

TRANSMIT ONLY FOR DENSE WIRELESS NETWORKS

by

BERNHARD FIRNER

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Yanyong Zhang

and approved by

New Brunswick, New Jersey

January, 2014

ABSTRACT OF THE DISSERTATION

Transmit Only for Dense Wireless Networks

by BERNHARD FIRNER

Dissertation Director:

Yanyong Zhang

As the size and cost of embedded wireless devices decreases, researchers and technology companies find an ever-growing number of applications in health and safety, environmental monitoring, and agriculture. The goal of many of these applications is to simultaneously monitor more things at finer time or spatial granularity than was previously practicable. Progress in the size and cost of ICs has made these applications possible, but new wireless protocols must be developed to leverage this modern hardware for multi-year, dense, “always on” deployments.

These applications have requirements that are not currently addressed, most importantly energy efficiency in mobile and dense deployments, because transmitter lifetime and system maintainability may be more important than the low latency or high reliability that are paramount in current wireless systems. Current wireless protocols, such as CSMA, favor channel throughput and packet transmission success without considering the energy efficiency and utility of energy hungry feedback mechanisms such as channel sensing and packet acknowledgements. A protocol that extends the lifetime of sensors by years at the cost of occasional sensing delays (measured in seconds) from packet loss is preferable in many emerging applications. In this dissertation one such approach, called Transmit Only, is introduced. Transmit Only is shown to achieve years of lifetime on current commercially available hardware by consuming just over 20mA

per year on a chip doing both sensing and data transmission.

We also show that Transmit Only can remain energy efficient even as transmitter density increases. Although packet collisions increase with the number of transmitters in a Transmit Only system, we show that careful placement of multiple receivers can maximize the capture effect and reduce actual packet loss. For example, in a sensor network with 1000 transmitters offering 100% channel load, with just five receivers the transmitters successfully convert 80% of their radio energy into successfully transmitted bits.

This dissertation will demonstrate that the drawbacks of Transmit Only, namely multiple receivers and no packet delivery guarantees, are easily worth the benefits in efficient energy usage and lifetime, transmitter simplicity, and high throughput at very high offered loads and transmitter densities.

Preface

Portions of this dissertation are based on work previously published or submitted for publication by the author [9, 10, 11, 42]. The F-EMBED algorithm was developed in collaboration with Junichiro Fukuyama and Robert Moore.

Acknowledgements

There are many people without whom I would not have ended up where I am now. Some of those people may have even had a positive impact upon my current state. The first of those people are Yanyong Zhang, my advisor, who gave me the opportunity to do this in the first place, and Richard Howard, whose tireless patience with new grad students allowed me to grow up at my own pace. I also need to thank all of my collaborators in Owl Platform who gave my research an application; namely Ilya Chigirev, Rob Moore, Rich Martin, Eitan Fenson, and Giovanni Vannucci. I would especially like to acknowledge how much Rob likes to argue. Junichiro Fukuyama is not part of the Owl team, but he was also integral in the completion of this thesis: kudos.

My thanks go out to Janne Lindqvist, who is always pushing me to publish more and come in on Sundays.

I also need to thank Dr. Roy Yates for actually reading one of my papers and finding a math error in it.

Finally, I would like to thank Alec Moran and Anthony Uzwiak for all of the misty mornings, merriment, and misery during our runs.

Dedication

I would like to dedicate this thesis to Dr. Harry Heffes, whose sense of style and genuine interest in his course material made me think that going back to school to get a PhD might be a cool thing to do.

Table of Contents

Preface	iv
Acknowledgements	v
Dedication	vi
List of Figures	x
1. Introduction	1
2. Motivations and Related Work	4
2.1. System Overview	4
2.2. Related Work: Existing Protocols	6
2.2.1. Low Power Wireless Standards	8
2.2.2. Low Power Medium Access Control (MAC) Protocols	9
Contention-Based Protocols	9
Low Power Listening	9
Asynchronous Duty Cycles and Dense Sensor Networks	10
Minimizing Overhead	11
3. Transmit Only	12
3.1. Overview of Transmit-Only	12
3.1.1. The Capture Effect	13
3.1.2. A Case Study: Capture in the CC1100 Radio	13
Radio Chip Limitations and Redundant Receiving	15
Evaluation of Collision Behavior	16
3.2. Predicting and Exploiting the Capture Effect	19
3.2.1. Capture Threshold Assumptions	24

3.3.	Analysis of TO	25
3.4.	Analytical Comparison with CSMA	28
4.	Maximizing the Effectiveness of Transmit Only	32
4.1.	An Important Metric in TO: Transmitter Contention	33
4.2.	Design Challenges with Transmit-Only	35
4.3.	Maximizing the Capture Gains	36
4.3.1.	Quantifying the Capture Gain	36
4.3.2.	Optimal Receiver Placement	37
	F Approximation for Receiver Placement	38
4.3.3.	Fast Receiver Embedding in a Grid	43
	Adaptive Receiver Embedding in a Grid	44
4.3.4.	Simulation Results: Reducing Contention with Receiver Placement	45
	Comparison of Different Receiver Placement Strategies	45
	Contention vs. Receiver Number	47
5.	Experimental Validation	50
5.1.	Validating a General Case	50
5.1.1.	Experimental Setup	51
	Hardware Description	51
	Test System Behavior	51
5.1.2.	Test Topology	52
5.1.3.	Results	53
5.2.	Validation Capture Predictions	57
5.2.1.	Capture Model Validation	57
5.2.2.	F-EMBED Placement in a Dense, High-Load Deployment	59
5.2.3.	Validation in a Real-Life Setting	60
6.	Use Cases	62
6.1.	Example 1: Asset Tracking	62

6.1.1.	Problem Statement	63
6.1.2.	Collision Prediction	65
6.1.3.	Robust Detection of Transmitter Presence From Packet Loss . .	67
	Ambient Packet Loss	67
6.1.4.	Detection Algorithms	67
	Detection Based on Historical Maximum Miss Chain Length (Detect-	
	MaxMiss)	69
	Detection Based on RSSI (Detect-RSSI)	69
6.1.5.	Detection Results	70
6.2.	Example 2: Monitoring and Notification	72
7.	Conclusions	75
	References	77

List of Figures

3.1. Collision regions in CC11xx radios	14
3.2. Reception rates in the five collision regions	17
3.3. BER Vs differences in signal strength	18
3.4. Packet decoding success Vs interference	21
3.5. Theoretical capture gains Vs. propagation coefficient	23
3.6. Theoretical packet loss	24
3.7. ALOHA without capture	28
3.8. Hardware performance and TO	29
3.9. Packet duration and non-perfect capture	30
3.10. TO and CSMA Throughput	30
3.11. TO and CSMA Radio Efficiency	31
4.1. Contention and duty cycles	33
4.2. Capture disk example	38
4.3. F-Embed algorithm	39
4.4. 2-Approximation examples	41
4.5. Grid-Embed Algorithm	43
4.6. Different placement strategies	45
4.7. Uniform transmitter simulation	46
4.8. Circumference transmitter simulation	46
4.9. Sine wave transmitter simulation	47
4.10. Contention reduction	48
4.11. Receiver growth	49
5.1. Topology of the Single Vs Multichannel Experiment	53
5.2. Experimental and theoretical packet loss	54

5.3. Performance of all permutations, 100 transmitter trial	55
5.4. Performance variations, Single Vs Multichannel	56
5.5. Outdoor uniform random test setup	57
5.6. Outdoor uniform random results	58
5.7. Outdoor sine topology	59
5.8. Outdoor uniform random results	60
5.9. Real-Life demonstration	61
6.1. Ambient packet loss	68
6.2. Observed ambient loss	70
6.3. T_{alarm} for each transmitter	71
6.4. Screenshot of the Owl Platform current sensor status	73
6.5. Screenshot of the Owl Platform historic sensor status	74

Chapter 1

Introduction

The size, cost, and flexibility of current low power radios and microcontrollers is enabling many new computing applications where wireless communication serves as the means to collect data at a finer spatial and temporal resolution than was previously practicable. Although mainstream wireless services, like cellular and WLAN, will continue to benefit from ongoing research into capacity enhancements at the physical layer (e.g. MIMO technologies have been responsible for the growth of data rates achieved by LTE and WiMAX), outside of a few limited cases [6] these advances do not benefit systems built from these size and energy constrained radio devices. Densely deployed systems of embedded wireless devices, for example in health, habitat, or infrastructure monitoring, smart grid and smart building control systems, and general Internet of Things (IoT) applications, do not require more per-device throughput, but instead require *some* connectivity for *all* devices, with an emphasis on the maintainability and lifetime of the sensors [5, 29]. The most important requirement for the wireless protocols used in these systems will be the need for low power consumption.

As an example we might consider a system that monitors the body temperature of chickens or other livestock to detect a viral response so that the spread of disease can be immediately halted. The transmitters on the sensors must obviously be small and will thus have limited power reserves. If an animal is removed we would like to be able to simply remove the sensor and attach it to new animals as they arrive. Obviously, we want these transmitters to run as long as is possible, otherwise the cost of system maintenance on a system with hundreds or thousands of transmitters would be impractical [20].

This low power requirement must be satisfied even as the network's traffic load

grows. Although the individual devices in these systems offer a low traffic load – perhaps 1ms of sensed data every second – the high density of sensors in these systems mean that potentially large amounts of data need to be collected and relayed. Companies expect that by the year 2020 there will be 50 Billion[8] networked devices in use, and growth in monitoring and IoT applications is expected to be a large part of future wireless sensing deployments. The bandwidth required by each individual device is small, but the bandwidth required by all of the devices taken together in a densely deployed area could be quite high. Thus these system require a high total amount of bandwidth in small spatial areas. This is similar to the constraints of cellular systems, but unlike in cellular systems, these wireless devices will not be recharged on a daily basis and throughput cannot be obtained at the cost of high energy consumption.

While these networks have some more stringent requirements than traditional wireless networks, they also loosen other requirements. Traffic patterns in these systems consist of periodic data, much of which is redundant. This means that 100% packet delivery guarantees are unnecessary. Instead devices only require that packets from individual devices are received periodically, with some minimum requirement for the time between successful packets from a device. The most important metrics of wireless protocols used in these setting are thus *Joules/bit successfully transmitted* and *percent of all devices able to achieve a minimum throughput*. Although some attempts have been made to address this problem, in the form of the Bluetooth Low Energy [3] standard and IEEE 802.15.4 [15], these standards still favor packet reliability and network feedback at the expense of energy consumption.

The goal of this dissertation is to describe a practical solution to this problem, in the form of a Transmit-Only MAC protocol. Transmit Only (TO) takes advantage of these relaxed requirements to support energy efficient, reliable wireless communication from large groups of transmitters to much smaller groups of receivers. In this scheme, transmitters never listen to the channel; instead, packet reliability comes from the deployment of multiple, spatially diverse receivers that increase the capture effect.

TO shifts the entire burden of channel conflict resolution from the transmitters to the receivers. TO will allow simplification of every activity taken by the transmitter,

which directly translates to decreased energy consumption and increased lifetime. The only cost paid is a need for multiple powered receivers. The growth in receiver number to transmitter number is low, so this approach is feasible even as we go to very dense deployments with many transmitters. The capture effect occurs in most frequency modulated radio hardware so this approach is also immediately available with existing hardware.

Chapter 2 will describe the theoretical foundations of Transmit Only, including the motivating deployments and an in-depth study of the capture effect. Chapter 3 presents preliminary work demonstrating the effectiveness of TO when compared to traditional MAC protocols. With that established, we will show that, unlike protocols such as 802.11, TO is most effective when all receivers operate on a single channel rather than spreading receivers and transmitters across multiple channels. Chapter 4 presents methods we have developed to maximize the effectiveness of TO and Chapter 5 will validate our approaches with experimental results.

Chapter 2

Motivations and Related Work

This dissertation assumes a certain kind of wireless system, so it is important to discuss the kinds of applications that would rely upon such a system before discussing the theoretical performance of a Transmit Only (TO) system. The motivation behind TO, namely the control overhead of existing protocols, will also be discussed, and TO will be compared to some existing protocols.

2.1 System Overview

The assumptions made in this dissertation reflect a focus on wireless applications where small, battery powered wireless devices are deployed for extended periods of time, possibly tens of years. Transmitters may be sparsely deployed, but any transmitter must be able to operate in a densely deployed network and must accommodate growth in the size of the network over time. Sensing, inventory, and tracking are the commonly predicted applications, which lead to transmitters sending packets at frequent, regular intervals, either to transmit sensed data or as heartbeat beacons for inventory management or tracking. High sensor density and power constraints make multi-hop networking within a cluster of nodes impractical so one or more powered sinks or cluster heads act as receivers and connect each cluster with the rest of the network. There are several systems that match this description in literature or in commercial deployment.

One such application in this category is agricultural monitoring [39]. For instance, the growing condition of plants can be determined by moisture sensors in the ground, temperature sensors, or light level sensors. 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 are important for the real-time components

of these systems, especially for behavior monitoring of livestock.

One particularly interesting of a livestock monitoring application is to prevent fighting between livestock, which is costly when injuries occur. In one project with bulls, 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 [40]. 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 sends 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[1]. 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[25] and Reliable Bursty Convergecast[45] 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[41] 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, without compromising the lifetimes of the transmitters. In order to overcome this challenge, the basic strategy of a Transmit Only system is to have multiple receivers operating at the same time to increase the capture effect and reduce packet loss. Of course one could argue that the increase in hardware may be better spent in a different approach, for instance by placing the receivers on multiple channels (such as in Y-MAC [19]) rather than in a single channel in multiple locations. Thus, the main thrust of this dissertation is first to show that TO is effective at providing high bandwidth to support these systems with a very low energy requirement, and then to explore the optimal use of receivers in terms of physical placement and also number of receivers per channel.

2.2 Related Work: Existing Protocols

The most important feature for low power protocols is the amount of time the radio spends on transmitting, receiving or listening compared to the amount of time the radio can spend in low power sleeping mode. In modern microcontrollers, low power sleep modes are available that consume tens of μ amps of energy. Current coin cell batteries store about 200mAh, so a sensor that consumed 1 μ amps of energy during sleep could run for over 8000 days in sleep mode. Entering transmit or receive in the radio obviously reduces this lifetime greatly.

The energy consumed by transmission and reception is dependent upon transmission power or receiver sensitivity and the data rate used, but typically ranges from 10 to 30 mA [35, 34]. For such a sensor, if transmission consumes 20 mA then every 6 minutes of radio use will reduce the energy of 200 mA-hour coin cell battery by 1%, reducing the lifetime by about 80 days. Even if we ignore other energy costs, such as energy spent calibrating the radio, running the microcontroller, or sensing, and just want to

run the given sensor for a year from the given battery it would have slightly more than 9.5 hours of transmission time *for the entire year*, which translates to about 2 minutes per day or a radio duty cycle of about 0.15%. This tight constraint cannot be easily met by a protocol that was not designed to conserve energy. For example, CSMA is a very effective method to achieve good throughput, but was not designed with low energy consumption as its primary goal.

The throughput of p-persistent CSMA in particular has received a great deal of attention for its good performance under different traffic loads with the right set of parameters [32]. In p-persistent CSMA, when the sender is ready to send data, it senses the channel to check if the channel is idle. If the channel is idle, the sender transmits a packet with a probability p . If the sender chooses not to transmit (the probability of this event being $1 - p$), it waits until the next available time slot and transmits again with the same probability p . By choosing a very low transmission probability such as $p = 0.01$, p-persistent CSMA can achieve a very high throughput even in dense deployments, while suffering from much delayed packet deliveries as well as poor energy efficiency.

A low transmission probability means that a large amount of energy will be wasted sensing the channel since even sensed idle channel may not lead to an actual packet transmission. With $p = 0.01$ the channel must be sensed idle an average of 100 times before a transmission occurs. When packets are short compared to the sensing duration, the energy spent carrier sensing is nontrivial relative to the energy spent in packet transmission, and can greatly decrease lifetime of the devices. Non-persistent CSMA systems will spend less time carrier sensing, but the time spent backing off will result in long packet latencies; if the transmission rate falls below the duty cycle of a device then eventually its transmission queue will fill and packets will be lost. Though p-persistent and non-persistent CSMA achieve good throughput, they do not balance throughput with energy efficiency and packet latency.

2.2.1 Low Power Wireless Standards

Power consumption is a well-studied problem in sensor networks and several low power protocols to support sensing systems have become commercial standards. ZigBee [47] and the 802.15.4 [15] standard have been widely used for low-power radio devices and recently 6LoWPAN [17] and 802.15.4 has received attention for ease of interoperability between sensing systems and existing Internet technology. Additionally, a Bluetooth Low Energy (BLE) [3] standard was published to support single hop communication from low power devices in home and body networks.

In indoor environments power is often available for a few receiving nodes so the Bluetooth Low Energy (BLE) standard attempts to take advantage of this by minimizing energy consumption for uplink traffic. The shortest BLE transaction consists of the following packets: an advertisement packet from the sensor, a data channel assignment packet from the receiver, a data packet from the sensor, an ACK packet from the receiver, a transmission complete packet from the sensor, and then another ACK from the receiver, lasting for a minimum of 3 milliseconds with no data transferred. Only the beginning of the transaction occurs on a shared channel – once the receiver responds the transmitter and receiver switch to another channel for data transfer so the chance of a collision is greatly reduced.

ZigBee with 802.15.4 was designed with mesh networking in mind, assuming a node can be both transmitter and receiver. Before transmitting data, the 802.15.4 standard uses carrier sensing to avoid collisions. For communications between two battery powered sensors, they need to spend extra energy to synchronize so that their duty cycles overlap. However, in an indoor sensor network with wall-powered receivers, this synchronization would not be needed. The 802.15.4 standard has many configuration options, allowing engineers to choose parameters that best fit their situation [26]. However, the simplest mode supported by the standard is unslotted CSMA/CA which requires carrier sensing intervals with durations on the order of hundreds of μ seconds [18]. This makes the standard suitable for small sensor networks (such as body area networks [36]), although networks of larger sizes require lower power solutions than the

802.15.4 supports on its own.

2.2.2 Low Power Medium Access Control (MAC) Protocols

To support low-power sensor networks some researchers, such as Park et al.[23], have focused on tuning the parameters of these low power standards. Many researchers have proposed variations or additions to 802.15.4 while others have proposed completely different protocols for specific deployment scenarios.

Contention-Based Protocols

A great deal of the energy spent on carrier sensing is *idle listening*, where the channel is sensed but there is no traffic. The power consumption of operating a radio as a receiver is on par with the power consumption spend while transmitting at high power, so carrier sensing is just as costly as sending an actual packet. When packets are short, the relative cost of carrier sensing is relatively high. Coordination between nodes can reduce this, so a strategy of shared sleep and contention periods is taken by T-MAC and S-MAC, which reduce the overhead of inter-node communication by scheduling sleep, contention, and transmission periods for groups of nodes[43, 37].

Contention-based protocols still require large amounts of energy during contention and listening windows so the radio duty cycles are usually above several percent. The windows must be large enough to accommodate large groups of nodes. Although nodes with no traffic to send can sleep during the transmission once they discover there is no traffic bound for them, they must still turn on during a large contention window. A technique called low power listening has been proposed to break these large windows into smaller ones where only a single node needs to turn on.

Low Power Listening

WiseMAC is an example of a MAC protocol that could achieve very low duty cycles and calculations showed that it outperformed ZigBee in terms of power consumption and latency[7]. In WiseMAC, transmission from a sensor to an access point (AP) consists of just sensing the channel, transmitting a packet, and receiving an ACK packet. To

support bidirectional traffic all nodes also sample the channel periodically to check for incoming packets. If the channel is busy the node remains on until the packet is fully received, but otherwise the node can return to sleep. If the AP or another node wishes to send the listening node a packet it sends the packet during this listening window with a long preamble to account for any drift between the clocks of the transmitter and receiver. If the receiver is early it will see the preamble and remain on until the data is sent, and if it is late it will see a short preamble and the data portion of the packet. This reduces the duty cycle of nodes to 1-2%, which is a significant improvement over previous protocols.

SCP MAC[44] uses synchronized polling intervals, which removes the need for long preambles from transmitters. Contention windows are thus made relatively short, consisting of carrier sensing, a brief tone to claim the channel, a second carrier sensing, and then data transmission. Polling intervals are infrequent so energy overhead is low and the protocol adaptively increases the polling interval when bursty traffic is detected so the protocol could be set to an initial rate of 0.1% duty cycle with polling every second.

Asynchronous Duty Cycles and Dense Sensor Networks

Special care needs to be taken in convergecast scenarios, where large amounts of data are being transmitted to a single sink node. MAC protocols like Crankshaft[13] were designed to reduce energy consumption through further synchronization to avoid contention. Y-MAC[19] goes one step further and divides traffic across multiple channels through a channel hopping scheme with a single control channel for coordination and TDMA for collision avoidance with contention windows and low-power listening in time slots.

Asynchronous protocols like RI-MAC[30] were designed to remove the necessity for synchronization except when there is data to send, while still accommodating high traffic loads. However, the energy consumed searching for a *rendezvous* time for an intended receiver could be very high. RI-MAC uses a short beacon from a receiver to indicate that it is ready to receive transmissions, but the transmitter must listen to

the channel and consume energy while it waits for this beacon. Although the channel utilization goes down compared to long preambles in other protocols, the energy consumption does not.

Another similar protocol called EM-MAC[33] uses a predictive mechanism to reduce energy overhead by predicting time slots in advance. Individual transmissions between a static link are made very efficient, but the time and energy spent during the discovery phase can be quite costly. It is thus unsuitable for a high dynamic sensor network.

Minimizing Overhead

The weakness of the discussed protocols is that they all require overhead in radio use. With small packet sizes this is especially bad, making these protocols inefficient in terms of Joules spent per successfully delivered bit of data. It is possible for a system that is using protocols to aggregate data and send it in a large packet, but this decreases many of the benefits of real-time monitoring systems.

Chapter 3

Transmit Only

A theoretical model of TO will be presented in this chapter. This model will be used throughout this dissertation to predict and describe performance obtained in experimental results. The first task here is to explain the details of the capture effect, and examine the theoretical limits of capture gains. With that information we will proceed to build a model of a multi-receiver Transmit Only system. This will also make an analytical comparison with CSMA possible.

3.1 Overview of Transmit-Only

Most of the control overhead in wireless MAC protocols is spent to deal with packet collisions, either avoiding collisions or reporting data loss, because the generally accepted wisdom is that collisions will always lead to data loss and thus must be avoided. In this dissertation, we take the viewpoint that collision avoidance is not only unnecessary, but also very wasteful, posing serious limits on channel utilization and energy efficiency. Instead, we adopt a fundamentally asymmetric architecture where the transmitters are only able to transmit. Thus, the overhead associated with all interference avoidance techniques, such as carrier sensing or time/frequency/code division, has been eliminated. We call this scheme Transmit-Only, or TO in short.

The key goal of TO is to shift the entire burden of channel conflict resolution to the receivers, which allows us to simplify the transmitters in hardware and functionality and drastically reduce their energy consumption. Specifically, we propose to resolve the collisions by having multiple receivers that can collectively reduce contention between transmitters and thus eliminate data loss due to collisions. In a TO network, we deploy receivers in such a way that each transmitter is within a single hop of one or more

receivers. The receivers work on the same channel and exploit the *capture effect* to reduce the effective contention between transmitters.

3.1.1 The Capture Effect

The capture effect occurs during a collision when one signal is strong enough compared to the others that the receiver can treat the other signals as noise and filter them out. The relative signal power required to capture a packet is dependent upon the sensitivity of the radio receiver and its tolerance to noise. For instance, a study by Firner et al.[11] showed that a relative power of 6dB is required for a Chipcon radio to capture packets, while a study by Lee et al.[21] found that just 1dB was enough for packet capture when the captured packet arrived before the interfering packet with an Atheros WiFi card.

The capture effect is well known, and the best use of network receivers with the capture effect was well studied as many of our current wireless protocols were being designed. For instance, Takagi and Kleinrock[31] considered the benefits of the capture effect when multiple receivers are present in multi-hop networks. They found that the ALOHA protocol benefits more from the capture effect than 1-persistent CSMA. More importantly, they found that with “perfect capture” where during any collision the one packet with the strongest signal power will be received regardless of its relative power compared to the other packets, ALOHA could outperform CSMA, while ALOHA with 1.5dB capture (i.e., packet capture occurs when the relative power is 1.5dB or greater) performed comparably to CSMA. These comparison results of CSMA and ALOHA hint that the capture effect can effectively reduce data loss even in the presence of collisions. In this dissertation, the methods to exploit the capture effect and maximize capture gains will be fully explored.

3.1.2 A Case Study: Capture in the CC1100 Radio

To analyze the capture effect further a case study will be shown for the popular CC1100 radio. The purpose of the case study is to give further details to a reader who is not well acquainted with the capture effect and to demonstrate the capabilities of modern, low-power radios.

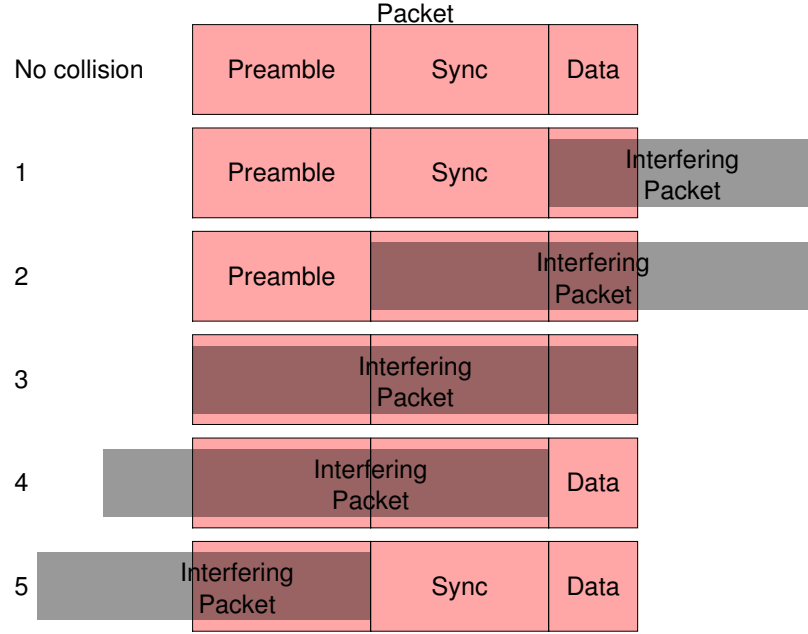


Figure 3.1: The five possible regions of packet interference with the three part packet used in the CC11xx line of radios. Packet regions with interference will have increased error rates depending upon the relative signal strengths of the two packets.

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 [35]. 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 3.1. 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 3.1.

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	$100 \times \frac{\delta_{sync}}{2\delta} = 20$
5	$100 \times \frac{\delta_{preamble}}{2\delta} = 20$

Table 3.1: The percent of collisions that fall into each collision region for a CC1100 radio transmitting a short, $288\mu\text{second}$ packet.

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 [27].

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.

Evaluation of 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[35]. 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 3.2 shows packet reception rates under different collision conditions. Figure 3.2(a) shows the packet reception rates when interference begins after the packet begins and Figure 3.2(b) shows reception rates when interference was present when the packet began transmitting.

In Figure 3.2(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 3.3, 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 because the radio locks on to the

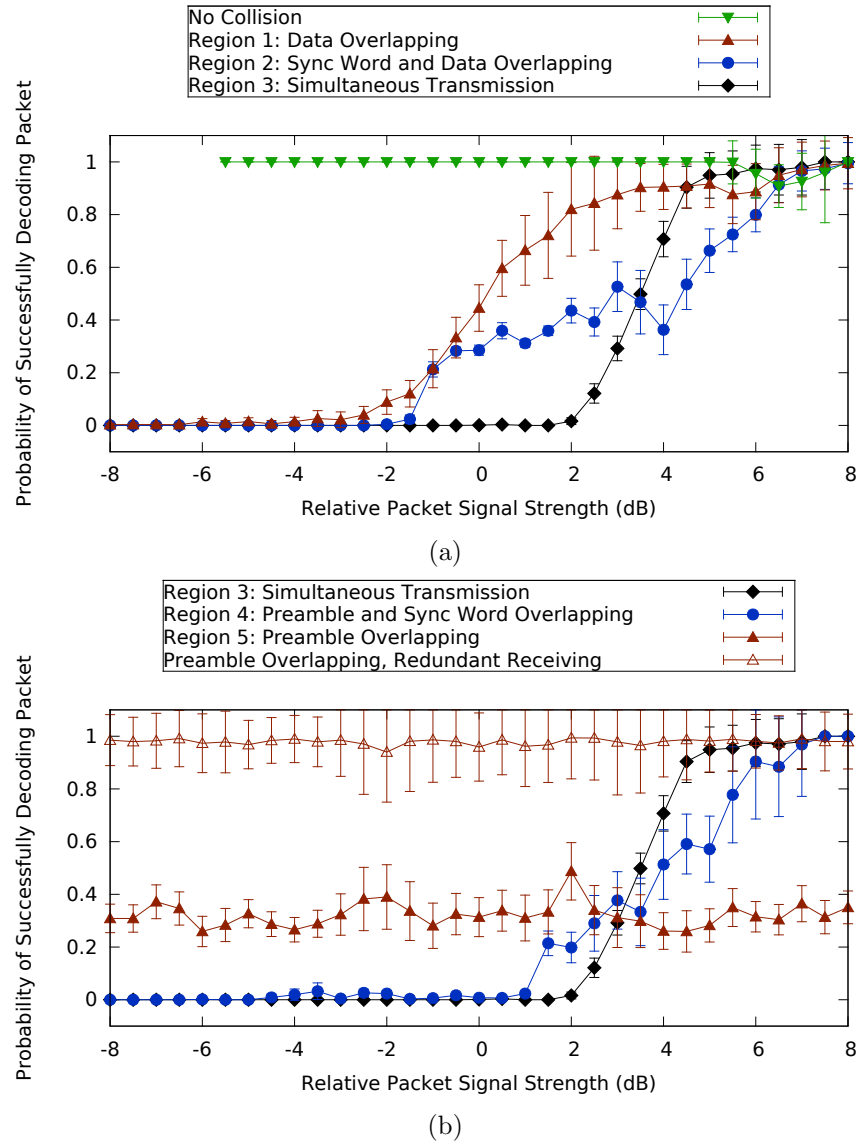


Figure 3.2: The packet reception rates for packets in collision regions 1-3 (a) and packets in collision regions 4-5 (b). The error bars are the estimated 95% confidence intervals.




Figure 3.3: 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 with several hundred collisions per point. A line generated from regression analysis is also shown, showing a close match between bit error and the complementary error function (erfc). The datasheet from the radio itself indicates that an error rate of as high as 1% should be expected even under ideal conditions, so once the bit error rate falls very low it is difficult to distinguish errors caused by interference from errors in either the transmitter or receiver hardware[35].

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 begin transmission simultaneously, as in collision region 3, successful packet decoding begins to occur when the packet is at about the same signal strength as the packet it collides with. As observed in [28] 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 3.2(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 *μseconds* 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. 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.

3.2 Predicting and Exploiting the Capture Effect

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[9], 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 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

$$P_{2-waycollision} = P_{capture} \frac{2 \times \delta}{\tau} \quad (3.1)$$

$$\begin{aligned} P(\text{collision} | N \text{ transmitters}) \\ &= 1 - P(\text{no collision})^{N-1} \\ &= 1 - \left(1 - P_{capture} \frac{2 \times \delta}{\tau}\right)^{N-1} \end{aligned} \quad (3.2)$$

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 3.1 the probability of a collision falling into that region is shown in Table 3.1. From the values in this table and the results from Figure 3.2 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 3.4.

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




Figure 3.4: 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.

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 3.3 and can be used to adjust the capture rate in Figure 3.4 for different packet packets.

We can calculate $P_{capture}$ from the curve in Figure 3.4 once we find a function to predict probabilities of different relative signal strengths between two signals. Power loss during signal propagation follows 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 relative dB amount, Δ , to a relative distance for two transmitters at distances l_1 and l_2 and call the conversion factor K .

$$\begin{aligned} \frac{1}{l_1^\alpha} &\geq \frac{10^{\Delta/10}}{l_2^\alpha} \\ l_1 &\leq l_2 10^{-\Delta/10\alpha} \\ l_1 &\leq l_2 K, \text{ where } K = 10^{-\Delta/10\alpha} \end{aligned} \tag{3.3}$$

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 K .

$$\begin{aligned} &\int_{a/K}^b \frac{1}{b-a} \int_a^{Kx} \frac{1}{b-a} dy dx \\ &= \frac{K}{(b-a)^2} \left(b - \frac{a}{K}\right)^2 \\ &= \frac{K}{2} \text{ if } a = 0. \end{aligned} \tag{3.4}$$




Figure 3.5: 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.

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 3.3 and 3.4. This relationship will be approximately correct for any receiver being used with uniformly deployed nodes. We can now predict the rate of the capture effect for our packets with attenuation $\alpha = 2.69$. This is shown in Figure 3.5.

$P_{capture}$ from Equation 3.2 can now be determined from Figure 3.4 and Equation 3.4. Figure 3.4 shows the cumulative probabilities of packet success, which are broken down into probabilities for each relative signal value. These are multiplied by the probability of that signal value occurring, obtained from Equation 3.4 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 3.2 and inserting $P_{capture}$ we can create the theoretical packet loss curves shown in Figure 3.6. 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

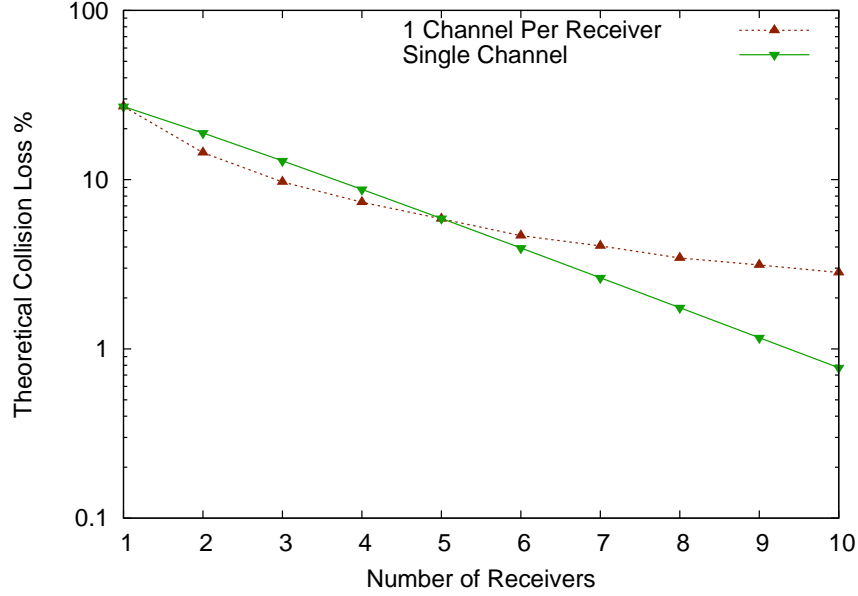


Figure 3.6: Theoretical packet losses with $\alpha = 2.69$, one $300 \mu second$ transmission per second, and 1500 transmitters, using Equation 3.2 and $P_{capture}$ from Figure 3.4 and Equation 3.3. Once of the number of receivers passes a threshold a single channel is better than using multiple channels.

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. Experimental validation of this theoretical model appears in Sections 5.1.1–5.1.3.

3.2.1 Capture Threshold Assumptions

In our analysis we do not sum the magnitudes or consider phase differences of multiple signals overlapping in time, and instead assume the capture effect occurs when the strongest signal is a given threshold above any of the other overlapping signals. This assumes that, for any n -way collision, interference remains Gaussian in nature so that the radio achieves its best SNR performance. If interference is not Gaussian, then the radio would need a higher SNR ratio to successfully ignore the interference and the

capture threshold would need to vary with the number of interfering signals.

With a single interfering signal, the data whitening of the signal makes it similar to Gaussian and the capture effect will perform at a good SNR. When there are many independent interfering signals, it is well known that by the law of large numbers the superposition of the signals can be modelled as a Gaussian random process. This is again the case where radio capture performance is good.

While it is possible that a small number of interfering signals will create non-Gaussian interference, with very small or large values of n in n -way collisions the interference will be ideal. There are cases when non-optimal SNR ratios are required for capture, but we do not expect these situations to occur often enough to make our model inaccurate.

3.3 Analysis of TO

For the sake of simplicity, we will assume that the signal power of a transmitter at each receiver is uncorrelated, meaning that the probability of any individual packet being captured during a collision is the same at each receiver. This allows us to show the possible benefits of having multiple receivers under a theoretically ideal situation and will motivate further chapters in this dissertation. We will revisit this assumption again in Section 4.3 and will discuss how to maximize potential capture gains with multiple receivers in real-world situations.

In TO the probability of two transmitters having a packet collision depends upon the duty cycle of transmitters. We will consider a simple case where traffic is periodic with interval τ and durations δ (a duty cycle of δ/τ). TO is unslotted, so a collision occurs when any overlap occurs, or $2\delta/\tau$. The probability of two transmitters colliding is

$$P_{2\text{-way-collision}} = \frac{2\delta}{\tau}. \quad (3.5)$$

With N transmitters, a transmitter's packet is received if no collisions occur, the

probability of which is

$$P_{succ} = (1.0 - \frac{2\delta}{\tau})^{N-1}. \quad (3.6)$$

In the presence of contention, a transmitter may collide with any of the $N - 1$ other transmitters, but due to the capture effect, its packet can be correctly received if it collides with transmitters that have lower signal strengths at a receiver. Let us assume that, even with the capture effect at the receivers, a transmitter's packet cannot be received when it collides with any of the C transmitters (a subset of $N - 1$ other transmitters). We say that this transmitter's *contention level* is C , and the packet success probability is thus:

$$P_{succ} = (1.0 - \frac{2\delta}{\tau})^C. \quad (3.7)$$

This equation gives a good intuition for packet success, but it does not include details of capture when there are multiple receivers or when it is difficult to determine the contention level of each transmitter. In this case we must model the system with a more exact equation that takes into account the probability of a transmitter's packet being captured in a collision involving any number of other transmitters, from 1 to $N - 1$. The probability of packet loss from a collision is simply a binomial random variable with the addition of the capture probability with each collision magnitude.

$$P_{loss} = \sum_{i=1}^{N-1} \left(\frac{2\delta}{\tau}\right)^i \left(1 - \frac{2\delta}{\tau}\right)^{N-i-1} \binom{N-1}{i} (1 - P_{capture}). \quad (3.8)$$

In the case of *perfect capture*, we assume that during any packet collision with any number of packets a receiver will always correctly decode one of the packets. In this case the probability of any transmitter involved in an n -way collision having its packet captured is simply $1/n$. Given n transmitters and r receivers the probability of a particular transmitter not having its packet captured is

$$1 - P_{perfect-capture}(n, r) = (1 - 1/n)^r. \quad (3.9)$$

In a non-perfect capture case we need to compute both the probability that the transmitter we are checking is the strongest transmitter and also that the strongest transmitter is actually captured over all other transmitters. From Equation 3.3 we saw that the threshold for capture could be modeled as K , which models both the capture threshold of the hardware being used, Δ , and the propagation coefficient of an environment, α , assuming a propagation model of $1/r^\alpha$. Further, Equation 3.4 showed that when transmitters are scattered in a uniformly random pattern from the receiver's location the likelihood of a transmitter having its packet captured is $K/2$, and thus the likelihood of any capture taking place is K . For simplicity we will assume that the probability of capture at different receivers is independent, although manually chosen receiver locations will usually have some relationship. The probability that the one of the packets in a collision is captured over all of the others is the probability of $n - 1$ captures taking place, or K^{n-1} . Combining this with the probability of a particular transmitter being that strongest signal and considering the possibility of capture at any of r receivers when n transmitters collide we find

$$\begin{aligned} 1 - P_{\text{capture}}(n, r) &= (1 - P_{\text{strongest}} P_{\text{strongest captures}})^r \\ 1 - P_{\text{capture}}(n, r) &= \left(1 - \frac{K^{n-1}}{n}\right)^r. \end{aligned} \quad (3.10)$$

Using Equations 3.8 and 3.9 we can calculate the theoretical gains of the perfect capture effect with multiple receivers, as shown in Figure 3.7. In Figure 3.7, we use *offered packets per packet interval* (i.e. the expected number of packets being transmitted at a given time) as the offered load. From this result we can see that the capture effect could greatly increase throughput in TO compared to a single receiver ALOHA system without capture.

Using Equations 3.10 and 3.3 we can also estimate performance under non-perfect capture conditions, for instance by using the 6dB capture threshold of a CC1100 radio or the 1dB [21] threshold of some Atheros WiFi cards. A very important difference between perfect and non-perfect capture is that performance degrades as the offered load increases in the non-perfect case. The rate of this change in performance is actually




Figure 3.7: ALOHA without capture has a low packet throughput, but when we consider the capture gains from multiple receivers very high packet throughputs are achievable. Throughput clears 100% because there are multiple receivers, each with its own chance to successfully receive a packet, so multiple packets may be simultaneously received successfully.

a factor of the packet duration and the interval between packet (the duty cycle), as can be seen by comparing Figures 3.8 and 3.9.

3.4 Analytical Comparison with CSMA

P-Persistent CSMA can be a very energy-efficient protocol while also achieving very good throughput. Although the value of “p”, the probability of transmitting when the channel is sensed empty, needs to be tuned for best performance, p-persistent CSMA is a good standard for us to compare with CSMA. There are of course many different low power protocols, but many of these were also built for special use-cases. A direct comparison between CSMA and TO is straightforward, and since so many other protocols have been compared with CSMA in the past and will be compared with it in the future, this comparison between CSMA and TO will allow future readers to draw comparisons between TO and any new protocols that are compared with CSMA.

Since we wish to compare TO with CSMA protocols we will make use of the radio efficiency analysis performed by Ramachandran and Roy[24] to model a p-persistent CSMA system. We first show the throughput values for different p-persistent CSMA




Figure 3.8: The hardware’s capture threshold has a large effect upon performance as shown in a comparison of the expected throughput of TO with different numbers of CC1100 radios Atheros WiFi cards. $\alpha = 3$ is assumed to find the capture rate from Equation 3.3.

and TO protocols in Figure 3.10. The results clearly show that when we have 3 or more receivers, the throughput of TO is much higher than the CSMA protocols. TO throughput improves with the number of receivers because with more receivers, more packets in a collision can be captured.

We show the corresponding radio efficiency results in Figure 3.11. Radio efficiency is the ratio between the time spent in successfully transmitting data packets and the total radio time. The energy requirements of the radio usually dominate the energy consumption of wireless devices, so this may be the single most important factor for device lifetime. Again, in TO systems having 3 or more receivers can greatly improve the radio efficiency, and thus device lifetimes.




Figure 3.9: In the non-perfect capture case the ratio of packet duration to packet interval, the duty cycle, affects when the system reaches peak performance. Longer packets cause the system to have smaller “good” ranges.




Figure 3.10: A comparison of throughput for TO and p-persistent CSMA with different p-values.




Figure 3.11: Although 0.01-persistent CSMA may give good throughput, it is not energy efficient. TO is able to achieve both good throughput and good energy efficiency.

Chapter 4

Maximizing the Effectiveness of Transmit Only

The transmit-only communication model is appropriate for a specific class of wireless networks where wireless embedded devices report self-contained units of data to a small number of receivers. The reporting devices are unconcerned with receiving information from another device, and the number of reporting devices is significantly greater than the number of receiving devices.

Instead of having complex transmitters that both transmit and receive, TO assumes a model where one or more dedicated receivers will coordinate data reception to reduce lost packets and maximize channel efficiency by providing redundancy against fading and packet collisions. A TO network adopts a single-hop architecture: every transmitter is within a single hop of one or more receivers. Receivers need to be strategically placed, maximizing the chances that packets suffering from a collision will be captured by at least one receiver. This means that the transmitter's radio hardware and software can be greatly simplified, reducing both the cost of the device and its development. Although the network of receivers is more complex than a single receiver, the ratio of receivers to transmitters is low so the overall improvements are substantial.

As an example, let us imagine a stadium or arena where each of the chairs reports whether it is occupied or empty (either via a weight sensor or its position). Every 5 seconds, the chair's sensor will transmit its unique identifier as well as its state. If we take New York's Madison Square Garden as an example, (listed as 20,976ft² with a capacity of 20,000 people), then this is a radio density of 1.0488 per square foot. Using retail WiFi base stations, we could hope for 50 devices per base station, requiring 400 access points and many more non-overlapping channels than the 3 available in the U.S. Clearly this type of protocol is impractical for such an application, while TO's simplicity




Figure 4.1: Packet loss at different contention levels for different radio duty cycles at each transmitter. The contention level seen by one device is the number of transmitters which prevent packets from this device from being captured by any receiver in the event of a collision.

can support these high device densities.

4.1 An Important Metric in TO: Transmitter Contention

The end goal of TO is to increase throughput, spectrum efficiency and energy efficiency in a dense wireless network. There are no control packets to change transmitter behavior or avoid collisions, so throughput must be increased by reducing packet loss due to collisions. Packet loss rate is not a clean metric to measure TO's effectiveness though, since packet loss rate is impacted by many other factors such as the offered load on each transmitter and the number of transmitters in the system. Instead of packet loss rate, we will use *transmitter contention* as the metric to describe the TO system since it is not affected by traffic load, but only by the network topology. It is especially advantageous for systems with a high density of embedded devices that frequently report their status. In this environment, TO provides more throughput than most of the conventional protocols.

In a TO system, the contention seen by a transmitter is the number of other transmitters which, should a collision occur, would prevent its packets from being captured by any receiver. Given two transmitters A and B, they are in contention with each other for the radio channel at a receiver if simultaneous transmission by A and B results in packet loss for both A and B. If A's signal is stronger than B's, then A's packet is captured by a receiver and B no longer contends with A (although A still contends with B for the channel).

In a dense network with N devices that can hear each other, when we say device A has a contention of n , we mean that among $N - 1$ possible collisions A may have (we only consider 2-way collisions here), for $N - n$ times we can successfully capture A 's packet. That is, A 's packet will be lost for the rest of n times. For instance, assume there were 100 transmitters (including transmitter A) in a TO network with 1 receiver and no capture effect. A may at most collide with all 99 other transmitters, and none of its packets can be captured. This indicates A 's contention is 99. If we added several receivers and each of the receivers had the ability to capture the strongest packets during a collision, then A 's packet can be captured 50 times out of the 99 possible collisions. Then A 's contention is cut down to 49, halved from the original setting.

A system's average packet loss rate is impacted by the contention level, as well as the traffic. To provide a visual explanation for these two metrics, we have plotted packet loss rate with different contention for different offered loads in Figure 4.1. Figure 4.1 shows that to achieve a low packet loss rate, a network needs to minimize both contention and offered load. This has very important implications to the design of TO networks. On one hand, the objective of TO is to use receivers to cut down the contention. On the other hand, we realize that TO is best suited to networks with relatively small packet sizes. If a network has ten devices transmitting 100ms packets every second then the packet loss rate would be over 90% – even if we could reduce the contention level by half packet loss would still be over 50%. On the other hand, if the network has 1000 devices transmitting 100us packets every second, then the packet loss rate is less than 20%. This packet loss rate can be further greatly reduced by the capture effect when we have multiple receivers.

4.2 Design Challenges with Transmit-Only

When designing a TO network, we need to pay attention to several detailed issues, both at the transmitters and at the receivers. Below we discuss these issues and our solutions.

First, since we rely on the capture effect to reduce packet losses from collisions, repeated collisions among the same set of transmissions need to be avoided since the capture effect is inherently unfair. If a certain receiver deployment does not have any receiver near to a particular transmitter, then packets from this transmitter are unlikely to be captured and will be lost if it repeatedly collides with the same stronger transmitters. Other systems would avoid this situation with coordination, but without listening, these transmitters cannot coordinate to achieve de-synchronization.

The solution to this problem is to ensure that the traffic pattern in TO networks is stochastic, making repeated collisions unlikely. Even if a deployment calls for a regular duty cycle each transmitter can add a random component to the timing of its transmission, or could use an inherently unpredictable local clock to schedule its duty cycle.

The second challenge faced by a TO network is its manageability. A network usually employs downlink traffic (from the receiver to the transmitter) to manage or re-configure the network. The parameter most likely to be reconfigured during the operation of a data reporting network is the reporting frequency. In a TO network, we solve this issue simply by over-provisioning. When we design a TO network, we decide on the maximum reporting frequency for each transmitter, and have the transmitter transmit at that frequency all the time. During the periods when the maximum frequency is not needed, the receivers will perform filtering. In this way, we can shift the burden of reconfiguring the network to the receiver. The cost of over-provisioning will be made up for by the efficiency of the system.

At the receiver side, several issues must be addressed. First, they require access to much more energy than transmitters because their radios need to be kept on most of the time. Second, they need to be connected to a powerful back-haul network.

Third, they must be able to cooperate, either through a centralized point or between themselves. Fourth, and the focus of this work, is that their placement can have large impacts on the quality of the reported data. A much reduced contention is the most important contribution of having multiple receivers in a TO system, so much of the rest of this chapter will be devoted to receiver placement strategies that can fully harvest the capture effect's benefits. We will show that, with good receiver placement, a small number of receivers are sufficient for low contention and low packet loss in a TO system.

4.3 Maximizing the Capture Gains

The capture effect's presence in wireless radio systems increases throughput by reducing contention for transmitters. If a given system has a desired packet loss rate then a TO deployment will use receivers to reduce the contention rate to achieve the target packet loss rate.

Given a certain number of receivers, the effective contention level for a transmitter is mainly determined by how the receivers are placed. Receiver placement has an important impact on how many packets can be captured during a collision. In this section, we will study several receiver placement strategies.

4.3.1 Quantifying the Capture Gain

Previously we assumed that receivers could achieve perfect capture, where a receiver can always successfully capture and decode the single strongest signal during any collision. Current radios do not achieve perfect capture and the ability of a radio to successfully decode the strongest signal during a collision is a function of the relative power of the strongest signal to the cumulative power of the interfering signals. There is not truly a hard threshold for successful packet decoding because the relative power of that signal against noise determines its bit error rate rather than directly determining the packet error rate. This behavior has been measured in different radio systems and yields similar curves of bit error rate to relative signal strength[11, 21].

To simplify modeling we will assume that packets will be successful once a threshold

is reached in the relative power of a transmitter to any colliding packets. Thus we can describe the probability of capture as a function of the threshold of a transmitter's signal power to the signal power of interfering transmitters, which we will call Δ , and the signal propagation coefficient of an environment, α .

The signal power observed by a receiver is commonly approximated as P/d^α , where P is the power of the signal at the transmitter, d is the distance between the transmitter and the receiver, and α is the attenuation factor of the radio signal. In free space $\alpha = 2$, but in most environments α ranges between two and four. If d_1 is the distance from the transmitter whose packet we want to capture to the receiver and d_2 is the distance from the interfering transmitter to the receiver, then the threshold for capture in terms of relative distance of the two transmitters to the receiver is described as the variable β in the inequality

$$\begin{aligned} 10\log_{10}\left(\frac{P/d_1^\alpha}{P/d_2^\alpha}\right) &\geq \Delta dB \\ \alpha 10\log_{10}\left(\frac{d_2}{d_1}\right) &\geq \Delta dB \\ (d_2/d_1) &\geq 10^{\Delta dB/10\alpha} \\ \beta &\geq 10^{\Delta dB/10\alpha}. \end{aligned} \tag{4.1}$$

With knowledge of β we can choose the best receiver placement to achieve high packet capture rates and achieve good fairness in the system.

4.3.2 Optimal Receiver Placement

We now formally define the optimal receiver placement problem. Consider two transmitters located at $t_1, t_2 \in \mathbb{R}^2$, and a receiver located at $r \in \mathbb{R}^2$. For the sake of simplicity, we will use t_1, t_2 and r as both their locations and identities. In case of a packet collision between t_1 and t_2 , the signal from t_1 can be captured by r if and only if

$$\|r - t_1\| \leq \beta \|r - t_2\|, \tag{4.2}$$

where $\|\cdot\|$ is the Euclidian norm of a vector in \mathbb{R}^2 , and $\beta \in (0, 1)$ describes the relative distance difference for the capture effect to take place. In this case, we say that the




Figure 4.2: The capture disk of transmitter pair (t_1, t_2)

ordered transmitter pair (t_1, t_2) is successfully *captured* by r . Our goal is to *find m receiver locations that maximize the captured transmitter pairs and thus minimizes the average contention*.

If (4.2) holds true with some receiver r , we say that the ordered transmitter pair (t_1, t_2) is successfully *captured*. For the sake of simplicity, we focus on 2-way collisions in this section though our algorithm can be easily extended to solve collisions that involve 3 or more parties. Formally, our *receiver embedding problem* is defined as follows:

Given: locations $t_1, t_2, \dots, t_n \in \mathbb{R}^2$ of n transmitters and the number of receivers, m .

Find: m receiver locations $r_1, r_2, \dots, r_m \in \mathbb{R}^2$ such that the number of captured transmitter pairs is maximum.

F Approximation for Receiver Placement

To solve the optimal problem, we need to first redefine the original optimization problem as a discrete problem such that we only need to examine a finite number of points in \mathbb{R}^2 . Let us start this process by finding a necessary and sufficient condition to (4.2). It is equivalent to

$$\begin{aligned} \|r - t_1\|^2 &\leq \beta^2 \|r - t_2\|^2 \\ \Leftrightarrow (1 - \beta^2) \|r\|^2 - 2r \cdot (t_1 - \beta^2 t_2) &\leq -\|t_1\|^2 + \beta^2 \|t_2\|^2, \end{aligned}$$

where \cdot takes the inner product of two vectors. It is then further transformed into

$$\left\| r - \frac{t_1 - \beta^2 t_2}{1 - \beta^2} \right\| \leq \frac{\beta}{1 - \beta^2} \|t_1 - t_2\|. \quad (4.3)$$

Algorithm F-EMBED**Inputs:**

1. n transmitter locations $t_1, t_2, \dots, t_n \in \mathbb{R}^2$.
2. Positive integer m .

Output: m receiver locations $r_1, r_2, \dots, r_m \in \mathbb{R}^2$.**begin**

1. Compute the center and radius of the capture disk of every ordered transmission pair $(t_1, t_2), (t_1, t_3), \dots, (t_n, t_{n-2}), (t_n, t_{n-1})$;
2. Compute all the solution points, *i.e.*, the centers and intersections of capture disks;
3. Construct a bipartite graph $G = (S, T, E)$ as follows;
 - 3-1. $S = \{s_1, s_2, \dots, s_l\}$ is the set of solution points;
 - 3-2. $T = ((t_1, t_2), \dots, (t_n, t_{n-1}))$ is the set of ordered transmitter pairs;
 - 3-3. E is the edge set such that $(s_j, (t_i, t_{i'})) \in E \Leftrightarrow$ the capture disk of $(t_i, t_{i'})$ contains the solution point s_j ;
4. **for** $k = 1$ to m **do**
 - 4-1. Find a solution point $s_j \in S$ connected to a maximum number of transmitter pairs $(t_i, t_{i'})$;
 - 4-2. Let the k^{th} receiver location r_k be this s_j ;
 - 4-3. Delete s_j and $(t_i, t_{i'})$ connected to s_j from G ;
5. **end for**

endFigure 4.3: Our F-EMBED Algorithm to Find m Receiver Locations

The region of possible receiver location r such that (t_1, t_2) is captured by r forms a disk whose center is $\frac{t_1 - \beta^2 t_2}{1 - \beta^2}$ and radius $\frac{\beta}{1 - \beta^2} \|t_1 - t_2\|$. We call it *the capture disk of (t_1, t_2)* illustrated in Fig. 4.2. This means that (t_1, t_2) is captured if and only if its capture disk contains a receiver.

We can define the discrete version of the receiver embedding problem by the following two steps. First, given n transmitters, we compute the capture disks of all the ordered transmitter pairs $(t_i, t_{i'})$ ($i \neq i'$). Second, we compute the corresponding *solution points* which include the center of each capture disk and the intersection point between the boundary circles of any two intersecting capture disks. After calculating the solution points, we can then just limit the possible receiver locations to these solution points.

The following lemma shows the above discrete problem is equivalent to the original receiver embedding problem.

Lemma 4.3.1. *For any given transmitters t_1, t_2, \dots, t_n for the receiver embedding problem, there exist optimal m receiver locations r_1, r_2, \dots, r_m such that every r_j ($1 \leq j \leq m$) is a solution point.*

Proof. Suppose that an optimal receiver location r_j captures p transmitter pairs. It means that r_j belongs to the intersection I of p capture disks. Any such region I contains an intersection between the boundary circles of two capture disks, or the center of a capture disk. Let $r'_j \in \mathbb{R}^2$ be such a solution point.

Given any optimal set $\{r_1, r_2, \dots, r_j\}$ of receiver locations, construct another by replacing each r_j by the above solution point r'_j . We have a solution consisting of only solution points. This proves the lemma. \square

Thus, in the discrete version, we just need to examine the solution points for receiver locations to find the optimal solution. Assuming finite digits are used to represent a vector in \mathbb{R}^2 , this discrete version is similar to classical **NP**-complete problems such as subset sum or vertex cover. To expand, in vertex cover we select a minimal subset of vertices such that every edge in the graph is incident to at least one of the selected vertices. In the receiver embedding problem we have a set of solution points to which one or more capture disks are incident and we want to select the minimal subset of solution points such that every capture disk is incident to at least one of the selected solution points. We thus believe that both the original and discrete versions of the receiver embedding problem are **NP**-hard. Therefore, an approximation algorithm of constant factor to the optimal solution will be a good alternative.

We will present a 2-approximation algorithm for the receiver embedding problem, *i.e.*, it always returns a solution that captures at least half as many collisions as the optimum solution. Our algorithm, called F-EMBED, has four steps:

- In step 1, we compute the capture disks for all the ordered transmitter pairs.
- In step 2, we compute the solution points.
- In step 3, we construct a bipartite graph G such that solution point s_j is connected to transmitter pair $(t_i, t_{i'})$ if and only if s_j is contained in the capture disk of $(t_i, t_{i'})$.
- In step 4, we go through m iterations and in each iteration we pick the solution point that has the maximum number of edges in the graph and remove the solution point as well as its edges.

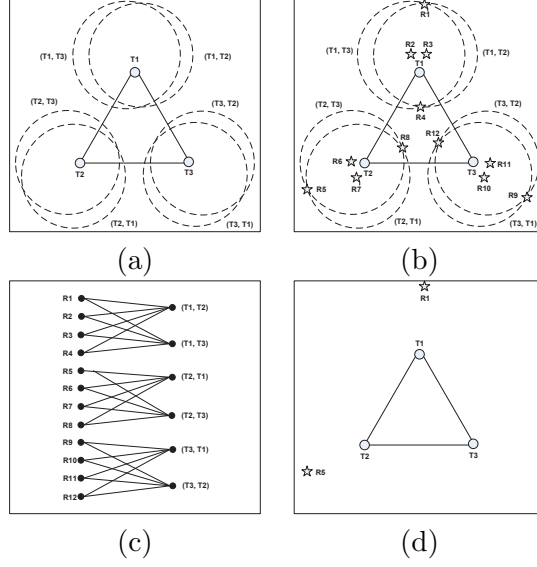


Figure 4.4: Walking through the proposed receiver embedding solution with a 3-transmitter 2-receiver example scenario. We calculate the capture disks for all six ordered transmitter pair in (a), and generate their solution points in (b). We show the bipartite graph between the solution points and the ordered transmission pairs in (c). (d) shows the resulting receiver locations.

A walk through of the algorithm using a 3-transmitter, 2-receiver example is illustrated in Figure 4.4.

The quality of the solution that this algorithm achieves is an important result though, so we confirm in the following theorem that its approximation factor is 2.

Theorem 4.3.2. *F-EMBED is a 2-approximation algorithm for the receiver embedding problem.*

Proof. Given transmitter locations t_1, t_2, \dots, t_n and an integer m for the embedding problem, let o_S and c_S denote the optimal and F-EMBED receiver locations, and o_T and c_T denote the ordered transmitter pairs captured by o_S and c_S respectively.

By Lemma 4.3.1, we assume that the optimal solution o_S consists of solution points only. This means that

$$o_S, c_S \subseteq S = \{s_1, s_2, \dots, s_l\}$$

$$\text{and } o_T, c_T \subseteq T = \{(t_1, t_2), (t_1, t_3), \dots, (t_n, t_{n-1})\}.$$

Here S and T are the vertex sets of the bipartite graph $G = (S, T, E)$ constructed by

Step 3 of F-EMBED.

We prove by induction on m that

$$|c_T| \geq \frac{1}{2} |o_T|. \quad (4.4)$$

The base case $m = 1$ is clearly true. Let us assume true for $m - 1$ and prove true for m .

Without loss of generality, assume that the solution point chosen by Step 1 is $s_l \in S$. Let $T_l \subseteq T$ be the (set of) ordered transmitter pairs captured by s_l . Step 3 removes them from G , so the remaining solution points are

$$S' = \{s_1, s_2, \dots, s_{l-1}\} \subset S,$$

and transmitter pairs

$$T' = T - T_l.$$

If $s_l \in o_S$, then (4.4) holds true. By induction hypothesis, F-EMBED returns c_S such that $|c_T \cap T'| \geq \frac{1}{2} |o_T \cap T'|$. Thus

$$|c_T| = |c_T \cap T'| + |T_l| > \frac{1}{2} (|o_T \cap T'| + |T_l|) = \frac{1}{2} |o_T|,$$

proving (4.4).

So we assume $s_l \notin o_S$. The optimal solution o_S chooses m solution points in $S' = S - \{s_l\}$. Let s_{l-1} denote an element in $o_S \subset S'$, and T_{l-1} denote the transmitter pairs captured by s_{l-1} . We have

$$\begin{aligned} |c_T| &= |c_T \cap T'| + |T_l| \geq \frac{1}{2} (|o_T \cap T'| - |T_{l-1}|) + |T_l| \\ &\geq \frac{1}{2} (|o_T| - |T_l| - |T_{l-1}|) + |T_l| \\ &\geq \frac{1}{2} (|o_T| - 2|T_l|) + |T_l| \geq \frac{1}{2} |o_T|. \end{aligned}$$

This completes the proof of the induction step. The theorem follows. \square

Algorithm GRID-EMBED**Inputs:**

1. n transmitter locations $t_1, t_2, \dots, t_n \in \mathbb{R}^2$.
2. Positive integer m .
3. Finite set S of points in \mathbb{R}^2 .
(S is typically a given set of grid points.)

Output: m receiver locations $r_1, r_2, \dots, r_m \in S$.**begin**

1. Compute the center and radius of the capture disk of every ordered transmission pair $(t_1, t_2), (t_1, t_3), \dots, (t_n, t_{n-2}), (t_n, t_{n-1})$;
 2. Construct a bipartite graph $G = (S, T, E)$ as follows;
 - 2-1. S is the given set of grid points;
 - 2-2. $T = ((t_1, t_2), \dots, (t_n, t_{n-1}))$ is the set of ordered transmitter pairs;
 - 2-3. E is the edge set such that $(s_j, (t_i, t_{i'})) \in E \Leftrightarrow$ the capture disk of $(t_i, t_{i'})$ contains the grid point s_j ;
 3. **for** $k = 1$ to m **do**
 - 3-1. Find a grid point $s_j \in S$ connected to a maximum number of transmitter pairs $(t_i, t_{i'})$;
 - 3-2. Let the k^{th} receiver location r_k be this s_j ;
 - 3-3. Delete s_j and $(t_i, t_{i'})$ connected to s_j from G ;
 4. **end for**
- end**

Figure 4.5: Grid-Embed Algorithm

4.3.3 Fast Receiver Embedding in a Grid

The F-EMBED algorithm slows down rapidly when the number of transmitters increases because the number of solution points grows rapidly with the transmitter count. Given n transmitters, there are $O(n^2)$ ordered transmitter pairs and capture disks. In the worst case, every single capture disk intersects with every other and there will be $O(n^4)$ solution points. We must check every solution point to see if it is contained in one or more capture disks, so with $O(n^2)$ capture disks the worst case running time of this algorithm is $O(n^6)$. In the average case, every capture disk does not necessarily intersect with every other, so the number of solution points will be lower than that in the worst case. Even so, we will have a vast number of solution points to consider in the F-EMBED algorithm.

To come up with a faster alternative, we approximate the receiver embedding problem in a grid: instead of searching all the solution points, we examine a list of preset grid points to see whether they are contained in one or more capture disks. In this way, we can avoid the calculation of solution points which incurs great computation complexities. Consider the following variant of the embedding problem.

Receiver Embedding Problem in Grid

Given: locations $t_1, t_2, \dots, t_n \in \mathbb{R}^2$ of n transmitters, the number m of receivers, and a finite set S of grid points on \mathbb{R}^2 .

Find: m receiver locations $r_1, r_2, \dots, r_m \in S$ such that the number of captured transmitter pairs is maximum.

The grid embedding problem differs from the original embedding problem in that we restrict receiver locations to be in S , which is a given set of grid points. The size of S is adjustable to achieve both good capture performance and fast computation of receiver locations. We have the algorithm GRID-EMBED in Fig. 4.5 to find a solution set to this grid-based problem. It is significantly faster than F-EMBED such that it is useful for large numbers of transmitters.

We have proved that the GRID-EMBED algorithm is a 2-approximation algorithm to the receiver grid embedding problem, and that when we have enough grid points (say $\geq 400,000$) the algorithm is also a 2-approximation algorithm to the original receiver embedding problem.

Adaptive Receiver Embedding in a Grid

While GRID-EMBED can significantly reduce the number of computations performed, estimating the number of grid points necessary for effective receiver placement is not straightforward. The number of grid points needed is highly related to the deployment area. For instance, the value provided above, of 100,000 points, would be inappropriate for a physical space of 4 m². As an alternative, we have also explored the use of an adaptive grid-based algorithm, ADAPTIVE-EMBED, which adaptively add grid points to those areas that yield the best capture gains.

The ADAPTIVE-EMBED algorithm works by initially creating a grid of points across the entire space, similar to Grid-Embed, but with much sparser grid points. After evaluating each point in the grid, we choose the one that is contained in the largest number of capture disks and create a new grid centered around the point, with dimensions of one-half the height and width of the previous grid. We continue these steps until we




Figure 4.6: The different approaches to finding solutions all yield very similar results so the fastest algorithm can be used without fear of lost performance.

select the same point in the consecutive two rounds.

Using this approach, we find that the results are comparable to F-EMBED and GRID-EMBED, but the required running time for the simulation is reduced.

4.3.4 Simulation Results: Reducing Contention with Receiver Placement

Comparison of Different Receiver Placement Strategies

Figure 4.6 shows the resulting contention level by all three algorithms, F-EMBED, GRID-EMBED and ADAPTIVE-EMBED when we vary the transmitter number from 20 to 100. We have only considered up to 100 nodes because the time needed to finish F-EMBED for over 100 nodes is too long to be realistic. The transmitters in these results are uniformly randomly distributed. We observe that though these three algorithms have very different completion times (with F-EMBED much slower than the other two), they have very similar results – the three lines are almost exactly on top of each other. The results show that having more receivers will help reduce the contention level, but the improvement becomes less pronounced as the receiver count increases. Since the

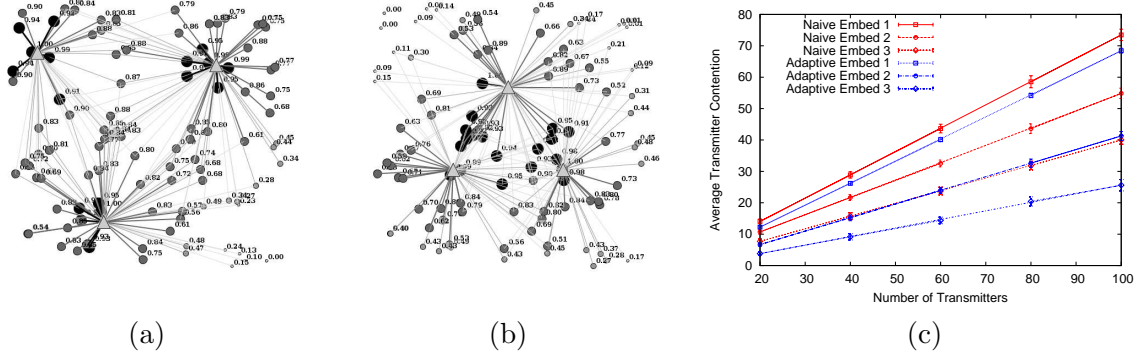


Figure 4.7: Simulation visualizations for three receiver locations calculated by ADAPTIVE-EMBED in a uniformly random distribution (a) and by naive placement (b). (c) shows an analytic comparison of their performance. Transmitters are represented by circles, larger radii reflecting a higher percentage of resolved collisions (displayed next to each transmitter). Receivers are represented by triangles, and lines are drawn between each receiver and the transmitters that it captures during collisions.

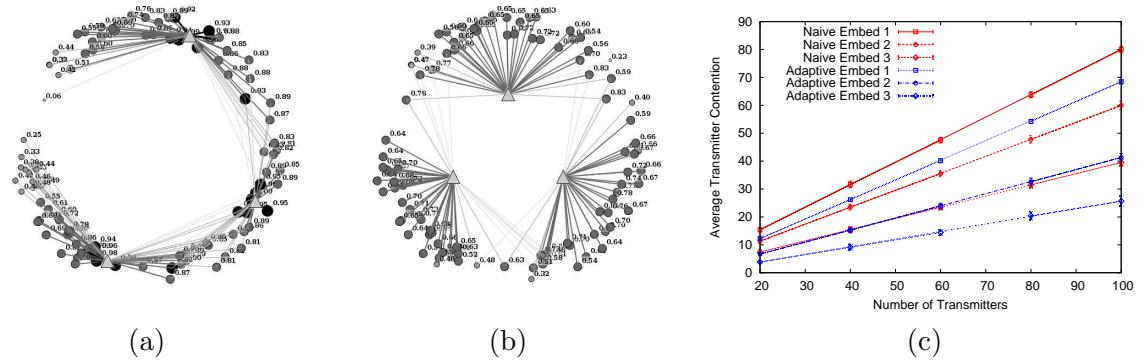


Figure 4.8: Simulation visualizations for receiver locations calculated by ADAPTIVE-EMBED along the circumference of a circle (a) and by naive placement (b). (c) shows an analytic comparison of their performance.

three algorithms give the same results, we will focus on ADAPTIVE-EMBED which is the fastest among the three.

We have used ADAPTIVE-GRID on many different deployment scenarios to calculate the receiver locations. Assuming we have 3 receivers and 100 transmitters we show the calculated receiver locations and naive locations in Figures 4.8 and 4.9 when the spatial distribution of the transmitters is along a circle or along a sine wave respectively. At the first glance, a naive placement of centrally located receivers arranged in a triangle may seem very intuitive – these receivers are in the “center” of the deployment and thus the distances between transmitters and each receiver are kept similar so transmitters achieve

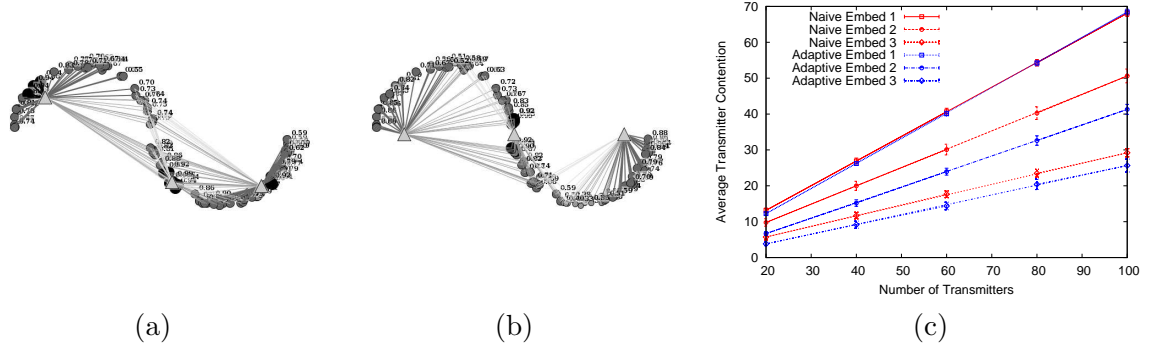


Figure 4.9: Simulation visualizations for three receiver locations calculated by ADAPTIVE-EMBED along a sine wave (a) and by naive placement (b). (c) shows an analytic comparison of their performance.

consistent and high signal strength values. Though this balanced receiver placement is beneficial for providing better signal strength coverage, it is not advantageous for maximizing the capture gains as different relative distances between transmitters and receivers are required for more captures.

The important lesson we have learned in this exercise is that receivers should be deployed where most transmitters are. Therefore, if during an actual deployment, receiver location calculation is not realistic, then we can simply place receivers where most transmitters are. This is consistent with all the calculated receiver locations that we have tried in the study.

Contention vs. Receiver Number

Next we look at the total contention reduction ratio by placing 1, 2, 3, 5, and 5 receivers using the ADAPTIVE-EMBED algorithm, and show the results in Figure 4.10. Here, we have increased the number of transmitters up to 1000 (following a uniform random spatial distribution). These results show that a small number of receivers are indeed sufficient to bring down the contention significantly. For example, 3 receivers can reduce the contention level by 71%, and 5 receivers can reduce the contention level by 87%. When we have more than 4 receivers, adding more receivers still yields benefits, but the benefit becomes smaller as we have more receivers. This observation provides the evidence that TO only needs a small number of receivers, which makes it a very practical




Figure 4.10: Our simulated results with the ADAPTIVE-EMBED algorithm show that a given number of receivers gives a predictable reduction in contention even at high numbers of transmitters. There are diminishing returns after four receivers, but at this level the contention of a 1000 transmitter network would be reduced to that of a 100 transmitter network, greatly improving performance.

and competitive approach for dense wireless systems.

We have also calculated the number of required receivers to achieve a desirable contention level, and shown the results in Figure 4.11. We have looked at three desired contention levels, 10, 30, and 100. The contention of 10 is a very strict requirement – with 1000 transmitters and .01% duty cycle, a contention of 10 will lead to a packet loss rate of .2%. To achieve this, we need 48 receivers for 1000 transmitters. On the other hand, the contention of 100 is rather relaxed. With 1000 transmitters and .01% duty cycle, a contention of 100 will lead to a packet loss rate of 2%, which can be achieved by having 7 receivers.

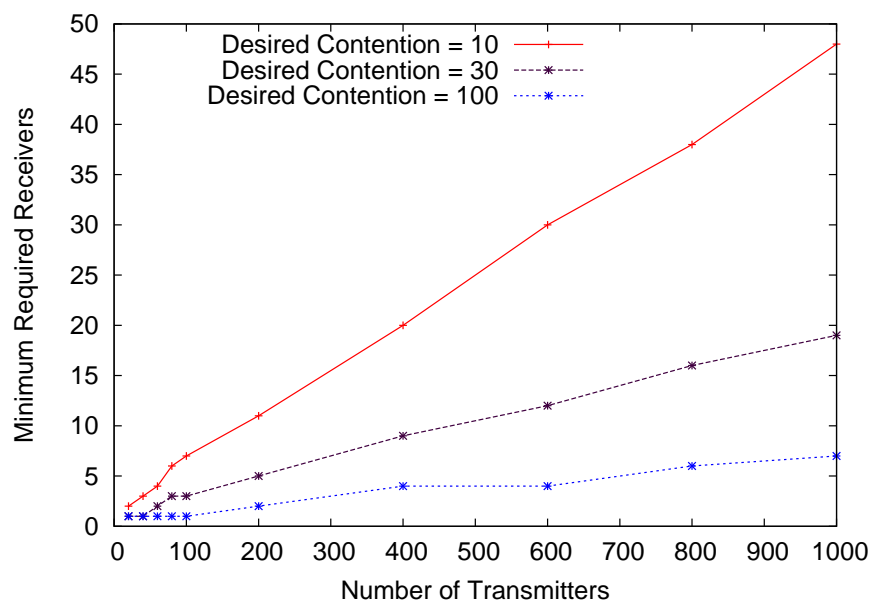


Figure 4.11: The growth of receivers to transmitters is slow even when maintaining a constant contention rate. A deployment with a specific packet loss rate requirement will not require a large number of receivers, even as the number of transmitters grows.

Chapter 5

Experimental Validation

The previous chapters of this dissertation have made a case for a *Transmit Only* (TO) communication technique over traditional wireless MAC protocols and other proposed protocols for wireless sensor networks. This section presents experimental validation of Transmit Only, first by showing agreement between experiments and the predicted capture effect and then by presenting several case studies of systems and applications that are supported by a TO deployment.

In TO, transmitters consume a constant amount of energy per transmission regardless of the number of transmitters and their offered load. The theoretical analysis shown in Chapter 3 showed that the energy efficiency of TO (as measured in energy spent to data delivered) depends upon the percent of packets successfully transmitted. Since packet success is directly determined by the collision rate (which increases with the number of transmitters) and the capture effect (a function of the number of receivers and their placement) the real-world performance of TO may differ greatly from the theoretical performance if the capture effect does not occur as predicted.

5.1 Validating a General Case

One of the first theoretical results presented was the analysis of capture the capture effect in an infinite plane of uniformly scattered transmitters that appeared in Equation 3.4. While building an infinite plane of transmitters is not possible (for this dissertation at least), it is simple to build a smaller-scale test bed of transmitters deployed in a uniform random manner. We will not use the F-Embed placement algorithm discussed in Chapter 4 because there is no structure to exploit in a uniform random distribution. Instead we will simply deploy transmitters in a way that maximizes their distance from

one another, according to the method described by Chen et al. [4].

The experiment described in this section will demonstrate two points. First, that the simple model of the capture effect in an infinite plane of transmitters is applicable to some real-life situations. Second, that the capture effect, given beyond some small number of receivers, reduces collisions more effectively than spreading traffic across multiple channels with the same number of receivers. The experiment is as follows.

5.1.1 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 3.1.2.), and 100 transmitters sending a 10-byte ($300\mu\text{second}$) packet ten times per second, creating an offered load of 30%.

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 attached USB hardware for loss-free data collection but are otherwise identical to the transmitters. More complete information can be found in [10].

The radio link will operate at 902.1 MHz. Transmitters will 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.

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 $300\mu\text{seconds}$.

We will perform tests to validate the predictions made in Figure 3.6. 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.

5.1.2 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 by Chen et al. [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 previous work 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 previous work [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

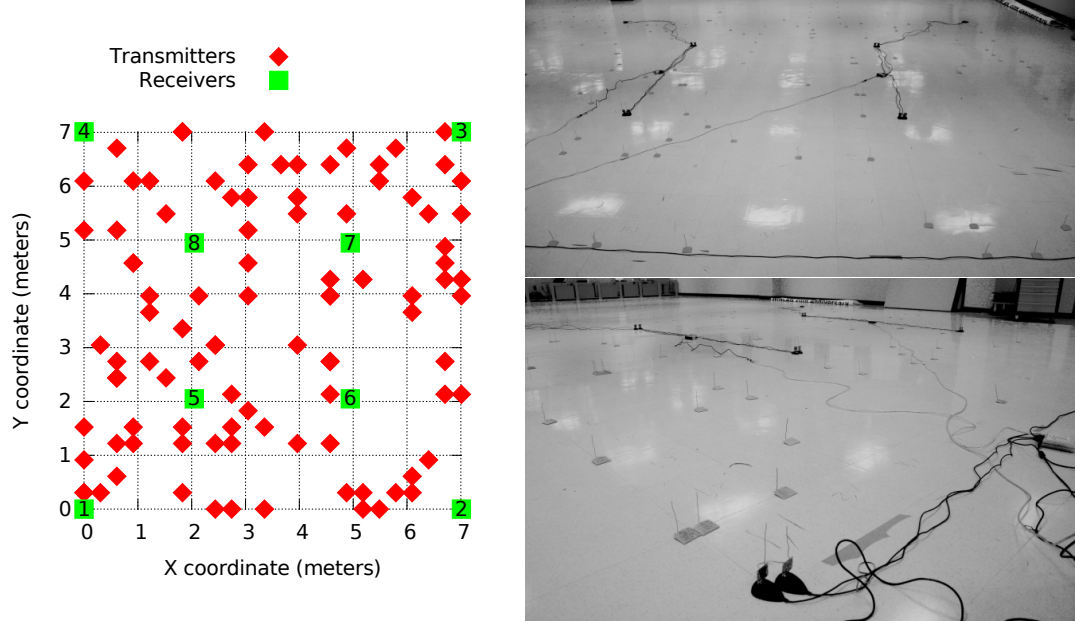


Figure 5.1: A gridded map of the experimental topology (a) and a photograph (b).

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 there are.

A map of the experimental topology appears in Figure 5.1. 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.

5.1.3 Results

Experiment results are shown in Figure 5.2. 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




Figure 5.2: 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 the number of receivers passes a threshold a single channel is better than using multiple channels, with or without redundant receiving.

supports message in message - Figure 5.2 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 $300\mu\text{seconds}$ so the offered load was 30%, of which 98.2% was achieved.

The theoretical single channel results in Figure 5.2 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 than they are, so their capture probabilities remain low. Larger deployment areas would probably achieve capture rates closer to the predicted values since they have

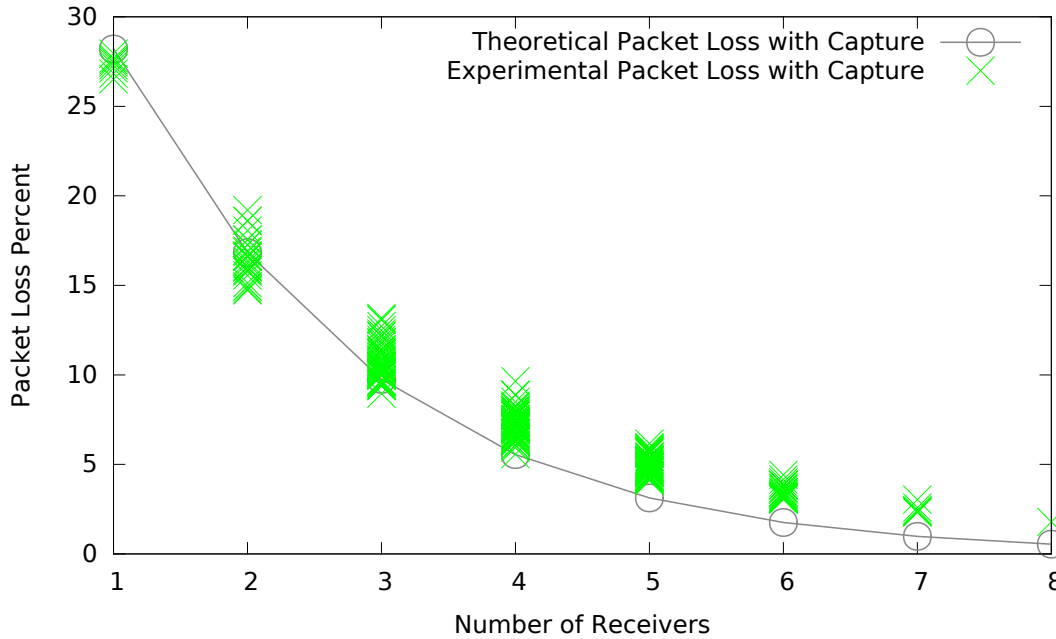


Figure 5.3: The packet loss of every single permutation of the 8 receivers sharing a single channel. We theorized that with transmitters deployed in a random uniform fashion the exact placement of the receivers would not have an order of magnitude impact. There is still something to be gained by choosing optimal, or near optimal, receiver placement. This issue will be explored in more detail in the next section.

smaller surface area to volume ratios.

Figure 5.2 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.

A hint of the effects of receiver placement can be seen in Figure 5.3. Some combinations of receivers are more effective than others, but the majority of receiver combinations have similar results. Differences are intuitive: highly asymmetrical groups of receivers give worse performance because some transmitters have no nearby receivers, and asymmetrical groups generally have more even (and better) performance because the

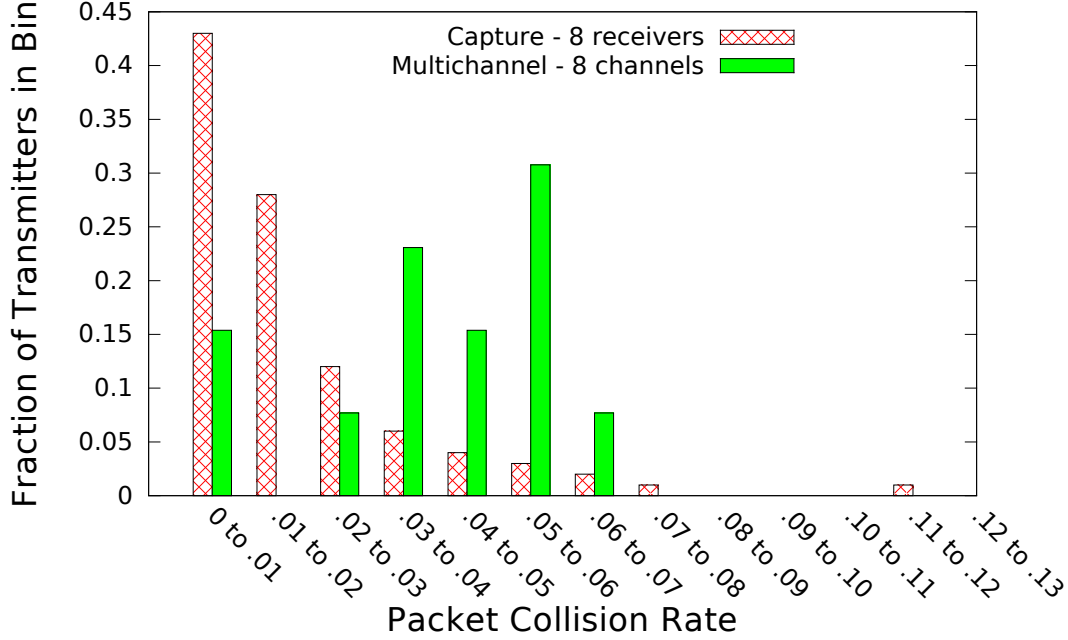


Figure 5.4: The frequency of different packet loss percentages for transmitters in the 8 receiver, single channel capture system and in the 8 receiver, 8 channel multichannel system.

distance to the closest receiver is similar across all of the transmitters. The impact of receiver placement will be discussed in more detail in the next section.

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 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, these problems are mitigated. Figure 5.4 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 while the packet loss rates in the capture system are mostly concentrated in a couple of bins and rapidly fall off. 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




Figure 5.5: High density experiments in an open area outdoors. Figure (a) shows the main set of 10 by 10 meters. (b) shows each bundle of 10 transmitters. (c) shows a receiver.

per channel without reducing the transmission rate of the transmitters closest to the receivers, which is also unfair in terms of transmission rate.

5.2 Validation Capture Predictions

This previous chapter showed that a general model of the capture effect is adequate when transmitters follow a simple, random uniform pattern. In such cases choosing receiver locations is also simple and a more complicated placement approach using the F-EMBED algorithm is unnecessary. However, this topology may not be typical in the majority of deployments. In this section we will demonstrate that we can predict the capture gains in more complicated topologies and that the F-EMBED algorithm can predict placement superior (in terms of capture effect) to simple symmetrical patterns when the deployment topology is not uniform random.

We conduct these experiments in an outdoor environment free of obstructions to confirm the theoretical models. We will also briefly show results for a dynamic (and uncontrolled) indoor environment simply to illustrate the effectiveness of TO in a real-life scenario.

5.2.1 Capture Model Validation

This experiment was conducted on a flat lawn, far (over 50m) from any buildings, in a 10m x 10m area. The purpose of this experiment is to show that the theoretical models match real hardware deployments at scale, in this case up to 500 transmitters.

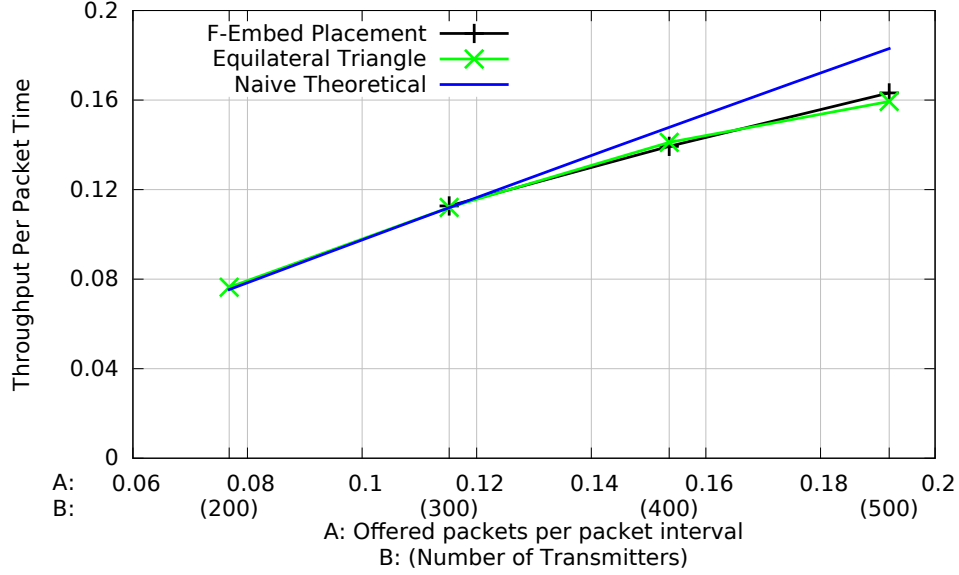


Figure 5.6: The throughput of the outdoor high-density experiment performs closely to our theoretical calculation. For simple deployment patterns, such as a uniform grid, simple receiver patterns will usually be as effective as a more complex approach, such as our F-EMBED. This is especially true when the offered load, and thus the probability of collisions, is low.

Figure 5.5(a) shows a map of the experiment for the 500 transmitter and 3 receiver case with the F-EMBED placement. Bundles of 10 transmitters are shown as black triangles.

Each transmitter was attached to a 12 inch wooden stake, as shown in Figure 5.5(b), forming a “flag” like configuration. In order to make the experiment tractable at scale, groups of 10 transmitters were bundled together in a circular arrangement, as shown in Figure 5.5(b). All three receivers are connected to a single laptop via USB extender cables.

Figure 5.6 shows the average number of packets received per packet transmitted as we scale the number of transmitters, each one beaconing a 24-bit ID once per second. It also compares the predictions of a naive theoretical model to the measured results of two receiver placements. The naive model assumes all the transmitters are equivalent to an average transmitter. The results show that even at 500 transmitters, predicted throughput closely matches the actual. A second result, also predicted by our models, is that for uniformly placed transmitters, simple uniform placement of receivers, i.e., and




Figure 5.7: The topology of the high density, high-load deployment. The F-EMBED algorithm was used to place four of the receivers, while a naive symmetrical placement was used for the naive placement.

equilateral triangle, is close to the optimal predicted by more sophisticated algorithms.

Figure 5.6 also shows that actual reception rates are good even with a modest number of receivers. With 200 transmitters 99.4% of the packets were received correctly for the equilateral triangle deployment. As we increase the number of transmitters with a constant beacon rate the idle throughput would scale along the 45 degree line. Figure 5.6 shows that 15% of the packets are lost at 500 transmitters using the F-EMBED algorithm, which is slightly better than the naive algorithm which loses 17%.

We do not expect the F-EMBED placement to significantly outperform the naive placement in this trial. The transmitters are deployed in a uniform random pattern so there is no underlying structure for F-EMBED to exploit and greatly increase the occurrence of the capture effect. The offered load is also not very high, so collisions are also not likely. When collisions become common we expect the F-EMBED placement to show increased performance beyond the naive approach.

5.2.2 F-EMBED Placement in a Dense, High-Load Deployment

In this experiment the offered load in the system was increased and transmitters were deployed in a more structured fashion. Transmitters increased their packets to 1ms durations and increased their transmission rates to two times per second. The offered load ranged from 20% with 200 active transmitters to 100% with 500 active transmitters.

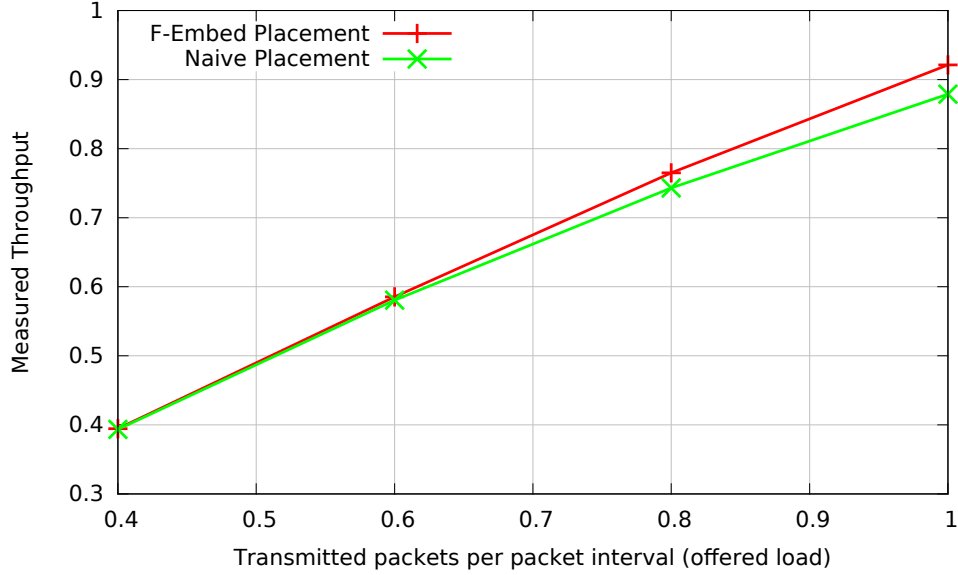


Figure 5.8: The throughput of the outdoor high-density experiment with a sine-wave deployment pattern. For simple deployment patterns, such as a uniform grid, simple receiver patterns will usually be as effective as a more complex approach, such as our F-EMBED.

Transmitters were also bundled in groups of five and deployed in a uniform random distribution along a sine-wave shaped line. This creates conditions similar to radios deployed along existing geographical or manufactured paths of interest to sensing systems, such as rivers, roads, and borders. An illustration of the deployment appears in Fig. 5.7.

At low levels of offered load the performance of the simple receiver placement and the F-EMBED placement were very similar, as shown in Fig. 5.8. However, as the offered load increased and packet collisions became more probable the F-EMBED placement begins to distinguish itself. While the simple placement only achieves 87.9% throughput with an offered load of 100%, the F-EMBED placement reduces packet by 35% compared to the simple placement, and achieves 92.1% throughput.

5.2.3 Validation in a Real-Life Setting

In this section we describe our experiences using TO in an indoor laboratory setting. We had 80 transmitters and 3 receivers. We placed the receivers using the F-EMBED algorithm.




Figure 5.9: Packet loss in our indoor deployment. Even though individual receivers saw a large variation in their hour by hour packet reception, the three together maintained a packet reception rate of about 99% throughout.

We found that we had an average 99.1% packet reception rate. We also found that our approach of over-provisioning the beacon rate worked well in practice because it allowed to quickly build a few applications. Additionally, the presence of multiple receivers provided redundancy even when one receiver experiences poor performance. For instance, in Figure 5.9 while receiver B experienced very poor reception during some hours, perhaps due to occlusion by a human, the three receivers in aggregate only suffered a mild degradation of performance.

Chapter 6

Use Cases

An important underpinning of this dissertation is the assertion that unidirectional communication is sufficient for many use-cases of wireless networks. To provide examples backing this assertion, this chapter will go over two important applications that can be built with unidirectional Transmit Only systems. The first is a simple inventory tracking system, where system users wish to be alerted as soon as possible when an item is missing but will tolerate some latency (on the order of seconds) to reduce false positives. The second example is of a sensing and monitoring system, with various logic dictating how the system responds to sensed events.

6.1 Example 1: Asset Tracking

Asset tracking is an important application domain for wireless sensor networks. However, continuous tracking of a large number of items at the individual item level over a significant period of time is still not feasible. There are two main obstacles. The first is the need for efficient, low-power communication protocols. Many current protocols employ energy-expensive methods to achieve reliable communication for arbitrary traffic situations. Such protocols are not suitable for continuous asset tracking applications. The second challenge is the lack of a robust presence detection algorithm that can differentiate packet losses caused by a missing item from packet losses caused by the ambient radio environment. This section will show that TO can be used to create a “heartbeat” protocol that supports two robust detection algorithms yielding low false alarm rates while achieving timely loss notification.

6.1.1 Problem Statement

Wireless sensor networks have enabled many applications that were impossible before, some of which are continuous, item-level asset tracking applications. A system that provides continuous item-level tracking capability can be useful in many scenarios. Imagine, such a system will allow the customers at a jewelry store to try on the merchandise and appreciate them from many angles while *freely* strolling around the store; such a system can help keep track of the whereabouts of patient charts so that they are *always* available when needed; such a system can also ensure a soldier's backpack is *constantly* equipped with all the required necessities. Compared to traditional asset tracking applications which can only report missing items at a coarse temporal (i.e. once a day) and spatial (i.e. not every item) granularity, continuous, item-level asset tracking applications demand much finer tracking granularity both in time and space. These applications must also be fail-safe, unlike many current inventory systems. Despite the much finer tracking granularity, these applications are not willing to compromise on the system lifetime – they usually require more than one year's lifetime using coin cell batteries. The finer granularity of this system lends itself to a new kind of asset tracking system - a *presence assurance system* (PAS). Whereas the status of an item in current asset tracking systems is expressed in the sentence “item X was last seen at location Y” the status of an item in a PAS is “item X is currently at location Y.” If a PAS is robust against sensor failure, tampering, and theft then we call it a fail-safe PAS.

Building a fail-safe PAS poses new challenges for the underlying system design, in both hardware and software. Pre-existing solutions do not adequately address these challenges[38]. For example, passive RFID tags, which are popular for traditional inventory tracking applications, suffer from poor range, difficulty avoiding transmission collisions in dense environments, and poor performance when attached to certain items. Sensor networks have also been deployed for asset tracking purposes [12, 22]. While these platforms provide good range and sophisticated communication protocols, they are designed for more general-purpose usage, and their protocols and algorithms may

not suit the specific requirements of these applications, thus leading to excessive energy consumption and much reduced lifetime.

Earlier work [46, 2] has argued for a “push” architecture for continuous asset tracking applications in which a sensor is attached to every item in the system, and the sensor periodically (e.g., once a second) transmits a packet announcing its presence. Once a sensor is not heard for a period of time, the system can declare that this sensor, or the item that it is guarding, is absent. Though this architecture has been proven effective through simulations, practical studies on how to build such an architecture are lacking because some of the previously discussed algorithms are too complex to be implemented on a simple and low-cost platform. The straightforward nature of TO is advantageous here, because it allows us to build applications with very inexpensive radio hardware.

To solve the asset tracking problem we will build a *uni-directional heartbeat* (Uni-HB) on top of the TO system. Each sensor will send a small radio beacon at predictable time intervals. The predictability of packet transmission time and the extremely short packets in our system (only a few bytes) make the uni-directional heartbeat protocol more suitable than more power hungry protocols while also providing reliable packet delivery. Also, we show that we can predict the packet collisions in this protocol in software because the packet arrival times from the same sensor are regular and can be easily anticipated, allowing us to reduce false positives during loss detection.

However, packets can also be lost due to the ambient radio environment, which can lead to incorrect detection results if the packet loss is assumed to correspond to an item loss. We will show that a good detection algorithm can assign probabilities to these events to minimize item loss response time. Our first detection algorithm is based on the length of packet miss chains: it discerns the sensor as missing when the current miss chain is much longer than what has been observed in the history. Our second detection algorithm checks the probability of a miss chain by calculating the probability of an ambient loss based upon RSSI values. Using a trace collected over the course of a week, we show that our detection algorithms can yield both low false alarm rates and short detection latencies.

6.1.2 Collision Prediction

In Uni-HB, each sensor sends out a beacon around the same time within transmission interval (called an epoch), allowing the back-end processing software to accurately calculate the anticipated packet arrival times and to predict the collisions that will occur in the future epoch. Based on the predicted collision occurrence, the system can correct the received packet sequence from each sensor by marking the unreceived collided packet as received, thus leading to a higher packet delivery ratio. In this section, we discuss how this collision prediction process works.

We imagine that an asset tracking system will consist of hundreds to thousands of transmitters, tens of basestations, and one central processing station (CPS). Data flows, in the form of packets, from the transmitters to the basestations, and from there to the CPS. Basestations are deployed to have overlapping coverage, allowing each transmitter to be heard by at least two basestations. Basestations listen for packets from the transmitters and place a timestamp on each packet before sending the received packet (together with the timestamp) to the CPS over an ethernet, WLAN, or USB connection. Timestamps from the basestation are most accurate when it merely places a time stamp on received packets and forwards them to the CPS. The timestamps should give twice the granularity of transmit events, so a 1/5millisecond resolution is recommended for the 2/5millisecond transmit times. The CPS will receive a stream of data from each basestation with each data entry consisting of the fields: source basestation ID, basestation timestamp, source transmitter ID, packet sequence number, RSSI value, and CRC status.

The CPS first tries to synchronize all the timestamps from different basestations, and merge all the data entries into a single *arrived* queue according to the adjusted timestamps with all the redundant packets from the same transmitter (but received by different basestations) removed. In addition to the arrived queue, the CPS also maintains the *expected arrival* queue, which includes the expected arrival time for each transmitter in the system. Both queues are sorted in the increasing temporal order. The actual processing involves taking the first packet in the arrived queue and searching

for its transmitter ID in the expected arrival queue. Usually, this search only needs to check the first few expected arrival events in the queue. If the transmitter is found, the CPS updates its expected arrival time, $T_{expected} = T_{now} + T_{epoch}$ and moves the event to the appropriate location in the queue (most likely the end). If the search has to go through the entire queue and cannot find the corresponding transmitter ID, it means that this is a new transmitter and the new ID is added to the expected arrival queue.

In such a system, collisions can be easily predicted. Every time when an updated expected arrival event is relocated to a new position in the queue, we check its expected arrival time with the expected arrival times of the neighboring events in the queue. If the gap between two expected arrivals is less than the *hazard window*, we mark these two packets as being potential collision hazards. If the packet with a hazard mark is not received around its expected time, the CPS assumes it was a collision and the item/transmitter is actually not missing. Each transmitter's epoch duration slightly varies from each other so no two packets will transmit at the same time for a large number of epochs. Here, the size of the hazard window is an important parameter. A hazard window which is too small will under-predict collisions, while a hazard window which is too large will over-predict and might mistakenly treat packet loss due to other sources (e.g. the item missing) as collisions and ignore them. The most reliable and adaptive way to determine the hazard window is to set it equal to a confidence level of the expected arrival time of a packet. Inter-packet times are estimated and thus have some error. By using the confidence interval of this estimation as the basis for the hazard we compensate for this error and achieve good hazard prediction. In our system we generally use the 99.9% confidence interval, depending upon the clock resolution of our basestations.

After removing the predicted packet losses from the total number of packet losses, unexpected packet loss was less than 1% of all packet losses. Note that the collision prediction technique cannot be applied to protocols like CS and ACK because their packet arrivals are not regular due to their feedback or collision avoidance mechanisms.

6.1.3 Robust Detection of Transmitter Presence From Packet Loss

The ultimate goal of a PAS system is to rapidly and reliably detect when an item is missing. Ideally every transmitter would send a packet at a fixed interval (called an epoch here), and a receiver would see one packet every epoch from every transmitter in the system. If a sensor is lost or broken, then it ceases beacon transmissions and the system would realize immediately that the sensor has been stolen or broken. In reality though, one cannot, and should not, simply equate a missing packet to a missing sensor because packet loss occurs for other reasons. As a result, a robust detection algorithm tries to differentiate packet losses due to missing sensors and packet losses due to other causes.

Ambient Packet Loss

There are two kinds of ambient packet loss in TO, those caused by low received signal intensity being indistinguishable from background interference and thermal noise, referred to as *ambient loss*, and those caused by direct collisions between packets in the system, referred to as *collision loss*. As discussed in Section 6.1.2, a large fraction of the collision loss can be identified by the CPS through the collision prediction technique. In this section, we will mainly focus our discussion on ambient packet loss.

Ambient loss defines the baseline packet loss for a sensor. Its magnitude varies signals from interfering RF devices (such as other wireless devices, microwave ovens, and electric motors) and radio propagation effects. The typical ambient loss rate observed during our trials and deployments is less than one percent. Even poorly placed sensors can communicate effectively if basestations are deployed so that the multiple basestations receive the sensor's packets.

6.1.4 Detection Algorithms

A good detection algorithm will identify that contact has been lost with a sensor quickly enough for a person to react to the problem but will not have so many false positives that the alarm is no longer taken seriously. In this chapter, we measure the effectiveness




Figure 6.1: Experimentally derived ambient miss percentage versus the RSSI, with the 95% confidence interval for each point.

of the detection algorithm of the RollCall system using two metrics: the probability of raising a false alarm – a false alarm means the system declares a sensor/item missing while the sensor is still present – P_f , and the time it takes for the system to raise an alarm after a sensor ceases transmitting packets, T_{alarm} . In this section, we consider detection based on the resulting packet sequence from the Uni-HB protocol with collision prediction. As a result, packet misses observed by the CPS will be either due to ambient loss or a missing transmitter.

There is a tradeoff between P_f and T_{alarm} . If we sought to minimize the detection latency T_{alarm} without considering P_f , then we would use a naïve technique where a sensor is considered lost whenever a single packet is missed. This gives the smallest possible T_{alarm} (1 epoch) at the expense of a very high P_f .

At the other extreme, if we only focus on minimizing P_f without regard to T_{alarm} , then we can simply wait very long time periods before reporting that a contact has been lost with a sensor. While waiting for five minutes to pass without seeing a single packet from a sensor will give a very low P_f which will be too long for anyone to respond to

an item's loss after it is detected. As a result, this extreme is not desirable as well.

In this study, we explored the following detection heuristics:

Detection Based on Historical Maximum Miss Chain Length (Detect-MaxMiss)

Ambient packet losses often occur in bursts because the ambient factor may last for a period of time. In *Detect-MaxMiss*, for each sensor, we keep track of the length of the longest miss chain. The baseline detection is simple: once the current miss chain is longer than the previously observed maximum miss chain length, we declare the sensor is missing. In order to avoid the high false alarm rates when the miss chains are short, we declare a sensor is missing if the current miss chain length is at least K misses more than the previous maximum chain length. K can be determined heuristically. We call this *Detect-MaxMiss+K*.

Detection Based on RSSI (Detect-RSSI)

In this algorithm, we first have a training phase to build a mapping between the RSSI value of the current received packet and the likelihood that the next packet from the same sensor will be missing. For example, we find that if the current packet's RSSI value is low, then the probability to have an ambient loss in the following epoch is high. Figure 6.1 shows the ambient loss ratio with RSSI histogram we collected in our lab.

In the test phase, we detect whether a sensor is missing based on the ambient loss ratio with RSSI histogram. Next, let us look at an example to understand the detection algorithm. Suppose we have the following packet sequence for sensor i :

`Recv(-97dB), Miss, Recv(-95dB), Miss, Miss, ...`

Then the detection algorithm works as follows. In the first epoch, since the packet is received, the probability of the sensor missing $P_{missing} = 0$. In the second epoch, the packet is lost, but from the last RSSI value (-97dB), we know that the likelihood of having an ambient loss in this round is roughly 0.8. As a result, we estimate $P_{missing} = 0.2$. Since $P_{missing}$ is below the preset threshold 0.8, we do not raise an alarm. In the third epoch, again we have $P_{missing} = 0$. In the fourth epoch, the packet is lost, and $P_{missing} = 1 - 0.3 = 0.7$ based on the previous RSSI value. In the fifth epoch, the

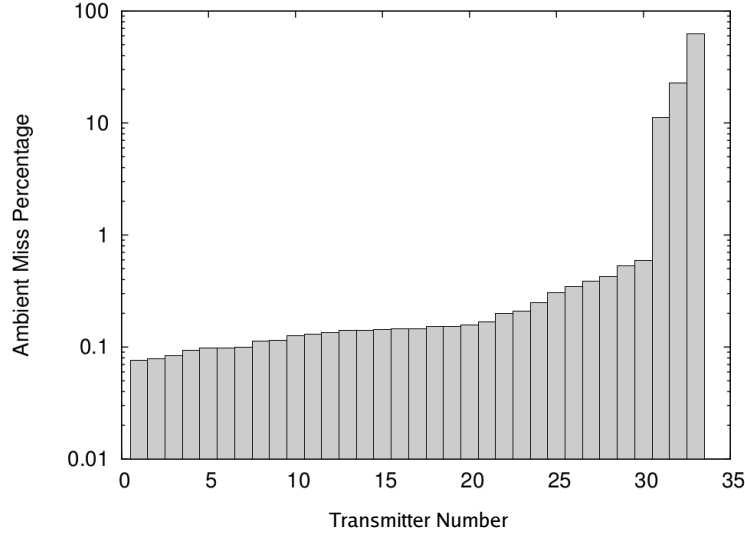


Figure 6.2: The ambient miss percentages seen by the transmitters during the detection trial.

packet is again lost, and we calculate $P_{missing} = 1 - (1 - 0.7)^2 = 0.91$. Here, since we do not have the RSSI value in the last round, we have to borrow the value two rounds ago to estimate $P_{missing}$. At this point, we will report the sensor being missing because $P_{missing}$ is above the threshold value. We note that we could also calculate $P_{missing}$ in the fifth epoch based on the likelihood of having 2 successive ambient losses following a received packet with RSSI value of -95dB. We will build such histograms in our future work.

6.1.5 Detection Results

We tested our detection algorithms using the trace that was collected using 31 transmitters over the period of a week. The transmitters were arranged in groups of three and were placed across several rooms in an office environment. Some of the transmitters were placed in metal bins and drawers. One transmitter was intentionally placed so far away from the basestations that the basestation was on the edge of its radio range and

Method	P_f (%)
<i>Detect - SingleMiss</i>	100
<i>Detect - MaxMiss + 5</i>	.013
<i>Detect - RSSI</i>	.011

Table 6.1: Average P_f for different detection algorithms.




Figure 6.3: T_{alarm} for each transmitter in the detection trial: (a) *Detect – MaxMiss*, and (b) *Detect – RSSI*.

its observed signal was extremely weak. During the testing week, those transmitters were moved, sometimes outside of the radio range of the basestation to test the system when an item is slowly removed, or had their batteries removed to test loss detection.

Table 6.1 shows the average P_f for the two proposed detection algorithm together with the naive algorithm *Detect-SingleMiss*. Here, P_f is defined as the ratio of the number of times the system declares a sensor is missing with respect to the number of miss chains in the trace. Since *Detect-SingleMiss* assumes a missing sensor every time there is a missed packet, the value of P_f is 100%. Compared to *Detect-SingleMiss*, both of our algorithms have substantially reduced the false alarm rate by correcting differentiating ambient and collision losses from actual loss events.

On the other hand, it is not meaningful to look at the average T_{alarm} over all the sensors as some of them were purposefully placed to have very poor signals. As a result, it is much harder to report that they are missing in a timely fashion – how can a police officer quickly detect that a person is missing if he only leaves his secret vault once a month? In this case, we look at T_{alarm} for each individual transmitter. Figure 6.2 plots the percentage of packet misses due to ambient loss for each of the 31 transmitters, showing that three of the transmitters had very poor transmissions. Figures 6.3(a) and (b) show the resulting individual T_{alarm} for the two detection algorithms. From

these results, we can clearly see that transmitters that have a poor ambient environment need a long period to be declared missing, while transmitters with a reasonable ambient environment can be reported missing in a much more timely fashion (around 10 epochs). This suggests that when deploying a tracking system, we need to make sure all the transmitters are covered by basestations. This will not cause any issue in the system we consider because the basestations share the same hardware as the transmitters, and as a result, share the same cost as well. Thus, having more basestations is easy to achieve. Further, we observe that the *Detect-RSSI* algorithm is more robust in rapidly changing environments and when sensors experience poor signal quality. The *Detect-MaxMiss+5* algorithm will adjust to a changing environment as its history window updates, but since radio environments can change rapidly the *Detect-RSSI* algorithm has an advantage.

6.2 Example 2: Monitoring and Notification

The TO architecture is inherently unidirectional, from the transmitters to a receiver. This is a natural fit for sensing and monitoring applications, where data is collected from many sensors before being disseminated to a user. There are some situations, such as Smart Grid applications or vehicular sensing, where sub-second reporting latency and reliability to several significant figures are required. However, the majority of existing and anticipated applications for wireless sensors do not require this kind of low latency. Using TO in smart homes, habitat monitoring, or data center monitoring would be perfectly reasonable because trading off a few seconds of latency in reported data (caused by missed packets) is worth the ease of maintenance for the users of these systems. Unfortunately, not all researchers would agree with this statement.

There is a movement that promotes running higher level networking on wireless sensor, such as IPv6 [14], allowing users to directly query each individual sensor. It is the opinion of this author though, that this approach is poorly matched to the idea of a sensor network. Users want information in one place, but so the information from the individually addressable sensors would need to be aggregated in any case. If a user wanted alerts when certain conditions in the sensors occurred, then the sensors would need to

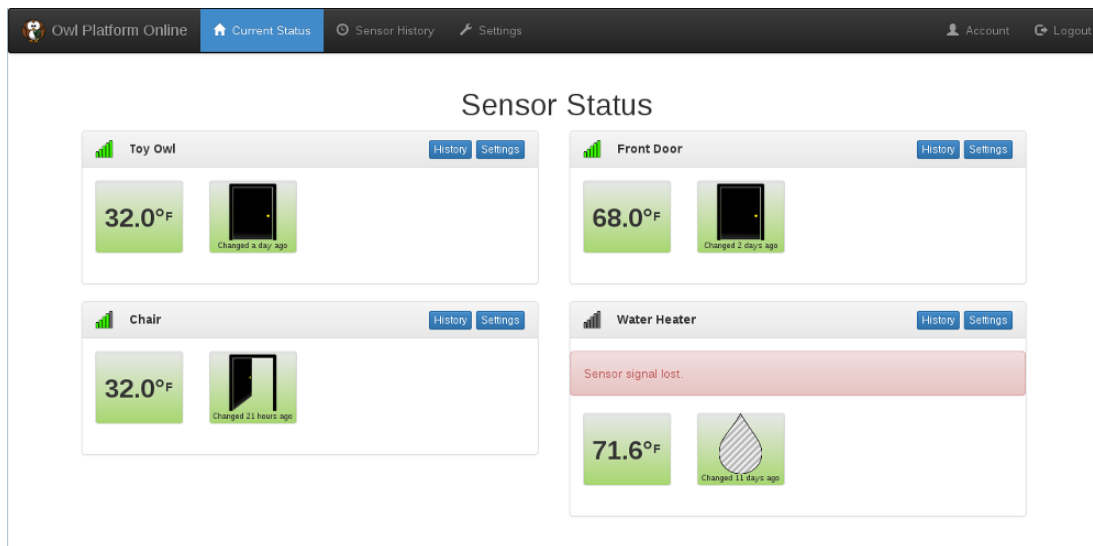


Figure 6.4: Screenshot of the Owl Platform current sensor status page. It is unnecessary to have individually addressable sensors since all of their information can be easily integrated into a back-end system.

communicate with one another to determine if a given state is reached, and would then need use a very high level communication system, such as SMS or e-mail, to actually contact the user. This would probably be handled by the same system that aggregated data, so why not just have an aggregation system and simplify the sensors? This is the rational behind TO.

An example of a system using this approach is Owl Platform [16], a website that allows users to view the status of their sensors and sign up for e-mail and SMS alerts. A screenshot of the current status of sensors in the system appears in Figure 6.4. Sensors used in the system are running TO, but the details of the sensors are hidden from (and unimportant to) users of the system. Historic data is easily stored in the back-end of the system, shown in Figure 6.5. Packet loss in the system may cause latencies, and hence inaccuracies, of less than a second per beacon missed.

While there are cases when individually addressable sensors are required, abstracting away the complexities of individual sensors is usually preferable for the user. Whenever this is the case, there is no disadvantage in using TO over more complex protocols that allow for bi-directional communication and the benefits of lifetime and ease of

Owl Platform Online
Current Status
Sensor History
Settings
Account
Logout

Sensor History

Sensor	Data Type
Toy Owl	Open Closed
Front Door	Open Closed
Chair	Open Closed
Water Heater	Wet Dry

Front Door

Past 24 Hours

1

Timestamp	Event	Duration
2013-10-08 03:22:22 pm	open	2 seconds
2013-10-08 03:21:34 pm	open	4 seconds
2013-10-08 03:20:56 pm	open	3 seconds
2013-10-08 03:20:12 pm	open	1 second
2013-10-08 03:19:17 pm	open	26 seconds
2013-10-08 03:17:57 pm	open	a minute

Copyright © 2013 Inpoint Systems Inc. All rights reserved.
[Terms of Service](#) | [Privacy Policy](#)

Figure 6.5: Screenshot of the Owl Platform historic sensor status page. Sensor history can be stored in the backend of the system, making it unnecessary for sensors to store data or respond to queries for historic data and allowing system designers to use TO in their deployments.

deployment work in TO's favor.

Chapter 7

Conclusions

In this work we demonstrated the advantages of Transmit Only in a variety of settings and network topologies. We demonstrated that with careful placement of receivers, we can leverage the capture effect to recover colliding packets and developed analytic models to approximate optimal receiver placement and predict their performance. We also showed that with simple deployment patterns (e.g. transmitters dispersed in random uniform patterns) complex analysis to predict optimal receiver placement is not required to achieve high capture rates and system performance. These models were validated via experiments scaling up to 500 transmitters.

Both our models and experiments are in close agreement. These support the feasibility of TO in real deployments by showing that the rate of growth of receivers to transmitters to maintain a desired level of performance is modest and may be linear in nature. We showed that with only five receivers, contention in a network can be reduced by 90%, although there are diminishing returns with each additional receiver so getting to extremely high reliability is not possible.

This makes TO suitable for many emerging applications where lifetime and maintainability are more important than latency and guaranteed reliability, for instance long-term environmental monitoring or inventory systems. The transmit-only, converge-cast system used in 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 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 to better take advantage of the capture effect [27]. 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.

The most important result of this dissertation is the idea that a highly asymmetric wireless protocol can be used to make the energy cost of wireless devices extremely small. This dissertation has shown that it is possible to run wireless radios for years on just a coin cell battery, even in deployments with high density and transmitter mobility. Transmit Only trades off packet delivery guarantees and feedback to the transmitters to make this low energy requirement possible. In return though, wireless devices using TO will have spare energy for sensing, and algorithms running external to the transmitters can detect location, mobility, or loss. This reduces the maintenance cost and effort for wireless systems to the point where large scale deployments with hundreds or thousands of sensors are feasible for large organizations and companies. The ease of maintenance may also make it possible for sensing deployments in private settings, in smart home and healthcare applications, where frequent visits by technicians would be unwelcome. TO is an enabling technology for many kinds of applications that were previously impracticable and fills an important niche among the many wireless protocols currently in use.

References

- [1] A. Arora, R. Ramnath, E. Ertin, P. Sinha, S. Bapat, V. Naik, V. Kulathumani, H. Zhang, H. Cao, M. Sridharan, S. Kumar, N. Seddon, C. Anderson, T. Herman, N. Trivedi, C. Zhang, M. Nesterenko, R. Shah, S. Kulkarni, M. Aramugam, L. Wang, M. Gouda, Y.-r. Choi, D. Culler, P. Dutta, C. Sharp, G. Tolle, M. Grimmer, B. Ferriera, and K. 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.
- [2] G. Bhanage, Y. Zhang, Y. Zhang, T. Wade, and R. Howard. Rollcall : The design for a low cost and power efficient active rfid asset tracking system. In *Eurocon*, 2007.
- [3] Bluetooth SIG. *Bluetooth Specification Version 4.0*, June 2010.
- [4] Y. Chen, J.-A. Francisco, W. Trappe, and R. 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] O. Chipara, C. Lu, T. C. Bailey, and G.-C. Roman. Reliable clinical monitoring using wireless sensor networks: Experiences in a step-down hospital unit. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, SenSys '10*, pages 155–168, New York, NY, USA, 2010. ACM.
- [6] S. Cui, A. J. Goldsmith, and A. Bahai. Energy-efficiency of mimo and cooperative mimo techniques in sensor networks. *Selected Areas in Communications, IEEE Journal on*, 22(6):1089–1098, 2004.
- [7] A. El-Hoiydi and J.-D. Decotignie. Wisemac: an ultra low power mac protocol for the downlink of infrastructure wireless sensor networks. In *Computers and Communications, 2004. Proceedings. ISCC 2004. Ninth International Symposium on*, volume 1, pages 244 – 251 Vol.1, june-1 july 2004.
- [8] Ericsson. More than 50 billion connected devices, Feb. 2011.
- [9] B. Firner, P. Jadhav, Y. Zhang, R. Howard, W. Trappe, and E. 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.
- [10] B. Firner, S. Medhekar, Y. Zhang, R. Howard, W. Trappe, P. Wolniansky, and E. Fenson. PIP Tags: Hardware Design and Power Optimization. In *Proceedings of the Fifth Workshop on Embedded Networked Sensors (HotEmNets)*, June 2008.

- [11] B. Firner, C. Xu, R. Howard, and Y. Zhang. Multiple receiver strategies for minimizing packet loss in dense sensor networks. In *Proceedings of the eleventh ACM international symposium on Mobile ad hoc networking and computing*, MobiHoc '10, pages 211–220, New York, NY, USA, 2010. ACM.
- [12] M. Gaynor, S. Moulton, M. Welsh, E. LaCombe, A. Rowan, and J. Wynne. Integrating wireless sensor networks with the grid. *IEEE Internet Computing*, 8(4):32–39, 2004.
- [13] G. P. Halkes and K. G. Langendoen. Crankshaft: an energy-efficient mac-protocol for dense wireless sensor networks. In *Proceedings of the 4th European conference on Wireless sensor networks*, EWSN'07, pages 228–244, Berlin, Heidelberg, 2007. Springer-Verlag.
- [14] J. W. Hui and D. E. Culler. IP is dead, Long Live IP for Wireless Sensor Networks. In *SenSys*, pages 15–28, 2008.
- [15] IEEE. IEEE 802.15.4-2006 Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANS). 2006.
- [16] Inpoint Systems. OwlPlatform. Retrieved October 8, 2013, from Owl Platform: <http://secure.owlplatform.com>, 2013.
- [17] E. Internet Engineering Task Force (IETF), J. Hui. Compression Format for IPv6 Datagrams Over IEEE 802.15.4-Based Networks. September 2011.
- [18] T. O. Kim, J. S. Park, H. J. Chong, K. J. Kim, and B. D. Choi. Performance analysis of ieee 802.15.4 non-beacon mode with the unslotted csma/ca. *Communications Letters, IEEE*, 12(4):238 –240, april 2008.
- [19] Y. Kim, H. Shin, and H. 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.
- [20] A. Lachenmann, P. J. Marrón, D. Minder, and K. Rothermel. Meeting lifetime goals with energy levels. In *Proceedings of the 5th international conference on Embedded networked sensor systems*, SenSys '07, pages 131–144, New York, NY, USA, 2007. ACM.
- [21] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. 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.
- [22] M. Malinowski, M. Moskwa, M. Feldmeier, M. Laibowitz, and J. A. Paradiso. Cargonet: a low-cost micropower sensor node exploiting quasi-passive wakeup for adaptive asynchronous monitoring of exceptional events. In *SenSys '07: Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 145–159, New York, NY, USA, 2007. ACM.

- [23] P. Park, C. Fischione, and K. H. Johansson. Adaptive ieee 802.15.4 protocol for energy efficient, reliable and timely communications. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, IPSN '10, pages 327–338, New York, NY, USA, 2010. ACM.
- [24] I. Ramachandran and S. Roy. Analysis of throughput and energy efficiency of p-persistent csma with imperfect carrier sensing. In *Proceedings of the Global Telecommunications Conference (GLOBECOM)*, page 5, St. Louis, MO, USA, Nov. 2005.
- [25] Y. ri Choi, M. G. Gouda, H. Zhang, and A. Arora. Routing on a logical grid in sensor networks. Technical report, 2004.
- [26] D. Rohm, M. Goyal, H. Hosseini, A. Divjak, and Y. Bashir. Configuring beaconless ieee 802.15.4 networks under different traffic loads. In *Advanced Information Networking and Applications, 2009. AINA '09. International Conference on*, pages 921 –928, may 2009.
- [27] 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.
- [28] D. Son, B. Krishnamachari, and J. 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.
- [29] F. Stajano, N. Hault, I. Wassell, P. Bennett, C. Middleton, and K. Soga. Smart bridges, smart tunnels: Transforming wireless sensor networks from research prototypes into robust engineering infrastructure. *Ad Hoc Netw.*, 8:872–888, November 2010.
- [30] Y. Sun, O. Gurewitz, and D. B. Johnson. Ri-mac: A receiver-initiated asynchronous duty cycle mac protocol for dynamic traffic loads in wireless sensor networks. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, SenSys '08, pages 1–14, New York, NY, USA, 2008. ACM.
- [31] H. Takagi and L. Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *Communications, IEEE Transactions on*, 32(3):246 – 257, mar 1984.
- [32] H. Takagi and L. Kleinrock. Throughput analysis for persistent csma systems. *Communications, IEEE Transactions on*, 33(7):627 – 638, jul 1985.
- [33] L. Tang, Y. Sun, O. Gurewitz, and D. B. Johnson. EM-MAC: A Dynamic Multichannel Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *The Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc11)*, May 2011.
- [34] Texas Instruments Incorporated. *CC2500 Low-Cost Low-Power 2.4 GHz RF Transceiver*, 2011. Data Sheet.

- [35] Texas Instruments Norway AS. *CC1100 Single Chip Low Cost Low Power RF Transceiver*, June 2006. Data Sheet (Rev.1.1).
- [36] N. Timmons and W. Scanlon. Analysis of the performance of ieee 802.15.4 for medical sensor body area networking. In *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*, pages 16 – 24, oct. 2004.
- [37] T. van Dam and K. Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the 1st international conference on Embedded networked sensor systems*, SenSys '03, pages 171–180, New York, NY, USA, 2003. ACM.
- [38] R. Want. An introduction to rfid technology. *Pervasive Computing, IEEE*, 5(1):25–33, Jan.-March 2006.
- [39] T. Wark, P. Corke, P. Sikka, L. Klingbeil, Y. 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.
- [40] T. Wark, C. Crossman, W. Hu, Y. Guo, P. Valencia, P. Sikka, P. Corke, C. Lee, J. Henshall, K. Prayaga, J. O’Grady, M. Reed, and A. 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.
- [41] J. Wilson and N. Patwari. Radio tomographic imaging with wireless networks. *IEEE Transactions on Mobile Computing*, 9(5):621–632, 2010.
- [42] C. Xu, B. Firner, Y. Zhang, R. Howard, J. Li, and X. Lin. Improving rf-based device-free passive localization in cluttered indoor environments through probabilistic classification methods. In *Proceedings of the 11th ACM/IEEE Conference on Information Processing in Sensor Networks*, IPSN '12, 2012.
- [43] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1567 – 1576 vol.3, 2002.
- [44] W. Ye, F. Silva, and J. Heidemann. Ultra-low duty cycle mac with scheduled channel polling. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 321–334, New York, NY, USA, 2006. ACM.
- [45] H. Zhang, A. Arora, Y.-r. Choi, and M. 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.
- [46] Y. Zhang, G. Bhanage, W. Trappe, Y. Zhang, and R. Howard. Facilitating an active transmit-only rfid system through receiver-based processing. *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON '07. 4th Annual IEEE Communications Society Conference on*, pages 421–430, 18-21 June 2007.

- [47] ZigBee Alliance. ZIGBEE 2007 SPECIFICATION.