WIRELESS DEVICE ORIENTATION ESTIMATION

An Undergraduate Research Scholars Thesis

by

DEREK HEIDTKE

Submitted to the Undergraduate Research Scholars Thesis program
Texas A&M University
in partial fulfillment of the requirements for the degree of

UNDERGRADUATE RESEARCH SCHOLAR

Approved by Research Advisor: Dr. Jean-Francois Chamberland

May 2017

Major: Electrical and Computer Engineering

ABSTRACT

Wireless Device Orientation Estimation

Derek Heidtke
Department of Electrical and Computer Engineering
Texas A&M University

Research Advisor: Dr. Jean-Francois Chamberland
Department of Electrical and Computer Engineering
Texas A&M University

Estimation of wireless device location (localization) has been studied extensively for its usefulness in network infrastructures. However, the related issue of determining device orientation has received less attention. Many contemporary electronic devices are equipped with inertial measurement units (IMUs) and, hence, they are aware of their own orientation. Yet in some situations, it may be valuable to estimate device orientation without relying on its IMU. This research uses the idea that a distributed array of sensing antennas can measure the differences in signal strength of a transmitting device from multiple directions. Combining the device's antenna characteristics with a simulation allows us to experiment with and predict the orientation estimation performance of different arrangements of sensing antennas. First, an anechoic chamber is used to ensure that the device characterization contains as little external noise as possible. Then, simulation software allows more freedom in the placement and testing of sensing antenna arrangements. Finally, an estimation system takes the previous data and returns an estimate that minimizes the mean squared error of the device orientation in question. Upon completion, we plan to answer whether or not this method is accurate or feasible, and what advantages it

may have over other methods of obtaining orientation information. As a side objective, we expect to learn if there is a general choice or method of choosing the sensing array used for orientation estimation. Though this project is concerned with orientation of antennas using the 2.4 GHz ISM band, the concepts could have application for a wide variety of devices that operate in various spectral bands.

ACKNOWLEDGMENTS

I would like to thank Dr. Jean-Francois Chamberland and his research group: Pranay Eedara, Hai Li, Mandel Oates, Austin Taghavi, Travis Taghavi, and Michael Bass for welcoming me into their team and helping me through the various problems that I faced. A special thanks also goes to Pranay Eedara and Abhay Shankar Anand.

TABLE OF CONTENTS

Page

LIST OF FIGURES

FIGURE Page

1. INTRODUCTION

1.1 Background

Contemporary smart phones almost universally contain Wi-Fi interfaces and they are designed to communicate with access points regardless of their positions. In reality, however, due to manufacturing compromises, the radiation pattern of a typical cellphone is not perfectly isotropic and therefore has subtle directional imperfections. This property makes them the ideal test devices for this research initiative. The concept of applying inference techniques to wireless devices is not new. It has received attention in the past, especially in the context of source localization. In this popular problem formulation, the goal is to determine the relative physical positions of devices on a wireless network. Typically, localization studies rely on methods like signal strength methods [1, 2], Bayesian networks [2], location fingerprinting [3], and algorithms based on angle-of-arrival [4]. Some authors consider the joint problem of localization and orientation estimation [5]. For example, Rohrig and Kunemund propose a system that estimate both location and orientation using radio signals [5]. Nevertheless, the literature on orientation estimation remains comparatively sparse with many opportunities for novel contributions.

1.1.1 Significance

Orientation has a role to play in network security. For instance, information about a device's location can be used by attackers to gain control over and corrupt its operation. Orientation also has the potential to be a valuable tool for developers working on future communication infrastructures. As mobile device technology advances, the need and potential benefits for orientation information continues to grow. While orientation information is typically supplied by inertial measurement units (IMUs), gyroscopes, accelerometers, and magnetometers [6], the ability to measure orientation in an exogenous

fashion is beneficial for validation and monitoring. As such, we aim to develop a system in which dedicated on-device hardware would become redundant. We plan to determine the orientation of a wireless device without any dedicated hardware attached to it. This is theoretically possible because consumer cellphones have unique antenna directionality.

As mentioned before, each distinct model has a unique radiation pattern that varies as a function of exit angle. Conceptually, identifying the cell phone's orientation can be separated into two phases: an initial training phase during which the sensing array learns the radiation pattern of the cellphone, and a testing phase where the cellphone orientation is unknown a priori and estimated based on measured data. These aims have the potential to provide alternative access to this information which can be beneficial from a cost, convenience, or security perspective. Also, if successful, this technology can be employed to infer the orientation of a wireless robot [7] or detect irregular behavior in a human operator [8]. As wireless sensing devices become more common, the potential and ease of applying this technology grows.

1.1.2 Techniques

The device's radiation pattern is characterized by exhaustively measuring the signal strength emitted from all angles. Once a characterization of the device is known, then it becomes appropriate to consider which types of algorithms are going to be used on the data. There are many choices, each with different strengths and goals. Is it better to use a supervised algorithm, unsupervised algorithm, decision tree, regression, instance-based, regularization, clustering, Bayesian, artificial neural network, or perhaps even a combination thereof?

Because the goal is to estimate the device's orientation (specified as an ordered pair of angles, azimuth and elevation), then the best course of action will be to use an algorithm that can make predictions about continuous possibility spaces. Since these angles can take

on a continuum of values, the ideal estimation algorithm might be some sort of regression method. Later we will discuss the method chosen for this research, the maximum likelihood algorithm.

1.1.3 Localization

As mentioned earlier, there are a few established ways of determining the location of a wireless device: using IMU's, angle-of-arrival, and fingerprinting.

The method of IMU localization relies on measuring the forces applied to a device. Any time an object moves, it experiences some acceleration. Acceleration is intimately related to location (its second derivative). For instance, if an IMU were to record all accelerations of a device, it would have enough information to guess where the device is. However, this method does not usually have a good way to correct for acceleration noise, and relies on knowing the initial position. This is not always practical.

Angle-of-arrival methods use antennas that are sensitive to incident transmission angle. In theory, this can be combined with signal strength and free-space path loss to locate the device. The system is able to measure both the angle and distance of the target device. Assuming a polar coordinate system, these are the only two pieces of information that are needed. Again, though this method suffers from measurement error (in both the angle and distance components), and the assumption that the device is transmitting at a constant intensity at all times. Multi-path effects may also play a role in confounding this method.

Location fingerprinting attempts to learn the effects that local objects have on the received signal strength at a given sensor; though, this will not work if the reflection characteristics of the location change.

After reviewing the common localization techniques, we decided to pursue this research by using a signal strength method. Since we are dealing directly with device radiation patterns, it makes sense to use signal strength as the main metric. Also, it is

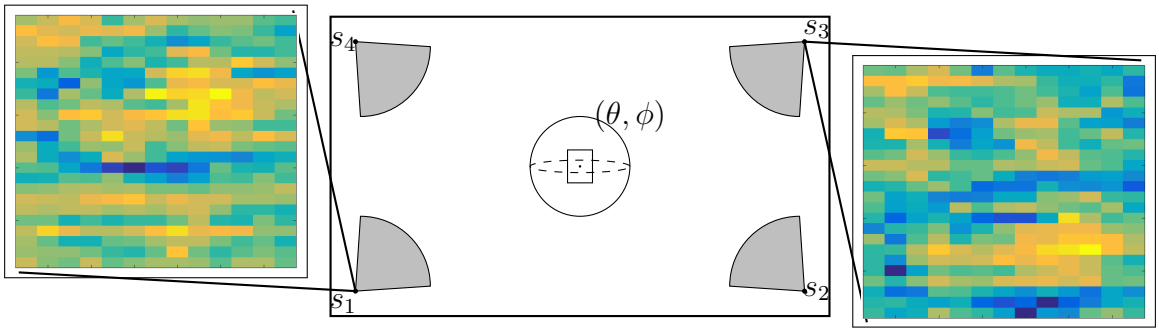

Figure 1.1: Diagram that depicts the general idea of a data collection system. System needs to be able to remotely measure signal strength of a target device from multiple angles.

straightforward to extract signal strength information from a Wi-Fi transmission.

1.2 Goals

In this section, we discuss the specific steps needed to complete the project. A functional system will require the discussed items. The next chapter explains, in detail, the specific choices that were made to accomplish these goals.

1.2.1 Data Collection

This project requires the use of both hardware and software. On the hardware side, it needs a platform that can easily take snapshot measurements of the target device from many different angles using some form of synchronization. Estimation is done wirelessly using transmission intensity; therefore, the system needs both an array of antennas and a way to measure the signal strength of the incoming transmission. In general, the system should look something like what is illustrated in Figure 1.1.

Briefly, the system diagrammed in Figure 1.1 was implemented as a network of computers including a server and a number of clients. The clients are equipped with antennas that allow them to gather signal strength data from the target device. The data is then relayed back to the server over the network. In the background, the server issues synchro-

nized commands over the network so that all antennas take measurements approximately at the same time.

1.2.2 Algorithm

The software component of this system is a routine that can estimate the orientation of the target device. The parameters of this routine are the radiation pattern of the target device, the number and locations of the sensing antennas, and the noisy measurements from those antennas. A working estimation system will do some calculations with this given information, and output a predicted orientation for the device.

This algorithm was implemented in Python, using the `NumPy` library. The estimation system computes the output orientation using an exhaustive maximum likelihood technique. Pytohn was chosen becasue it is simple, quick to develop with, and easy to use.

1.2.3 Simulation

Once the radiation pattern of a device has been obtained, we would then like to use it in simulations. This will allow us to rapidly assess the predicted performance of a given arrangement of subject antennas. Simulations are incredibly helpful because setup for a single real experiment takes a long time, and produces a small amount of data. Ultimately, the goal is to find an optimal solution through the use of simulations so that valuable time is only spent on the sensor arrangements that have been predicted to perform well.

Simulation through virtualization was implemented using Python, similar to the above goal. It has the advantage of being very readable and easy to debug, while also promoting rapid development. One drawback to Python is, it can run between ten and one hundred times slower than the same program written in C++ for computationally intensive tasks because it's an interpreted language.

1.3 Theory

1.3.1 Antenna Radiation Patterns

Consider the following hypothetical situation: an immobile access point provides signal coverage for some number of hosts. The simplest way that this access point can reach all of the devices is to have an antenna that radiates isotropically. In an intuitive way, we can see that this is also a fair setup; no host attempting to receive from the access point will experience a weaker signal than any other host that is the same distance away from the access point. If a host is also immobile, but requires a stronger signal because it is too far away, the access point need only to increase its transmitting power.

However, this situation is neither practical nor ideal. In the real world, there are a myriad of constraints that an antenna must balance. For instance, it is very common to place power usage limits on a design. In the previously mentioned scenario, the access point would be wasting most of the energy that it transmits. A directional antenna would be more efficient if one could assume the host's direction relative to the access point. Another example, would be a shape constraint. In mobile devices, it is absolutely critical that the antenna fit within the device. Moreover, there is increasing pressure to shrink that envelope smaller and smaller as technology improves. In a situation like this, the final design might be mostly isotropic, with the exception of being highly irregular. That is, the amount of energy transmitted as a function of exit angle would not be smooth and uniform. For the designers that are forced to balance an irregular antenna design with these other constraints, this might seem like a bad thing; however, from the point of view of this research, the irregularity is what allows us to deduce information about which way the object is pointing.

The important concept to glean from the previous discussion is symmetry. When antenna design compromises are made, the resulting antenna design would go from speci-

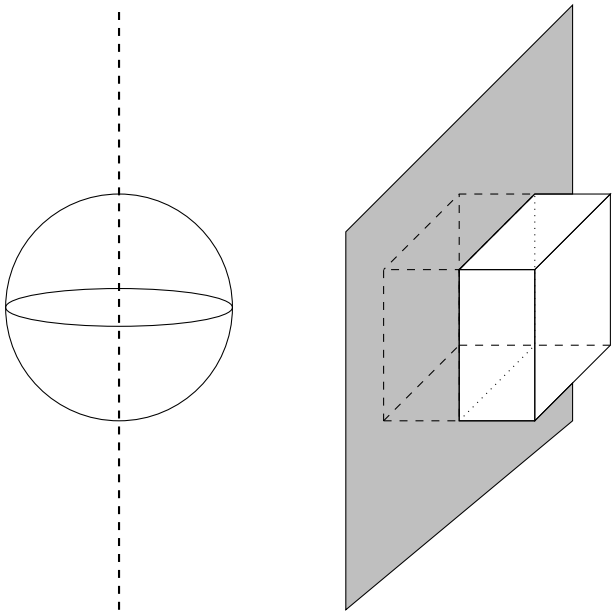

Figure 1.2: This figure demonstrates the two kinds of symmetry that are important to this research. Depending on the existence and type of symmetry of a wireless device's radiation pattern, it may be more or less difficult to determine its orientation. The sphere on the left has an infinite number of symmetric rotations. The cube on the right has only finitely many symmetric rotations.

fying an antenna with many symmetries (an infinite number in the isotropic case) to an antenna that may have only one or two symmetries, if any at all. In the simplest case of symmetry, consider a function of one variable, $f(x)$ (to represent energy radiated for given x). In this example, x is analogous to the angle of transmission in our antenna model, and f represents the signal strength corresponding to a particular x. To say that this function is symmetric is to also say that $f(x) = f(-x)$ for all x. Now, given that $f(x)$ is symmetric, consider blindly evaluating $f(x)$ at one of two particular points, x_0 or $-x_0$ (neither equal to 0). I.e. we do not know which of the two evaluations was made. If we also require that the choice of evaluation be left to a fair coin flip, then, upon learning the evaluation of the randomly chosen point, we are not able to deduce with any certainty which of the two points was picked.

Notice that the uncertainty described above will emerge not only when the antenna radiation pattern is symmetric but also, more generally when the radiation pattern $f(x)$ is not unique-valued. In practice, this non-unique-valuedness makes the estimation system susceptible to uncertainty due to the presence of noise.

Despite symmetry being detrimental to the certainty of an estimation system, it is still important to make distinctions between two different types of symmetry that dictate the available options when placing subject antennas around the object device. The key to solving this problem lies in the experimenter being able to choose the number and positions of the subject antennas. The most optimal situation is when there is no symmetry. Another possible case is if the device is rotationally symmetric. In this context we mean that there exists an axis through which the object device can be rotated, and a stationary observer (subject antenna) will not observe a change in signal strength. This type of symmetry cannot be "broken" by placing an additional observer in a strategic position. The third case is when the object device has plane symmetry. When this occurs, uncertainty due to symmetry *can* be overcome with the placement of another sensor, giving the estimation system more information.

So, given that the estimation system has prior knowledge of the object device's radiation pattern, estimators can choose the optimal arrangement of subject antennas to minimize the amount of uncertainty based on the observed symmetries of the device. Now that the subject antennas have been positioned, the system is ready to run the estimation algorithm.

1.3.2 Maximum Likelihood Estimation

There are many different mathematical methods that could be used to deduce the orientation of the object device. However, we will make do with a simple, classical method.

Optimization problems are historically very difficult. It is often required that a function

be continuous, smooth, and convex before anything of value can be said about the existence or uniqueness of local and global maxima or minima. For this problem, we cannot guarantee any of these properties apply to the object device antenna. Instead, we will rely on the fact that characterizations of object devices can be discretized into a finite amount of points. The result is a matrix representing radiation strength as a function of orientation. In this form, the estimation system can resort to an exhaustive numerical search. The orientation that is most likely the orientation which produced the observed vector of signal strengths is chosen by,

$$\text{argmin } ||v - v_{\theta,\phi}||^2 = (\hat{\theta}, \hat{\phi}). \tag{1.1}$$

As mentioned before, the object device will have some number, n, of subject antennas taking simultaneous measurements from different angles around it. The signal strengths resulting from the measurement are thought of as an n-dimensional vector, v. Because the estimation system has knowledge of the device's radiation pattern, it can compare v with all possible discretized shifts by θ and ϕ of the n sensors, $v_{\theta,\phi}$. In this situation, comparison will be done by the L^2-norm, or Euclidean distance. Once the norm of the residual error for every entry of $v_{\theta,\phi}$ has been calculated, the smallest one is picked and the corresponding angles represent the device orientation at which this minimum is achieved.

The advantages of this method are it is exhaustive and relatively simple compared to other methods. In these preliminary experiments, it is important to be able to know with certainty that the entire system works, without doubting whether or not the computational complexity is the source of the problem. To that end, the maximum likelihood implementation used in this project is easy to understand, easy to verify, and will guarantee that the global minimum is found.

On the other hand, there are some downsides to using this method. The most obvious is that it necessitates a strong trade-off between computation speed and accuracy. In

theory, one could implement a feature to down-sample the radiation pattern to decrease the amount of time it takes to search through the entire $v_{\theta,\phi}$ matrix. This is not ideal. This method also doesn't use previous orientation predictions to narrow down the possible current predictions. For example, if the estimation system outputs a prediction $\hat{v}$; then, assuming the device's angular velocity has a small upper bound in between measurements, it is more likely that the next orientation will be close to its previous position, $\hat{v}$. An additional downside to this method is the uncertainty of whether or not small changes in orientation result in small changes in $||v - v_{\theta,\phi}||^2$, which stems from not being able to guarantee that the object device is smooth.

The remainder of this thesis will proceed as follows: Chapter 2 will detail the steps needed to accomplish the above-stated goals. Chapter 3 will organize the results and findings of the experiments conducted. Then finally, Chapter 4 will conclude with the analysis of the various methods and significance of the findings.

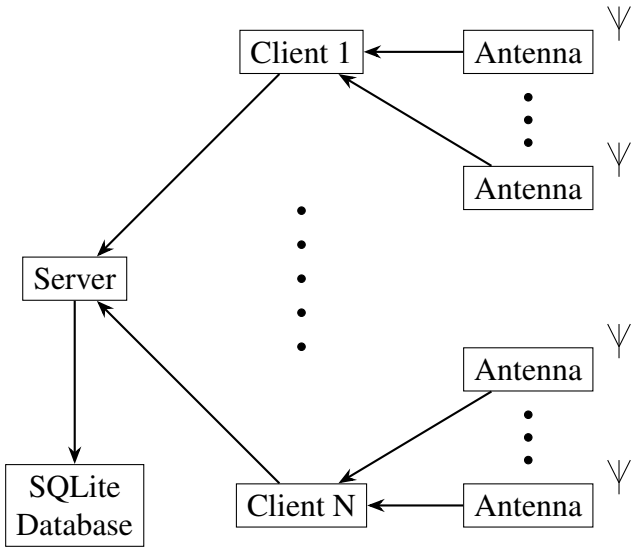

Figure 2.1: Block diagram showing the major components of the packet capture system. Generally consists of server and networked clients. Clients use connected Wi-Fi antennas to measure device, and send data back to server. Server puts data in database.

2. METHODS

2.1 Packet Capture

2.1.1 Packet Capture Array

The platform which mediates the collection of data from an object device is described in this section. This research is focused on leveraging existing Wi-Fi infrastructure; so, the portion of the system that measures signal strength does so using Wi-Fi protocol and equipment. Figure 2.1 depicts the major components of the measurement system. Most importantly, the system runs on the interconnection of a server with a number of clients. Each client has a number of omni-directional Wi-Fi antennas to make measurements with. Synchronized data is communicated back to the server and stored in a SQLite database.

The main roles that the server and clients play are to provide communication, synchro-

nization, and storage. They are implemented as a pair of C programs that can communicate through the use of POSIX network sockets over an arbitrary interface (wired or wireless). Once connected, the server enters a mode whereby it can issue commands to the clients. Each client has some number of Wi-Fi antennas *for measurement* and at least one ethernet or Wi-Fi interface for network communication. In general, when the user issues a "measure" command from the server, it sends a synchronized message to all connected clients to take a measurement or series of measurements. Synchronization is implemented by the POSIX multi-threading library, `pthread`. Using multiple threads and queues, the server and client are able to balance the individual tasks that need to be completed for a successful measurement.

When the time comes for a particular client to make a measurement, many low-level aspects must be taken into consideration. First, before it is even possible to use the connected antennas as a measurement device, they must first be switched and be capable of switching into "conspicuous mode." This mode is different from their normal mode of operation in that they are now able to receive packets not originally intended for their MAC address. This is not entirely necessary, but simplifies the procedure a bit. You could conceive of a system where the transmitting device knows each and every sensor that is measuring it, and sends data to each one. However, this method would be rigid and inflexible and not support the dynamic exchange of sensors. So, to get around this, conspicuous mode is used to eavesdrop on the messages that are broadcast by the target device. The actual interpretation of messages is done within the client program using the `pcap` library, or packet capture library. With this tool, the Wi-Fi frames of the intercepted packets can be analyzed. When intercepted by the conspicuous antenna, a signal strength field, among others, is populated. Then, the data is extracted by the client, compiled with other information about system-wide antenna identification, and sent back to the server over the appropriate socket.

When the measurement data finally arrives at the server, it is forwarded to a database. For each window of measurements, the server will store data from all clients in one self-contained database file. The main difficulty with this method is it requires careful naming of all clients and antennas. It is crucial that, in the final database file, each antenna is unambiguously attributed to the measurements that it made.

At this point, the data is ready to be analyzed by a prepared script or matrix computation package like MATLAB. The data will either need to be carefully pieced into a full characterization of the target device, or the data is fed, in real-time, to the estimation software running on the server. Although this project did not see the full implementation of a real-time estimation system, such a system would only require a few extra steps.

2.2 Anechoic Chamber

The anechoic chamber is a room specialized for taking noise-free electromagnetic measurements. The chamber available at Texas A&M University is able to block external electromagnetic signals. Perhaps more importantly, it also has a special material (Figure 2.2) coating the interior to prevent internal reflection of signals. The attenuation of outside noise and reflected signals makes this chamber a good approximation of free space and an ideal location to characterize an antenna. If the characterization of an antenna is not performed in an environment like this, the data will be skewed by multi-path interference and interference from nearby devices.

The diagram in Figure 2.2 illustrates the type of insulation used to protect the inside of the chamber from noise and multi-path interference. The shape and material of the padding tends to cause any incident electromagnetic waves to get lost in the crevices; it promotes maximal energy absorption so that less energy is available for reflection.

The properties of this room make it ideal for narrowing down experimental uncertainties. It is also valuable for testing estimation in a noise free environment. Isolation allows

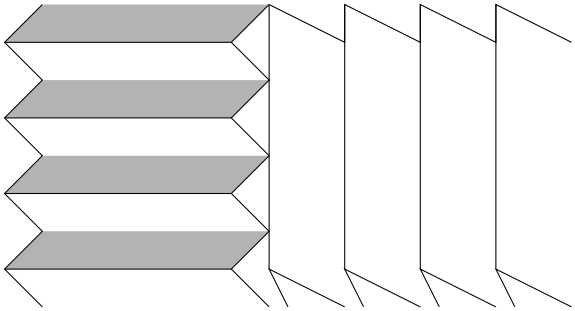

Figure 2.2: Simplified illustration of electromagnetic shielding used in anechoic chamber. Unique shape maximizes absorption of incident radiation to prevent reflection.

estimation algorithm to ensure that results are as close to theory as possible. When noise and multi-path interference are involved, more complicated models must employed.

To facilitate consistent rotation of the object device, the anechoic chamber is also equipped with a motorized positioning system. The system can be controlled externally, and was configured to allow two degrees of freedom for rotating the phone: azimuthal angle and elevation. The system allows precise control over the orientation of the object device. By rotating the device through 360° of azimuthal angle and 180° of elevation at each azimuth, a full characterization was obtained.

2.3 Characterization

2.3.1 HTC Smartphone

The HTC smartphone device was chosen out of convenience. Not only this, but it satisfies the original goal of wanting to perform tests on an actual, real antenna. The entire goal of this research is focused on determining the orientation of real objects, using only the signal strength that they give off. Another benefit to using a smartphone is that it is very feasible to design applications for. For instance, characterization required that the device continually be broadcasting packets so that the sensors have something to sense. This issue was handled by a small application that did just that. Though this particular

characterization used a smartphone, Wi-Fi protocol, and the 2.4 GHz spectrum, these same ideas could be applied to other similar technologies.

To characterize the object device's antenna using the packet capture system, we used a compilation of six independent characterizations. Given that the equipment we were working with was not precise, it makes sense to take many different measurements and average them together. This gives us confidence that the average of many sample radiation patterns is close to the actual radiation pattern.

For the characterization procedure, the smartphone was placed on the motorized apparatus (described in previous section) in one end of the anechoic chamber. In three of the corners, six sensors were placed, connected to omni-directional antennas. Figure 2.3 shows the exact positioning of the six sensors relative to the smartphone in the middle of the anechoic chamber. As a result, sensors zero through three measured the smartphone from approximately 6.5 m, and sensors four and five measured from approximately 1.6 m away. The geometry of the placement means that compared to the (0,0) reference shown in the figure, sensors zero and one had a viewing position of (13°,0°), sensors two and three had a viewing position of (347°,0°), and sensors four and five had a viewing position of (131°,0°).

Once all sensors were in place, the characterization could begin. Azimuth was varied through twenty-four equally spaced angles ranging from 0° to 360°, corresponding to a full rotation. Elevation was varied through thirteen, equally-spaced angles ranging from 0° to 180°. Notice that it is only necessary to vary elevation by a half-rotation to get the full characterization. If elevation is ranged from 0° to 360°, then a reflection of the entire characterization is measured as well.

Figure 2.4 shows an illustration of how the varying θ and ϕ angles change the position of the object device. You can think of the device as having a handle extending through the (0,0) reference, and adjustment of the angles corresponds with moving the handle to

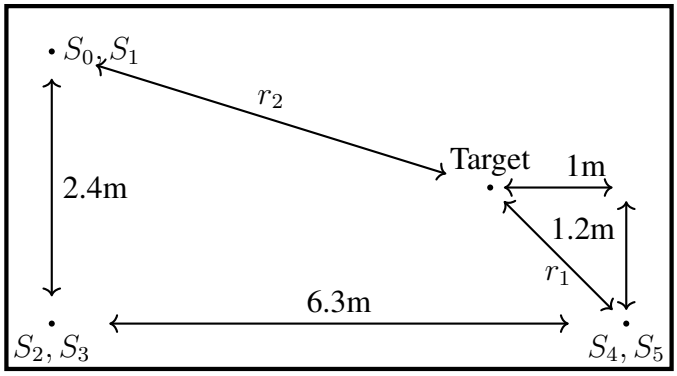

Figure 2.3: Schematic layout of the placement of the device and sensors during radiation pattern characterization. Notice that not all sensors are positioned at the same distance from the target, and that sensors view the target from many angles.

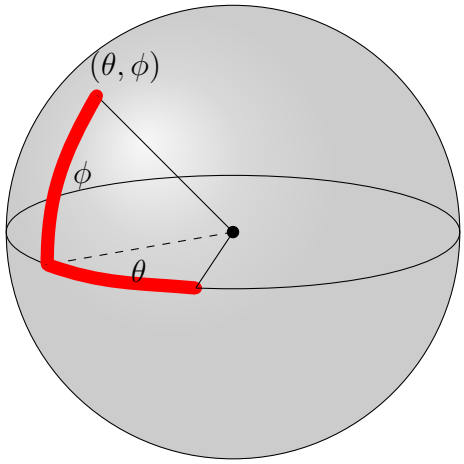

Figure 2.4: Diagram that illustrates how azimuthal angle and elevation are mapped to specific points on the unit sphere. Azimuth ranges from 0° to 360° and elevation ranges from 0° to 180°.

adjust the orientation of the device.

2.4 Simulation

The purpose of the simulation software is twofold: to speed up the engineering design process by lowering the amount of time successive design iterations take and to quickly

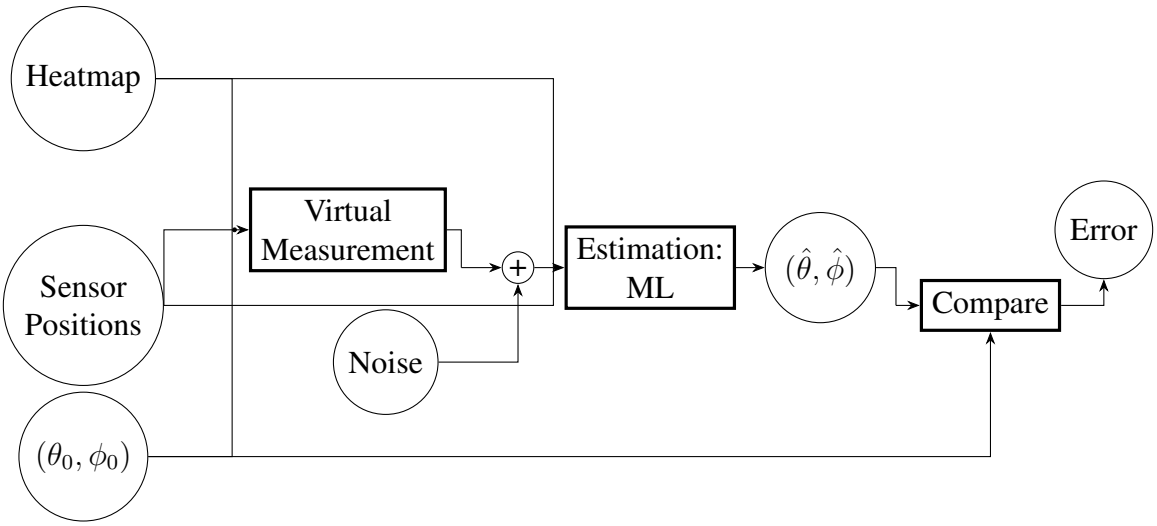

Figure 2.5: This block diagram depicts simulation software, from left to right. Inputs supply scenario parameters to the virtual measurement block. Noise is added before estimation block. Resulting predicted orientation is compared with actual orientation for performance prediction.

learn the characteristics of certain individual antenna patterns. We do not know ahead of time which algorithmic technique is best to estimate the orientation of our device, it is valuable to be able to test a series of proposed methods on our test data and compare their performance. This also saves us the time and hassle of having to make many real-life measurements, when simulated measurements will work just fine. Second, the relationship between the characteristic pattern of the given antenna and the number of antennas and arrangement of those antennas strongly affect the performance of the estimator. By virtually adjusting the positions and numbers of sensing antennas around the device, we can learn which orientation allows the most discerning measurements to be made. Then, in a real-life situation the sensors will be placed in the optimal positions for the estimation of the device.

The simulation system consists of three components implemented in Python: the char-

acterization or radiation pattern, the estimator, and the measurement system. See Figure 2.5 for a block diagram describing the simulator. The characterization can be obtained in two ways; either by physically measuring an antenna's radiation pattern or by generating a radiation pattern artificially. This data serves as the input to both the measurement and estimation blocks. The antenna characterizations are described by a matrix of values with corresponding angles of measurement, θ and ϕ, the two angles that describe the orientation of the radiation pattern.

For this research, the estimation was done with two degrees of freedom; but, the algorithm works equivalently with arbitrary dimension (must consider performance, though). For example, only two parameters were used to describe the orientation of the object device; but, to be complete, one should assign a triplet of numbers describing orientation (something like roll, pitch, and yaw).

2.4.1 Virtual Measurement

One of the advantages of using a simulation is that it allows the user to create any setup that they want. Those setup preferences are the main inputs to the first block of the simulator. In order to fully specify the scenario, the virtual measurement block needs to know the radiation pattern of the object device, the number and locations of sensors, and orientation of the device. With this information, the virtual measurement block will show what those sensors would have measured if they were actually setup and measuring. This is possible because of the relativity of reference frames between object and sensor.

That is, the same action is interpreted differently by each viewpoint. Consider a sensor at the (0,0) reference point, measuring the object device. The following actions are equivalent:

- Leaving the sensor where it is, and measuring the device at an orientation of (θ_i,ϕ_i).
- Moving sensor to ($-\theta_i$,$-\phi_i$), and measuring the device at orientation (0,0).

For a sensor virtually placed at (θ_i,ϕ_i), the measurement block simply has to lookup the value of the radiation pattern at ($-\theta_i$,$-\phi_i$). This is repeated until all sensors have a virtual measurement, resulting in a vector of signal strengths.

The last input to the block specifies the orientation of the device when a measurement occurs. For the purposes of predicting estimation performance of a particular sensor arrangement, the angle given to the device should be random. Then, after many trials are done, the user will have a reasonable idea of the effectiveness of that arrangement.

To determine the robustness of a particular sensor arrangement, it is useful to add noise to the resulting vector. This allows the experimenter to see how performance degrades with increasingly suboptimal levels of noise (with no noise being the ideal amount). This simulator used a random Gaussian number generator to create a vector of noise. Noise was then added to the output of the measurement block before passing to the estimation block.

2.4.2 Estimation

The most important block of the simulator is the estimation block. As mentioned earlier, we focus only on the maximum likelihood algorithm for estimation. The basic idea behind this method is that the estimator has perfect knowledge of the radiation pattern of the object device. Therefore, it can compare some observed vector of signal strengths, v, with *all possible* orientations that could have resulted in that measurement, $v_{\theta,\phi}$. In this context, "compare" refers to computing the squared Euclidean distance between the two vectors (see Equation 2.1). Assuming that the noise component is zero-mean, the most likely guess for orientation will be the pair (θ_i, ϕ_i) that minimizes that Euclidean distance:

$$\text{argmin } ||v - v_{\theta,\phi}||^2 = (\hat{\theta}, \hat{\phi}) \tag{2.1}$$

This method is ideal for our situation because we can supply the estimator with perfect

knowledge of the object radiation pattern. However, if this were not the case, this method would not be possible; guesses would be based purely on real-time measurements with nothing to compare them to. Another downside to this method is that it is quite slow. An exhaustive search will guarantee that the most likely orientation is found but, depending on how detailed the target radiation pattern is, this search could take a long time. Pseudo-code for the estimation procedure is displayed in Algorithm 1.

Algorithm 1 Algorithmic description of Maximum Likelihood estimator. v is the measured n-dimensional signal strength vector corresponding to the n sensors; p is a list of the positions of the n sensors; R is the radiation pattern of the device. $vTemp$ represents the hypothetical measurement when device is at orientation $i, j, \cdots, p$.

procedure ESTIMATEANGLEML(v, p, R)
  $numSensors \leftarrow length(p)$
  $squaredError \leftarrow initialize(dim(R))$
  for $i\ in\ first\ dimension,\ d_i,\ of\ R$ **do**
    for $j\ in\ second\ dimension,\ d_j,\ of\ R$ **do**
      $\vdots$
      for $p\ in\ last\ dimension,\ d_p,\ of\ R$ **do**
        $vTemp \leftarrow initialize(numSensors)$
        for $n\ from\ 1\ through\ numSensors$ **do**
          $vTemp[n] \leftarrow R[\ (i, j, \cdots, p) + p[n])\ ]$
        end for
        $squaredError[i, j, \cdots, p] \leftarrow elementwiseSum(\ (v - vTemp)^2\ )$
      end for
      $\vdots$
    end for
  end for
  $vHat \leftarrow argmin(squaredError)\ over\ i, j, \cdots, p$
  $return\ vHat$
end procedure

2.4.3 Error

By the end of the block diagram, the estimation block has provided its output orientation guess. In a practical scenario where the orientation is actually unknown, this would be the end result. But, in these preliminary stages, it is necessary to describe performance in a quantifiable way. Error quantification is done with the help of Figure 2.4. This figure describes the mapping between the two adjustable parameters, θ and ϕ, and the points on a unit sphere. It is a mapping between a two-dimensional space, and a three-dimensional space. By expanding the coordinate system in this way, we can create an intuitive measure of distance.

It is well-known that the angle between two vectors, ψ, can be expressed as the product of the lengths of each vector and the cosine of the angle between them:

$$\begin{aligned} \langle x, y \rangle &= ||x||\,||y|| \cos\psi \\ \implies \psi &= \arccos \frac{\langle x, y \rangle}{||x||\,||y||}. \end{aligned}$$

But, since our entire set of points lies on the unit circle,

$$\psi = \arccos \langle x, y \rangle.$$

Therefore, since the angle between two vectors is exactly what we want to define as our distance, the distance between two points on our unit sphere is:

$$\begin{aligned} D(P_1 = (\theta_1, \phi_1), P_2 = (\theta_2, \phi_2)) = \arccos(&\cos\theta_1 \cos\phi_1 \cos\theta_2 \cos\phi_2 \\ &+ \sin\theta_1 \cos\phi_1 \sin\theta_2 \cos\phi_2 \cdot \\ &+ \sin\phi_1 \sin\phi_2) \end{aligned} \tag{2.2}$$

Notice that this result implies that the maximum attainable distance would be π (or 180°).

This is consistent with our formulation, and is an intuitive, useful measure of distance for this project.

3. RESULTS

3.1 Device Characterization

The plots in Figure 3.1 show the HTC smartphone characterization data recorded from the six sensors placed in the anechoic chamber. Refer to the setup description in Figure 1.1 for a detailed account of the positioning of the sensors. The smartphone was rotated with two degrees of freedom, so the resulting datasets are two-dimensional matrices. Each plot corresponds to a different sensor. Along the horizantal axis, elevation angle is measured (ranging from 0° to 180°), and the vertical axis measures azimuthal angle (ranging from 0° to 360°).

The characterizations in Figure 3.1 shows a clear pattern. The sensor placement corresponds exactly to the amount of shift needed to align the views with a reference view at (0°,0°). This is a strong indication that the data collected by each one of these six antennas is consistent. The reasoning behind this is geometrical: Assuming that the measurements in the anechoic chamber are indeed as if they were in free space, then we know that the rotating device has a *constant* radiation pattern. This implies that if the device were spun on a single axis (like a globe) forever, a continuous series of measurements from a subject antenna would measure a periodic signal. Similarly, this also implies that subject antennas separated by a constant angular shift will experience a delayed view of the radiation pattern as the device rotates. The concept is very similar to how the Dirac delta function behaves for the convolutional sifting property, i.e., a function, $f(x)$, convolved with a shifted impulse, $\delta(x - x_0)$,

$$f(x) \star \delta(x - x_0) = f(x - x_0), \tag{3.1}$$

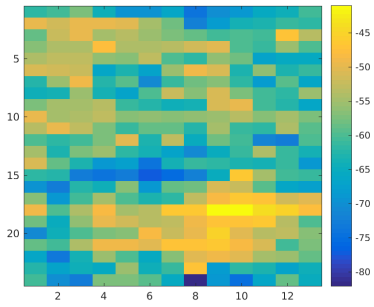

(a) Sample of HTC smartphone from antenna S_1 placed at approximately (13°,0°). Compare with subfigure b).

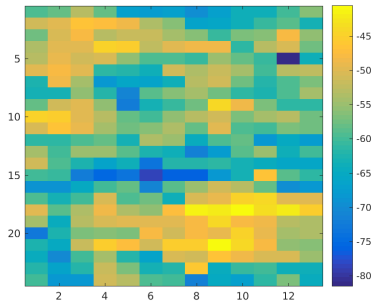

(b) Sample of HTC smartphone from antenna S_0 placed at approximately (13°,0°). Compare with subfigure a).

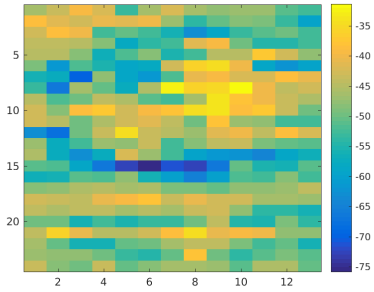

(c) Sample of HTC smartphone from antenna S_4 placed at approximately (229°,0°). Compare with subfigure e).

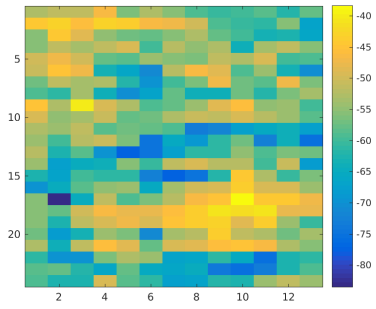

(d) Sample of HTC smartphone from antenna S_2 placed at approximately (347°,0°). Compare with subfigure f).

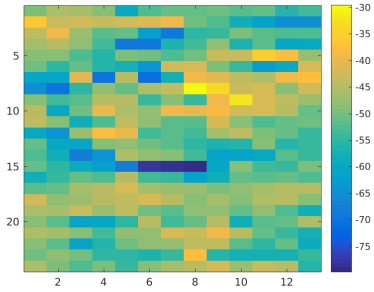

(e) Sample of HTC smartphone from antenna S_5 placed at approximately (229°,0°). Compare with subfigure c).

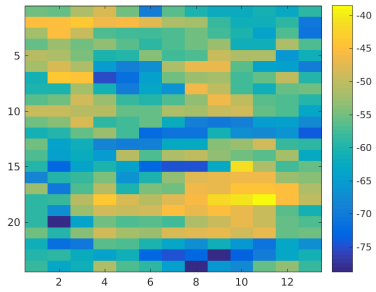

(f) Sample of HTC smartphone from antenna S_3 placed at approximately (347°,0°). Compare with subfigure d).

Figure 3.1: Characterization of HTC smartphone from point of view of six differnet antennas. 13 varying elevations (ranging from 0° to 180°) for each azimuthal angle. 24 azimuthal angles (ranging from 0° to 360°) for each elevation angle. dB values closer to zero represent stronger signal strength.

equals the a shifted $f(x)$ by x_0. We can think of the subject antennas as sampling the object device's radiation pattern at certain shifted points corresponding to their placement around the device.

Notice in Figure 1.1 that two antennas were placed at one radius, r_1, relatively close to the object device, and the other four antennas were placed at a more distant, r_2. Though we know that the relative positioning *angle* has an effect on the shifting of the views; the relative *distances* also play a role. When an electromagnetic wave passes through a medium other than a vacuum, it experiences attenuation. This attenuation can be modeled in the following way:

$$\Phi(x) = Ae^{-\alpha x} \tag{3.2}$$

where $\Phi(x)$ is the wave intensity as a function of distance.Since the antennas on the sensors automatically convert signal strength to a logarithmic dB value, we can compare two distant measurements,

$$\begin{aligned}
\Delta\Phi_{dB} &= 20\log_{10}\Phi(r_1) - 20\log_{10}\Phi(r_2) \\
&= 20\log_{10} Ae^{-\alpha r_1} - 20\log_{10} Ae^{-\alpha r_2} \\
&= 20\log_{10} A - \alpha r_1 20\log_{10}\frac{1}{e} - 20\log_{10} A + \alpha r_2 20\log_{10}\frac{1}{e} \\
&= 20\log_{10} e\alpha(r_1 - r_2)
\end{aligned} \tag{3.3}$$

which is a constant for two stationary sensors. Thus, logarithmic measurement in dB has converted a multiplicative factor into a constant additive difference. For our six antennas, this means that we simply need to pick either r_1 or r_2 as a reference distance, and then add a constant to the views that are from different distances so that their views have the same average value. This modification will make it seem like all sensors were at the same distance away from the object device.

Now that we know that the six sample characterizations are simply shifted versions of

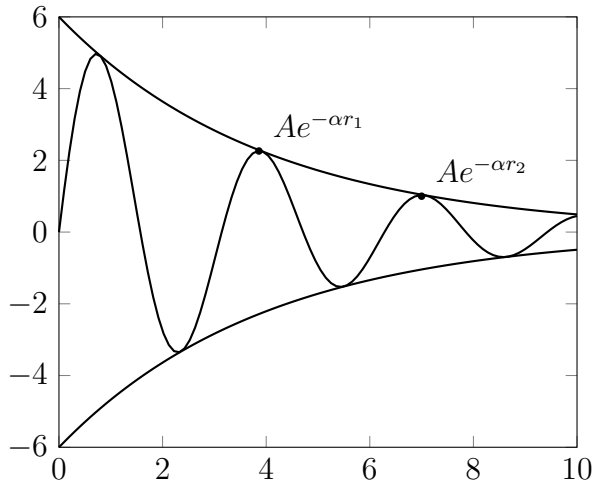

Figure 3.2: Plot to illustrate the effect of free-space path loss on measured signal strength. When signal intensity is measured in dB, exponential maps to linear function. Using dB, Δr is proportional to $\Delta\Phi_{dB}$.

one another, and have been adjusted so that all measure from the same distance, it makes sense to average them together into one mean characterization. First, there is an important step in between alignment and averaging: the placement of each sensor during the characterization was not according to any grid of angular shifts, so the sample characterizations will not align perfectly. To fix this, we interpolate each sample characterization so that all six can be aligned. In this case, a bi-cubic interpolation (implemented by MATLAB, using cubic splines) was used to increase the resolutions of the views. Now, the views can be properly aligned and averaged for the final, interpolated mean characterization (shown in Figure 3.3).

The most obvious takeaway from the resulting characteristic is that it is irregular and asymmetric. On one hand, as discussed in the Theory section of Chapter 1, the asymmetry assures us that there will be no uncertainty between two orientations that symmetries of each other. On the other hand, the radiation pattern is not monotone, uniform, or smooth. When noise is introduced, this will cause problems for the estimation system.

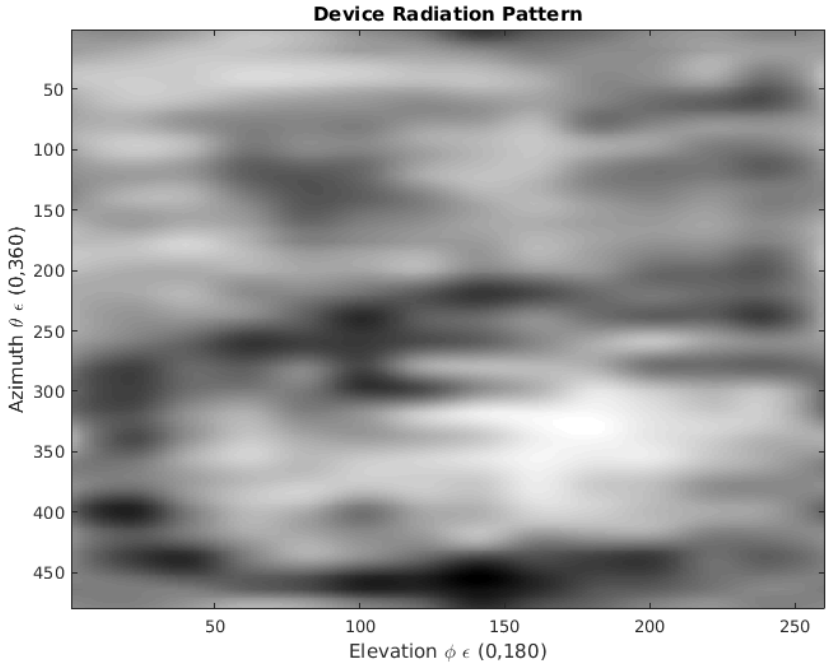

Figure 3.3: Antenna characterization of HTC smartphone after interpolation, distance correction, alignment, and averaging of six sample characterizations.

3.2 Simulation

In this section, the results of simulation with various device radiation patterns will be displayed and discussed.

3.2.1 Ideal Directional Test Antenna

To verify the theoretical discussion about symmetry, we've used an ideal directional antenna (generated by MATLAB's `mvnpdf()` function). The shape is simply a bell-shape with a global maximum along one direction and a gradual slope off to the global minimum at antipode. So, because the maxima and minima of the radiation pattern lie on the same axis, and the transition between the two points is the same along all directions, the radiation pattern a few planes of symmetry. Because of the way that the radiation pattern

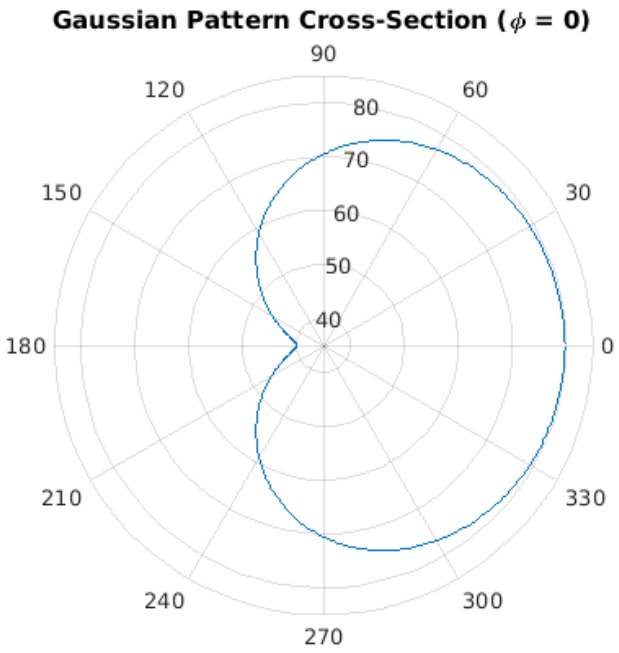

Figure 3.4: Cross-section of virtual antenna provides idealized directional antenna. Allows simulation with radiation pattern that has a single line of symmetry. Strong signal peaks at -40 dB and weak signal peaks at -80 dB

is mapped to a spherical device, the radiation pattern has two planes of symmetry (each at 90° from each other, their intersection coinciding with the min/max pole). A cross-section is used to verify theoretical predictions for patterns with plane symmetry, and is shown in Figure 3.4

The plots in Figure 3.5 show four different subject antenna placement scenarios when using the cross-section of the idealized directional antenna. Figure 3.5 a) clearly illustrates the need for multiple "viewing angles." This scenario represents the most trivial setup that can be used. The only situation where using one sensor would be feasible would be if you only had one degree of freedom, and if the radiation pattern were monotone. This doesn't make sense for a circular radiation pattern, though; because, monotonicity on a periodic function is a contradiction. There would necessarily have to be jump discontinuities in

the radiation pattern. Figure 3.5 b) shows the estimation accuracy when two sensors are placed on the axis where the two extrema are located. Again, as expected, the estimation is poor even at low noise levels. Theoretically, we know this is because there is ambiguity between points on the "equator" (i.e. point midway between extrema). The low-noise error seems to converge to approximately 45°. Notice that the estimator is essentially guessing between two points that are between 0° apart and 90° apart; so, the average should be around 45°, as measured. Monotonicity on a periodic function is a contradiction. There would necessarily have to be jump discontinuities in the radiation pattern. Figure 3.5 c) is the first of the group to show significant estimation accuracy around the 0 to -1 dB range. This arrangement of sensors is similar to that of b), however, the sensors are placed so that they break the symmetry of the radiation pattern. Two sensors placed 90° apart will almost always be able to deduce the orientation, in sufficiently low noise conditions. Because of their placement, it is exceedingly unlikely that each sensor will measure similar signal strength values. Monotonicity on a periodic function is a contradiction. There would necessarily have to be jump discontinuities in the radiation pattern. Figure 3.5 d) is included to show that further increasing the number of sensors only increases the 0 dB accuracy.

3.2.2 Smartphone Antenna

From Figure 3.3 we can see that the smartphone characterization has many intricate features. While it is not symmetric, the irregularity tends to make the estimation difficult. Despite these setbacks, we should still expect to see the 0 dB estimation improve as the number of sensors increases. In general, we'd expect that each added subject antenna gives more information (since this radiation pattern doesn't have any symmetry).

Figure 3.7 shows the progression of estimation performance as the number of equally-spaced sensors increases from one, to two, to four, to eight. Figure 3.7 a) starts off with one

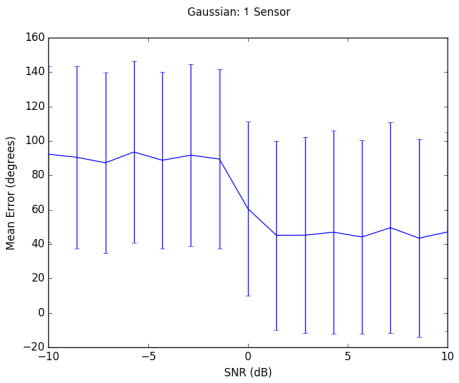

(a) Estimation is expectedly poor even at low noise levels due to symmetry issues (only one sensing antenna).

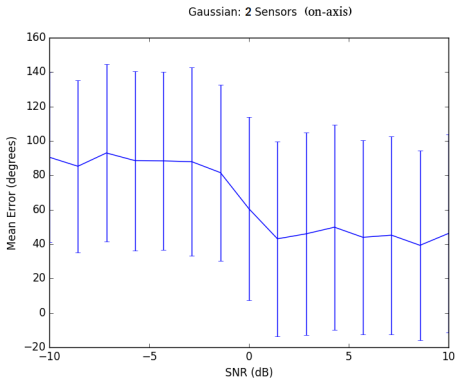

(b) This arrangement of two antennas on the axis of symmetry also expectedly gives poor estimation results. Ambiguous views result from combination of radiation pattern symmetry and sensor placement.

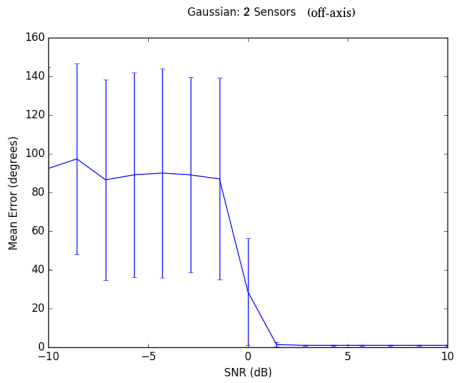

(c) Estimation improvement over two symmetrically placed antennas. Symmetry is broken.

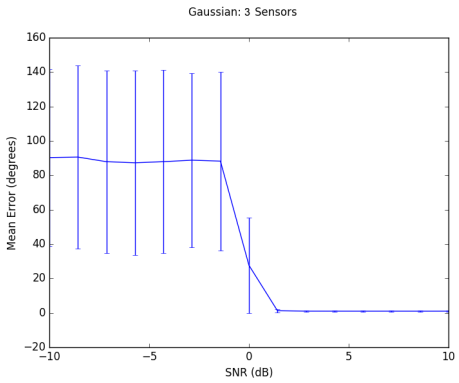

(d) As number of sensors increases, the 0-dB error drops. Performance projected to be even better than two off-axis sensors, as expected.

Figure 3.5: Limited simulation of idealized directional antenna. Directional cross-section, $\phi = 0$.

sensor, predicting poor accuracy with wide variability. Again, this is expected from such a small amount of information. Figure 3.7 b) doubles the number of sensors. Performance is still quite bad, but there is slight correlation with decreasing noise, which is typical of successful estimation scenarios. Figure 3.7 c) doubles the sensors again to four. Now there is a clear jump in accuracy as noise level becomes comparable to radiation pattern

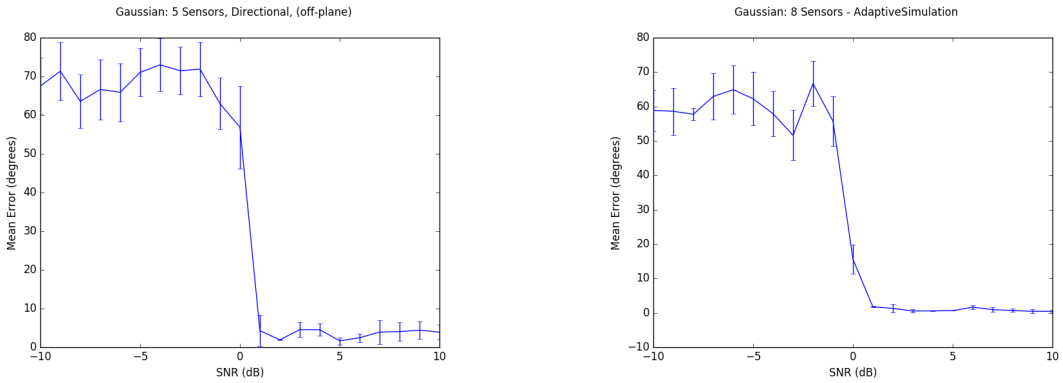

(a) Caption. (b) Caption.

Figure 3.6: Full simulation of idealized directional antenna.

at the 0 dB mark. On the right side of the 0 dB mark, we see performance improve as noise continues to drop. Figure 3.7 d) predicts performance with eight sensors. In this simulation, there is a clear threshold between poor and accurate performance. This plot strongly resembles the constrained simulation of the idealized directional antenna, indicating that 8 sensors is sufficient for low-level orientation estimation.

Up until now, the main guiding principle for positioning sensors around the objective device has been to avoid placing them on axes or planes of symmetry. But, in the case of the measured smartphone antenna, there is no symmetry and still a lot of irregularity. In this situation, there is not really any obvious reasoning for positioning of sensors. However, the plots in Figure 3.9 demonstrate the implications of a certain idea: if the estimation problem becomes hard when multiple sensors tend to record similar signal strengths, then it might seem reasonable to position the sensors so that they could potentially measure a large number of extreme points at the same time. This would ensure that the values measured by the sensors are as different as possible. A graphical representation of this is shown in Figure 3.8. By placing antennas so that they align with the high and low points of the radiation pattern, the hope is that it will be very unlikely for the combined knowledge

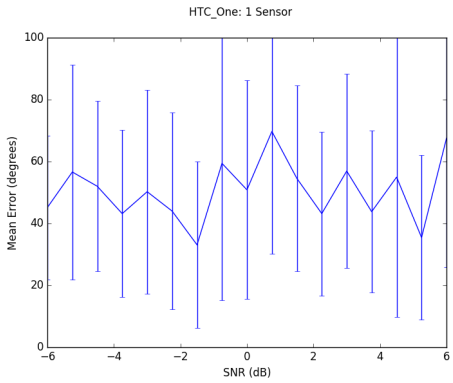

(a) Estimation simulation with measured smartphone antenna, using one sensor.

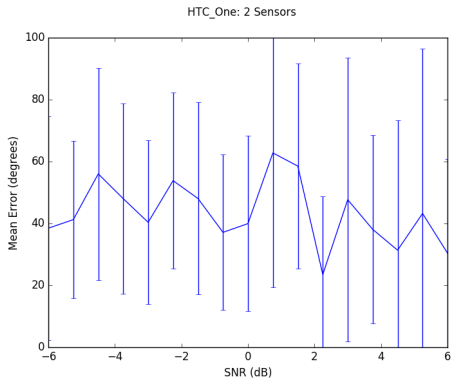

(b) Estimation using two sensors (equally spaced). Error tends to trend downward as noise decreases. Estimation still has a lot of variability.

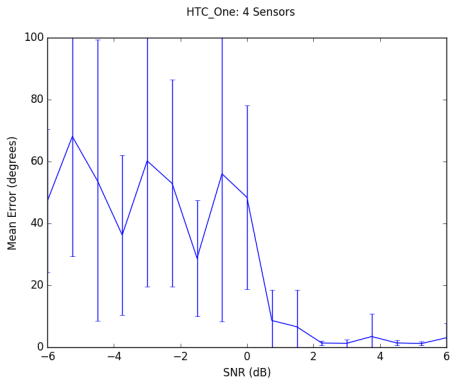

(c) Estimation using four sensors (placed on the vertices of a tetrahedron). Shape of plot looks more like the desired thresholded step.

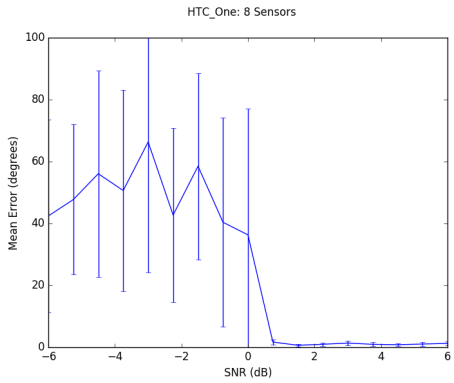

(d) Estimation using eight sensors (placed on the vertices of an octahedron). Notice the characteristic drop-off.

Figure 3.7: Simulations using measured smartphone antenna. Notice trend of increasing low-noise accuracy as the number of antennas increases.

to result in ambiguity. The results of this type of simulation with four and eight sensors is shown in Figure 3.9.

As the plots show, the peak-trough placement of the sensors performed much worse than the corresponding sensor arrangements using the earlier placement methods. Perhaps there is a flaw in the earlier reasoning: in the peak-trough placement, the sensors are arranged so that they would line up with the extreme points of the radiation pattern.

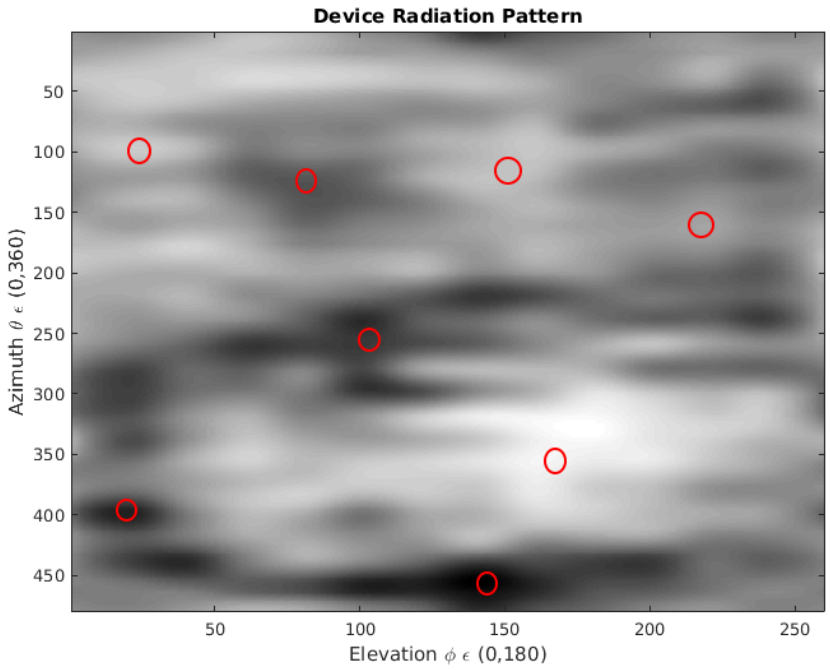

Figure 3.8: Antenna characterization of the smartphone with select extreme points highlighted. These points are the interest of the peak-trough sensor placement discussion.

However, the fatal assumption is that if you consider the orientation of the device to be random (as they are in the simulations), then it is exceedingly unlikely for the extreme points to *actually* align with the corresponding sensors. That is, each individual sensor is most likely going to measure a value lower than it's corresponding max point, or higher than its corresponding min point. This behavior tends to cause the sensors to read similar values of signal strength, except during the rare case that the device is in the exact right orientation. Notice how this is the exact opposite of what we intended.

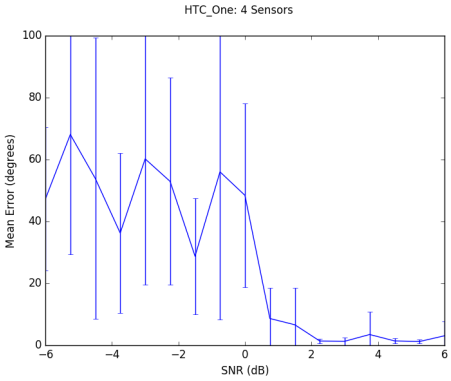

(a) Four sensors equally-spaced using earlier method. Same results as earlier; reasonable drop-off when noise falls below 0-dB mark.

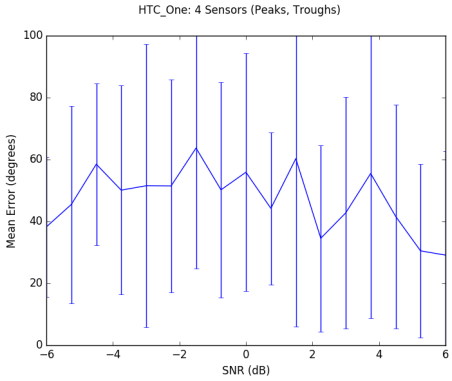

(b) Peak-Trough placement of four sensors. This arrangement performs much worse than it's corresponding counterpart. There is no semblance of any increase in accuracy as noise decreases.

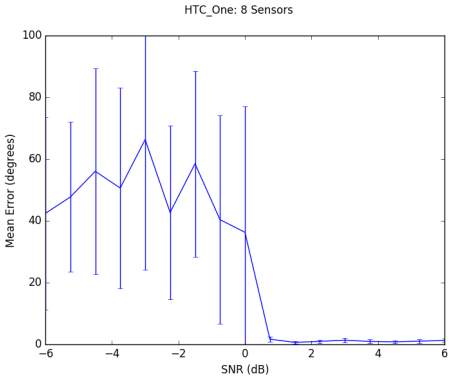

(c) Eight sensors equally-spaced using earlier method. Strong characteristic drop-off.

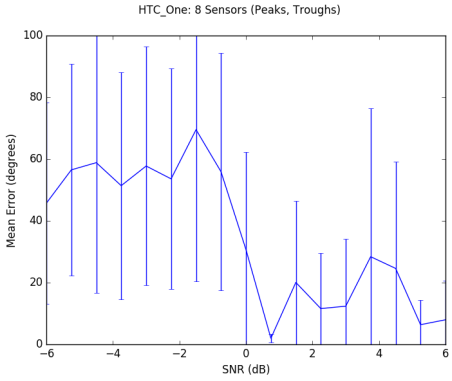

(d) Peak-Trough placement of eight sensors. Low-noise regions see small drop-off in error, but variability still much higher than corresponding counterpart.

Figure 3.9: Points in smartphone characterization that correspond with maxima and minima. They are used to test the idea that sensor placement at these points results in better estimation accuracy.

4. CONCLUSION

The goal of this research is to provide theoretical and simulation support for a technique of wireless orientation estimation that uses externally gathered signal strength. The model assumes that the target device is at a single location in a noise-free environment that does not suffer from multi-path interference. The number and positions of the external sensors are chosen to optimally take advantage of the features in the target's radiation pattern.

Antennas like the one possessed by the smartphone used here do not radiate isotropically. With exact knowledge of a target's radiation pattern, an external array of sensors can, under certain circumstances, extract enough information about the device's orientation for the estimation system to make a reliable guess. The optimal placement of these sensors is dependent on the symmetry and complexity of the target's radiation pattern.

The tests and simulations performed here are consistent with the theory. Simulations with an irregular smartphone antenna demonstrate the dependence of accuracy on number of sensors, in the absence of symmetry. As more sensors are placed around the target, the estimation system is able to deduce orientation with more certainty. Simulations of antennas with symmetry predict uncertainty that may be preventable, depending on the type: if the radiation pattern has an axis of rotational symmetry, then the uncertainty cannot be helped; if the radiation pattern has plane symmetries, then both number and position of sensors become critical predictors of success.

The amount of noise also strongly affects performance. When a sensor arrangement is properly positioned to break any radiation pattern symmetries, the error-noise plot has a characteristic shape. Noise above a certain threshold results in widely unreliable estimation, noise below that threshold results in near-perfect estimation, and the boundary

transition between the two results is very sharp. When the sensor arrangement is not optimal, the estimation performance is mostly constant across all noise levels.

Compared to localization, orientation is a relatively small field and has yet to be widely adopted in practice. Some potential applications were mentioned in Chapter 1, but new technologies are not always predictable. Like anything else, it has the potential to be a valuable tool and its uses are only limited by the user.

There are still many questions left to be answered, though: Does this technique work with many targets at once? Are there more effective mathematical methods? Can this be combined with localization? Are other frequencies of wireless communication able to benefit from this?

Experiments and simulations done here imply that orientation estimation be done in highly controlled situations. Further research is necessary to confirm whether the techniques are still viable in more realistic scenarios. This means researching the effects of multi-path interference, external noise, etc. on the performance of the algorithms. Part of why Wi-Fi is such a good exploratory tool for this is because the existing protocol already assigns unique identifiers to each connected device. Combined with this protocol, the strategies here should theoretically be capable of discerning orientation of multiple targets.

If it is indeed feasible to estimate the orientation of multiple targets simultaneously, then it will be necessary to increase the number of known target antenna characterizations.

In the case that the simple maximum likelihood algorithm is not sophisticated enough to handle the issues that a real-life application demands, research into more powerful statistical and mathematical tools will be necessary. Perhaps the applications require a an algorithm that is faster, and more efficient than maximum likelihood's exhaustive search. Or, maybe practical scenarios cause the method to become less accurate. An algorithm that factors past orientation states into prediction, rather than a simple instance-by-instance es-

timation, could increase relative performance.

It is also worth testing whether this technique is compatible with existing localization techniques. The development here assumes a stationary target, but localization information may be able to provide enough information to get around that issue. A combined localization and orientation that uses only information gathered externally surely has potential.

Lastly, other frequencies of wireless communication, not just Wi-Fi, should be able to use these ideas to their advantage. Of course, higher frequencies that experience greater path loss might not be suitable, but the range of applicable frequencies remains to be determined.

REFERENCES

[1] A. M. Ladd, K. E. Bekris, A. Rudys, L. E. Kavraki, and D. S. Wallach, "Robotics-based location sensing using wireless ethernet," *Wirel. Netw.*, vol. 11, pp. 189–204, Jan. 2005.

[2] P. Castro, P. Chiu, T. Kremenek, and R. Muntz, *A Probabilistic Room Location Service for Wireless Networked Environments*, pp. 18–34. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001.

[3] A. Taheri, A. Singh, and A. Emmanuel, "Location fingerprinting on infrastructure 802.11 wireless local area networks (wlans) using locus," in *Local Computer Networks, 2004. 29th Annual IEEE International Conference on*, pp. 676–683, Nov 2004.

[4] R. P and M. L. Sichitiu, "Angle of arrival localization for wireless sensor networks," in *2006 3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks*, vol. 1, pp. 374–382, Sept 2006.

[5] C. Rohrig and F. Kunemund, "Estimation of position and orientation of mobile systems in a wireless lan," in *Decision and Control, 2007 46th IEEE Conference on*, pp. 4932–4937, Dec 2007.

[6] F. Li, C. Zhao, G. Ding, J. Gong, C. Liu, and F. Zhao, "A reliable and accurate indoor localization method using phone inertial sensors," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, (New York, NY, USA), pp. 421–430, ACM, 2012.

[7] T. Harada, H. Uchino, T. Mori, and T. Sato, "Portable absolute orientation estimation device with wireless network under accelerated situation," in *Robotics and Automa-*

tion, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on, vol. 2, pp. 1412–1417 Vol.2, April 2004.

[8] A. D. Young, M. J. Ling, and D. K. Arvind, "Orient-2: A realtime wireless posture tracking system using local orientation estimation," in *Proceedings of the 4th Workshop on Embedded Networked Sensors*, EmNets '07, (New York, NY, USA), pp. 53–57, ACM, 2007.