DISTANCE ESTIMATION BETWEEN TRANSCEIVERS OVER SHORT DISTANCES

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Computer Engineering

---

by
Karsten Lowe
December 2007

Accepted by:
Dr. Adam W. Hoover, Committee Chair
Dr. Eric R. Muth
Dr. Richard R. Brooks

ABSTRACT

Three methods for determining distances between a user and a fixed coordinate system are considered. One system is based on 802.11g packet communications, another on HF radio carrier frequency interference, and the final system on a signal transmitted over a radio channel. Emphasis is placed on the use of these systems in indoor, building-sized environments. Tests are performed to examine the effectivness and potential of these methods. None of the suggested methods are determined to be usable. Reasons for their failure are examined.

TABLE OF CONTENTS

Page

LIST OF FIGURES

Figure Page

List of Figures (Continued)

Figure Page

List of Figures (Continued)

Figure Page

LIST OF TABLES

Table Page

CHAPTER 1

INTRODUCTION

People have always needed to know where other people and objects of importance are located. From ancient Roman milestones to modern-day GPS systems, the speed and accuracy of location measurement has improved. Despite all the advances in outdoor location technologies, there is still a deficiency in modern tracking capability: the inability to accurately locate entities indoors. This thesis is interested in tracking people and resources with a high degree of accuracy (on the order of 1 cm or less), particularly in indoor environments of any size using as little instrumentation and *a priori* calibration as possible.

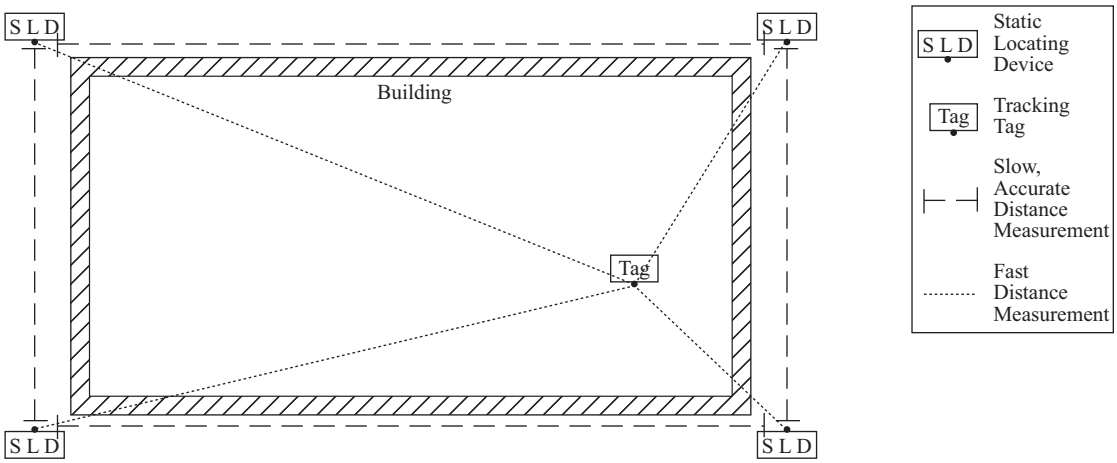

Figure 1.0.1 Diagram of goal system

A final implementation of a system using this tracking system is envisioned as shown in Figure 1.0.1. A small team of technicians or engineers sets up a number of static locating devices around a building in which tracking is desired. These devices do not move after placement, and they can be placed inside the edges of the building or spaced around it externally. The static devices are capable of estimating the relative positions of all the other static devices. It is desired that tracking could be performed by this system in building sizes ranging from that of a house up to a hospital, so that a small number (e.g., between 4 and 8) of the static positioning devices are required. All devices communicate via carefully chosen radio frequencies. Once the static locating devices have been placed, objects to be tracked are assigned small transceiver tags. These tags are able to estimate

the distance between themselves and the static locating devices. The tags should be small enough to be worn by a person or affixed to equipment. The eventual goal would be to create tags small enough that a person could wear several of them at various locations on their body, allowing for highly detailed tracking capability.

This leaves two problems which need to be solved. The first is the process of finding the relative positions of all the static positioning devices after they have all been placed. This can be a rather time-intensive process if necessary; on the order of a few minutes up to an hour. The critical goal here is for the system to be able to identify the relative positions of all the static locating devices as accurately and precisely as possible. The second problem is then to estimate the distances between the static devices and the tracking tags. This process should be very quick and highly accurate. It is desired that each tracking tag can determine its position several times per second (perhaps up to 50 or 100 Hz).

1.1 Uses of indoor location

The ability to accurately locate the relative positions of objects indoors could enable a variety of possible location-aware applications. One such use would be full-body tracking of an individual as they move through a building or other space. This system would employ multiple high-precision tracking tags to record motion and orientation of every limb. It could potentially be used in such varied areas as movie and video game motion tracking, athlete analysis and training, medical procedures, remote surgery, physical therapy, dance teaching, and other types of training. An indoor/outdoor version of the system could be implemented that assisted in construction, machinery automation, and surveying. Such a system would operate as a multi-environment version of GPS, but with much greater accuracy.

A further potential use of an indoor tracking system is that of asset tracking in offices or warehouses. Such a system could enable a person to locate a needed coworker. It would also be capable of giving turn-by-turn directions to a user that was searching for equipment in an office or goods in a warehouse. This type of system could be deployed in hospitals, so that administrators and personnel could efficiently locate doctors, nurses, patients, and crucial medical supplies.

Another potential application for a local positioning system would be for emergency responders or military building-clearing. A system used in this way would allow tracking of personnel as they move through buildings. It would help eliminate any potential redundant searches of rooms by

indicating which rooms have already been examined. Should a team member be incapacitated inside the unfamiliar building, the system could direct backup to the room in an efficient manner. The application in a military building-clearing exercise would be almost identical; allowing the fire team to check for enemies quickly and efficiently without repeatedly clearing rooms. Furthermore, individuals on the fire team could easily locate teammates and reduce the chances of mistaken identity that lead to friendly fire. In either of these applications software could be written to analyze the teams' movements through the unfamiliar buildings, and thus generate a map of the buildings on-the-fly. Such a map would be useful for after-action-reviews and could be used to generate scenarios for use in training new recruits.

A final set of feasible uses for accurate indoor location systems is that of entertainment and advertising. In the entertainment segment, new forms of user interfaces involving 3D motion tracking could be put to use in video gaming systems or the computer industry. In advertising, a mall property could send targeted ads to a user, based on where the user is located in the mall. For example, clothing stores could advertise specials as users entered competing stores. More potential applications for an accurate location system have been compiled in [8].

1.2 Approaches to accurate positioning

Modern methods for positioning first arose in the 1920's with the advent of radionavigation. These systems used radio waves broadcast from standard land-based antenna towers to allow users to get a rough estimate of the direction of shore-based transmitters when they were in range. After finding the headings of several transmitters, it would be possible for a user to calculate his approximate location. These types of navigation systems were the earliest forerunners to fast and accurate positioning; they offered much greater accuracy and speed than the more traditional use of compasses and sextants. The remainder of this section discusses more advanced methods of locating.

1.2.1 GPS

One well-known location system is Navstar GPS, the Global Positioning System. The concept of using satellites for navigation was developed shortly after the launch of Sputnik 1 in 1957, when scientists discovered a Doppler shift in the satellite's telemetry messages. This shift varied from blue to red (while the satellite was in the sky above the observing antenna) in a predictable pattern based on where in the sky the satellite was located. Beginning in the 1960's, branches of the U.S. military

began working on competing designs of space-based navigation systems. In the early 1970's the separate branches pooled their research to create a concept for the modern GPS system. The rest of the 1970's and most of the 1980's were spent in testing, launching, and upgrading the satellites. By 1995, the GPS satellite constellation was completed with the launch of the 24th Block II satellite [10].

1.2.1.1 How GPS works

GPS operates on a pair of frequencies in the L-Band: L1 (1575.42 MHz) and L2 (1227.6 MHz). Civilian GPS receivers use only the signal on the L1 frequency, since the data on L2 is encrypted. Encoded on the L1 frequency via binary phase-shift-keying is a 37,500-bit data message, transmitted at a data rate of 50 Hz. Transmission of one copy of this message takes exactly 12.5 minutes. The data in the message includes, but is not limited to: the current time as indicated by the atomic clock on the satellite, correction information to allow a GPS receiver to synchronize the time signal with UTC time, ephemeris data used to accurately locate the position of the satellite in its orbit, data to help correct inaccuracies caused by disruptions in the ionosphere, and an almanac that tells where all of the GPS satellites should be. The length of time required to receive this message is one of the factors that causes a GPS receiver that has been newly purchased or turned off for several weeks to take so long when first acquiring its position.

The 50 Hz data message used by GPS is not fast enough for use in accurate positioning. To create a signal with enough accuracy to be used in positioning, the 50 Hz message is cut into "chips" at a rate of 1.023 MHz. The chipping is performed by XOR-ing a 1023-bit "chipping code" (generated at the "chipping rate") with the 50 Hz data signal. This implies that each data bit is cut into 20,460 chips. The chipping code is a pseudo random code carefully selected to be different for each GPS satellite. They are referred to as Gold codes, and there are 37 unique codes selected so that all of the satellites can transmit on the same frequency without interfering with each other. A GPS receiver knows all of the Gold codes, and it monitors the L1 frequency for signals that are encoded with any of them. Essentially, the GPS receiver generates its own internal data message bit and its own copy of the chipping code for each satellite and then runs the received signal through several parallel correlators (one for each satellite). This allows the receiver to use a single antenna and radio to monitor all the satellites at once.

Since the chipping code runs at 1.023 MHz, each chip lasts approximately 0.978 μsec. Assuming that a user's GPS receiver is perfectly synchronized with the GPS satellites, the receiver then measures how many chips it must delay its own chipping code in order to get a match. Taking into account the speed of light (approximately 300 million meters per second), each chip is approximately 300 meters long. At first glance, this does not appear to be accurate enough to allow the accuracy seen in modern GPS receivers. In actuality, receivers are capable of calculating the signal shift down to 1% of a single chip. Once the GPS receiver knows the time offset between its internal clock and the clocks of 4 or more GPS satellites, it is a relatively simple matter to calculate the distance to each satellite, and then to use geometry to calculate its position [3].

1.2.1.2 GPS Strengths and Weaknesses

GPS performs well in its operating domain. Using low-power generated from antennas that are more than 20 km away from the receiver, basic GPS receivers are capable of finding their position almost anywhere on the globe with an accuracy of $\pm$15 m. With enhancements such as WAAS (Wide Area Augmentation System) and EGNOS (Euro Geostationary Navigation Overlay Service), this can be reduced to an accuracy of $\pm$5 m. This provides enough resolution for navigation of anything on the land, sea, or air. It can be useful in surveying as well. No other system to date has had this much accuracy over such a broad area.

However, like any real-world system, GPS has its weaknesses. The largest problem for this thesis's application is that GPS does not work indoors. Typical roofing materials weaken the signals from GPS satellites to an unusable power level. Another problem is that the electromagnetic properties of the ionosphere change over time, due to interactions with the ionizing forces of the sun. This atmospheric layer causes the paths of the signals from all GPS satellites to bend, thus making it take longer for the signals to reach any receivers. Yet more inaccuracy is caused by unforeseen drifts in satellite orbits and errors in the atomic clocks on the satellites. Finally, the effects of multipath reflections can cause signals to take longer to reach the receiver, thus introducing even more error. GPS is not a viable solution to the problem of indoor location.

1.2.2 RFID

Another well known tracking technology is radio-frequency identification (RFID). In an RFID system, every object to be tracked has a tag affixed to it. A scanner is then used to locate this tag.

The scanner sends out a query over a specified radio frequency, and the tag then responds with an answer (if it is in range). The tags may be actively powered by a source such as a battery or a wall outlet, or they can be passively powered. A passively powered RFID tag receives all the energy it requires to send a response by leeching the power from the signal that was used to query the tag. A typical response to an RFID query contains information that describes the object to which the RFID is attached.

There are many types of RFID systems in use today. One simple example is the "1-bit transponder" seen in many electronics retailers. Tags used in these systems can only contain enough information to indicate whether or not they are positioned within a scanner's range of operation. Tags are affixed on high-value products, and scanners are placed by the exits to the retailer. Anyone who attempts to take an unpurchased product through the door sets off an alarm. The tags are typically deactivated or removed by the clerks at the cash registers.

More typically, RFID tags can contain anywhere from a few bytes up to several kilobytes of data. Data capacities in this range allow RFID systems to store more information about the items to which they are affixed: data such as a short description, a serial number, or a time stamp of when the tag was attached. Different manufacturers create RFID systems using a variety of tag power sources, modulation schemes, and frequency ranges. Some tags have read/write capability, while others are read-only. The wide selection of RFID systems adds complexity to choosing hardware that could work well as a position tracking system [2].

RFID systems are typically deployed in a barcode-like manner. Large retailers require vendors to affix RFID tags to crates of their goods. These crates can then be wirelessly examined by distributors, and the information contained within the tags can be used to speed the distribution process [16]. RFID scanners would typically be seen at entries and exits to rooms. Centralized computer software would then be able to track the contents of every room by incrementing and decrementing counters as crates passed through the entrance and exit portals.

The use of RFID systems for accurate position tracking is a more recent development. One example system, LANDMARC, is described in [12]. In the LANDMARC system, RFID scanners are installed at regular intervals throughout the tracking area. These scanners regularly poll tags that are affixed to tracked objects. After each scanner receives a response from a tag, a data packet containing the tag ID, the scanner ID, and the signal strength of the received signal is sent to a

central server for processing. In more primitive systems, this information would be used to determine an estimate of the tracked objects' position, based on trilateration. The LANDMARC system adds another layer of complexity, however. In addition to the RFID scanners, there are a number of fixed *reference tags* placed throughout the tracking area in known locations. These reference tags also respond to scanner queries with their ID and position. Information about the reference tags is also sent to the central server to aid the positioning algorithm in accurately locating the tracked objects. In the LANDMARC system, 4 RFID scanners and approximately 1 reference tag per square meter enable the system to track objects with a maximum error of 2 meters and an average error of approximately 1 meter.

RFID systems work well in portal-type tracking systems. Placing networked scanners around doorways allows computer systems to track the location of tracked objects to a large granularity. RFID scanners can be adapted for use in more precise tracking systems by measuring the received signal strength of query responses. However, these systems do not scale well to having large numbers of devices. Each tracked tag must be polled and respond separately. With every tag added to the system, the amount of time required to calculate the position of all tags increases. Also, the tags must be increasingly complex as more are desired in a tracking area. More bits per tag are required if the system is to be able to distinguish them all. These issues combine to make RFID tracking technology ill-suited for use in a highly-scalable positioning system.

1.2.3 Ultrasound-based location systems

There have been a number of researchers looking into the problem of indoor position finding in recent years. Some of these researchers have taken to using ultrasonic emitters for rangefinding indoors. One system of this type is the "ORL ultrasonic location system", designed by staff of the University of Cambridge Computer Laboratory and the Oracle Research Laboratory [18]. This system requires that the trackable area be instrumented with a wireless radio transmitter and a network of ultrasonic detectors, whose locations are known. Each object that is desired to be tracked is fitted with a radio receiver and an ultrasonic transmitter. Every 200 msec, the master radio transmitter addresses a single receiver and instructs it to emit an ultrasonic pulse. All of the ultrasonic receivers then time how long it takes for the ultrasonic pulse to reach them. With enough readings, it is possible to trilaterate the location of the ultrasonic transmitter. This system only

allows a 5 Hz update rate, spread amongst all of the tracked devices. Therefore, a system tracking 5 devices would only have a 1 Hz update rate per device.

Another system using ultrasonic emitters and detectors is the "Cricket location-support system", developed at the MIT Laboratory for Computer Science [14]. MIT's ultrasonic location system uses a different strategy for position finding. The Cricket employs an infrastructure of beacons placed at very specific locations in each room where tracking is to occur. These beacons *simultaneously* emit a radio signal (containing information identifying which room the beacon is located in) and an ultrasonic pulse. Each object to be tracked, then, has both a radio receiver and an ultrasonic detector. The objects receive the radio signal, turn on their ultrasonic receivers, and then time how long it takes to detect the ultrasonic pulse. The tracked objects then calculate the distance to the beacon. However, instead of placing a large network of beacons in the environment, the Cricket system relies on as few beacons as possible. The granularity of the Cricket system is room-sized, and therefore it only tries to determine what room, or as a best case, what part of a room each object is located in. It compares the distance between different beacons, and decides it is in the room with the beacon that had the shortest distance. This system does not have the desired accuracy, and its update rate is too low for the suggested applications.

Yet another ultrasonic system is the Bat system, implemented by the AT&T Laboratories in Cambridge and the University of Cambridge [1]. This system works on the same premise as the ORL ultrasonic location system. Each person or object to be tracked carries a small device, called a "Bat", that emits an ultrasonic pulse when it is addressed over a wireless connection. The Bat system is larger than the ORL system, with a grid of 750 receiver units, 3 separate tracked areas, and a total population of 200 Bats being tracked. Each tracked area has a radio transceiver that can simultaneously track 3 Bats at an update rate of 50 Hz.

Another ultrasonic location system is the Dolphin system [6]. This system's original implementation was initially more interested in user privacy: allowing the object being tracked to determine its own position, and only sharing that information with the infrastructure if the user wants to. To that end, it operated similar to an ultrasonic version of the GPS system. A grid of ultrasonic emitters was laid out in the ceiling of the area to be tracked. Each emitter sent its ranging signal at a well defined time, and all of the emitters were synchronized over some type of data link. Each tracked object could then receive the ultrasonic pulses and compute its position, based on the time-of-arrival

of each signal and the known locations of the transmitters. To allow more emitters to transmit simultaneously, each emitter transmitted its ranging signal encoded with a Gold code, much like GPS. This required broadband ultrasonic transmitters and receivers, but allowed the system to achieve a much higher update rate. Further research on the Dolphin system has implemented an additional non-private mode that functions similarly to the Bat system [5]. The main innovation of the Dolphin systems is the use of broadband ultrasonic devices rather than narrowband equipment.

A final ultrasonic tracker is the system implemented by the University of Bristol [11]. Their system is much more advanced, allowing for a small grid of ultrasonic transmitters (typically, 4 in each room-sized tracking area). In the University of Bristol implementation, setup and calibration is reduced to a simple process of placing transmitters in random, fixed positions; moving a receiver throughout the tracked space for a minute or two; and then running a simultaneous localization and mapping algorithm adapted from robotics research on the data from the tracked receiver. Also, with advanced use of Kalman filters and precise definitions of system parameters (such as the time intervals between the chirps from each transmitter), the researchers have been able to eliminate the use of radio waves from their system entirely; enabling a very small, ultrasound-only receiver to be placed on objects to be tracked. This is probably the most successful ultrasound system to date, with an accuracy of 1.6 cm. Processing on the mobile units takes 33 percent of the CPU cycles on a 200 MHz StrongARM processor, and the autocalibration procedure requires approximately 5 minutes on a 1.6 GHz single-core CPU.

Ultrasonic positioning systems show promise for single-room tracking. They can provide high-accuracy tracking of multiple objects with a fast update rate. However, the amount of instrumentation and calibration increases dramatically as the size of the tracked area increases. This makes ultrasonic tracking unfeasible for large spaces. Another problem with ultrasonic waves is that they reflect off of walls and floors. Because of this, a system placed in one room cannot track objects in another room. These factors contribute to making ultrasound a technique which cannot be used as a fast-deploy tracking system on a building-sized scale.

1.2.4 Radio-based location systems

Research into radio-based methods for indoor positioning has not been as active as that in ultrasound. Most researchers cite problems of mulitpath interference and unpredictability of radio waves

around indoor obstacles as reasons why indoor location technologies should focus on ultrasound. However, there have been a few notable implementations of radio-based location systems.

1.2.4.1 Region / cell-based

The first type of system is a pico-cell method, exemplified by an implementation named P-POINTS (Personal Positioning INformaTion System) that was developed in Japan [15]. It uses small transmitters called E-markers that are placed on objects of interest. Each E-marker broadcasts its position and some text describing what it is that is being marked. The transmission range of these markers is very small (less than 10 meters). Unfortunately, this system is designed more to help objects of interest broadcast their availability (printers, restaurants, office equipment), than to allow precise positioning.

1.2.4.2 RSS trilateration

Another radio-based indoor location system is exemplified by researchers at the Universiti Technologi Malaysia [9]. This type of system requires that an infrastructure of radio transmitters be set up around the space to be tracked. These can be transmitting any type of signal: AM, FM; Analog, or Digital. Implementations typically use 802.11x access points, due to their ease of use. Each object to be tracked measures the Received Signal Strength (RSS) of all of the transmitters. Then, using a model of the environment (that includes the signal strength loss due to distance, number of walls, and the thickness of walls), the receivers attempt to calculate the distance to all transmitters. After estimating the distances, it is possible to locate the receiver through trilateration.

This system fails to accurately track indoors because radio waves emitted in indoor areas typically do not follow simple models. Effects such as attenuation, reflection, diffusion, and multipath all cause radio waves to differ from their expected behavior. Due to the abundance of obstacles indoors, it is impossible to accurately model the expected behavior of a radio wave. As a result, the accuracy of these systems is not high enough to meet the stated design goals.

1.2.4.3 RSS location fingerprinting

A third type of tracking system improves on the accuracy of the RSS trilateration method, but at the cost of increased calibration difficulty. These methods involve the use of "fingerprinting", where a technician carefully walks throughout an entire tracking area with a laptop, recording the

signal strength from all visible wireless transmitters. A database of these signal strengths (and their associated positions) is built. To calculate location, trackers measure the signal strengths to all transmitters and search through the database to find the closest recorded match to the live readings. One such system is commercially available as the Ekahau Real-Time Location System (www.ekahau.com). These systems require an extensive site survey and calibration procedure before they can be operated. Also, their accuracy is cited as 1-3 meters, which is not accurate enough for precise tracking. Finally, the granularity of the positioning is heavily dependent on the granularity of the calibration. Attempts have been made to increase the accuracy of these types of systems by adding Kalman filters and other techniques, but they do not address the chief concern of intensive calibration [4].

Typically, a location system using location fingerprinting uses 802.11x (also known as WiFi) signals. However, other recent research has investigated the use of GSM cell phone signals as an alternative [13]. Researchers cite the often-changing layout and equipment of 802.11x transmitters versus the relatively fixed configurations of cell phone transmitters as reasons to opt for GSM instead. The less often the transmitters change, the less often a site survey is required. Furthermore, since cell phones operate in a licensed frequency band (compared to 802.11x which is in an unlicensed band), they are less likely to receive interference from neighboring transmitters on the same frequency. GSM fingerprinting methods still suffer from the problems of calibration and inaccuracy. They are also not suited for simple indoor tracking systems.

1.2.5 Other location systems

Other subclasses of indoor positioning systems exist as well, but they share many of the same drawbacks as methods already examined. One such system is the Active Badge [17], which uses infrared emitters and detectors to determine what room a user is located in. More accuracy is needed than simply knowing in which room a tracked object is located. Another class of systems uses a network of video cameras to compute which areas of a floor are occupied by objects [7]. While these are effective at achieving the accuracy we desire, they do not generally uniquely identify specific objects. Problems can also occur when objects move close together, as they become hard to discriminate. Furthermore, such camera systems require difficult installation and calibration procedures, and are vulnerable to changes in lighting conditions. Other types of position tracking systems and methods are mentioned in [8].

1.2.6 Summary of tracking systems

Table 1.2.1 shows a list of the best tracking technologies of each type reviewed in this introduction. The last entry in the table summarized the design goals for the envisioned tracking system. As can be seen here, none of the previously researched methods for position tracking works well enough to satisfy these goals. The ultrasound and camera-based tracking methods are closest to the desired accuracy, but the instrumentation and calibration requirements of these systems along with their single-room restrictions leave them unsuitable for the goal system.

System	Coverage	Accuracy	Deployment Difficulty
GPS	Worldwide (outdoor only)	10 m	Already available
RFID	Several Rooms	2 m	Extensive instrumentation
Ultrasound	Single Room	2 cm	Extensive instrumentation
Radio (fingerprinting)	Whole Building	3 m	Extensive calibration
Infrared	Whole Building	room-sized	Extensive instrumentation
Camera	Single Room	2 cm	Extensive inst. and cal.
LPS	**Whole Building**	**$<$ 1 cm**	**Minimal inst. and cal.**

Table 1.2.1 Summary of indoor tracking methods

The radio fingerprinting approach yields the highest accuracy of the radio-based technologies, but the time-consuming calibration of systems using this approach also makes it inappropriate for the desired application. New concepts for indoor rangefinding methods must be investigated if the goal system is to be achieved.

1.3 New approach

A different approach to this problem is to be examined. To date, most of the research for indoor position tracking has focused on using time-of-flight measurements from ultrasound devices or received signal strength of radio transmitters. Neither of these approaches are optimal: ultrasound requires a large amount of instrumentation if a highly accurate system is to be designed (and typically only covers small areas), and signal-strength based approaches are either inaccurate (due to hard-to-model interference due to walls and other obstacles) or they require large amounts of calibration time. This thesis focuses on using time-of-flight measurements of radio waves to perform accurate, high-rate positioning with a minimum amount of instrumentation or calibration. In short, it examines results from testing of three potential radio-based methods for indoor location. The first method uses a digital, packet-based approach in an attempt to measure the time-of-flight. The remaining

methods both attempt to measure the phase difference of a periodic wave in the radio spectrum. The first radio method uses the phase of the raw carrier wave, while the second sends a wave over the carrier (any modulation, be it AM, FM, or PSK, can be used) and measures the phase difference of the encoded signal.

CHAPTER 2

METHODS

This chapter covers in greater detail the actual implementation of each test. Each of the three sections will describe the reasoning for attempting one of the three methods. They will each also give detailed explanations of the theory behind each proposed system. The sections delve particularly into the model of each system, the software written for each test, and the method by which each test was performed. In relation to the ultimate goal of the project, the first method (802.11g rangefinding) would be a candidate for the slow, high-accuracy distance estimation between static locating devices. The other two methods were tested as candidates for the high-speed position estimation of tracking tags.

For each system, it was assumed that tracking an object in one dimension would be sufficient to determine the feasibility of a system. Once it is possible to measure the distance from the tracked object to a reference point, it would then be easy to measure distances to multiple reference points and thus be able to generate two- or three-dimensional tracking data. Therefore, each of these experiments was designed to track only the one-dimensional distance between a tracked object and a reference point.

2.1 Packet-based 802.11g rangefinding

802.11x networking is a set of mature packet-based wireless communication protocols. 802.11g is one member of this set which has become the standard wireless networking protocol for many organizations. It was chosen for this test due to availability of hardware and ease of programming. The 802.11g protocol operates in the 2.4 GHz ISM[1] frequency band, dividing these frequencies into 14 overlapping channels. In the United States, only the lowest 11 of these channels can be used, due to FCC regulations. This test was performed over a WiFi link from a laptop (Toshiba Portégé R200) to a standard consumer-grade access point (Linksys WRT54G v.3) on channel 1 (center frequency of 2.412 GHz).

[1]See abbreviations and definitions in Appendix A

2.1.1 System model

It was theorized that as the distance between the laptop and the access point increased, the time required to successfully complete a data transmission between the two devices would also increase. For the small differences in distance being studied, it was predicted that the differences in time would be small. Therefore, it was determined that a packet should be sent from the laptop to the access point, and the laptop should then wait for a response. Multiple iterations of this process were completed (on the order of 1,000's), and they were timed as a group. That is, a timer was started before the first packet was sent, and it was stopped after the last response was received.

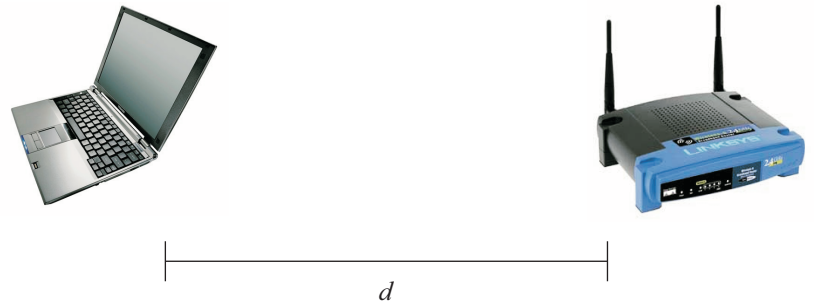

Figure 2.1.1 Example WiFi test setup

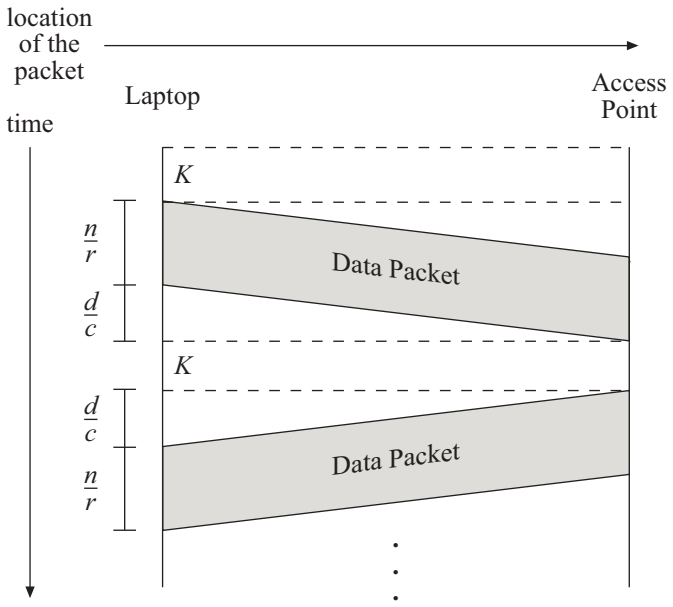

Figure 2.1.2 Signal timing between laptop and access point

The diagrams shown in Figures 2.1.1 and 2.1.2 will help illustrate the model of the WiFi system. Each packet was 1000 bytes ($n = 8000$ bits) long and was transmitted at the 802.11g bitrate of $r = 54$ Mbit/sec. Ignoring protocol overhead, this means that each packet took at least 148 μsec to finish transmitting. The packet then traveled the distance between the two devices. This took $\frac{d}{c}$ sec, where c is the speed of light ($3 \cdot 10^8 \frac{\text{m}}{\text{sec}}$). Under a maximum distance estimate of 100 m, the travel time of the packet was estimated to be $\frac{1}{3}\mu$sec. Once the packet arrived at the access point, it was then retransmitted back to the laptop. This transmission should have taken the same amount of time. This process was then repeated. In addition, there was an approximately constant processing time required at both the laptop and access point between the time a packet arrived and it was retransmitted (K). By assembling these segments of time required for a signal to travel, Equation 2.1.1 shows the estimated time required for a packet to travel from the laptop to the access point and back. This process is hereafter labeled a "bounce".

$$t_{bounce} = 2 \cdot \left(\frac{n}{r} + \frac{d}{c} \right) + K \tag{2.1.1}$$

Plugging in values, Equation 2.1.1 becomes:

$$t_{bounce} = 2 \cdot \left(\frac{8000}{54000000} + \frac{d}{3 \cdot 10^8} \right) + K$$

$$t_{bounce} = \frac{2 \cdot d}{3 \cdot 10^8} + K + \frac{1}{3375} \quad \text{sec} \tag{2.1.2}$$

Solving for d:

$$d = \frac{3 \cdot 10^8}{2} \cdot \left(t_{bounce} - K - \frac{1}{3375} \right) \quad \text{m} \tag{2.1.3}$$

Equation 2.1.3 shows the anticipated relationship between the distance between the two WiFi devices and the amount of time required for one bounce to complete. In practice, K is hard to estimate. Also, neither K nor r is actually a constant. To counteract the randomness of these variables, 1,000 to 100,000 bounces were sent back-to-back and timed as a group. The mean time per bounce was then calculated and labeled t_{bounce}.

2.1.2 Software

The software used in this experiment was a command-line program for Windows written in Visual C++ 2005 using the .NET framework. First, a 1000-byte memory block is filled with data to be sent in the bounce packet. A stopwatch is then created and started. The program next enters a

loop that sends the bounce packet to the wireless access point as an ICMP Echo Request (Ping) for the specified number of times (generally 1,000 - 100,000). Any bounce that fails is noticed after a 5-second timeout period and an error message is printed to the screen. A status report listing the total number of bounces sent, the total time, total number of failed bounces, and number of bounces per second is printed to the screen periodically, for user feedback. After the specified number of bounces has been sent, these statistics are printed again and the program exits.

2.2 Radio carrier-wave rangefinding

It was hypothesized that a simple analog radio system could be capable of providing the desired positioning capability. The proposed system would detect the phase difference between carrier waves sent from 2 antennas. It would use this measurement to determine the position of the receiver with respect to one of the transmitting antennas, using a time-of-flight approach. A frequency of 21 MHz was chosen for these experiments. A 21 MHz radio wave has a wavelength of approximately 15 m, which was the desired size of the tracking area when the experiments were performed.

2.2.1 System model

This set of experiments was designed to test the feasibility of using 2 transmitters and 1 receiver to measure how much closer the receiver was to one transmitter, rather than the other. The goal was to achieve a system where the 2 transmitting antennas were no more than 1 wavelength apart and the receiver was free to roam between them. By combining two of these systems on different frequencies and at right angles, it was thought possible to achieve a 2-dimensional tracking system. The actual position detection would be performed by detecting the phase difference between the two transmitted carrier waves. Since both transmitters are influencing the received signal on the same frequency, the effects of the two carrier waves add together and the receiver only picks up a single wave. This effect is shown for two signals that are in phase and out-of-phase in Figures 2.2.1 and 2.2.2. As can be seen in the figures, the receiver will see a carrier wave that is anywhere from twice as strong as the original wave to nonexistent, depending on the phase of the two received waves. In theory, this phase difference could be measured as a calculation of the RMS power of the received signal. When the receiver is equally distant from the two transmitters, the received waves will be perfectly in-phase and should add. When the receiver is closer to one transmitter than the other, the carrier waves will interfere with each other. When the receiver is as close as possible to one

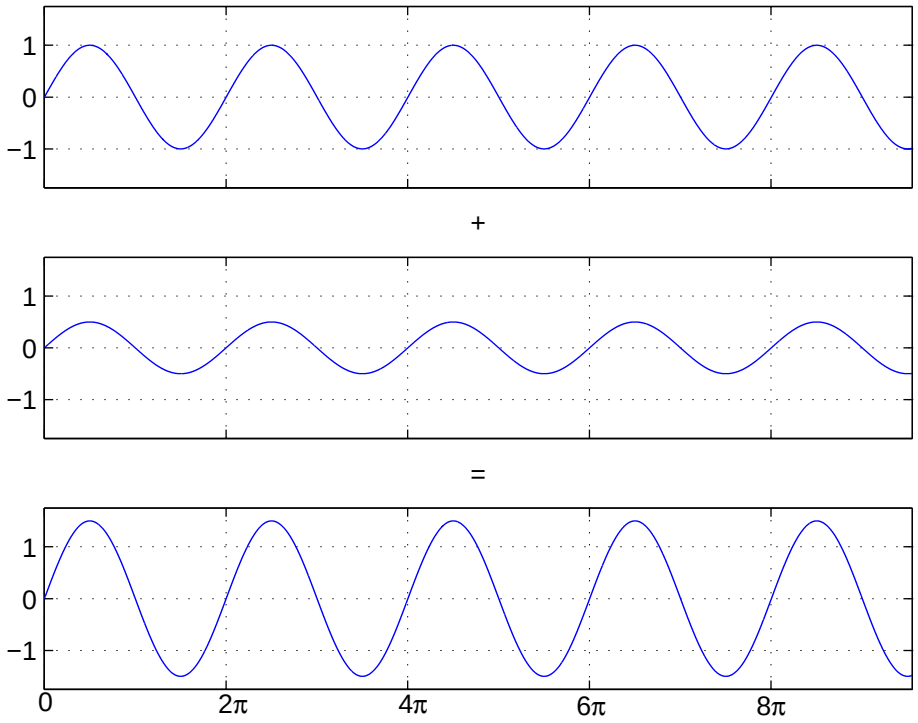

Figure 2.2.1 Addition of two in-phase sine waves

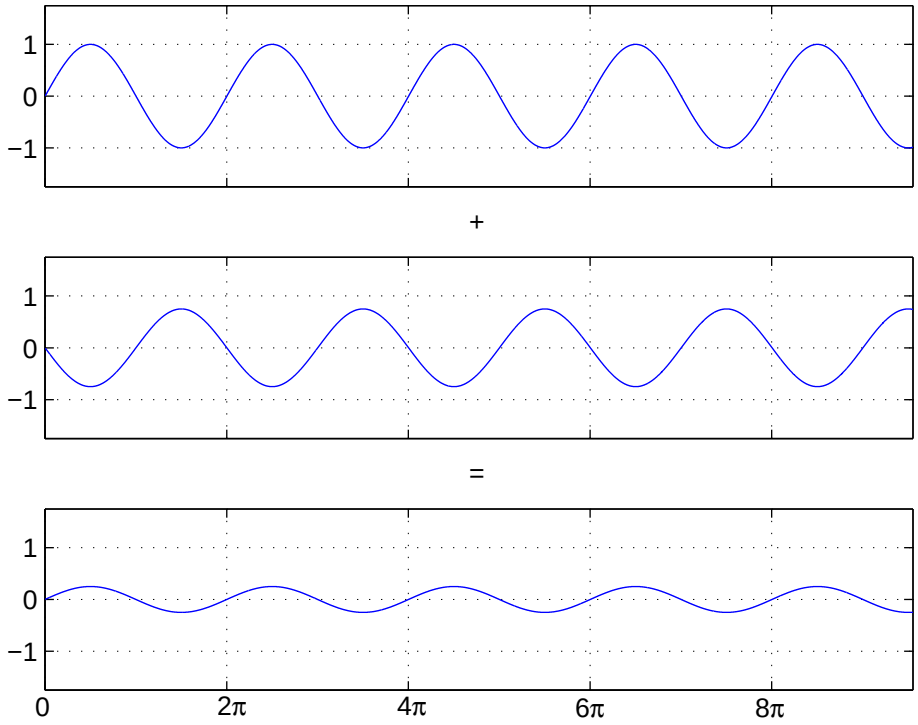

Figure 2.2.2 Addition of two out-of-phase sine waves

transmitter but as far away as possible from from the other, the received signal will be at its lowest strength.

2.2.2 Software

The software component of this experiment was written in LabView 8.0. A virtual instrument was defined that contained all of the RS-232 serial port commands required to control the 3 radios used in this experiment. Yaesu FT-817 transceivers were chosen due to their frequency range and ease of control via serial port commands. The software allowed remote power up, frequency selection, modulation selection, and transmitter keying. A section of the user interface was also designed to allow the user to feed a sine wave at any audible frequency from the computer's sound card to the microphone input of the transmitters. Since the computer did not have any hardware capable of analyzing radio signals, a digital oscilloscope was used as the receiver. This scope had several built-in math functions and a serial port interface. The software written for this experiment contained a section of code capable of triggering the oscilloscope's math functions and reading back the results over the data port. The function most used was the RMS power calculator, which determined the RMS of the sampled waveform. The final element of the user interface was a chart that showed a live view of the last 100 RMS measurements. When the software had performed a test, it saved the collected data to an output file.

2.3 Radio carried-signal rangefinding

A second radio system test was devised. It was theorized that a sound wave signal generated from a single source could be transmitted simultaneously over two radio frequencies from two antennas located at opposite ends of a test range. Both sound waves would then be received at a pair of co-located receiving antennas. Each copy of the sound wave could then be sampled with an A/D converter card and they could be compared digitally. It was hypothesized that as the receivers moved between the two transmitters, the phase difference between the received sound waves would vary monotonically.

2.3.1 System model

The basic formula for a sine wave is $V(t) = A \cdot sin(\theta t - 2\pi k)$. A represents the amplitude of the sine wave, or its maximum height above 0. This function assumes that θ, the frequency, is given in $\frac{\text{rad}}{\text{sec}}$, and its period is 2π radians. k in this function is generally a number between 0 and 1 that

represents the amount of phase shift of the signal. For example, Figure 2.3.1 shows two sine waves with k of 0.0 and 0.1. The θ value for these sine waves is 1. To model the system used in these experiments more accurately, the equation was generalized as:

$$V(t) = A \cdot sin\left(2\pi ft - 2\pi\frac{df}{c}\right) \tag{2.3.1}$$

where f is the frequency of the wave signal (in hertz), t is the number of seconds the signal has been running, d is the distance between the transmitter and the receiver (in meters), and c is the speed of light, $3 \cdot 10^8$ $\frac{\text{m}}{\text{sec}}$. $2\pi f$ to replaces θ from the basic equation, which converts the frequency of the sound wave from Hz to $\frac{\text{rad}}{\text{sec}}$. Figure 2.3.2 demonstrates the effect of varying the frequency of the sine waves. The parameter k from the basic formula is replaced with $\frac{df}{c}$, which yields a fraction from 0 to 1 that scales the phase shift of the signal according to the distance between the transmitter and the receiver. Figure 2.3.3 shows this.

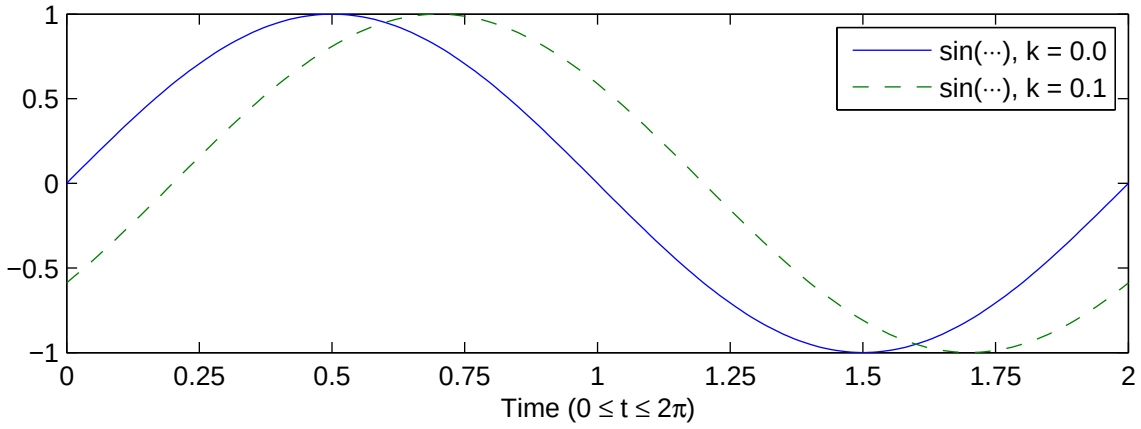

Figure 2.3.1 Two out-of-phase sine waves

In the test setup, one transmitter was placed at either end of a hallway. A pair of receivers was then moved to several positions between the two transmitters. An audio tone of a 3 kHz sine wave was transmitted over both transmitter / receiver pairs. This tone was generated on a computer's sound card, split, and then fed into the audio input jack of each transmitter. The audio output of both receivers were then hooked up to two channels of a 12-bit A/D converter card and each channel was sampled at 25,100 Hz[2] for 1,024 samples per channel. This yielded an average of close to 8.35

[2]The sample rate reported by the A/D card was 24,096 Hz, yet it was observed by experimentation that the actual sample rate for this A/D card was approximately 4.166% higher than the rate reported by the software.

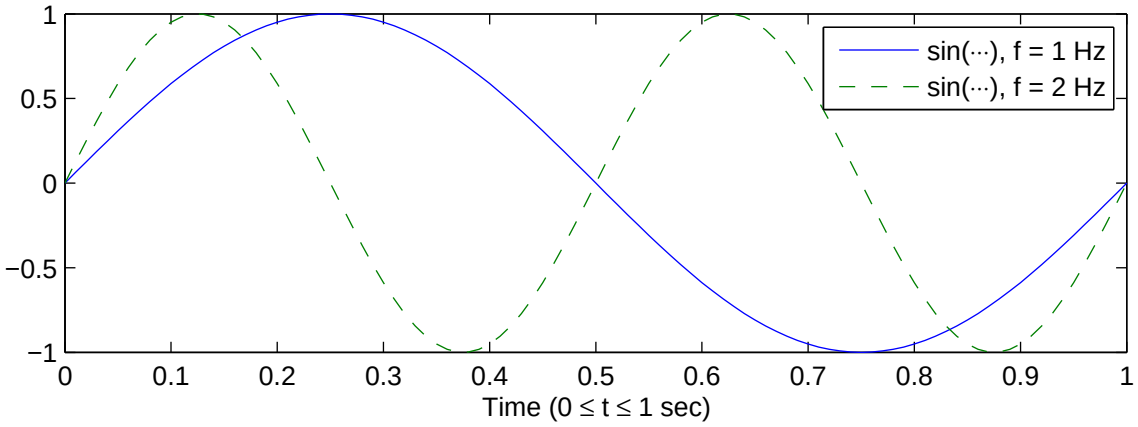

Figure 2.3.2 Sine waves of differing frequency

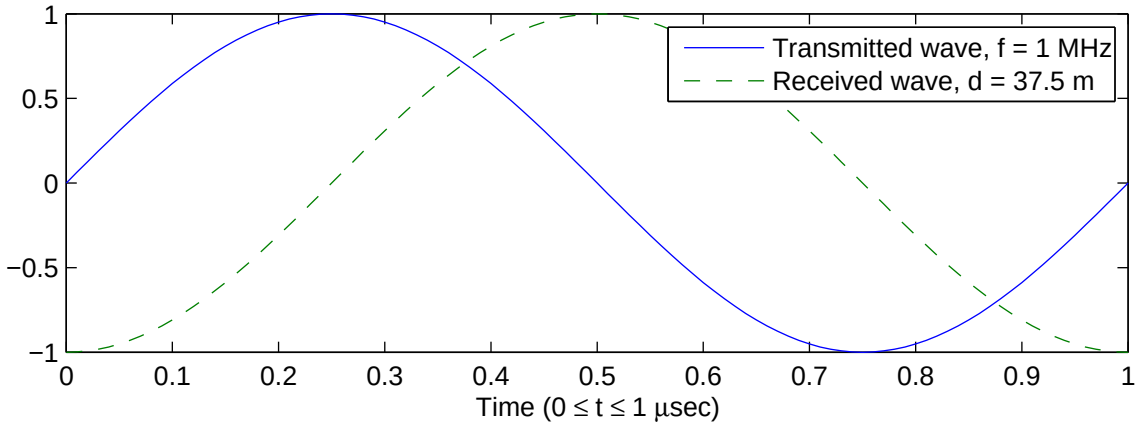

Figure 2.3.3 Phase offset due to transmission distance

samples of each audio waveform per cycle, for a total of approximately 122 cycles recorded from each channel.

It was observed that the amplitude of the received sound waves varied significantly. Figure 2.3.4 shows an example of the voltages of the received sound wave. After a cursory investigation, it was determined that this variation was caused by inherent system noise, and was unavoidable. In order to determine the phase offset between the two sound waves, a method that was minimally effected by the varying amplitude was developed. It was decided that a way to do this would be by finding the phase offset relative to the start of the sampling time for each wave, separately. To find the phase offset of a single wave, every peak of the waveform excluding the endpoints was found. In order to get a better approximation of the peak of the sound wave, a parabola was fit to the local

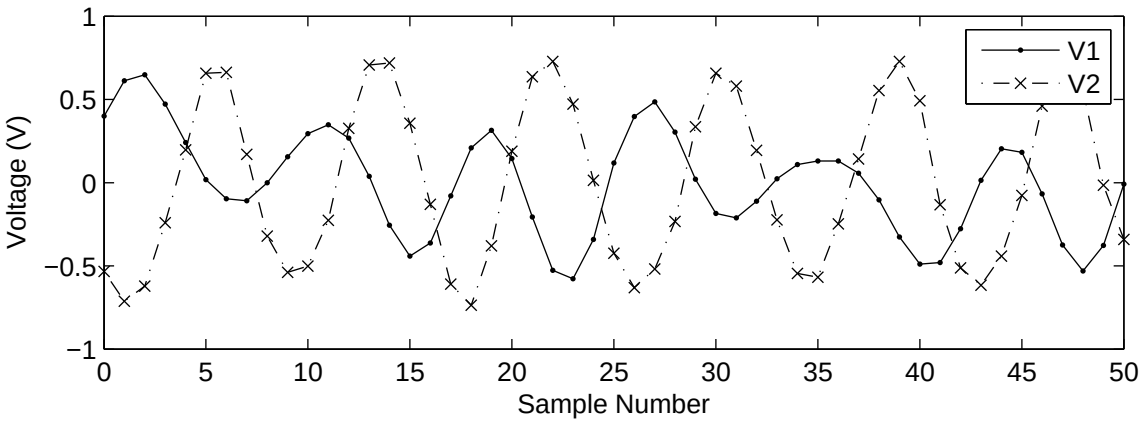

Figure 2.3.4 Example of received sound waves

maximum and the samples located on either side of the maximum. The x-coordinate of the peak of this parabola was then recorded as the location of the peak of the sound wave. The period of the wave was then determined by taking the average of the difference in x-coordinates between adjacent peaks. Next, the phase offset of each peak is calculated and then they are all averaged together. This yields the phase offset of one waveform. The same process is applied to the other sampled wave, and then the two offsets are subtracted.

It was desired to predict the range of phase offsets that would be found in the actual tests. To perform this prediction, equation 2.3.1 was used to derive the waveforms that were expected to be received and sampled. Assuming that the two sound waves are synchronized as they leave the transmitters, the waves as observed by the receiver would be described by:

$$V_1(t) = A \cdot \ sin\left(2\pi ft - 2\pi\frac{d_1 f}{c}\right) \tag{2.3.2}$$

$$V_2(t) = A \cdot \ sin\left(2\pi ft - 2\pi\frac{d_2 f}{c}\right) \tag{2.3.3}$$

where d_1 is distance d between the first transmitter and the pair of receivers in meters, d_2 is the distance $27.7 - d$ (in this experiment) between the receivers and the second transmitter in meters, A is the amplitude of the audio output of 1 V, f is the frequency of the sound wave 3000 Hz, and c is the speed of light, approximately $3 \cdot 10^8 \ \frac{\text{m}}{\text{sec}}$. Substituting in values, this gives:

$$V_1(t) = sin\left(6000\pi\left(t - \frac{d}{3 \cdot 10^8}\right)\right) \tag{2.3.4}$$

$$V_2(t) = sin\left(6000\pi\left(t - \frac{27.7 - d}{3 \cdot 10^8}\right)\right) \tag{2.3.5}$$

The range of d used for these experiments was from 0.3 m to 27.4 m. At the range of 0.3 m, the phase offset of V1 was calculated to be 83.334 μsec and the phase offset of V2 was calculated to be 83.425 μsec. The calculated relative phase offset between the two sound signals was then -0.91 μsec. In the case of the receivers being positioned at the 27.4 m mark, the phase offsets were swapped and the phase offset was 0.91 μsec.

2.3.2 Software

To perform these experiments, two separate programs were written. The first program was written in LabVIEW and allowed control of the radio transmitters via RS-232 serial ports. It provided a user interface that allowed an experimenter to choose frequency and modulation of transmission, turn on or off the transmitter, set the frequency of the sound wave, and turn the sound wave generator on or off. The second program was written in C to control the A/D card and save the data. This program takes the filename where data is to be saved as input, instructs the A/D card to gather 1024 samples from each of two of its input channels at a rate of 24,096 Hz, saves the data to the file, and exits. The number 1024 was chosen because that was the maximum number of samples per channel that could fit in the A/D card's sampling buffer. The sample rate of 24,096 Hz was chosen because that was the closest sample rate to the desired rate of 24,000 Hz that the A/D card would allow. In practice, it was found that even though the sample rate of 24,096 Hz was reported by the A/D card programmatically, the actual sample rate was 25,100 Hz. The cause for this 4% increase in sample rate is unknown.

CHAPTER 3

EXPERIMENTAL RESULTS

The results of the three sets of experiments are presented below. First, the results of the packet-based 802.11g rangefinding method are shown. Next are the results of the carrier-wave method. The final set of results are those from the carried-signal rangefinding method. Each section begins with a view of how the method being tested showed a trend of working as anticipated. However, all implementations have data samples that do not support their trends. After the trends are shown, series of experiments designed to determine the causes of the aberrations are detailed.

3.1 Packet-based 802.11g rangefinding

To perform the experiments based on the 802.11g rangefinding method, a number of individual tests were set up and performed in a site that was relatively free of competing 2.4 GHz radio waves. The site chosen was a residence in a neighborhood that had few neighboring houses or other 802.11g wireless networks. No 2.4 GHz portable phones were known to be operating in the area. Competing Bluetooth transmissions (which also operate at 2.4 GHz) were minimized.

In the packet-based method, it was decided that the distance between the laptop and the access point would be estimated by measuring the number of "bounces" per second and relating that measurement to the distance between the devices. A higher number of bounces per second would relate to a shorter distance, while fewer bounces per second would indicate longer distance.

A test was performed that shows the overall trend of increasing distance between the laptop and router causing a decrease in the bounce rate. This test involved placing the router on the exterior of the residence and carrying the laptop down the street. Data points were collected approximately every 4 m away from the access point. Each data point consisted of timing 1,000 bounces. A diagram of this test is shown in Figure 3.1.1.

Figure 3.1.2 shows the result of this test of the WiFi system. Samples 1-6 of this graph correspond with the segment of the path between the residence and the street. Samples 7-12 correspond with the segment along the first half of the street. Sample 12 was collected with the laptop approximately 50 m from the router. Samples 13-18 correspond with the last segment along the street. Sample 18 was gathered with the laptop positioned approximately 80 m from the router. After sample 18, the

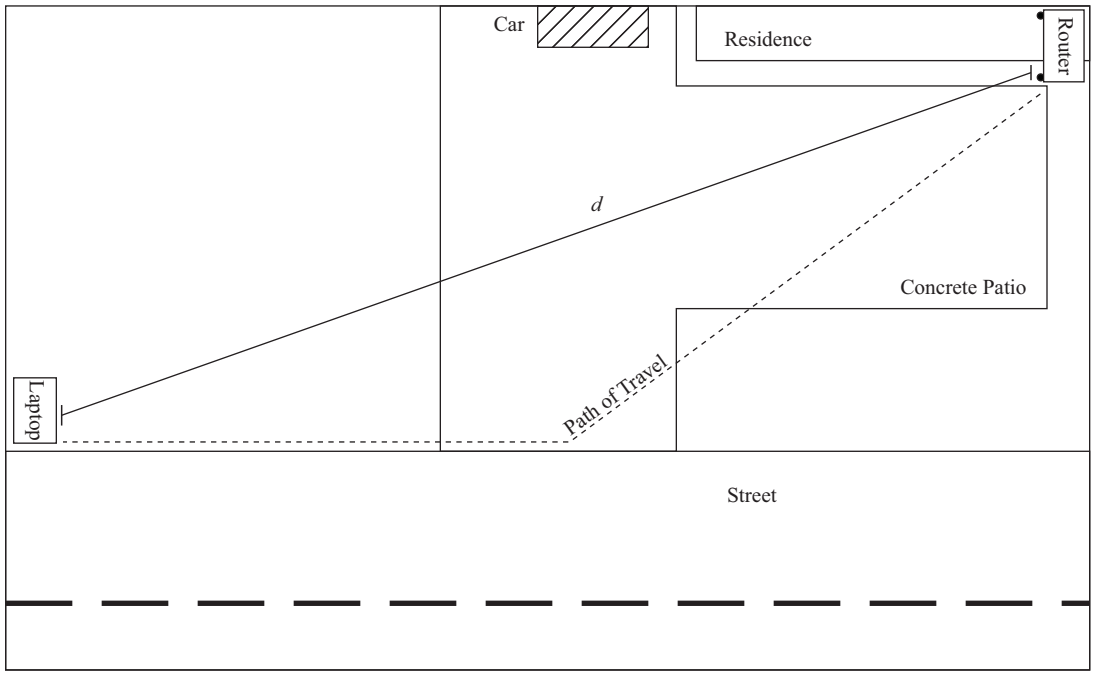

Figure 3.1.1 WiFi trend experimental setup

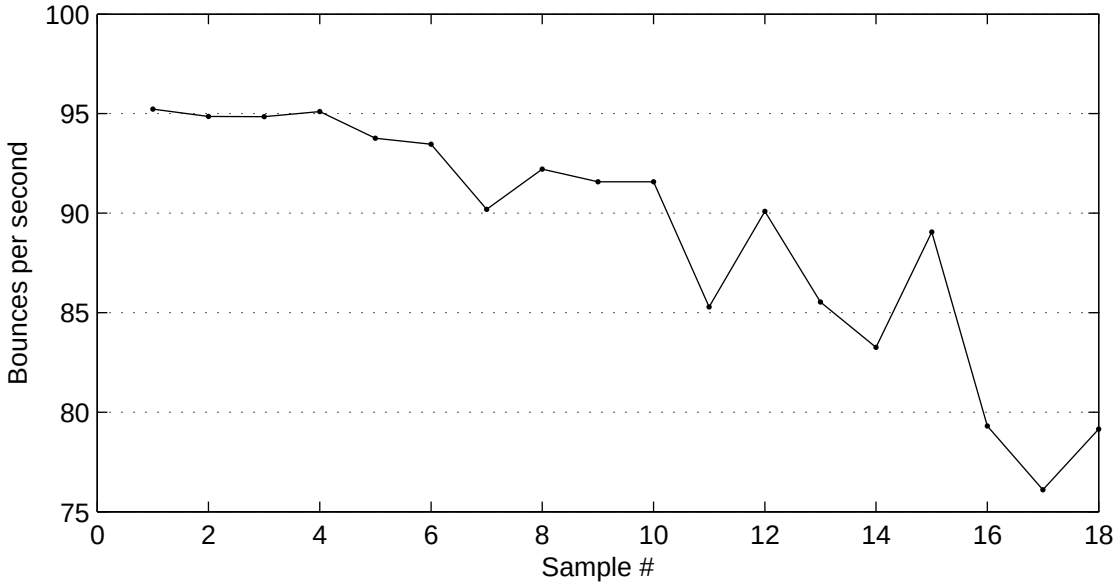

Figure 3.1.2 Results from WiFi trend experiment

two radios suddenly disconnected and no further readings were possible. As seen in the graph, the bounce rate steadily decreases as the laptop moves farther from the router.

Potential Causes of Error in WiFi Rangefinding

Outdoor vs. Indoor Behavior
Single Wall
Multiple Walls
Metallic Masses

Table 3.1.1 Error sources in packet-based rangefinding method

Table 3.1.1 shows a list of potential sources of error in the carried-signal rangefinding method. These error sources are investigated in the following experiments.

An experiment was performed to show how the bounce rate changed between outdoor and indoor locations. In this procedure, the access point and laptop were positioned outdoors 1.5 m apart, increasing by increments of 1.52 m until reaching 7.6 m. At each position, data was collected for five 100,000-bounce runs. This was then repeated indoors. This can be thought of as a close up of the first 3-4 data points of Figure 3.1.2. The positioning of equipment in the indoor/outdoor test is illustrated in Figure 3.1.3.

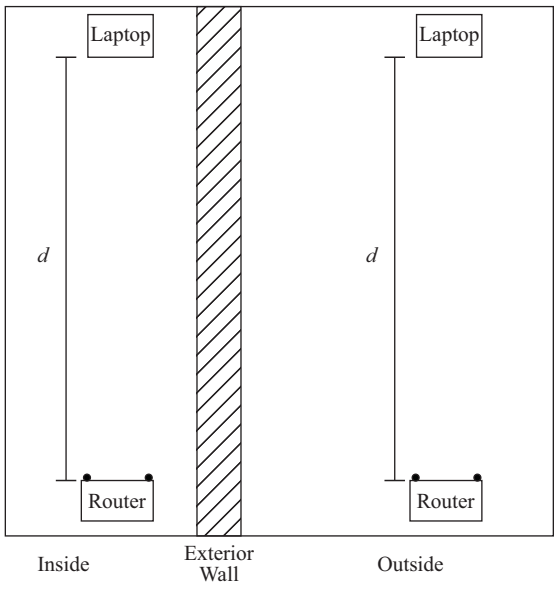

Figure 3.1.3 WiFi outdoor/indoor experimental setup

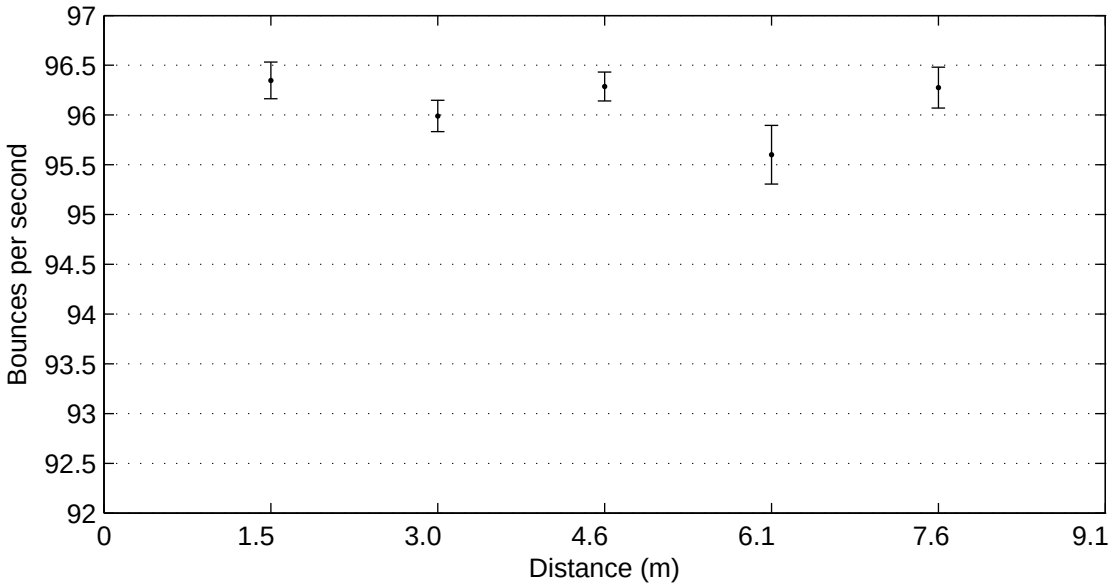

Figure 3.1.4 Results from WiFi outdoor experiment

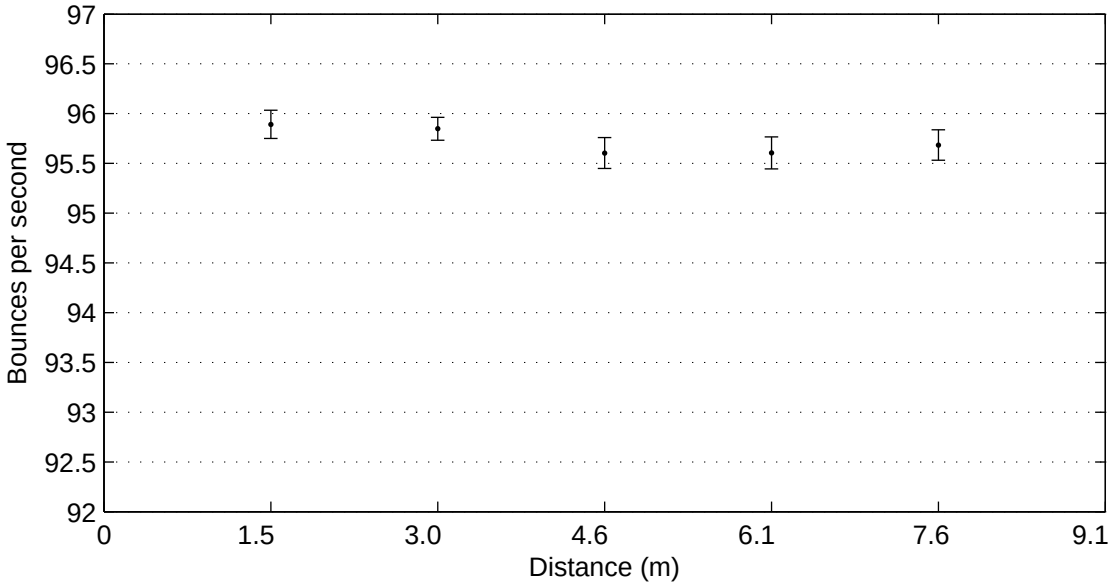

Figure 3.1.5 Results from WiFi indoor experiment

Figures 3.1.4 and 3.1.5 show the results of the outdoor/indoor test. Each data point is representative of the average and standard deviation of 10 samples of 10,000-packet calculations. It can be seen that the trend of bounce rate reduction with increasing separation distance continues in this test, but is not drastic enough to be reliably detected. By comparing the results of the outdoor and indoor tests, it can be seen that the bounce rate was consistently 0.5 bounces per second higher for the outdoor test than the indoor test. This indicates that changes in the operating environment introduces noise and error into the WiFi rangefinding system.

A test was then performed to observe the effect of placing a single interior wall between the access point and the laptop. The access point was positioned with its antennas flush against the wall, and the laptop's starting position was 1.5 m from the opposite side of the wall as shown in Figure 3.1.6. In this configuration, data was collected at increments of 1.52 m, up to a maximum distance of 7.6 m. Data was collected for 2 hours, or approximately 680,000 samples.

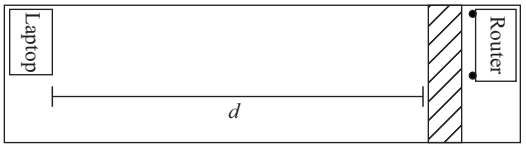

Figure 3.1.6 WiFi single-interior-wall experimental setup

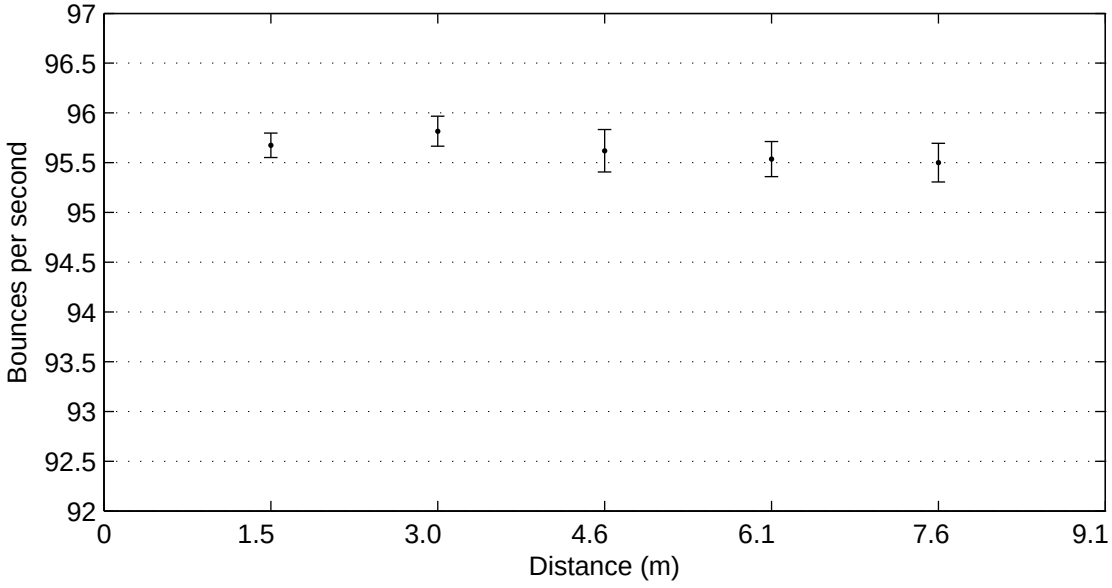

Figure 3.1.7 Results from WiFi single-interior-wall experiment

Figure 3.1.7 shows the count of bounces per second between the laptop and router. Each point in this plot again represents the mean and standard deviation of 68 10,000-packet samples (excluding any samples with dropped packets). By comparing these results with those in Figures 3.1.14 and 3.1.5, it can be seen that the addition of a single interior wall did not drastically affect the bounce rate between the laptop and the router.

A further set of experiments was performed to determine the effects of multiple walls on the system. The access point was placed outdoors, 6.1 m away from the exterior wall of the residence. The laptop program was run at distances of 0.0 m, 1.5 m, 3.0 m, and 6.1 m away from the inside of the exterior wall. Interior walls were positioned approximately 3.0 m and 5.0 m away from the exterior wall. Exterior walls were standard vinyl siding, and interior walls were drywall with wooden studs. Figure 3.1.8 illustrates this setup.

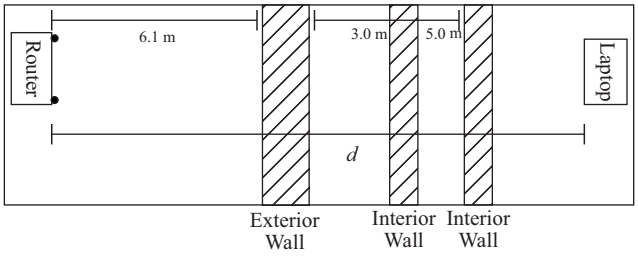

Figure 3.1.8 WiFi multiple-wall experimental setup

Figure 3.1.9 shows the results of the multiple-wall experiment. Each point and error bar represents the mean and standard deviation of 10 blocks of 10,000 bounces. The scale of the "Bounces per second" axis of this plot is much greater than the scales of the previous plots, due to the data point collected with a separation distance of 12.2 m being so low. It can be seen here that the bounce rate did drop slowly with increasing distance for the first three data points. Though the graph is too small to see the actual values, these were approximately 95.8 bounces per second at 6.1 m, 94.6 bounces per second at 7.6 m, and 94.2 bounces per second at 9.1 m. This drop is more significant than the change detected in the line-of-sight test (Figure 3.1.14). This indicates that multiple interior and exterior walls can have a cumulative effect on the bounce rate.

After the first three data points of Figure 3.1.9, there was a sharp drop off before the next point. This result was verified with repeated testing. The cause for this sudden drop in bounce rate was unknown, but it was suspected that the change was influenced by the presence of a large metal shelf

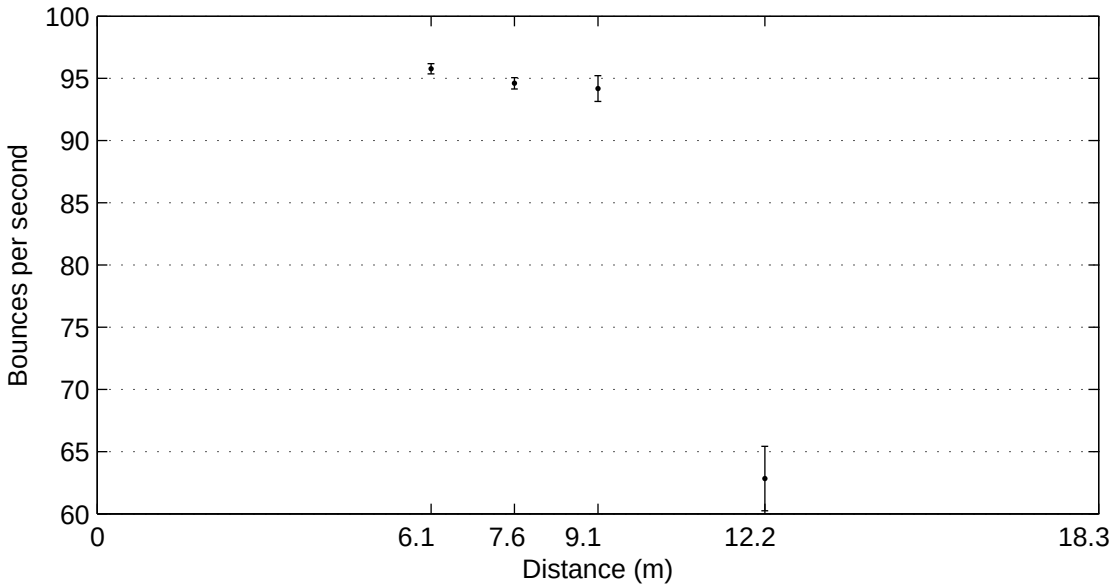

Figure 3.1.9 Results from WiFi multiple-wall experiment

located next to the wall between the test points at 9.1 m and 12.2 m of separation between the laptop and the router.

Another experiment was conducted that again showed an effect of large metallic masses in the test range. Measurements were collected outdoors with separation distances of 3.0 m, 7.6 m, 15.2 m, 22.9 m, 30.5 m, and 38.1 m. 10 blocks of 10,000 bounces were collected at each distance. A car was positioned next to the 22.9 m test point. Figure 3.1.10 shows the setup for this experiment. This experiment was then repeated in a different location with no nearby metallic objects.

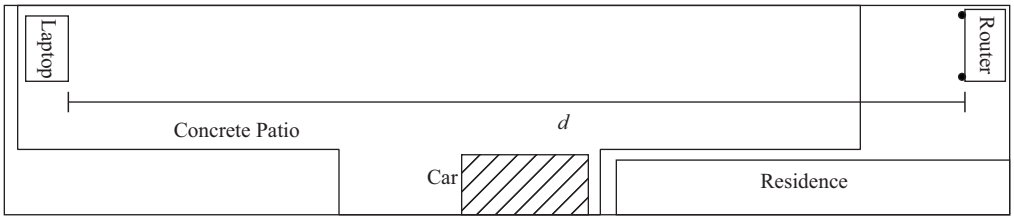

Figure 3.1.10 WiFi metallic object experimental setup

Figure 3.1.11 plots the measured bounce rate relative to the separation distance between the laptop and router for the test performed next to the car. Each point is the mean and standard deviation of 10 average bounce-per-second calculations, with each calculation incorporating 10,000 bounces. In this graph, the rate of bounces per second obviously drops with increasing separation.

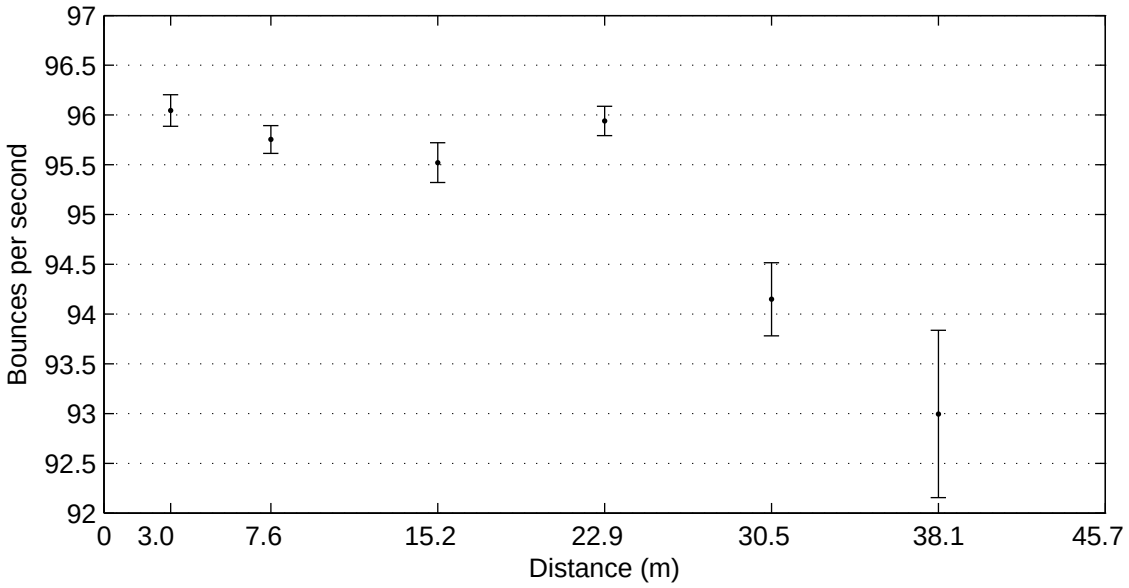

Figure 3.1.11 Results from WiFi metallic object experiment

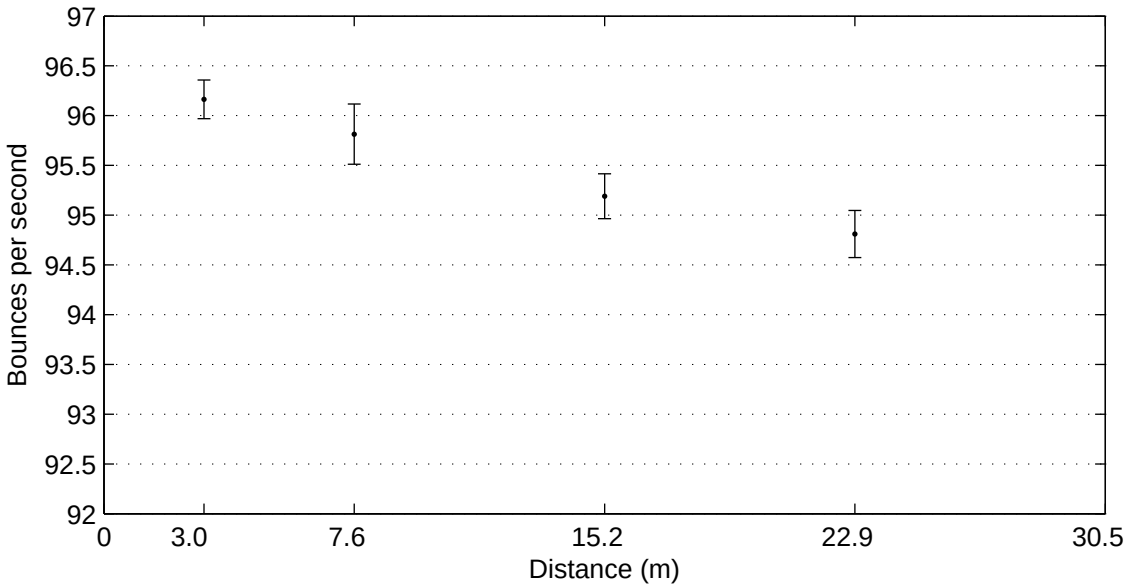

Figure 3.1.12 Results from WiFi no metallic object experiment

The one discrepancy exists at the separation distance of 22.9 m where the car was located. In these results, it can bee seen that the presence of the metallic object significantly increased the bounce rate.

Figure 3.1.12 shows the bounce rate results from the test with no car. Each point in this data set represents the mean and standard deviation of the 10 samples, and each of those being the average count of bounces per second for 10,000 packets. In this graph it can be seen that the bounce rate did not vary from the trend when there were no nearby metallic objects.

It was observed during testing that some experiments when repeated did not yield precisely the same results. It was suspected that many of the previously tested error sources contributed to this non-repeatability. To test the repeatability of the bounce rate measurements, the laptop and the access point were set up on a short, indoor, line-of-sight experiment. The two devices were positioned first at 1.5 m apart, increasing in 1.52 m increments until the final test point at 6.1 m of separation. This test can be thought of as a "zoom in" of the first three points of the test results presented in Figure 3.1.2. The data packet bouncing program was run in each position for 2 hours (approx. 680,000 bounces). This process was then repeated, starting at 6.1 m and moving in to 1.5 m. Figure 3.1.13 illustrates this setup.

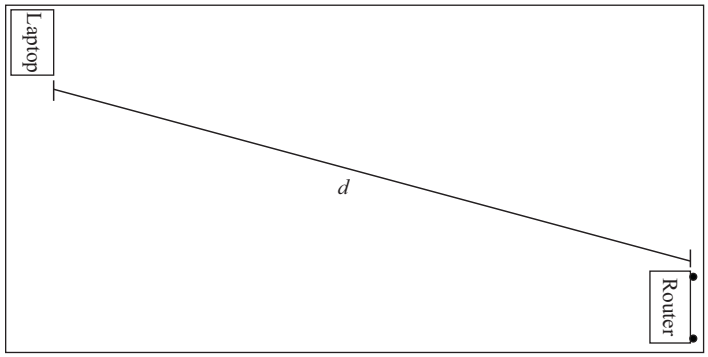

Figure 3.1.13 WiFi repeatability experimental setup

Figures 3.1.14 and 3.1.15 show the recorded number of bounces per second between the laptop and the router at each tested separation distance. Both of these figures show the average and standard deviation in the recorded number of samples. The approximately 680,000 bounces taken in 10,000-bounce sections yielded 68 samples per point on the figure. Thus, each point in both figures represents the mean and standard deviation of the 68 samples. Any samples that included a dropped

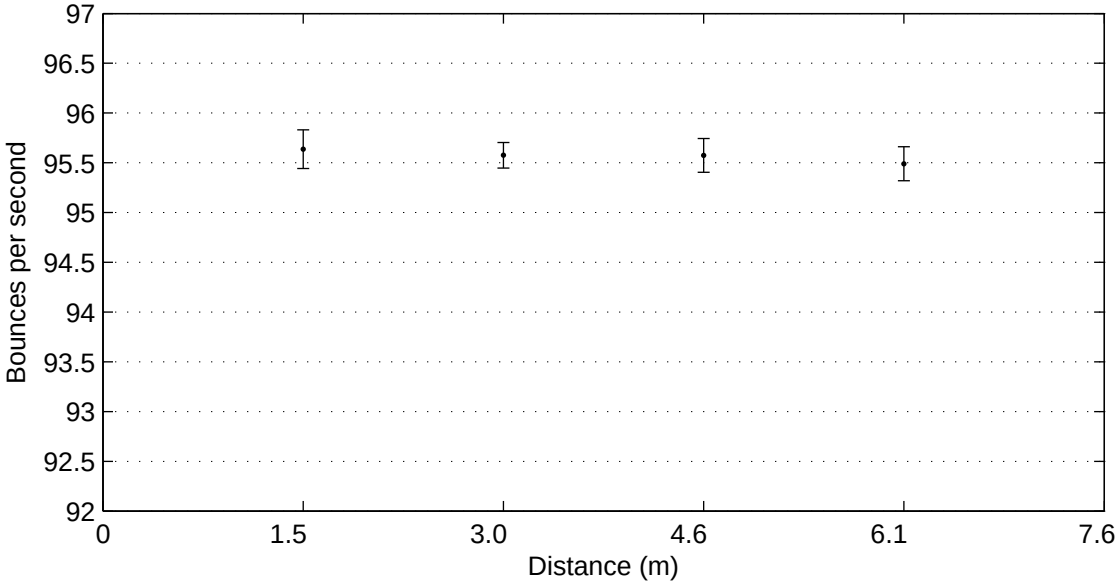

Figure 3.1.14 Results from WiFi repeatability experiment 1

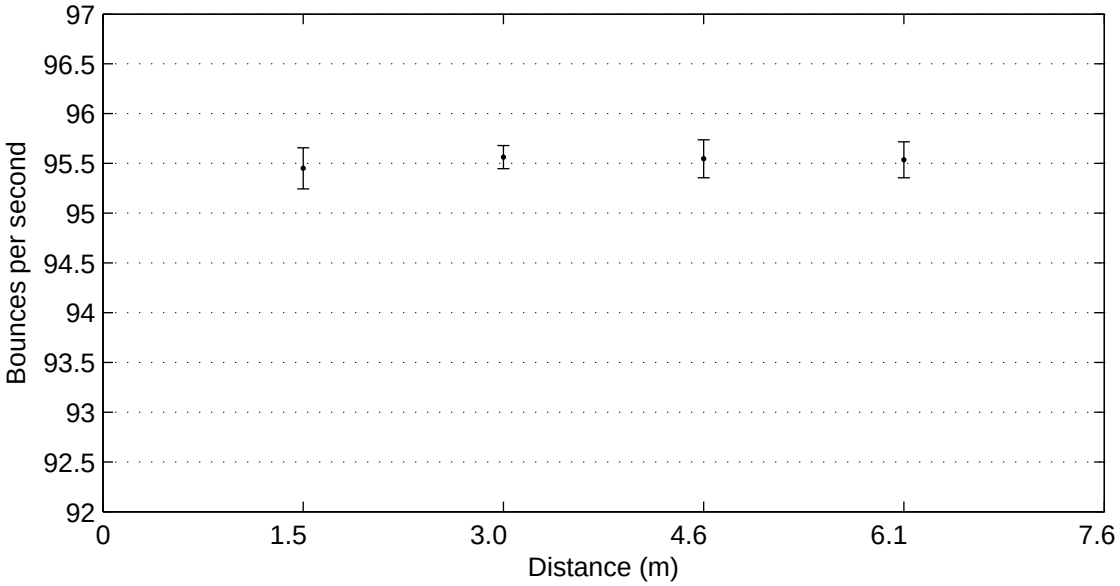

Figure 3.1.15 Results from WiFi repeatability experiment 2

packet were excluded from the calculations. Figure 3.1.14 shows a slight drop in the bounce rate as separation distance increases, but Figure 3.1.15 shows a slight increase. This non-repeatability was not drastic enough to affect measurements over long ranges, but was a significant source of error in the short range experiment.

Expected Change of Bounce Rate Across 50 m Range				
± 10.0 Bounces per Second				
Source of Error	**Amount of Error**			**Derived From**
Outdoor vs. Indoor Behavior	±	0.5	Bounces per Second	Figures 3.1.4 & 3.1.5
Single Wall	±	0.25	Bounces per Second	Figure 3.1.7
Multiple Walls	±	1.5	Bounces per Second	Figure 3.1.9
Metallic Masses	±	30.0	Bounces per Second	Figures 3.1.9, 3.1.11, & 3.1.12
Non-Repeatability of Experiments	±	0.3	Bounces per Second	Figures 3.1.14 & 3.1.15

Table 3.1.2 Estimated amount of error in packet-based rangefinding method

Table 3.1.2 shows the expected amount of variation in the bounce rate as the separation distance between the laptop and router increases from 0 m to 50 m. It also shows an estimate of the amount of error that can be attributed to the different sources mentioned at the beginning of the section. While most of the error sources do not individually contribute enough noise to greatly affect the distance measurement over the range of 50 m, they are strong enough to skew results when attempting to measure shorter distances. Multiple walls and large metallic masses have a much greater impact on the amount of error. These error sources are strong enough to cause false distance measurements that could make a WiFi-based system unusable.

3.2 Radio carrier-wave rangefinding

The goal of the carrier-wave rangefinding system experiments was a test involving one transmitting antenna positioned on either side of the test range and the receiving antenna moving to various locations between the two transmitters. When either transmitter was broadcasting on its own, the receiver would pick up a relatively constant power level across the entire test range. However, when both radios were transmitting simultaneously, the power of the received signal would vary as the receiver was moved across the test range.

An experiment was performed that showed these results. Two custom quarter-wave dipole antennas were built, attached to transmitters, and mounted to opposite walls of the laboratory. The two dipole antennas were separated by approximately 9.5 m of lab space. The oscilloscope acting as

the receiver was fitted with the longer of the two "rubber duck" antennas included with the Yaesu FT-817 radios. This setup is shown in Figure 3.2.1. In the first test, only the antenna at the left edge of the figure was activated. Only the antenna at the right edge of the figure was activated for the second test. In the final test, both antennas transmitted simultaneously. For each test, 100 measurements were taken by the oscilloscope at positions starting at 50 cm away from the left antenna and moving in 50 cm increments to 900 cm away. 100 measurements of the RMS voltage on the oscilloscope's antenna were taken at each location.

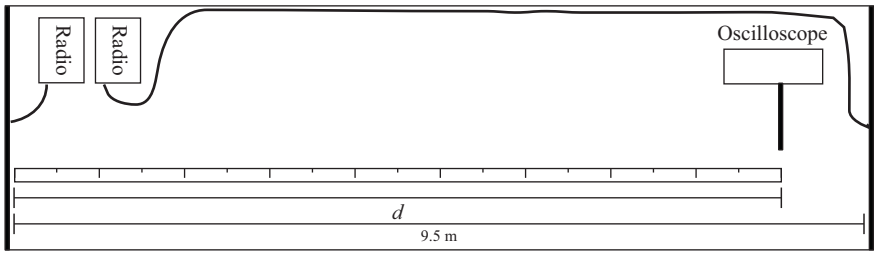

Figure 3.2.1 Received power from two dipole transmitters experimental setup

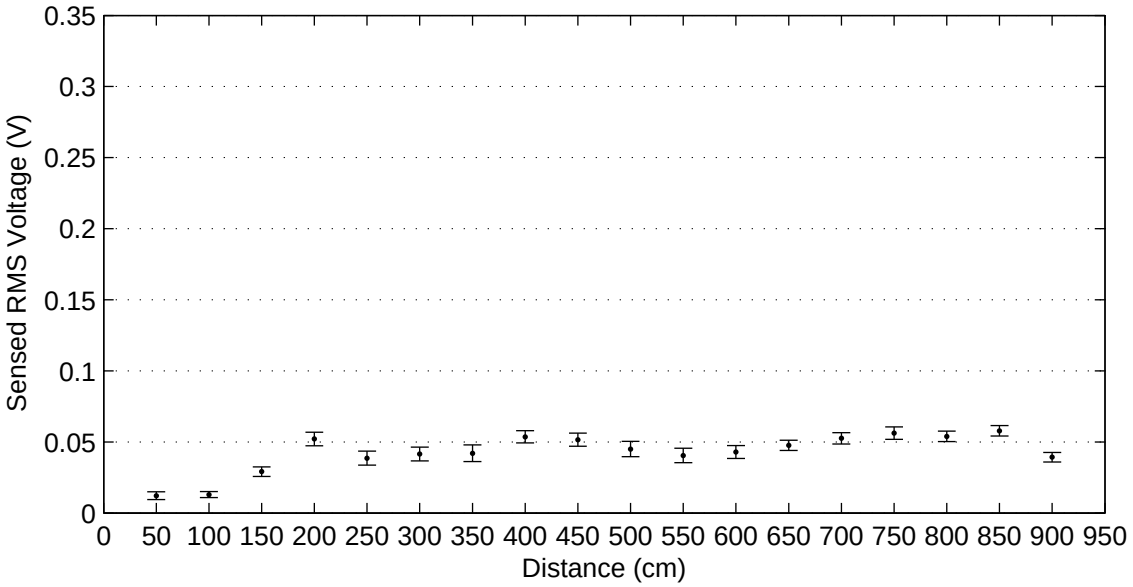

Figure 3.2.2 Results of received power from two dipole transmitters experiment, left antenna active

Figure 3.2.2 shows the mean and standard deviation of the samples of the RMS voltage at the receiver antenna with only the left transmitting antenna active. Figure 3.2.3 shows the same type of results, but with only the right transmitting antenna. Figure 3.2.4 shows the results of the test

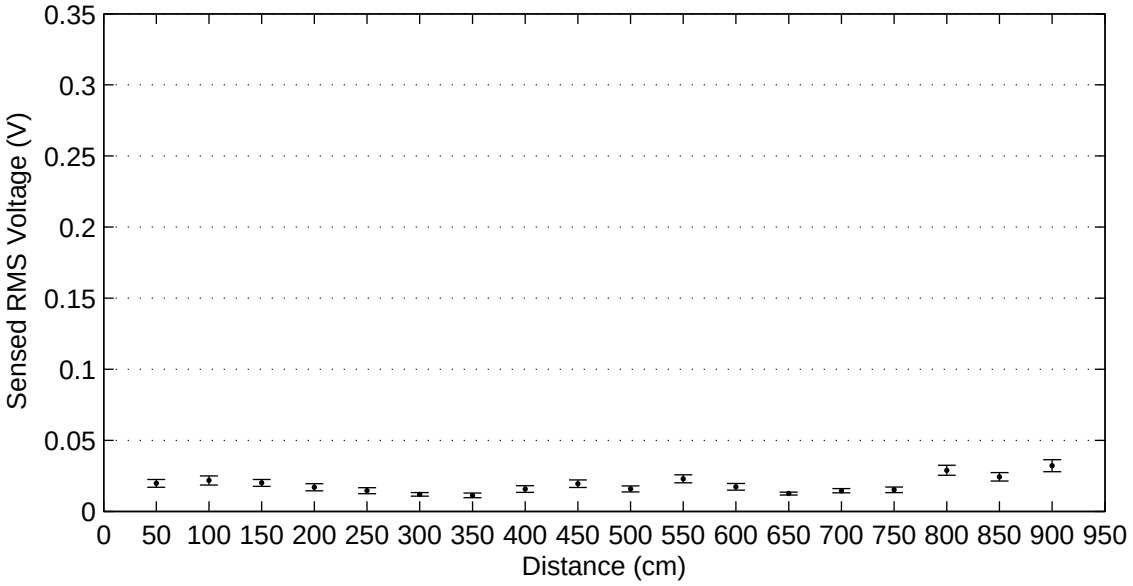

Figure 3.2.3 Results of received power from two dipole transmitters experiment, right antenna active

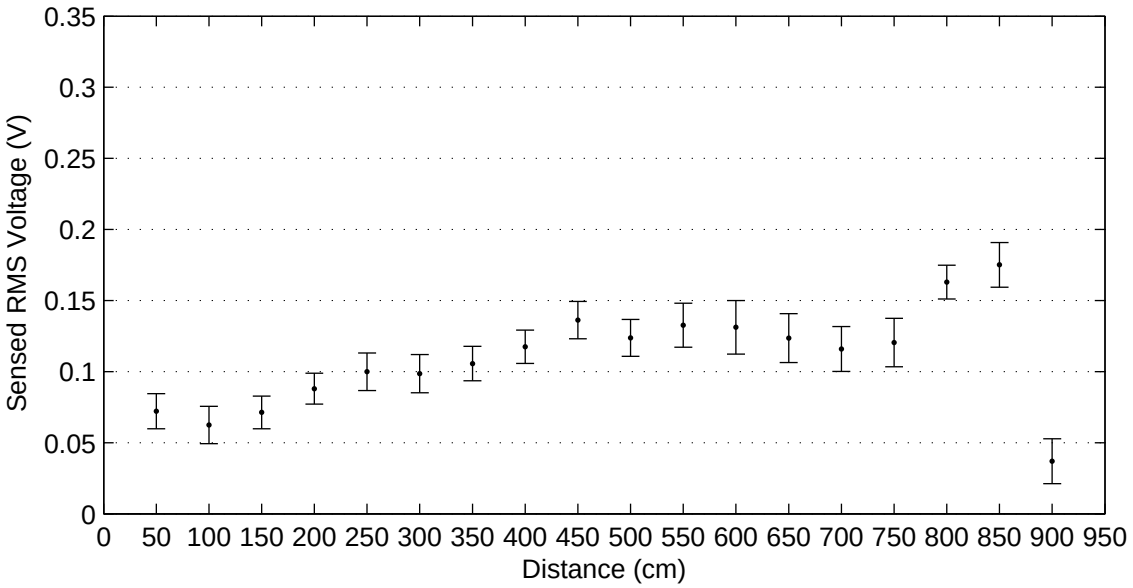

Figure 3.2.4 Results of received power from two dipole transmitters experiment, both antennas active

with both antennas activated. It was expected to see in this graph a single peak or valley of received power, with the other points descending or rising monotonically from both sides of this feature. If the peak or valley was not present in the plot, it was expected to see a monotonically rising or falling function. It can be seen here that the observed data somewhat matched the predicted results. The data was not as supportive as hoped though, so more experiments were performed in order to isolate the causes of the errors.

Potential Causes of Error in Carrier-Wave Rangefinding

Antenna Design
Antenna Orientation
Receiver Proximity to Transmitter
Sound Transmitted over Carrier Wave
Wall Effects

Table 3.2.1 Error sources in carrier-wave rangefinding method

Table 3.2.1 shows a list of potential sources of error in the carrier-wave rangefinding method. Each of these sources was investigated in turn as a potential cause for the noise observed in Figure 3.2.4.

An experiment was performed that helped verify the amount of error caused by different antenna designs. This experiment was a repeat of the previous experiment, but with the dipole antennas

replaced by the rubber duck antennas that were included with the radios. Also, the separation between the two transmitting antennas was reduced to 850 cm. This test setup is illustrated in Figure 3.2.5.

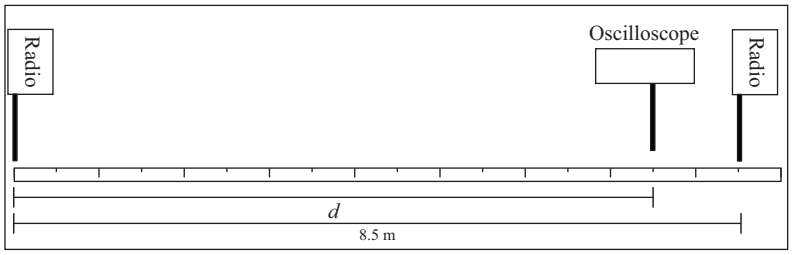

Figure 3.2.5 Received power from two rubber duck transmitters experimental setup

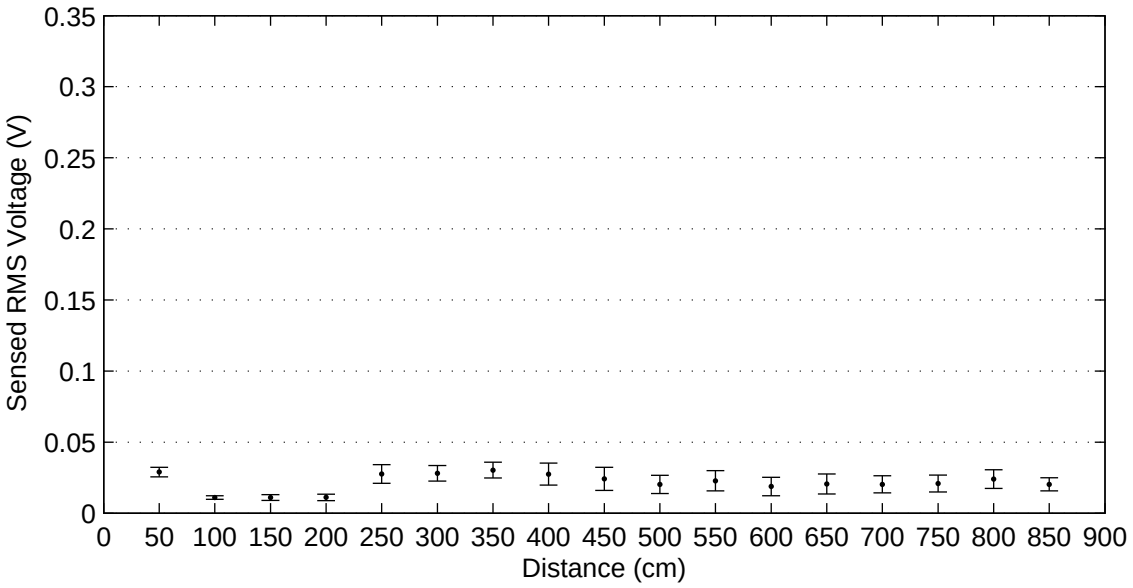

Figure 3.2.6 Results from received power from two rubber duck transmitters experiment, left antenna active

Figures 3.2.6 through 3.2.8 show the results of using rubber duck antennas instead of the quarter-wave dipoles. The three plots show the observed power first with only the left antenna active, next with only the right antenna active, and finally with both antennas active. By comparing Figures 3.2.6 – 3.2.8 with Figures 3.2.2 – 3.2.4, it can be seen that transmitters using the rubber duck antennas consistently produced approximately 0.025 less V_{rms} power at the receiver individually. When both transmitters were activated, however, the dipole antennas were stronger by almost 0.1 V_{rms}.

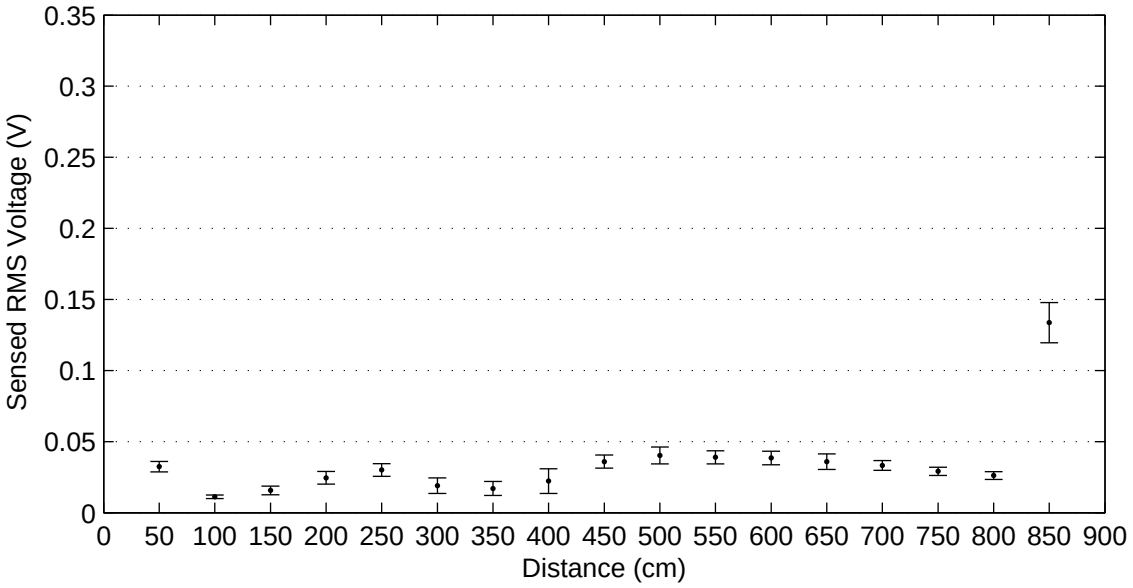

Figure 3.2.7 Results from received power from two rubber duck transmitters experiment, right antenna active

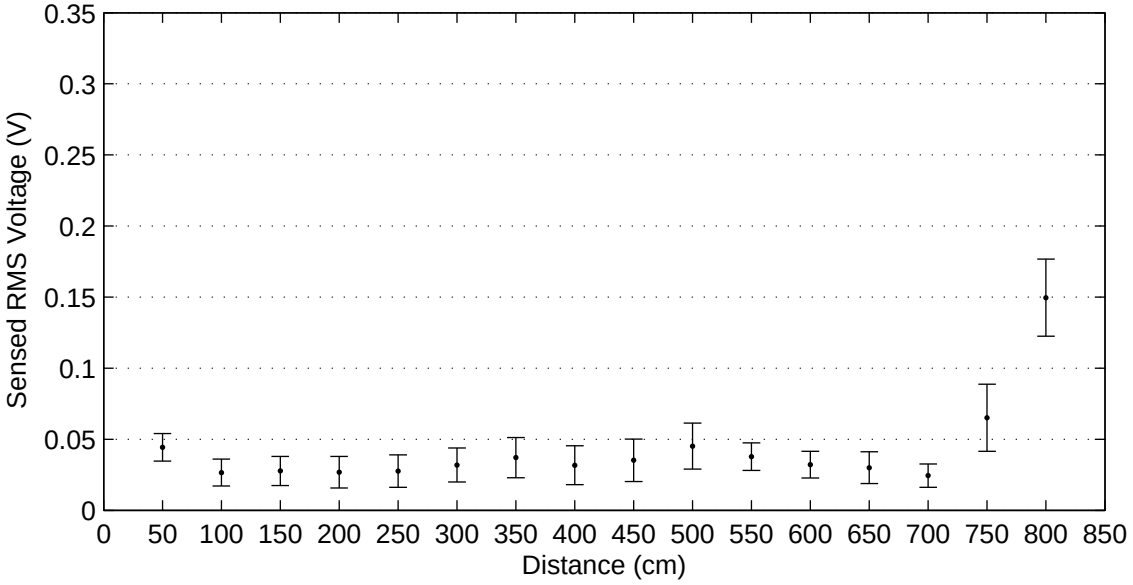

Figure 3.2.8 Results from received power from two rubber duck transmitters experiment, both antennas active

Another experiment was designed to test the effects of using different receiver antenna orientations. One dipole antenna was attached to a transmitter and mounted to one wall of the laboratory, as shown in Figure 3.2.9. The receiving oscilloscope used the rubber duck antenna, and data was gathered with the receiver moving from 1 m to 6 m away from the transmitting antenna in 1 m increments. 100 measurements of the RMS power of the carrier wave were gathered at each position, and these measurements were gathered at each position 5 times. A 1500 Hz sound wave was modulated on the carrier during this experiment. This experiment was repeated a second time, but this time with the receiving antenna positioned in a vertical orientation as shown in Figure 3.2.10.

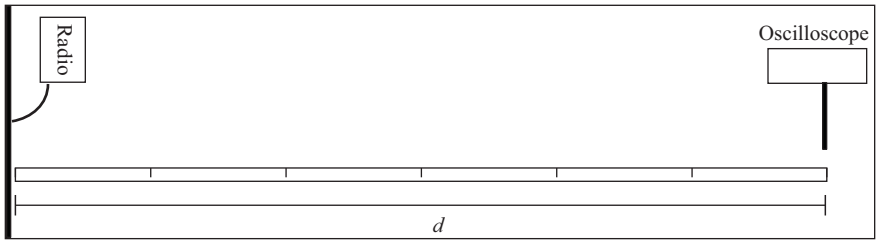

Figure 3.2.9 Horizontal receiving antenna experimental setup

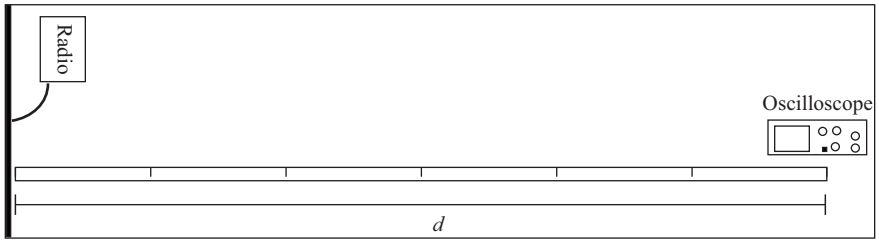

Figure 3.2.10 Vertical receiving antenna experimental setup

Results of the test with the receiving antenna in a horizontal orientation are shown in Figure 3.2.11, while results of the test with the receiving antenna vertically oriented are presented in Figure 3.2.12. The change in antenna orientation increased the received RMS power by more than 0.1 V_{rms}. However it also substantially increased the measurement noise, as evidenced by the larger error bars in Figure 3.2.12. Changing the orientation did not alter the effectiveness of the system in a positive manner.

An experiment was performed that showed that received power increased substantially if the receiver was positioned too close to the transmitter. In the test, the transmitter and oscilloscope were

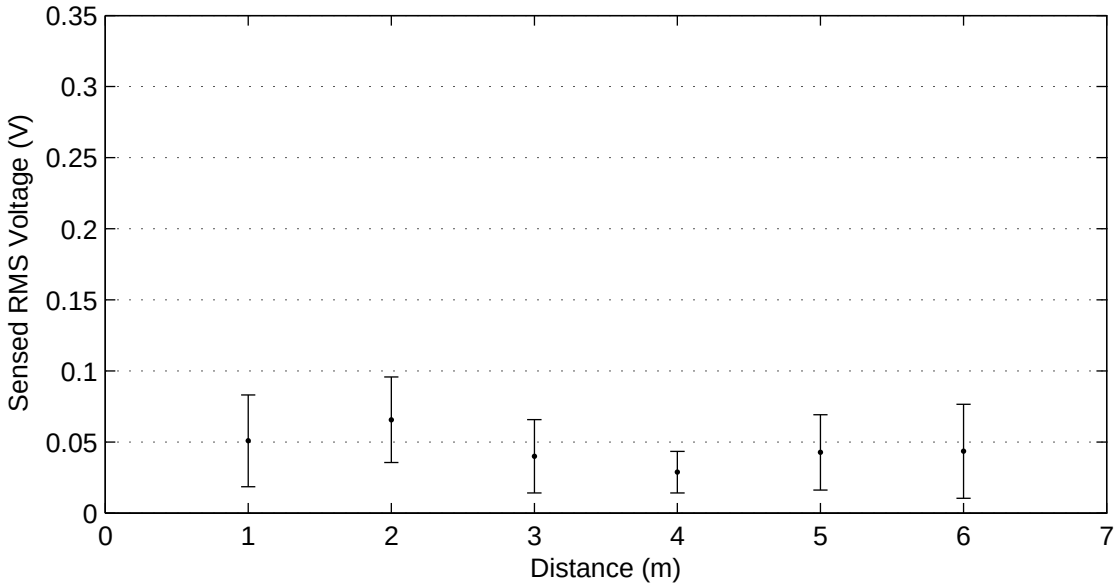

Figure 3.2.11 Results from horizontal receiving antenna experiment

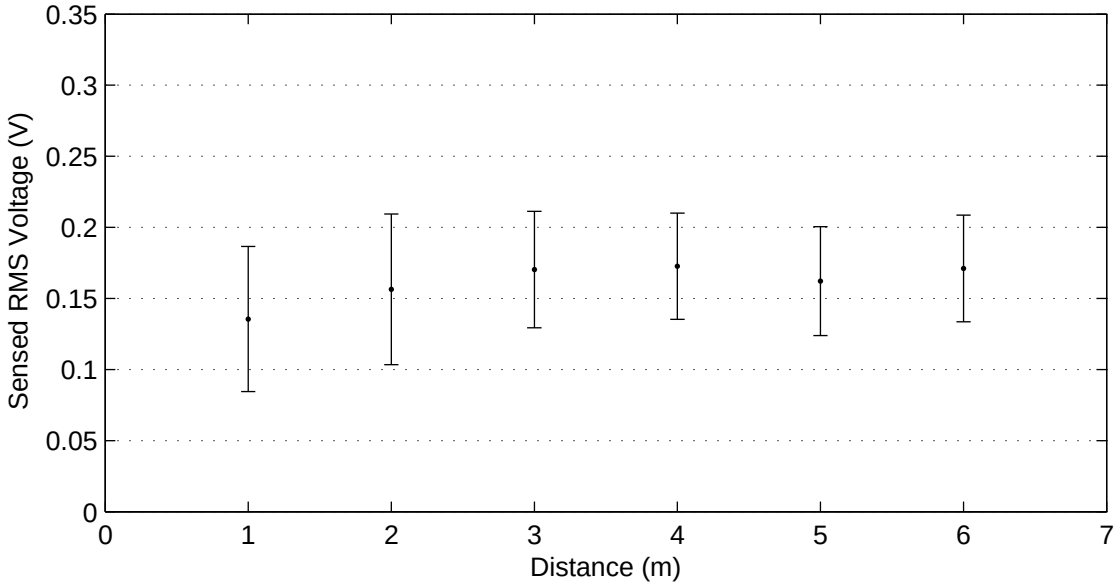

Figure 3.2.12 Results from vertical receiving antenna experiment

placed along a 1 meter test platform. Both devices were fitted with rubber duck antennas. Signal strength measurements were gathered at 10 cm increments of separation between the transmitter and the oscilloscope, moving from 10 cm to 1 m of separation. At each position, 100 measurements were taken. This process was repeated 5 times for each distance increment. This setup is shown in Figure 3.2.13.

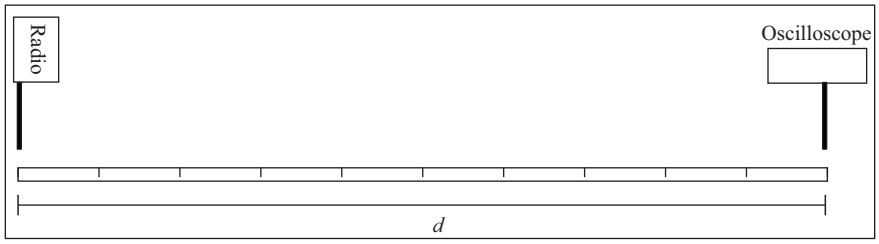

Figure 3.2.13 Close-proximity experimental setup

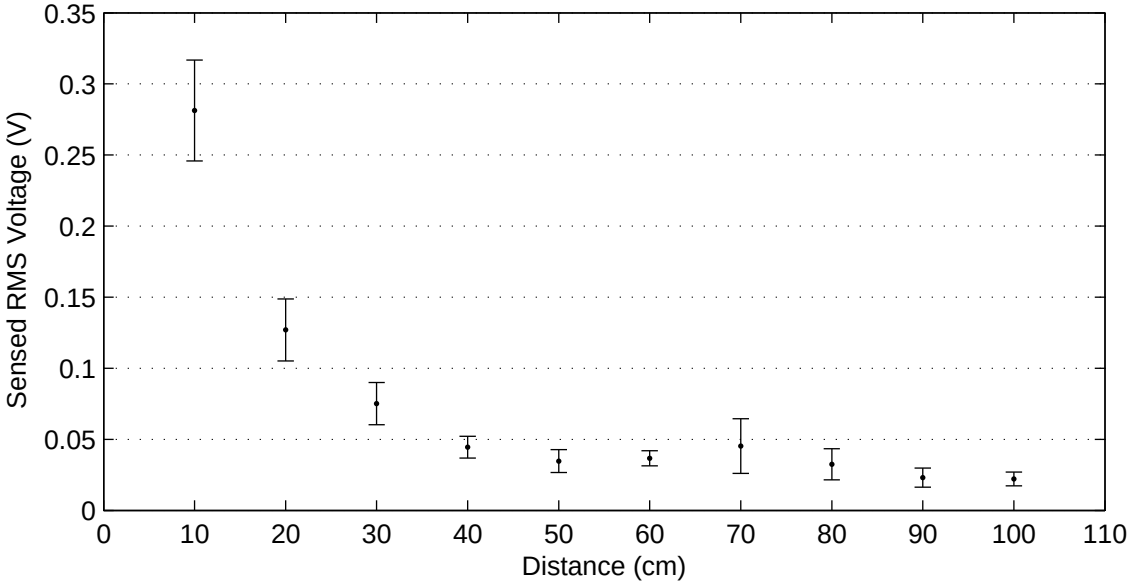

Figure 3.2.14 Results close-proximity experiment

Figure 3.2.14 shows the results of the close-proximity test. It can be seen that the power drops off significantly while the receiving oscilloscope moves from 10 cm away from the transmitter to 40 cm away from the transmitter, but the power levels stabilize somewhat after that. This shows that the system is affected negatively when the receiver moves too close to transmitting antennas.

An experiment was performed to observe if modulating a sound wave on the carrier affected the received signal strength. To test this, the close-proximity test was repeated, but this time the transmitter was set to send an AM-modulated sinusoidal sound wave at 1500 Hz. Signal strength measurements were again gathered at 10 cm increments of separation between the transmitter and the oscilloscope, moving from 10 cm to 1 m of separation as shown in Figure 3.2.13. At each position, 100 measurements were taken. This process was repeated 5 times for each distance increment.

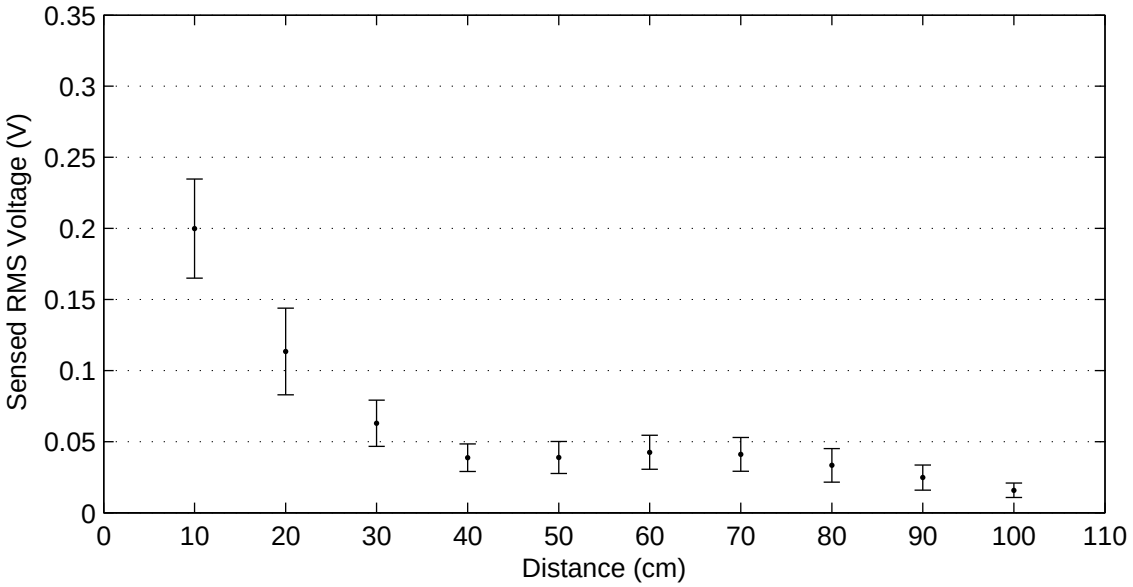

Figure 3.2.15 Results from close-proximity with sound experiment

Figure 3.2.15 shows the results of the close-proximity experiment with the sound wave generator sending a 1500 Hz tone. By comparing this graph with Figure 3.2.14, it can be seen that at ranges less than 40 cm the sound wave does increase received power. At usable ranges of the carrier-wave system, however, the addition of a sound wave has very little impact on the received signal strength.

A final experiment was performed to gauge the effects of walls and other indoor obstacles on the system. For this test, only one dipole antenna was used for transmitting. The receiving oscilloscope was constantly in motion throughout this test. Beginning approximately 1 m from the transmitting antenna, the receiver moved toward the exit from the laboratory. At approximately 30 seconds after the test begins, the receiver moved through the doorway from the lab into a hallway. It then progressed down the hall, moving away from the transmitting antenna. At approximately 60 seconds, the receiver stopped moving away and progressed down the hall in the opposite direction. 90 seconds

after the test had started, the receiver stopped moving down the hall and began to move back toward the lab entrance. 105 seconds after the test began, the receiver re-entered the lab. The test ended approximately 120 seconds after the start with the receiver returning to its starting location. This test was repeated 5 times. The path traveled by the oscilloscope is illustrated in Figure 3.2.16.

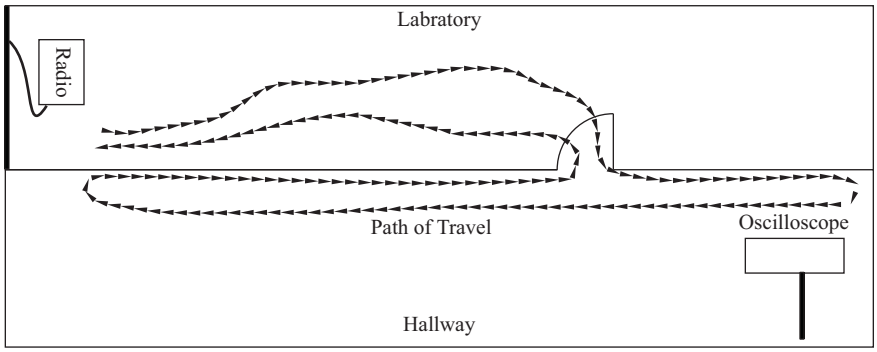

Figure 3.2.16 Wall effects experimental setup

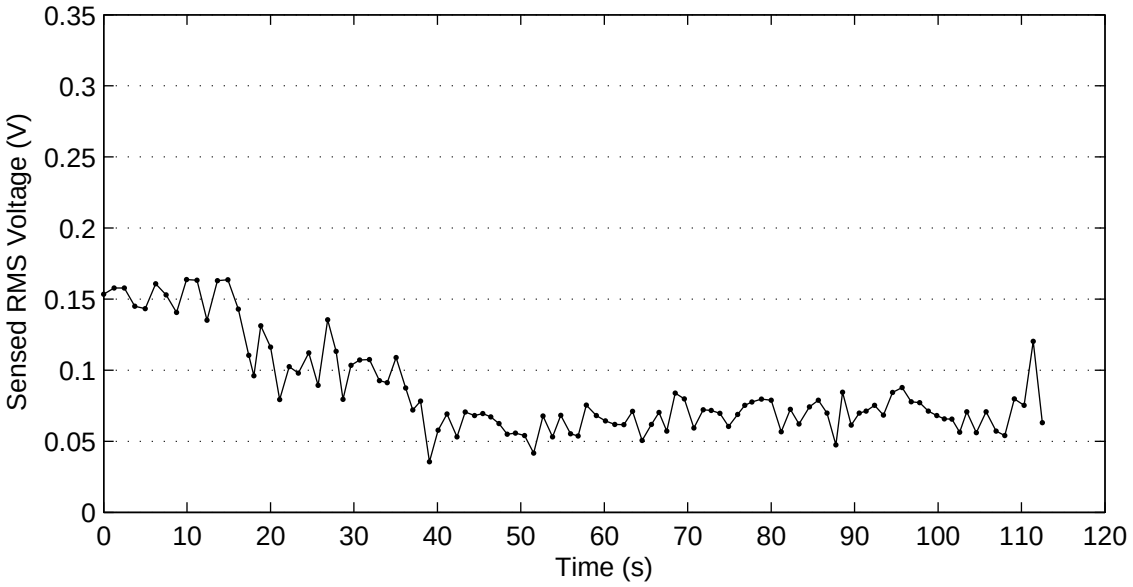

Figure 3.2.17 Results from wall effects experiment, all runs averaged together

Figure 3.2.17 shows the results of averaging together the 5 repetitions. It was hoped that walls would have a negligible effect on the system, but as can be seen in the figure, this is not the case. The RMS voltage varies significantly depending on the location of the receiver. It should also be noted that the power received at the start of the test is not as strong as the received power at the

end of the test, even though the receiver was returned to its original position. This behavior was not well understood, and as such no attempt is made here to explain it. The results of this experiment show that since the walls greatly affect the received power from the transmitter, the carrier-wave rangefinding system cannot be used to reliably measure distances indoors.

Expected Change of V_{rms} Across 9.5 m Range
$\pm$ 0.12 V_{rms}

Source of Error	**Amount of Error**	**Derived From**
Antenna Design	$\pm$ 0.10 V_{rms}	Figures 3.2.2, 3.2.3, 3.2.4, 3.2.6, 3.2.7, & 3.2.8
Antenna Orientation	$\pm$ 0.13 V_{rms}	Figures 3.2.11 & 3.2.12
Receiver Proximity to Transmitter	$\pm$ 0.24 V_{rms}	Figure 3.2.14
Sound Transmitted over Carrier Wave	$\pm$ 0.01 V_{rms}	Figures 3.2.14 & 3.2.15
Wall Effects	$\pm$ 0.10 V_{rms}	Figure 3.2.17

Table 3.2.2 Estimated amount of error in carrier-wave rangefinding method

The preceding experiments show the strength of error sources in the carrier-wave rangefinding method. As can be seen in Table 3.2.2, most of the error sources are just as strong as, if not stronger, than the signal being looked for. This presents a compelling reason to look for a better alternative for an indoor rangefinding system.

3.3 Radio carried-signal rangefinding

For most of the tests using this system model, FM frequencies of 52.000 MHz and 52.500 MHz were used to transmit the sound waves. One Yaesu FT-817 transceiver was located at each end of the test range. These were used as transmitters and were located 27.7 m apart down a straight hallway, and one transmitter was tuned to each FM frequency. Another pair of Yaesu FT-817 transceivers was then moved to various points between the two transmitters. These two radios were used as receivers, and one was tuned to each FM frequency. In all tests, the transmitters were controlled via an RS-232 serial port connection to the test computer, and sound was sent to the transmitters over 2 conductors of a standard category-5 ethernet cable (due to the sound input connection jack on the transmitters). The received sound waves were returned to the A/D converter card in the computer by a pair of 2-conductor audio cables. This test set up is show in figure 3.3.1.

The first experiment was a simulation to predict the expected values for these phase differences over the test range. The possible range for the phase offset measurement is from approximately -166.67 μsec to 166.67 μsec. A negative phase offset on this scale shows that the wave from the antenna

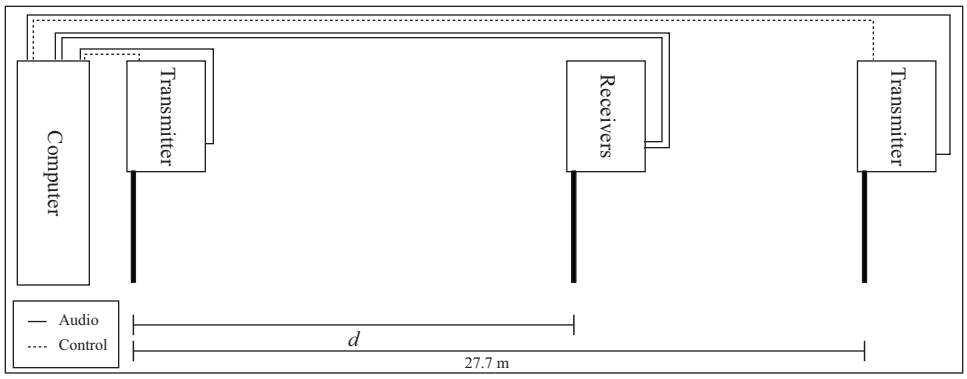

Figure 3.3.1 Diagram of the VHF carried-signal rangefinding method

closest to the computer arrived at the receiver first, whereas a positive phase offset indicates that the sound wave from the transmitter at the other end of the test range was earliest. The point and error bars for each distance represent the mean and standard deviation of the phase difference between 20 measurements of the phase offset between the waveforms. Each of the 20 offset measurements is calculated as described in the Methods chapter of this thesis. Figure 3.3.2 shows the results of the first experiment. It can be seen in this graph that the phase offset between received waves was expected to rise gradually, but they are all grouped around the 40 μsec range.

A test was performed to determine if this system would be able to achieve rangefinding indoors. This was done by placing the two transmitters on opposite ends of the hallway with 27.7 m of separation and moving the receivers to several points between the two transmitters. Data was collected with the receivers positioned at 0.3 m, 4.7 m, 7.8 m, 10.8 m, 13.9 m, 16.9 m, 20.0 m, 23.0 m, and 27.4 m away from the transmitter located next to the controlling computer. At each position, 20 data sets consisting of 1024 samples of the sound wave from the receivers were collected. This was done first with the transmitter that was located next to the computer set for 52 MHz and the transmitter at the opposite end of the hallway set to 52.5 MHz. This test was then repeated with the transmitter frequencies swapped, to ensure that frequency choice did not influence the test results.

Figure 3.3.3 shows the calculated phase difference between the two sound waves for the test with the transmitter placed at the computer end of the test range (the left edge of the graph) set to 52 MHz. This plot shows the expected trend (the phase offset is rising from left to right), but the actual offset value is lower than anticipated and the degree of offset change is higher than expected.

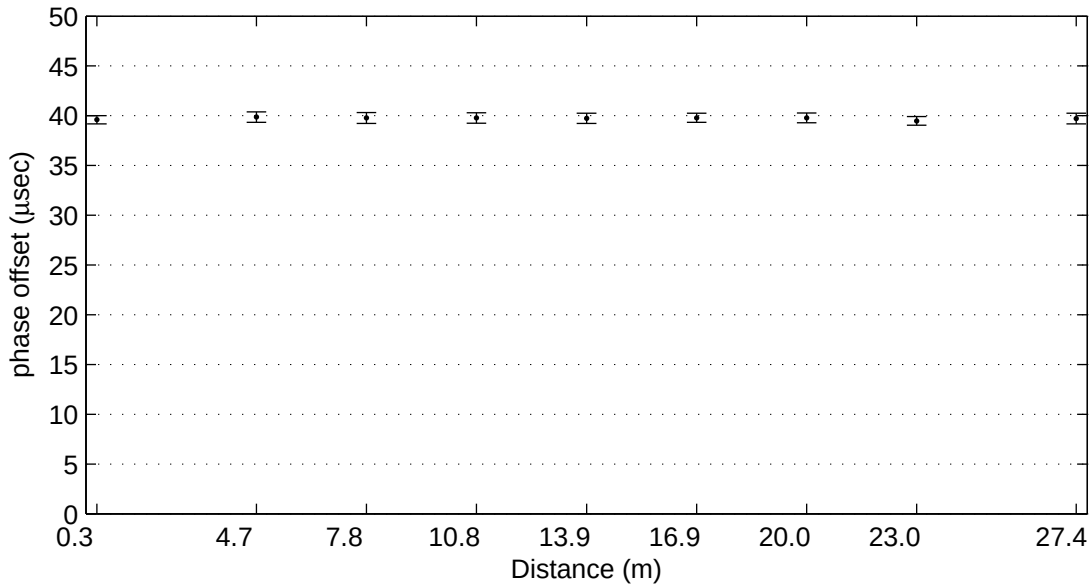

Figure 3.3.2 Expected phase differences, simulated

Figure 3.3.4 shows the calculated phase difference when the transmitter located at the computer end of the test range set to 52.5 MHz. The data gathered was expected to mirror the graph in Figure 3.3.3 across the horizontal axis. This indicates that the expected phase offset would have been approximately 130 μsec with the receivers positioned at the left side of the test range, lowering to approximately 160 μsec at the right side. As can be seen in the figures, this behavior was observed during the experiment. For all further experiments related to this system, only the graph corresponding to the transmitter at the computer end of the test range set to 52 MHz (Figure 3.3.3) is presented. All data with the transmitter at the computer end of the range set to 52.5 MHz supports the results presented.

Potential Causes of Error in Carried-Signal Rangefinding

Modulation Circuitry Delay
Phase Offset Measurement
Sound Cables of Varying Length
Sound Cable Coil Shape
Transmitter Frequency
Walls / Doorways
Environmental Clutter (e.g. Desks)

Table 3.3.1 Error sources in carried-signal rangefinding method

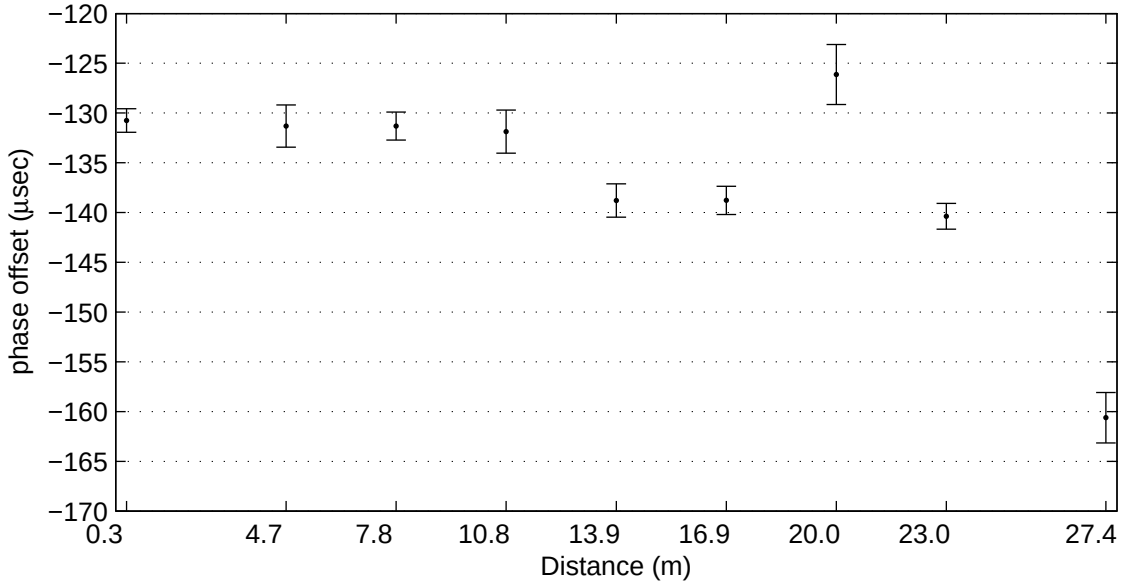

Figure 3.3.3 Measured phase difference in hallway, transmitter at computer end set to 52 MHz

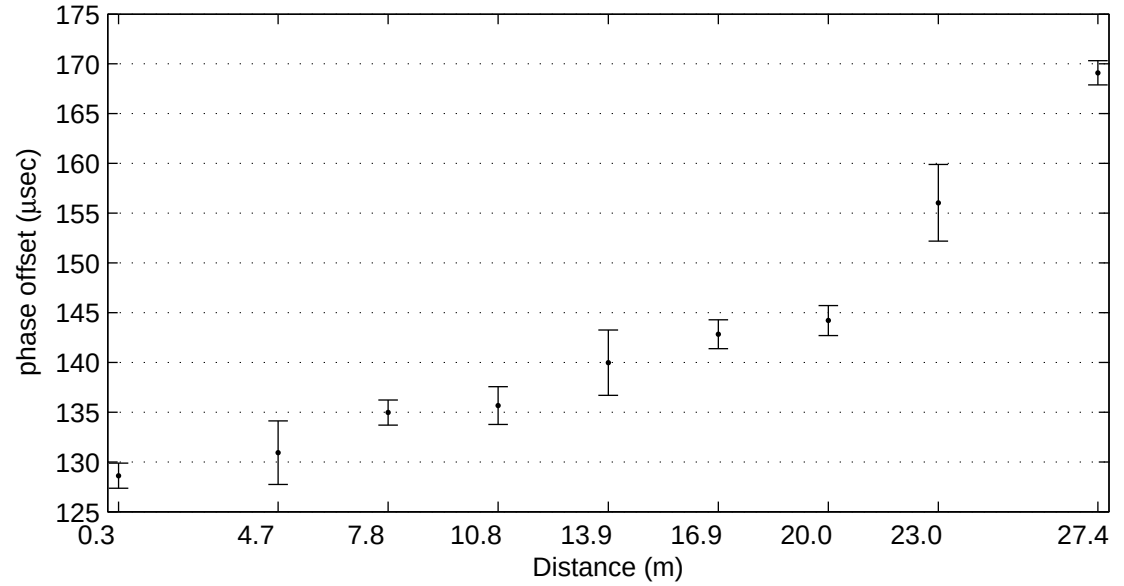

Figure 3.3.4 Measured phase difference in hallway, transmitter at computer end at 52.5 MHz

Table 3.3.1 shows a list of potential sources of error in the carried-signal rangefinding method. Each of these sources was investigated as a potential cause for the difference between Figures 3.3.2 and 3.3.3.

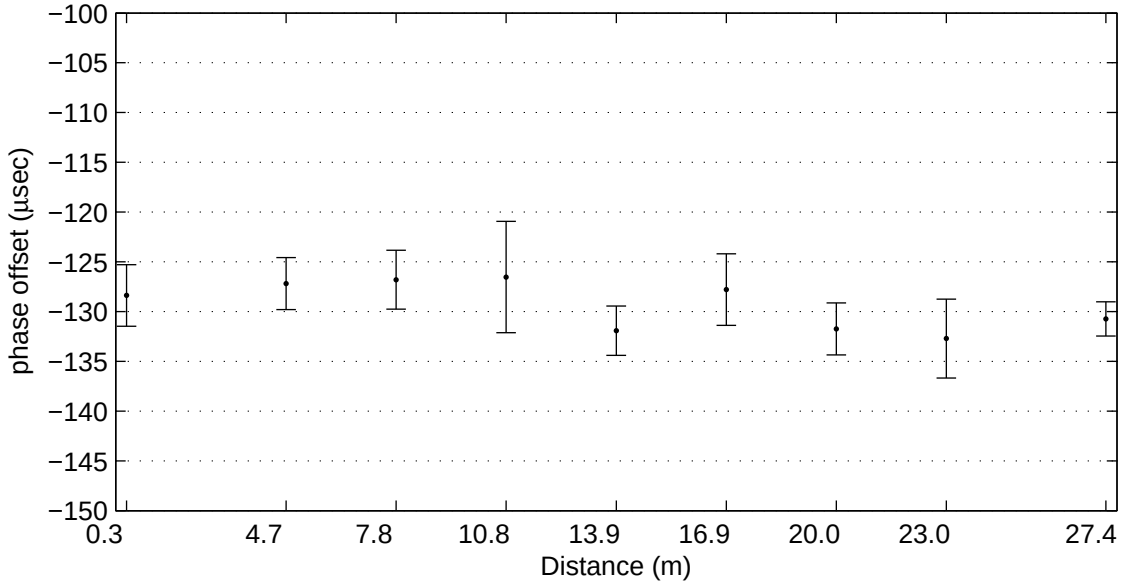

Figure 3.3.5 Measured phase difference in outdoor environment

To observe the impact of delay cause by variations between the modulation circuitry of the radios used, the hallway test was repeated outdoors along a sidewalk. The results of this experiment are shown in Figure 3.3.5. The data points in this graph again show the unexpected phase offset of approximately -128 μsec. By comparing these results to the results in Figure 3.3.3 and the results of the remaining experiments described in this section, it was determined that this phase offset was a result of the modulation circuitry of the radios used. The data being sent from the transmitter located next to the computer system consistently arrived at the receivers 128 μsec ahead of the data from the other transmitter, regardless of a drastic change in operating environment.

The sidewalk test also illustrates another potential source of error in the carried-wave rangefinding system. By comparing the results of this test with the results of the in-hallway test, it can be seen that the phase offset did not change as drastically with distance outdoors as it did indoors. This supports the hypothesis that environmental factors such as walls and hallways affect the phase offset of the transmitted sound waves.

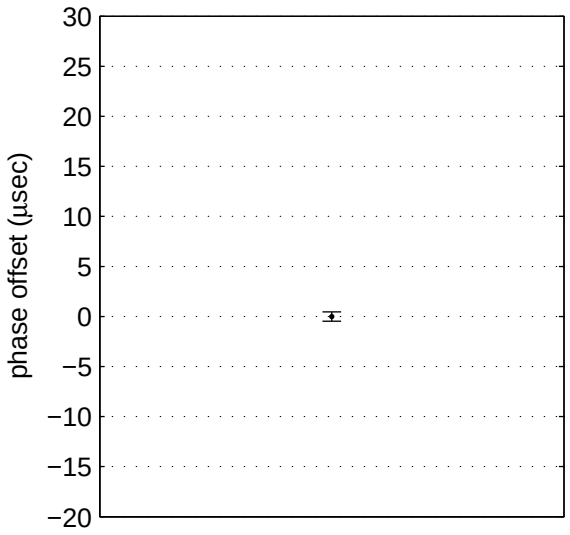

Figure 3.3.6 Measured phase difference of non-transmitted sound waves

A test was performed to evaluate the error introduced by the method for detecting phase differences between the two sound waves. To observe this, sound cables were run directly from the sound card output to the A/D converter input. 20 data collections were performed; each collection consisting of turning on the sound wave generator, sampling the wave for 1024 samples at 25,100 Hz, and then turning off the sound wave. This data was then analyzed to calculate the means and standard deviations of the resulting phase differences. It was expected that there would be no phase difference between the two sound waves. Figure 3.3.6 shows the plot of the resulting data. It can be seen in this graph that the measured phase difference was as expected, and the standard deviation of the 20 measurements was small. This evidence shows that the phase difference measurement technique was not a significant source of error to the system.

Another simulation was run to test the effectiveness of the phase detection method. This was done in order to have samples of "perfect data" that could be used to gauge the error caused by the phase detector. This test used MATLAB to generate two pairs of sine waves of the expected frequency. This simulator also allowed the user to control the phase delay between the generated waves. The program generated data files similar to those created by the A/D controlling software. The first pair of sine waves was set to be 120 μsec apart. A second pair of sine waves was then generated with a phase difference of -154 μsec. As can be seen in Figure 3.3.7, the phase difference calculation worked well. The phase offset detection method correctly identified the phase difference,

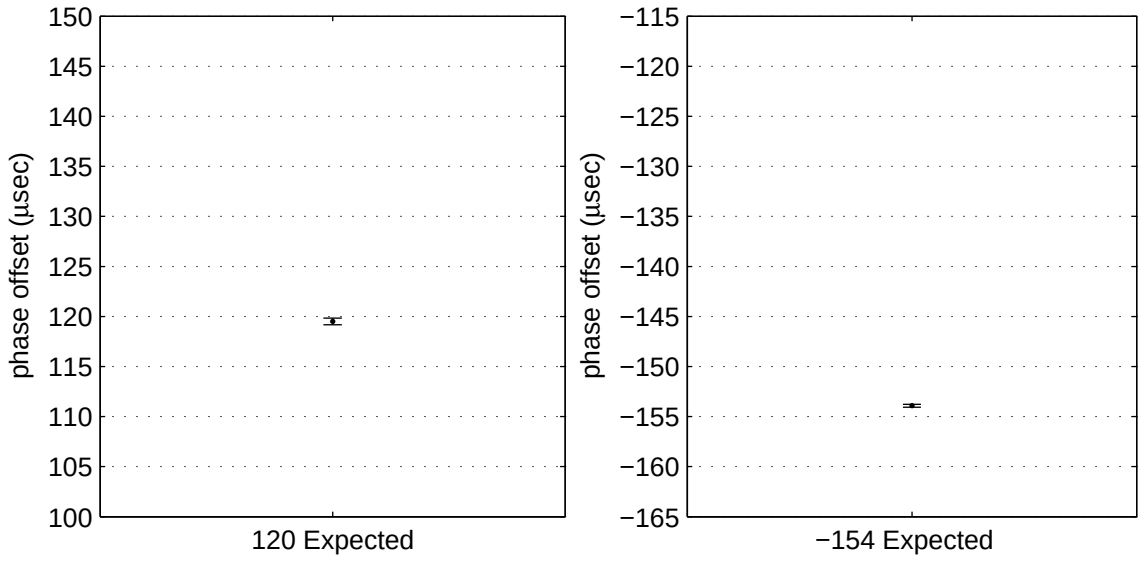

Figure 3.3.7 Measured phase difference of simulated waves

and there was a low standard deviation. This also supports the hypothesis that the phase difference measurement technique did not significantly contribute to the error of the carried-signal rangefinding method.

It was theorized that another possible cause of noise in this system was a difference in the lengths of the sound cables used to carry the 3 kHz sound wave to the transmitters. To test if the cable length affected the phase difference of the sound waves, an experiment was performed without the transmitter/receiver pairs. Sound cables of differing length were run from the computer's sound card straight to the A/D converter card. One cable was 38 m long, and the other cable was 2 m long. A second test was performed with both sound cables being 38 m long. The results of these two tests are shown in Figure 3.3.8.As can be seen here, the length of the sound cables did not affect the phase offset by a significant amount. This cannot then be a significant cause of error in this rangefinding system.

During the testing as the receivers were moved up and down the hallway, the excess cabling of their sound cords was not closely managed. It was hypothesized that perhaps the shape of the coil of this cabling somehow affected the measured results. A test was performed to observe whether changing the tightness of the coil of the receivers' sound output cables changed the phase delay of the received sound signals. For this experiment the two transmitters were again positioned 27.7 m

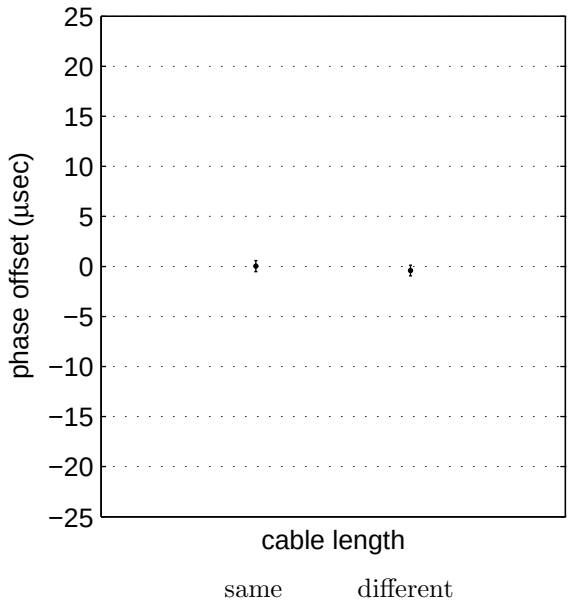

Figure 3.3.8 Measured phase difference of non-transmitted sound waves with different-length cabling

apart, and the receivers were located 13.9 m away from the reference transmitter. First, the phase offset was measured with no coil in the sound cables. Phase offset measurements were then taken for each of the following scenarios: forearm-sized coil placed next to receivers, hand-sized coil placed next to receivers, forearm-sized coil placed next to test-control computer, hand-sized coil placed next to test-control computer. The results of this test are shown in Figure 3.3.9. As can be seen in this graph, the coil shape and position both cause error in the phase offset, but not in any predictable pattern. It was suspected that the coils did not change the phase of the received sound waves being carried down the wires inside them. Instead, it was thought that the coils reflected the transmitted 52 and 52.5 MHz radio waves in ways that caused the travel time taken by these waves to drift. This hypothesis, however, is speculatory based on the erratic results of this test.

A test was performed to see if the FM frequency choice affected the observed phase difference. In this set of experiments, the radio at the end of the hallway was set to 52.5 MHz and the radio next to the computer was set in turn to 52, 53, 53.5, and 54 MHz. The receivers were located halfway between the two transmitters, or 13.9 m away from the reference transmitter. For each transmitter frequency, the phase difference between received sound waves was recorded. Figure 3.3.10 shows the results of this test. Note that the 52.5 MHz frequency was not tested because that was the frequency

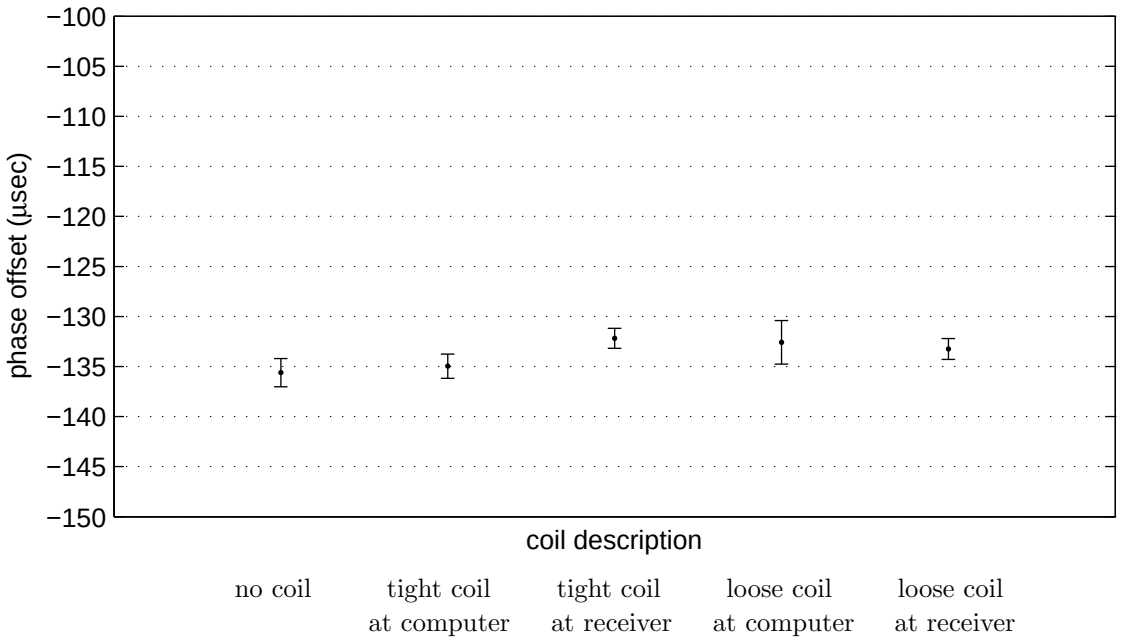

Figure 3.3.9 Measured phase difference with different coil shapes and positions

used by the radio at the end of the hall opposite the test computer. A larger frequency range was not tested because the FT-817 radios would not transmit at other nearby frequencies outside this range. It can be hypothesized that the use of different nearby transmitter frequencies does not cause significant error to the measured phase offset.

An experiment was done to see how strongly the environment affected the measured phase delay. In this test, the transmitters were placed in their standard locations, and the receivers were placed 7.8 m away from transmitter that was colocated with the test computer. The phase delay of the sound wave was recorded. The receiver pair was then moved approximately 5.5 meters perpendicular to the hallway into an adjacent room. Phase delay was calculated with the receivers in this position both with the door open and then the door closed. As can be seen in Figure 3.3.11, placing the receivers inside a room did cause some error in the phase offset. The position of the door, however, did not noticeably affect the measured results. These results indicate that walls in a building cause strong aberrations in the phase offset, while wooden doors do not.

Another set of experiments was performed to test how environmental factors affect the detected phase offset. These tests were designed to observe the noise caused by indoor objects. First, a standard measurement was taken in the hallway with the receivers located 13.9 m away from the

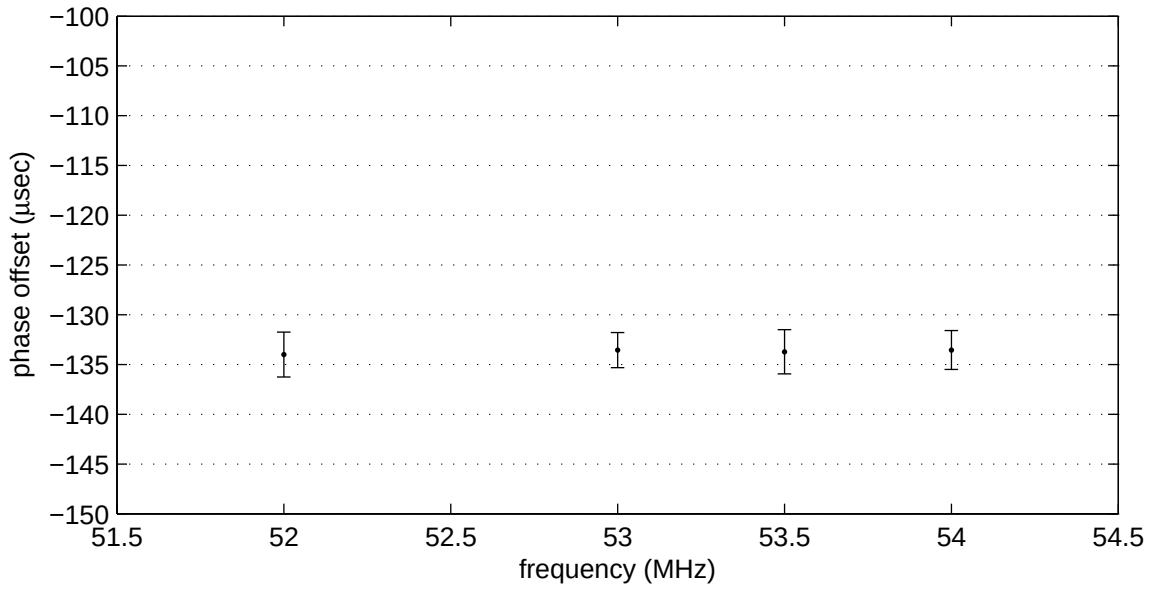

Figure 3.3.10 Measured phase difference with different transmitter frequencies

reference transmitter. Then, 4 school desks and a large metal furniture pushcart were positioned between the receivers and the 52.5 MHz transmitter at end of the hall opposite from the test computer. The phase difference was measured again. The results of this experiment are shown in Figure 3.3.12. The addition of desks and the metal pushcart had a large impact on the detected phase offset. The lowering of the phase offset from -136 μsec to -147 μsec coincides with an increase in the relative time taken for the signal from the 52.5 MHz radio to reach the receivers. This shows that the presence of metallic clutter in the tracking area introduces significant error into the carried-signal rangefinding method.

As shown in the experiments regarding the carried-signal rangefinding method, significant sources of error are present in typical indoor environments. Table 3.3.2 details the amount of error evidenced by the data collected during these tests. The combination of these errors presents a substantial barrier to the use of this system for indoor rangefinding.

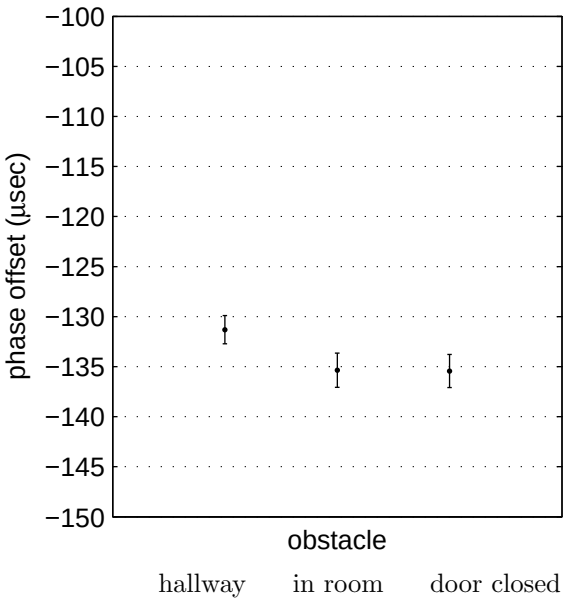

Figure 3.3.11 Measured phase difference with indoor obstacles

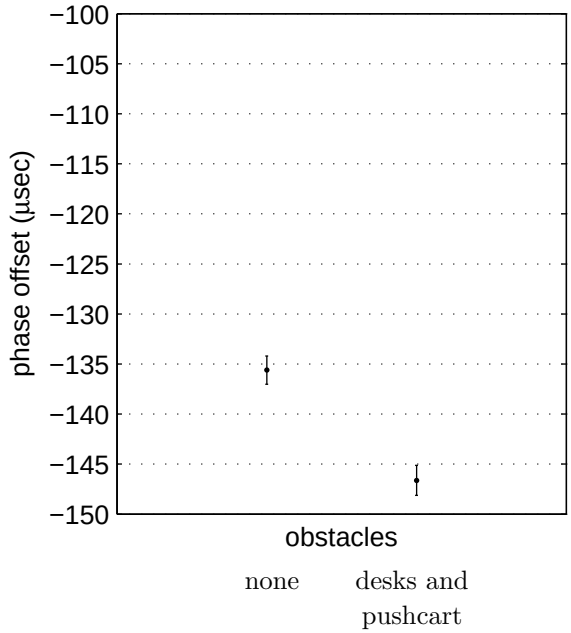

Figure 3.3.12 Measured phase difference with environmental clutter

Expected Change of Phase Across 27.7 m Range
± 1.0 μsec

Source of Error	**Amount of Error**		**Derived From**
Modulation Circuitry Delay	± 130.00	μsec	Figure 3.3.5
Phase Offset Measurement	± 0.25	μsec	Figures 3.3.6 & 3.3.7
Sound Cables of Varying Length	± 0.50	μsec	Figure 3.3.8
Sound Cable Coil Shape	± 7.00	μsec	Figure 3.3.9
Transmitter Frequency	± 0.50	μsec	Figure 3.3.10
Walls	± 4.50	μsec	Figure 3.3.11
Wooden Doors	± 0.10	μsec	Figure 3.3.11
Environmental Clutter (e.g. Desks)	± 11.50	μsec	Figure 3.3.12

Table 3.3.2 Estimated amount of error in carried-signal rangefinding method

CHAPTER 4

CONCLUSIONS

None of the methods tested worked as expected or as they were hoped for. However, two of the methods showed encouraging trends. The 802.11g packet-based rangefinding method worked well for long distances outdoors. The carried-signal rangefinding method also showed a definite correlation between the receivers' position in the hallway and the measured phase offset. Unfortunately, none of the three methods tested is suitable for creating a positioning system.

The 802.11g packet-based rangefinding method did not sufficiently handle indoor obstacles. The signals were able to penetrate through most of the walls in the residence where testing occurred without issue. However, the wall with the metal shelf proved to cause a significant unexpected drop in the bounce rate. Any building with large metal masses (a classification that includes many locations where tracking is desirable) would prove problematic for a tracking system based on this process. Even though no supporting data was shown in this thesis, it was also noted by the researchers that the 802.11g method was particularly fragile when competing devices were using the 2.4 GHz frequency band. Devices in this band include microwave ovens, portable telephones, Bluetooth devices (including the controller for a new video game console), and other WiFi networks. The prevalence of these interfering devices in modern environments makes the use of this frequency range in the design of a new system rather prohibitive.

The carrier-wave rangefinding method proved to be ineffective as a positioning technology. The only repeatable method for detecting distance using this method was the raw power level from a single transmitting antenna at close range. The expected interference between two transmitters was not observed. It was theorized that this was due to the transmitters not being perfectly synchronized. Since actual oscillators inside radios are not perfect devices, the frequencies they generate drift slightly depending on ambient temperature and other factors. Oscillators must therefore be synchronized using a device such as a phase-locked loop (PLL) if they are to be used in systems where signals generated from multiple sources must interact. Unfortunately it was not possible to synchronize the oscillators in the test equipment that was used.

The carried-signal rangefinding method also failed to work as intended. One significant problem was the discrepancy between the expected phase offset and the measured phase offset. After many experiments tried to isolate the cause of this discrepancy, it was finally decided that the cause was a difference in the amount of time required by the separate radio pairs to modulate and demodulate the audio tone. It was also observed that objects in the environment such as walls, desks, and other obstacles had a much greater impact on the observed phase offset that anticipated. It was originally thought that even though reflections of the desired radio wave may be present in the environment, the wave that took the direct path to the receiver from the transmitter would be sufficiently stronger to "swamp out" the competing reflections. This would tend to leave the desired straight-path signal relatively unaltered in phase as it reached the receiver. In practice, however, the reflections add to the desired signal in unpredictable ways, causing shifts in phase and amplitude that cannot be easily modeled.

Future researchers into this area should look away from the methods tried here. Technologies are needed that can pierce through indoor obstacles without affecting the signal characteristics that are used for the rangefinding process. Ultra wideband transceivers could potentially enable this type of system. These systems send signals spread across a very large portion of the electromagnetic spectrum. Some of the frequencies present in this type of transmission are likely to travel through obstacles with impunity. This could enable the system to ignore walls and obstructions so that it could accurately measure distances indoors. Another possible route for future researchers could be more complicated signaling schemes. The scheme used by the GPS system helps it to ignore reflected signals that may be just as strong as the desired ones. More research on this topic is required in order to find a definitive answer.

APPENDIX

ABBREVIATIONS AND DEFINITIONS

Dipole – A type of antenna that radiates equally well in all directions and is very simple in design. Commonly recognizable as an indoor, wall-mounted TV or FM radio antenna.

GPS – [NAVSTAR] Global Positioning System

GSM – Global System for Mobile communications

HF – High Frequency radio waves; between 3 and 30 MHz

ISM – Industrial, Scientific, and Medical frequency bands. Internationally reserved frequencies for purposes other than communications. Often used for license-free experimental equipment.

LPS – Local Positioning System

RSS – Received Signal Strength

UHF – Ultra High Frequency radio waves; between 300 MHz and 3 GHz

VHF – Very High Frequency radio waves; between 30 MHz and 300 MHz

WiFi – Licensed brand name used by the Wi-Fi Alliance to designate wireless networking equipment certified to meet the IEEE 802.11 standard.

REFERENCES

[1] M. Addlesee, R. Curwen, S. Hodges, J. Newman, P. Steggles, A. Ward, and A. Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8):50–56, Aug 2001.

[2] Klaus Finkenzeller. *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards and Identification*. John Wiley and Sons, 2003.

[3] GPS NAVSTAR. *Global Positioning System Standard Positioning Service Signal Specification*, 2nd edition, June 2 1995.

[4] I. Guvenc, C.T. Abdallah, R. Jordan, and O. Dedeoglu. Enhancements to RSS based indoor tracking systems using kalman filters. In *International Signal Processing Conference and Global Signal Processing Expo*, Dallas, TX, March 31 – April 3 2003.

[5] M. Hazas and A. Hopper. Broadband ultrasonic location systems for improved indoor positioing. *IEEE Transactions on Mobile Computing*, 5(5):536–547, May 2006.

[6] M. Hazas and A. Ward. A high performance privacy-oriented location system. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, pages 216–223, Tokyo, Japan, Mar 23–26 2003.

[7] A. Hoover and B. D. Olsen. A real-time occupancy map from multiple video streams. In *IEEE Conference on Robotics & Automation*, pages 2261–2266, May 10–15 1999.

[8] K. W. Kolodziej and J. Hjelm. *Local Positioning Systems: LBS Applications and Services*. CRC Press, 2006.

[9] L.A. Latiff, A. Ali, O. Chia-Ching, and N. Fisal. Development of an indoor GPS-free self-positioning system for mobile ad hoc network (MANET). In *Proceedings of IEEE International Conference on Networks*, pages 1062–1067, Kuala Lumpur, Malaysia, Nov 16–18 2005.

[10] S. L. Lewis. An examination of the history, expansion, and technology of the navstar global positioning system. *Massachusetts Institue of Technology*, Sep 2005.

[11] M. McCarthy, P. Duff, H. L. Muller, and C. Randell. Accessible ultrasonic positioning. *IEEE Pervasive Computing*, 5(4):86–93, Oct–Dec 2006.

[12] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil. LANDMARC: Indoor location sensing using active RFID. *Wireless Networks*, 10(6):701–710, Nov 2004.

[13] V. Otsason, A. Varshavsky, A. LaMarca, and E. de Lara. Accurate GSM indoor localization. In *Lecture Notes in Computer Science*, pages 141–158, Tokyo, Japan, Sep 11–14 2005. 7th International Conference on Ubiquitous Computing, UbiComp.

[14] N. Priyantha, A. Chakraborty, and H. Balakrishnan. Cricket location-support system. In *Proceedings of the Annual International Conference on Mobile Computer and Networking, MOBICOM*, pages 32–43, Boston, MA, Aug 6–11 2000.

[15] A. Shibuya, M. Nakatsugawa, S. Kubota, and T. Ogawa. A high-accuracy pedestrian positioning information system using pico cell techniques. In *Vehicular Technology Conference Proceedings*, pages 496–500, Tokyo, Japan, May 15–18 2000.

[16] C. Silwa. Wal-mart suppliers shoulder burden of daunting RFID effort. *Computerworld*, Nov 2003.

[17] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. In *ACM Transactions on Information Systems*, pages 91–102, Jan 1992.

[18] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, Oct 1997.