobile Computing*, 9(5):621–632, 2010.

[21] Hongwei Zhang, Anish Arora, Young-ri Choi, and Mohamed G. Gouda. Reliable bursty convergecast in wireless sensor networks. In *MobiHoc '05: Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, pages 266–276, New York, NY, USA, 2005. ACM.