

Synchronization in an Indoor Precision Location System

by

Vincent T. Amendolare

A Thesis
Submitted to the Faculty
of the
WORCESTER POLYTECHNIC INSTITUTE
in partial fulfillment of the requirements for the
Degree of Master of Science
in
Electrical and Computer Engineering
by

May 2007

APPROVED:

Professor David Cyganski, Major Advisor

Professor R. James Duckworth

Professor John A. Orr

Abstract

This thesis was conducted as part of the efforts related to WPI's Precision Personnel Location (PPL) project, the purpose of which is to locate emergency personnel in hazardous indoor environments using radio location techniques. A unique signal processing algorithm, σ ART, developed within the PPL project provides means to determine precise position estimates of a wideband transmitter from multipath corrupted signals captured by distributed receivers. This algorithm has synchronization requirements that can not be met without extraordinary expense and complexity by direct means. This thesis develops digital signal processing that achieves the necessary synchronization to satisfy the σ ART algorithm requirements without additional implementation complexity. The mathematical underpinnings of this solution are introduced and the results are evaluated in the context of experimental data.

Acknowledgements

To my family: I am eternally grateful of my family: Dad, Mum, Nick and Rhiannon, for the support in my academic efforts and everything else I do.

To my sponsor: I would like to thank the Department of Justice, National Institute of Justice for funding this research, both for giving me this research opportunity and also striving to protect the lives of personnel in hazardous situations.

To my fellow team members: I would never have achieved this goal without the wisdom and knowledge of my fellow labmates, David Holl and Benjamin Woodacre. Also thanks to every student on this project who has struggled with me through many days of system testing: Jack Coyne, Hauke Daempfling, Hemish Parikh, Shashank Kulkarni, Jay Farmer and Vivek Varshney. And thanks to Bob Boisse for his clever gizmos and driving us around.

To my committee: I also extend my appreciation to my thesis committee members, Professor Duckworth and Professor Orr, for taking the time out of their busy schedules to review this thesis.

To my advisor: Lastly I would like to thank my advisor, Professor Cyganski. His immense knowledge, guidance, patience, and sense of humour have made the last three semesters working for him a great pleasure and very rewarding.

Contents

List of Figures	vi
List of Tables	viii
1 Introduction	1
1.1 Precision Personnel Location Problem	1
1.2 RF-based Indoor Positioning	3
2 A Multicarrier Approach to Precision Indoor Location	7
2.1 Multicarrier Signal Considerations	7
2.1.1 Mathematical Signal Model	7
2.1.2 Aliasing Issues	9
2.1.3 Channel Response	10
2.1.4 Digital Signal Processing	11
2.2 Singular Value Array Reconciliation Tomography	13
2.2.1 Rephasing	14
2.2.2 The First Singular Value	16
2.2.3 Ideal Performance	18
2.2.4 Synchronization Required	18
2.3 Single Pole Frequency Estimation	21
2.3.1 Direct State Space Pole Estimation: Observability Method	22
3 System Hardware	26
3.1 Transmitter Unit	26
3.2 Antennas	27
3.3 Receiver Units	29
4 Synchronization Problems	32
4.1 Sample Clock Drift	32
4.1.1 Sample Clock Drift Problem	33
4.1.2 Reference Array Solution	36
4.2 Constant Time Offset Between Arrays	39

5	Performance	42
5.1	Performance of Sample Clock Drift Tracking	42
5.1.1	Drift Tracking Proof of Concept	42
5.1.2	Drift Tracking Performance in a 60 MHz System	48
5.2	Array Synchronization Performance	54
5.3	Positioning Performance	57
5.3.1	Improved Algorithms	59
5.3.2	Effects of Array Synchronization	59
6	Conclusion	61
6.1	Future Implications	62
	Bibliography	63

List of Figures

1.1	Precision Personnel Location System	4
1.2	Multipath	5
2.1	Example Multicarrer Magnitude Spectrum	8
2.2	Sinusoid Periodicity	10
2.3	Magnitude of DFT Output for Frequencies on and off DFT bins	12
2.4	Rephasing Procedure	15
2.5	Basic σ ART Simulation	19
2.6	Position Errors vs. Timing Errors	20
3.1	Transmitter Block Diagram	26
3.2	Transmit Antenna	28
3.3	Receive Antenna with Ground Plane	28
3.4	Receiver Block Diagram	29
3.5	Graphical User Interface	31
4.1	Temporal Symbol Alignment	34
4.2	Time Estimate Error Performance	38
4.3	Inaccurate Synchronization Between ADCs	39
4.4	Configuration for Synchronization Scheme	41
5.1	AK317a Laboratory	43
5.2	Steel Studded walls in Atwater Kent	43
5.3	AK317a Test Layout 6/16/2006	44
5.4	Tracking $\tau(t)$ at Transmitter Location 2 6/16/2006	45
5.5	$\tau(t)$ Error at Transmitter Location 2 6/16/2006	46
5.6	$\tau(t)$ Estimate Standard Deviation 30 MHz 6/16/2007	47
5.7	Tracking $\tau(t)$ at Transmitter Location 4 6/16/2006	48
5.8	AK317a Test Layout 3/6/2007	49
5.9	Symbol Captures on 4 Elements 3/6/2007	50
5.10	AK317a Test Layout 3/6/2007 First Element Only	51
5.11	Tracking $\tau(t)$ at Transmitter Location 1 on Element 1 3/6/2007	52
5.12	$\tau(t)$ Error from mean at Transmitter Location 1 on Element 1 3/6/2007	52
5.13	$\tau(t)$ Estimate Standard Deviation 60 MHz 3/6/2007	53

5.14 $\tau(t)$ Estimate Standard Deviation 30 MHz 3/6/2007	54
5.15 σ ART with ADC Time Offsets	55
5.16 σ ART with ADC Time Offsets Undone	56
5.17 σ ART with ADC Time Offsets Undone (Zoomed)	56
5.18 σ ART Result Transmitter Location 1 3/6/2007	57
5.19 Position Error Vector Plot 3/6/2007	58
5.20 Position Error Vector Plot 3/6/2007 with Advanced Algorithm	59
5.21 Position Error Vector Plot 3/6/2007 with No Array Synchronization	60

List of Tables

1.1 Acronyms	2
------------------------	---

Chapter 1

Introduction

This thesis was written in support of the Precision Personnel Location (PPL) project which is being conducted by the Electrical and Computer Engineering Department of Worcester Polytechnic Institute (WPI). This project is funded by the Department of Justice's National Institute of Justice. The goal of the project is to design a precision location system to locate first responders in indoor environments. This thesis deals specifically with solving some of the synchronization problems that arose during the development of this system.

Throughout this document several acronyms are used. Each are defined in the text, but are also provided in Table 1 as a quick reference.

1.1 Precision Personnel Location Problem

On December 3, 1999 a veritable tragedy occurred in Worcester, Massachusetts. An aging brick warehouse, the former Worcester Cold Storage and Warehouse Co. building, became host to a fierce inferno. What at first seemed like a routine response by the Worcester Fire Department turned out to be deadly. Two firefighters entered the building initially, concerned that people might be inside. The two men soon found themselves in trouble and called for help over their radios. Four more men were sent in to search for them, but soon became missing themselves. Several more firefighters searched for their comrades for some time, but to no avail. The fire was too fierce, and they were forced to evacuate. None of

ADC	Analog to Digital Converter
AK	Atwater Kent building
DAC	Digital to Analog Converter
DAQ	Data Acquisition Unit
DFT	Discrete Fourier Transform
DSS	Direct State Space
FFT	Fast Fourier Transform
GPS	Global Positioning System
PPL	Precision Personnel Location
RF	Radio Frequency
RMS	Root Mean Square
SNR	Signal to Noise Ratio
σ ART	Singular Value Array Reconciliation Tomography

Table 1.1: Acronyms

the six men made it out of the building alive. [4, 15]

This depressing event left many questions in people’s minds. How could this happen? What could have been done to prevent such a tragedy? If only the firefighters had better knowledge of where to look for their fallen companions, they may have been able to all get out safely. Could this have been possible? Does the technology exist to keep track of individuals in such an environment? Sadly, it does not. It is for this reason that the Electrical Engineering Department of Worcester Polytechnic Institute has assembled a team of researchers, funded by the United States government’s Department of Justice, to create such a technology.

The goal that our team is trying to achieve is a location system for personnel such as firefighters that requires no preexisting infrastructure and can perform accurately in hazardous indoor environments. This project has been titled the Precision Personnel Location project. In our proposed system, an incident commander would be constantly updated with the knowledge of the location of his/her personnel. The system needs to be able to be put into effect quickly, without any preexisting knowledge of the site, since it is never known when or where a fire response (or similar event) is needed. [17]

The desired level of accuracy is that our solution estimates are within 1 foot of the true position. This is so that there is as little ambiguity as possible as to where the personnel

are located, whether on one side of a wall or another for example. The system should also record position estimates over time so that the paths taken by personnel are available in case they are needed to direct personnel back the way they came or send someone after them. Other features of the system will include a wireless data channel for sending information such as atmospheric and physiological conditions. [17]

There are many avenues that may lead to a solution to the indoor location problem. The existing Global Positioning System (GPS) uses radio frequency electromagnetic waves to perform positioning, but suffers from poor accuracy indoors in its standard operation [12]. Other methods include inertial navigation and even ultrasound based systems [11].

1.2 RF-based Indoor Positioning

Our team decided to use radio frequency electromagnetic waves (RF) as our means for positioning. In the proposed scheme, a person inside a building to be located would wear a transmitter device with an antenna to generate a signal. Outside the building this signal would be received by units with antennas on them.

It is known that in air electromagnetic waves at radio frequencies travel at very close to the speed of light in a vacuum, 299,792,458 meters per second. If the time the signal took to propagate through the air from the transmit antenna to the receive antennas could be determined, then the distance between the antennas, called the range, can be deduced.

$$range = \frac{c}{t}, \tag{1.1}$$

where c is the speed of light in meters per second, and t is the time of propagation from the transmit antenna to the receive antenna.

If the range from the transmit antenna to several receive antennas outside the building can be measured, then the position of the transmitter may be determined geometrically. This assumes that the locations of the receive antennas are known. Thus some sort of procedure must be taken when the system is put into effect to determine these receive antenna positions. This may be performed manually by surveying with measuring devices, although in an emergency situation this would not be practical. Thus the locations of the

receive antennas should be measured automatically by the system using a radio ranging approach similar to how the transmitter is located. In the firefighter application the receive antennas would most likely be fixed to firetrucks parked outside the building. This is illustrated in Figure 1.1.

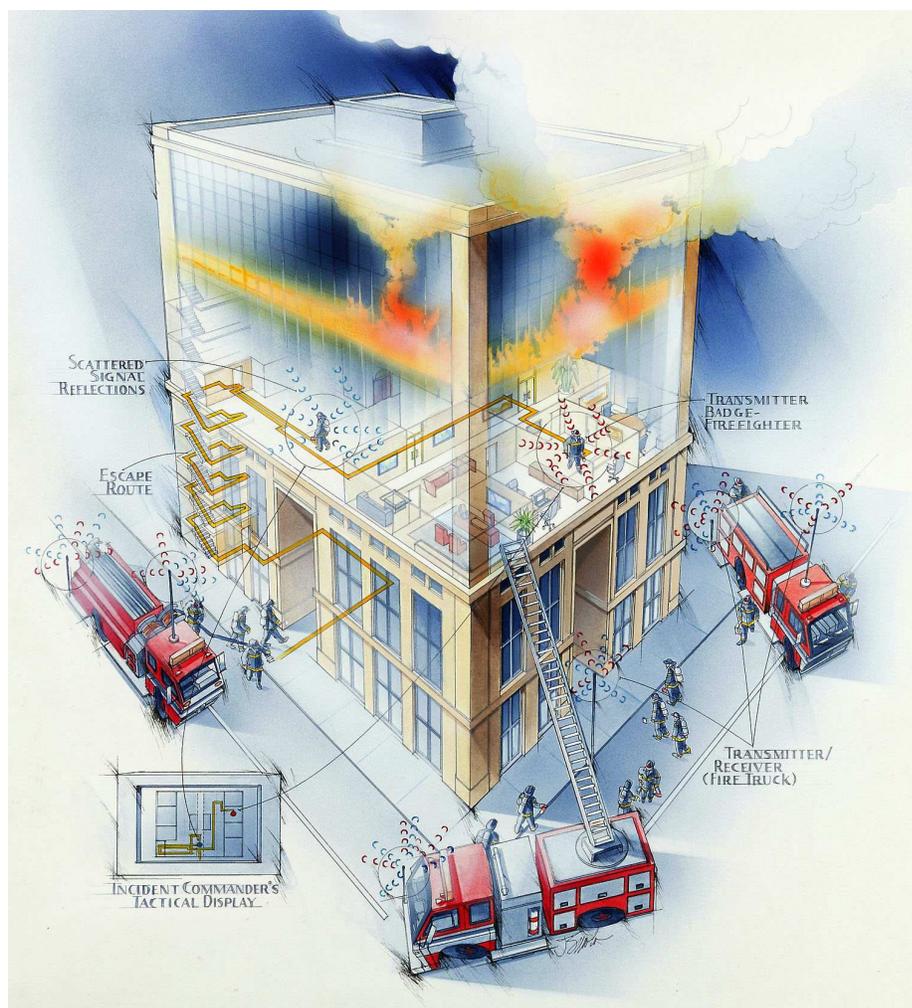


Figure 1.1: Precision Personnel Location System

The largest challenge with performing indoor location with radio frequency electromagnetic waves is the complexity of the radio propagation environments involved. Radio waves are reflected by metal objects, which are plentiful in indoor environments. This results in the signals received being a combination of the so called direct path signal, and reflected

signals, called multipath. This behavior is illustrated in Figure 1.2. The direct path is

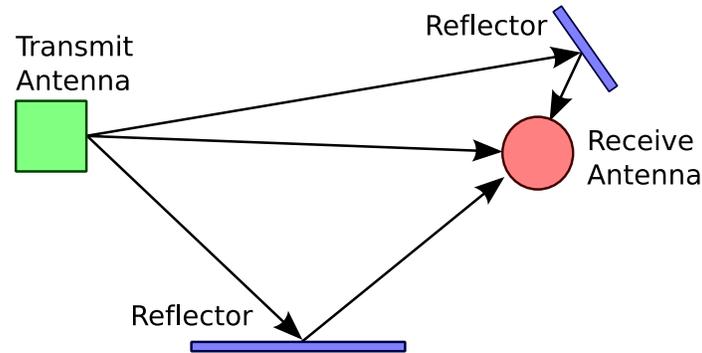


Figure 1.2: Multipath

always the shortest path, since it travels directly from the transmit antenna to the receive antenna. But before one could hope to make such a distinction the direct path signal and multipath signals must be disentangled somehow. This problem contributes strongly to why current GPS systems do not function accurately indoors. A successful RF indoor positioning system must find a way to take the multipath problem into account and mitigate its effects.

There are different possible RF signal structures that may be used to solve the indoor positioning problem. One notable approach is known as Impulse Ultra-Wideband (UWB) which uses a series of pulses that are very narrow in time, and thus very broad in frequency. This causes the signal to occupy frequencies in use by other services so such sources are regulated to use low power levels so as to not interfere with these other services. The low power levels make it difficult to perform location at larger distances. [16]

The Impulse Ultra-Wideband approach attempts to perform one dimensional ranging between the transmit antenna and each receive antenna. Then from these one dimensional ranges a position estimate is deduced. In practice it is generally not possible to determine the absolute distance from the transmitter to receiver. Instead the relative differences of the ranges is determined, adding an extra degree of freedom that is resolvable with an additional receive antenna position. This is known as the Time Difference Of Arrival (TDOA) approach, as opposed to absolute Time Of Arrival (TOA). From these range

estimates (or relative range estimates) the position can be deduced. One algorithm that can achieve this is the Bard algorithm [2].

Because of the issues inherent with the Impulse Ultra-Wideband approach, the early efforts in the WPI Precision Personnel Location project considered alternate positioning systems and associated ranging signals. The next chapter will introduce the method that was chosen, which uses a so called multicarrier signal.

Chapter 2

A Multicarrier Approach to Precision Indoor Location

The WPI PPL project group decided to use an RF-based positioning system with a signal structure quite different from GPS and Impulse Ultra-Wideband, a multicarrier approach [7]. Using this method we interpret received signals to obtain position estimates by frequency domain analysis where methods have been determined to accept direct path signals and reject multipath signals. The chosen multicarrier signal consists of several unmodulated sinusoids evenly spaced in frequency.

2.1 Multicarrier Signal Considerations

In this section we will discuss the behavior of the signal structure we have selected, tradeoffs to be considered, and some of our decisions regarding the use of the signal.

2.1.1 Mathematical Signal Model

Consider our signal, which is a sum of several sinusoids. For now we will assume that they are evenly spaced in frequency and have an initial phase angle of zero.

$$g(t) = \sum_{n=0}^{m-1} e^{-j2\pi(f_0+n\Delta f)t} \quad (2.1)$$

where we have m sinusoids spaced Δf apart in frequency, and the lowest frequency sinusoid has a frequency f_0 . Taking the Fourier transform of $g(t)$ we get

$$G(f) = 2\pi \sum_{n=0}^{m-1} \delta(2\pi(f - (f_0 + n\Delta f))), \quad (2.2)$$

which is a series of impulses in the frequency domain. Figure 2.1 shows an example of what the magnitude spectrum of a Multicarrier signal with 10 carriers looks like in the frequency domain.

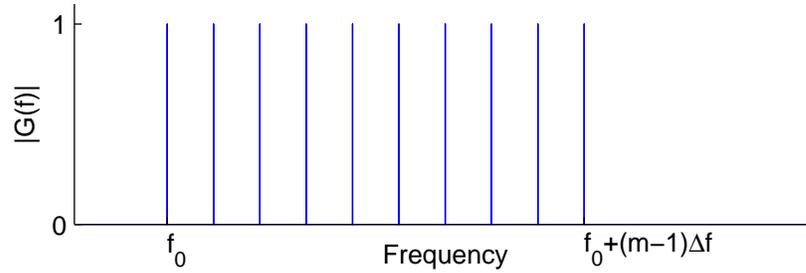


Figure 2.1: Example Multicarrier Magnitude Spectrum

What happens when we time delay our signal? This effect is of course what we're interested in exploiting to perform our positioning. In the time domain a signal $g(t)$ delayed by t_0 seconds becomes $g'(t) = g(t - t_0)$. In the frequency domain this translates into the relationship [13, p. 265]

$$g(t - t_0) \iff G(f)e^{-j2\pi ft_0}. \quad (2.3)$$

We see that when a time delay is introduced the phase spectrum is changed but not the magnitude spectrum. In the case of our delayed multicarrier signal we have

$$G'(f) = 2\pi \left(\sum_{n=0}^{m-1} \delta(2\pi(f - (f_0 + n\Delta f))) \right) e^{-j2\pi ft_0}. \quad (2.4)$$

The phase angle changes linearly with frequency by a factor determined by t_0 . In theory then we should be able to determine time delay using only one carrier and measuring its received phase angle.

It is also important to note that our carriers initial phase angles may be arbitrary. So

we can reconsider the definition of our transmitted signal

$$G(f) = 2\pi \sum_{n=0}^{m-1} \delta(2\pi(f - (f_0 + n\Delta f)))e^{j\theta_n}. \quad (2.5)$$

Where θ_n are arbitrary initial phase angles. Generally speaking, these phase angles are known and under our control, and may be removed by division in further processing.

2.1.2 Aliasing Issues

Any signal consisting of a sum of sinusoids must be periodic. Let us consider first the simplest case, a single sinusoid with period T

$$f_1(t) = \cos\left(\frac{2\pi t}{T}\right). \quad (2.6)$$

Since the signal repeats itself every time one period elapses, we have ambiguity when trying to determine its time delay. We can only determine the time delay modulo T . This problem is not just true of sinusoids, but any periodic function. Any Multicarrier signal will be periodic. For example, with a two-carrier signal, if one of our sinusoids repeats itself every two seconds and the other every three seconds, then their sum will repeat itself every six seconds, see Figure 2.2. The period or aliasing window of a Multicarrier signal is determined by the least common multiple of the periods of the individual carriers.

Alternatively we can interpret this aliasing issue in the same way one interprets aliasing with sampling signals in the time domain. The Nyquist-Shannon sampling theorem states that from samples of a continuous-time signal, the signal can be exactly reconstructed if the signal is bandlimited and the sampling frequency is greater than twice the signal highest frequency in the signal. If this condition is not met, then frequency content above twice the sampling frequency is aliased back into the lower part of the band. [13, p. 321]

Our situation is exactly the same, except instead of sampling signals in the time domain we are sampling signals in the frequency domain at each of our carrier frequencies. We are sampling the complex exponential term $e^{-j2\pi ft_0}$. If t_0 is large relative to our carrier spacing then we will have aliasing. Thus our carrier spacing is what determines the aliasing window of our Multicarrier signal.

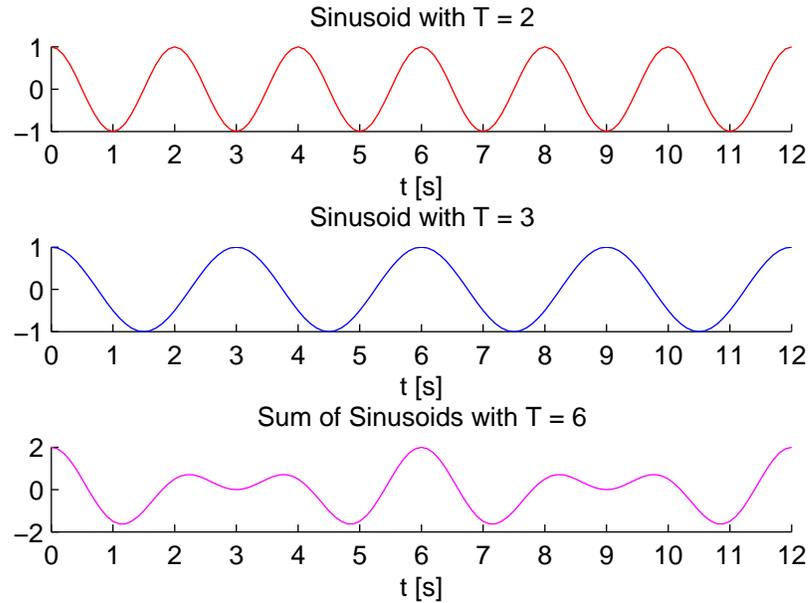


Figure 2.2: Sinusoid Periodicity

We need to ensure that this aliasing window is large enough to avoid any ambiguity. For radio waves, time delays translate into ranges via Equation 1.1. If we assume that the range from the receive antenna to the transmit antenna is less than a certain amount (several hundred meters) then we can determine how large our aliasing window needs to be.

2.1.3 Channel Response

In a real situation the channel through which our signal propagates induces effects on our signal as well. We assume that the channel is passive (there are no active transmitters other than our own) and linear. The channel is time varying however as the transmitter and objects in the building may move over time. We'll denote the transfer function of the channel as $H(f)$, the effect the channel has on our transmitted signal as a function of frequency. Thus a received signal

$$R(f) = G(f) \cdot H(f). \quad (2.7)$$

Our carriers effectively sample the channel response at their respective frequencies.

How do we expect our channel to behave? Even with no multipath, our signal will be attenuated by some factor $\alpha(f)$. The farther the transmit antenna is from the receive antenna, the less of the radiated energy is captured. Objects such as walls can attenuate the signal as well, and this attenuation can even be a function of frequency. Thus a received signal $R(f)$ with one signal source can be written as

$$R(f) = G(f)\alpha(f)e^{-j2\pi ft_0}. \quad (2.8)$$

Now we will introduce the effects of multipath. Our received signal $r(t)$ is a sum of multiple delayed versions of the original signal.

$$R(f) = \sum_{i=1}^k G(f)\alpha_i(f)e^{-j2\pi ft_i}, \quad (2.9)$$

where we sum over i , referencing the direct path signal and $k - 1$ additional reflected signals.

Our task is to interpret such signals from several receive antennas and deduce from them a position estimate. This is performed by our σ ART algorithm to be described in Section 2.2.

2.1.4 Digital Signal Processing

All of our signal processing is done discretely and in the frequency domain. For this reason our multicarrier signal was chosen with the Discrete Fourier Transform (DFT) in mind. We use the Fast Fourier Transform algorithm (FFT) to compute the DFT. The DFT takes a discrete time signal of some number of samples and converts it into a frequency domain representation. The spectral energy of the signal is distributed into frequency bins, where there are the same number of bins as the number of samples.

Consider a sampled signal of n samples at a sampling frequency f_s . If the DFT is performed on this signal then n frequency bins are created. These bins represent frequencies evenly spaced from $-\frac{f_s}{2}$ to $\frac{f_s}{2}$, with a spacing of $\frac{f_s}{n}$. The complex values in these bins represent the magnitude and phase of positive and negative frequencies up to half of the sampling frequency. [13, p. 319]

If a sinusoid in the time signal has an exact frequency corresponding to one of the frequency bins, then all of its energy will be present in that bin. If the frequency is not precisely a bin frequency, then most of the energy will be present in the closest bin but some

of the energy will reside in neighboring bins. This is known as leakage, and can reduce signal to noise ratio and also cause sinusoids of different frequencies to interfere with each other, muddling the underlying information. Figure 2.3 shows this behavior. We can see one case where the DFT was taken of a signal with the exact frequency of a DFT bin; this contrasts with another signal with the same amplitude wherein the frequency is in between two bins. We can see the energy has leaked into neighboring bins.

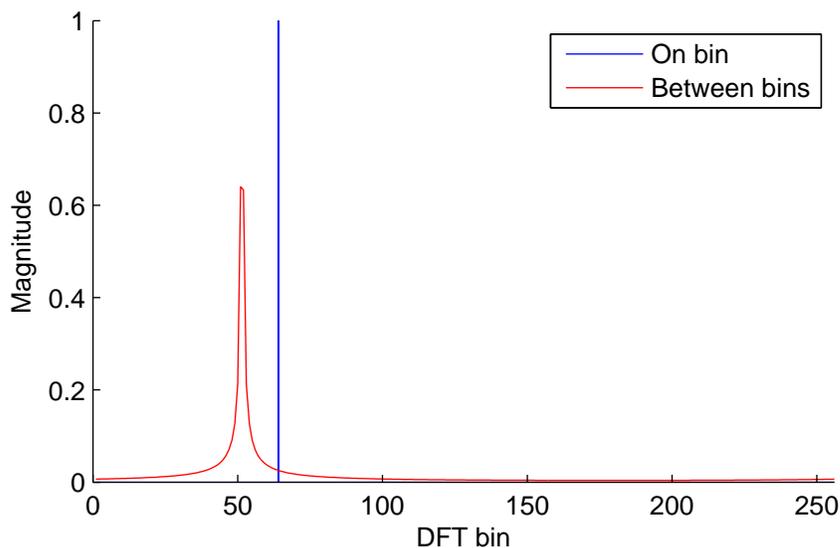


Figure 2.3: Magnitude of DFT Output for Frequencies on and off DFT bins

Our signal travels through the air at radio frequencies, centered about some center frequency. After the signal is received at the antennas the signal is downconverted to baseband frequencies so that it can be sampled. We chose our carrier spacing carefully such that it is an integer number of DFT bins, and the local oscillator frequency by which we downconvert, such that our multicarrier signal frequencies are centered on the DFT bins to minimize leakage. That means that after we take the DFT of our sampled signal the magnitude and phase of our received signal carriers are directly represented in the appropriate bins. This reduces our carrier representation effectively to a list. Our analog received signal is discretized and becomes

$$R(n\Delta f) = e^{j\theta_n} H(f_0 + n\Delta f). \quad (2.10)$$

As stated, our carriers essentially sample the channel response, which is returned as the vector we will denote as $R(n\Delta f)$. To obtain $R(n\Delta f)$ we index the output of our DFT at the appropriate bins.

In our system we have decided to use a digital sample clock of 200 MHz. We take our DFTs on blocks of 8192 samples at a time, which are known as symbols. This means that our symbols have a time duration of 40.96 μ s. We also use the same symbol definition with our transmitter as we do with our receiver. The transmitter uses a software radio approach wherein it plays a digitally stored version of the desired multicarrier signal through a digital to analog converter. Since our signal is periodic with respect to the symbol window, our transmitter can simply play the same 8192 samples repeatedly to transmit a constant multicarrier signal.

It is also important to note that since we chose our multicarrier signal to have its carriers centered on our FFT bins that this means our signal is periodic with respect to the symbol window. This fact gives us another considerable benefit. Suppose we capture a symbol of data at one of our receive antennas. What if we capture a second symbol immediately after it? If we assume that the channel response has not changed, then the second symbol will be the same as the first (with different noise). This means it may be valid for us to analyze symbols captured at different times if necessary. We will explore such an analyses in Chapter 4.

2.2 Singular Value Array Reconciliation Tomography

The precision location algorithm developed in the WPI PPL project that is currently used in our system was named by this group Singular Value Array Reconciliation Tomography (σ ART). This novel algorithm obtains a position estimate directly with received data from all of the receive antennas, rather than determining the ranges from the transmitter to each receive antenna independently followed by a multilateralization solution of the indicated source position.

2.2.1 Rephasing

σ ART is an exhaustive algorithm. The entire space that the transmitter may reside in is discretized as a grid with some spatial resolution. This grid is scanned and a metric is evaluated at each scan location in the grid. Our metric is chosen such that it should be maximized at the transmitter location. The sampling density of the grid is chosen based upon the bandwidth of our multicarrier signal. The spatial resolution may also be increased using interpolation techniques after the completion of the σ ART scan.

We start with our frequency domain data $R_l(n\Delta f)$, the received complex values at each of our carrier locations from the l th receive antenna, which can be considered a vector. These vectors form the columns in our raw data matrix \mathbf{R} , with one column for each receive antenna.

$$\mathbf{R} = \begin{bmatrix} R_0(0\Delta f) & R_1(0\Delta f) & \dots & R_p(0\Delta f) \\ R_0(1\Delta f) & R_1(1\Delta f) & \dots & R_p(1\Delta f) \\ \vdots & \vdots & \ddots & \vdots \\ R_0(m\Delta f) & R_1(m\Delta f) & \dots & R_p(m\Delta f) \end{bmatrix}, \quad (2.11)$$

where we have p receive antennas and m carrier frequencies.

At each point in space that is scanned, the distance from that location to each of the receive antennas is determined from the known locations of the receive antennas. Figure 2.4 illustrates this procedure. The left side of Figure 2.4 shows the ideal behavior of our transmitted signal, which simply gets delayed as it propagates from the transmit antenna to the receive antennas, with a different delay t_0, t_1, \dots, t_p determined by the range from the transmitter to each receiver via Equation 1.1. The space that the transmitter may be located within is scanned in a grid pattern as depicted on the right side of Figure 2.4. At each scan location the ranges from the scan location to the receive antennas is computed, and then the corresponding time offsets are inversely applied to the received data.

Thus, the received data matrix \mathbf{R} has a negative time offset applied to all its columns. This undoes the time delay that would have been applied to the data by propagating through free space if the transmitter was at that location. So for every point in space we have a

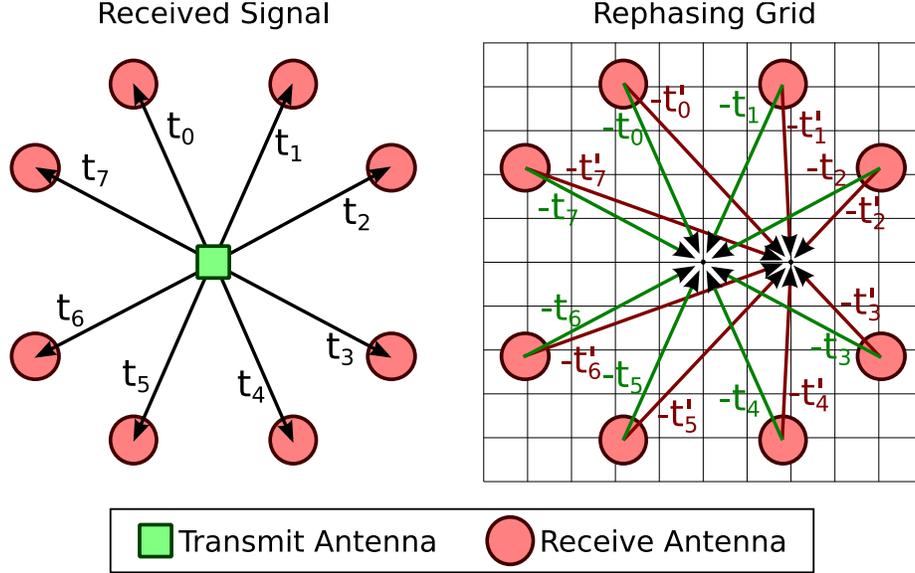


Figure 2.4: Rephasing Procedure

different rephased version of \mathbf{R} we'll call \mathbf{R}' , where

$$\mathbf{R}' = \begin{bmatrix} R_0(0\Delta f)e^{j2\pi 0\Delta f t_0} & R_1(0\Delta f)e^{j2\pi 0\Delta f t_1} & \dots & R_p(0\Delta f)e^{j2\pi 0\Delta f t_p} \\ R_0(1\Delta f)e^{j2\pi 1\Delta f t_0} & R_1(1\Delta f)e^{j2\pi 1\Delta f t_1} & \dots & R_p(1\Delta f)e^{j2\pi 1\Delta f t_p} \\ \vdots & \vdots & \ddots & \vdots \\ R_0(m\Delta f)e^{j2\pi m\Delta f t_0} & R_1(m\Delta f)e^{j2\pi m\Delta f t_1} & \dots & R_p(m\Delta f)e^{j2\pi m\Delta f t_p} \end{bmatrix}. \quad (2.12)$$

The time delay terms t_0, t_1, \dots, t_p correspond to the RF propagation times from the scan location to each receive antenna.

An important point to make is that this rephasing does not change the “energy” or Frobenius Norm of the matrix. The Frobenius Norm of a matrix \mathbf{A} with b elements, regardless of shape is defined as [14]

$$E = \sqrt{\sum_{a=1}^b |\mathbf{A}_a|^2}. \quad (2.13)$$

Since the rephasing only affects the phase, and not the magnitude of the elements of \mathbf{R} , we can conclude that the energy in the matrix remains constant as \mathbf{R} is rephased to create \mathbf{R}' .

2.2.2 The First Singular Value

At each location an operation is performed on \mathbf{R}' to obtain a metric indicating how strongly that location is judged as a potential position estimate, meaning the consistency of the rephased data with the ideal signal structure for that location. The metric is the first singular value of the singular value decomposition of \mathbf{R}' . Consider the ideal case where our transmitted signal is the function

$$G(f) = 2\pi \left(\sum_{n=0}^{m-1} \delta(2\pi(f - (f_0 + n\Delta f)))e^{-j\Theta_n} \right). \quad (2.14)$$

so all of our carriers have some arbitrary phase angle Θ_n . After the signal propagates through space it gets delayed by t_0 and becomes

$$G'(f) = 2\pi \left(\sum_{n=0}^{m-1} \delta(2\pi(f - (f_0 + n\Delta f)))e^{-j\Theta_n} \right) e^{-j2\pi ft_0}, \quad (2.15)$$

assuming for now that there is no attenuation to consider. So, for our raw data matrix \mathbf{R} , in this case

$$G'_l(f) = G(f)e^{-j2\pi ft_l} \Rightarrow G(n\Delta f)e^{-j2\pi n\Delta ft_l}, \quad (2.16)$$

where l is an index of our p receive antennas. Thus

$$\mathbf{R} = \begin{bmatrix} G(0\Delta f)e^{-j2\pi 0\Delta ft_0} & G(0\Delta f)e^{-j2\pi 0\Delta ft_1} & \dots & G(0\Delta f)e^{-j2\pi 0\Delta ft_p} \\ G(1\Delta f)e^{-j2\pi 1\Delta ft_0} & G(1\Delta f)e^{-j2\pi 1\Delta ft_1} & \dots & G(1\Delta f)e^{-j2\pi 1\Delta ft_p} \\ \vdots & \vdots & \ddots & \vdots \\ G(m\Delta f)e^{-j2\pi m\Delta ft_0} & G(m\Delta f)e^{-j2\pi m\Delta ft_1} & \dots & G(m\Delta f)e^{-j2\pi m\Delta ft_p} \end{bmatrix}. \quad (2.17)$$

When we perform our scan and are at the correct transmitter location the time delay term is undone, so the phases are again all zeros.

$$\mathbf{R}' = \begin{bmatrix} G(0\Delta f)e^{-j2\pi 0\Delta f(t_0-t_0)} & G(0\Delta f)e^{-j2\pi 0\Delta f(t_1-t_1)} & \dots & G(0\Delta f)e^{-j2\pi 0\Delta f(t_p-t_p)} \\ G(1\Delta f)e^{-j2\pi 1\Delta f(t_0-t_0)} & G(1\Delta f)e^{-j2\pi 1\Delta f(t_1-t_1)} & \dots & G(1\Delta f)e^{-j2\pi 1\Delta f(t_p-t_p)} \\ \vdots & \vdots & \ddots & \vdots \\ G(m\Delta f)e^{-j2\pi m\Delta f(t_0-t_0)} & G(m\Delta f)e^{-j2\pi m\Delta f(t_1-t_1)} & \dots & G(m\Delta f)e^{-j2\pi m\Delta f(t_p-t_p)} \end{bmatrix}. \quad (2.18)$$

$$\mathbf{R}' = \begin{bmatrix} G(0\Delta f) & G(0\Delta f) & \dots & G(0\Delta f) \\ G(1\Delta f) & G(1\Delta f) & \dots & G(1\Delta f) \\ \vdots & \vdots & \ddots & \vdots \\ G(m\Delta f) & G(m\Delta f) & \dots & G(m\Delta f) \end{bmatrix}. \quad (2.19)$$

All of the columns of our matrix are identical, which means they are linearly dependent and \mathbf{R}' has a rank of 1. This means that all of the energy of the matrix is in the first singular value.

Even if the columns were scaled by constants because the carriers underwent different attenuations and phase delays, the rank would still be 1. Another important property of σ ART is that it is agnostic of the magnitude and phase angles of $G(f)$, the transmitted signal. This means that the carrier phase angles θ_n of our transmitted waveform can be arbitrary. Now consider the case where we have a global time offset τ_g applied to each column of received data. When the data is rephased to the correct location in the ideal case, we have

$$\mathbf{R}' = \begin{bmatrix} G(0\Delta f)e^{-j2\pi 0\Delta f\tau_g} & G(0\Delta f)e^{-j2\pi 0\Delta f\tau_g} & \dots & G(0\Delta f)e^{-j2\pi 0\Delta f\tau_g} \\ G(1\Delta f)e^{-j2\pi 1\Delta f\tau_g} & G(1\Delta f)e^{-j2\pi 1\Delta f\tau_g} & \dots & G(1\Delta f)e^{-j2\pi 1\Delta f\tau_g} \\ \vdots & \vdots & \ddots & \vdots \\ G(m\Delta f)e^{-j2\pi m\Delta f\tau_g} & G(m\Delta f)e^{-j2\pi m\Delta f\tau_g} & \dots & G(m\Delta f)e^{-j2\pi m\Delta f\tau_g} \end{bmatrix}. \quad (2.20)$$

This matrix still has a rank of 1 since all columns are identical. This means that an arbitrary time offset on $G(t)$ will not affect the results, as long as the time delay is global and applied uniformly to all columns of the received data matrix. This is an important property that we will discuss later.

Now suppose that our signals have been rephased to an incorrect location. In this case the variable τ_l will be used to represent a rephasing time delay that does not match the true time delay.

$$\mathbf{R}' = \begin{bmatrix} G(0\Delta f)e^{-j2\pi 0\Delta f(t_0-\tau_0)} & G(0\Delta f)e^{-j2\pi 0\Delta f(t_1-\tau_1)} & \dots & G(0\Delta f)e^{-j2\pi 0\Delta f(t_p-\tau_p)} \\ G(1\Delta f)e^{-j2\pi 1\Delta f(t_0-\tau_0)} & G(1\Delta f)e^{-j2\pi 1\Delta f(t_1-\tau_1)} & \dots & G(1\Delta f)e^{-j2\pi 1\Delta f(t_p-\tau_p)} \\ \vdots & \vdots & \ddots & \vdots \\ G(m\Delta f)e^{-j2\pi m\Delta f(t_0-\tau_0)} & G(m\Delta f)e^{-j2\pi m\Delta f(t_1-\tau_1)} & \dots & G(m\Delta f)e^{-j2\pi m\Delta f(t_p-\tau_p)} \end{bmatrix}. \quad (2.21)$$

Now each of our columns has a complex exponential term with a different frequency with respect to row index, thus each of our columns should be linearly independent. So our matrix \mathbf{R}' now has full rank and the first singular value as a result must be smaller. Only at the correct location does the metric take on its maximum value.

The performance of σ ART is limited by bandwidth however. We are dealing with windowed sinusoids since we have finite bandwidth, thus our columns will never be completely orthogonal. The closer in frequency two columns are the more linearly dependent they will become. Thus σ ART like other approaches benefits from having more signal bandwidth.

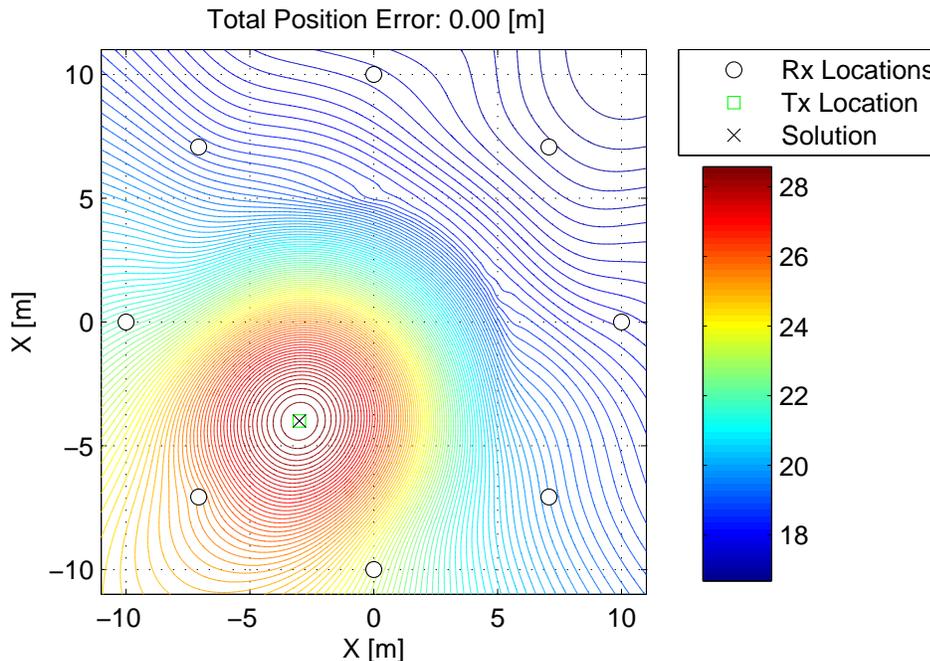
2.2.3 Ideal Performance

We contrived a basic test based upon a two dimensional scan to simulate the behavior of the σ ART algorithm. A symmetric radial geometry was chosen for our receive antenna configuration with 8 antennas. Pristine synthetic frequency domain data was used, free of noise and multipath, to simulate received data from a transmitter at an arbitrary location with a bandwidth of 30 MHz. Thus our data matrix \mathbf{R} captures the features of the ideal case previously discussed. Our spatial scan resolution was 0.5 meters. This resolution was increased with interpolation by a factor of ten to 0.05 meters. Figure 2.5 shows the outcome of σ ART in this case. A contour map is laid over the scan region indicating the strength of the σ ART metric. We see that it is maximized at exactly the location of the transmitter. Hence the position solution, the location where the metric is maximized, is at the correct location.

The behavior of σ ART in regards to bandwidth, number of carriers, geometry, presence of reflectors, etc. is important for us to understand for the implementation of our system, but is beyond the scope of this thesis. We will focus on the synchronization requirements of σ ART.

2.2.4 Synchronization Required

Since the phase angles of $G(f)$ are arbitrary, a global time offset applied to all of the columns of \mathbf{R} does not alter the results obtained with σ ART. What certainly would not

Figure 2.5: Basic σ ART Simulation

be acceptable however would be the condition wherein columns had phase offsets due to differing additional time offsets.

Such differing offsets can arise, for example, after our signal is received by an antenna, as the travels down a cable before it can be downconverted and sampled. These cables act as transmission lines and introduce additional delay to our signal before it is digitized. If all of the cables induce the same delay, then there is no issue with σ ART because it would be a common time offset across all columns of \mathbf{R} . If the cables differ in length/delay slightly however, the solution can be perturbed. For this reason, preparation of our PPL system requires that procedures be followed to calibrate-out differences in cable delay times. This has been done in the past by using a network analyzer to measure the transfer function of the cables and deducing their time delays.

There are other means by which these erroneous time offsets can creep in: errors in receive antenna position knowledge, retarded signal propagation through walls, and failure to capture data from different receive antennas at exactly the same time. One issue that

must be addressed is how much error of this nature is tolerable. A simulation was executed to determine this.

The purpose of this simulation was to determine the general behavior of σ ART faced with these timing errors. We used the symmetric radial receive antenna geometry with 4, 8 and 16 antennas. A transmitter location was chosen at a location uniformly random within the circle enclosed by the receive antennas. Then each antenna's pristine synthetic data had a random time offset applied to it with some standard deviation. The σ ART algorithm was executed, and the position error was determined. Our position error is defined as the Euclidean distance from the true transmitter position to the estimated transmitter position. We performed this analysis with 50 random transmitter positions for several different standard deviations of timing errors. The root-mean-square (RMS) position errors is plotted versus the standard deviation of the timing errors in Figure 2.6.

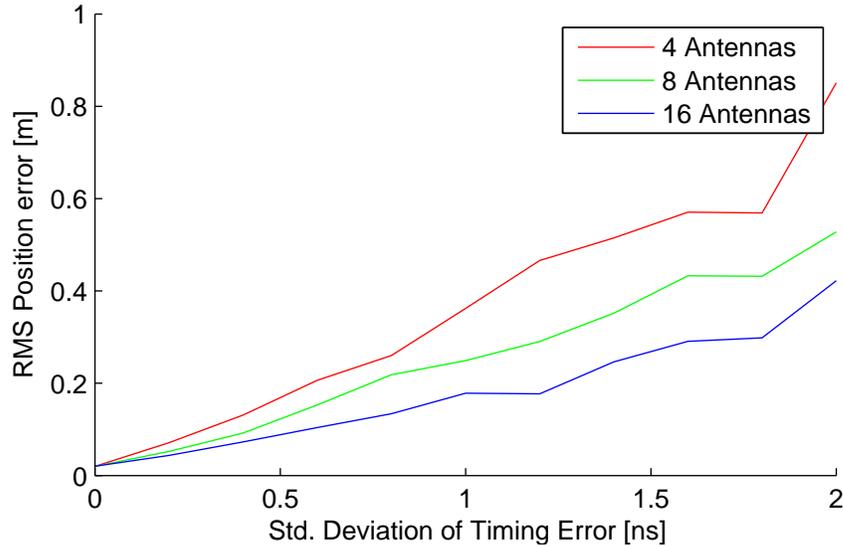


Figure 2.6: RMS Position Errors vs. Standard Deviation of Timing Errors

We observe that as the timing errors get larger, so does the σ ART position error. We also observe that with more receive antennas the position errors are smaller as well. We wish to ensure the degree to which our position errors are perturbed by these timing errors are small enough. Assuming the use of approximately 16 receive antennas, we have decided

that timing accuracy with a 0.5 ns standard deviation (or smaller) is generally acceptable for our system. Thus the position errors due to timing errors are generally less than 0.1 meters.

The RMS position error, σ_p , can be expressed as an approximate function of the standard deviation of the timing error, σ_t .

$$\sigma_p = \frac{a \cdot \sigma_t}{\sqrt{p}}, \quad (2.22)$$

where a is a proportionality constant and p is again the number of antennas. This relationship coincides with previous analytic work conducted for this project [8].

2.3 Single Pole Frequency Estimation

For reasons to be discussed in Chapter 4, a tool used in our signal processing is direct state space (DSS) single pole frequency estimation. This is a model-based estimation technique that can determine the frequency of a discrete signal with a single complex sinusoid in noise more accurately than traditional methods such as the discrete Fourier transform.

The operation basis for state-space estimation can be grasped via a much simpler method than the full DSS frequency estimation algorithm by considering only the case of a single sinusoid. We shall examine this simplification in the following section.

Suppose we have a discrete signal that is a complex sinusoid of frequency f_0

$$f(n\Delta t) = e^{-j2\pi f_0 n\Delta t}. \quad (2.23)$$

Consider the product of a sample's conjugate and the subsequent sample

$$f(n\Delta t)^* \cdot f((n+1)\Delta t) = e^{j2\pi f_0 n\Delta t} \cdot e^{-j2\pi f_0 (n+1)\Delta t}. \quad (2.24)$$

$$pole = f(n\Delta t)^* \cdot f((n+1)\Delta t) = e^{-j2\pi f_0 \Delta t}. \quad (2.25)$$

The result is a complex number with an angle that corresponds to the difference in angle $f_0 \Delta t$ of the original two samples. This complex number is known as the pole. Since Δt is known we can find the desired frequency

$$f_0 = \frac{\angle(f(n\Delta t)^* \cdot f((n+1)\Delta t))}{-2\pi \Delta t}. \quad (2.26)$$

This is valid in the ideal case, but what if our signal is in noise? Instead of basing our frequency estimate on only two samples we can use many more to get the most accurate answer possible.

Say we have m total samples of $f(n\Delta t)$. We can form from this two vectors

$$\mathbf{H}_0 = \begin{bmatrix} f(0\Delta t) \\ f(1\Delta t) \\ \vdots \\ f((m-1)\Delta t) \end{bmatrix} \quad \mathbf{H}_1 = \begin{bmatrix} f(1\Delta t) \\ f(2\Delta t) \\ \vdots \\ f((m)\Delta t) \end{bmatrix}. \quad (2.27)$$

If we take the dot product

$$\langle \mathbf{H}_0^*, \mathbf{H}_1 \rangle = \sum \begin{bmatrix} e^{-j2\pi f_0 \Delta t} \\ e^{-j2\pi f_0 \Delta t} \\ \vdots \\ e^{-j2\pi f_0 \Delta t} \end{bmatrix} = m \cdot e^{-j2\pi f_0 \Delta t}. \quad (2.28)$$

Each element in the sum ideally the same, with the desired phase angle. After the sum we normalize by the magnitudes of the elements of \mathbf{H}_0 .

$$pole = \frac{\langle \mathbf{H}_0^*, \mathbf{H}_1 \rangle}{\sum_{n=1}^{m-1} |H_0(n)|^2}. \quad (2.29)$$

This sum and normalization effectively takes the mean of the poles yielded by the different samples. This averaging improves accuracy in the presence of noise, since any zero-mean errors should tend to average themselves out over many samples.

2.3.1 Direct State Space Pole Estimation: Observability Method

A more elaborate approach to pole estimation exists that uses a state-space model based approach. We will focus on one implementation of this approach that exists known as the observability method. Unlike the simpler version discussed, this method can actually determine the values of multiple poles (frequencies). While for the purposes of this thesis we only need to determine a single pole, this enhancement yields improved pole estimation in the presence of interference and noise.

Consider our signal as a sum of several sinusoids;

$$f(n\Delta t) = \sum_{i=1}^k \alpha_i e^{-j2\pi f_i n\Delta t}, \quad (2.30)$$

where i is an index of k sinusoids with amplitudes α_i and frequencies f_i . We can remodel this as a state-space matrix equation:

$$f(n\Delta t) = \underbrace{\begin{bmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_k \end{bmatrix}}_{\mathbf{C}} \underbrace{\begin{bmatrix} e^{-j2\pi f_1 \Delta t} & 0 & \dots & 0 \\ 0 & e^{-j2\pi f_2 \Delta t} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e^{-j2\pi f_k \Delta t} \end{bmatrix}}_{\mathbf{A}}^n \underbrace{\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}}_{\mathbf{B}}. \quad (2.31)$$

Thus we define the matrices \mathbf{C} , \mathbf{A} , and \mathbf{B} as the three matrices in the above expression, allowing us to simplify the equation as

$$f(n\Delta t) = \mathbf{C}\mathbf{A}^n\mathbf{B}. \quad (2.32)$$

What we wish to determine ultimately are the poles from the matrix \mathbf{A} . To do this we first construct a Hankel matrix, \mathbf{H} .

$$\mathbf{H} = \begin{bmatrix} f(0\Delta t) & f(1\Delta t) & \dots & f(b\Delta t) \\ f(1\Delta t) & f(2\Delta t) & \dots & f((b+1)\Delta t) \\ \vdots & \vdots & \ddots & \vdots \\ f(a\Delta t) & f((a+1)\Delta t) & \dots & f(m\Delta t) \end{bmatrix}. \quad (2.33)$$

where $a + b = m$. The values of a and b can be arbitrary, however it has been determined that the optimum Hankel size for performance occurs when

$$a \approx \frac{2}{3}m, \quad (2.34)$$

$$b \approx \frac{1}{3}m. \quad (2.35)$$

for this Observability method of DSS pole estimation [18]. Our matrix \mathbf{H} can be rewritten as

$$\mathbf{H} = \begin{bmatrix} \mathbf{C}\mathbf{A}^0\mathbf{B} & \mathbf{C}\mathbf{A}^1\mathbf{B} & \dots & \mathbf{C}\mathbf{A}^b\mathbf{B} \\ \mathbf{C}\mathbf{A}^1\mathbf{B} & \mathbf{C}\mathbf{A}^2\mathbf{B} & \dots & \mathbf{C}\mathbf{A}^{b+1}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}\mathbf{A}^a\mathbf{B} & \mathbf{C}\mathbf{A}^{a+1}\mathbf{B} & \dots & \mathbf{C}\mathbf{A}^m\mathbf{B} \end{bmatrix}. \quad (2.36)$$

This matrix can be factored

$$\mathbf{H} = \underbrace{\begin{bmatrix} \mathbf{CA}^0 \\ \mathbf{CA}^1 \\ \vdots \\ \mathbf{CA}^a \end{bmatrix}}_{\mathcal{O}} \underbrace{\begin{bmatrix} \mathbf{A}^0\mathbf{B} & \mathbf{A}^1\mathbf{B} & \dots & \mathbf{A}^b\mathbf{B} \end{bmatrix}}_{\mathcal{C}}. \quad (2.37)$$

We'll label these two resultant matrices \mathcal{O} and \mathcal{C} , which fit the definitions of the Observability and Controllability matrices of a state-space system [5, pp. 145,156]. So we have

$$\mathbf{H} = \mathcal{O}\mathcal{C}. \quad (2.38)$$

We can obtain these two matrices by factoring \mathbf{H} . Deducing them directly from \mathbf{H} is not possible. Whatever method of factorization is chosen however, we know that the resultant matrices $\tilde{\mathcal{O}}$, $\tilde{\mathcal{C}}$ can be related to the original matrices \mathcal{O} , \mathcal{C} by a similarity transformation.

$$\mathbf{H} = \mathcal{O}\mathcal{C} = \tilde{\mathcal{O}}\tilde{\mathcal{C}}. \quad (2.39)$$

For some arbitrary transformation matrix \mathbf{T}

$$\mathcal{O}\mathcal{C} = \tilde{\mathcal{O}}\mathbf{T}^{-1}\mathbf{T}\tilde{\mathcal{C}}. \quad (2.40)$$

To within an allowed degree of freedom we can identify [18]

$$\mathcal{O} = \tilde{\mathcal{O}}\mathbf{T}^{-1}. \quad (2.41)$$

Thus $\tilde{\mathcal{O}}$ corresponds to the Observability matrix of an equivalent state-space system with parameters $\tilde{\mathbf{A}}$, $\tilde{\mathbf{B}}$, and $\tilde{\mathbf{C}}$ where [5, p. 95]

$$\begin{aligned} \tilde{\mathbf{A}} &= \mathbf{T}\mathbf{A}\mathbf{T}^{-1}, \\ \tilde{\mathbf{B}} &= \mathbf{T}\mathbf{B}, \\ \tilde{\mathbf{C}} &= \mathbf{C}\mathbf{T}^{-1}. \end{aligned} \quad (2.42)$$

There are many methods of factorization possible. In our implementation we use the singular value decomposition for its numerical stability. With the singular value decomposition [9, p. 109]

$$\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H = \underbrace{\mathbf{U}\sqrt{\mathbf{\Sigma}}}_{\tilde{\mathcal{O}}} \underbrace{\sqrt{\mathbf{\Sigma}}\mathbf{V}^H}_{\tilde{\mathcal{C}}} = \tilde{\mathcal{O}}\tilde{\mathcal{C}}. \quad (2.43)$$

So we have

$$\tilde{\Theta} = \mathbf{U}\sqrt{\tilde{\Sigma}}. \quad (2.44)$$

At this point another step is performed called rank truncation, which we will not discuss in detail here. See Reference [18].

From our matrix $\tilde{\Theta}$ we construct two smaller matrices:

$$\tilde{\Theta}_- = \begin{bmatrix} \tilde{\mathbf{C}}\tilde{\mathbf{A}}^0 \\ \tilde{\mathbf{C}}\tilde{\mathbf{A}}^1 \\ \vdots \\ \tilde{\mathbf{C}}\tilde{\mathbf{A}}^{a-1} \end{bmatrix}, \quad (2.45)$$

$$\tilde{\Theta}_+ = \begin{bmatrix} \tilde{\mathbf{C}}\tilde{\mathbf{A}}^1 \\ \tilde{\mathbf{C}}\tilde{\mathbf{A}}^2 \\ \vdots \\ \tilde{\mathbf{C}}\tilde{\mathbf{A}}^a \end{bmatrix}. \quad (2.46)$$

Thus we have the relationship

$$\tilde{\Theta}_- \tilde{\mathbf{A}} = \tilde{\Theta}_+. \quad (2.47)$$

We can solve this equation for $\tilde{\mathbf{A}}$

$$\tilde{\mathbf{A}} = \tilde{\Theta}_-^\dagger \tilde{\Theta}_+. \quad (2.48)$$

Note that \dagger denotes the Moore-Penrose pseudo-inverse [10, p. 257]. We know that $\tilde{\mathbf{A}}$ is related to \mathbf{A} by the similarity transformation \mathbf{T} , and also that \mathbf{A} is diagonal. Thus we can find the original \mathbf{A} by taking the eigenvalue decomposition,

$$\mathbf{A} = \mathbf{T}\tilde{\mathbf{A}}\mathbf{T}^{-1}. \quad (2.49)$$

Thus we have our matrix of poles from which we can deduce the frequencies of the sinusoidal components of the signal.

Chapter 3

System Hardware

Our system consists of two major hardware sub-systems, the transmitter unit that is to be worn by personnel being located, and the receiver units outside the building receiving the transmitted signals.

3.1 Transmitter Unit

Our transmitter units are standalone devices to be worn by personnel, thus they must be completely wireless and battery powered. Figure 3.1 shows general functionality of the transmitter units.



Figure 3.1: Transmitter Block Diagram

We use a software defined radio approach for signal flexibility. Our signal is stored digitally and is played through an digital to analog converter (DAC). This way we can reprogram the digital transmit waveform to change our signal without any hardware modifications. Our signal is output from the DAC one sample at a time at a rate of 200

Megasamples per second. This rate is determined by a local 200 MHz sampling clock. Based on the Nyquist sampling theorem a 200 MHz sampling rate can support signals with frequencies only as high as 100 MHz. This lets us generate a baseband signal with content from 0 to 100 MHz.

The baseband signal is then upconverted with a mixer. This raises the frequency content from baseband up to our desired radio frequencies over the air, which is generally centered about 440 MHz. This signal is then driven into our transmit antenna, which converts the electrical signal into an electromagnetic wave that propagates through space to our receive antennas.

The center frequency of 440 MHz and power level of -12 dBm per carrier were chosen for our transmitted signal based on an allotment from the Federal Communications Commission. Previously we were allotted 30 MHz of bandwidth (420-450 MHz) but more recently we have been allotted 60 MHz (410-470 MHz). System tests using the 30 MHz allotment used 103 carriers, providing a range aliasing window of 1,023 meters. More recent tests using a 60 MHz bandwidth also use 103 carriers, halving our range aliasing window to 512 meters. This is sufficient as our current testing scale is no more than 30 meters between our farthest separated antennas.

3.2 Antennas

The antennas currently in use in our system are vertical dipole antennas optimized for response in about 440 MHz. These antennas emit vertically polarized radiation. Figure 3.2 is a picture of the vertical dipole currently used as our transmit antenna. This antenna is also omnidirectional in the horizontal plane. In a final system a different antenna more suitable to be worn by personnel would likely be chosen.

Our receive antennas are also vertical dipoles but with ground planes attached shown in 3.3. The ground planes stop signals from being picked up by the antenna that come from behind. The antennas are faced towards the building in which we are doing location, with the ground planes on the back of the antennas to block any reflections from the back which we know can not be the direct path.



Figure 3.2: Transmit Antenna

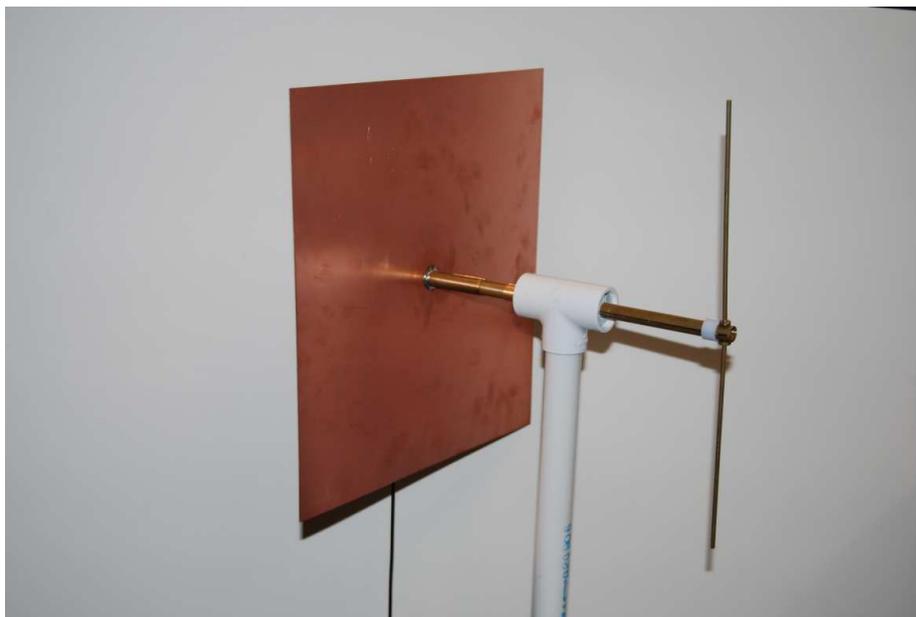


Figure 3.3: Receive Antenna with Ground Plane

3.3 Receiver Units

Receiving and processing our radio signals are our receiver units. Figure 3.4 shows general functionality of the receiver units in an example configuration. Our signal is received

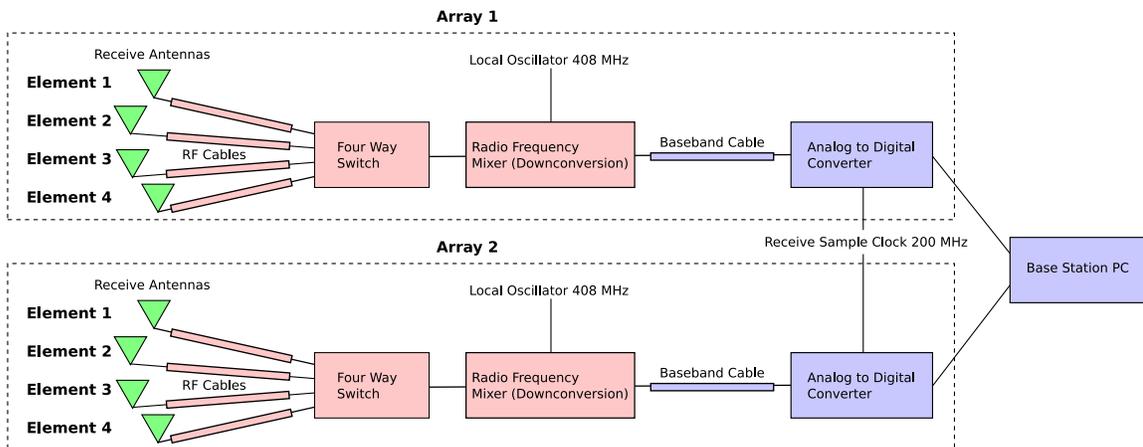


Figure 3.4: Receiver Block Diagram

by our receive antennas and then propagates down a cable at its radio frequencies still centered at 440 MHz. We call these cables our RF cables. Each antenna has its own RF cable. As mentioned previously, it is very important for σ ART to work that there are no differences in time delay in the different cables. Thus we need to undo any differences in cable delays in software. These differences were captured during special calibration testing and should remain constant.

Next the RF cables may connect to an RF switch. The switches enable us to time multiplex the captures from the different antennas, so we do not need a complete receive hardware chain for every antenna. These switches can support up to four antennas.

Next our signal is downconverted from radio frequencies centered at 440 MHz to baseband frequencies by a mixer and amplified, which we'll call an RF front-end. This brings all of our frequency content down below 100 MHz so it can be then sampled by a 200 MHz analog to digital converter (ADC). Before the ADCs however, the signal propagates over a signal at baseband. We call these cables Baseband Cables, and it is also important for us to calibrate out any delay differences in these cables. Figure 3.4 shows two main receiver

chains ending with the two ADCs. Our current system supports up to five of these chains. These chains are commonly called “arrays”. The antennas on each array are called “elements” and labeled by which port they are connected to on their respective switch. This scheme will be discussed further in Section 4.1.1.

The ADCs are controlled by customized Field Programmable Gate Arrays (FPGAs), which organize the data and prepare it for transmission over Ethernet to our Base Station Computer. Collectively the ADC and FPGA pairs are called data acquisition units, or DAQs. Once the data is received at the Base Station Computer, our signal processing is applied. Our signal is converted into the frequency domain and analyzed with σ ART. Once a position estimate is determined it is displayed in a real-time graphical user interface. An example image of this graphical user interface is shown in Figure 3.5, which displays the multicarrier spectrum from the receive antennas, the σ ART metric image, position error and other information useful for our system testing. Figure 3.5 shows the graphical user interface running in a simulation mode, though it also is the interface used when collecting real data during our system tests.

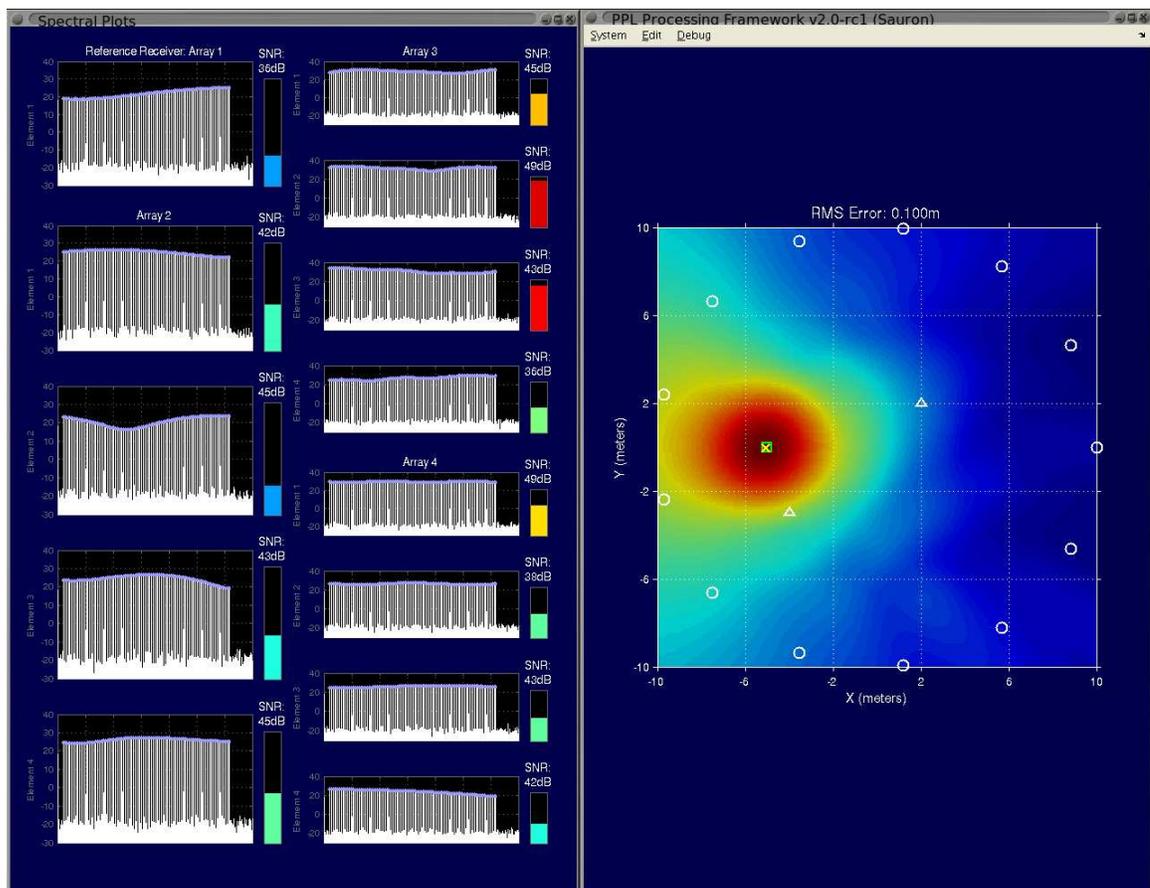


Figure 3.5: Graphical User Interface

Chapter 4

Synchronization Problems

The σ ART algorithm that was introduced earlier was described with the simplifying assumption of perfect synchronization. More specifically this introduction assumed that the transmit DAC and receive ADCs have sampling clocks at exactly the same frequency, and that the signals at the receive antennas are recorded at exactly the same time. For practical reasons neither of these conditions is true.

4.1 Sample Clock Drift

For the precision location problem our transmitter needs to be wireless, so it can not be connected directly to the sampling clock that drives the receive ADCs. That is not the only way to share a common frequency reference however; a master frequency beacon could be transmitted over the air for the transmitter to use. Multipath causes problems with this approach however, as it is possible for the transmitter to be at a location where the direct path and a reflected path happen to superimpose destructively resulting in signal loss. This is unacceptable for our application. Transmitting a single frequency over the air has the problem just discussed, but it may be possible to transmit a more complicated signal, such as a multicarrier signal, and assure that there will always be frequency content for the transmitter to lock onto. Such approaches require that our transmitter has the capability of receiving such a signal and processing it, which would add to the hardware requirements. In this section we will discuss why the lack of a common frequency reference is a problem,

and the solution that we found which required relatively little modification of our hardware requirements.

4.1.1 Sample Clock Drift Problem

We must first more fully understand the problem that unsynchronized sample clocks causes. In the current hardware configuration that our system uses, we have our software radio transmitter using one sample clock and our receiving ADCs running on another independent sample clock. The ideal sample clock frequency chosen for our system is 200 MHz. The transmitter's sample clock comes from a crystal oscillator (CSX750ACB200.000) with 50 ppm of frequency drift [6]. The receive ADCs have a sampling clock from an Agilent E4426B signal generator with less than 5 ppm of frequency drift [1]. The relative drift between the two clocks is therefore 55 ppm in the worst case.

Consider what happens when we generate our waveform with a different sampling frequency than with which we receive it. If our transmitter's sample clock frequency f' is faster or slower than the desired frequency f , then the signal generated will be accordingly compressed or stretched in time. Similarly on the receiving end, if the ADCs' sample clock is faster or slower than desired then the recorded signal will appear stretched or compressed accordingly. What ultimately determines the stretch/compression of the received signal is the difference in frequency between the transmit sample clock and the receive sample clock.

Since our concern is with the relative frequency drift between the two clocks, we will treat the receive sample clock as the reference frequency about which the transmit sample clock drifts. Our worst case difference in frequency leads to a ratio of

$$\frac{f'}{f} = 1 \pm 55 \cdot 10^{-6}. \quad (4.1)$$

This means that the duration of the symbols transmitted and recorded by the ADCs will be slightly different. The reference symbol duration of the receive ADCs

$$T = \frac{a}{f}, \quad (4.2)$$

where a is a proportionality constant. The period of the symbols from the transmitter DAC

is

$$T' = \frac{a}{f'}. \quad (4.3)$$

Hence in the worst case we have

$$\frac{f'}{f} = \frac{T}{T'} = 1 \pm 55 \cdot 10^{-6}. \quad (4.4)$$

For each symbol of data we use 8192 samples at 200 MHz so $T = 40.960\mu s$. The transmitted symbol duration is then at most

$$T' = T(1 + 55 \cdot 10^{-6}) = 40.962\mu s, \quad (4.5)$$

and at least

$$T' = T(1 - 55 \cdot 10^{-6}) = 40.958\mu s. \quad (4.6)$$

The difference in symbol duration is at most 2 ns. This difference is small enough not to affect the output of our FFT and thus can be ignored. This was proven in a previous thesis conducted and written in support of this project [3].

Assuming a transmit symbol and receive symbol start at the same time, the drift between the two sample clocks does not cause a problem since they end at the same time within 2 ns which we have accepted as negligible. The next pair of symbols however will not start at exactly the same time. These errors build up after several symbol periods. Figure 4.1 illustrates this behavior, though in a more extreme case where $\frac{T}{T'} = 1.2$, so the symbols become misaligned after just a few symbol periods. In our case the symbols can become

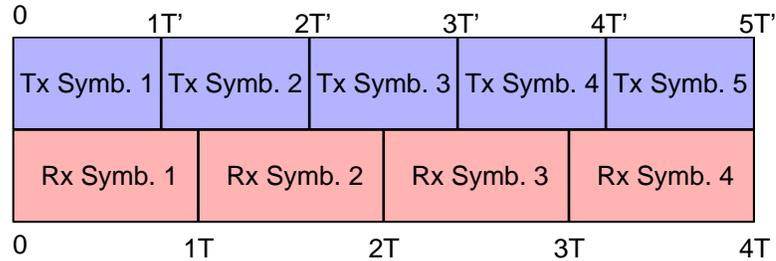


Figure 4.1: Temporal Symbol Alignment

misaligned by at most 2 ns every symbol period, so that our symbols can be completely

misaligned (offset by $\frac{T}{2}$) after as few as 10,240 symbol periods, which is 0.4194 seconds. The drift in the relative sample clock differences is caused by unpredictable factors such as temperature fluctuations. In Chapter 5 we will observe how long it takes for considerable drift to occur in our system.

Since the drift is unknown, we record symbols at our receive ADCs with unknown time offsets, while our hope for determining our position estimates is based upon the relative time of arrival of our signals. Is this really a problem? As explained in Section 2.2.4, σ ART is not affected by an arbitrary time offset applied to its input data as long as the same offset is applied to every antenna's recorded data. So our algorithm will still work in spite of our symbols being uniformly misaligned. However, this is only true when all of our receive ADCs capture symbols at the same time.

Consider again Figure 4.1. Suppose several receive ADCs (with a common sample clock) capture symbols during the second receive symbol window. They will not be aligned with a transmit receive symbol window, thus the received symbols will have some unknown time offset, but it will be uniform for all of the receive ADCs since it was captured at the same time. If the receive ADCs then capture symbols during the third receive symbol slot, there would be a new unknown time offset thus making it invalid to use symbols from receive capture 2 and receive capture 3 together in the same σ ART scan without correcting that time offset first somehow.

This imposes some undesirable restrictions on our system. If the symbols must all be captured simultaneously, then we must have a separate ADC for every antenna. The principle that σ ART exploits is direct path presence across several antennas, so generally speaking the more antennas the better. Our current system uses as many as 17 antennas. It would be very costly for us to implement a dedicated RF front-end and ADC for each antenna, not only monetarily, but also with respect to the time commitment for our team to assemble this hardware and in terms of the complexity of the system. A far more elegant solution would involve multiplexing our antennas in time using electronically driven RF switches and only a few ADCs and RF front-ends. This is the idea of using several "elements" on a few "arrays" discussed in Section 3.3 and shown in Figure 3.4. A solution to the problem imposed by the sample clock drift, and its implications in a time multiplexed

symbol acquisition based system, needed to be found for us to get around these undesirable hardware requirements.

4.1.2 Reference Array Solution

The solution to this problem that was found allows us to use antenna switching at the cost of having one unswitched antenna with its own dedicated array that we will call the reference array. We can use this single array, which captures data at the same time as the other array, to track the sample clock drift so that its effect can be undone in software.

Consider a single antenna connected to an array which is sampled by an ADC without any other antennas or switching. Recall that the signal we wish to record is ideally

$$R(f) = \sum_{i=1}^k G(f)\alpha_i(f)e^{-j2\pi ft_i}, \quad (4.7)$$

where we have i as an index of k signal sources. But we have an unknown time delay $\tau(t)$ applied to the signal that varies over time due to the drift between the sample clock of the transmitter and receiver. This means that the signal received at the antenna is actually

$$R'(f) = R(f)e^{-j2\pi f\tau(t)}. \quad (4.8)$$

If we made successive data captures from this array, how would the recorded symbols differ? If we assume the data captures were taken over a short enough period of time, then we can assume $R(f)$ does not change. This is because $R(f)$ is determined by the channel itself. If we assume the channel has not changed in a small period of time (our transmitter, any reflectors, or anything else that can affect our signal has not moved) then $R(f)$ remains the same. The only way our signal does change is due to the undesirable time delay term.

Suppose we have two captured data symbols at times t_0 and t_1 that have been processed with the FFT such that we now have vectors of the complex amplitudes of each of our carriers

$$R'_0(n\Delta f) = R(n\Delta f)e^{-j2\pi n\Delta f\tau(t_0)}, \quad (4.9)$$

$$R'_1(n\Delta f) = R(n\Delta f)e^{-j2\pi n\Delta f\tau(t_1)}. \quad (4.10)$$

If we divide the first symbol's data by the second we get

$$\frac{R'_0(n\Delta f)}{R'_1(n\Delta f)} = \frac{R(n\Delta f)e^{-j2\pi n\Delta f\tau(t_0)}}{R(n\Delta f)e^{-j2\pi n\Delta f\tau(t_1)}}. \quad (4.11)$$

Canceling $R(n\Delta f)$ then

$$\frac{R'_0(n\Delta f)}{R'_1(n\Delta f)} = \frac{e^{-j2\pi n\Delta f\tau(t_0)}}{e^{-j2\pi n\Delta f\tau(t_1)}}. \quad (4.12)$$

We can invert the denominator and multiply it by the numerator;

$$\frac{R'_0(n\Delta f)}{R'_1(n\Delta f)} = e^{-j2\pi n\Delta f\tau(t_0)}e^{-(-j2\pi n\Delta f\tau(t_1))} = e^{-j2\pi n\Delta f(\tau(t_0)-\tau(t_1))}. \quad (4.13)$$

We obtain a complex exponential whose frequency is solely dependent on the difference between the two time delays. If this frequency is determined then we can solve for $\tau(t_0) - \tau(t_1)$. This is a perfect application for the single pole DSS frequency estimator described in Section 2.3.1, since we are guaranteed the presence of only one complex exponential component.

Once $\tau(t_0) - \tau(t_1)$ is determined, we can correct the undesired time delays on our input data. We do not need to correct the first symbol's data since we are using that as the reference, but for the second symbol let

$$\widetilde{R}'_1(n\Delta f) = R'_1(n\Delta f)e^{-j2\pi n\Delta f(-\tau(t_0)+\tau(t_1))}, \quad (4.14)$$

$$\widetilde{R}'_1(n\Delta f) = R'_0(n\Delta f). \quad (4.15)$$

making the second symbol equivalent to the first, undoing the effects of sample clock drift.

So far we have assumed that there is no noise in our received signal, which would certainly not be true of real data. What we need is for our noise to be small enough relative to our signal whose frequency we need to estimate, based on the performance of the DSS single pole frequency estimator.

To determine the performance of the DSS single pole frequency estimator more precisely, a simulation was conducted. The data analyzed was a synthetic multicarrier signal with 103 carriers and some time delay trend on it, that is a phase shift $e^{-j2\pi n\Delta f t_0}$, plus white Gaussian noise. For varying amounts of bandwidth and signal to noise ratio, the time delays were chosen at random and then estimated. For each combination of bandwidth

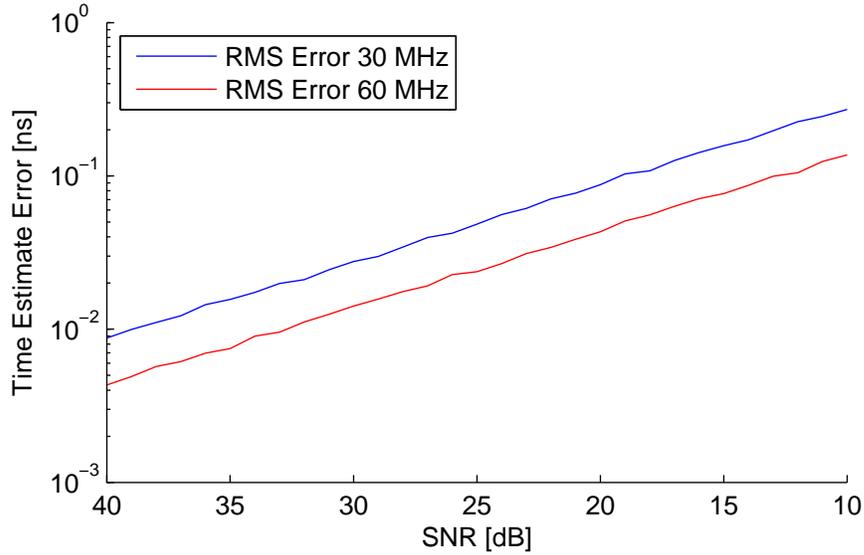


Figure 4.2: Time Estimate Error from DSS Single Pole Estimator vs. Signal to Noise Ratio

and signal to noise ratio, solutions were estimated 1000 times. The accuracy of these estimates are shown in Figure 4.2. We can see that for better signal to noise ratios that the performance improves as expected. Furthermore at 60 MHz bandwidth the time estimate error is approximately one half than that at 30 MHz. This coincides with analyses previously conducted in this project.

We are now able to determine the value of the random function $\tau(t)$ at any time that our unswitched reference array captures a symbol. If another antenna on another array (whether being switched or not) captures a symbol at the same time as the reference array, it also is subjected to a delay of $\tau(t)$, but since we have estimated that time delay with the reference array, we can undo it and generate valid data to use with our σ ART algorithm.

So, assuming our signal to noise ratio is large enough, we now have a means of tracking the values of $\tau(t)$ relative to some arbitrary initial value $\tau(t_0)$. This accomplishes our goal, since the σ ART algorithm can have an arbitrary global time offset applied to its input data without affecting its operation.

This technique solves the problem of sample clock drift with minimal hardware requirements. In Chapter 5 we will explore the practical implementation and performance of this technique.

4.2 Constant Time Offset Between Arrays

There is another significant synchronization issue that exists due to the behavior of our digital sampling hardware. In our system we have several DAQs running in parallel, each with an ADC sampling a signal for a different array. Our DAQs are currently co-located and run off of a common sample clock, so sample clock drift between them is not an issue. When the DAQs are first powered on, a signal is sent from one of the DAQs, the master, to indicate to the others the beginning of the symbol window. Unfortunately this alignment is only guaranteed to an accuracy of about 20 ns, due to hardware constraints. This means that if the same signal was sampled by the ADCs on two different DAQs, one may have a time offset relative to the other by as much as 20 ns. This behavior is illustrated in Figure 4.3. These time offsets are guaranteed to be fixed however, until the DAQs are powered on and

ADC 1	Symbol 1	Symbol 2	Symbol 3 ...
ADC 2	Symbol 1	Symbol 2	Symbol 3 ...
ADC 3	Symbol 1	Symbol 2	Symbol 3 ...
ADC 4	Symbol 1	Symbol 2	Symbol 3 ...
ADC 5	Symbol 1	Symbol 2	Symbol 3 ...

Figure 4.3: Inaccurate Synchronization Between ADCs

off again.

This is unacceptable for σ ART processing. Signals received by antennas from one array would have different time offset from another array. With σ ART it is acceptable to have a global time offset on the data for all receive antennas, not different time offsets for different antennas. Thus a solution to this problem was needed.

The solution that we implemented involved a synchronization procedure during system start-up, in which the different ADCs sample a reference signal. We have the different ADCs sample the reference signal during the same symbol cycle, then analyze the recorded data to determine the amount of unwanted time delay for each ADC so it can be later undone on all subsequent symbols.

The reference signal is another multicarrier signal. We use the baseband components of a second transmitter unit, creating a baseband signal that can be sampled by each of the ADCs after being passed through a splitter. These signals are then analyzed to determine the time offsets between them by converting them into the frequency domain and using the DSS single pole estimator, like in the sample clock drift solution. Again, since σ ART is agnostic to a global time offset we only need to undo the relative time offsets from the different ADCs.

So, our reference multicarrier $G(f)$ is generated, passed through a splitter, and received by the different ADCs. We'll show two in this example.

$$R(f)_0 = G(f)\alpha e^{-j2\pi f\tau_0} \quad (4.16)$$

$$R(f)_1 = G(f)\alpha e^{-j2\pi f\tau_1} \quad (4.17)$$

Any attenuation α in the path should be the same for all ADCs since the paths are symmetrical. Each ADC has a different undesired time delay τ on its received signal. When we divide the second signal by the first

$$\frac{R(f)_1}{R(f)_0} = \frac{G(f)\alpha e^{-j2\pi f\tau_1}}{G(f)\alpha e^{-j2\pi f\tau_0}} = e^{-j2\pi f(\tau_1 - \tau_0)}. \quad (4.18)$$

We obtain a complex exponential with a periodicity determined by the time difference $\tau_1 - \tau_0$. We can analyze this signal using the DSS Single Pole Frequency Estimator to determine $\tau_1 - \tau_0$. The time offset for each ADC can be determined relative to the first. These values are recorded and used throughout the session that the system is active to repair incoming data. Since this is a completely cabled solution, the signal to noise ratio of the received data will be much larger than the signals received from over the air, so the time estimates should be much more accurate than in the sample clock drift solution.

So our system must be able to operate in two modes, a synchronization mode in which the ADCs sample the reference signal to determine time offsets, and the normal mode in which the ADCs sample data from the receive antennas. In order to be able to switch modes easily we use a multiplexer to switch which signals are passed to the analog to digital converters. The configuration for this scheme is shown in Figure 4.4. In this example five

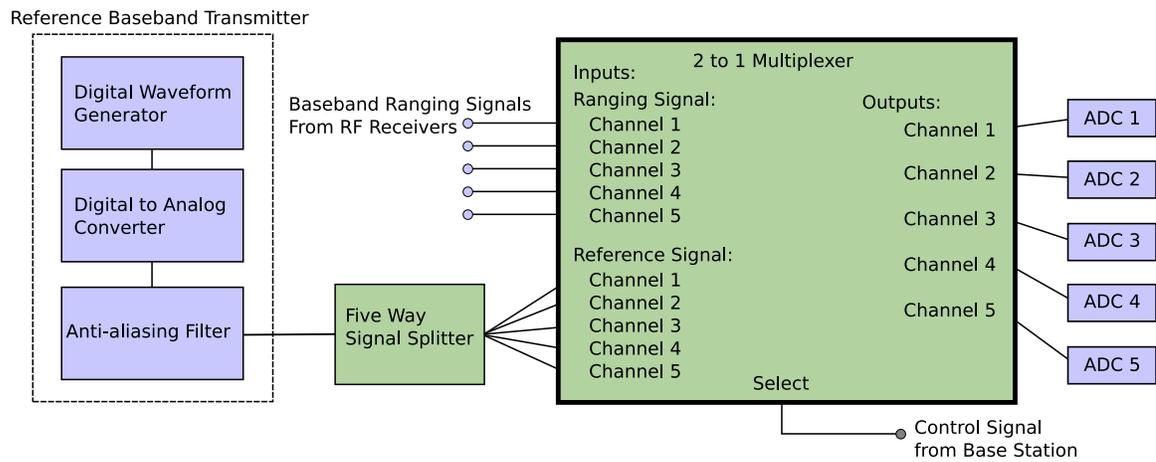


Figure 4.4: Configuration for Synchronization Scheme

data acquisition units are used so we have a five channel 2 to 1 multiplexer. The current mode is selected by a control signal from the Base Station Computer.

Chapter 5

Performance

5.1 Performance of Sample Clock Drift Tracking

In this section we will discuss the performance of the reference array solution to the sample clock drift problem. Initially the basis for the proposed technique of drift tracking was verified, then the drift compensated data was tested in σ ART processing.

5.1.1 Drift Tracking Proof of Concept

On June 16, 2006 a test of the Precision Personnel Location system was conducted. This test took place on the third floor of the Worcester Polytechnic Institute Electrical and Computer Engineering building, Atwater Kent (AK). Specifically this test was conducted in and around the laboratory AK317a, shown in Figure 5.1. This location was chosen as a high multipath indoor environment. The most challenging feature of this environment are the steel studs in the walls which are space 16 inches apart, shown in Figure 5.2. Also contributing to the difficulty there is: metal equipment and cabinetry, metal shelving, and a corrugated metal roof in testing area. These walls limit the penetration of our signal as well as reflect it.

For this test we used one transmitter which was completely free running with its own sample clock and which was placed in various location inside the room. We also used three arrays, each with one unswitched antenna and one ADC. Figure 5.3 shows the layout for



Figure 5.1: AK317a Laboratory

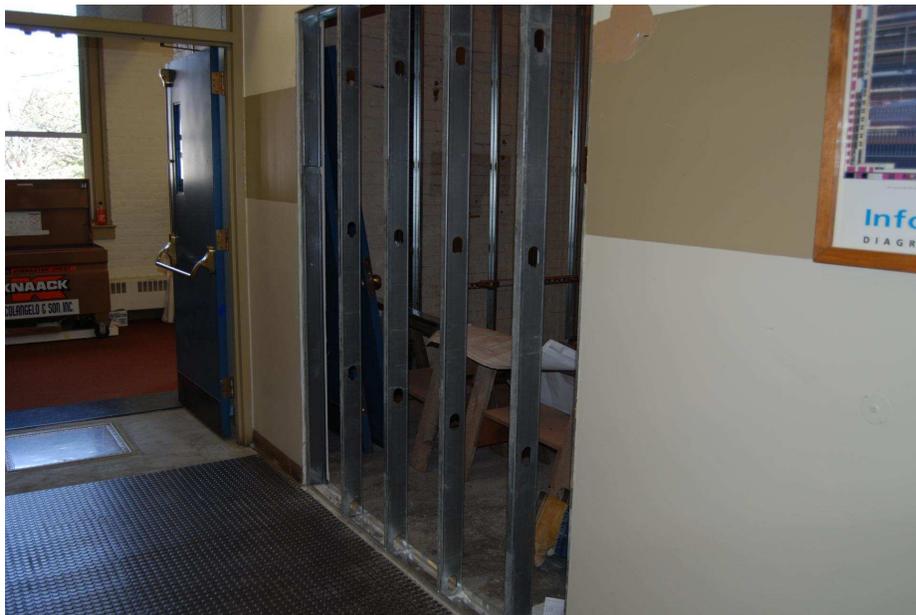


Figure 5.2: Steel Studded walls in Atwater Kent

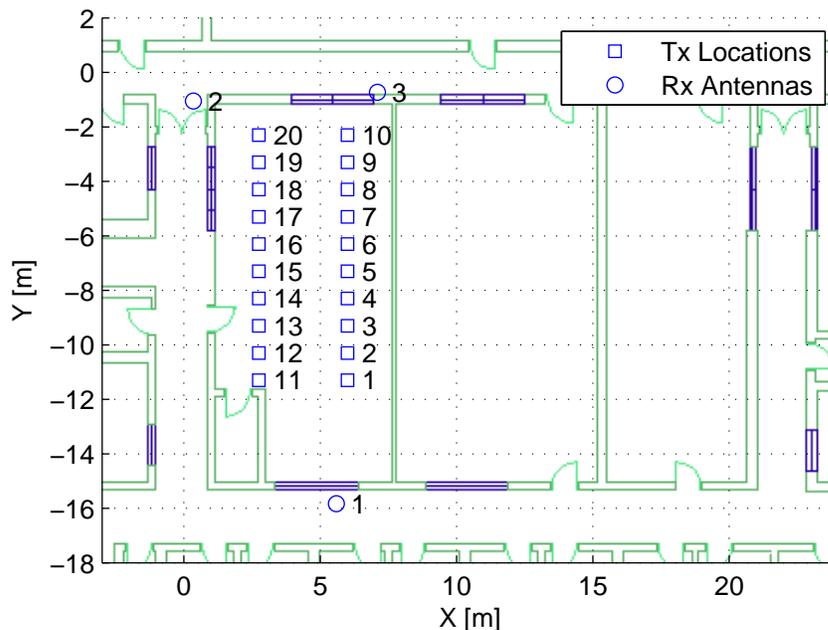


Figure 5.3: AK317a Test Layout 6/16/2006

this test, displaying the floor plan as well as the locations of the receive antennas and the locations of the transmit antenna at which we recorded data. All transmit and receive antenna locations for this test, and all of our test are surveyed manually. Automated receive antenna positioning is currently not a capability of our system. The multicarrier signal used for this test consisted of 103 carriers from 420-450 MHz (30 MHz span) at a power of -12 dBm for each of the carriers.

As discussed in Chapter 4, the reference array approach for using σ ART involves using one unswitched reference array and several other switched arrays. In this test we had three arrays that could have been reference arrays since they were all unswitched, and no switched arrays. This was a convenient configuration for the purposes of testing our ability to perform drift tracking. In theory we should have been able to track the drift at any one of the three arrays, and the function $\tau(t)$ that we measured should have been the same for all three.

For each transmitter location we captured 128 symbols of data at each array simultaneously. It is noteworthy that the symbols that were captured were not contiguous in

time. The sampling hardware keeps a count of how many symbol times have elapsed since the device was powered on. This is known as the global symbol index. Due to hardware constraints, the actual recording of symbols was quite sparse, capturing one symbol approximately every 2000-3000 symbol times. This means the actual data capture spanned almost 20 seconds. This was acceptable because the channel remained constant during the capture; we were not attempting to perform any real-time tracking for this test.

For each array we attempted to track the time offsets due to sample clock drift, $\tau(t)$. The first symbol from each array was used as a reference where we assumed there was no offset, $\tau(t_0) = 0$. Then for each subsequent symbol the single pole DSS estimator was used in order to compute $\tau(t) - \tau(t_0)$ for each symbol. Figure 5.4 shows the outcome of this

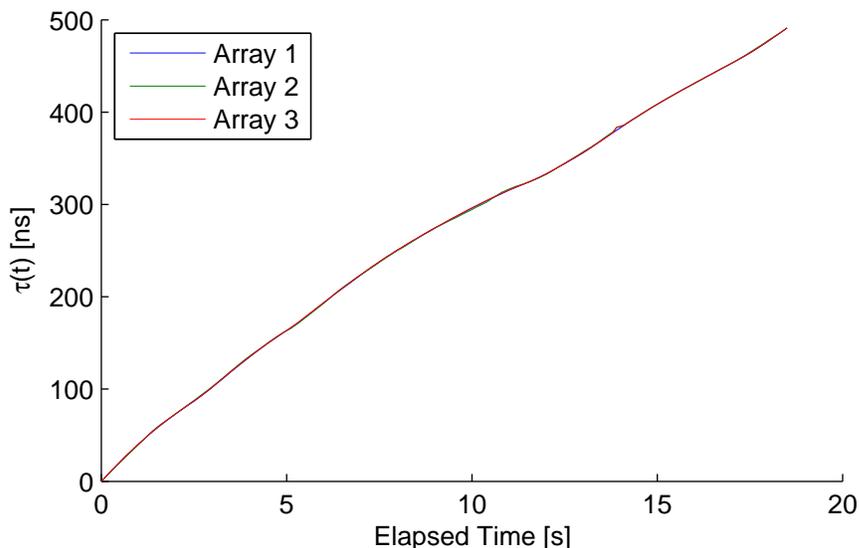


Figure 5.4: Tracking $\tau(t)$ at Transmitter Location 2 6/16/2006

analysis for transmitter location 2. As desired, the function $\tau(t)$ estimated at each of the three arrays was approximately equal. This implies that if we calculate the drift using one unswitched reference array, we can use those values to undo the effects of sample clock drift at other arrays. It was also an interesting that the curves found were continuous. The drift of the sample clocks were therefore relatively slow. On the other hand, after a few seconds we can see that the time offsets due to sample clock drift have already climbed to hundreds

of nanoseconds, whereas our tolerance for these time offsets is a standard deviation of 0.5 ns. Leaving these effects uncorrected is therefore clearly unacceptable.

We do not have direct access to the true $\tau(t)$ of course, but we can get an idea of how accurately we are tracking it by considering how close the three tracking curves are to each other. We'll consider the mean of the three curves as the true $\tau(t)$ and the difference between the curves from the mean to be the error, shown in Figure 5.5. In this case the standard

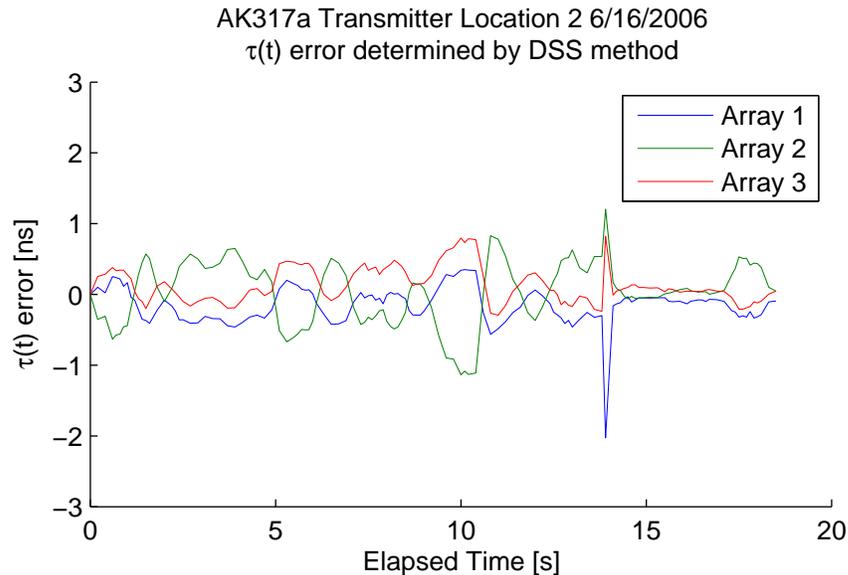


Figure 5.5: $\tau(t)$ Error at Transmitter Location 2 6/16/2006

deviation of the three error curves was 0.3461 ns. This is smaller than the maximum desired standard deviation of 0.5 ns, established in Section 2.2.4. Also note that for transmitter location 2 the estimated signal to noise ratio of the three arrays were respectively 33 dB, 33 dB and 39 dB.

These results are worse than predicted by Figure 4.2. According to the figure, for 30 MHz and signal to noise ratios of 30-40 dB the standard deviation of our time offset estimates should have been less than 0.1 ns. This is most likely due to the simulation being ideal, in the sense that the noise was uniform across the band. Also there was no multipath introduced and no attenuation as a function of frequency (to simulate hardware magnitude responses). With the real signals received on the third floor of Atwater Kent, there is high multipath.

At some frequencies the signals from the different sources may add constructively, and at others they may add destructively. This coupled with real hardware responses results in varying signal to noise ratios for the different carriers. For this reason, the simulation is considered a general guide of the DSS algorithm's performance trends with bandwidth and signal to noise ratio, but not necessarily an expected measure of performance in our real system.

The above referenced transmitter location 2 yielded one example of the drift tracking performance for this test. The standard deviation for the rest of the transmitter locations is shown in Figure 5.6. In most cases the standard deviation is under 2 ns. There is one case

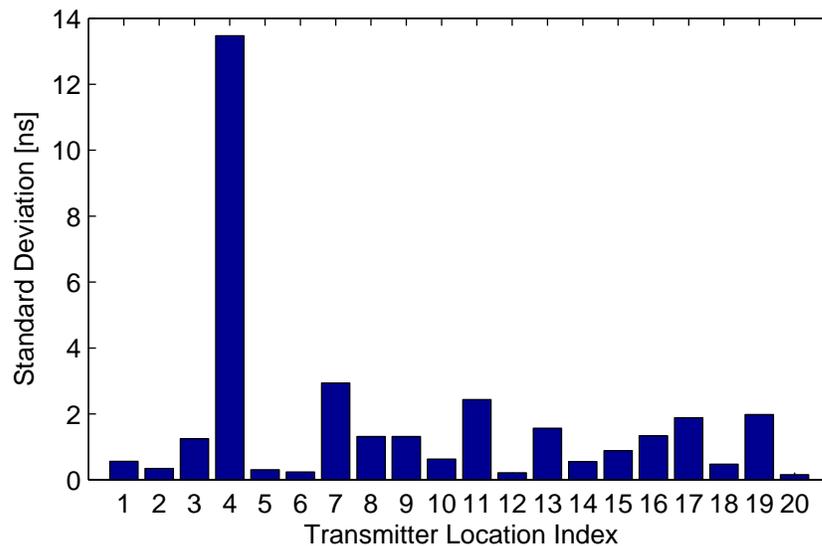


Figure 5.6: $\tau(t)$ Estimate Standard Deviation for 20 Transmitter Locations 6/16/2006

in particular, that of transmitter location 4, where the standard deviation is quite large. We look more closely at what happened in this case in Figure 5.7. Clearly Array 2 had large errors in this case for much of the capture time. The estimated signal to noise ratio of the three arrays for this location were respectively 34 dB, 31 dB and 41 dB. The 31 dB signal to noise ratio that Array 2 experienced at this location was relatively low. It was the lowest that Array 2 had for the entire test.

The 6/16/2006 test was considered an overall success for proving the concept of our

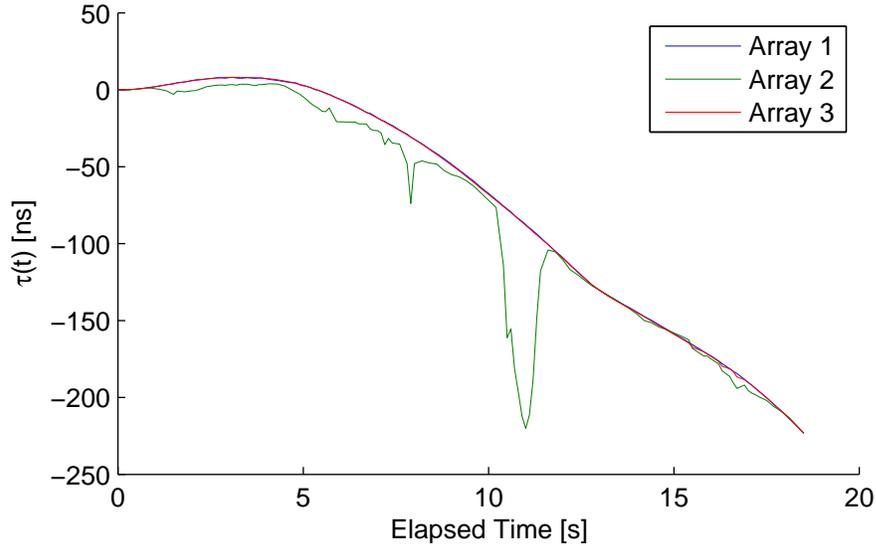


Figure 5.7: Tracking $\tau(t)$ at Transmitter Location 4 6/16/2006

sample clock drift tracking approach. However, the overall standard deviation of the drift tracking was 1.69 ns, which is significantly larger than our desired 0.5 ns. As expected from the theory, the threat to our ability to track the sample clock drift is noise. For this reason, during subsequent tests attempts were made to place the antenna for the reference array in a location that would have the best signal to noise ratio possible. The other way to improve our drift tracking performance however, is to increase our signal bandwidth.

5.1.2 Drift Tracking Performance in a 60 MHz System

We continued to use our 30 MHz location system for several months after the introduction of the drift tracking approach. In the fall of 2006 we were allotted more bandwidth by the Federal Communications Commission, 60 MHz from 410-470 MHz. As mentioned in Section 2.2, the more bandwidth we have the better we expect σ_{ART} to perform.

Once our system was modified to use 60 MHz, we performed another test on the third floor of Atwater Kent on 3/6/2007. Figure 5.8 shows the layout for this test. For this test we attempted to locate the transmitter at 20 locations inside of AK317a. We used five arrays: one reference array with one antenna, a second array with three antennas, and three more

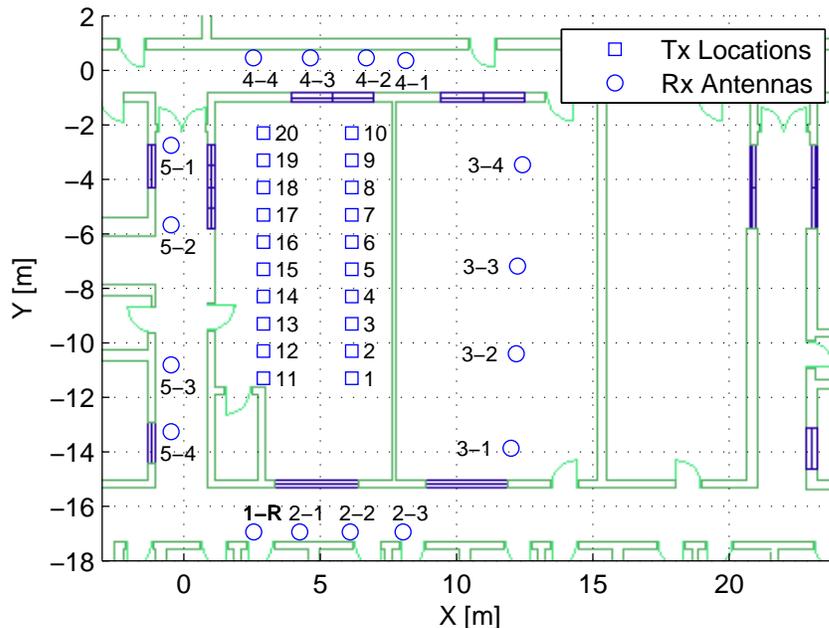


Figure 5.8: AK317a Test Layout 3/6/2007

arrays with four antennas each. The antennas in Figure 5.8 are labeled by array number and element number. The label 1-R indicates the antenna on the reference array. We used the same signal power level for this test as the 6/16/2006 test, -12 dBm per carrier.

Now that we had twice as much bandwidth as the 6/16/2006 test, we expected our drift tracking to be more accurate, since a DSS single pole frequency estimator performs better with more bandwidth. In the 6/16/2006 test we were able to observe the performance of our drift tracking because we had three unswitched arrays, any one of which was qualified to be a reference array. For this test however we only had one reference array, so we could not simply track the drift on multiple arrays and observe how close the $\tau(t)$ curves were as was done previously.

Consider our antenna switching scheme used during this test. Each switched array has up to four switched antennas connected to it, referred to as elements 1-4. We cycle through the various antenna elements capturing symbols for each. So when we are capturing symbols for element 1, symbols will be captured simultaneously at the reference array's antenna as well as element 1 antennas on each of the other arrays.

Figure 5.9 shows how symbols were captured from the four elements for the first transmitter location from this test. For this data capture we recorded 64 symbols from each

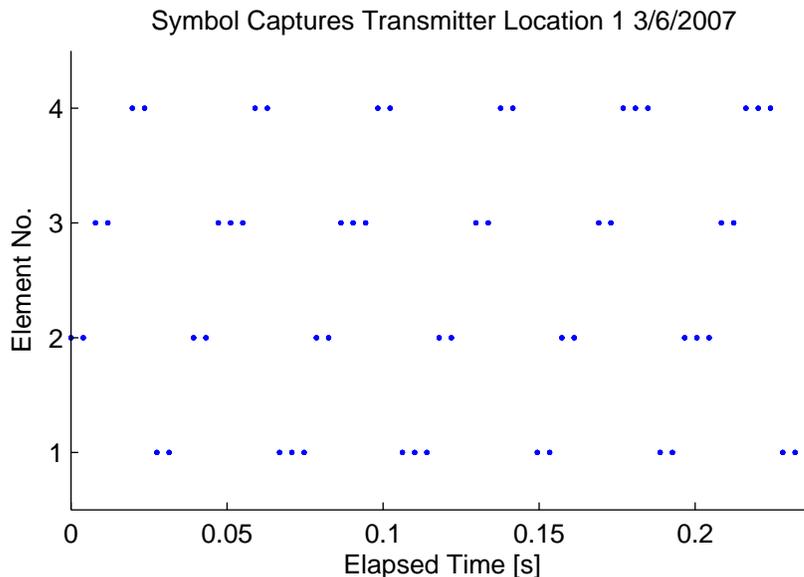


Figure 5.9: Symbol Captures on 4 Elements 3/6/2007

element. Our switching scheme cycles through the different elements for durations of approximately 10 ms, during which time we capture 2-3 symbols. After several cycles (not all shown in Figure 5.9) our quota of 64 symbols was met in this case. We also captured symbols at a faster rate during this test than the previous 6/16/2006 test. We captured symbols approximately every 96 symbol times. A full capture of 64 symbols from all four elements took approximately 1 second in this case.

The reference array, as it is not switched, is qualified to track sample clock drift for every element, unlike the other arrays. This is not true for the other antennas, so we can not directly perform the same analysis as we did in the 6/16/2006 test where we compare the tracking of $\tau(t)$ on the independent unswitched antennas. Instead we must make an abstraction. Suppose we can throw away all symbols from all elements but element 1, as shown in Figure 5.10. This means that the remaining data consists of only symbols captured simultaneously on all antennas, effectively making them act as unswitched antennas. Now the data can be analyzed in the same way that we analyzed the data from the 6/16/2006

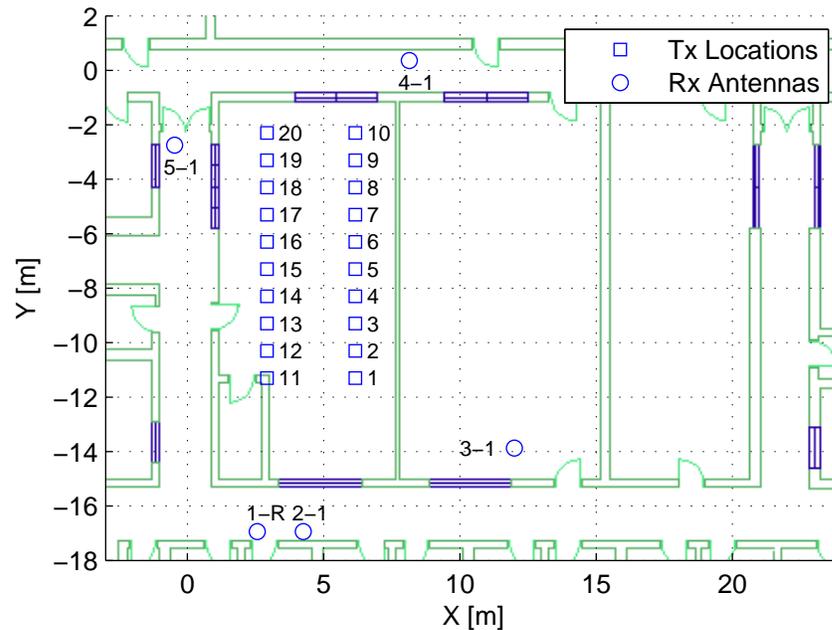


Figure 5.10: AK317a Test Layout 3/6/2007 First Element Only

test, and allows us to establish how well the independent antennas track $\tau(t)$. This same analysis can also be repeated keeping only data from elements 2,3 and 4.

So, we'll consider first only symbols captured from the first elements at the first transmitter location. Using each antenna to track the drift independently we generate Figure 5.11. At this scale any differences between the drift tracking from the five different arrays is not visible. Figure 5.12 shows the test result wherein we plot the error of the different arrays where the error is again defined as the difference from the mean of the five curves. The error is generally less than 0.1 ns, which is very good. The standard deviation was calculated as 0.066 ns.

This analysis was also performed with elements 2,3 and 4 for this transmitter location, and the standard deviations were 0.052 ns, 0.224 ns and 0.110 ns respectively. On an even larger scale, all four elements at all twenty transmitter locations were analyzed in the same manner. The performance was very good, with an average standard deviation overall of 0.232 ns. Figure 5.13 shows the performance for each transmitter location and each element. Generally speaking the performance was very good, and much improved over the 6/16/2006

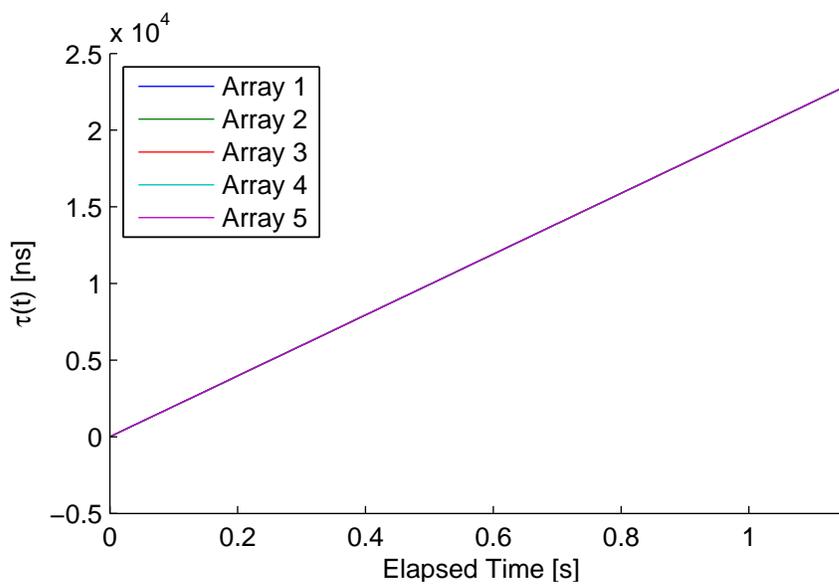


Figure 5.11: Tracking $\tau(t)$ at Transmitter Location 1 on Element 1 3/6/2007

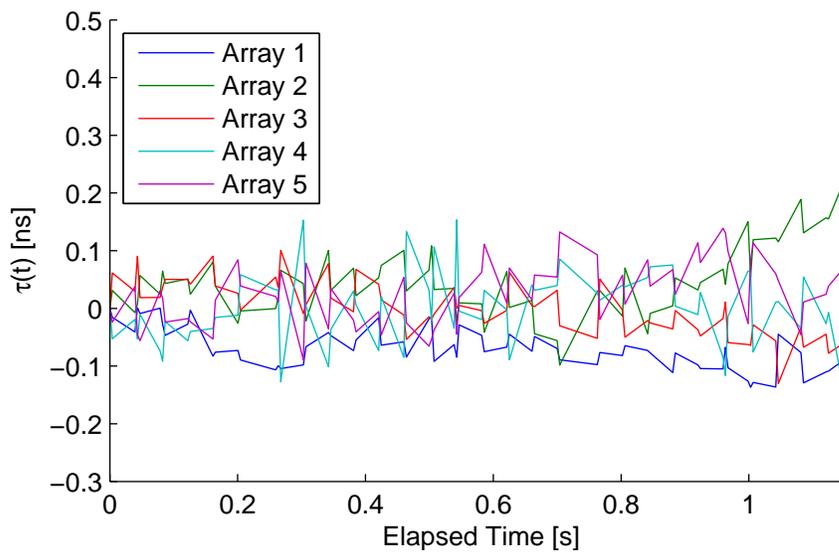


Figure 5.12: $\tau(t)$ Error from mean at Transmitter Location 1 on Element 1 3/6/2007

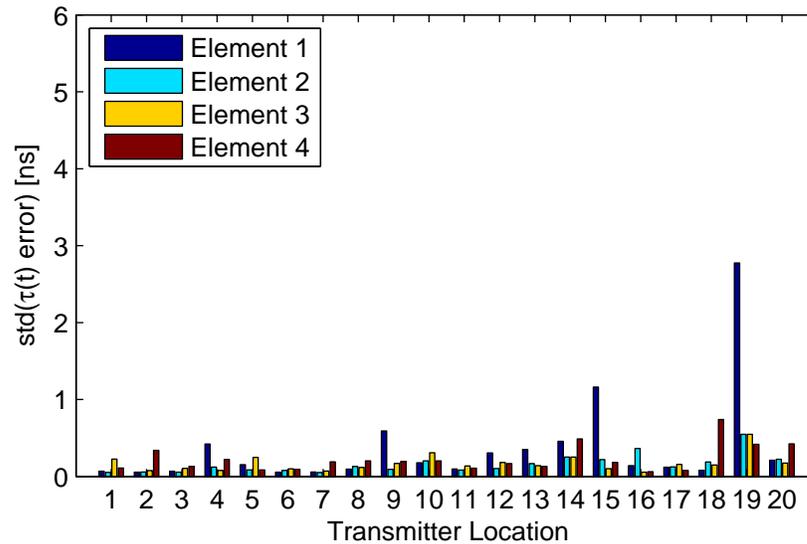


Figure 5.13: $\tau(t)$ Estimate Standard Deviation 60 MHz 3/6/2007

test.

The data from this test was also analyzed in the same manner using only the lower 30 MHz. This yielded a larger overall standard deviation of 0.458 ns. Figure 5.14 shows the performance for each transmitter location and each element in the 30 MHz case. We can see that the performance is better with 60 MHz than with 30 MHz as predicted. This also shows us that we are able to undo the effects of sample clock drift more and more accurately as we increase our bandwidth.

Recalling our general rule of thumb from Section 2.2.4, we determined that the standard deviation of our timing errors should be less than 0.5 ns. We can conclude from this analysis that this condition has been met. Even the 30 MHz case generally performs better than this constraint; the 60 MHz case performing well beyond this. Thus we conclude that we have solved the synchronization problem imposed by sample clock drift with the reference array solution.

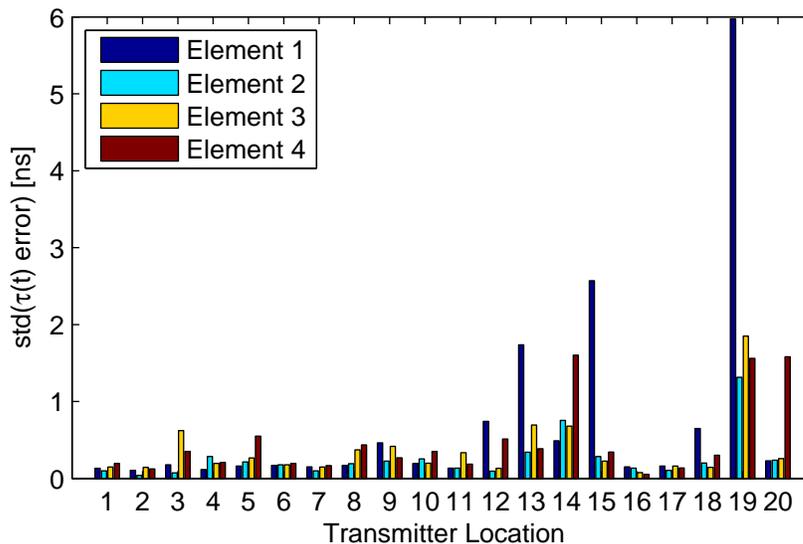


Figure 5.14: $\tau(t)$ Estimate Standard Deviation 30 MHz 3/6/2007

5.2 Array Synchronization Performance

In Section 4.2 we discussed our proposed scheme for determining the time offset between ADCs by using a reference transmitter. This transmitter's output is sent through a splitter and then connected directly to the ADCs (through a multiplexer).

We conducted a test where we used this approach in practice. The transmitted signal was a baseband multicarrier signal with 27 carriers spread between 10 MHz and 52 MHz. We captured the signal on five different ADCs simultaneously, and estimated the time offsets with the previously described method. The relative to the first ADC the remaining four had time offsets 10.2 ns, 15.1 ns, 15.2 ns and 15.2 ns.

What would the effects of these time offsets had been if we attempted to use this data in σ ART processing? We analyzed this data in a pseudo-simulation, where we pretend the data from the five ADCs came from five antennas symmetrically spaced in a circle 10 meters from the origin. If the data from the five ADCs was the same then the maximum of the σ ART metric should have been at the origin, since that is the one location equidistant from the antennas. Since we did not compensate for the time offsets between the ADCs, this was not the case, as shown in Figure 5.15 We see that there is an error of 1.90 meters, which

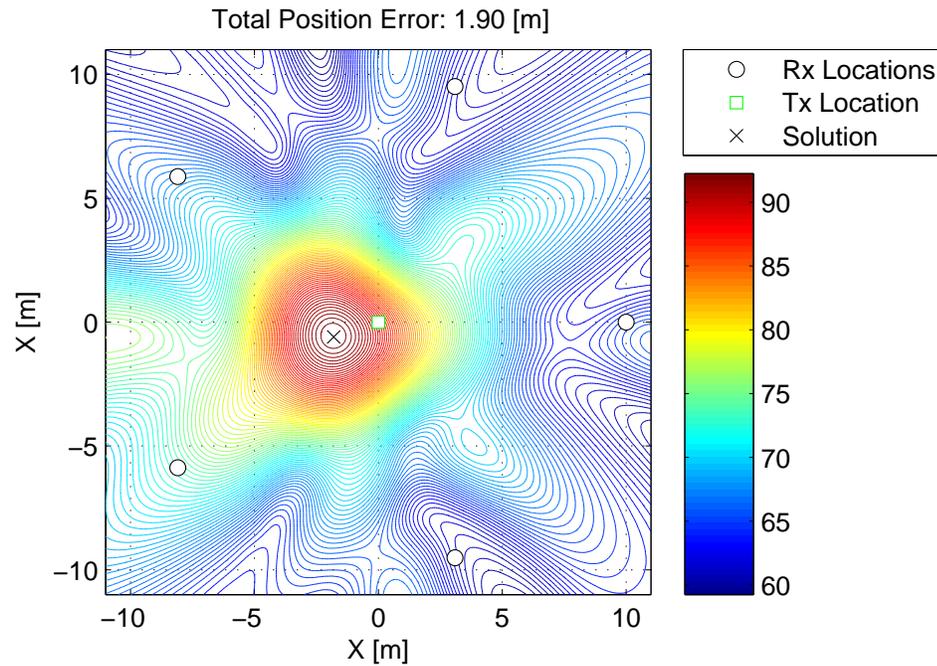
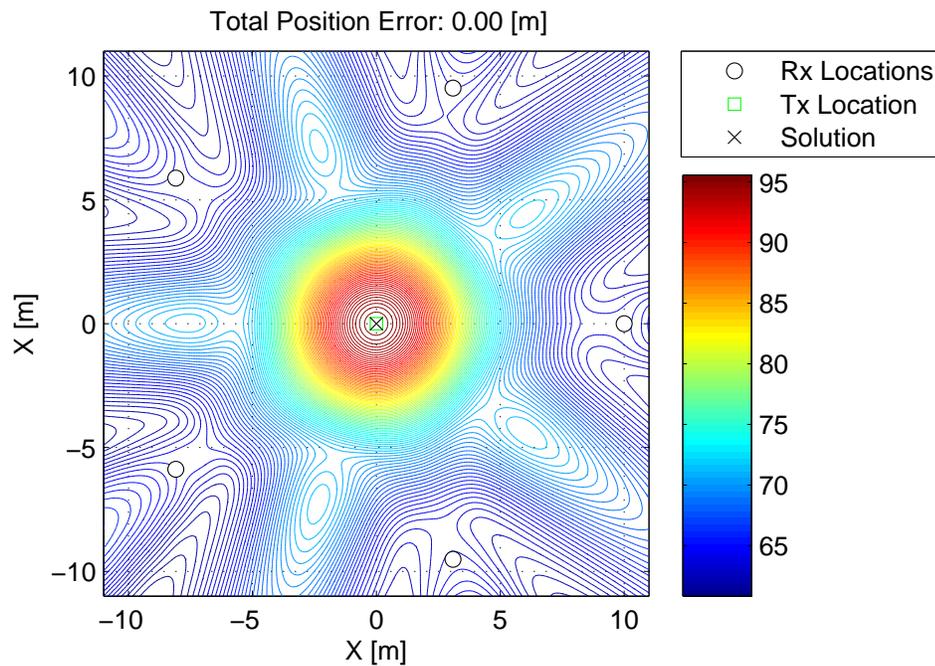
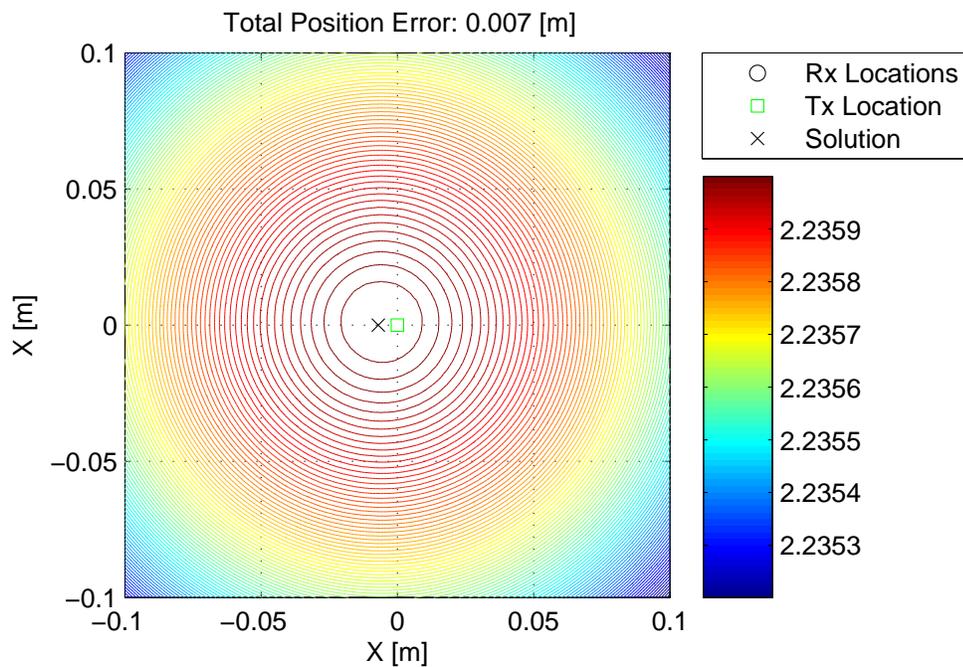


Figure 5.15: σ ART with ADC Time Offsets

is unacceptable. Since we were able to determine these time offsets however, we can undo them in software. Figure 5.16 shows the σ ART result for the same case after these time offsets are undone. We can see that the error goes down to approximately zero meters. This particular scan had a resolution of 10 cm, and the error is undetectable in this case. If we zoom in and scan a tighter region, we can see how large our error really is. Figure 5.17 shows the same case with a resolution of 1 mm. Our actual error is 7 mm, which is much better than our desired level of accuracy.

These tests let us observe the amount of error that σ ART is subject to when the ADCs are not synchronized properly. We also observed that the proposed method of determining the time offsets and undoing them works very accurately, meeting our needs.

Figure 5.16: σ ART with ADC Time Offsets UndoneFigure 5.17: σ ART with ADC Time Offsets Undone (Zoomed)

5.3 Positioning Performance

With our synchronization solutions in place, can we finally use σ ART to do positioning? Lets consider the results from the 3/6/2007 test with 60 MHz. Figure 5.18 shows the outcome of σ ART for the first transmitter location. For this analysis all analyses to follow

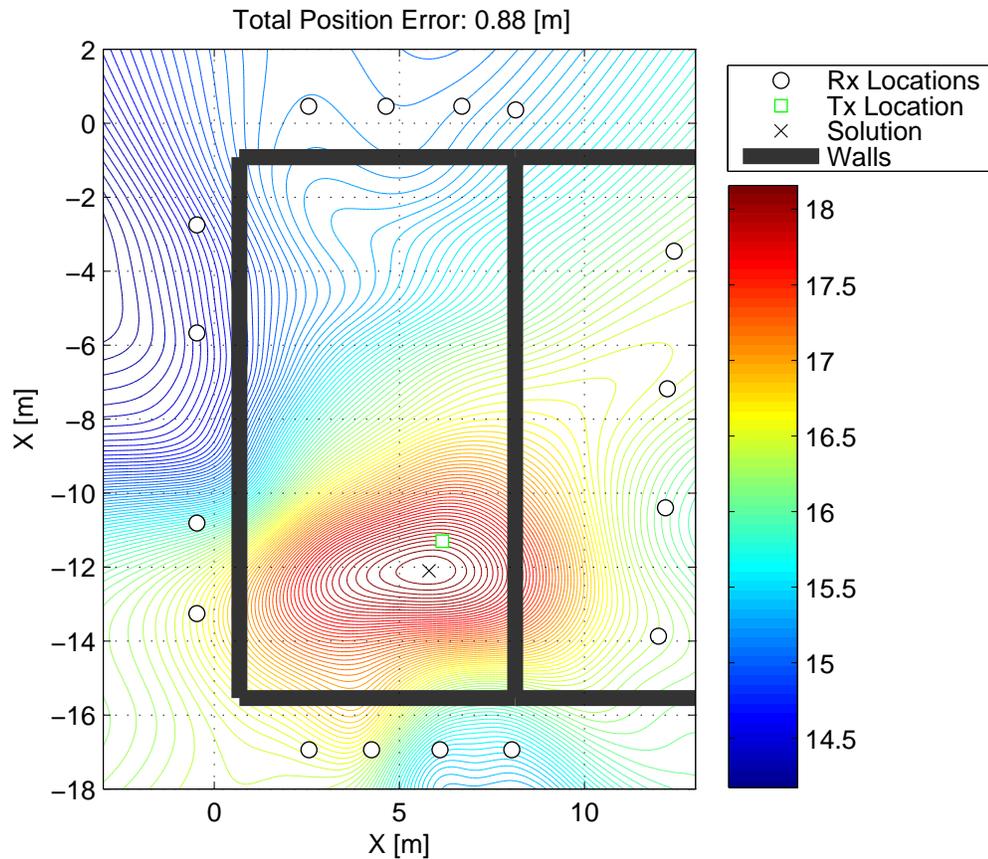


Figure 5.18: σ ART Result Transmitter Location 1 3/6/2007

a two dimensional scan in the horizontal plane was performed, where the spatial resolution of the σ ART scan was 1 meter, subsequently increased to 0.1 meters with interpolation. We can see that the maximum of the σ ART metric is 0.88 meters away from the true position. This result is not as accurate as we would like, but it does seem to indicate that our system is working overall.

Figure 5.19 shows the outcomes for all 20 transmitter locations in this test. The vectors

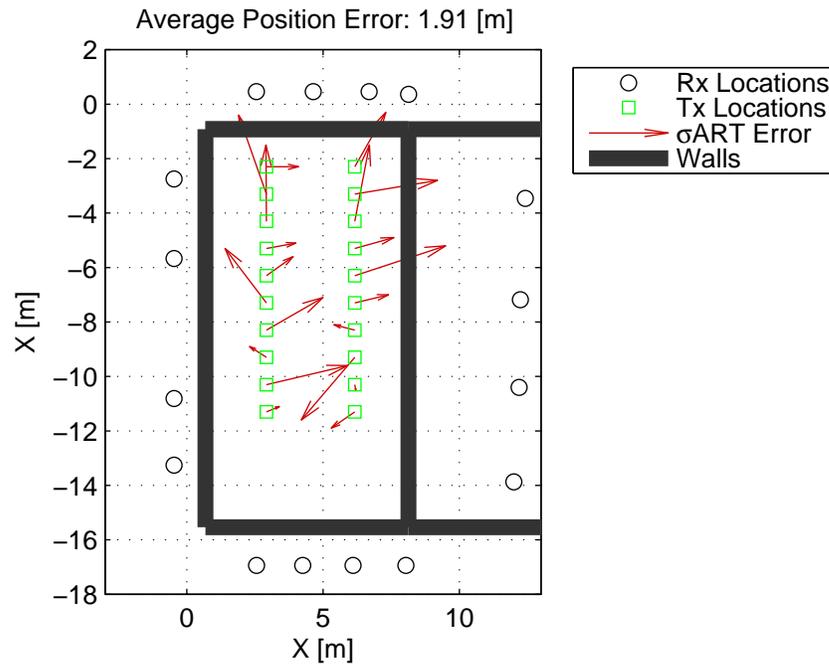


Figure 5.19: Position Error Vector Plot 3/6/2007

in the plot show where the algorithm concluded the solution was, relative to the true position. The average error was 1.91 meters. These solutions are not as accurate as they need to be ultimately, but actually were the best we had seen in that location to date. As we have discussed, the third floor of Atwater Kent is a very high multipath environment, which has always been extremely challenging. In some cases we have even been able to discern the contribution of large specular reflectors (such as metal-backed blackboards) by analyzing σ ART metric images.

Again, Figure 5.19 shows that the average error for this test was 1.91 meters using σ ART. These errors are on the same order as the error shown in Figure 5.15 of 1.90 meters, when the synchronization between arrays was not compensated. Thus the improvements from array synchronization are not obvious in this context. In the next section we will discuss improved positioning algorithms that will clearly reveal the improvement from array synchronization.

5.3.1 Improved Algorithms

In the course of the PPL project, effort is always being made to improve the performance of our algorithms. As of writing of this thesis, newer algorithms exist that improve upon σ ART. We will not discuss any details regarding these newer algorithms here, as they are beyond the scope of this thesis. It is important to note however, that like σ ART these newer algorithms use an exhaustive scanning approach. Also, these newer algorithms have the same synchronization requirements as σ ART, so the synchronization solutions discussed in this thesis are applicable with them as well. Using these newer algorithms, as well as varying some parameters associated with them, we have seen for this same 3/6/2007 test average errors as low as 0.72 meters, which is shown in Figure 5.20.

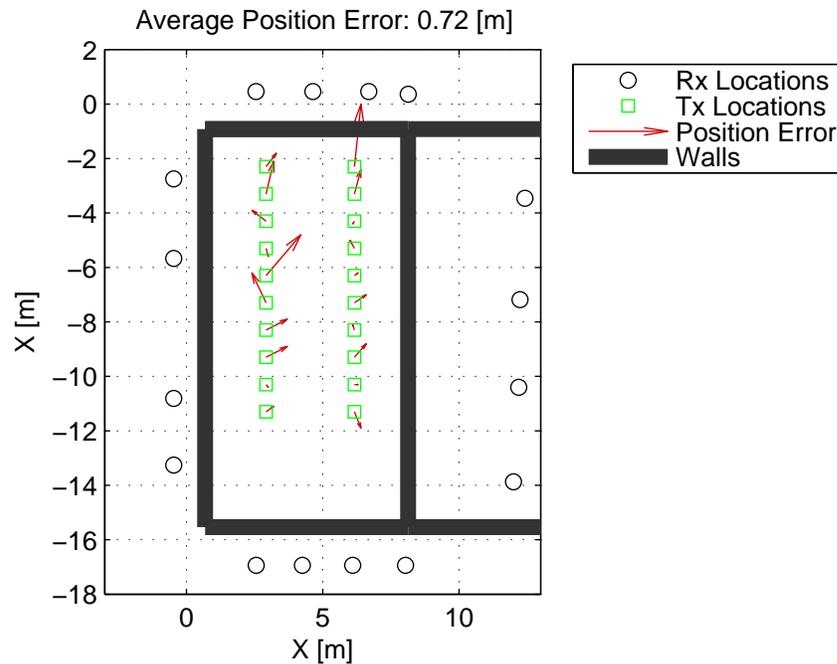


Figure 5.20: Position Error Vector Plot 3/6/2007 with Advanced Algorithm

5.3.2 Effects of Array Synchronization

With this improved result the effect of the constant time offsets between arrays can be shown here. Using the same case of the 3/6/2007 test, with the improved algorithm

that yielded 0.72 meters average error, the data was analyzed again. This time however the constant time offsets between arrays was not compensated for. This result is shown in Figure 5.21. We can see that the average error gets significantly larger, now 2.07 meters. This punctuates the need for the solution to this problem that has been discussed.

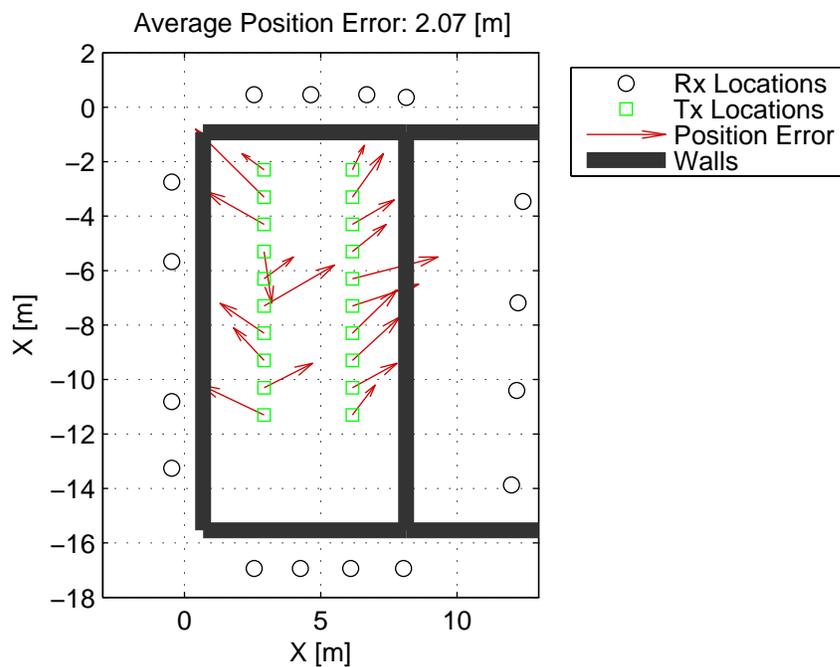


Figure 5.21: Position Error Vector Plot 3/6/2007 with Advanced Algorithm and No Array Synchronization

Chapter 6

Conclusion

This thesis explores solutions to synchronization issues that arise in the implementation of a precision indoor location system with a target goal of resolving the location of a radio transmitter to within 1 foot. By understanding the particular synchronization requirements of the σ ART algorithm we observed two problems impeding its usability.

The first was the issue caused by sample clock drift, which would have required us to have unreasonable hardware requirements had we not found a way to undo its effects. An elegant solution was found using a reference antenna to track the effects of the sample clock drift, that enabled us to use time multiplexed antennas to reduce our hardware complexity. The precision of the DSS single pole periodicity estimator allowed us to track the sample clock drift effects so they could then be removed from the data. This was a perfect application of this super-resolution spectral estimation technology since our theory guaranteed the presence of a single sinusoidally varying phase component in our frequency domain data.

We also solved the problem of constant time offsets between ADCs by introducing a synchronization procedure on the startup of our system. The constant time offsets between ADCs are estimated with the DSS single pole periodicity estimator. These offsets are recorded and then used to repair incoming data.

The σ ART algorithm was simulated to determine how synchronization errors affected its position estimates. We established that both of these synchronization solutions could be performed accurately enough for our desired specifications. Also since the DSS single

pole estimator solution was shown to improve with additional bandwidth, we know that if we move to a higher bandwidth system our synchronization techniques should become even more accurate.

6.1 Future Implications

The synchronization solutions described in this thesis were implemented to bring our current system to full functionality with σ ART processing. The fundamental question we are trying to answer first and foremost is whether we are able to perform indoor positioning to the desired degree of accuracy. Our current hardware configuration was chosen to let us answer that question without all of the requirements of a final system. One example of this is the fact that our ADCs are currently co-located and running from the same sample clock. In a final system, several fire trucks may arrive at a scene, each with several antennas mounted to it. There must then be at least one ADC on each truck to sample the data on its antennas.

This separation of ADCs would cause problems with the current scheme. If the ADCs do not all run from the same sample clock, then they will drift relative to each other. This takes the problem of sample clock drift correction to another level of difficulty. This problem has not yet been considered in depth. It may be possible to distribute a master frequency reference wirelessly so that the sample clocks for the different ADCs drift together. The signal for this frequency reference would need to be chosen appropriately so it does not succumb to problems from multipath.

The problem of constant time offsets between ADCs is also worsened in this new case. Since the ADCs are not co-located it is not possible to ensure they start their symbol acquisition windows at the same time. This issue has been temporarily bypassed in the current system by directly wiring a common signal to each of the ADCs so we could measure and correct the time offset. It would be undesirable and possibly impractical to have to connect cables between different trucks in an emergency situation. This is another significant synchronization issue that needs to be overcome in order to avoid additional burdens on the users.

Bibliography

- [1] AGILENT TECHNOLOGIES. Agilent ESG-A/AP and ESG-D/DP RF signal generators, 2007.
- [2] BARD, J., HAM, F., AND JONES, W. An algebraic solution to the time difference of arrival equations. *IEEE Transactions on Signal Processing* 47, 2 (Feb. 1999), 313–319.
- [3] BREEN, D. E. Characterization of multi-carrier locator performance. Master’s thesis, Worcester Polytechnic Institute, 2004.
- [4] CANTY, E. J. Six firefighters missing in blaze at vacant building. *Worcester Telegram and Gazette* (Dec. 1999).
- [5] CHEN, C.-T. *Linear System Theory and Design*. Oxford University Press, 1999.
- [6] CITIZEN AMERICA CORPORATION. Oscillator part numbering system, 2003.
- [7] DAVID CYGANSKI, JOHN ORR, W. R. M. A multi-carrier technique for precision geolocation for indoor/multipath environments. Tech. rep., Worcester Polytechnic Institute, Sept. 2003.
- [8] DAVID CYGANSKI, JOHN ORR, W. R. M. Performance of a precision indoor positioning system using a multi-carrier approach. Tech. rep., Worcester Polytechnic Institute, Jan. 2004.
- [9] DEMMEL, J. W. *Applied Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.

- [10] GOLUB, G. H., AND VAN LOAN, C. F. *Matrix Computations*. The Johns Hopkins University Press, 1996.
- [11] KENNEDY, T. J. Workshop on precision personnel location and tracking August 7-8, 2006. Tech. rep., Worcester Polytechnic Institute, Aug. 2006.
- [12] LACHAPELLE, G., KUUSNIEMI, H., DAO, D. T. H., MACGOUGAN, G., AND CANNON, M. E. HSGPS signal analysis and performance under various indoor conditions. Tech. rep., University of Calgary, Sept. 2003.
- [13] LATHI, B. P. *Signal Processing and Linear Systems*. Berkeley-Cambridge Press, 1998.
- [14] LLOYD N. TREFETHEN, D. B. I. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.
- [15] NANGLE, R. Worst building fire disaster in 21 years. *Worcester Telegram and Gazette* (Dec. 1999).
- [16] OPSHAUG, G. R., AND ENGE, P. GPS and UWB for indoor navigation. Tech. rep., Stanford University, Sept. 2001.
- [17] ORR, J., AND CYGANSKI, D. Firefighter and other emergency personnel tracking and location technology for incident response. Tech. rep., Worcester Polytechnic Institute, July 2001.
- [18] WRIGHT, C. R. *Multidimensional Direction of Arrival Performance Bounds and Optimization for Non-Stationary Noise*. PhD thesis, Worcester Polytechnic Institute, 1994.