INDOOR LOCALIZATION IN WIRELESS SENSOR NETWORKS

BY AMAR H PATEL

A thesis submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Master of Science

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Dr. Marco Gruteser

and approved by

____

____

____

____

New Brunswick, New Jersey

October, 2008

ABSTRACT OF THE THESIS

Indoor Localization in Wireless Sensor Networks

by Amar H Patel

Thesis Director: Dr. Marco Gruteser

We considered the issue of indoor localization through the use of wireless sensor networks (WSN). We value not necessary the algorithm that provides the best accuracy but the one that provides a good enough level of accuracy in a simple and efficient manner. In the first part of our work we examined some state-of-the-art localization techniques that are deployable on wireless sensor motes. These techniques are evaluated to a set of criteria that an indoor WSN-based localization application must consider. In our investigation, we considered not only accuracy but many other factors that determine a suitable indoor WSN-based localization system. These factors among other things determine energy efficiency. We broadly separate the criteria list into two categories: efficiency-based and accuracy-based. We discovered that one of the techniques, Ecolocation, evaluates quite well to the efficiency-based criteria, but the evaluation of the accuracy-based criterions is not as promising. However, the inherent simplicity and potential for good performance (based on open environment results) make the algorithm quite attractive. We proceeded to modify the algorithm to improve its accuracy while maintaining its positive qualities. Our Weighted-Constraints algorithm, named as such due to the nature of the modification, performs in terms of average error 13.1% better than the original Ecolocation algorithm in an open environment. Furthermore, our modified algorithm shows it is more robust to noise compared to the original algorithm by perform on average 21.2% better in a noisy environment.

List of Tables

List of Figures

Acknowledgements

I would like to thank my thesis adviser Dr. Marco Gruteser without whom this thesis may never have seen its completion. I am grateful for his continued support and guidance throughout the whole thesis process. I wish to thank Dr. Predrag Spasojevic and Dr. Roy Yates for being on my thesis committee and providing me with comments from which to improve my work.

I want to express my gratitude to Dr. Roy Yates for his understanding and patience during the complicated first year of my thesis work.

I would also like to thank Dr. Ram Sastry for offering me a wonderful internship position at Siemens Corporate Research (SCR) for the past two years while I worked on my thesis.

I give great thanks to my friends for being there for me when I needed them the most. I am grateful to the Fourth Generation interns at Siemens that provided me some relief at stressful moments and to the Second Generation ones that supported and encouraged me throughout the whole process. I give special thanks to those who believed in me even when I doubted myself. I am very appreciative to my dry-run committee who provided helpful suggestions and much needed confidence during the presentation preparation phase. Finally, I would like to thank my lifetime advisor for always being a great listener.

Dedication

To all those who have helped me along the way with their love and support.

Table of Contents

Chapter 1

Introduction

In this thesis we study whether an efficiency-based localization technique for wireless sensor networks will be sufficient to provide accurate estimation in indoor environments. Efficiency is defined in terms of energy, time, and mote usage and accuracy in terms of best-case and robustness to corruption. Both terms are defined in more detail later.

1.1 Localization in Wireless Sensor Networks

Wireless Sensor Networks (WSNs) are characterized by the use of a large number of resource-constraint devices, called motes, deployed autonomously to monitor the environment ([7], [24]). The use of wireless device makes mobility an issue. That is, unlike wired sensors, the wireless motes are not tethered to a physical location and hence can be moved unknowingly. As a result, the monitoring data acquired from such a mote believed to be at Location A will be meaningless and destructive since it will in fact be coming from Location B. Localization is important to make the sensor data valuable when the position of the mote is in doubt. In a building, it is conceivable that certain sensors are in known static locations as they are monitoring for a specific reason (a temperature sensor monitor a particular section of a room) [16]. However, the nature of wireless sensors is that they can be easily added anywhere. As a result, it may be necessary to quickly deploy additional sensors in a room temporary to monitor additional section of the room. In such a scenario, it doesn't make sense to do time-consuming measurement to determine the location of these temporary non-static sensors. Furthermore, some wireless sensor motes are designed to be used in a mobile fashion. For example, a pulse rate sensor is attached to a patient and hence will move anywhere a patient moves. In such a case, localization can be used to determine where the patient currently is.

Although, much research has been done in localization [14], using wireless sensor networks presents the following unique limitations. Wireless sensor motes are not powerful devices in the sense that they are constrained in terms of computation power and memory. Additionally, we must consider energy efficiency to avoid depleting the battery life of the motes too quickly with cumbersome signal processing on the motes. However, there are also advantages that wireless sensor networks provides to localization. The high density, large mote deployment is an attractive feature. For example, in a building deployment, a large number of the deployed motes will be in known locations. As a result, these motes provide us with reference points we can exploit in the localization of the mobile motes.

1.2 Problem Statement and Contributions Summary

This Section offers a brief description of the main contributions the thesis aims to provide.

Indoor localization requires foot-level accuracy because of the confined nature of indoor environments. A typical office room may not be bigger than 100 square feet. As a result, location inaccuracies of more than a couple of feet would be very unacceptable in certain scenarios; at the same time, accuracies in the realm of a few inches may not be extremely necessary. Take for example a surveillance monitoring station in which rooms are equipped with sensors. In such a scenario, it is not necessary for the camera to point directly at the sensor when an event is detected since the station worker would have the ability to pan and tilt the camera, but it is important that the camera points in the general direction of the event given the camera's peripheral vision. As wireless sensor networks become more of a commonplace in indoor environments, such as homes and office buildings, they present an attractive deployment platform for localization services [10]. A small wireless sensor mote can be attached to any moveable object and use existing known location motes as reference points in the localization process.

To this extent, we investigate existing WSN localization techniques to find ones that can provided fine-grain accuracy yet are efficient in their nature. We identify seven state-of-the-art localization techniques that can be classified as WSN-compatible: Cricket, Active Bat, Centroid, A.P.I.T., Ecolocation, RF Time-of-Flight, and R.I.P.S. We evaluate simplicity based on a set of criteria we develop, outlined in Chapter 4. By wanting to use an infrastructure (reference motes) whose sole purpose is not localization, a non intrusive technique is necessary to avoid compromising the infrastructures primary function. Several of the techniques measure up favorably to one or two criteria, but evaluate poorly to the remaining criteria. Ecolocation is the only one of the techniques we studied that conforms well to most of the criterions and claims to have the potential to provide fine-grain accuracy. However, in the process of testing the accuracy of the algorithm in an indoor environment, we find that accuracy is reasonably good in a clutter-free environment, but the performance drops off dramatically as the environment is made noisier. An acceptable level of accuracy would be highly dependent on the localization situation being used. As a result, for testing purposes we define a reasonable accuracy level based on our reference mote deployment.

In the experiments, the reference motes are deployed in a grid fashion where the minimum separation between two motes in 4 feet. Therefore, we use an average error 2 feet as a rough guideline for determining if the error is reasonable. While some of the techniques we study claim to provide better indoor accuracy than Ecolocation, none of them can provide without requiring complex signal exchanges and processing. Such complex algorithms would require dedicated reference motes, since too much of a signaling and computation burden would be placed on motes that any other function would be disturbed. The simplicity of the algorithm and its ability to use reference motes passively is a very attractive feature of Ecolocation. Furthermore, in the best case conditions performance evaluation, Ecolocation performed reasonably well which means that the algorithm has the potential to provide the level of accuracy we are seeking. As a result, we modified the Ecolocation algorithm to improve accuracy in environments were data corruption is an issue without comprising the inherent simplicity of the algorithm.

1.3 Report Outline

The subsequent chapters of the thesis will proceed as follows.

In Chapter 2, we will give an overview of localization from the point-of-view of how techniques can be classified into different categories and an explanation of select prior localization techniques that paved the ground work for the field of study will be presented.

In Chapter 3, we discuss localization techniques that were developed with Wireless Sensor Networks in mind, by introducing techniques that work with existing commercial motes and those that require additional hardware integration.

In Chapter 4, we establish a set of criteria that indoor WSN-based localization techniques should possess and explain how existing solutions evaluate to the criteria.

In Chapter 5, we present the details that were involved to setup the experimental testbed.

In Chapter 6, we present the experimental performance analysis of the Ecolocation algorithm.

In Chapter 7, we introduce the design details for our modified algorithm and compare the experimental performance analysis with the original algorithm.

Chapter 8 presents the conclusions for the thesis and we close out the thesis by imparting some future work thoughts in Chapter 9.

Chapter 2

Localization Background Overview

Since location estimation is needed for wide ranging applications, much research has gone on in this field. As a result, a multitude of different approaches for location estimation have become known. These approaches vary from the type of signal used to the actual information extracted from the signal. The most common types of signals used are radio frequency (RF), acoustic, infrared (IR), or magnetic. Now given the type of signal used, there are a variety of parameters that can be determined to use in the localization method such as the received signal strength (RSS), time-of-flight (TOF), angle-of-arrive (AOA), or phase difference. Additionally, localization techniques differ in the way they use the extracted signal information. While some techniques use the information to determine the 'absolute' distance to a reference node (range-based), others make no assumption that the absolute distance can be determined by the information provided (range-free). The exact approach one chooses depends highly on the application involved. We will discuss these classification groups in turn in the subsequent Sections of this Chapter. Please note that, unless otherwise mentioned, we will refer to transmitters as the nodes which have known location and receivers as the nodes with unknown location.

Following our discussion about the different approaches to localization, we will discuss how a few historically relevant techniques that laid the ground work for localization utilized these approaches. The techniques we selected to highlight are GPS, RADAR, and Active Bat.

2.1 How Are The Metrics Are Used?

In this Section, we will discuss the classification group that divides localization techniques in to two mutually exclusive categories based on how they handle the gathered data: range-based or range-free. Range-based techniques try to strictly calculate distance values while range-free techniques try to strictly avoid calculating distance.

2.1.1 Range-Based

In range-based localization techniques (e.g. [5], [15], [6]), the goal is to try to use the observed signal metric (whether time, signal strength, angle, or other) of the transmitter's signal to estimate what the true distance between the receiver and each of the transmitters is. Once all estimated distances between receiver and transmitters are determined, some method of triangulation is used to determine the location of the receiver. In its simplest and perhaps most well-known form, the triangulation method determines the location of a receiver by finding the intersection of three circles [25]. By knowing the distances to at least three transmitters, these distances can be used as the radii of circles centered at the location of the transmitters. The location of the receiver can then simply be determined to be the intersection of the circles (see Figure 2.1).

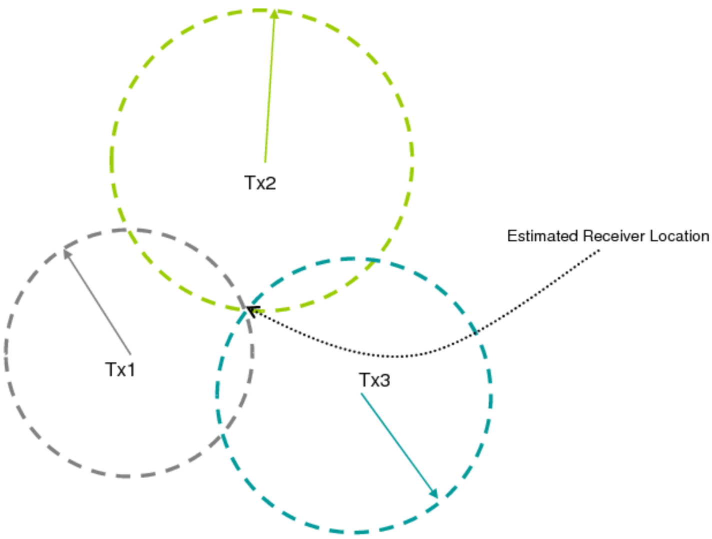

Figure 2.1: Simple triangulation of receiver based on three idea distance measurements to transmitters.

Now of course for this simple triangulation method to work, the distance measures must be perfect. That is, the circles will not intersect at one unique point when the radii are not of correct length. Naturally, there are many different methods available to solve such a problem. One common method to estimate position given a set of "noisy" range method is the non-linear least squares method [25]. The non-linear least squares solution can be found be solving the equation:

$$(\hat{x}_0, \hat{y}_0) = \arg \min_{(x_0,y_0)} \sum_{i=1}^{N} (\sqrt{(x_i - x_0)^2 + (y_i - y_0)^2} - d_i)^2 \tag{2.1}$$

where

$$\begin{aligned} d_i &= \text{“noisy” range estimate} \\ x_i, y_i &= \text{transmitter coordinates} \end{aligned}$$

That is, the estimate position is the location that minimizes the sum of the squared differences of the predicted range $\sqrt{(x_i - x_0)^2 + (y_i - y_0)^2}$ and observed range, d_i.

2.1.2 Range-Free

In range-free localization techniques (e.g. [9], [29], [19]), the algorithm is developed so that the receiver location can be estimated without the need to determine the distances to each transmitter. The motivation for such types of algorithms is that using that technique's observed signal metric it may not be possible to determine such a distance or if it is possible the determined distance may not be a reliable estimate. These types of techniques try to infer something like a 'pattern' in the location space. That is, they determine location by finding proximity to particular transmitters or determine what an expected signal environment should look like in certain places in the location space. The absence of a requirement of channel parameters is an attractive feature for those localization systems that desire low set-up or configuration overhead and ease of deployment.

2.2 Signal Metric Used to Localize

Now that we have seen how signal metrics are used in localization, let's take a deeper look at which of the signal metrics are typically extracted from a signal. Among the many parameters that can be extracted from a signal, the ones commonly used in localization techniques are time-of-flight, angle-of-arrival, and received signal strength (e.g. [15], [21], [8]).

2.2.1 Time-of-Flight

The time-of-flight ranging technique is a quite simple concept at its heart. To determine the distance between two points, one must just find the time it takes to travel from one point to the other. That is, given the velocity of propagation for the signal (typically a RF or acoustic), the distance between the points is simply the propagation velocity multiplied by the propagation time (assuming the receiver is not moving). However, one quickly learns that the problem of determining this propagation time is not quite as simple as the concept behind it. The complications that arise range from the ability to accurately determine the arrival time of the line-of-sight signal (LOS), since non-LOS paths take longer and do not represent the true distance between the transmitter and receiver, to the necessitate of precisely synchronized clocks. The actual time-of-flight value is typically determined via the use of a matched filter or correlation receivers. The time is estimated as the instant that the output of the match filter is peaked or the time-shift needed for a template signal to produce the largest cross-correlation (auto-correlation if no error) with the received signal [22]. Time-of-flight based techniques can be further broken down into two different categories depending on the level of synchronization among the devices. These categories are time-of-arrival (TOA) and time-difference-of-arrival (TDOA).

Time-of-Arrival

In the TOA technique, the transmitters and receivers must share a common clock. Therefore, a transmitter can send a time-stamped signal to the receiver which can determine the propagation time by calculating the difference between the arrival time and the time-stamp. If the transmitter and receiver do not maintain highly synchronized clocks, the time calculated by the receiver would be arbitrary and meaningless since it would be offset by the difference in the 'time' kept by each of them. Now multiplying the propagation time with the velocity, we determine the absolute distance to the transmitter (see Figure 2.2).

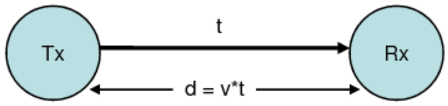

Figure 2.2: Distance calculation based on synchronized Time-Of-Arrival

In [22], they showed that the TOA ranging can achieve the inequality (in multipath-free channel):

$$var(TOA) \geq \frac{1}{8\pi B T_S F_c^2 SNR} \tag{2.2}$$

where

$$\begin{aligned} B &= \text{bandwidth (Hz)}, \\ T_s &= \text{signal duration (s)}, \\ F_c &= \text{center frequency (Hz)} \end{aligned}$$

Therefore, we can expect to receive higher accuracy with systems that use larger bandwidths or have higher SNR.

For completeness, it should also be pointed out that for fully unsynchronized systems, a round-trip TOA method can be used. In this method, both devices are transmitters and receivers; therefore, we will refer to them as Device One and Device Two. Device One starts by sending a signal to Device Two. After an internal delay, Device Two transmits its own signal to Device One. Device One measures the time it took from when it sent the signal to when it received the response. This time is twice the propagation time between the devices plus the internal delay time of Device Two (see Figure 2.3).

Therefore, once Device Two tells Device One its internal delay (whether it is a known

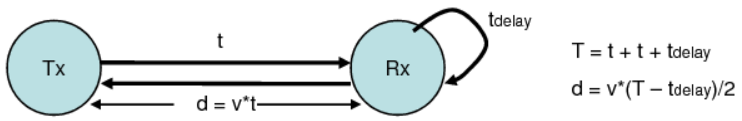

Figure 2.3: Distance calculation based on roundtrip Time-Of-Arrival

quantity or a measured), Device One can determine the distance between the two devices.

Time-Difference-of-Arrival

The TDOA technique can be used when the receiver cannot calculate the actual time it took a signal to travel from a single transmitter because there is an absence of synchronization between the transmitter and receiver. However, this technique requires that all the transmitters are synchronized to a common clock. Unlike with the time-of-arrival technique where the receiver determines the distance between itself and a transmitter, in time-difference-of-arrival the receiver tries to determine the difference in distance in relation to itself and two transmitters. That is, the receiver determines the signed value of the difference between the distance to one transmitter with the distance to another transmitter (see Figure 2.4).

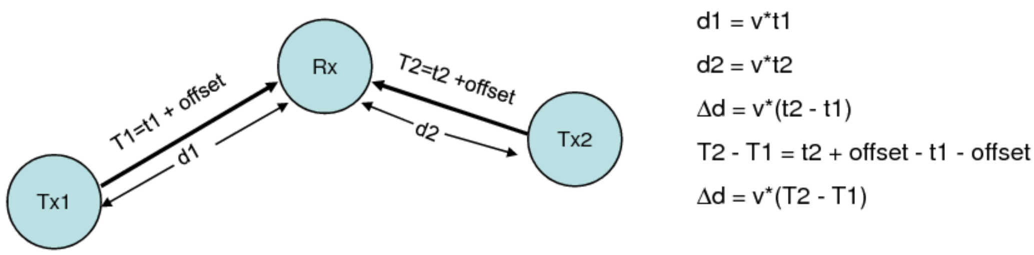

Figure 2.4: Distance calculation based on Time-Difference-Of-Arrival

Since the transmitters are synchronized, the receiver can simply subtract the 'perceived' time-of-arrival from each of the transmitters and effectively eliminate the clock offset in the unsynchronized transmitter-receiver pair to determine the time difference. As stated before, the distance difference can be determined by multiplying the time difference with the propagation velocity. Therefore, we now know that the receiver lies on a locus of points which have a constant difference in distance to two fixed points. This curve is simply a hyperbola with the two transmitter locations as the foci. The estimated location can be determined by the intersection of several such hyperbolic curves.

2.2.2 Angle-of-Arrival

The AOA technique exploits the fact that a triangle can be defined by one of its side's length and the two angles at the ends of that side. Therefore, using two transmitters as two point of a triangle (the receiver will make up the third point), we can know one side length of the triangle. Now if we know the angles between the two transmitters and the receiver, we can determine the location of the receiver uniquely (see Figure 2.5).

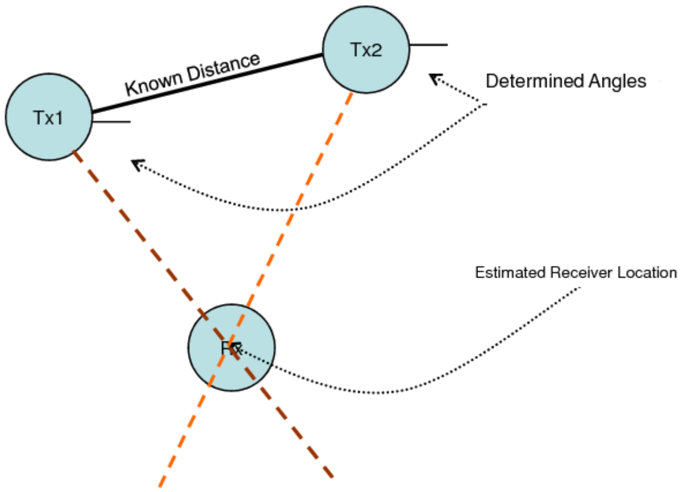

Figure 2.5: Estimating receiver location using Angle-Of-Arrival

As a result, at the root of the AOA technique is the method to determine the direction of the signal propagation incident on the receiver. To determine the angles, there are two common methods used [22]. Antenna arrays are the most popular solution for AOA. A transmitter sends out a single signal which is picked up by the antenna array. Each antenna in the array determines the arrival time of the signal. By using the time difference in the arrival times, the array knows the relative phase differences observed at each antenna (which are in known locations relative to the device). Therefore, with the known positions of the antennas and the difference in phase, the incident angle can be determined.

2.2.3 Received Signal Strength

The RSS method works on the principle that the energy of a signal decreases as it move farther away from its source. Therefore, if one knows the relationship between the decrease in energy and the distance traveled (i.e. a path-loss equation), the distance between the receiver and transmitter can be calculated from the signal strength observed at the receiver. One simple form of the path-loss equation is:

$$P(d) = P_0 - 10n \log \frac{d}{d_0} \tag{2.3}$$

where

$$\begin{aligned} n &= \text{path-loss exponent,} \\ P_0 &= \text{power at distance } d_0 \text{ (dBm)} \end{aligned}$$

However, a major problem with RSS is that, unlike with the simple form presented above, the power attenuation in realistic conditions is not monotonically decreasing . Instead, the observed signal strength at the receiver suffers from problems in multipath and shadowing, just to name a few. Therefore, in environments with many obstructions, it becomes quite difficult to accurately infer distance based on RSS. As a result, for indoor propagation, most people added an extra uncertainty variable into the simple path-loss model called the shadowing effect. The model is expressed as:

$$P(d) = P_0 - 10n \log \frac{d}{d_0} + X_\sigma \tag{2.4}$$

where

$$X_\sigma = N(0, \sigma^2)$$

Environments that suffer from large shadowing effects are expressed with a higher variance Gaussian random variable. In [11], they show that distance estimation based on RSS achieve the inequality:

$$\sqrt{var(\hat{d})} \geq \frac{\ln 10}{10} \frac{\sigma_s h}{n} d \tag{2.5}$$

where

$$\sigma_{sh} = \text{standard deviation of shadow effect,}$$
$$d = \text{distance between devices}$$

Therefore, we can expect lower accuracy for larger ranging distance and in environments with more obstructions (represented by a higher standard deviation due to shadowing). Despite the short comings of a RSS based method, much research has gone into finding effective ways to use it in localization. The main reason for its appeal is the relative easy of use and inexpensive cost. That is, most wireless devices already have a Received Signal Strength Indicator (RSSI) circuit build into them. Therefore, no additional hardware is required for such systems and because RSS can be measured from any ordinary communications packet, you may not even require additionally communication bandwidth.

2.3 Signals Types Used in Localization

After having an understand about the primary metrics preferred by localization techniques, it is important to know from which type of signals these metrics are extracted from. We will focus on the two main types of signals used in localization techniques: radio frequency and sound. Although there are techniques that use other signals, such as infrared or magnetic, they are quite rare and as a result, we do not focus on them. However, in Section 2.4.3, we will discuss the use of infrared in Active Badge.

2.3.1 Radio Frequency

Radio Frequency (RF) signals are electromagnetic waves in the frequency range from 3 Hz to 300 GHz . There are several properties about RF propagation that are pertinent to their use in localization techniques. Firstly, the signal power weakens as it traverses in distance away from the source because the wavefront of the signal energy expands as it traverses. The free space loss in signal power of a wavelength λ signal at a distance d can be represented as:

$$L_{fs} = (\frac{4\pi d}{\lambda})^2 \tag{2.6}$$

In addition to the free space loss, signal power also attenuates as it passes through objects such as walls because the objects absorb some of the signal power. One very important problem that is critical in relation to localization is multipath. Multipath is the process by which a signal can interfere with itself. That is, a signal can travel in a non-line-of-sight path by bouncing off of objects and combine with the line-of-sight signal. The result of the latter two properties is that signal attenuation is no longer a deterministic relationship with distance. Another important characteristic of RF propagation is the bandwidth of the signal. The bandwidth of a signal is basically the difference between the highest frequency component and the lowest frequency component in the signal.

Narrowband vs. UWB

Narrowband signals are generally regarded as those whose message bandwidth is quite trivially compared to the channel's full bandwidth. Conversely, when the message bandwidth is much larger than the channel bandwidth, the signal is referred to as wideband. Ultra-Wideband (UWB) signals have bandwidth in excess of 500 MHz or 20% of their center frequency. Narrowband signals typically transmit message data by modulating it onto a carrier wave (normally a sinusoid). UWB transmission doesn't require any carrier. Instead, the message is sent via extremely short pulses spaced accordingly. From the point of view of localization, the type of RF signal used can provide some unique advantages.

The use of short duration pulses by UWB allows the combating of multipath effects since any reflected pulse arriving after a duration period from the arrival of the line-of-sight pulse can be properly detected as multipath components. For example, when a line-of-sight pulse is able to be detected, UWB is better equipped to be able to determine the time-of-flight since the direct path can be resolved from the multipath components. That is, since the time-of-flight is determined as the peak position of the correlation of the transmitted signal and the received signal, any corruption of the received signal will shift the peak away from the true position. Narrowband signals are highly susceptible to multipath interfere which causes the received signal to not only be the time-shift version of the transmitted signal, but also the summation of all the multipath components. Additionally, as we saw in Section 2.2.1, we can expect higher time-of-flight accuracy with larger bandwidths. However, one problem with UWB signals is that communication range is significantly shorter than with narrowband signals.

2.3.2 Sound

A sound signal travels through the air as a wave of mechanical energy disturbance caused by alternating pressure. These signals can vibrate at various frequencies. Typically humans can hear sounds in the frequency range from 20 to 20,000 hertz which is referred to as acoustic sound. Sound at a frequency greater than 20,000 hertz is referred to as ultrasound. Since ultrasound is inaudible to humans, it is the typical form of sound signal used in localization so as to remain stealthy and non-hostile to humans. The speed of sound in air is about 343 meters/s. Considering the relatively slower speed of sound compared to RF signals, it makes sound signal propagation time calculation less complex and a more attractive solution than using RF signals. However, the use of high frequencies causes the range of the signals to diminish compared to the audible range frequencies. Additionally, the high frequency causes the sound waves to travel in a narrow beam causing a directionality dilemma.

2.4 Primordial Localization Techniques

The following select historic techniques attempted to use the different principles outline above in their own manner to determine location. Although these techniques are not fully feasible for wireless sensor networks, their underlying methodologies provide valuable insight into the problem of localization.

2.4.1 GPS

The Global Positioning System (GPS) is probably the most popular and best known localization technique in the industry [5]. GPS is a radio frequency-based time-of-flight system that is primarily suited for outdoor localization using earth orbiting satellites as transmitters. By determining the distance to the satellites via a time-difference-of-arrival calculation, the receiver can determine its location through triangulation. The system is made up of three separate segments: a space segment, a control segment, and a user segment. The space segment consists of at least twenty-four orbiting satellites which are highly synchronized via atomic clocks. Each satellite follows a well-defined orbit; therefore, their absolute locations can be determined and the orbits are defined such that at least four of them are always reachable at any location and at any time. The satellites continuously transmit a unique code which consists of the orbital location and clock information (called Navigation Message). The control segment consists of four monitoring stations and one master control station. The monitoring stations continuously receive data from the satellites and pass them to the master control station. The master control station analysis the data and sends correction information about the orbital and clock to the satellites. The user segment is simply the receiver trying to find its location. By using the orbital information, the receiver can determine the location of each satellite it communicates with. The satellites transmit a pseudo-random noise code of which the receiver has a locally generated copy. To find the time-of-arrival, the receiver just determines how much time shift is required to match the code. Time-difference-of-arrival technique is used to determine the location because the receivers are not as highly synchronized to the satellites. One major problem GPS faces is the time delay cause by the atmosphere because the signal slows down. The accuracy is also affected by the number of satellite and the relative position of them. That is, accuracy is better when more satellites are present and when the satellites are spatially diverse at wide angles. And of course error can be cause by incorrect orbital and clock information (the job of the control segment is to minimize this source of error). With these sources of error in consideration, the accuracy range is typically 6-to-12 meters.

Differential GPS

Differential GPS (DGPS) was developed as an enhancement to GPS to reduce some of the sources of error. Conventional GPS use a model to estimate the atmospheric delays; however, DGPS can provide actual delay information. DGPS works by using a network of reference stations (GPS receivers) at fixed location on the earth. The system works by calculating the ranging error associated with satellites by comparing the range determined by the received signal to the range determined by the actual position location. This correction information is transmitted to the DGPS receivers who can use the information to improve their location calculation. DGPS reduces the position error down to a range of 1-to-5 meters.

2.4.2 RADAR

One of the earliest and most popular localization techniques for indoor location estimation based on signal strength called RADAR is presented in [8]. The technique is deployed using a buildings' existing wireless LAN network. In the paper, the authors present two implementation methods for RADAR. There are three base stations on the floor that together cover the entire location space. In both methods, the receiver collects an RSS triplet for the basestations. However, one of the implementations uses "empirically-determined" data to determine the most likely location of the receiver based on the triplet, where as, the second implementation uses "theoretically-computed" data to determine the most likely match for the triplet. The two methods represent a trade-off between accuracy of location estimation and ease of deployment.

Pattern Matching Method

The pattern matching method of RADAR requires two phases. In the initial phase, a database of RSS values is created by going to various points in the location space and storing the signal strength of each basestation at the point (averaging multiple beacons). Additionally, at each location point, the orientation of the receiver (north, south, west, and east) is changed and RSS values at each orientation are stored. The authors generate the database based on all four directions at seventy location points. The final phase of the method is to determine the location of the receiver by searching the database created in the initial phase. Basically, after determining the RSS values observed from each of the basestations, a nearest neighbor search is done on the database to determine the most likely location. The authors claim that the method provides sub 3 meter location error 50 percent of the time.

Triangulation Method

The authors propose the triangulation method to eliminate the need for the initial phase of the pattern matching method. In this method, the database is created using an indoor signal propagation model. Using this model, the signal strength data in the database at each location is theoretically calculated. In essence, they are estimating distance based on signal strength. That is, once the receiver obtains the RSS values of each basestation, it searches the theoretical database to determine the closest match. Furthermore, the receiver discovers the location it should be at if the observed RSS values were used to determine the ranging from the basestation based on the path-loss model. They study several model and adapt Floor Attenuation Factor Model to the Wall Attenuation Factor model:

$$P(d) = P_0 - 10n \log \frac{d}{d_0} - \begin{cases} nW * WAF & \text{if } nW < C \\ C * WAF & \text{if } nW \geq C \end{cases} \tag{2.7}$$

They define C to be the maximum number of walls up which the attenuation makes a difference, nW is the number of walls between receiver and transmitter, and WAF is the wall attenuation factor (which is empirically determined for the environment). This theoretically method provides an error of 4.3 meters 50 percent of the time.

The triangulation method results in an error of more than a meter as compared to the pattern matching method. However, the less accurate results are compensated by the fact that an extensive empirical data collection phase is not required. Only the empirical estimation of a channel parameter (WAF) is necessary.

2.4.3 Active Badge

The Active Badge system proposed in [26] is one of the earliest attempts at performing indoor localization. The system is proximity based and allows room level granularity without the real possible of achieving fine localization if desired. The system makes use of infrared technology to achieve the receiver to transmitter communication. The unknown object in this system is a small badge that transmits an infrared signal while the known reference location is a sensor that detects the signal. Therefore, in this section we will explicitly refer to the unknown location as a badge and the known location as a sensor (since the receiver/transmitter roles are reversed in reference to our previous stipulation). In the system, each badge periodically transmits a unique infrared signal (every badge has its own code) to identify itself. The signals are picked up by fixed location infrared sensor around the building. All the sensors are networked together to a central server that aggregates the sensor readings to provide location information about each badge in the system. Since infrared signal don't penetrate through walls, the badge can be accurately located to within room boundaries.

2.5 Chapter Summary

In this Chapter, we have given a broad overview of the different factors the one can consider for a localization technique. We have seen that some techniques try to make pair-wise distance estimations for the end goal of determine location, whereas, others try to infer location by processing the gathered data without making distance estimations. Additionally, we have observed some of the helpful metrics that can be extracted from signals and also some of the challenges that these metrics face when trying to be useful, such as synchronization or corruption. Finally, we looked into the common signal types used by localization techniques along with there inherent issues when used in practice, like obstructions or audibility.

GPS's requirement for sophisticated hardware prelude direct use of in wireless sensor localization. However, the essential time-of-flight methodology for determining location can be universally used if a less stringent method, in terms of hardware, for determining time was implemented. Similarly, it is conceivable to implement the RADAR's RSS methodology on wireless sensor motes given we can shift the computationally expensive database and search algorithm to a central server. Additionally, with wireless sensor networks dense population it is reasonable to think with more known location basestations, the accuracy could be improved. Active Badge use of infrared prevents it from direct use in wireless sensor networks. However, coarse localization can be implemented in wireless sensor networks using RSS data instead. If one could strategically install motes in every room that sent periodic beacons to that room only (adjusting power level to avoid spillover into next room), a mobile mote could be localized to the room that corresponds to the high strength RSS receiver (highest strength since it is reasonable to assume the any spillover would have low strength).

Chapter 3

State-of-the-Art Wireless Sensor Networks Localization Techniques

In this Chapter, we will discuss some of the state-of-the-art techniques proposed for localization in wireless sensor networks. We have split these techniques up into two categories depending on their level of hardware requirements. One category of techniques requires no additional hardware modification to existing mote technologies. That is, the hardware already equipped in commercially available wireless sensor motes will be sufficient for deployment of the localization techniques. The other category of techniques requires that additional specialized hardware is necessary for use of the techniques. We first present the category of techniques that require additional hardware.

3.1 Require additional hardware

We first investigate two novel techniques that are in theory suitable to be used with Wireless Sensor Networks: Cricket and Active Bat. However, both techniques make use of acoustic signaling via hardware not typically found on commercial wireless sensor motes.

3.1.1 Cricket

The Cricket Location Support system proposed in [23] is a time-of-flight based system that makes use of both RF and acoustic signals. Cricket was developed with a specific emphasis on user privacy and scalability. To achieve the scalability goal, the system was designed in a decentralized manner in which the individual components in the system need not be kept accounted for by a central unit. That is, a transmitter (or receiver) can be added into the existing system without the need to be configured and installed by a special administrator of the overall system. The system is such that a newly installed transmitter need not follow strict guidelines to avoid causing chaos to existing transmitters. To insure the privacy of the receiver, the system is designed such that the receiver is the only one who has all the necessary information to determine location. In this way, the receiver performs all the localization calculations necessary and determines its own location, and as a result, a central database cannot store the receiver's location unless the receiver chooses to publish its own location.

Cricket exploits the orders of magnitude difference between radio frequency and acoustic signals to calculate the time-of-flight range. The transmitters will broadcast an RF signal while at the same time transmit an ultrasonic pulse. Basically, the RF signal with its much faster speed will act as a time synchronizing mechanism while the ultrasonic pulse is the signal used to determine time-of-flight. The receiver will hear the RF signal and wait for the ultrasound pulse. The time difference between the first reception of the RF signal and the ultrasonic pulse can be estimated as the time-of-flight of the ultrasonic pulse. Also, to define the time interval the receiver should expect the ultrasonic pulse to arrive, the authors used a low data rate message. The message is design so that any receiver in range to hear the RF signal should hear the ultrasonic pulse before the completion of reception of the RF signal. This eliminates the possibility that the valid ultrasonic pulse is lost and a stray pulse from another transmitter is not falsely considered. In addition, to alleviate the interference resulting from the decentralized and uncoordinated transmitters, the authors implement a randomized algorithm for broadcasting by the transmitters.

The Cricket system employs both a proximity method and allows for fine-grain localization through triangulation. The proximity method is implement by having the transmitters determine regions in the location space. When a transmitter is configured, it is coded with a specific string that the owner determines as referring to the region.

> Whenever a beacon is placed to demarcate a physical or virtual boundary corresponding to a different space, it must be placed at a fixed distance away from the boundary demarcating the two spaces. [23]

In this manner, a receiver determines the region it is in by using proximity information about the transmitters. Simply put the receiver associates with the transmitter it determines is the closest in distance (the smallest time-of-flight value). Furthermore, if the receiver chooses to determine its fine-grain location in terms of coordinates, it can simply use the ranging information (distances calculated by time-of-flight) it has about the transmitters, and their position information, to perform a triangulation to determine the absolute location.

3.1.2 Active Bat

The Bat System proposed in [12] and [27] is a time-of-flight based system that makes use of both RF and acoustic signals. Like with Cricket, the Bat system exploits the great orders of magnitude difference between the speeds of propagation of a radio signal and an acoustic signal. By using the much faster RF signal to synchronize the transmitter and receiver, the system is able to use the acoustic signal to perform a ranging calculation. In this section, we will explicitly refer to the unknown location objects as Bats and the fixed location reference objects as Ultrasound receiver units. The Bats are constructed from a radio transceiver, microprocessor, and a couple of ultrasonic transducers. Each Bat is also uniquely identified by an ID. The Ultrasound receiver units consist of an ultrasonic detector and a serial network interface. A grid of Ultrasound receiver units are mounted on the ceiling and connected to each other through the serial interface. There is a central computer that has a base station attached to it and is also connected to the receiver unit network. Periodically the base station sends out a RF signal with the unique code of the Bat being localized and at the same time sends a reset to all the Ultrasonic receiver units through the wired serial network. Upon the Bat receiving the RF signal, it transmits an ultrasonic pulse to the grid of receivers. The Ultrasonic receiver units will determine the time-of-flight by measuring the difference in time from the reset signal to the detection of the first ultrasonic pulse. To eliminate the use of reflected ultrasonic pulses, the authors implemented a "statistical outlier rejection algorithm." Once the time-of-flight measurements are determined the Ultrasonic receiver units send the information to the central computer for processing.

The location estimate is computed by a process of multilateration.

Cricket and Active Bat are quite similar in their underlying methodology as both use some form of time-of-flight ranging using RF and acoustic signaling. However, Active Bat relies on a central computer for determining location, while Cricket focus on a decentralized approach. Additionally, Active Bat has the unattractive feature of requiring a wired network of receiver units.

3.2 Work with existing mote hardware

We will now explore five techniques that can be readily deployed using the existing mote hardware: Centroid Method, Approximate Point-In-Triangle, Ecolocation, RF Time-of-Flight, and RIPS. These techniques make use of the radio hardware typically found on motes along with varying degree of complexity in their algorithms.

3.2.1 Centriod

The Centroid method proposed in [9] is a RSS based method which is range-free. It use a simple proximity based algorithm that can provide coarse-grain location estimates while requiring almost no setup. The estimation technique is simply that the receiver localizes itself to the centroid of the transmitters that are in its “proximity range.” The “proximity range” is determined by a connectivity metric which is the number of beacons received by the receiver divided by the number of beacons sent by the transmitter. If the connectivity metric is above a certain threshold, then the transmitter is determined to be in “proximity range.” Once the N “proximity range” transmitters are determined, the receiver localizes itself to the location:

$$(X_{est}, Y_{est}) = \left(\frac{X_1 + \cdots + X_N}{N}, \frac{Y_1 + \cdots + Y_N}{N} \right) \tag{3.1}$$

Because the algorithm relies on a spherical radio propagation model to setup boundaries in the connectivity range, the authors test the validity of the connectivity in indoors and outdoors environment. They reach the conclusion that the model is not valid for indoor, but performs reasonable well for outdoor environments. As a result, they perform outdoor test with four reference nodes at the corners of a 10 meter-by-10 meter area (nodes of range about 9 meters). Taking measurements at 1 meter intervals in both directions, the results were an average error of 1.83 meters and a maximum error of 4.12 meters.

3.2.2 Approximate Point-In-Triangle

The approximate point-in-triangle technique proposed in [13] is a RSS based method which is range-free. The basic procedure is to determine if the receiver is in a triangle formed by a set of three transmitters. This "in-triangle" determination is done for all set of three transmitters possible. The estimated location is the center of gravity of the maximum region of all overlapping "in-triangles." The theory the authors use for determining whether a receiver is in a triangle is based on what they call the "Perfect P.I.T. Test Theory" which is stated as:

> If there exists a direction such that a point adjacent to M is further/closer to points A, B, and C simultaneously, then M is outside of ΔABC. Otherwise, M is inside ΔABC [13].

Although the perfect P.I.T. would guarantee the correctness of the decision, the authors develop an approximation of the test because in practice the perfect P.I.T test would not be possible. They base their approximation on the weak assumption that signal strength is monotonically decreasing for a particular direction. That is, between two neighboring receivers, the receiver closer to a transmitter should observe higher signal strength (though no assumption is made that the values will follow a well-defined path-loss model). The "Approximate P.I.T Test" is defined as:

> If no neighbor of M is further from/closer to all three anchors A, B, and C simultaneously, M assumes that it is inside triangle ΔABC. Otherwise, M assumes it resides outside this triangle [13].

In essence, the neighbors are acting like the movement of the receiver (in finite directions) trying to determine if it is in the triangle in the perfect P.I.T. test. Therefore, the neighbors communicate their observed RSS values from the transmitters with the receiver, the one trying to localize itself (though the neighbors are also localizing themselves via communications with their neighbors, for simplicity we refer to the receiver as the only one localizing itself). Based on the comparison with its own observed RSS values, the receiver determines which neighbors are closer to/farther from which transmitter relative to its own location. The authors point out that the APIT test suffers from the fact that, unlike with the perfect P.I.T. test, only a finite number of directions can be tested, represented by the direction of the neighbors. The location estimation is determined by the maximum triangle overlap. The method they propose for this determination is making a location space grid and incrementing the grid points which correspond to an "in-triangle" and decrementing the ones that correspond to an "out-triangle." The center of gravity of the grid points that have the maximum value is estimated to be the location of the receiver.

In the paper, the authors don't provide any experimental results for the algorithm. However, they do present simulation results based on the varying of different system parameters. They perform the simulations with three other range-free schemes, namely, Centroid [9], DV-Hop [20], and Amorphous Positioning [18]. They conjure that A.P.I.T is most desirable when there is random transmitter placement and if there is an irregular radio pattern in addition to low communication overhead.

3.2.3 Ecolocation

The Error COntrolling LOCAlizaTION (Ecolocation) algorithm proposed in [29] is an RF based range-free localization technique. The key idea behind the algorithm is that the localization space can be divided into distinct non-overlapping regions based on the ordered sequence of transmitters. The ordered sequence of transmitters is defined as the sequence of the transmitters ordered from closest to farthest. That is, every point in the localization space can be associated with one and only one ordering of the transmitters based on each transmitters distance from that point.

Each transmitter sends out beacons and the receiver determines the RSSI of each received signal. An ordered sequence is generated based on the RSSI. Now for each point in the location space, a theoretical distance based ordered sequence is generated.

Both ordered sequences are converted into a constraint matrix. The constraint matrix is essentially an n-by-n (n is number of transmitters) matrix of -1's, 0's, and 1's where the elements obey:

$$element(i,j) = \begin{cases} -1 & \text{if transmitter i is farther than transmitter j} \\ 0 & \text{if transmitter i and transmitter j are equidistance} \\ 1 & \text{if transmitter i is closer than transmitter j} \end{cases} \tag{3.2}$$

Equation 3.2 implies that the RSSI constraint matrix is generate by the sign function as, $element(i,j) = sign(RSSI(i) - RSSI(j))$ and the theoretical constraint matrix is generated based on the negative sign function as, $element(i,j) = -sign(dist(i) - dist(j))$.

For every location space point, a theoretical constraint matrix is generated and compared with the RSSI constraint matrix. The number of constraint matches is determined and the locations with the maximum number of matches are found. The centroid of these locations is calculated and used as the estimated location. The algorithm is able to overcome minor errors in the RSSI because it uses the relative comparison rather than absolute values. That is, even if the absolute values are corrupted, as along as the relative rank is still correct the algorithm is unaffected. Furthermore, the technique allow for correcting errors through the redundancy in constraints. Hence, the name Error COntrolling LOCAlizaTION because of the similarity to error controlling codes.

The authors perform experiments in indoor and outdoor environments. In the outdoor environment, they randomly placed eleven Crossbow Mica2 [1] motes in a 12 meter-by-12 meter open parking lot. They localized each mote by using the other ten as reference transmitters. To quantify the results, they used a metric called percentage of average inter-reference nodes distance. The maximum error observed was about 80%; however, the other ten localization points resulted in error of sub 25% with four of those sub 5%. An indoor office building scenario representing a cluttered RF environment was also tested. They placed twelve reference nodes along a corridor and in a room. They performed the test at five location points. The results showed the maximum error was about 50% with three others in a 30-40% range and the lowest at about 10%. As expected, the technique works better in a relatively open RF channel with low interference.

Sequence-based

A similar technique based on ordered sequences was proposed by the same authors in [28] called Sequenced-Based Localization (SBL). Again, the transmitters send out beacons which the receiver uses to generate an ordered sequence based on the RSSI. However, the determination of location is not performed through a constraint matrix method. Instead, a pre-defined "location sequence table" is used to determine the closest match sequence to the RSSI determined sequence. The location sequence table is a mapping of all feasible sequences in the location space to a unique location. Basically, the location space (in two-dimensions) is divided into distinct regions by the perpendicular bisector of each line joining pairs of transmitters. As a result, each region is defined by a unique ordered sequence of the transmitters by distance. The centroid of the region is taken to be the unique location defined by that ordered sequence. The authors provide a pseudo-code of determining the location sequence table. The technique exploits the fact that the number of feasible ordered sequences is much less than all possible ordered sequences. Given all ranking of the transmitters is possible the number of possible sequences is $O(n^n)$; however, the authors prove that the number of feasible sequences is $O(n^4)$. Therefore, an erroneous RSSI based ordered sequence can be corrected by finding the nearest feasible sequence in the location sequence table. The authors use two methods to determine the nearest feasible sequence: Spearman's Rank Order Correlation Coefficient and Kendall's Tau. They notice through simulation results that Kendall's Tau performs slightly better than Spearman's correlation as the number of reference nodes increases.

3.2.4 RF Time of Flight

In [15], the authors propose a system of measuring the time of flight in sensor motes based on radio frequency signals. Their design is for narrowband signals in the 2.4 GHz ranges and based on future 2nd generation motes because they require access to the signal at the physical layer. Because tight clock synchronization is not possible in ad hoc wireless sensor networks, the approach proposed uses roundtrip time-of-flight measurements. Since the choice is to use narrowband signals, they conclude that for accurate time-of-flight calculation they should determine the phase offset of the signal. As a result, they choose to use pseudorandom noise (PN) code signals because of their property to have a large peak at the phase offset.

The system works in two phases: online measurement phase and offline range extraction phase. The system uses a clock which is the same as one period of the PN code. Because both devices will act as receivers and transmitters in the roundtrip scheme, we will refer to them explicitly as Device One and Device Two. First the devices are synchronized to the level of one clock period. Then Device One transmits multiple copies of the PN code modulated onto the carrier signal. The system requires that Device One start transmitting after waiting one clock period after synchronization. Device Two starts receiving these copies after waiting two clock period from synchronization and averages the demodulated signal and creates a local copy. This averaged signal is then transmitted in multiple copies back to Device One using the same principle of waiting. Device One creates the received signal by averaging the copies. Now both devices have finished the first phase and the second phase can begin. The second phase is simply determining the cross-correlation between the estimated received signal by Device One and the initial transmitted signal by Device One. The peak of the function is found to determine the round trip time of flight. To combat multipath errors, they propose taking multiple measurements at multiple carrier frequency since multipath interference is highly dependent on the frequency of the signal.

The authors also discuss some of the impacts that the system parameters have on ranging performance. Since the correlation peak (effective SNR) of the PN code is linearly dependent on the code length, the ranging accuracy can be improved by increasing the code length at the cost of computational complexity. However, this only improves the error caused by noise; as a result, the code length should not be increase without bound. Multipath induced errors can be combated by increasing the code chip rate, effectively increasing the bandwidth. That is, by reducing the chip period, the multipath signals with large path delays can be rejected. However, the authors point out that increasing the chip rate has little affect until the chip period is similar to the multipath signal delay. They show through simulation that signals with 1 Mchips/s and 5 Mchips/s have almost the same performance and a further increase to 10 Mchips/s has a very slight improve. Only when the rate is increased to 20 Mchips/s and 50 Mchips/s does the performance error improve. As a result, very large bandwidth will improve results; however, increasing bandwidth of low bandwidth signals will have negligible affect on performance while needlessly increasing energy consumption. The authors state the Cramer-Rao lower bound on the accuracy of RF time-of-flight is:

$$\sigma^2_{TOF} = \frac{1}{8\pi \cdot SNR \cdot \sqrt{\alpha} \cdot BW^2 \cdot N} \tag{3.3}$$

where

$$\begin{aligned} SNR &= \text{average SNR}, \\ \alpha &= \text{\# of copies averaged}, \\ BW &= \text{bandwidth}, \\ N &= \text{\# of chips} \end{aligned}$$

The authors provide experimental results done in different indoor and outdoor settings. The signal parameters chosen for the experiments are 1 Mchips/s code chip rate, 8 chips code length, and 8 code copies averaging. To combat multipath interference, they average results using multiple frequencies. Furthermore, averaging multiple measurements at each frequency is done to reduce noise interference. They claim to observe typically better than one meter error outdoors and better than three meters error indoors.

3.2.5 RIPS

In [17], the authors present a radio interferometry based technique for localization which they call Radio Interferometric Positioning System (RIPS). The technique has its origins in the fields of physics, geodesy, and astronomy. The location of a radio source can be determined by two receivers at know locations by cross correlating the radio signal received resulting in an interference signal that can be analyzed to determine location. However, since radio interferometers are complex and expensive devices, the technique can't be directly use in its traditional sense. Instead, the authors propose to create the interference signal directly by using two transmitters. By using two slightly offset frequencies, the resulting signal has a low frequency envelope that can easily be measured by the typical radio chips on sensor motes. The authors propose making multiple measurements using different carrier frequencies using a minimum of 8 motes in the network to avoid tight synchronization requirements. That is, using only the single interference signal to determine the relative position of the two transmitters and receiver would require additional hardware to meet the synchronization constraints. However, they determine that the phase offset of the signal associated with two receivers would be a function of the relative location of the four motes (two transmitters and two receivers) and the carrier frequency. Furthermore, they determine that to solve for all relative positions a minimum of 8 motes is required.

Using two transmitter motes, A and B, the relative phase offset of the received signal at the two receiver motes, C and D, is determined. As noted before, the two transmitters transmit pure sine waves simultaneously at slightly different frequencies, f_A and f_B, resulting in an interference signal with a low beat frequency of $|f_A - f_B|$. The phase of the interference signal is determined at each receiving mote. The authors proved that the relative phase offset between the signal detected at C and at D is constrained by the equation:

$$\Delta\phi = \phi_C - \phi_D = 2\pi \frac{d_{ABCD}}{\lambda_{carrier}} (mod\ 2\pi) \tag{3.4}$$

where

$$\begin{aligned} d_{ABCD} &= d_{AD} - d_{BD} + d_{BC} - d_{AC}, \\ \lambda_{carrier} &= \frac{2c}{f_A + f_B} \end{aligned}$$

Since the phase offset is modulo the carrier frequency, multiple measurements at different carrier frequencies are necessary to determine the range value d_{ABCD}. Furthermore, the authors state that at most there can be $n(n-3)/2$ linearly independent set of measurable ranges in an n-mote network. Coupled with the fact that in 3D the number of unknowns is $3n - 6$, a minimum of 8 motes are needed in the network for there to be more measurements than unknowns.

The authors implemented there algorithm on Crossbow Mica2 [1] motes because it uses the Chipcon CC1000 [2] radio chip which is very configurable. They required a custom driver for the radio chip so they could transmit unmodulated pure sine wave at distinct frequencies. Additionally, to ensure that the relative phase offset was measured correctly, the author design a scheduling scheme for transmission and reception times to create a common time instant between the participating motes. Furthermore, a fair amount of tuning is required to set the two frequencies of the transmitting motes.

The phase offset equation can be recast and for each carrier frequency used we obtain the set of equations (for their tests the author used 10 different frequencies):

$$d_{ABCD}(i) = \lambda_i n_i + \gamma_i \tag{3.5}$$

where

$$\begin{aligned} \gamma_i &= \lambda_i \frac{\Delta\phi_i}{2\pi}, \\ n_i &= \text{integer} \end{aligned}$$

Because of error introduced by the noisy environment, the equations can't be solved exactly. Instead the range value d_{ABCD} is determined by finding the minimum of the error function:

$$error = \sqrt{\sum_i \left(d_{ABCD} - d_{ABCD}(i)\right)^2} \tag{3.6}$$

Since the ranging values are derived for distance separation among four motes, instead of pair-wise, the estimation of location can't be determined via simple triangulation or some other common method. As a result, the authors decide to develop a method based on genetic algorithms to determine the relative positions of motes based on a set of range values. We refer the reader to [17, Section 6] for a more detailed description of the genetic algorithm. The author noticed that interference patterns could be observed up to twice the communication range. The experiments were conducted using 16 motes, in a 4-by-4 grid, deployed in an 18-by-18 meter open field with no obstructions. They state that the entire localization process, including the required frequency tuning, multiple measurements, etc, takes roughly 80 minutes. However, they believe that after further optimization, and in a smaller scale setup, the process should be able to finish in about 5 minutes. The observed accuracy on average was 3 cm; however, they note that the mote placement was only to an accuracy of 5 cm so they simply claim this indicates the RIPS can achieve high precision.

3.3 Chapter Summary

In this Chapter, we presented localization techniques that where design with Wireless Sensor Networks in mind or are directly applicable to them. We first explored two similar techniques that require acoustic signaling and receiving hardware and found they differ in the principle of centralization. We next shifted our focus to techniques that could be easily implemented on commercial wireless sensor motes. We found that they ranged from providing coarse-grain localization using ordinary signals and simple algorithms to promising more fine-grain localization using specialized signals and complex algorithms.

Chapter 4

Evaluation of Localization Techniques

This Chapter focuses on the set of design criteria that an indoor Wireless Sensor Networks based localization technique should possess in many applications and why they are important. The criteria can be broadly classified as efficiency-based and accuracy-based. We explain how the localization techniques from the previous Chapter evaluate to these criteria and in particular, we show that only one of the techniques examines well to the efficiency-based criteria. However, the published test cases do not cover some of the scenarios we are seek to have performance details for and so we conclude the section by indicating which test scenarios remain unanswered to evaluate the accuracy-based criteria.

4.1 Design Considerations

The techniques examined in the previous Chapter are merely techniques that have the potential to be deployed on wireless sensor motes but may not necessarily be fully appropriate for many indoor applications. Indoor localization with wireless sensor networks in such applications must meet the requirements that follow. The first five criteria are referred to as efficiency-based and the last two are considered accuracy-based. Therefore, any technique that evaluates well to the first five criteria will be referred to as 'efficient' and those that evaluate well to the last two criteria will be considered 'accurate'.

1. **Deployable on Commercial Motes With No Hardware Integration**

   A major benefit to using wireless sensor motes in localization is that they are cheap devices deployed in a large number. As the use of sensor motes becomes more and more prevalent, it is possible to use some of the motes deployed for purposes such as monitoring temperature (we refer to them as non-localization motes) as reference motes in a localization technique. This is an attractive development for localization systems since it allows the use of many reference points without any monetary or otherwise investment from the localization system. However, it is important to remember to not interfere with the primary function of these non-localization motes.

A technique that requires specialized hardware whose sole purpose is for localization adds extra monetary cost to the solution. In additional, it precludes the use of non-localization motes which means the large deployment benefits of wireless sensor motes is lost. Any motes used in the localization system will have to have specialized hardware integrated. This criterion demonstrates cost efficiency and efficient use of motes in a wireless sensor network.

2. **Low Communication Overhead**

   A low communication overhead is desirable to converse battery power and reduce interfere within the wireless sensor network. Wireless sensor motes operate on batteries so energy efficiency needs to be one of the utmost concerns. A technique that uses the mote's radio extensively for transmission would deplete a mote's battery quicker than a technique which requires fewer signal exchanges. In addition, extensive signaling would cause a large communication overhead which would interfere with non-localization bandwidth. This criterion demonstrates energy efficiency and efficient use of the channel in a wireless sensor network.

3. **Low Complexity Algorithm**

   Although an algorithm's complexity can be carried out by the limited capabilities of the motes, it may be at the expensive of energy-efficiency. That is, the microcontroller would be able to complete all the computations, but the lengthy computations would result in the microcontroller drawing a lot of power from the batteries. Additionally, this would be a time consuming process that could preclude the use of the algorithm in real-time localization. This criterion would also imply that technique avoid using too many complex signals, since the generation and reception of them is energy demanding. This criterion demonstrates energy efficiency and efficient use of localization time.

4. **Ease of Initial Setup**

   Given that it is an indoor deployment, there are many factors to consider. Firstly, an indoor environment is transient so the initial deployment setup should not be too cumbersome because you may need to redeploy as things change. Also, wireless sensor networks are meant to be autonomous so a requirement to place motes in carefully decided locations defeats the purpose. This criterion demonstrates deployment time efficiency and efficient use of motes in a wireless sensor network.

5. **Flexible to Deployment Environment**

   The localization system should be flexible to different environments. That is, localization may be required in a room where only a few reference motes are present. Even if it means poorer accuracy, the localization system should be able to work in such conditions. This criterion demonstrates efficiency in adapting to the localization environment.

6. **Potential Accuracy Level**

   Potential accuracy level remains important because a fully 'efficient' technique that doesn't provide a reasonable level of accuracy would be completely useless. However, many applications don't require absolute accuracy. For such applications, a technique that gives needless pinpoint accuracy at the expense of efficiency considerations would be highly undesirable. As a result, we don't immediately discard techniques that provide foot level accuracy as uninteresting. Here we looking for accuracy in an environment without obstructions. The accuracy coverage are must also be considered. That is, the technique shouldn't just provide high accuracy in certain regions of the localization area while providing substandard accuracy in the remaining locations. This criterion demonstrates accuracy in low multipath and shadowing environment and coverage.

7. **Robustness to Multipath and Shadowing**

   Finally, above all, since we are working in an indoor environment that can have many obstructions, the algorithm must be robust enough to not be significantly affected by noise (multipath and shadowing effects). It is expected that the accuracy suffers more as there is an increase in multipath interference and shadowing effects; however, the level of drop off should not be such that accuracy is complete lost, otherwise, it means that the technique is not appropriate for indoor use. This criterion demonstrates accuracy in dealing with harsh environments.

Many times localizations techniques are judge solely on the accuracy-based criteria. Although accuracy is a very important criteria, we must remember that for wireless sensor networks the efficiency-based criteria are equally important. Tables 4.1 and 4.2 summarize how each technique evaluated to the efficiency-based criteria.

	Commercial motes	Communication Overhead	Algorithm Complexity	Initial Setup	Flexibility to Environment
Cricket	No	Moderate	Low	Moderate	High
Active Bat	No	Moderate	Low	High	High
Centroid	Yes	Low	Low	Low	High
APIT	Yes	Moderate	Low-moderate	Low	Low
Ecolocation	Yes	Low	Low	Low	High
RF TOF	No	High	Moderate	Low	High
RIPS	Yes	High	High	Low-moderate	Moderate

Table 4.1: Evaluation of localization techniques to efficiency-based criteria listed in Section 4.1

With the efficiency-based criteria in mind, the most suitable techniques out of those discussed in the previous Chapter are the Centroid method and Ecolocation/Sequenced-Based method. The Cricket and Active Bat techniques require specialized hardware in the form of ultrasonic transducers and detectors and a dedicated channel for ultrasound communication; as a result, they are not so attractive solutions and were eliminated from our list of possible choices. While A.P.I.T. requires relatively low communication overhead for exchanging RSSI measures between unknown location receivers, its requirement of many receivers to accurately determine the in-triangle decision is highly undesirable. It means that the technique is not too flexible to systems where a single mote is trying to localize itself. With this in consideration, A.P.I.T. doesn't seem too appropriate. In addition to specifically being designed in mind for second generation motes, the RF time-of-flight technique is also precluded because it requires a

	Commercial motes	Communication Overhead	Algorithm Complexity	Initial Setup	Flexibility to Environment
Cricket	Ultrasound transducers and speakers	Require dedicated channel for Ultrasound	Triangulation based on RF-Ultrasound time-of-flight data	Beacon placement in strategic locations	Can Localize with few reference motes
Active Bat	Ultrasound transducers and speakers	Require dedicated channel for Ultrasound	Triangulation based on RF-Ultrasound time-of-flight data	Require wired network of reference motes	Can Localize with few reference motes
Centroid	N/A	can use any communication packet for RSSI data	Centroid of reference locations in "proximity" range	Just need reference motes	Can Localize with few reference motes
APIT	N/A	can use any communication packet for RSSI data, but requires much neighbor node communication	Classifying all triangles as "in-or-out-triangles" based on neighbors	Just need reference motes	Require neighbor nodes for APIT test
Ecolocation	N/A	can use any communication packet for RSSI data	Comparing RSSI constraint matrix with grid point constraint matrices	Just need reference motes	Can Localize with few reference motes
RF TOF	2nd Generation motes	Multiple special signals	Correlating PN signals for time-of-flight data	Just need reference motes	Can Localize with few reference motes
RIPS	None	Multiple special signals	Determining phase offset for multiple interference signals and solving a system of equations	Need motes with special signaling	Require minimum eight nodes

Table 4.2: Explanation of evaluations from Table 4.1

large communication overhead of multiple special signals. Finally, not only does RIPS require specialized signaling using different frequency sinusoids, but it requires inter-node communication to resolve the ranging. This process is both time-consuming and energy-consuming. The technique claims to provide 3 cm, but it comes at too great an expense.

With the use of wireless sensor motes indoors in a relatively small room, all transmitters would be in radio range of the unknown location receiver and therefore, would make the connectivity calculation of the Centriod method meaningless. Additionally, the nature of the simple technique is to provide only coarse-grain proximity location and our desire to achieve a reasonable level of accuracy steers us away from the use of this method. As a result, Ecolocation/Sequence-Based seems to be the only techniques able to evaluate successfully to the first five criteria while claiming to provide an acceptable level of accuracy. However, as we will discuss in the following section, the published results are insufficient to evaluate criteria 6 and 7. The technique would be unsatisfactory if it doesn't have the potential to achieve a reasonable level of accuracy or if it is able to perform well in the presence of obstructions.

Ecolocation and Sequence-Based Localization are quite similar techniques. However, Sequence-Based Localization requires the generation and storage of a large sequence table while Ecolocation doesn't. As a result, from this point on we will only consider Ecolocation.

4.2 Ecolocation Performance Evaluation

We have seen that Ecolocation provides the 'efficient' technique that we were trying to find. However, for it to be a suitable solution, it must be 'accurate'. To determine the potential level of accuracy a technique can obtain, the experimental test must be performed in a low obstruction scenario. In [29], the authors have not addressed such a scenario. Furthermore, by repeating the same procedure used in those tests with the addition of obstructions, the technique's robustness to multipath and shadowing can be inferred with a direct comparison to the obstruction-free scenario's results. The performance of Ecolocation in such a case also remains unknown.

The authors neglect to test the performance in a low error scenario. That is, the authors in [29] claim via simulations that a low error environment is with low standard deviation, high reference mote density, and grid-style deployment. However, neither of the real world experiments looked at in [29] try to show such an environment. The outdoor example provides a low standard deviation environment, but reference motes are deployed in a low density random fashion. Additionally, the nature of the outdoor implementation meant that each unknown node observed a different reference node deployment. The disjoint nature of the indoor and outdoor testing doesn't allow for a direct comparison to be made to infer robustness to multipath and shadowing. By testing in just few locations (11 points outdoors, 5 indoors), the localization coverage area also remains an open issue.

As a result, we will test the performance of the Ecolocation algorithm ourselves. Chapter 6 presents the result of this performance evaluation, and Chapter 5 outlines the systematic approach we took in determining accuracy to the criteria's standards.

4.3 Chapter Summary

In this Chapter, we first explained what aspects an indoor wireless sensor networks technique should feature and our reasoning for them. After eliminating those techniques that failed to meet efficiency-based criteria, we identify that the Ecolocation algorithm is an efficient solution. However, we determine that the original paper didn't provide some of the performance results we were looking for in evaluating accuracy.

Chapter 5

Implementation and Experimental Setup

In this section, we explain our procedure to evaluate the Ecolocation algorithm in terms of potential accuracy level and robustness to multipath and shadowing. First, we investigate which of the commercially available mote platforms is most suitable for use with the Ecolocation technique. We then give a more comprehensive description of the algorithm and the details of how we implemented it. The Chapter finishes off by explaining the experimental testbeds for a low-error scenario to determine potential accuracy level and a high-error scenario to determine robustness.

5.1 Choice of Mote Platform

Now that we have determined that we must conduct experiments, we must decide which mote platform to use. During the decision process we must keep in mind which platform is most suitable for Ecolocation. In this section, we will present an overview of the commercially available wireless sensor motes and how they different in respect to their effectiveness in use with Ecolocation. Additionally, we perform experiments with each of the motes to determine which it the best choice for us.

5.1.1 Motes of Consideration

With the commercially available wireless sensor motes in the market on mind, our decision was to choose among the Crossbow Mica2 [1] and the MicaZ [3], and the Moteiv Tmote Sky [4] sensor motes (see Table 5.1). As a result, we compared the sensors motes based on their differences in relation to the chosen localization technique. That is, for our chosen technique of Ecolocation, the determination (and as a consequence the accuracy) of the RSSI values between exchanged message is of the utmost importance.

	Crossbow Mica2	Crossbow MicaZ	Moteiv Tmote Sky
Microcontroller	Atmega128L (8MHz, 4K RAM, 128K Flash)	Atmega128L (8MHz , 4K RAM, 128K Flash)	TI MSP430 (8MHz, 10K RAM, 48KFlash)
Transceiver	Chipcon CC1000 (38.4 Kbaud, 868/916 MHz)	Chipcon CC2420 (250 kbps, 2.4 GHz)	Chipcon CC2420 (250 kbps, 2.4 GHz)
RSSI	16-bit analog	8-bit digital	8-bit digital
Outdoor Range	75 to 100 m	75 to 100 m	125 m
Indoor Range	20 to 30 m	20 to 30 m	50 m
Integrated Antenna	No	No	Yes
Integrated Sensors	No	No	Yes
I/O Connector	51-pin expansion connector	51-pin expansion connector	USB

Table 5.1: Data specifications of three commercially available Wireless Sensor Networks motes

Hence, we look at how each mote determines the RSSI value of the received packet. We discovered that the Crossbow Mica2's provide the best precision for RSSI values and so in theory the technique should provide the most accurate results using these motes. Both the MicaZ's and Tmote Sky's provided their digital RSSI values with eight bit signed precision while the Mica2's provide an analog RSSI value with sixteen bits of precision. It should be noted that based on this determination, the authors in [29] also chose to test with Mica2.

5.1.2 Experimental Evaluation

Although Mica2 motes provide the most precise RSSI value, the Ecolocation algorithm doesn't make use of the absolute RSSI value and so it is not so obvious that Mica2 motes would be the best to use. The Ecolocation algorithm implicitly assumes that there is a correlation between RSSI and distance between transmitter and receiver. Therefore, the mote that provides the best correlation relationship between RSSI and distance is the best mote platform for the algorithm. As a result, we test the three mote platforms of interest. We place a receiver mote at the center of a fourteen feet-by-twelve feet rectangle and take measurements from a transmitter which is moved in one foot intervals in the x- and y-directions starting at the origin. We average ten RSSI measurements from each location point. To quantify the amount of distance to RSSI value correlation present in the motes, we calculate a correlation percentage. The correlation percentage is simply the percent of times two location points exhibit a positive correlation (the further location point has a lower RSSI value) for all 18,915 (195 choose 2) combinations.

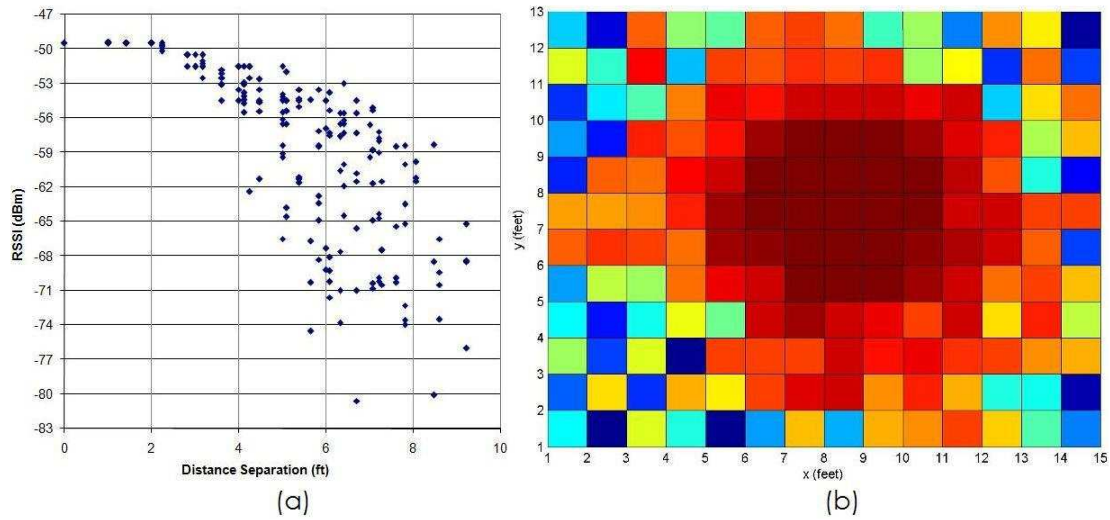

Figure 5.1: (a) Scatter-plot and (b) Surface-plot of RSSI data for Mica2 motes

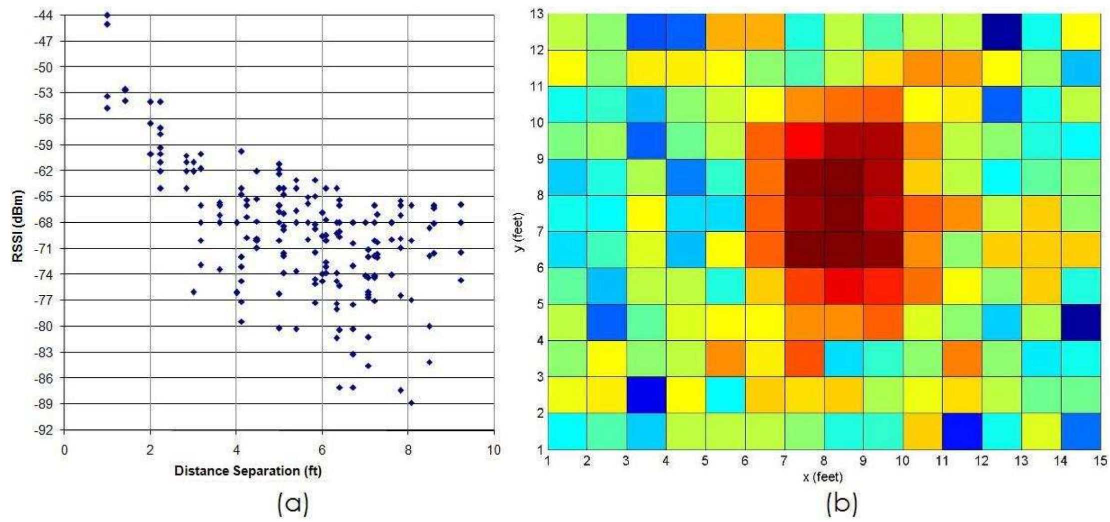

Figure 5.2: (a) Scatter-plot and (b) Surface-plot of RSSI data for MicaZ motes

Figures 5.1a, 5.2a, and 5.3a are scatter-plots for RSSI versus distance for Mica2, MicaZ, and Tmote Sky, respectively. Figures 5.1b, 5.2b, and 5.3b are surface-plots of the same data. From the figures, it is quite obvious that Mica2 outperforms MicaZ and Tmote Sky. Figure 5.1a shows Mica2 has a clear inverse relationship between

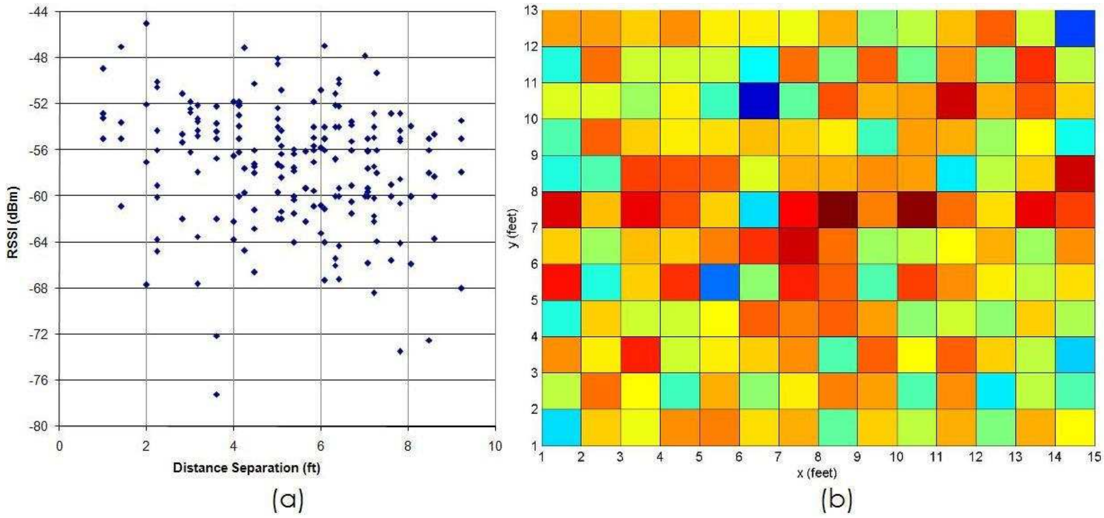

Figure 5.3: (a) Scatter-plot and (b) Surface-plot of RSSI data for Tmote Sky motes

RSSI and distance. Furthermore, the surface-plot in Figure 5.1b shows how the RSSI drops off radial from the position of the receiver. It can be seen that MicaZ seems to perform somewhat adequately. Figure 5.2a shows that there is a downward trend for RSSI-distance relationship with MicaZ; however, we can see from the surface-plot in Figure 5.2b that the radial drop off is not as good as with the Mica2. The Tmote have been found to be not adequate for the algorithm. From Figure 5.3a, we see that the Tmote RSSI and distance are not very uncorrelated. Figure 5.3b shows that there is no radial trend. The correlation percentages are 81.1%, 55.2%, and 67.9% for Mica2, Micaz, and Tmote Sky, respectively. As a result, we have chosen to use the Crossbow Mica2 for the implementation of the Ecolocation algorithm.

5.2 Ecolocation: In-depth Look

Before we proceed with the description of the experimental testbed, we will provide a more thorough explanation of the Ecolocation technique. First the localization area is defined to some rectangular region defined by an x-maximum and y-maximum. At fixed known location in the localization space, a certain number, n, of reference nodes are placed (see Figure 5.4). After the unknown location node collects all RSSI values for the reference node, the RSSI constraint matrix is formed according to the element

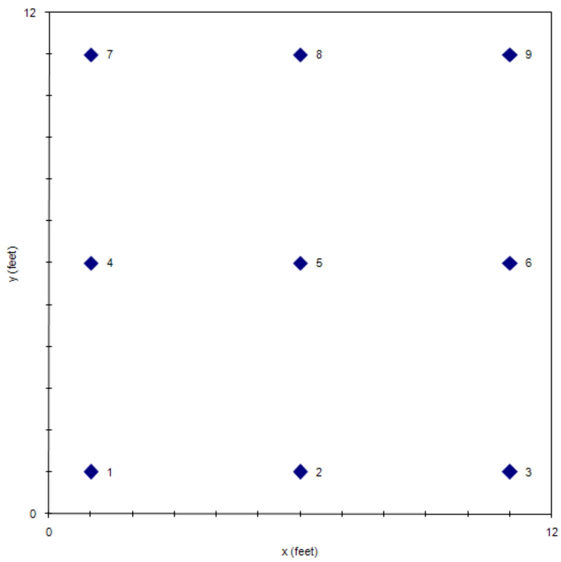

Figure 5.4: Shows a localization region defined 12 feet-by-12 feet with 9 reference nodes.

definition:

$$R_{n\times n}(i,j) = sign(RSSI_i - RSSI_j) \tag{5.1}$$

Figure 5.5 shows an example RSSI collection from nine reference nodes and the corresponding RSSI constraint matrix.

RSSI Ranking

Ref Id	RSSI
1	-50
2	-47
3	-54
4	-49
5	-45
6	-53
7	-57
8	-55
9	-60

RSSI Constraint Matrix

	1	2	3	4	5	6	7	8	9
1	0	-1	1	-1	-1	1	1	1	1
2	1	0	1	1	-1	1	1	1	1
3	-1	-1	0	-1	-1	-1	1	1	1
4	1	-1	1	0	-1	1	1	1	1
5	1	1	1	1	0	1	1	1	1
6	-1	-1	1	-1	-1	0	1	1	1
7	-1	-1	-1	-1	-1	-1	0	-1	1
8	-1	-1	-1	-1	-1	-1	1	0	1
9	-1	-1	-1	-1	-1	-1	-1	-1	0

Figure 5.5: A sample RSSI collection for all nine reference nodes and the corresponding RSSI constraint matrix.

Now a scanning resolution is chosen to determine the grid points at which a constraint match test should be taken. The scanning resolution should be small enough to insure enough grid points are used in the algorithm, but it should be kept in mind that the finer the scanning resolution, the larger computational cost. In Figure 5.6, we see an example with a scanning resolution of three inches. The grid points are represented by every intersection of the crossing lines.

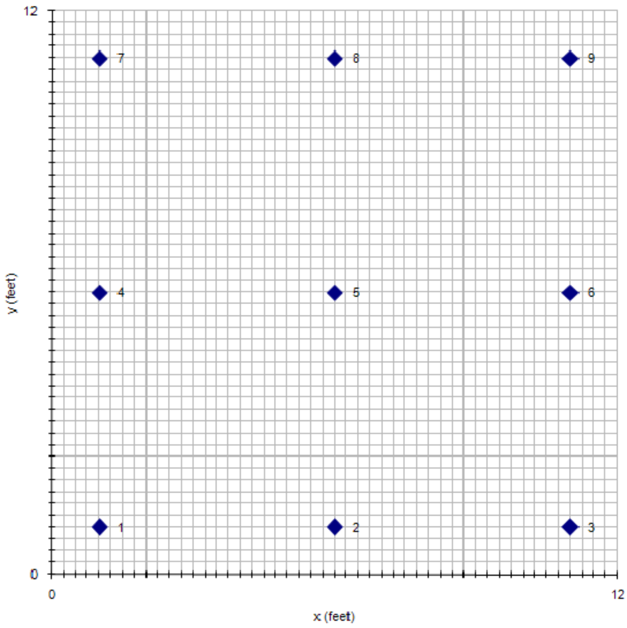

Figure 5.6: The localization region with scanning resolution of 3 inches. Every grid point is represented by the intersection of the crossing lines.

For each grid point (m, n), the distance to each reference node k is determined as:

$$d_k^{(m,n)} = \sqrt{(x_k - m_x)^2 + (y_k - n_y)^2} \tag{5.2}$$

Figure 5.7 shows an example for grid point (10, 35) which represents location (2.5, 8.75) feet in real units.

Now the distance calculations are converted into a grid point constraint matrices for each grid point according to the element definition:

$$C_{n\times n}^{(m,n)}(i,j) = -sign\left(d_i^{(m,n)} - d_j^{(m,n)}\right) \tag{5.3}$$

Figure 5.8 shows the grid point constraint matrix for the grid point (10, 35).

Now the RSSI constraint matrix is compared with each of the grid point constraint matrix to determine which grid point is the most likely location of the unknown node.

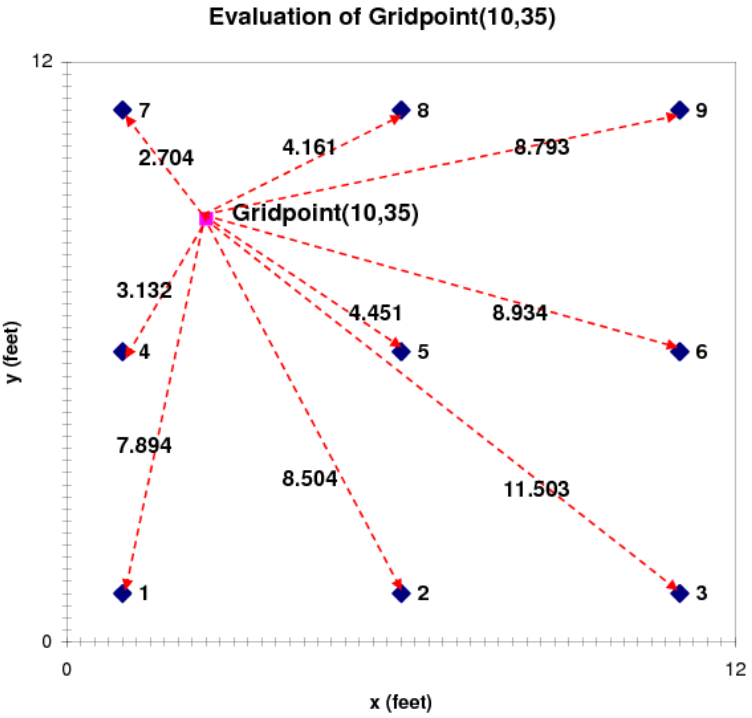

Figure 5.7: An example grip point (10, 35) which with 3 inch resolution represents location (2.5, 8.75) feet and the distances to all the reference points.

Distance Ranking

Ref Id	Distance
1	7.894
2	8.504
3	11.503
4	3.132
5	4.451
6	8.934
7	2.704
8	4.161
9	8.793

Grid Point Constraint Matrix

	1	2	3	4	5	6	7	8	9
1	0	1	1	-1	-1	1	-1	-1	1
2	-1	0	1	-1	-1	1	-1	-1	1
3	-1	-1	0	-1	-1	-1	-1	-1	-1
4	1	1	1	0	1	1	-1	1	1
5	1	1	1	-1	0	1	-1	-1	1
6	-1	-1	1	-1	-1	0	-1	-1	-1
7	1	1	1	1	1	1	0	1	1
8	1	1	1	-1	1	1	-1	0	1
9	-1	-1	1	-1	-1	1	-1	-1	0

Figure 5.8: The distance values from grid point (10, 35) to all the reference nodes and the corresponding grid point constraint matrix for that grid point.

For each element (i, j) in the RSSI constraint matrix and grid point constraint matrix that match, a counter is increment; otherwise, the counter is decremented. The counter value for each grid point is calculated based on:

$$counter^{(m,n)} += \begin{cases} 1 & \text{if } C^{(m,n)}_{n\times n}(i,j) = R_{n\times n}(i,j) \\ -1 & \text{otherwise} \end{cases} \tag{5.4}$$

For example, for grid point (10, 35) the counter value is 13 because there are 47 constraints matched out of a possible 81 (see Figure 5.9).

RSSI Constraint Matrix

	1	2	3	4	5	6	7	8	9
1	0	-1	1	-1	-1	1	1	1	1
2	1	0	1	1	-1	1	1	1	1
3	-1	-1	0	-1	-1	-1	1	1	1
4	1	-1	1	0	-1	1	1	1	1
5	1	1	1	1	0	1	1	1	1
6	-1	-1	1	-1	-1	0	1	1	1
7	-1	-1	-1	-1	-1	-1	0	-1	1
8	-1	-1	-1	-1	-1	-1	1	0	1
9	-1	-1	-1	-1	-1	-1	-1	-1	0

Grid Point Constraint Matrix

	1	2	3	4	5	6	7	8	9
1	0	1	1	-1	-1	1	-1	-1	1
2	-1	0	1	-1	-1	1	-1	-1	1
3	-1	-1	0	-1	-1	-1	-1	-1	-1
4	1	1	1	0	1	1	-1	1	1
5	1	1	1	-1	0	1	-1	-1	1
6	-1	-1	1	-1	-1	0	-1	-1	-1
7	1	1	1	1	1	1	0	1	1
8	1	1	1	-1	1	1	-1	0	1
9	-1	-1	1	-1	-1	1	-1	-1	0

Figure 5.9: Comparison of the RSSI constraint matrix with the constraint matrix of grid point (10, 35). The 47 constraint matches out of a total 81 constraints are outlined.

The estimated location is determined as the centroid of all the grid points that have the maximum counter value. In this example, the maximum counter value is 81 and 13 grid points correspond to this value (see Figure 5.10). The centriod of these values is (5.096, 3.904) feet and that is the estimated location determine by the Ecolocation algorithm for the receiver who obtained the RSSI values in Figure 5.5.

5.3 Actual Implementation

Given the details about the Ecolocation algorithm, we must decide how we want to implement the algorithm in a wireless sensor network. The implementation is done for testing purpose but in a way that it can easily be converted for realistic deployment.

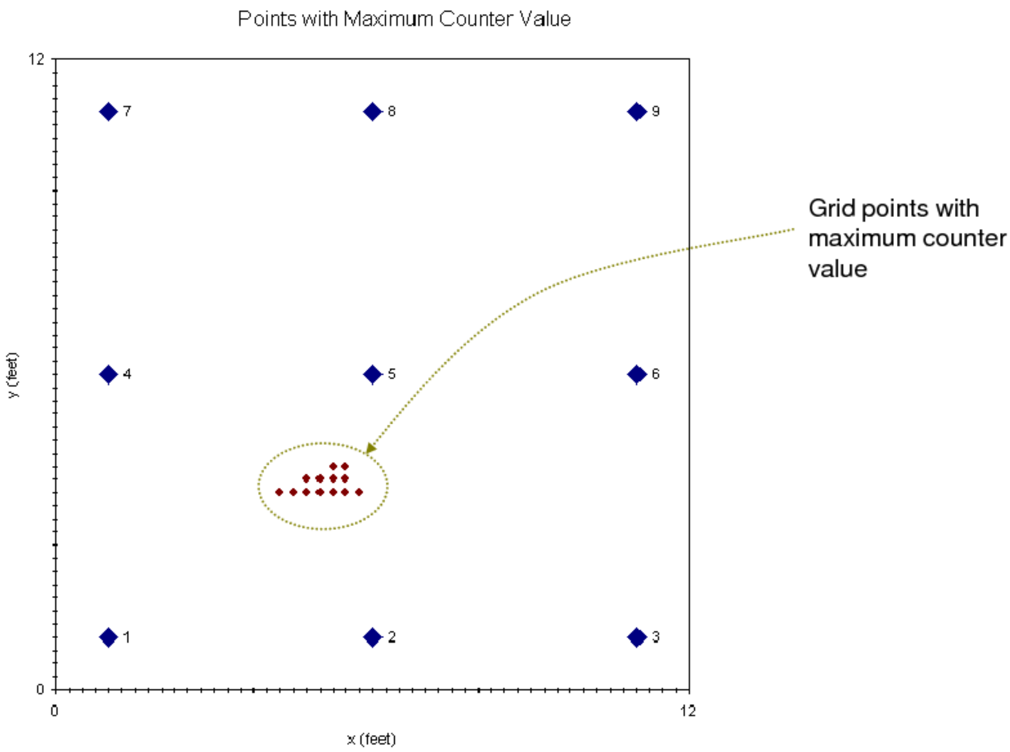

Figure 5.10: The 13 grid points that correspond to the maximum counter of 81 for the sample RSSI values in Figure 4.2.

5.3.1 Coding

We begin the description of the experimental implementation by first explaining the code that runs on the wireless sensor motes and a PC. In our implementation, there are three designations of wireless sensor motes: basestation motes, reference motes, and mobile motes. The basestation motes are attached to the PC and facilitate the acquisition of the data packets. The reference motes are wireless sensor motes in the network that have known locations while the mobile motes are the motes that require the localization service to determine their location. The PC is used as a location engine that aggregates the data from the sensor motes and analyzes the acquired data to determine the location. In our implementation, the mobile motes will transmit periodic beacons to the reference motes. The reference motes will simply determine the RSSI value of the beacon and forward it to the basestation. By 'passively' using the known location motes, we avoid putting a localization burden on reference motes. We feel the known location motes have their own purpose, such as monitoring temperature of a room or vibration of a conveyor belt, and by using them in this manner, we can use them in our localization system without causing problems in their primary function. In addition, both the reference motes and mobile motes are uniquely identified by an ID number. The wireless sensor motes are programmed in NesC using TinyOS 1.1.15. On the PC the application and Ecolocation algorithm are programmed in C# with .NET Framework using Microsoft Visual Studio 2005.

We acknowledge that a reference mote would suffer a substantial burden if there are many beacons being received (either from a single mobile mote or multiple ones) and it has to forward everyone individually as in our implementation. However, this is only for testing purposes; in a realistic deployment, the RSSI information could be piggybacked on communication packets generated by the reference mote.

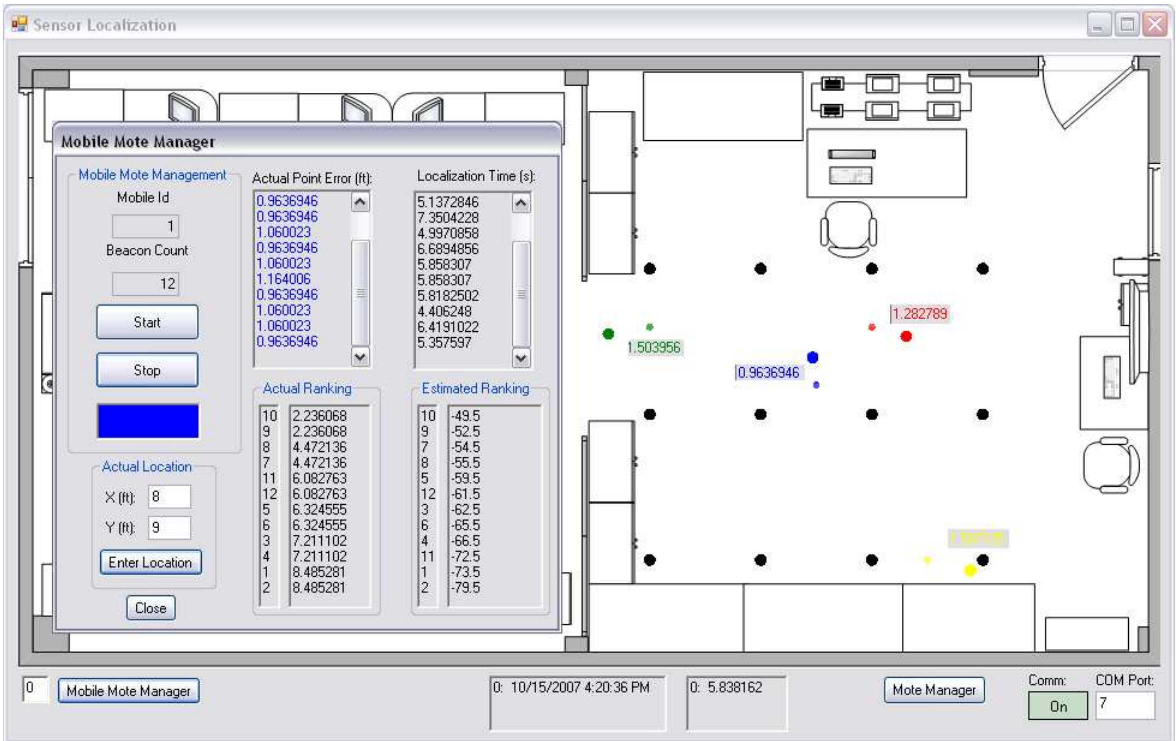

Figure 5.11: Screenshot of Sensor Localization Application. The black dots are the reference motes, the small colored dots are the actual location, and the big colored dots are the estimate locations.

On the PC, we designed an interactive application called Sensor Localization (see Figure 5.11), which will start and stop the localization process and display the estimated locations of the mobile motes along with the reference motes on a floorplan of the room. To start the localization process, the user enters the MoteId of the mobile mote he wishes to locate and also the quantity of reference motes to use (for our simple implementation we simple use the reference motes numbered from 1 to the quantity specified).

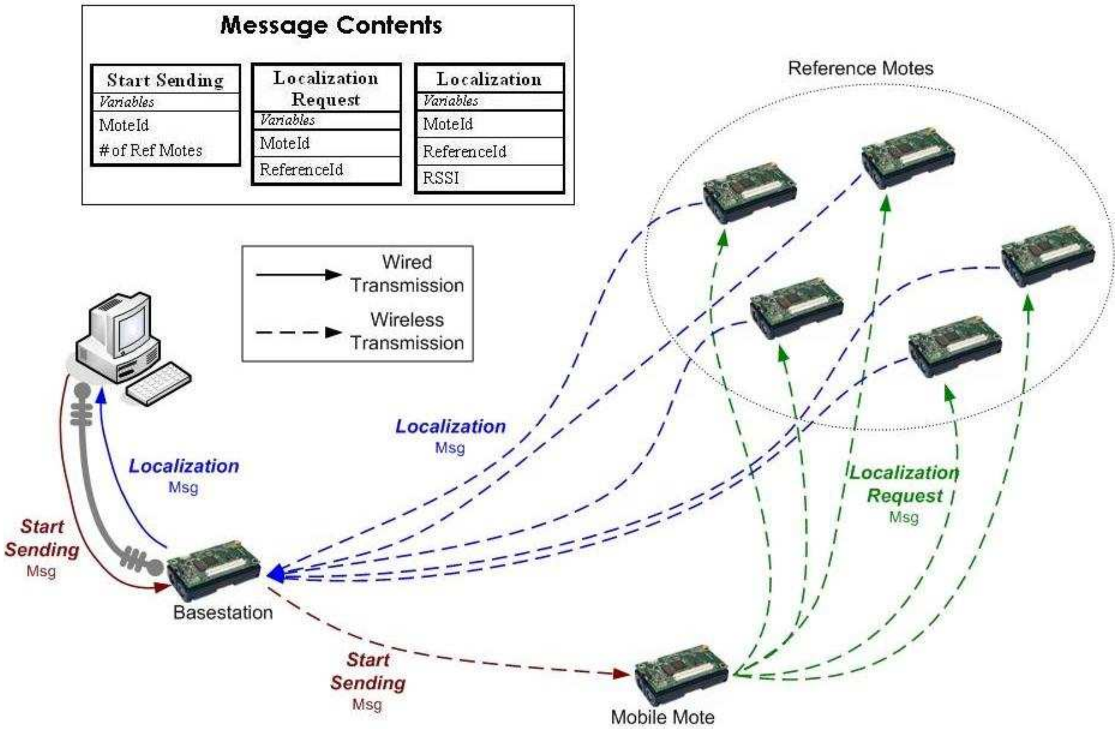

Figure 5.12: The communication process for determining the unknown location of the Mobile Mote based on the know locations of the reference motes. A basestation is attached to a PC via wired connection (typically USB or serial) and is responsible for forward message from the PC to the WSN and vice-versa.

The Sensor Localization application will give these values to the basestation and the basestation will transmit a special "Start Sending" message. When the mobile mote with the corresponding MoteId receives this message, it will enter the "Localization State" and begin periodically transmitting beacons in the form of "Localization Request" message with its MoteId embedded in the message. To minimize collision and insure fast localization, we time-division multiplex the transmission of the beacons. That is, we send the "Localization Request" messages with a ReferenceId and as a result, only the reference node with the corresponding ReferenceId will acknowledge the packet. The mobile mote starts with ReferenceId equal one and cycles through to the quantity number extracted from the "Start Sending" message in intervals of half a second and an then starts back at one. When the corresponding reference mote receives the beacon, it extracts the MoteId from the packet and determines the RSSI and transmits a "Localization" message with these two values along with its ReferenceId to the basestation. The basestation forwards "Localization" messages to the application. The application extracts the MoteId, ReferenceId, and RSSI and creates a list for each MoteId containing a mapping of ReferenceId and RSSI. After a complete list is formed (contained all ReferenceIds based on quantity specified), the list is passed to the Ecolocation algorithm. We refer to this complete list as a RSSI sequence. Figure 5.12 illustrates this sequence of events.

5.3.2 Experimental Implementation Open Area

For our first experiment, we wanted to test indoors but with the minimum obstructions possible (least interference and multipath affects). Therefore, we decided to use an open area on the floor of the RFID Lab at Siemens Corporate Research in Princeton, NJ (see Figure 5.13). This environment gets us a low-error scenario to determine a baseline for potential accuracy level. That is, if accuracy is not reasonable in this environment then the algorithm is not worth considering.

Although there are various tables and cabinets in the room, we use the area in the middle of the room so that between all motes there are no obstacles. The localization area is 14 feet-by-12 feet and we deploy 12 reference nodes in a grid of 4-by-3 with 4 feet separation in the x-direction and 5 feet in the y-direction (see Figure 5.14). The reference motes are deployed in a grid-style because the authors determined in [29] that provide better results over a random deployment. Starting at the origin of the localization space, location estimation measurements are taken in one foot intervals in both the x- and y-directions for a total 195 location points. We systematically test at many points over the entire localization area to get a more conclusive accuracy level and not just at a random 5 or 11 points as in [29]. Additionally, by covering the whole area, we can observe if there are any deadspots in the localization area. At each point, we collected several independent complete RSSI sequences and calculate the location

Figure 5.13: Photograph of deployment area showing that all motes are in line-of-sight with no obstructions in between.

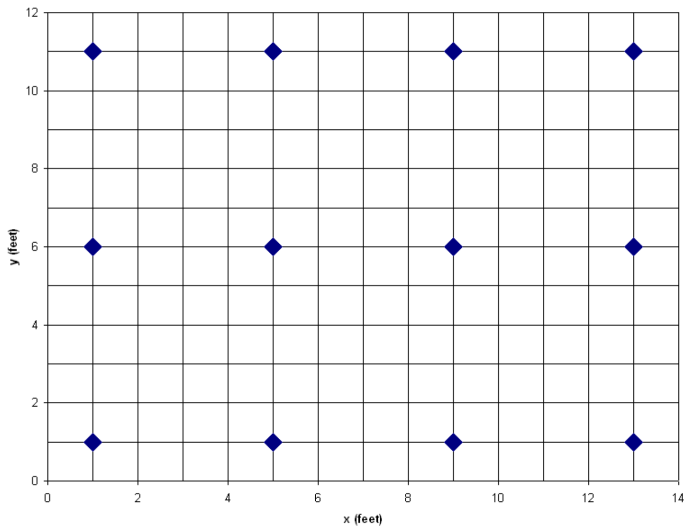

Figure 5.14: Configuration of reference mote deployment (blue dots are reference motes)

estimation and subsequent location error. We use the mean of the location errors at the point as the error for the point. It should be noted that we observed relatively small standard deviation in the error at a particular point (typically less than three percent) and therefore, average only about ten samples at each point.

5.3.3 Experimental Implementation Noisy Environment

For our second set of experimental results, we wanted to test performance in a more realistic, harsher environment (suffering from reasonable interference and multipath). As a result, we add barriers into the first experiment's area (from Section 5.3.2). The barriers are constructed from quarter inch thick masonite hardboard cut into 2 feet wide and 8 inches high pieces with one side of the barriers covered with aluminum foil. Figure 5.15 shows the configuration of the barriers in the 14 feet by 12 feet area and Figure 5.16 is a photo of the area. We again systematically test at the 195 location points from the first experiment. By conducting the experiments in the same place and at the same locations, we are able to actually compare the algorithms robustness to multipath and shadowing since the two tests only differ in the level of noise present.

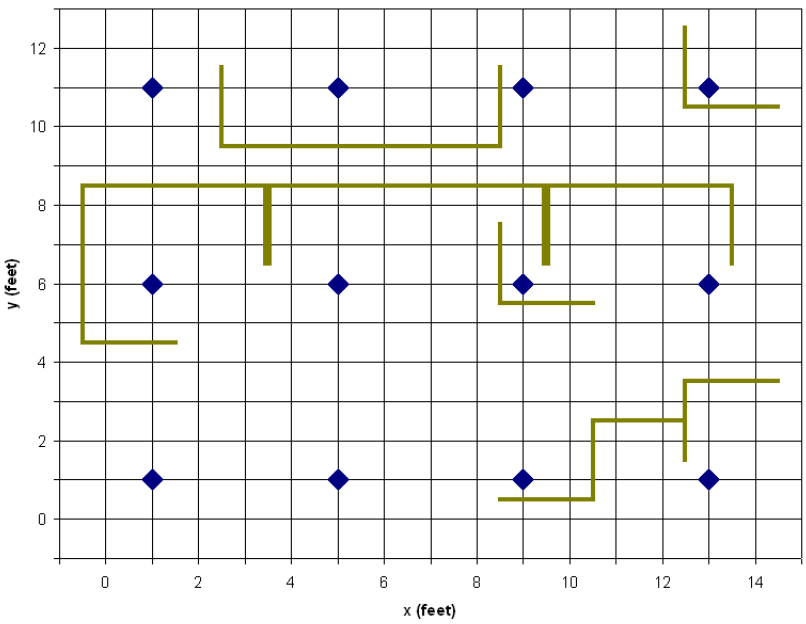

Figure 5.15: Configuration of barriers within mote deployment area (Gold lines are walls and blue dots are reference motes)

Figure 5.16: Photograph of deployment area showing that barriers create obstructions between motes preventing them from being in line-of-sight with each other.

5.4 Chapter Summary

First, we focused on choosing which commercially available wireless sensor motes to deploy the Ecolocation technique on. We found that Crossbow Mica2 motes provided the best RSSI to distance correlation. We then presented an in-depth look at the Ecolocation algorithm which we followed up by explaining how we implemented the algorithm with the coding on the motes and PC. Finally, we discussed the testbed we used to determine potential accuracy and robustness to multipath and shadowing.

Chapter 6

Best-case Indoor Performance of Ecolocation and Robustness to Noise

In this Chapter we will test the performance accuracy of the Ecolocation algorithm. As we mentioned in Section 4.2, the authors of [29] failed to test Ecolocation in what they claim to be a low-error scenario. With our initial experiment we aim to create a testbed that emulates such an environment. Additionally, we will observe how the results are affected when a more higher-error scenario is tested to determine the algorithms robustness to multipath and shadowing.

6.1 The Ecolocation Results

6.1.1 Results in Open Environment

We analysis the performance results of Ecolocation and determine its validity and usefulness in terms of accuracy in a low-error scenario. The reason for this is to get baseline best accuracy possible results which will allow us to determine whether this algorithm is even worth considering.

The Cumulative Distribution Function (CDF) for the 195 locations point with respect to error in feet is shown in Figure 6.1. The mean error observed over all the location points is 1.559 feet with nearly 60% of the points falling below this mean. A maximum error of 6.917 feet is observed. Knowing that the minimum separation between two adjacent reference motes is 4 feet, it is important to note that 75% of the location points have error below 2 feet.

Uniform error distribution over the entire localization space is an important feature for a localization technique to have. This shows that the technique has good coverage

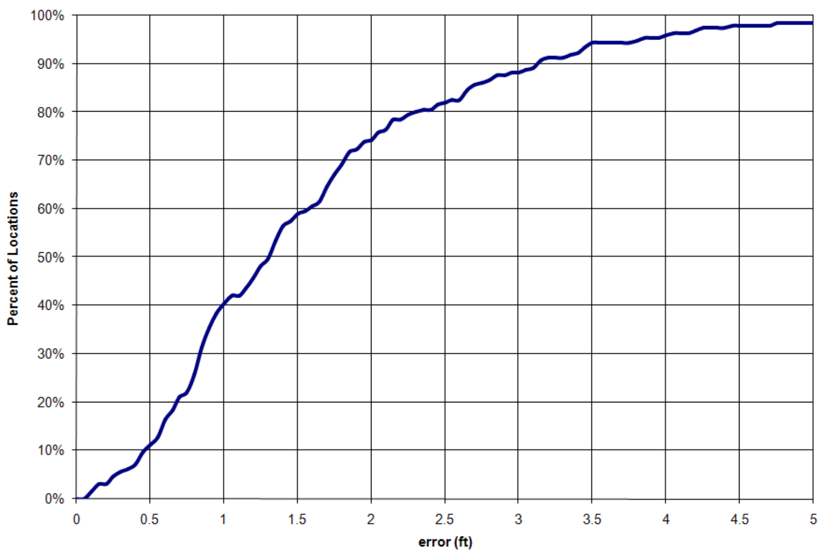

Figure 6.1: The CDF of estimation error in feet for Ecolocation. The y-axis is the percent of the 195 points that have error equal to or below the x value.

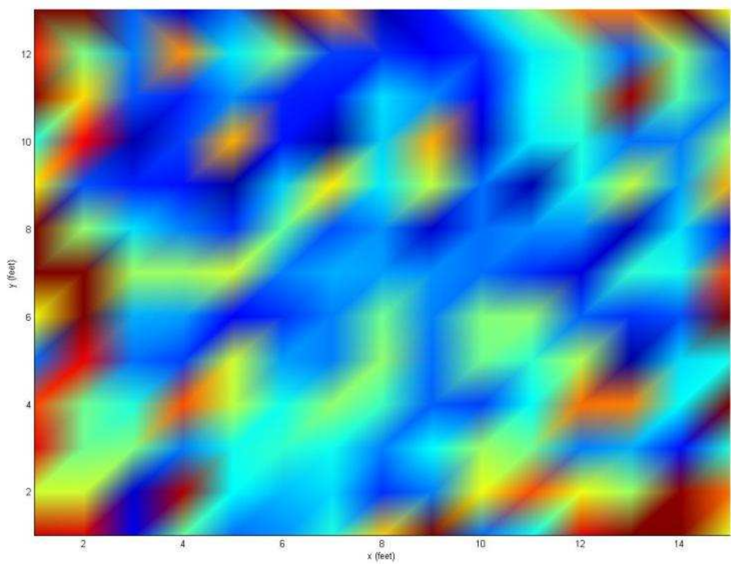

Figure 6.2: The surface-plot of the error observed at each location point. Blue low error, red high error (see Figure 6.3)

Figure 6.3: Color Bar representing error in feet for surface plots

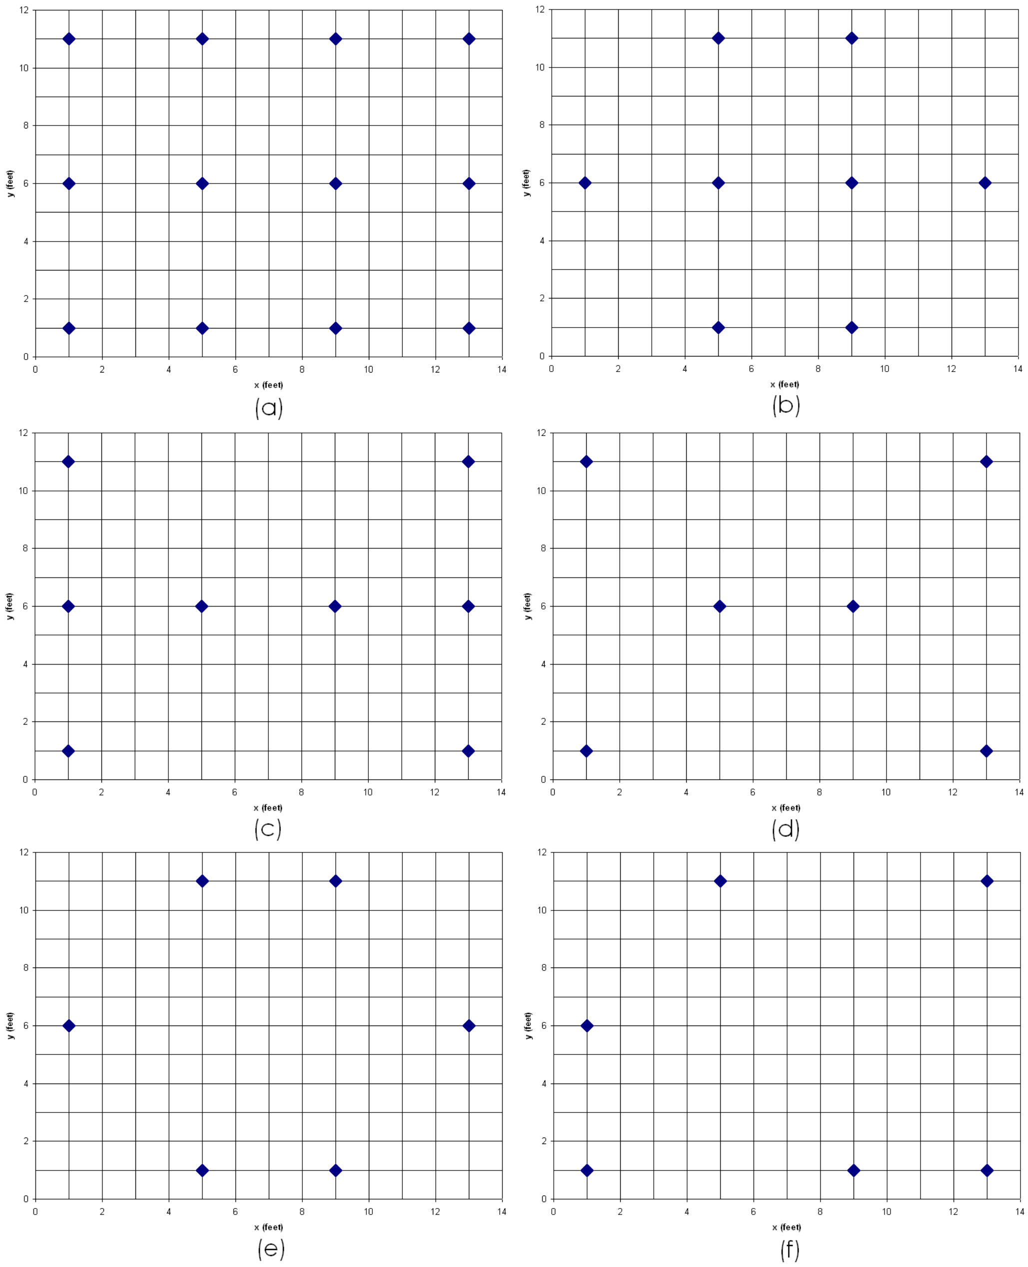

Figure 6.4: Configuration of reference mote deployment (blue dots are reference motes): (a) Scenario 1 (12 reference motes), (b) Scenario 2 (8 reference motes), (c) Scenario 3 (8 reference motes), (d) Scenario 4 (6 reference motes), (e) Scenario 5 (6 reference motes), and (f) Scenario 6 (6 reference motes).

range. Figure 6.2 shows the surface plot of the location errors. We can see that the algorithm doesn't suffer from any real deadspots and that high estimation accuracy is not isolated to one spot. This results seem to show that Ecolocation can provide the reasonable level of accuracy we are looking for.

Sensitivity to Reference Node Deployment

We stored RSSI sequences collect at each location points so we could examine perform in several deployment scenarios. That is, by using only a subset of the RSSI values in the sequence, we are able to create sequences based on fewer reference motes and in different arrangements. The five additional Scenarios (along with Scenario 1 being the full 12 reference motes) are depicted in Figure 6.4. Scenarios 2 and 3 have 8 reference motes while Scenarios 4-6 have 6 nodes each. The CDF of all the Scenarios is shown in Figure 6.5. As logically expected, with fewer reference motes, the localization

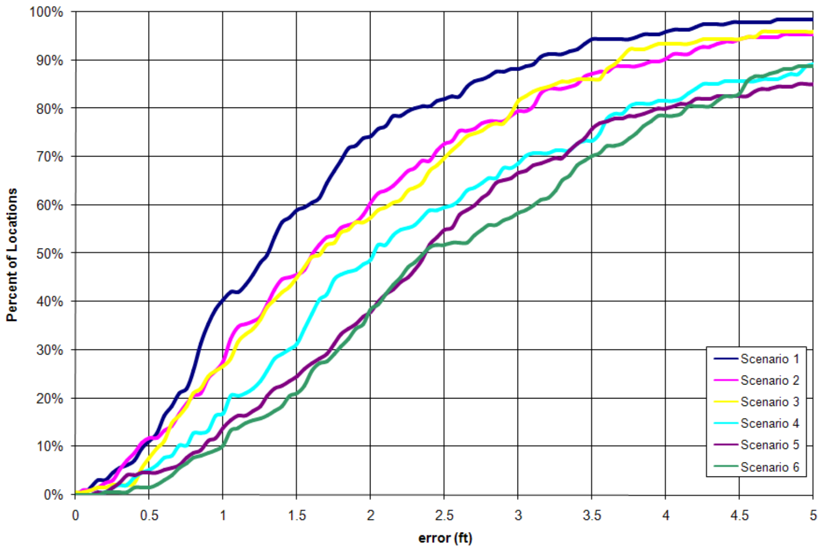

Figure 6.5: The CDF of Ecolocation for all six Scenarios.

performance in terms of error is worse. The mean error in Scenarios 2 and 3 are 1.987 and 2.013 feet, respectively, resulting in an average of 2.000 feet error in the 8 reference mote scenarios. The mean error in Scenario 4, 5, and 6 are 2.522, 2.765, and 2.898 feet, respectively, resulting in an average of 2.728 feet error in the 6 reference mote scenarios. One possible reason for Scenario 4 outperforming the other two 6 node scenarios is by the inclusion of the two interior reference motes. Scenarios 5 and 6 perform reasonably similar with both make use of only exterior reference motes.

Although Scenario 1 gave favorable results, we have observed a considerable fall off in accuracy with the use of fewer reference motes. Only about 40% of the time is error below 2 feet in the 6 reference mote scenarios compared to the 75% observed with 12 reference motes. To judge the significance of this fall off, in the following sub-section we focus on performing a baseline comparison.

6.1.2 Comparison with Baseline Techniques

To further test the performance of the Ecolocation algorithm, we compare its results with other simple RF-based localization techniques, namely, nearest reference and weighted 3-centroid. The nearest reference technique simply estimates the unknown location as the location of the nearest reference mote (the mote with the highest RSSI measurement). In weighted 3-centroid method, we calculate the location as the weighted centroid of the reference locations of the three highest received RSSI values (each reference location is weighted by its RSSI value – higher the RSSI value, larger the weight). We choose to compare Ecolocation against these two techniques because they are simple, low computation cost techniques that provide a baseline for localization estimates and can be used to provide a 'sanity' check. That is, if Ecolocation can't outperform these techniques, than its usefulness is not justified. We perform the comparison using the same data, and in the six scenarios outlined from the previous Section.

The CDF and surface plots (with Color Bar from Figure 6.3) for the error in Scenario 1-6 is shown in Figures 6.6- 6.11. For Scenario 1, the mean error for the Ecolocation is 1.559 feet, as previously noted, while the mean error for the nearest reference and weighted 3-centroid are 1.625 feet and 1.832 feet, respectively.

Observing Figure 6.6a, we can see that Ecolocation's CDF is for the most part higher than the weighted 3-centroid's with it being an average of 10.33% points greater in the error range 0.75-2.25 feet (with a maximum of 13.85% at 1.3 feet). As a result, it seems that Ecolocation outperforms weighted 3-centroid rather nicely. The analysis based on the nearest reference technique is not as easy. For starters, the nearest reference

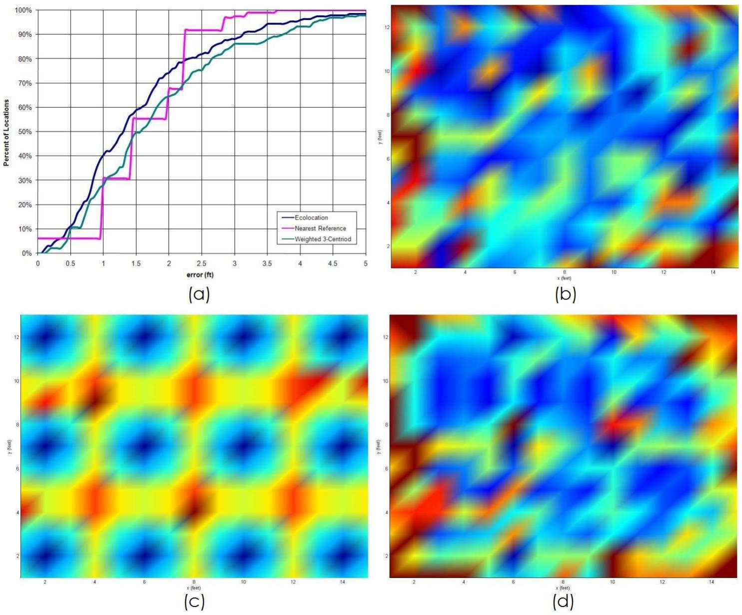

Figure 6.6: For Scenario 1: (a) The CDFs of the three techniques, and the surface plots for (b) Ecolocation, (c) Nearest Reference, (d) Weighted 3-Centroid (see Figure 6.3 for Color Bar)

technique expectedly exhibits an undesirable staircase CDF because of the nature of the point measurements made. Additionally, at the beginning the nearest reference CDF is below Ecolocation, however, at 2.25 feet, it becomes higher. Because of the near-ideal conditions, the nearest reference is almost always able to be correctly picked out. As a result, the location points corresponding to the reference mote location provide an estimate with zero error. On the other hand, Ecolocation provides an average error of 1.834 feet at those locations. The surface plots in Figures 6.6b-c show that nearest reference work well near the reference points whereas the performance of Ecolocation is distributed across the localization area.

For Scenario 2, it is easier to see that Ecolocation outperforms the other two simple techniques. The mean errors for nearest reference and weighted 3-centriod are 2.338 feet and 2.835 feet. In the case of weighted 3-centriod, Ecolocation's CDF is on average greater than 20% point for a spanning of the range 0.9-3.8 feet (with a maximum of 25.13% at 2.35 feet). Additionally, we see that with fewer reference motes, nearest neighbor's performance drops off much more than Ecolocation's performance. Similarly, for the other 8 reference mote scenario, Scenario 3, Ecolocation outperforms the other two techniques (see Figures 6.7- 6.8).

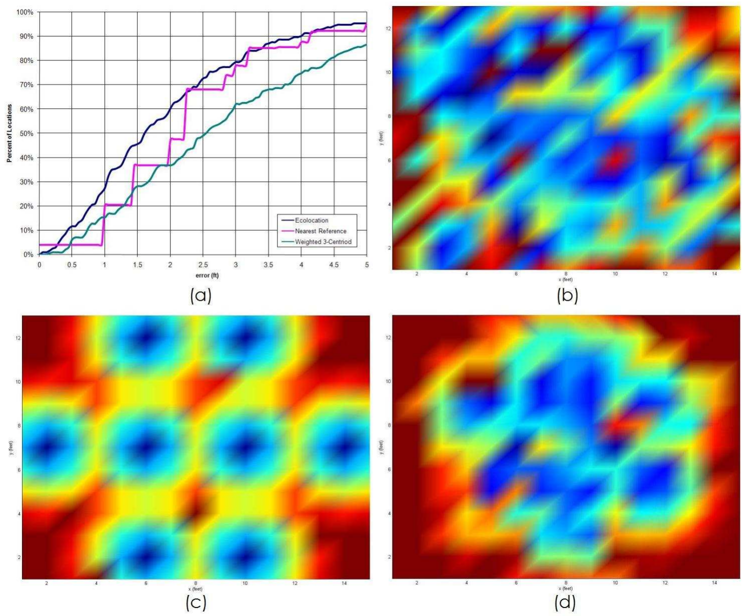

Figure 6.7: For Scenario 2: (a) The CDFs of the three techniques, and the surface plots for (b) Ecolocation, (c) Nearest Reference, (d) Weighted 3-Centroid (see Figure 6.3 for Color Bar)

For the three 6 reference mote scenarios, Scenario 4-6, Ecolocation's outperformance is not quite as significant. It clearly outperforms weighted 3-centriod. Based on the CDF curves, the Ecolocation algorithm seems to be no better than simply choosing the closest reference mote as the location estimate. However, if we look at the surface plots in Figure 6.9- 6.11b-c, we see a different story. Nearest reference performs well near the reference motes, however, at all other locations it performs quite badly. On the other hand, Ecolocation performs adequately across the localization area and provides few dead spots.

Table 6.1 summarizes the comparison of Ecolocation with the two simple techniques for all six Scenarios.

6.1.3 Results in Noisy Environment

Now that we know Ecolocation can provide a good level of accuracy in obstruction-free environments, we want to know how much accuracy is affected with a more realistic environment. Using again the 195 test points outlined in previously, we generate a

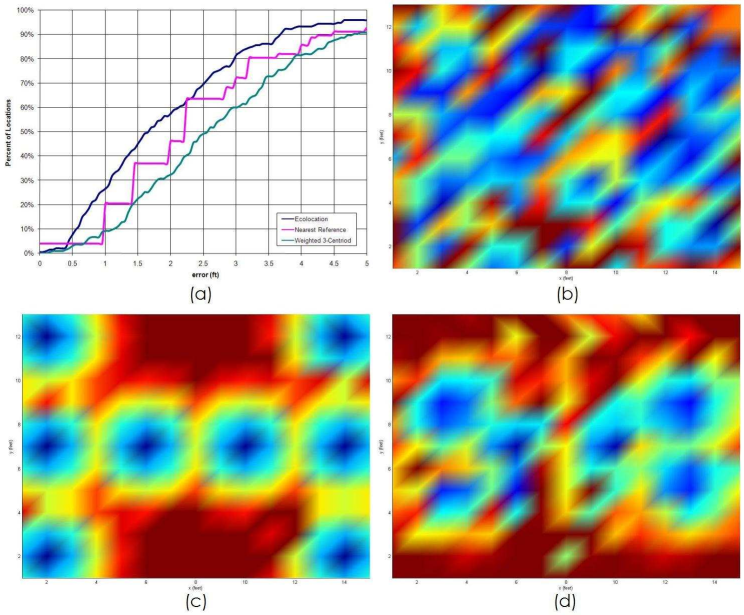

Figure 6.8: For Scenario 3: (a) The CDFs of the three techniques, and the surface plots for (b) Ecolocation, (c) Nearest Reference, (d) Weighted 3-Centroid (see Figure 6.3 for Color Bar)

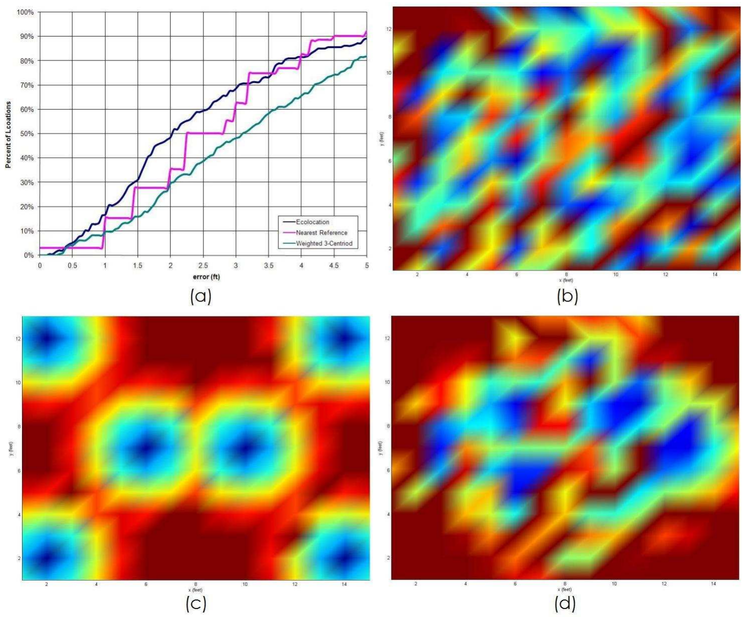

Figure 6.9: For Scenario 4: (a) The CDFs of the three techniques, and the surface plots for (b) Ecolocation, (c) Nearest Reference, (d) Weighted 3-Centroid (see Figure 6.3 for Color Bar)

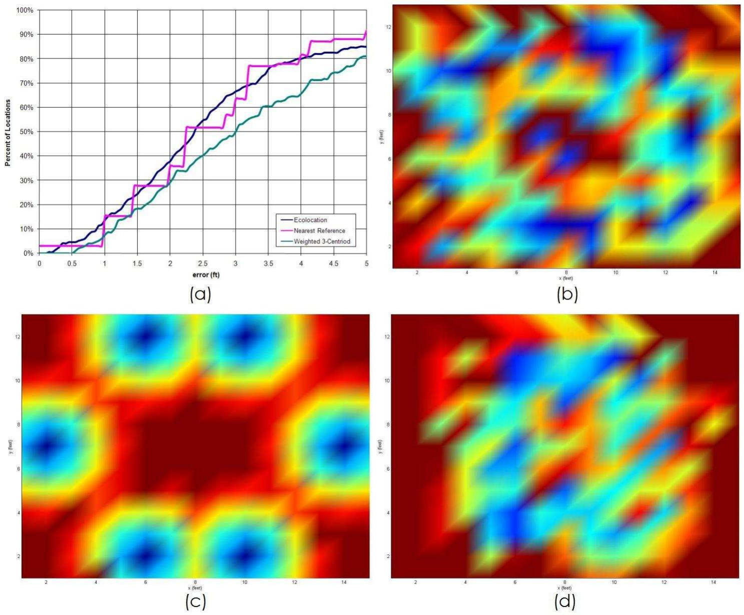

Figure 6.10: For Scenario 5: (a) The CDFs of the three techniques, and the surface plots for (b) Ecolocation, (c) Nearest Reference, (d) Weighted 3-Centroid (see Figure 6.3 for Color Bar)

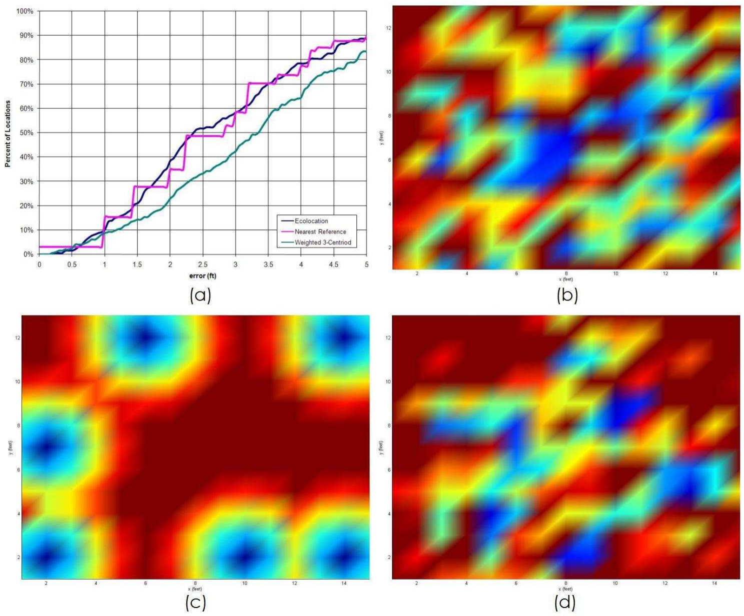

Figure 6.11: For Scenario 6: (a) The CDFs of the three techniques, and the surface plots for (b) Ecolocation, (c) Nearest Reference, (d) Weighted 3-Centroid (see Figure 6.3 for Color Bar)

	Ecolocation	Nearest Reference	Weighted 3-centriod
Scenario 1	1.559	1.625	1.832
Scenario 2	1.987	2.338	2.835
Scenario 3	2.013	2.413	2.809
Scenario 4	2.522	2.683	3.253
Scenario 5	2.765	2.745	3.269
Scenario 6	2.898	2.817	3.393

Table 6.1: Summary comparison of mean estimation error in feet for all six scenarios

CDF curve for the error associated with the Ecolocation algorithm. In Figure 6.12, we can see that the performance is severely impacted. The average jumps from 1.559 feet in the open area to 2.465 feet in the area with barriers which is a 58.1% increase. Now only 42% of the test points have error of 2 feet or less compared to the near 75% exhibited in the open area. We see that Ecolocation is not robust enough to multipath and shadowing and so it cannot be considered and adequate solution for the problem discussed in Section 4.1.

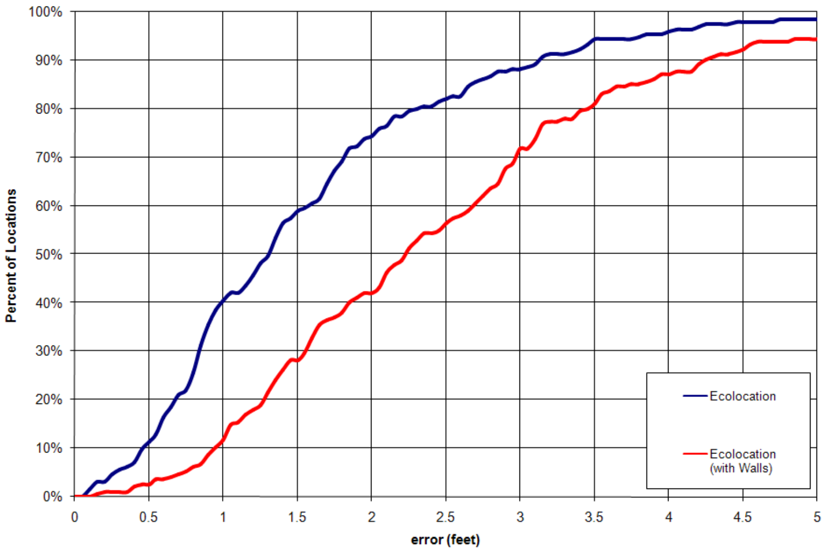

Figure 6.12: The CDFs of esimation error in feet for Ecolocation in open area and with noise.

6.2 Chapter Summary

Following a description of how we implement the coding on the motes and PC, we illustrated our experimental open environment. The algorithm was tested over several different reference mote deployments. We expectedly found that more reference motes meant better accuracy. To validate the performance of Ecolocation, a comparison between two simply techniques was done. From the comparisons between the three techniques, it is apparent that in most cases Ecolocation outperforms the other two. From the surface plots, we have seen that Nearest Neighbor technique can provide areas of great accuracy if the mobile mote is right next to one of the reference motes, but other locations provide much worse results. On the other hand, Ecolocation provides few areas with great accuracy while still possessing a wide coverage region with above adequate performance. This presents a trade-off between a desire to have great accuracy but only in a few regions and having slightly less accurate but much wider coverage region.

A hybrid technique many be possible to get the best of both worlds. That is, when a mobile mote is near a reference mote use Nearest Reference and in other cases use Ecolocation. One possible implementation could be to determine if a single dominant RSSI value is present in the sequence since this might imply that the mobile mote is near that reference mote. In this situation, the hybrid localization algorithm could use the Nearest Neighbor technique. However, when a single dominant RSSI value is not present, we can assume that the mobile mote is not near a reference mote and as a result, the hybrid localization algorithm would decide to use the Ecolocation technique for its wider coverage.

However, when barriers were added to the open area, the performance of Ecolocation was severely impacted negatively. As a result, it was shown that the algorithm was not robust to multipath interference and shadowing effects and so none of the existing localization solutions satisfy the requirements listed in Section 4.1.

Chapter 7

A Weighted Algorithm and its Robustness

In Chapter 4, we showed that the Ecolocation algorithm possessed many of the 'efficiency' qualities necessary for in an indoor WSN-based localization system. However, the issue of accuracy remained a critical open topic. Our performance analysis presented in the previous Chapter showed that Ecolocation performed well in environments where the RSSI values could be trusted and, as a result, had the potential to provide a reasonable level of accuracy. However, as RSSI values became more corrupted, the performance deteriorated quite significantly showing that it was not very robust to multipath and shadowing.

Nevertheless, Ecolocation has the ability to provide good accuracy in obstructionless situations while being efficient. As a result, our task was now to explore if we could improve the algorithm's accuracy in noisy environments while still maintaining the 'efficiency' of the original algorithm.

7.1 Introduction

In the Ecolocation algorithm, the constraint matrices are based on a balance ternary system consisting only of -1's, 0's, and 1's. The constraints are chosen depending on whether a RSSI value (or distance) is less than, equal to, or greater than the RSSI value (or distance) it is being compared to. Note that the raw values are never utilized except for the sign function evaluation involved in the comparison to generate the constraint matrices. The reason for avoiding raw values is logical in that to make sense of them it is required, whether by estimation or other means, some sort of channel parameters. In addition, raw values are easily corrupted by multipath fading and shadowing, and as a result cannot be fully trusted to be accurate. Hence, Ecolocation relies on them to a very minimal extent in making the sign comparison. However, it is our premise that the raw values also should not be entirely ignored. That is, some valuable information can be extracted from them even if they cannot be fully trusted. Our reasoning is simple, if the raw values cannot be trusted, then even comparing them is meaningless. The fact is the raw values can be trusted to some level (depending on the harshness of the environment) and we should try to use them as much as possible rather than as little as possible, which is currently the case in the Ecolocation algorithm.

This Chapter will present our suggested improvement to the Ecolocation algorithm along with some insight to the modification.

7.2 The Modified Algorithm

In this Section, we will discuss the implementation details of our modified algorithm which we refer to as Weighted-Constraints do to the nature of the modification involved. We will also present some rationale behind why the algorithm was designed as such and the affect some parameters have on it. Additionally, we will point out a key effect the results from the modification.

7.2.1 Weighted-Constraints Algorithm

Our proposed modification is to in some way consider the amount a pair of RSSI values differ by in addition the sign of their difference for determining the constraint values. Since the raw values cannot be fully trusted, we do not use them as is, but we associate a certain significance to the amount a pair of RSSI values differ by. Thus, instead of just -1, 0, and 1 representing the constraints, we fix the constraints to range from -x to +x (x is some positive integer parameter, a design choice), where the sign again signifies less than or greater than comparison and zero mean equality and the magnitude representing the amount the value differ by, with 1 meaning they differ only slightly and x meaning they differ a lot. We will refer to x as the scale value and our algorithm as Weighted-Constraints because we weighted every element in the constraint matrix based on its relative magnitude.

Let us illustrate the Weighted-Constraints algorithm with the same example setup used to demonstrate the Ecolocation algorithm in Section 5.2. In this example, we will take x = 5. The RSSI constraint matrix is populated according to the element definition:

$$R_{n\times n}(i,j) = sign(RSSI_i - RSSI_j) * Ceiling\left(\frac{|RSSI_i - RSSI_j|}{RSSI_{\Delta MAX}} * x\right) \tag{7.1}$$

where

$$RSSI_{\Delta MAX} = RSSI_{max} - RSSI_{min}$$

which simplifies to:

$$R_{n\times n}(i,j) = Ceiling\left(\frac{RSSI_i - RSSI_j}{RSSI_{\Delta MAX}} * x\right) \tag{7.2}$$

The first part of the equation is the same sign function used in Ecolocation; the second part is where our suggested modification is introduced. The second part takes into account the degree of difference between the two RSSI values. Using the example RSSI data, Figure 7.1 shows the RSSI constraint matrix generated.

RSSI Ranking

Ref Id	RSSI
1	-50
2	-47
3	-54
4	-49
5	-45
6	-53
7	-57
8	-55
9	-60

RSSI Constraint Matrix

	1	2	3	4	5	6	7	8	9
1	0	-1	2	-1	-2	1	3	2	4
2	1	0	3	1	-1	2	4	3	5
3	-2	-3	0	-2	-3	-1	1	1	2
4	1	-1	2	0	-2	2	3	2	4
5	2	1	3	2	0	3	4	4	5
6	-1	-2	1	-2	-3	0	2	1	3
7	-3	-4	-1	-3	-4	-2	0	-1	1
8	-2	-3	-1	-2	-4	-1	1	0	2
9	-4	-5	-2	-4	-5	-3	-1	-2	0

Figure 7.1: Using the RSSI data from Section 5.2, the corresponding RSSI constraint matrix for the Weighted-Constraints algorithm is generated. $RSSI_{\Delta MAX} = 15$.

As with the Ecolocation Algorithm, we chose a scanning resolution and determine the grid point constraint matrix for each grid point. The Weighted-Constraints grid point constraint matrix is populated according to the element definition:

$$C_{n\times n}^{(m,n)}(i,j) = -sign\left(d_i^{(m,n)} - d_j^{(m,n)}\right) * Ceiling\left(\frac{\left|d_i^{(m,n)} - d_j^{(m,n)}\right|}{d_{\Delta MAX}} * x\right) \tag{7.3}$$

where

$$d_{\Delta MAX} = d_{max} - d_{min}$$

which simplies to:

$$C_{n\times n}^{(m,n)}(i,j) = -Ceiling\left(\frac{d_i^{(m,n)} - d_j^{(m,n)}}{d_{\Delta MAX}} * x\right) \tag{7.4}$$

Notice again that the first part of the equation is simply the negative sign function used by the Ecolocation algorithm in determining the grid point constraint matrix and the second part is where the weighted-constraints modification is introduced. Based on the two element definitions, the special case of x = 1 leads to the Ecolocation algorithm. Figure 7.2 shows the grid point constraint matrix for the example grid point (10, 35).

Distance Ranking

Ref Id	Distance
1	7.894
2	8.504
3	11.503
4	3.132
5	4.451
6	8.934
7	2.704
8	4.161
9	8.793

Grid Point Constraint Matrix

	1	2	3	4	5	6	7	8	9
1	0	1	3	-3	-2	1	-3	-3	1
2	-1	0	2	-4	-3	1	-4	-3	1
3	-3	-2	0	-5	-5	-2	-5	-5	-2
4	3	4	5	0	1	4	-1	1	4
5	2	3	5	-1	0	3	-1	-1	3
6	-1	-1	2	-4	-3	0	-4	-3	-1
7	3	4	5	1	1	4	0	1	4
8	3	3	5	-1	1	3	-1	0	3
9	-1	-1	2	-4	-3	1	-4	-3	0

Figure 7.2: The distance values from grid point (10, 35) to all the reference nodes and the corresponding grid point constraint matrix based on the Weighted-Constraints Algorithm.

For determining which location constraint matrix most closely matches the RSSI constraint matrix, we compare like in the Ecolocation the cells of the matrix to see which ones match and which ones don't. As before, we add one if the constraints match; however, for non-matching constraints, we subtract based on there amount of difference. That is, the count is updates based on the equation:

$$counter^{(m,n)} += 1 - \left| C^{(m,n)}_{n\times n}(i,j) - R_{n\times n}(i,j) \right| \tag{7.5}$$

As a result, the count is incremented by a maximum 1 and a decremented by a maximum $2 * x - 1$. For the example grid point of (10, 35) the counter has a value -139. The estimated location is the centroid of the location points with the maximum counter value. For this example, the maximum counter value is 73 and it occurs at the single grid point (20, 16) representing location (2.5, 8.75).

7.2.2 Algorithm Discussion

The idea behind using weighted-constraints is simply to get a sense of how well matched the constraints are. The original algorithm hinges on the assertion that RSSI and distance separation are correlated. As a result, we simply want to know how strongly or weakly matched the constraints are based on this assertion. That is, if a mobile mote sees a large RSSI difference between two reference motes, it is reasonable to think that the there is a large difference between the distances to these reference motes. Therefore, if we have a RSSI constraint matrix element that has a value +10 and a grid point constraint matrix element that has a value +2 based on reference motes 'One' and 'Two', the Weighted-Constraints algorithm tells us that this constraint is only weakly match whereas Ecolocation would just say it is a match. We are saying that a grip point constraint matrix element with a larger positive value would be a better match for this RSSI constraint matrix element since both would convey that 'One' is closer to the grid point than 'Two'; however, Ecolocation doesn't allow for such a determination to occur. Figure 7.3 illustrates this point. Likewise, if we have the scenario where the RSSI constraint matrix element that has a value +2 and a grid point constraint matrix element that has a value -1, the Weighted-Constraints algorithm can tell us that the constraint is reasonably matched whereas Ecolocation would simply dismiss it as a no match.

The strength/weakness of the comparison between constraints is taken into account

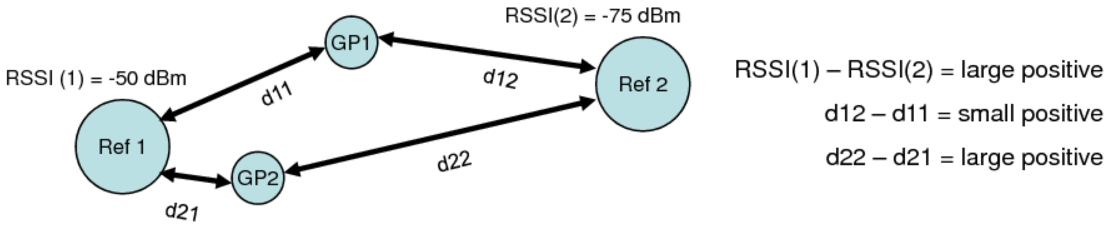

Figure 7.3: We will compare two grid points to see which is a better estimate. We have received a high RSSI value from Reference Mote 1 and a low RSSI value from Reference Mote 2 causing a large positive constraint element. From grid point 1, Reference Mote 1 is only slightly closer then Reference Mote 2 causing a small positive constraint element. From grid point 2, Reference Mote 1 is much closer then Reference Mote 2 causing a large positive constraint element. Based on this single constraint, the Weighted-Constraints algorithm would consider grid point 2 to be a better estimate than grid point 1, while Ecolocation would consider them equally good.

by how the count is incremented. As stated before, when the comparison results in a perfect match, the count value is increment; otherwise it is decremented by the amount of difference. The institution behind decrementing it as such is that a weak match implies an incorrect grid point (this grid point is not a good estimated location based on the RSSI values) and the difference implies the level it is incorrect. As a result, a higher difference means a worse estimate, so the count is decremented more.

The design parameter ‘x’ directly relates to how much you trust the RSSI values and the environment. In the extreme case where the environment is ideal and RSSI values are perfectly correlated with distance, ‘x’ should be infinite.

7.2.3 Smooth-out Effect

By using weighted constraints, we are able to smooth out the estimated regions determined by the count values. Let us illustrate what we mean by take a look at the surface plot of the count values for a particular test point (using the RSSI data gathered in Section 6.1.1). Location points with high count values are represent by red and those with low are represent by blue. In Figure 7.4, we show the surface plots of the original and weighted-constraints algorithm for test point (9, 5) feet (the location of the mobile mote). The Ecolocation surface plot shows that the regions are rigidly defined. This observation is to be expected considering that the localization space is divided into regions based on the perpendicular bisectors. The counter for all location points in those regions would be the same and hence the surface plot has rigid, protruding boundaries. However, observing the Weighted-Constraints' surface plot, we see that the boundaries are smooth and fluid with a gradual transition. The benefit of this is that not only is the maximum region more well defined, but also allows the use of lesser regions to fine tune the estimation since the regions fall off radially.

7.3 The Weighted-Constraints Results

We will now compare the performance of Weight-Constraints with Ecolocation in terms of accuracy. We will use the same experimental environments used in the analysis of Ecolocation in the previous Chapter.

7.3.1 Results in Open Environment

We will first present results using real data from the relatively non-harsh environment. For analyzing the performance of the algorithms, we generate cumulative distribution functions (CDFs) for the estimation error based on the 195 test points outlined in Section 5.3.2 and the stored RSSI sequences. In Figure 7.5 , the CDF with the scale value of 10, along with the Ecolocation algorithm, is shown and it can be seen that the Weighted-Constraints algorithm outperforms the original algorithm. For the low error range between 0.5 feet and one foot, the performance difference is not very significant with the weighted-constraints algorithm performing on average only 2.84% points better (the height difference between the curves). However, for the range of 1 foot of error to 2.6 feet, weighted-constraint algorithm performs on average 7.43% points better with a maximum of 11.79% at 1.65 feet. Furthermore, the original algorithm provides location estimates of the tests point with an average error of 1.559 feet. On the other hand, the Weighted-Constraints algorithm exhibits a 13.1% improvement resulting in an average error of 1.355 feet.

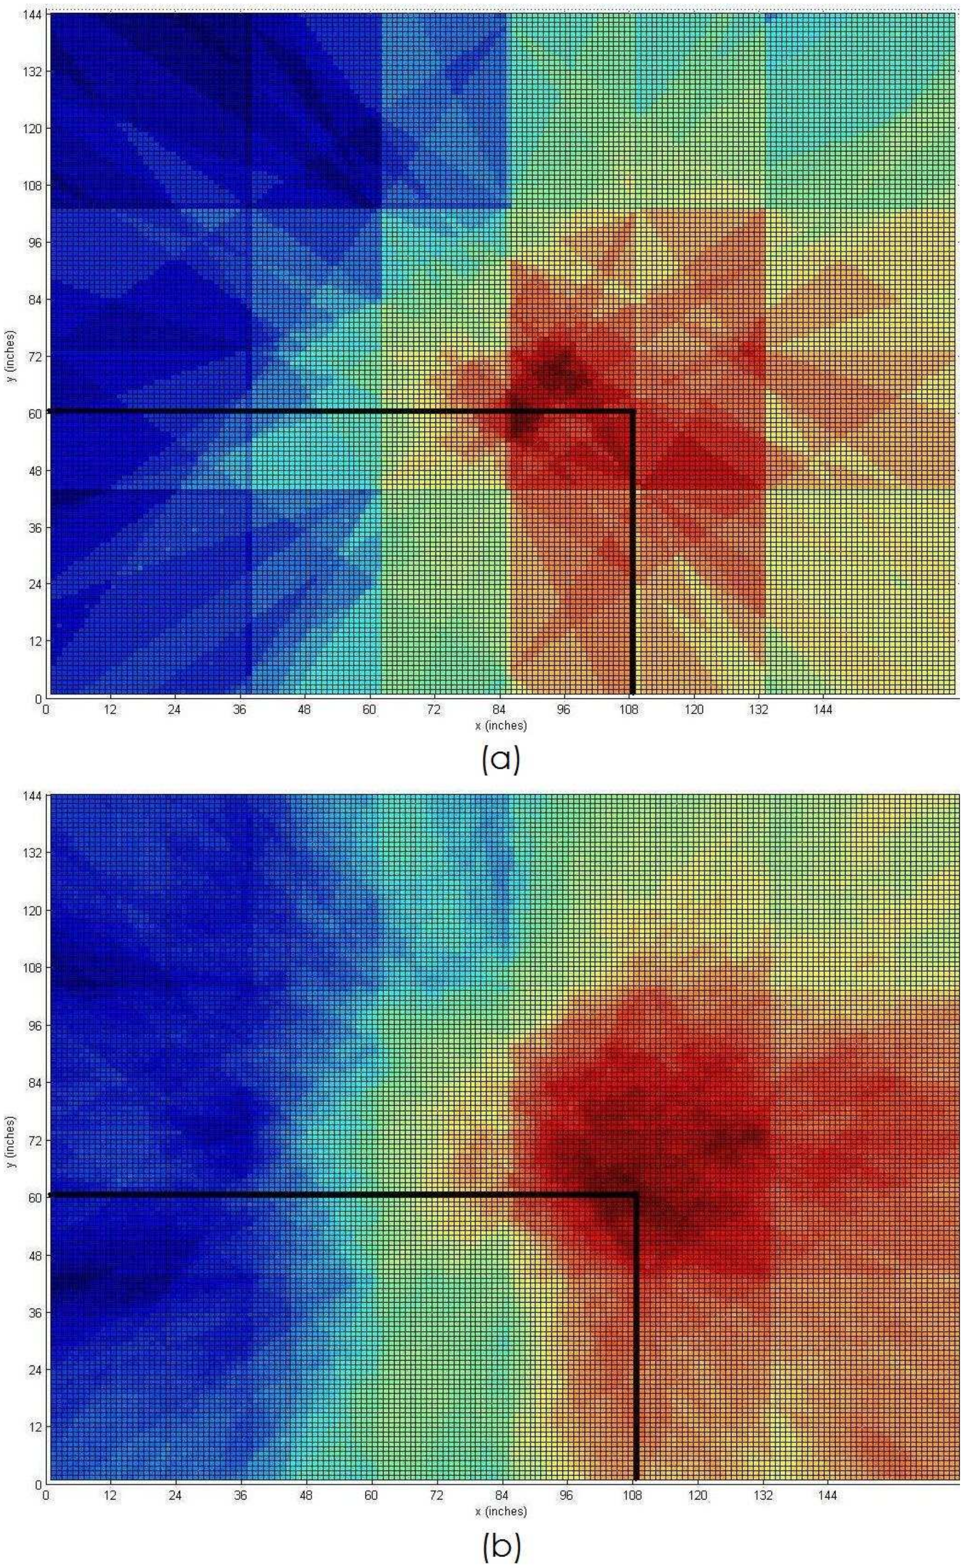

Figure 7.4: (a) Ecolocation and (b) Weighted-Constraints: The surface-plot of count value over all grid points for a mobile mote located at (11, 8) feet. Red signifies high values, blue signifies low values

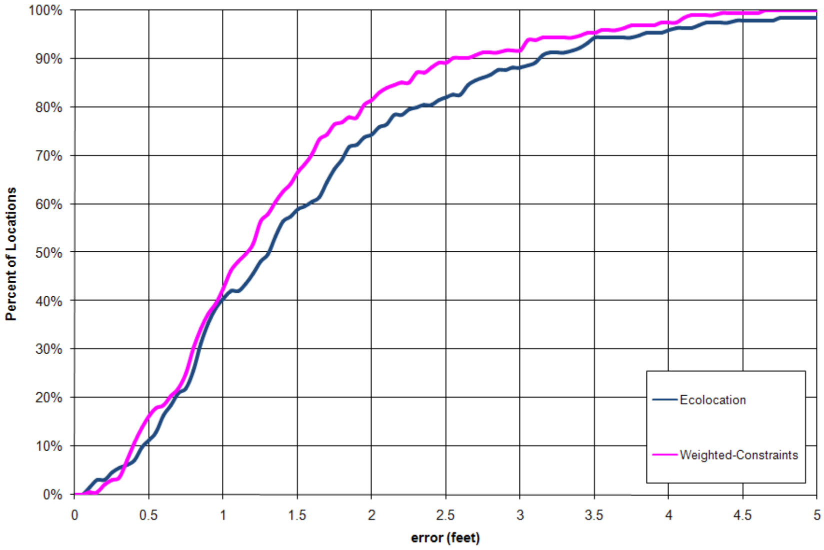

Figure 7.5: The CDF of Ecolocation and Weighted-Constraints (with x=10)

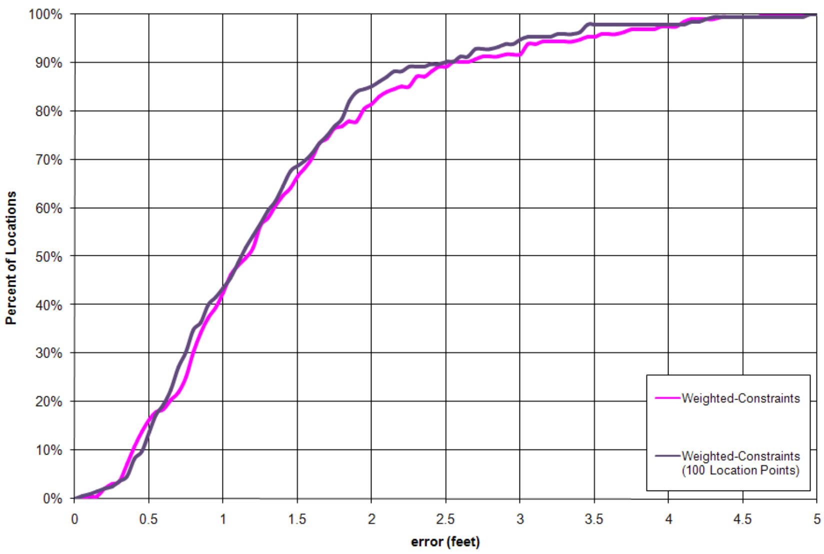

Figure 7.6: The CDF of Weighted-Constraints and Weighted-Constraints with minimum 100 location points used in estimation. Both with x=10.

Fine-tuning

As we mentioned in Section 7.2.3, the Weighted-Constraints has the feature that the count value forms a radial pattern. As a result, it is conceivable that accuracy could be improved my including grid point with count value below the maximum. The another reason for doing this is that Weighted-Constraints' count values are more unique than Ecolocation which can mean only one grid point could be with the maximum value (as was the case with the example in Sections 5.2 and 7.2.1 where Ecolocation produced 13 grid points whereas Weighted-Constraint had only 1 grid point). In Figure 7.6, we show the performance obtain when we require the Weighted-Constraints algorithm to use a minimum 100 location points in the centroiding (the 100 highest count values are used). The average error with fine-tuning improves on Weighted-Constraints by 4.5% resulting in 1.295 feet of error.

7.3.2 Results in Noisy Environment

The true benefits of the Weighted-Constraints algorithm will be fully recognized when compared with the performance of the original algorithm in a noisy environment. In this Section, we will see that Weighted-Constraints' performance is not as severely deteriorated when the RSSI values become corrupted. Furthermore, the Weighted-Constraints algorithm continues to get a reasonable level of accuracy in the presence of multipath and shadowing.

Using the stored RSSI sequence from Section 6.1.3, we generate the CDF curve associated with the Weighted-Constraints algorithm. In Figure 7.7, we see the comparison of the performance in the open area and the area with the barriers. Although, the accuracy takes a performance hit, it is not as severe as was the case with Ecolocation in Figure 6.12. In the noisy environment, Weighted-Constraints performs with an average error of 1.942 feet compared to the average error of 1.355 feet in the open area which is a 43.3% increase.

Finally, in Figure 7.8, we see the performance of both algorithms in the noisy environment. We see that with a harsher environment Weighted-Constraints more greatly

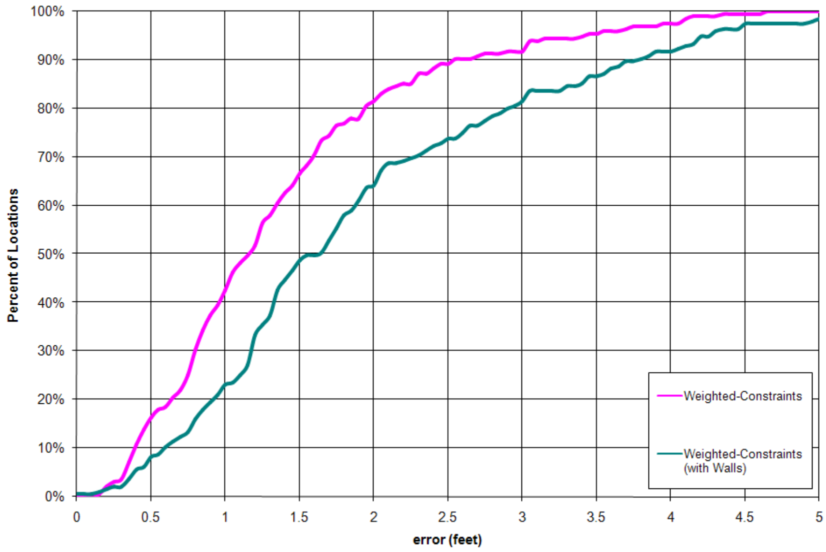

Figure 7.7: The CDFs of Weighted-Constraints in open area and in noise.

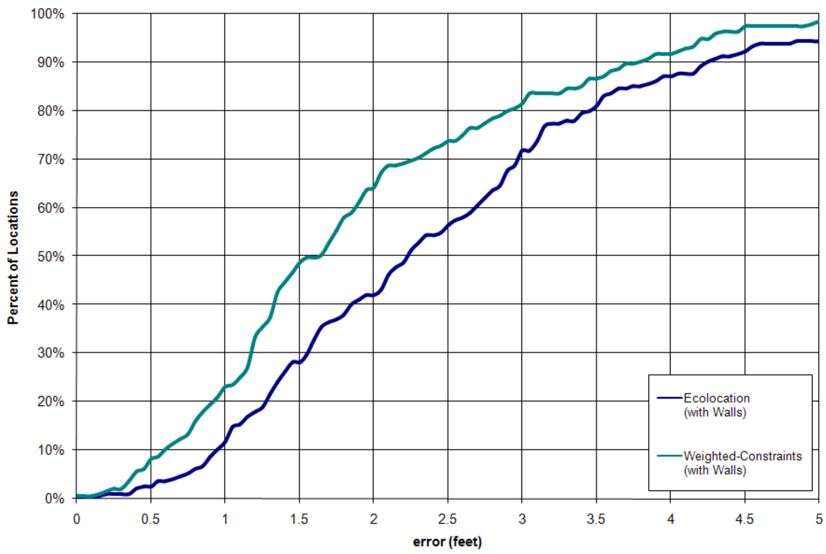

Figure 7.8: The CDFs of Ecolocation and Weighted-Constraints in noise.

outperforms Ecolocation than it did in a relatively ideal environment. That is, Ecolocation provides on average 2.465 feet of error while Weighted-Constraints improves on Ecolocation by 21.2% resulting in an average error of 1.942 feet. Compared to the 13.1% improve Weighted-Constraints showed in the open area, this result shows that Weighted-Constraints can even more greatly outperform Ecolocation in harsh environments which means it is more robust to multipath and shadowing. Furthermore, with Weighted-Constraints increased performance, nearly 65% of the location points have error less than 2 feet in the noisy conditions which is only a 10% drop observed with Ecolocation in the obstruction-less case.

7.4 Chapter Summary

In this Chapter, we have described our suggested improvement to the Ecolocation algorithm, Weighted-Constraints, to provide higher accuracy when environment conditions are not ideal. The concept of strongly/weakly match constraints was introduced and how they can produce more insightful constraint comparisons was shown. We revealed the inherit smooth-out effect the modification has to the count value for the grid points and how this feature can be used to our advantage by allowing the grid point of lower count values to help in the location estimation. Next, we displayed how the Weighted-Constraints algorithm compared in the same two test environments used to evaluate Ecolocation performance. We see that Weighted-Constraints outperforms Ecolocation in both cases and we see that its improvement is even more prevalent in the noisy environment. Futhermore, Weighted-Constraints satisfies all the criteria from Section 4.1.

Chapter 8

Conclusions

The problem of whether an efficient solution is feasible to provide accurate position estimates in wireless sensor networks based localization has been solved as shown in the performance results of Section 7.3.2. In Chapters 3-4, we found that only one of the existing state-of-the-art WSN-based localization techniques, Ecolocation, satisfied the efficiency-based criteria. We discovered in Section 6.1.1 that Ecolocation had the potential to provide a necessary level of accuracy; however, in Section 6.1.3, we saw that accuracy performance diminished greatly in the presence of multipath fading and shadowing effects. As a consequence, Ecolocation was determined not to be an acceptable solution to our problem and so none of the existing solutions meet the entire criteria necessary for an indoor WSN-based localization system.

However, Ecolocation's ability to be efficient and have potential for necessary accuracy made it an attractive possibility. As a result, we enhanced the algorithm to keep the efficient qualities while improving the accuracy, specifically accuracy in a noisy environment. Our revised algorithm is called Weighted-Constraints and Section 6.1.1 showed that it provided an 13.1% improve in terms of average error in an open environment compared to the original Ecolocation algorithm. Furthermore in Section 7.3.2, we found that Weighted-Constraints is more robust to noise by observing that the addition of noise to the environment causes an accuracy loss of only 43% compared to the 58.1% that Ecolocation suffered. As a result, Weighted-Constraints greatly out-performed Ecolocation in terms of average error in a noisy environment by 21.2%. We created a localization solution that was efficient and provides accurate results.

Chapter 9

Future Work

The performance of the algorithm using less correlated motes remains an open issue. We identified that Mica2 provide better distance-to-RSSI value correlation than Micaz and Tmote Sky. However, the amount and the manner in which the less correlation will effect performance is unclear. This type of error is not the same we experience with multipath and shadowing where RSSI values are randomly corrupted. In this case, the RSSI value themselves can't be trusted to begin with.

We believe a study into whether Ecolocation and as an extension Weighted-Constraints can be deployed in 3-dimensions is necessary. That is, currently we have studied performance with the unknown mote and all reference motes in a single plane. By requiring all reference motes be mounted on a single plane, the reference motes have placement restrictions which make them dedicated localization motes. That is to say, motes being used around the building for other functionality (such as monitoring ambient temperature) may not be used passively as reference motes because they may not be on the same plane and it is unfeasible to require them to be placed in such a manner. A 3-dimensional algorithm provides a more realistic deployment and that non-localization motes can be as reference motes by the localization system. It is unclear whether the accuracy gains in the z-direction will cause accuracy losses in the x- and y-directions and if it does will the gains outweigh the losses.

A study into reference mote placement would also be useful. This includes two parts. First, we need to see the affects a non-grid style reference mote deployment has. Unless we use dedicate localization reference mote, it may not be feasible to for non-localization motes to be located in a grid. The second part is to determine if their can be benefits had with a clustering of reference motes. That is, if there are several motes in close proximity to each other, is it beneficial to treat them as a single cluster. Multipath and shadowing vary greatly in direction so each individual reference mote would experience a different level of corruption. Using the different RSSI values in some combination and treating the individual reference motes as a single cluster at a single location might help alleviate some of the multipath and shadowing effects.

Finally, it would be valuable to see if the algorithm can be modified such that fixed location reference nodes are not needed and relative locations are determined. The algorithm can work with few reference motes but we have see that accuracy suffers quite a lot in those case. For example, using 10 unknown location motes and having each of them determine their RSSI sequence based on the others. Can we use the 10 RSSI sequences together to determine the relative location of the 10 motes?

References

[1] October 1th, 2008. http://www.xbow.com/products/product_pdf_files/wireless_pdf/mica2_datasheet.pdf.

[2] October 1th, 2008. http://focus.ti.com/lit/ds/symlink/cc1000.pdf.

[3] October 1th, 2008. http://www.xbow.com/products/product_pdf_files/wireless_pdf//micaz_datasheet.pdf.

[4] October 1th, 2008. http://www.sentilla.com/pdf/eol/tmote-sky-datasheet.pdf.

[5] GPS guide for beginners, October 1th, 2008. http://www8.garmin.com/aboutGPS/manual.html.

[6] J. Albowicz, A. Chen, and Lixia Zhang. Recursive position estimation in sensor networks. In *Network Protocols, 2001. Ninth International Conference on*, pages 35–41, 2001.

[7] A. Arora, P. Dutta, S. Bapat, V. Kulathumani, H. Zhang, V. Naik, V. Mittal, H. Cao, M. Demirbas, M. Gouda, Y. Choi, T. Herman, S. Kulkarni, U. Arumugam, M. Nesterenko, A. Vora, and M. Miyashita. A line in the sand: A wireless sensor network for target detection, 2004.

[8] Paramvir Bahl and Venkata N. Padmanabhan. Radar: An in-building RF-based user location and tracking system. In *INFOCOM (2)*, pages 775–784, 2000.

[9] N. Bulusu, J. Heidemann, and D. Estrin. Gps-less low cost outdoor localization for very small devices, 2000.

[10] Diane J. Cook and Sajal K. Das. How smart are our environments? an updated look at the state of the art. *Pervasive Mob. Comput.*, 3(2):53–73, March 2007.

[11] S. Gezici, Zhi Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *Signal Processing Magazine, IEEE*, 22(4):70–84, 2005.

[12] Andy Harter, Andy Hopper, Pete Steggles, Andy Ward, and Paul Webster. The anatomy of a context-aware application. In *Mobile Computing and Networking*, pages 59–68, 1999.

[13] Tian He, Chengdu Huang, Brian M. Blum, John A. Stankovic, and Tarek Abdelzaher. Range-free localization schemes for large scale sensor networks. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 81–95, New York, NY, USA, 2003. ACM Press.

[14] Jeffrey Hightower and Gaetano Borriella. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, 2001.

[15] S. Lanzisera, D. Lin, and K. Pister. Rf time of flight ranging for wireless sensor network localization. In *Workshop on Intelligent Solutions in Embedded Systems*, 2006.

[16] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson. Wireless sensor networks for habitat monitoring, 2002.

[17] Maroti, Volgyesi, Dora, Kusy, Nadas, Ledeczi, Balogh, and Molnar. Radio interferometric geolocation. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 1–12, New York, NY, USA, 2005. ACM Press.

[18] Radhika Nagpal. Organizing a global coordinate system from local information on an amorphous computer. In *A.I. Memo 1666*. MIT A.I. Laboratory, 1999.

[19] Neal. Using proximity and quantized rss for sensor localization in wireless networks. In *WSNA '03: Proceedings of the 2nd ACM international conference on Wireless sensor networks and applications*, pages 20–29, New York, NY, USA, 2003. ACM Press.

[20] D. Niculescu and B. Nath. Ad hoc positioning system (aps). In *Global Telecommunications Conference, 2001. GLOBECOM '01. IEEE*, volume 5, pages 2926–2931 vol.5, 2001.

[21] D. Niculescu and Badri Nath. Ad hoc positioning system (aps) using aoa. volume 3, pages 1734–1743 vol.3, 2003.

[22] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero, R. L. Moses, and N. S. Correal. Locating the nodes: cooperative localization in wireless sensor networks. *Signal Processing Magazine, IEEE*, 22(4):54–69, 2005.

[23] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43, New York, NY, USA, 2000. ACM Press.

[24] C. S. Raghavendra. *Wireless Sensor Networks*. Springer, July 2005.

[25] Andreas Savvides, Chih-Chieh Han, and Mani B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *MobiCom '01: Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 166–179, New York, NY, USA, 2001. ACM Press.

[26] Roy Want, Andy Hopper, Veronica Falcão, and Jonathan Gibbons. The active badge location system. *ACM Trans. Inf. Syst.*, 10(1):91–102, January 1992.

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

[28] Kiran Yedavalli and Bhaskar Krishnamachari. Sequence-based localization in wireless sensor networks. *Transactions on Mobile Computing*, 7(1):81–94, 2008.

[29] Kiran Yedavalli, Bhaskar Krishnamachari, Sharmila Ravula, and Bhaskar Srinivasan. Ecolocation: a sequence based technique for rf localization in wireless sensor networks. In *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on*, pages 285–292, 2005.