VERY LOW FREQUENCY REMOTE SENSING OF THE LOWER IONOSPHERE

By

MICHAEL F. MITCHELL

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2015

For my father, Daniel L. Mitchell.

ACKNOWLEDGMENTS

The completion of this dissertation marks the end of a long and interesting Ph.D. adventure. A major component of the work described in this dissertation involved extensive travel to unusual places to deploy and upgrade receiver hardware, while simultaneously satisfying ridiculously stringent field-work requirements. I did not complete this journey alone, and I have a great many people to thank. First and foremost, I would like to thank my Ph.D. advisor, Dr. Robert Moore, who helped keep me on track over the years and who has provided sound wisdom and sage advice. I also thank the members of my Ph.D. committee: Dr. Jian Li, Dr. Vladimir Rakov, Dr. Martin Uman, and Dr. Steven Detweiler.

I offer special thanks to specific members of the Ionospheric Radio Laboratory: Dr. Shuji Fujimaru, Daniel Kotovosky, and Neal Dupree. Few things are more binding than sharing the difficulties of field work. These three students helped to build a successful research group, not only in terms of performance, but also in terms of character, and I hope they will forever be remembered for their efforts. To all the members of the research group not explicitly mentioned above, I thank you for your help both in the laboratory and in the field. I would also like to thank Dr. Evans Paschal for his help with receiver design and theory. His tutelage has helped to ensure that the current receiver designs are the best that they can be. I also extend my thanks to the members of the University of Florida department of Electrical and Computer Engineering staff, who proved to be particularly helpful over the years. And many thanks to all of the people, especially Norma and Doyle Traw and Jay Scrimshaw, who have helped with Florida field work during HAARP campaigns in Alaska.

Finally, I would like thank my family and Kym Mason for their unwavering support while I completed the arduous task of scaling “Ph.D. Mountain.” It has been a long time coming, and I am proud to have finally reached my goal after five long years.

The work presented in this dissertation is supported in part by the following grants and contracts awarded to the University of Florida: National Science Foundation grants PLR-1246275, ANT-0944639, and AGS-0940248, DARPA grant HR0011-10-1-0061, BAE Systems contract 709372 (prime: DARPA HR0011-09-C-0099), White Oak Technologies/Novetta Solutions Co. contract WHITEOAK2013 (prime: Air Force 11D1300S-AF-005 01), US Air Force grant FA9453-12-1-00246, and ONR grant N000141010909.

TABLE OF CONTENTS

page

LIST OF TABLES

LIST OF FIGURES

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

VERY LOW FREQUENCY REMOTE SENSING OF THE LOWER IONOSPHERE

By

Michael F. Mitchell

May 2015

Chair: Robert C. Moore
Major: Electrical and Computer Engineering

Very Low Frequency (VLF, 3–30 kHz) transmitters are located around the world with the primary purpose of communicating with submarines and other submerged assets. VLF signals propagate long distances (>10,000 km) around the globe with relatively low attenuation within the waveguide formed by the conducting Earth at one boundary and the conducting lower ionosphere (∼85 km altitude) at the other. The narrowband (∼200-Hz bandwidth) VLF signals radiated by these transmitters are used as "signals of opportunity" to remote-sense the electrical properties of the lower ionosphere both under ambient ionospheric conditions and during transient ionospheric disturbances (i.e., relatively short duration periods of changing ionospheric conductivity).

This dissertation presents significant improvements to the VLF remote sensing method. First, the specifications and designs for a new VLF receiver are discussed. Results are presented to demonstrate the successful implementation of a VLF receiver with high (>90 dB) linear dynamic range, high (>90 dB) spurious-free dynamic range (SFDR), and high (>90 dB) signal-to-noise and distortion ratio (SINAD).

Secondly, a new method for quantifying the "sum-of-modes" amplitude and phase of a received narrowband VLF transmission is presented and characterized. The performance of the new method is analyzed and compared with other contemporary narrowband VLF signal processing techniques. Under ambient conditions, the new method out-performs contemporary methods and is used to experimentally demonstrate the regular detection of frequency-dependent VLF signal propagation. It is determined, however, that the performance of the new method is sensitive to the

specific impulsive noise environment present at the receiver, highlighting the importance of using a high dynamic range VLF receiver.

The response of the new method to transient VLF signal perturbations is also evaluated and compared with the responses of other analysis techniques. Typically, the different techniques produce very similar *changes* in amplitude and *changes* in phase during transient events, and the observed differences are attributed to multi-mode VLF propagation. The new method excels in multi-mode propagation environments, such as the transition between daytime and nighttime.

Lastly, experimental observations are used to demonstrate that although the amplitudes and phases observed during transient events can exhibit frequency dependence, the percent change in amplitude and the change in phase are essentially constant with frequency. This result implies that, while the frequency dependence of the amplitude and phase contribute significantly to the analysis of ambient conditions, bandwidth-averaged amplitude and phase values are sufficient to fully characterize transient events.

CHAPTER 1
INTRODUCTION

This Ph.D. dissertation analyzes the excitation, propagation, detection, and processing of narrowband (200 Hz) very low frequency (VLF, 3–30 kHz) signals within the Earth-ionosphere waveguide. While the specific focus is on the signals produced by VLF transmitters operating in the ∼18–30 kHz band, a similar analysis will apply to any narrowband signal that propagates between the Earth and ionosphere. VLF waves propagate to great distances around the globe with relatively little attenuation, and penetrate relatively deeply into sea water. Due to these unique propagation characteristics, VLF signals are useful for communications with submarines and other submerged assets, as well as for global navigation systems. Prior to the implementation of the satellite network that supports the current global position system (GPS), the United States operated the OMEGA navigation system. The OMEGA system used a network of transmitters located around the world and could determine position with an accuracy of a few kilometers [*Kasper and Hutchinson*, 1978; *Pierce*, 1989; *Warren et al.*, 1992]. Tens of VLF transmitters are currently maintained and continuously operated by different countries around the world. The amplitudes and phases of the VLF signals that propagate around the globe depend sensitively upon the electron density and temperature of the *D*-region ionosphere (∼50–100 km altitude) between the transmitter and the receiver. As a consequence, the detection and analysis of subionospherically propagating VLF signals provides a sensitive means for the remote sensing of the lower ionosphere with important scientific and geophysical applications.

This introduction begins by discussing modal VLF propagation within a simple parallel plate waveguide. The complications inherent to the Earth-ionosphere waveguide are then discussed. The focus then shifts to modulated signals which are generated by VLF transmitters, propagate within the Earth-ionosphere waveguide, and are detected by a network of receivers. Contemporary signal processing methods used to analyze VLF transmitter signals are then discussed, followed by a summary of the scientific contributions of this dissertation.

Chapter 2 details the specifications and characteristics of a new high dynamic range VLF receiver. Chapter 2 then shifts to focus on typical application problems associated with VLF receivers with a focus on clipping and capacitive loading. A calibration technique is then detailed and its importance as a means to equate the received voltage with a corresponding Tesla value is discussed. Prior systems have focused on an amplitude calibration system but the presented method provides a means to quantify the amplitude and phase response of a receiver in the field.

Chapter 3 details a new spread spectrum processing technique and compares it with other contemporary signal processing methods. Performance metrics for the new technique are evaluated for a variety of noise environments. The performance analysis focuses on ambient signal measurements at multiple receiver sites.

In Chapter 4, the application of the new method to transient events is considered and a statistical analysis that evaluates the spectral content observed during transient events is undertaken in order to determine the importance of retaining the spectral content for subsequent event analysis. Different types of VLF events, such lightning-induced electron precipitation (LEP) events, early/fast events, and solar X-ray flare events are evaluated using experimental data and lend themselves to statistical analysis due to the large number of occurrences in a given year. The events included in the statistical analysis are limited to large-amplitude events that are clearly not affected by other ionospheric events occurring closely in time. Chapter 4 also compares and contrasts the performance of the signal processing methods for individual transient events.

Finally, Chapter 5 summarizes the scientific contributions of this dissertation, and potential future research is discussed.

1.1 Parallel-Plate Waveguide: Mode Theory

A simple parallel plate waveguide can be described as two parallel metallic plates of infinite extent in the yz plane separated by a distance h in the x direction, as shown in Figure 1-1. The time harmonic form of Maxwell's equations, assuming the propagation medium is simple, source

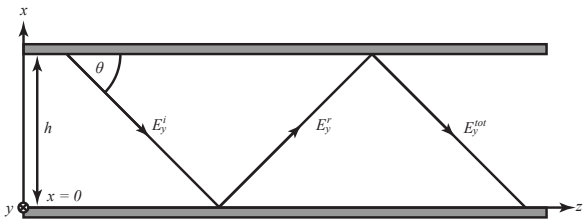

Figure 1-1. A simple parallel plate wave guide.

free, and lossless, can be expressed:

$$\nabla \times \mathbf{E} = -j\omega\mu\mathbf{H} \tag{1–1a}$$

$$\nabla \times \mathbf{H} = j\omega\varepsilon\mathbf{E} \tag{1–1b}$$

$$\nabla \cdot \mathbf{E} = 0 \tag{1–1c}$$

$$\nabla \cdot \mathbf{H} = 0 \tag{1–1d}$$

where $\mathbf{H}$ is the magnetic field, $\mathbf{E}$ is the electric field, $(\varepsilon = \varepsilon_r\varepsilon_0)$ is the permittivity of the medium, $(\mu = \mu_r\mu_0)$ is the permeability of the medium, and $(\sigma) = 0$ is the conductivity of the medium.

The solution to these equations, including the application of perfect electrical conductor boundary conditions, yields an infinite number of propagating and evanescent modes. Assuming propagation in the $+z$ direction, the modal solutions take the form of Transverse Electric (TE) solutions (with $E_z = 0$), Transverse Magnetic (TM) solutions (with $H_z = 0$), and Transverse Electromagnetic (TEM) solutions (with $E_z = H_z = 0$). Assuming no variation in the y direction, the solutions may be expressed:

$$\left.\begin{aligned} E_y &= A_m \sin\left(\frac{m\pi}{h}x\right) e^{-\gamma z} \\ H_x &= -\frac{\gamma}{j\omega\mu} A_m \sin\left(\frac{m\pi}{h}x\right) e^{-\gamma z} \\ H_z &= -\frac{m\pi}{j\omega\mu h} A_m \cos\left(\frac{m\pi}{h}x\right) e^{-\gamma z} \end{aligned}\right\} \quad m^{\text{th}} \text{ Transverse Electric (TE) Mode} \tag{1–2}$$

$$\left.\begin{aligned} E_x &= \frac{\gamma}{j\omega\varepsilon} B_m \cos\left(\frac{m\pi}{h}x\right) e^{-\gamma z} \\ E_z &= \frac{jm\pi}{\omega\varepsilon h} B_m \sin\left(\frac{m\pi}{h}x\right) e^{-\gamma z} \\ H_y &= B_m \cos\left(\frac{m\pi}{h}x\right) e^{-\gamma z} \end{aligned}\right\} \quad m^{\text{th}} \text{ Transverse Magnetic (TM) Mode} \tag{1–3}$$

$$\left.\begin{aligned} E_x &= \frac{\gamma}{j\omega\varepsilon} C e^{-\gamma z} \\ H_y &= C e^{-\gamma z} \end{aligned}\right\} \quad \text{Transverse Electromagnetic (TEM) Mode} \tag{1–4}$$

In the equations above, A_m, B_m, and C are arbitrary complex constants. γ is the wave propagation constant that is modified by application of the boundary conditions and expressed:

$$\gamma = \sqrt{\left(\frac{m\pi}{h}\right)^2 - \omega^2\mu\varepsilon} = \beta\sqrt{\left(\frac{f_{cm}}{f}\right)^2 - 1} = j\beta\sqrt{1 - \left(\frac{f_{cm}}{f}\right)^2} \tag{1–5}$$

where f_{cm} is the cutoff frequency of the m^{th} mode. In this case, γ is either purely real-valued (evanescent modes, $f < f_{cm}$) or purely imaginary-valued (propagating modes, $f > f_{cm}$). Each mode also propagates with a different phase velocity and group velocity:

$$v_p = \frac{c}{\sqrt{1 - \left(\frac{f_{cm}}{f}\right)^2}} \tag{1–6}$$

$$v_g = c\sqrt{1 - \left(\frac{f_{cm}}{f}\right)^2} \tag{1–7}$$

For propagation over large distances, these equations imply that each mode will arrive at the receiver with a different propagation delay and with a different phase. At steady-state, the total field detected at the receiver can be expressed as the sum of all modes. If one expresses the amplitude and phase of each (TE, TM, and TEM) mode as $Ae^{j\phi}$, the total resulting field may generally be expressed as the sum over all TE, TM, and TM components:

$$Fe^{j\psi} = \sum_n A_n e^{j\phi_n} \tag{1–8}$$

In this context, F and ψ are referred to the "sum-of-modes" amplitude and phase of the received signal component.

1.2 Propagation in the Earth-Ionosphere Waveguide

1.2.1 A Brief History

Understanding of VLF propagation in the ionosphere has served both science and industry. In the early 1900's, the discovery of radio and the resulting boom in VLF radio stations help to further the understanding of propagation due to the increase in the number of receivers and transmitters. *Round et al.* [1925] analyzed several experimental measurements of electric field strength for a variety of a transmitters as both functions of distance and time of day. The results showed diurnal variation, differences in the direction of propagation, and "around the world" propagation from different directions. The evidence of naturally occurring phenomena also led to publications associated with causes of 'atmospheric' or natural sources of VLF waves propagating in the ionosphere, [*Appleton and Chapman*, 1937]. Appleton's observations of electric field changes in proximity to lightning storms and the resulting waveforms aided in understanding the mechanics of lightning and the nature of natural wave propagation in the ionosphere. While research in to the VLF propagation continued, it was not until the 1950's and 60's that a large amount of research was published about the nature of mode propagation and radio atmospherics. Several publications in the early fifties detailed experimental observations that helped to establish the underlying mechanics associated with VLF wave propagation [*Bracewell et al.*, 1951; *Chapman and Macario*, 1956; *Chapman and Pierce*, 1957; *Helliwell et al.*, 1951; *Holzer and Deal*, 1956].

Budden in 1955 published two seminal papers [*Budden*, 1955a,b] that defined a numerical solution for VLF wave propagation in the ionosphere where the ionosphere is treated as a parallel plate waveguide. Wait [*Wait*, 1957] published a similar solution, but for spherical coordinates, and treated the lower boundary as having a finite conductivity compared to Budden's perfect conductor. Both Budden and Wait continued to revise their theories over the course of several publications

[*Budden*, 1957; *Wait*, 1959; *Wait and Spies*, 1960] with the help of other publications on the subject that dealt with the variety of the complications associated with the Earth-ionosphere waveguide. These effects include the reflection of propagating waves at a sharp boundary [*Crombie*, 1961; *Yabroff*, 1957], the presence of the Earth's magnetic field and its affect on wave propagation [*Budden*, 1961; *Crombie*, 1959; *Wait and Spies*, 1960], and direction-dependent propagation (i.e., East-to-West or West-to-East) [*Dobrott and Ishimaru*, 1961; *Taylor*, 1960]. The research in the fifties and early sixties culminated in a variety of books that detailed VLF wave propagation and mode theory [*Budden*, 1961; *Galejs*, 1972; *Wait*, 1961].

1.2.2 The Earth-Ionosphere Waveguide

With a working understanding of modal VLF propagation, the structure of the ionosphere and its resulting effects on VLF propagation can be considered. The ionosphere spans hundreds of kilometers, from ∼50 km altitude to ∼1000 km altitude. The *D*-region ionosphere (50–100 km altitude) is of greatest interest for VLF observations, and can often vary with both space and time based on atmospheric activity or on the solar-terrestrial environment. The primary engine for the *D*-region ionosphere is the sun, resulting in quiet days, turbulent nights, and a day-to-night terminator. The solar cycle and seasonal changes produce variations in the ionosphere, but the differences are not so dramatic that data from one year cannot be compared to data from the next year. Electron and neutral molecule densities are strong functions of altitude and the densities will shift based on the time day or in the presence of an ionospheric event. Additional ionization occurs through a variety of different events will affect the *D*-region ionosphere by shifting the VLF reflection height and potentially causing perturbations in the received amplitude and phase of VLF signals.

The ambient ionosphere also deserves consideration for its day-night cycle and the associated day-night terminator. While no two days are the exactly the same, there will often be similar characteristics. Figure 1-2 shows an example of a typical spectrogram at Palmer Station, Antarctic (a) and 23 hours worth of data received at San Antonio (b). Several distinct phenomena can be seen in the spectrogram, such as lightning-generated radio impulses (radio atmospherics, or sferics)

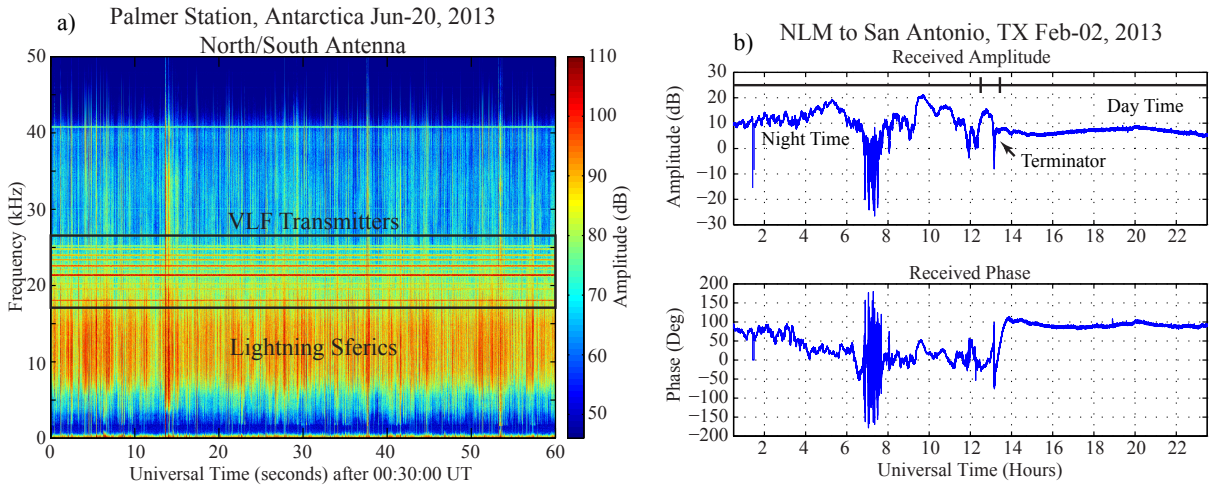

Figure 1-2. A one-minute duration spectrogram from Palmer Station, Antarctica (a) and the received amplitude and phase from the NLM to San Antonio over a 23 hour period (b).

which take the form of vertical lines with a wide range of frequency content over a short time scale. VLF transmitters can be seen as horizontal lines in the 20–25 kHz range. A typical day, as shown by the 23 hours of data, is broken up in to the quiet day and the turbulent night where each section is separate by a day-night or night-day transition region (terminator).

The propagation between a VLF transmitter and receiver is often perturbed by a transient disturbance. As Figure 1-3 illustrates, a lightning storm in proximity to the transmitter-to-receiver propagation path changes the VLF reflection height. The change in ionization manifests as a perturbed amplitude and/or phase at the receiver that will exhibit different onset delays, the time for an event to begin relative to its causative source, and recoveries, the *e*-folding time scale for relaxation to the ambient condition, depending on the type of event. Depending on the strength of an event and on the geometry of the transmitter/disturbance/receiver setting, it may be possible to see the perturbation on multiple VLF transmitters at a single receiver. Solar X-ray flares will often have a strong effect on multiple transmitters signals because half of the entire Earth is bathed in X-ray producing ionization while lightning induced electron precipitation will often be more localized and path dependent.

In order to model the effects of these transient disturbances in the ionosphere on VLF signal propagation within the Earth-ionosphere waveguide, a computational model that properly accounts

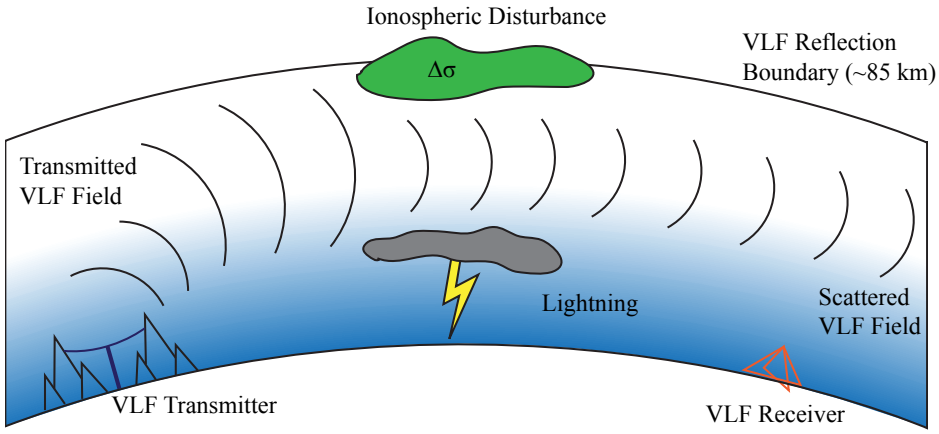

Figure 1-3. An example of the propagation between a VLF transmitter and receiver when the received wave is perturbed by at some point along the path.

for the fundamental physics involved is required. Such a model would account for the effect of Earth's magnetic field, which makes the ionosphere anisotropic and results in direction-dependent wave propagation. The porous nature of the ionospheric boundary will also prevent the propagation of pure TE and TM modes. Instead, quasi-TE (QTE) and quasi-TM (QTM) modes propagate with an (typically small) field component in the direction of propagation. Additionally, these modes will propagate with different attenuations rates, propagation constants, and group velocities. Mode-coupling will occur at sharp changes in boundary conductivity. Such discontinuities occur where the ionospheric conductivity changes rapidly with space (e.g., at the day/night terminator) or where the ground conductivity changes rapidly with space (e.g., land/sea interface). The curvature of the Earth also plays an important role because it will reduce the VLF signal amplitude on the order of a few dB.

Despite the complications of a more realistic Earth-ionosphere waveguide, however, the fundamental problem remains the same. The VLF signal will propagate in a multi-mode environment, and the signal detected at the receiver may simply be expressed as the sum of waveguide mode field values, each with different amplitudes and phases. It is noted at this point that the vast majority of numerical modeling methods for VLF propagation in the Earth-ionosphere waveguide output the sum-of-modes amplitude and phase (or the equivalent) as a prediction. It is this sum-of-modes

solution that the vast majority of VLF signal processing methods are intending to produce for comparison with theoretical modeling predictions.

1.3 MSK-Modulated VLF Transmitter Signals

The US military continuously operates several high power VLF transmitters in order to communicate with its submarine force. By utilizing the advantages of the Earth-ionosphere waveguide the transmitted signals are able to propagate around the globe with attenuation on the order of a few dB/Mm and penetrate sea water to a depth of several meters. Most of the transmitters use a 200-Hz bandwidth minimum shift keying (MSK) modulation. The bit sequence is guaranteed to have a Bernoulli(1/2) distribution that is independent and identically distributed (iid), giving an equal chance of receiving a 1 or 0 bit independent of the rest of the bit sequence. For each bit, the phase of the signal increases or decreases by 90° (or π/2 radians) over 5 ms [*Gronemeyer and Mcbride*, 1976; *Pasupathy*, 1979], and the phase is forced to be continuous at bit transitions. Processing the MSK modulation is more complicated than processing a simple CW signal.

The MSK-modulated transmission can be described as:

$$x(t) = A \sum_{k=-\infty}^{\infty} \cos(\omega_c t + \phi_k + \omega_\Delta a_k t) p(t - k/r_b) \tag{1–9}$$

where:

$x(t)$	is the transmitted signal
A	is the transmitted amplitude
$\omega_c = 2\pi f_c$	is the carrier frequency
ϕ_c	is the carrier phase
$\omega_\Delta = 2\pi f_\Delta = \pi r_b/2$	is the frequency offset from the carrier
a_k	is the bit rate
$p(t) = u(t) - u(t - r_b^{-1})$	is a one-bit pulse of unit amplitude
$\phi_k = \frac{\pi}{2}\sum_{j=0}^{k-1} a_j$	is the carrier phase

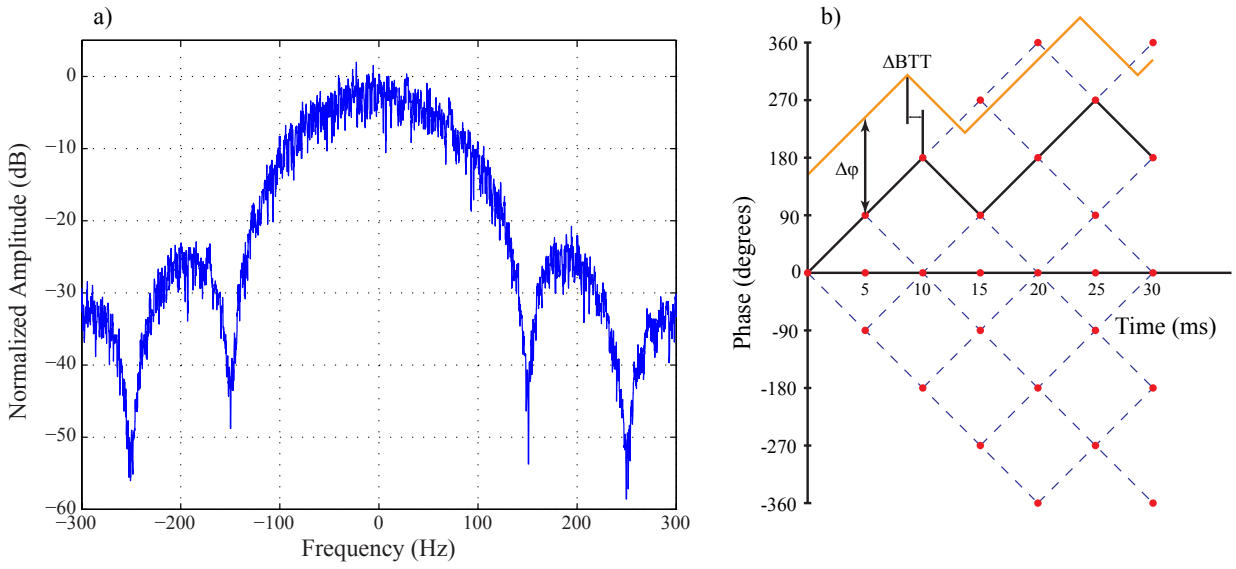

Figure 1-4. A typical MSK modulation spectrum with a random bit sequence (a). A random bit sequence (b). While an ideal bit sequence is plotted in black, it can also be expected the initial phase, green, to vary as well as the bit transition time, orange. There is no guarantee that at time zero there will be zero initial phase and be at the start of a bit.

An alternative way to express the signal is simply:

$$x(t) = A\cos\left(\omega t + \phi_{\mathrm{MSK}}(t) + \phi_0\right) \tag{1–10}$$

where $\phi_{\mathrm{MSK}}(t)$ expresses the phase introduced to the signal by the MSK modulation (i.e., the MSK phase trellis).

Figure 1-4 is an example of the MSK modulation spectrum for a random bit sequence (a), and a phase trellis where a random sequence has been plotted (b). The propagation from transmitter to receiver will result in a variation of the initial phase, $\Delta\phi$, and the bit transition time, ΔBTT, due to propagation effects. Similar to the phase measurement for a CW signal, MSK phase measurements are referenced to an MSK transmission that starts with zero phase. Phase measurements of MSK-modulated signals require an additional assumption about the timing of the bit sequence: in our case, it is assumed that, at transmission, the start of each bit (modulo 5 ms) is synchronized with the beginning of a second. While the signals received often have a high SNR (20 dB+), bit errors still occur due to complications when fitting a bit sequence to the received phase due to the presence of

Table 1-1. A list of VLF transmitters. The United States military operates all of these transmitters and they are the highest power VLF transmitters in the world, resulting in high SNR signals received as far away as Antarctica.

Call Sign	Location	Frequency	Power	Modulation
NWC	North West Cape, Australia	19800 Hz	1000 kW	MSK
NPM	Lualualei, Hawaii	21400 Hz	424 kW	MSK
NAA	Cutler, Maine	24000 Hz	1000 kW	MSK
NLK	Jim Creek, Washington	24800 Hz	192 kW	MSK
NLM	LaMoure, North Dakota	25200 Hz		MSK
NAU	Aguado, Puerto Rico	40750 Hz	100 kW	MSK

impulsive noise, long path propagation interference, transient events, and multi-mode interference (during modal nulls, for instance).

Each frequency component of the transmitted MSK spectrum propagates within the Earth-ionosphere waveguide in a multi-mode environment. Each frequency component of the VLF signal that is detected at the receiver is thus a sum waveguide modes. In order to provide a significant simplification for comprehension, one can assume that each frequency within the 200-Hz band of the VLF transmission exhibits the same modal structure, and the signal at the receiver can be expressed:

$$y(t) = \sum_n A_n \cos\left(\omega t + \phi_{\mathrm{MSK}}(t-\tau_n) + \phi_{0n}\right) \tag{1–11}$$

where A_n is the amplitude of each mode (which are different due to different modal attenuation rates), τ_n is the propagation delay associated with each mode (which are different due to different modal group velocities), and ϕ_{0n} is the phase of each mode (which are different due to different modal phase velocities). This perspective is valid in many cases, and it is certainly useful for thought experiments, but it is not perfectly correct.

An abbreviated list of US-operated VLF transmitters is listed in Table 1-1. These transmitters are located in the United States or one of its territories, except for NWC, which is located in Australia. There are other VLF transmitters located around the world, but for the purposes of this work, the focus will primarily be on paths from these transmitters to receivers around the world. According to the US military, these transmitters operate with precise timing (the bit sequence is

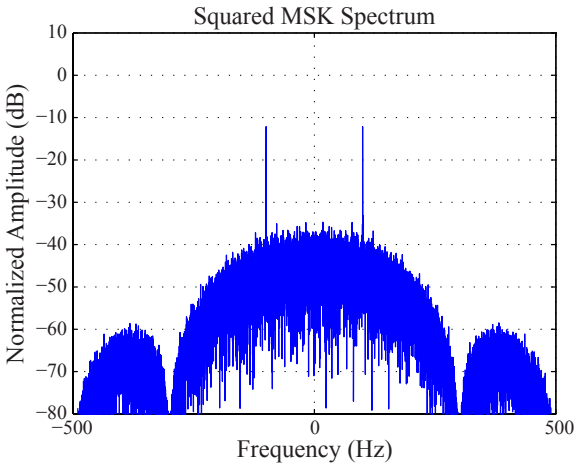

Figure 1-5. The MSK spectrum used by the Dowden method. Two frequencies, ± 100 Hz, are generated by squaring the spectrum making it similar to Frequency Shift Keying (FSK) modulation .

synchronized with the start of a second) and with a precise transmission frequency. Per usual, cases will be found occasionally where neither of these characteristics are true.

1.4 MSK Signal Processing Techniques

Several techniques have been used in the past to analyze VLF signals. *Dowden and Adams* [1989] present a technique to perform a two-frequency measurement of the MSK-modulated signal, Stanford University has implemented an unpublished technique to calculate the CW-equivalent amplitude and phase of the MSK-modulated signal, and *Shafer* [1994] presents a method to calculate the ionospheric transfer function using the entire 200-Hz bandwidth of the signal. As will be shown in subsequent chapters, each of these methods produces a reasonable approximation of the amplitude and phase of the propagating VLF signal; this new method out-performs each of these methods, however. In order to understand the distinctions between the signal processing methods, a summary of each method is provided here.

1.4.1 The Dowden/Adams Method

The method presented by *Dowden and Adams* [1989] starts by mixing down and low-pass filtering the MSK-modulated VLF signal. They then proceed to square the complex-valued result, producing a signal whose spectrum is very similar to that of a frequency shift-key (FSK) modulated signal. The resulting spectrum can be seen in Figure 1-5, with impulses at ±100 Hz. Using

the simplified version of the MSK-modulated signal described above, the resulting signal can be expressed as:

$$y_{lp}^2 = \left[\sum_n A_n e^{j(\phi_{\text{MSK}}(t-\tau_n)+\phi_{0n})}\right]^2 \tag{1–12}$$

In the single-mode case, this results in:

$$y_{lp}^2 = A^2 e^{j2(\phi_{\text{MSK}}(t-\tau)+\phi_0)} \tag{1–13}$$

implying that the amplitude and phase detected at ± 100 Hz can be converted to the amplitude and phase of the MSK-modulated signal. In the general (multi-mode) case, however, this squaring technique does not have a physical interpretation and cannot be related to either the sum-of-modes amplitude or the sum-of-modes phase of the MSK-modulated signal. Nevertheless, it will be shown in subsequent chapters that this technique provides a reasonable approximation of the received amplitude and phase.

1.4.2 The Stanford Single-Frequency Method

In contrast with the dual-frequency method presented by *Dowden and Adams* [1989], the Stanford method is a single frequency measurement at a VLF transmitter's center frequency. This unpublished method essentially converts the MSK-modulated signal to a CW signal and then calculates the resulting amplitude and phase. Similar to the *Dowden and Adams* [1989] method, the Stanford method starts by mixing down and low-pass filtering the MSK-modulated VLF signal. The amplitude is calculated directly from this baseband version of the signal. The method then attempts to reproduce the MSK phase trellis, first calculating the perceived bit-transition-time Υ and then the bits, and generating a model version of the phase: $\Phi_{\text{MSK}}(t-\Upsilon)$. The method then subtracts the model phase from the received phase, resulting in the phase measurement. The amplitude can thus be expressed:

$$A = \left|\sum_n A_n e^{j(\phi_{\text{MSK}}(t-\tau_n)+\phi_{0n})}\right| \tag{1–14}$$

and the phase can be expressed:

$$\phi_{lp} = \angle\left[\sum_n A_n e^{j(\phi_{\text{MSK}}(t-\tau_n)+\phi_{0n})}\right] e^{-j(\Phi_{\text{MSK}}(t-\Upsilon))} = \angle\sum_n A_n e^{j(\phi_{\text{MSK}}(t-\tau_n)-\Phi_{\text{MSK}}(t-\Upsilon)+\phi_{0n})} \tag{1–15}$$

In the single-mode case, the amplitude will be correct, and the phase, assuming accurate assessment of the bit transition time, will also be correct. In the multi-mode case, however, the amplitude (square root of the sum of the real and imaginary components squared) is not the sum-of-modes amplitude. Additionally, the perceived bit transition time is likely to be different from the bit transition times associated with each mode (they are all expected to be different). This introduces error into the phase calculation as well. Similar to the *Dowden and Adams* [1989] method, it will be shown in subsequent chapters that, despite these philosophical drawbacks of the method, the Stanford method is also a good approximation of the sum-of-modes amplitude and phase.

1.4.3 The Shafer Spread-Spectrum Method

Shafer [1994] presented an algorithm to calculate the transfer function of the Earth-ionosphere waveguide. This is a very powerful method as it allows the full use of the MSK spectrum in order to analyze events and their corresponding frequency components. This algorithm forms the basis for the new algorithm presented in this dissertation. The Shafer spread spectrum processing technique transfer function is simply expressed:

$$\hat{G}_N(e^{j\omega}) = \frac{\hat{\Phi}^N_{yu}(\omega_0)}{\hat{\Phi}^N_u(\omega_0)} \tag{1–16}$$

where

$\hat{\Phi}^N_{yu}(\omega_0)$ is the weighted FFT of the received signal

$\hat{\Phi}^N_u(\omega_0)$ is the weighted FFT of the modeled transmission

$\hat{G}_N(e^{j\omega})$ is the weighted raw estimate at frequency ω

Shafer's method also relies upon developing an accurate model of the transmitted MSK phase trellis and the bit transition time. The model phase trellis implemented by Shafer, however, lined up the bit transition time (of the model signal) with the perceived bit transition time (of the received signal). Shafer interpreted this process as accounting for speed-of-light propagation delay. Due to

multi-mode interference, however, the perceived bit transition time can be shifted (either advanced or retarded) in time relative to the true speed-of-light propagation delay. As a result, Shafer's phase calculations could not be used to perform site-to-site transfer function measurements (i.e., using observations at multiple receivers). Of lesser importance, Shafer's method does not account for the filtering employed or the windowing employed in the signal processing. In the performance assessment in subsequent chapters it will be shown that accounting for these details affects the accuracy of both the amplitude and the phase measurements. Shafer also included a clipping step in the algorithm in an attempt to mitigate the effects of strong radio atmospherics. The performance analysis included in Chapter 3 demonstrates that any clipping of the signal introduces significant errors to both the amplitude and phase measurements.

1.5 Scientific Contributions

The scientific contributions of this Ph.D. dissertation are as follows:

1. A new spread spectrum VLF signal processing technique is developed and implemented. Previous difficulties with propagation delay, signal filtering, and signal windowing are successfully addressed. The new technique allows for direct comparisons of the amplitude and phase of MSK-modulated signals observed at different receiver sites.

2. The new signal processing algorithm is used to demonstrate the existence of frequency dependent VLF signal propagation under ambient conditions (within the 200-Hz bandwidth of the signal). Comparative performance analysis demonstrates the superiority of the method over other contemporary methods under ambient conditions.

3. Amplitude and phase measurements during individual transient events are shown to exhibit frequency dependence. It is demonstrated that this is *not* the typical case, however. All signal processing methods considered produce very similar changes in amplitude and phase during transient events. The differences between the measurements are attributed to the effects of multi-mode signal propagation.

4. A statistical analysis demonstrates that the *percent change* in amplitude and the *change* in phase observed during transient events do not exhibit significant frequency dependence. This result implies that transient event analysis can be completed successfully using only bandwidth-averaged signal measurements.

5. The successful design and implementation of a VLF receiver with high (>90 dB) linear dynamic range, high (>90 dB) spurious-free dynamic range (SFDR), and high (>90 dB) signal-to-noise and distortion ratio (SINAD) is demonstrated.

CHAPTER 2
VLF RECEIVER HARDWARE

Understanding the received VLF signal (by analyzing the VLF receiver hardware) goes hand-in-hand with understanding the signal processing. Much of Stanford's work regarding receiver hardware has been published under the ELF/VLF Radiometer project [*Fraser-Smith and Helliwell*, 1985] and the purported AWESOME system [*Cohen et al.*, 2010]. Other Universities have also contributed to the advance of VLF receiver hardware with the OmniPAL system from the University of Otago, New Zealand, [*Dowden et al.*, 1998], the receivers used by the University of Washington for their World Wide Lightning Location network [*Lay et al.*, 2004], and a wideband low-frequency receiver developed at the University of Bath, United Kingdom [*Füllekrug*, 2009]. While receivers at the block level can be similar, there are variations between systems and it is important to understand the architectures used because design choices directly impact the performance of the system and the integrity of its associated data.

This chapter highlights several important choices when designing a VLF receiver system and analyzes the performance of the latest version of the University of Florida VLF receiver. The chapter then shifts to analyze some of the application problems that affect VLF hardware design with a focus at the end of the chapter on an amplitude and phase calibration scheme. While calibration is a simple concept – inject a known value and measure the output – practical considerations must be made when performing calibrations at a receiver already deployed in the field. A 'laboratory' calibration is useful but cannot be considered an accurate calibration of a system in the field for an extended length of time. Variations due to temperature and elemental wear degrade a system's performance over time and require that calibration is repeated periodically. With a full field calibration, recorded data can be correctly calibrated to account for system degradation.

Receivers operated by the University of Florida's Ionospheric Radio Laboratory (IRL) are located around the world with locations in Antarctica, Greenland, and the United States. Figure 2-1 shows a map of currently operating VLF receivers and some important VLF transmitters around the world. A list of locations has been tabulated in Tables 2-1 and 1-1. Different sites offer different

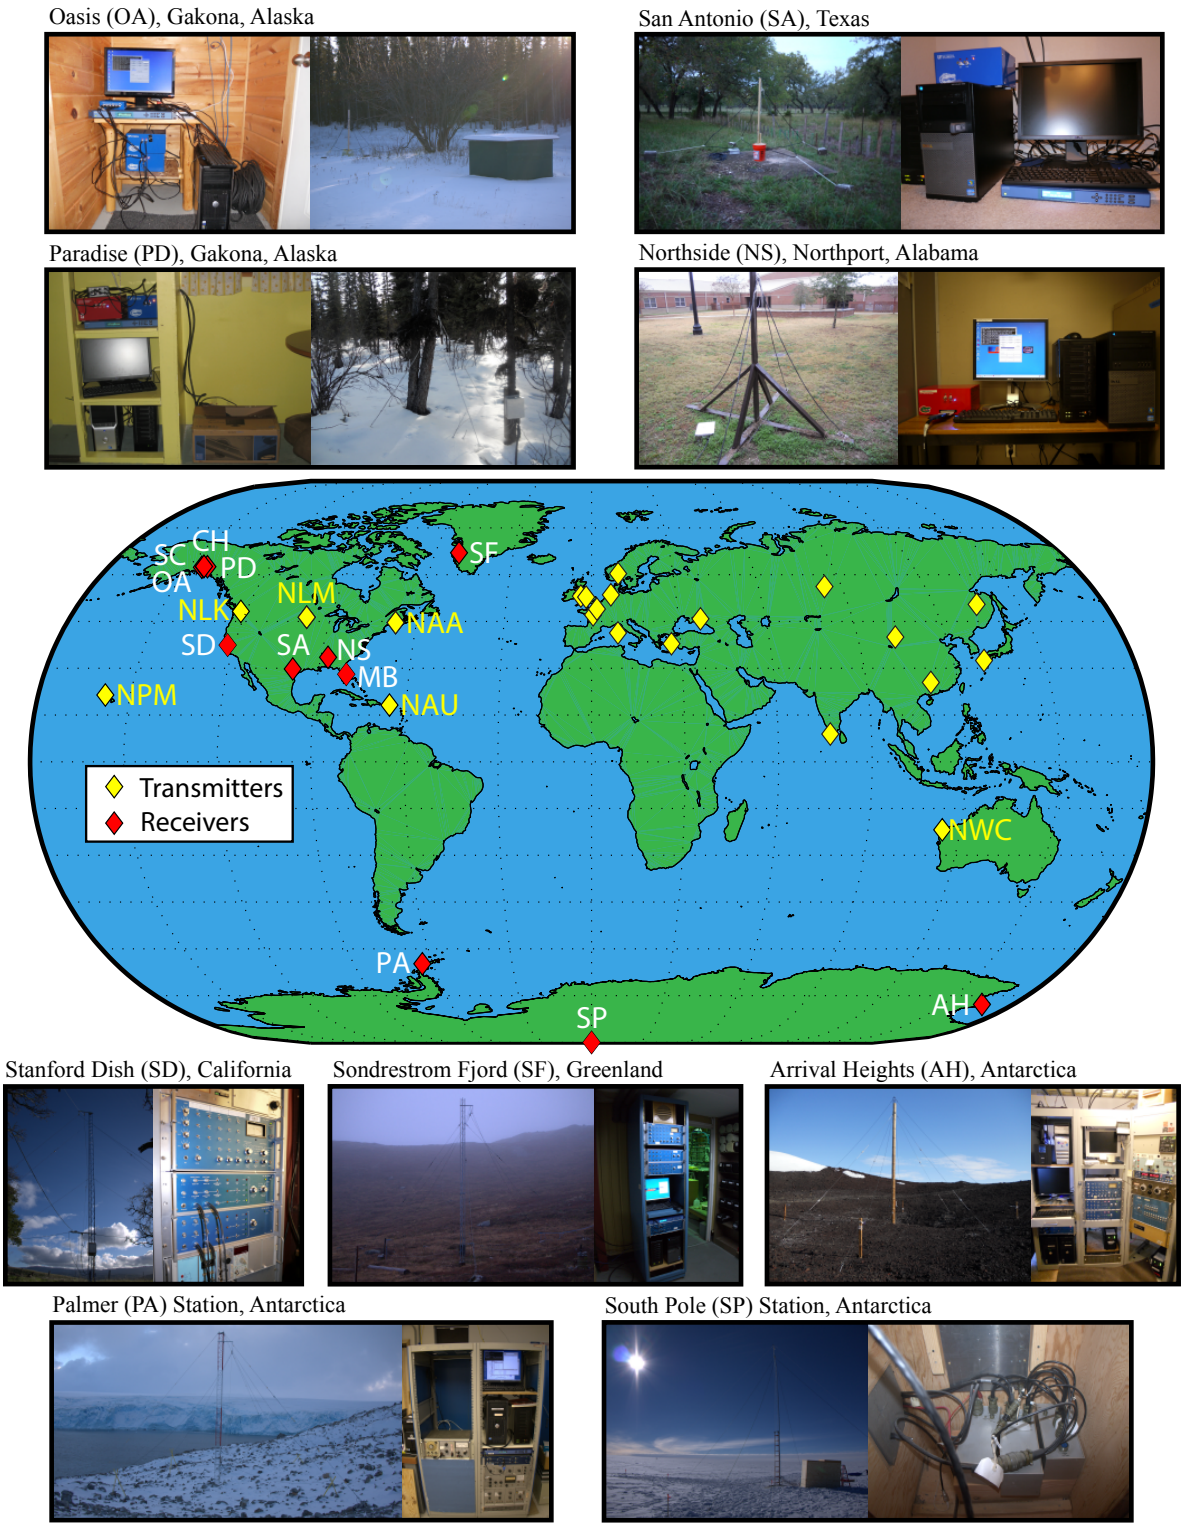

Figure 2-1. A map of transmitters, yellow, and receivers, red, around the world. While there are tens of transmitters the focus will only be on transmitters that are labeled. Several pictures have been included to showcase the wide variety of receiver locations. Photos courtesy of Michael Mitchell and Daniel Kotovsky.

Table 2-1. A list of receivers located around the world. Receivers located in Alaska are in close proximity to Gakona, Alaska. Some locations have Extremely Low Frequency (ELF) receivers in addition to Very Low Frequency (VLF) receivers.

Site Code	Site Name	Location	Type
OA	Oasis	Gakona, Alaska	VLF and ELF
SC	Sinona Creek	Gakona, Alaska	VLF
CH	Chistochina	Gakona, Alaska	VLF
PD	Paradise	Gakona, Alaska	VLF and ELF
SD	Stanford Dish	Stanford, California	VLF and ELF
SA	San Antonio	San Antonio, Texas	VLF
NS	Northside	Northport, Alabama	VLF
MB	Melbourne	Melbourne, Florida	VLF
PA	Palmer	Palmer Station, Antarctica	VLF
AH	Arrival Heights	McMurdo Station, Antarctica	VLF and ELF
SP	South Pole	South Pole Station, Antarctica	VLF

background noise environments as well as different propagation paths from specific transmitters. Certain types of ionospheric events are more likely to be detected at different locations. Different paths also lead to different mode structures due to the different conductivity profiles produced by sea water, land, and ice. For instance, the NPM to Palmer Station, Antarctica path is typically a single mode propagation path given that the lower boundary of the Earth-ionosphere waveguide is sea water for the nearly the entire path. Paths with water-to-land or land-to-water boundaries can induce mode conversion resulting in multi-mode propagation. Using observations from multiple receiver locations also provides a diverse data set with which to analyze the spread spectrum processing method.

2.1 VLF Receiver Architecture

The typical VLF receiver can be broken up in to several components: antennas, preamplifier, cable, line receiver, digitizer, and storage. Figure 2-2 shows a more detailed block level diagram of a VLF receiver and includes the relevant locations for calibration injection. A single set of two antennas are positioned with one inside the other and oriented along magnetic field lines or cardinal directions (grid north, for example), based on geographic locations. Each antenna connects to the preamplifier which is composed of three, stages: a custom low noise amplifier using matched PNP transistors, a frequency compensation stage, and a line driving output stage.

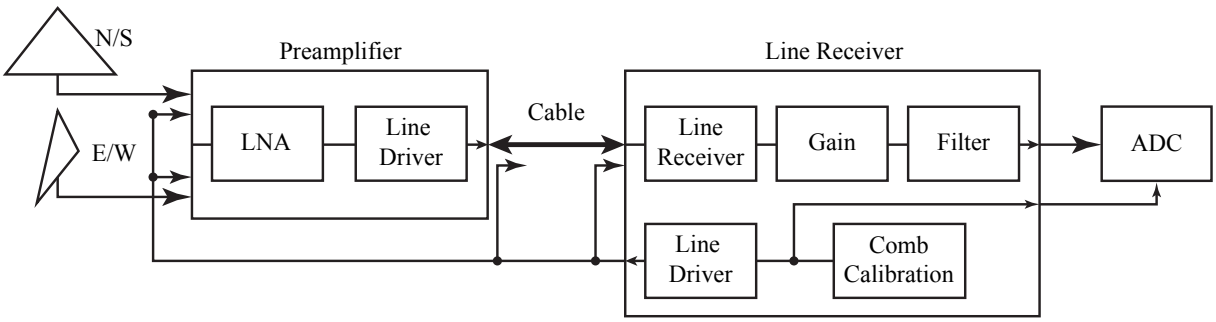

Figure 2-2. A typical VLF receiver at the block level. Note that the label, line receiver, has been used twice to denote both the enclosure containing several elements and the element used to terminate the cable. While potentially confusing both terms are correct but the terminology has become muddled over the years across several system iterations.

The preamplifer is separated from the line receiver by thousands of feet (with a typical distance being 1000 feet). Receivers in remote geographic locations will often have a separation of 2000 to 3000 feet. A line receiver will contain the line-matching electronics, gain stages, and filter stages for each channel. There is also a calibration stage and its associated line driver stage which together deliver calibration tones to the preamplifier and other points in the system's signal path for the purposes of full calibration. There are slight differences in some of the receivers that are part of the network and it is convenient to think of the major blocks in terms of their individual components. Receivers located in Antarctica, Greenland, and California are legacy VLF system from Stanford's ELF/VLF Radiometer project and have a separate enclosure that houses the filter block. The filter enclosure is often refereed to as the signal conditioning unit or the anti-alias filter. There are other facets to VLF Radiometer systems that will be detailed in later subsections as the systems differ greatly in some areas compared to the light weight and easily deployable systems that make up the rest of the IRL receiver network.

2.1.1 VLF Antennas

Designing antennas for the VLF range is dependent on several factors, such as the number of turns, the wire diameter, and the loop area. Ultimately, the antenna sensitivity (or lowest detectable signal level) depends on the inverse product of the antenna's mass and area [*Paschal*, 1988]. A

normalized sensitivity is typically used, S_0, defined such that the sensitivity $S = S_0/f$ has units of Tesla/Hz$^{1/2}$. The normalized sensitivity is:

$$S_0 = \frac{(4kT\rho c_1)^{1/2}}{2\pi(AM)^{1/2}} \tag{2–1}$$

where

- k is Boltsmann's constant
- T is the temperature in Kelvin
- ρ is the resistivity coefficient of copper, 1.724E-8
- c_1 numerical constant used to calculated loop inductance
- δ is the density of copper, 8.5E-3
- A is the antenna loop area
- M is the mass of the wire in the antenna

Thus, the antenna sensitivity depends on the inverse product of AM and the best (lowest) sensitivity comes from antennas with a large area and mass product. A large loop will result in better (lower) sensitivity but is somewhat unwieldy. Conversely a smaller loop is less sensitive but easier to manage. Loops for VLF antennas will be designed in order to fit constraints of sensitivity, S_0, and impedance (resistance and inductance) with choices of the shape, wire gauge, and size of the loops [*Paschal*, 1988]. The VLF radiometer loops, also known as IGY antennas after International Geophysical Year where they were deployed during a flurry of research in Antarctica, are a single turn of 6 gauge welding cable with an area of 81 m^2, an inductance of 65 μH, a resistance of 0.0615 Ω, and a normalized sensitivity of 1.866×10^{-5} T-Hz$^{1/2}$.

The pure size of the antenna, a right isosceles triangle with a height of 29 feet, makes the system difficult to deploy in all but the best locations, which is why these systems are located in especially remote areas, such as Greenland and Antarctica, to take advantage of the sensitivity offered by such a large sized antenna.

For comparison, the portable VLF system has right isosceles triangle antennas that are made out of 12 turns of 16 gauge wire with an area of 1.695 m^2, an inductance of 1 mH, a resistance of 1 Ω, and a normalized sensitivity of 2.97×10^{-5} T-$Hz^{1/2}$. IGY antennas are more sensitivity by an order of magnitude but at the price of mass and size. The portable VLF antennas are a good compromise in terms of size and sensitivity and it can often be the case where the sensitivity limitation will be the environment that the antenna is located at and not the antennas themselves. However, it is also possible to create larger antennas that matches the inductance and resistance for a portable system. A receiver located in Chistochina, Alaska uses a larger antenna but the site has been permanently installed for a number of years and is located within a forest with large trees making the process of guying the antenna easier. In designing a recent set of antennas with a square shape, it was found that the theoretical derivations [*Paschal*, 1988] were 15-20% high on inductance and 10-15% low on resistance.

The smaller loops are typically oriented parallel to the North/South and East/West magnetic fields with accuracy on the order of a few degrees. Due to physical the location of the receivers located in Antarctica, it is difficult to orient antennas along magnetic lines. It is also difficult to identify a distinct direction associated with North, South, East, or West (e.g., at the South Pole, every direction is North). At these locations, grid North, which points toward the prime meridian, is used as a reference.

2.1.2 VLF Preamplifier

Despite the age of the preamplifier used in the original VLF radiometer, its performance has weathered the test of time. A few important changes have been made since its initial inception, the most notable is the use of matched pair PNP transistor on a single chip as opposed to matching individual PNP transistors. The (so called) AWESOME receiver and the University of Florida variant use a similar design with a few differences. The preamplifier stage starts with a custom wound matching transformer that matches to the 1 Ω, 1 mH impedance of the antennas and transforms the impedance to match the input impedance of two common base amplifiers. This is followed by a differential pair and then an emitter follower stage which provides a low output impedance [p. 65

Paschal, 1988]. While the transistor stage is differential, it is important to note that the antennas are pseudo-differential with only positive and negative terminals. It is possible to add a common terminal for a multiple turn antennas, but it is never done in practice due to the difficulty in taking the center point of the antenna. For a single turn antenna, such as the IGY antennas it is not practical given the antenna's structure. And the additional benefit of an input transformer is that the input terminations do not have an additional DC bias and can be referenced directly to ground, negating the need for an antenna common.

While the output of the preamp stage is differential it has typically been taken as single ended in order to apply frequency compensation for the lower frequencies. The latest University of Florida version takes the differential output of the transistor stage, high pass filters it with a simple RC filter and then converts the differential signal to single ended with a high end instrumentation amplifier. It is possible to keep the outputs differential and provide frequency compensation differentially but component variation, even with small tolerances, will produce slight differences that will degrade the performance of the preamp stage. This method also requires a second set of components compared to the single ended version. The conversion from differential to single ended version also has its own set of problems, mainly the increased signal range. Two differential signals with a voltage amplitude $\pm V_0$ will give a total signal of $2V_0$ when converted to a single ended system. However, careful selection of the power rails, choosing a voltage rail twice the saturation limit of the preamp stage, and components, amplifiers that operate near the power rails, helps to mitigate potential clipping and a loss of dynamic range.

2.1.3 VLF Line Driver and Line Receiver

A line driver and receiver pair is an important component when driving a cable of any length. In the case of the VLF radiometer and portable systems, the preamplifier and line receiver are separated by 1000s of feet of cable. While the cable will have either a characteristic impedance of 50 or 75 Ohms, there is a large amount of capacitance, on the order of nanofarads, that needs to be compensated.

There are two methods to drive cable of such a great length correctly. One is to use a pair of transformers to match to the impedance of the cable. Transformers lessen the requirements on the output and input amplifiers located in the preamplifier and line receiver while providing excellent common mode rejection. The downside is that they are large and difficult to produce. The purported AWESOME system and a Florida variant have used this method in order to drive cable on the order of 1000 feet. The VLF radiometer and the newest version of the Florida VLF receiver uses a special amplifier pair in order to drive the cable. Amplifiers have the advantage of being smaller and more manageable leading to an overall more modular system. The disadvantages include a degraded common mode rejection ratio compared transformers and a decrease in bandwidth due to the necessary capacitance compensation needed to maintain stability. A more detailed analysis of frequency compensation techniques for amplifiers will be covered in a later section but they can be summarized as sacrificing performance in order to drive a capacitive load.

2.1.4 Variable Gain Stage

The variable gain stage of a receiver is important in order to take full advantage of the dynamic range of the analog to digital converter. The radiometer and Florida system use a variable gain amplifier located in the line receiver. Conversely, the AWESOME receiver has the variable gain located in the preamplifier. The primary advantage of having the gain located in the line receiver is to quickly change the gain of the system, with the gain located in preamplifier, it can be arduous to change the gain as the trek to and from the preamplifier; which is at least a 1000 foot trek in the snow, each way. From a systems theory point of view it is advantageous to have the highest gain and lowest noise amplifiers as close to the antennas as possible in order to minimize the noise components associated with later stages of the system. The calculation for noise figure, the amount of noise power from input to output greater than the signal gain, of a cascaded system takes the form:

$$NF_N = NF_1 + \frac{NF_2 - 1}{G_1} + \frac{NF_3 - 1}{G_1 G_2} + \cdots + \frac{NF_N - 1}{G_1 G_2 \ldots G_{N-1}} \tag{2–2}$$

where NF is the noise figure for a specific cascaded stage and G is the gain for a specific cascaded stage. From equation 2–2 it is easy to see that the noise figure and gain for the first stage have the greatest effect on the entire system's noise performance. The noise figure for the first stage is unmodified and should be as low as possible. Each stage following the first stage has its noise figure divided by the product of the gain for each previous stage making the choice of a high gain stage ideal because it will limit the noise associated with each following stage. Therefore, it is not necessary to place the variable gain stage in the preamplifier instead of the line receiver because the gain of the first stage will have the most dominant effect in the overall system noise figure. The typical variable gain stage for a radiometer system are 0, 10, 20 and 30 dB while a portable system will have either the same gain or a linear gain of 1, 3, 10, and 30. Given the environments for a radiometer system it is better able to take advantage of a quieter background environments.

2.1.5 Anti-Alias Filter

Filtering the signal prior to digitization is necessary in order to prevent aliasing. Sampling theory [*Oppenheim and Schafer*, 1975, p. 28-29] states that in order to prevent higher frequency components from aliasing, the input signal's frequency should not exceed half of the sampling rate, often referred to as the Nyquist rate. A typical VLF system will have an elliptical filter with a passband from 0 - 40 kHz and a stop-band from 40 - 60 kHz reaching an attenuation in excess of 90 dB at 60 kHz. The sharp attenuation comes at a cost of a ripple in the pass and rejection bands but the cost is justifiable given the prevention of aliasing and the rapid transition from the stop band to the pass band. All VLF systems operated by the University of Florida use some form of elliptical filter in order to prevent aliasing; the older radiometer systems have had a filter box added due to the changes in sampling systems over time.

2.1.6 Analog to Digital Converter and Timing

The current digitizer used by the University of Florida in conjunction with a VLF receiver is an off the shelf multifunction Data Acquisition (DAQ) card. It offers 8 differential 16-bit inputs giving a spurious free dynamic range (SFDR) of 98 dB at 10 kHz. Unfortunately, most of the systems are not recorded differentially. Due to the number different systems in the field and changing ADC

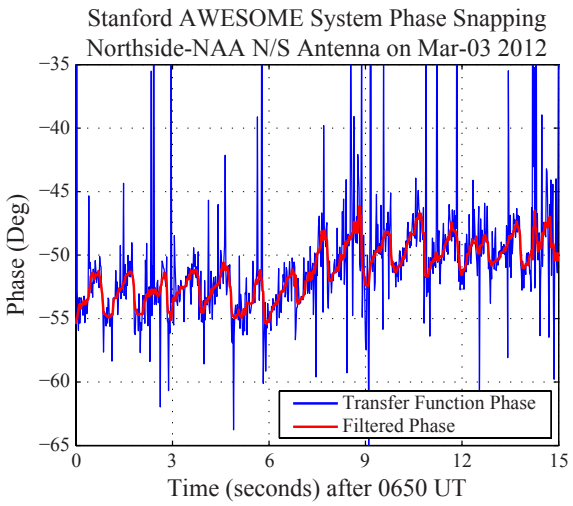

Figure 2-3. An example of phase snapping. The transfer function phase is plotted in blue and the filtered phase in red in order to highlight the distinct nature of the phase snapping.

connectors, a break out box has been implemented that takes a single ended BNC input and then converts to the correct ADC channels inputs.

Another important component of the acquisition system is the timing used. As opposed to the Stanford system which uses an internal GPS clock, the Florida systems use an external GPS receiver. The receiver is expensive but the benefit is the absolute GPS timing accurate to withing 30 nanoseconds RMS of UTC and a reference frequency stability better than 1×10^{-12}. Both the sampling frequency, 100 kilosamples per second, and a 1 pulse per second pulse train are supplied to the DAQ card by the GPS receiver for timing purposes. Accurate timing is not something that Stanford can claim for its integrated GPS system due to the off the shelf embedded system that is used in the AWESOME system. Numerous instances have been found where the phase associated with Stanford systems will ‘snap’ as the 1 pulse per second timing signal is synchronized with the embedded GPS time.

Figure 2-3 is an example of phase snapping in the transfer function phase for the Northside-NAA receiver-transmitter pair on March 3rd, 2012 received on the North/South antenna. The transfer function phase is sampled at 50 Hz in this case. The 50-Hz sampling frequency is important to detect this ’snap’ because lower output frequencies averaged over the effect. In the figure, the

phase snaps once per second as the AWESOME sample clock is updated/corrected by the GPS-derived 1-PPS reference. The result is a saw tooth like function for the calculated phase. This effect likely varies with time, as the specific frequency of the AWESOME oscillator likely depend on temperature and other environmental parameters. In contrast, the University Florida systems use an external GPS system which is properly synchronized to GPS timing and does not produce phase snapping. In the past, some of the Florida systems used the integrated GPS of the AWESOME system but over time, all of the receivers have been upgraded to use the external GPS unit.

In conjunction with the DAQ card, a custom program is used to interface with the card and provide functionality to record both broadband data sampled at 100 kHz and multiple narrowband VLF channels. A typical system will record nearly continuously for 23 hours a day, every day of the year. The amount of data generated by a site is about 30 GB/day or approximately 12 TB/year for two channels of broadband data and several narrowband channels.

2.1.7 System Performance

The most recent Florida VLF system was built in late 2013 and represents the state-of-the art in VLF system design. Several changes were made from the previous version to enhance performance such as upgraded components and newer components, new implementations of the line drivers and line receivers, external controlled gain, and a better designed anti-alias filter. One of the primary benefits of this system was that is was built to work in conjunction with a 24-bit ADC. The typical ADC used in conjunction with VLF receivers have 16-bit resolution and are limited to a dynamic range of 91.5 dB compared to the lower limit of 118 dB for the 24 bit ADC.

A performance analysis of this VLF system uses figures of merit typically associated with ADCs which are almost standard and have been detailed by [*Kester*, 2008]. Quantifying the system was accomplished by injecting a 1 kHz sinusoid in to the preamp and increasing the amplitude to the point just before the output starts to clip. With the aid of the 24 bit ADC, a noise floor of -121.9 dB and an SNR of 118.7 dB was achieved. (See Figure 2-4 for a graphical definition of terms.) The Signal-to-Noise Ratio (SNR) is calculated by taking the ratio of the root-mean-squared (rms) signal amplitude, f_0, compared to the mean of the root-sum-square (rss) of the noise level

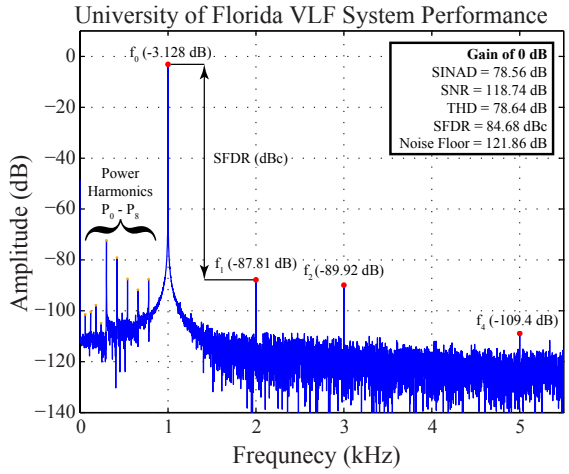

Figure 2-4. The output spectrum associated with a 1 kHz input sinusoid and the resulting performance parameters associated with the system.

excluding harmonics. For a VLF system in particular, a distinction is made between the harmonics associated with the test signal, f_{1-4} and the power harmonics, P_{0-8}. The 60 Hertz noise and its associated harmonics are not generated by the system, but couple in through the input due to their near omnipresence in a laboratory environment. Steps can be taken to minimize their effects but it is difficult to remove power noise completely and for the purposes of quantifying the system performance, power harmonics have been excluded for the associated harmonic calculations.

The Signal-to-Noise-and-Distortion (SINAD) is calculated as the ratio of the rms signal to the rss of the other spectral components including harmonics giving a value ranging from 78.6 dB to 93.54 dB, depending on gain setting. The SINAD performance is largely impacted by the presence of the first, second, fourth, and fifth harmonics of the input signal. Similarly, the Spurious Free Dynamic Range (SFDR) relative to the carrier is defined as the ratio of the rms signal to the rms value of the worst spurious signal which in this case happens to be the first harmonic resulting in a

Table 2-2. The University of Florida's VLF system performance at different gain settings.

Gain (dB)	SINAD (dB)	SNR (dB)	THD (dB)	SFDR (dBc)	Noise Floor (dB)
0	78.6	118.7	78.6	84.7	-121.9
10	87.9	104.0	89.4	89.4	-113.0
20	93.5	93.5		93.5	-103.1
30	90.3	90.30		92.4	-93.3

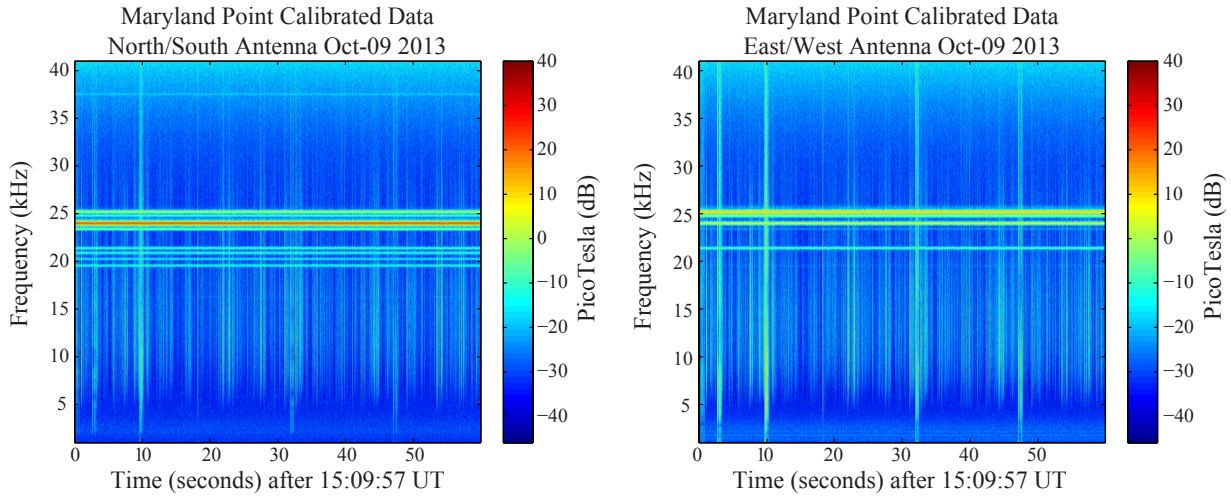

Figure 2-5. Calibrated one minute spectrograms for the north/south and east west antennas at Maryland Point on October 9th, 2013.

SFDR ranging from 84.7 dBc to 93.6 dBc, depending on gain setting. While both the SINAD and SFDR do not fully utilize the available range for a 24-bit system, they come close to the absolute limit of the 16-bit system, which has a dynamic range of 91.5 dB. A tabulated list of the receiver performance at difference gain settings is shown in Table 2-2.

Testing using the 24-bit system was performed in the field at Maryland Point, Maryland. Figure 2-5 shows the calibrated one minute spectrograms for the north/south and east/west antennas oriented along the north/south and east/west magnetic field lines respectively. Both spectrograms clearly exhibit lightning-generated sferics as vertical lines and VLF transmitters as horizontal lines. NAA (24 kHz) is strongest on the north/south antenna due to receiver's position relative to Cutler, Maine and NLM (25.2 kHz) is strongest on the east/west antenna.

2.2 VLF Hardware Applications

Building a fully functional prototype from a schematic should be a simple processes but there are numerous pitfalls that can prevent a circuit from working or lead to degraded performance. These issues are often referred to as applications problems and are not typically covered in a formal sense. Often the subject is relegated to application notes, white papers, or sections within a datasheet. One of the better texts on the subject is appropriately subtitled, *A Handbook of Black Magic* [*Johnson and Graham*, 1993]. Several of the larger integrated circuit manufacturers have

their unique take on applications notes such as Texas Instruments' *Infrequently Asked Questions* or Analog Devices' *Analog Dialogs*. Both companies also have detailed handbooks [*Carter and Brown*, 2001; *Jung*, 2006] on operational amplifiers that delve in to the finer details associated with several of the problems that are not typically covered in a classroom. The topics covered include component placement, grounding, power distribution, power supply decoupling, clipping, and capacitive loading. Poor component placement for feedback resistors and power supply decoupling capacitors can lead to additional parasitics which result in degraded performance at higher frequencies. Ground, while often considered a omnipresent zero voltage source, in practice will vary slightly at different locations. The presence of long ground traces or different paths can lead to ground loops and ground bounce, which degrade performance. Power distribution networks can also lead to problems where feeder networks distort or attenuate regulated power, resulting in a performance loss. Power supply decoupling goes hand in hand with feeder networks because there are strict requirements (within 100 mils of a power pin) for the placement of ceramic decoupling capacitors in order to prevent instability. While all of these considerations are important for the design of a VLF receiver, despite the low frequencies involved, clipping will be touched on because of its effects on the signal processing and capacitive loading will receive a thorough analysis. Detailing all of the applications problems associated analog circuits is beyond the scope of this work.

2.2.1 Clipping

A system clips when the received voltage is higher than the ADC can handle, typically above its power rail. The ADC used in the Florida systems are limited to ± 5 Volts and the received voltages are typically well below this limit. Often, the gain of the system is specifically chosen to prevent clipping. However, due to the high-power nature of lightning, which dominates the impulsive noise environment in the VLF band, there are cases where especially strong lightning sferics will produce a received voltage at the ADC in excess of the voltage supply rails. Figure 2-6 is an example of several clipped sferics at Palmer Station, Antarctica on March 3rd, 2013. Antarctic sites do not usually clip during strong lightning strikes, but Palmer is an exception because it

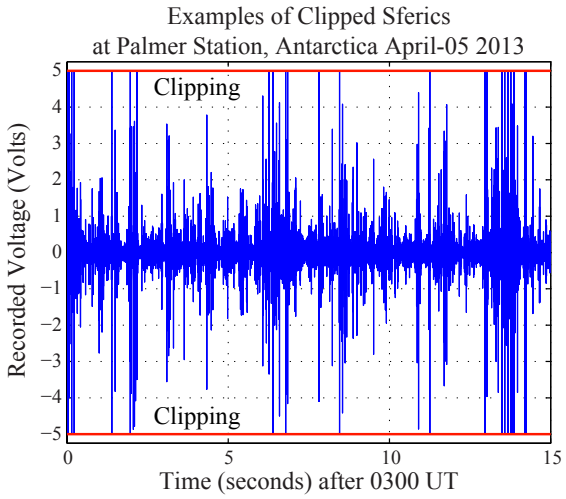

Figure 2-6. The receiver voltages at Palmer Station, Antarctica over a 15 second period. Several strong sferics occur which clip at the maximum allowable input to the ADC, the red line at ± 5 V.

is located much closer to typical thunderstorm locations than McMurdo Station and South Pole Station. Nevertheless, it is unusual for the VLF receiver at Palmer Station to clip as often as is shown in the figure. The problem of clipped sferics is more of an issue in the contiguous 48 states where lightning storms can occur in close proximity, or on top of, VLF receivers, resulting in clipped sferics and degraded date sets.

From a hardware perspective, it is important to protect the ADC from over voltages through the use of diodes or limiter chips, because surpassing an ADC's maximum rating can result in damage or degraded performance. Often, diodes will be used at the output of the preamp and the input of the line receiver to limit the effects of potential pick up caused by the instrumentation cable between the pre-amp and line receiver. And then a limiter chip will be placed in the signal path of the the line receiver in order to further protect the ADC after the gain stage. The cost of protection diodes or limiter chips is cheap compared to potential damage that the ADC might sustain over years deployed in remote locations.

The effect of clipping on signal processing is profound because a signal that clips cannot be trusted and results in an unreliable measure of narrowband VLF signal amplitude and phase. The detailed effects on the spread spectrum method will be discussed in Chapter 3.

2.2.2 Capacitive Loading

While several the applications issues are associated with high frequency effects, beyond the VLF range, it is important to note that the amplifiers have operation ranges well in to the MHz range and some times the GHz range. To neglect the high frequency effects would be disastrous for system performance. Poor layout and component choice can lead to instability and degrade performance in the frequency band of interest. Capacitive loading is an important topic because there is capacitance associated with everything, from driving a printed circuit board trace to the parasitic capacitance associated with a surface mount component's pad. A designer must be ever vigilant to account for common and uncommon capacitive sources to avoid destabilizing an amplifier. Analog Devices covers the issue well with two applications notes in their *Analog Dialogue*'s "Ask The Application Engineer" section [*Bendaoud and Marino*, 2004; *King*, 1997].

Often times an amplifier will be compensated to handle light capacitive loads in order to prevent poor performance but lacks the ability to drive a large capacitance due to the large performance penalty. With external compensation it is possible to use an amplifier that would otherwise be unable to drive a larger capacitive load by itself.

Figure 2-7 (a) shows a simple inverting amplifier with a capacitive load. An additional pole is formed buy the load capacitance, C_L, and the amplifiers output resistance, R_0, which degrades the amplifiers performance and lowers the Gain Bandwidth Product (GBP). The additional pole takes the form:

$$A_{loaded} = A\left(\frac{1}{1+j\frac{f}{f_p}}\right), \text{ where } f_p = \frac{1}{2\pi R_0 C_L} \tag{2–3}$$

where A is the open-loop gain. The additional pole increases the amplifier's rate of closure (ROC) as shown in Figure 2-7 b). A typical amplifier will reach its unity gain frequency, or GBP, with a -20 dB/decade slope for amplitude and a 90 degree phase lag after the dominant pole. In this case, the pole f_p adds another -20 dB/decade to the slope and 90 degrees phase lag which reaches the unity gain frequency sooner and causes instability in the form of oscillation.

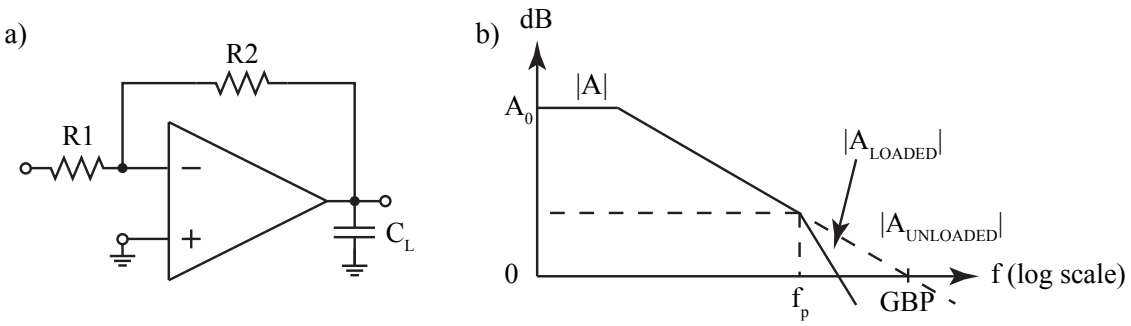

Figure 2-7. An example of a simple capacitive load, (a), and the resulting effects on the Bodé plot, (b). It is often the case that the capacitive load, C_L will not be a fixed component and instead take the form of a parasitic capacitance.

Capacitive loading is a problem, but there are several techniques available to correctly compensate an amplifier, as shown in Figure 2-8. Often, the correct solution is dependent on the design, given that there are advantages and disadvantages for each of the given techniques. In-loop compensation (2-8a) uses a small resistance R_x to decouple the output of the amplifier from the capacitive load, C_L, and a feedback capacitance, C_F, as a high frequency bypass of C_L. By solving for the poles and zeros associated with each capacitance individually and then equating each pole and zero pair, two equations can be generated that allow for R_x and C_F to be calculated:

$$R_x = \frac{R_0 R_{in}}{R_f} \tag{2–4a}$$

$$C_F = \left(1 + \frac{1}{|A_{cl}|}\right)\left(\frac{R_F + R_{in}}{R_f^2}\right) C_L R_0 \tag{2–4b}$$

The addition of the $1/A_{cl}$ term was determined experimentally [*Bendaoud and Marino*, 2004, p. 2]. Values in 2–4 are often listed in a components data sheet or can be inferred from plots. While this method prevents instability, it comes at the cost of bandwidth as the dominant pole is now set by two external components C_F and R_F.

In-loop compensation is a good method when amplifier characteristics are well known and high bandwidth is not needed. Because the VLF spectrum is relatively small and close to DC,

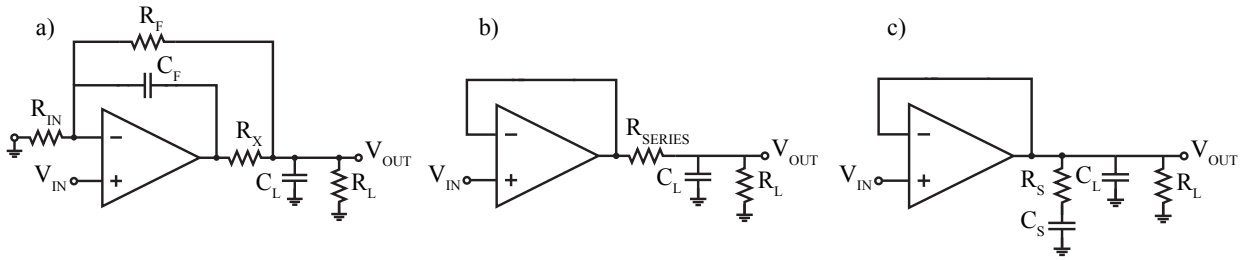

Figure 2-8. Different techniques for compensating voltage-feedback amplifiers with a capacitive load (a) in-loop compensation, (b) out-of-loop compensation, and (c) a "snubber" network.

bandwidth can be traded for capacitive drive ability without a loss of performance in the band of interest. There is a simpler method that involves out-of-loop compensation (Figure 2-8b) and uses a resistor R_{SERIES} between the output of the amplifier and the capacitive load. This isolates the capacitive load from the amplifier and creates a zero in the transfer function of the amplifier which helps to negate the effect of the load capacitance. Typical values for R_{SERIES} can range from a few Ohms to 50 Ohms and the value is often chosen so that the zero created is at least one decade below the unity-gain frequency. However, the performance penalty is great and the output signal will be attenuated by the voltage divider formed by resistance R_{SERIES} and R_L.

Another method (Figure 2-8c) uses a "snubber" network to compensate for the capacitive load. The resistor, R_s, loads down the output of the amplifier when peaking occurs and the capacitor, C_s, decreases the loading effect at lower frequencies by functioning as an AC short. This technique effectively snubs down the amplifiers gain at higher frequencies where the capacitive load has the greatest effect while minimizing the effect at lower frequencies. Unfortunately, the values for C_s and R_s are often determined through trial and error by observing the peaking frequency caused by an uncompensated load capacitance, choosing a resistance to limiting the peaking, and then picking a capacitance based on the resistance and peak frequency given by the equation:

$$C_s = \frac{3}{2\pi f_p R_s} \tag{2–5}$$

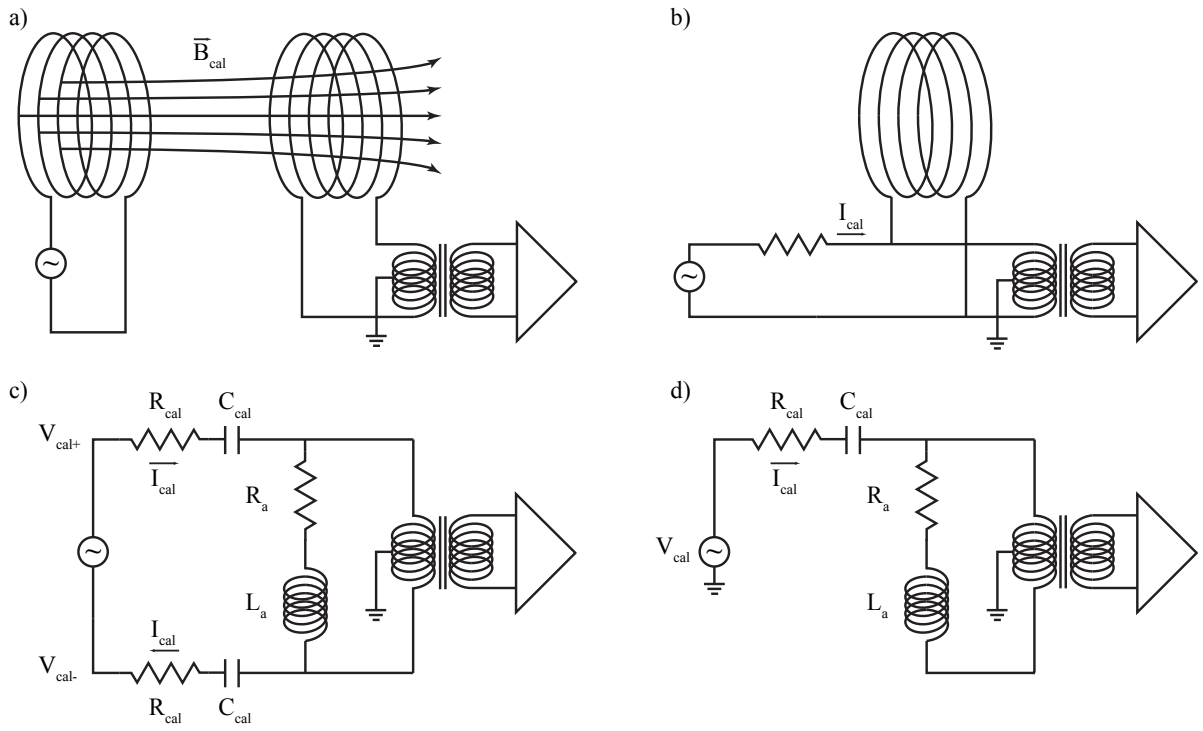

Figure 2-9. Different methods for calibrating a VLF receiver by either (a), inducing a known flux in the receiving antenna or (b), injecting a known current through a "dummy" antenna. For injecting a known current, there are two configurations, balanced, (c), and unbalanced, (d).

While this methods combines the advantages of both types of loop compensation, it is often difficult to experimentally measure the peak frequency and then choose resistor values in a realistic hardware environment. Compensation techniques used in VLF receivers are often in-loop or out-of-loop compensation, with a preference for in-loop compensation in a general case and out-of-loop compensation when driving the cable between the preamplifier and line receiver.

2.3 Calibration

The method for calibrating a system is incredibly important because it allows for the measured ADC voltage to be equated with a B-field signal strength at the antenna (in Tesla, or Tesla/$Hz^{1/2}$). There are two methods for calibration: inducing a known magnetic flux density using an external antenna, or injecting a known current through a "dummy" antenna at the preamplifier input. Figure 2-9 (a) and (b) show the two different configurations: inducing a known flux and injecting a known current, for calibration. While both methods should produce the same result [*Paschal*, 1988, p. 61],

injection proves to be the superior method for VLF receivers due to the logistics associated with inducing a VLF field and the large number of variables that adversely affect creating a uniform induced flux.

Figures 2-9 (c) and (d) show two injection circuit configurations for injecting a known current through a "dummy" antenna: balanced and un-balanced injection respectively. The "dummy" antenna is represented by a series resistance R_a and inductance L_a that have the same values as the actual antenna. While a real antenna can be used, it is beneficial to use a "dummy" antenna so that the input to the preamplifier is only the injected current (i.e., no VLF transmitter signals, sferics, or hum-lines interfere with the measurement). By carefully choosing the values for R_{cal} and C_{cal}, the relationship between the calibration voltage, V_{cal}, and its associated magnetic field, B_w, can be simplified to a few variables. For the balanced case, the equation takes the form:

$$V_{cal} = \frac{NAR_{cal}}{L_a} B_w \tag{2–6}$$

where

N	is the number of turns of the antenna
A	is the area of the antenna
R_{cal}	is a calibration resistor
L_a	is the inductance for the antenna and "dummy" loop
B_w	is the effective magnetic field

while for the unbalanced case, the equation takes the form:

$$V_{cal} = \frac{2NAR_{cal}}{L_a} B_w \tag{2–7}$$

Both cases, balanced and un-balanced, are useful for calibration but the method choice often depends on the application. A balanced configuration is advantageous when the calibration voltage is applied differentially, usually from a twister pair or by the differential output of a signal generator.

However, it can be difficult to supply a differential input signal in the field, and the un-balanced configuration allows for a relative measurement without having to deal with the common mode associated with the balanced configuration [*Paschal*, 1988, p. 62-64].

By using a series of simultaneously-generated tones, equally spaced in frequency and approximately constant in amplitude (often refereed to as a comb), for the calibration voltage, a system can be calibrated over its usable frequency band. A series of comb calibration tones can be generated using several digital components to create a pseudo random shift register [*Fraser-Smith and Helliwell*, 1985; *Golomb*, 1967; *Paschal*, 2013], which generates tones in the frequency domain with a fixed spacing. Given the advances with embedded processors, the digital components have been replaced with a single micro-controller that produces the same output. The frequency components can be described in terms of the shift clock, f_c, of a sequence with length m (of the form $2^n - 1$) where V_p is the peak value for a shift register. Given a square wave with a value of either $\pm V_p$ the output spectrum will have components every f_c/m Hz and the RMS amplitude of the k^{th} component is defined as:

$$V\left(k\frac{f_c}{m}\right) = \frac{V_p\sqrt{2}\sqrt{m+1}\,\operatorname{sinc}\left(\frac{k}{m}\right)}{m} \tag{2–8a}$$

$$\text{where,} \quad \operatorname{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{2–8b}$$

With $V_p = 2.5$ V, $n = 10$, $m = 1023$, and $f_c = 250$ kHz, tones are spaced 244.3793 Hz apart with an amplitude $V(244\text{ Hz}) = 0.1106\ \text{V}_{\text{rms}}$ and a variation in amplitude of only $V(34946\text{ Hz})/V(244\text{ Hz}) = 0.9682 = -0.28$ dB, giving only a slight variation over the band of interest.

With a known input, it is relatively simple to build a system that injects the calibration tones into the VLF preamplifier using a set of "dummy" antennas and measure the resulting output to calibrate the system. In practice, the actual amplitudes of the comb components are often assumed to be a fixed value, due to the inability to simultaneously measure the calibration tones as they are being injected at the preamplifier (e.g., when the source of the calibration tones is located within the preamplifier). In this case, one cannot realistically approximate the phase of each component,

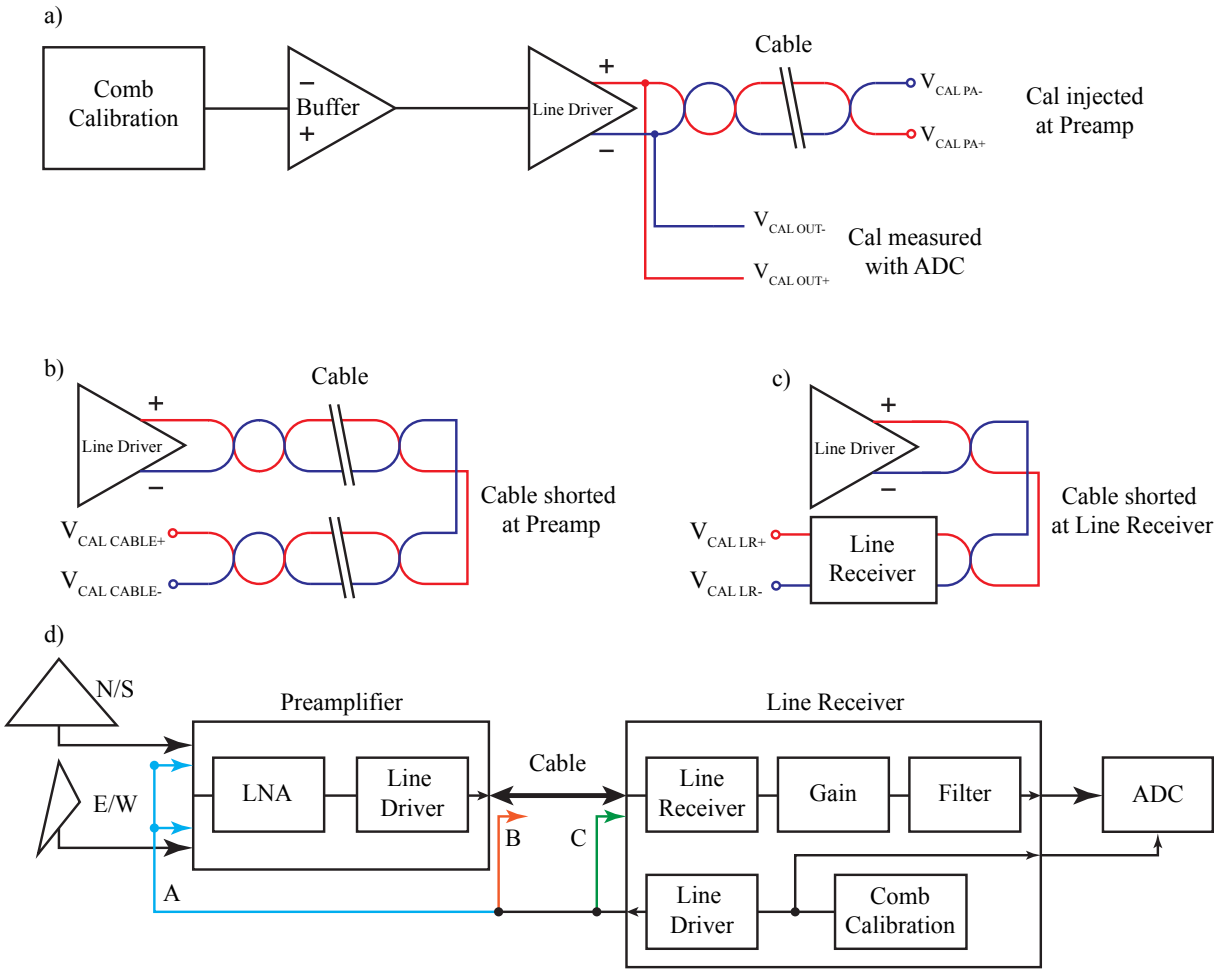

Figure 2-10. One level of abstraction below Figure 2-2 for three different calibration cases. The calibration tones are injected at the preamplifier , (a), down the cable and back , (b), and injected directly in to the line receiver, (c). The block diagram from Figure 2-2 has been replicated, (d), to indicate the different injection paths. Light blue refers to the path in (a), orange to the path in (b), and green to the path in (c).

making the calibration an amplitude-only calibration. In order to account for phase, great care must be taken to record the calibration tones as they are injected in to the preamp. It is therefore necessary to place the calibration tone generator in close physical proximity to the ADC at the line receiver.

Figure 2-10 shows a block diagram for a three step calibration. In all three steps, the calibration tones are simultaneously recorded by the ADC at the line receiver (not shown). The output of the calibration tone generator is buffered, due to the limited drive capability of the microcontroller,

and then driven differentially down a twisted pair to the pre-amplifier as shown in Figure 2-10a. Intuitively, it may seem best to buffer the output of the calibration tones for the ADC as well, but the loading effect caused by a fixed length of RG–58 cable and the input impedance of the ADC (typically 100 MOhms in parallel with 10 pF) has a negligible effect on the driver in the VLF band. The effects of the twisted pair connecting the line receiver to the pre-amplifier, a typical DC resistance on the order of tens of Ohms and a capacitance on the order of tens of nanofarads, has a much greater effect on line driver than tapping off to measure the calibration tones.

The output of the calibration tones can be represented by the initial calibration voltage defined in Equation 2–8, V_{comb}, multiplied by the transfer functions associated with the buffer, H_B, and line driver, H_{LD}, giving the equation for $V_{CAL\,OUT}$ as:

$$V_{CAL\,OUT}(f) = V_{comb}(f)\, H_B(f)\, H_{LD}(f) \tag{2–9}$$

Similarly, the voltage applied to the input of the preamplifier can be defined as the multiplication of Equation 2–9 and the transfer function associated with the cable, H_C. The equation for the calibration voltage applied at the input of the pre-amplifier takes the form:

$$V_{CAL\,PA}(f) = V_{comb}(f)\, H_B(f)\, H_{LD}(f)\, H_C(f) \tag{2–10}$$

The cable response, H_C, can be quantified by using the configuration in Figure 2-10b, where the measured output takes the form:

$$V_{CAL\,CABLE}(f) = V_{comb}(f)\, H_B(f)\, H_{LD}(f)\, 2H_C(f) \tag{2–11}$$

By shorting the cable at the preamp and taking the output at the line receiver the transfer function associated with the cable is quantified. Substituting Equation 2–9 into Equation 2–11 gives the cable transfer function, H_C. It is an approximation to simply use $2H_C$ in these equations. However, when the line driver is able to remain stable when driving both loads, single and double cable lengths, the approximation is valid. The approximation would be invalid if the line driver became

unstable, and the output of an unstable amplifier in any case would not be acceptable for calibration. In a laboratory setting, it is possible to compare the difference between H_C and $2H_C$, but they should be proportional. It is also possible to measure the cable attenuation using multi-meters and measuring the AC voltage at specific points throughout the system. However, this method only gives the amplitude response and does not provide a measure of the phase phase response. These types of measurements are also subject to the the variation between different types and styles of multi-meters making the measurements a good approximation at best.

With two steps, the tones injected into the preamplifier have been determined together with the transfer function associated with the cable. While injecting into the preamp, Figure 2-10a, the output of the line receiver is also recorded (not shown) in order to calibrate the entire system. This output voltage is now defined as $\mathrm{V_{CAL\ TOTAL}}$ and is described by the following equation:

$$\mathrm{V_{CAL\ TOTAL}}(f) = \mathrm{V_{CAL\ PA}}(f) H_{SYSTEM}(f) \tag{2–12}$$

Both the input to the system and the output are now known, and the transfer function (both amplitude and phase) associated with the entire system can be calculated. This is a superior method compared to the amplitude-only calibration method. For completeness, and to better calculate the performance of each individual section of the receiver, it is also of interest to determine the individual transfer functions associated with the preamplifier and the line receiver. By using the third configuration in Figure 2-10c, injecting directly in to the line receiver, and measuring the output, the transfer function associated with the line receiver, H_{LR} is determined. The resulting equation takes the form:

$$\mathrm{V_{CAL\ LR}}(f) = V_{comb}(f) H_{LR}(f) \tag{2–13}$$

and given that the system transfer function can be defined as:

$$H_{SYSTEM}(f) = H_{PA}(f) H_C(f) H_{LR}(f) \tag{2–14}$$

all three individual transfer functions can now be determined. It is important to note that this is a basic theory for calibration that has been used before on other system and will continue to be used in the future as it follows the simple method of a measured output referenced to a measured input. The distinction is the application to VLF receivers in the field which have not had a thorough calibration scheme, given the distributed nature of a system where the preamplifier and line receiver are separated by thousands of feet. In addition to the theory, the implementation benefits immensely if the system is designed with foresight for an amplitude and phase calibration scheme. The ability to tap off of certain points in the system is aided by well placed BNC connectors when the alternative is temporary wires placed at critical junctions. Extensive field work has shown that it is better to implement these features prior to production as opposed to attempting to add functionality to an already deployed system.

CHAPTER 3
A NEW SPREAD SPECTRUM VLF SIGNAL PROCESSING METHOD

In this chapter, a new spread spectrum MSK signal processing technique is described in detail and the performance of the method is evaluated in comparison with contemporary VLF signal processing techniques. First, an overview of the signal processing method is provided. Signal condition details are then discussed, followed by a detailed description of the signal modeling process. Lastly, the calculation of the output transfer function is described in detail, and the performance of the method is characterized as a function of SNR and compared with the results of other methods. One of the strongest metrics for comparison is the site-to-site transfer function. This transfer function can be directly calculated using broadband data (i.e., providing a reference "answer"), and the accuracy of each method can be directly evaluated by comparison with this "answer."

3.1 Signal Processing Overview

An example transmitted MSK phase trellis was shown in Chapter 1, Figure 1-4, and it is repeated here in Figure 3-1 (black trace) for ease of reference. The transmitted signal is constant in amplitude, while the phase trellis varies with time. The black trace is synchronized to the start of a second, both in terms of its initial phase and in terms of the bit start time. The yellow bit sequence in Figure 3-1 includes an initial phase offset ϕ, presumably due to the phase velocity of the signal, and it also exhibits an offset bit transition time, due to the group delay of the signal.

The spectrum of the transmitted MSK-modulated signal is shown in Figure 3-2, together with those for a continuous waveform (CW) signal, and a "squared" MSK signal for reference. CW waves are the simplest case with all of the frequency content located at a single frequency. Because of the simplicity of CW, it will be considered an upper limit for some of the statistics associated with this new method. MSK modulation is used by all of the transmitters analyzed in this work. "Squared" MSK, while not a standard form of modulation, is used in the Dowden/Adams method to generate two frequency measurements, at $\pm$50 Hz.

For the CW and the MSK spectra, the spectral content at each frequency excites propagating modes within the Earth-ionosphere waveguide, and each mode is excited with a different excitation

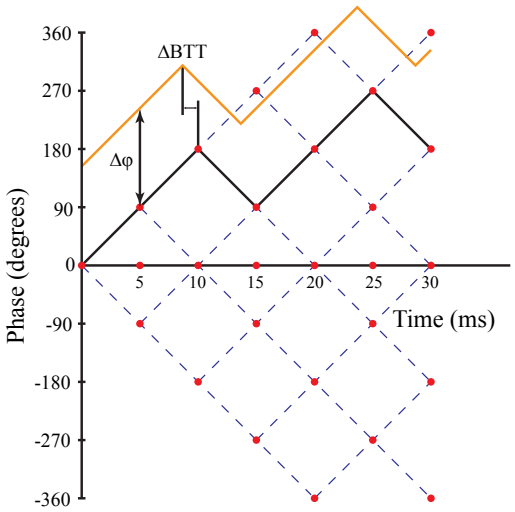

Figure 3-1. An example of the MSK phase trellis repeated from Chapter 1. The phase will vary due to ionospheric and transmission conditions resulting in a bit sequence that does not start at time zero with zero phase.

factor that depends on the geometry of the transmitter and the eigenangle of the mode. Each of the modes (for each of the frequencies) propagates within the Earth-ionosphere waveguide with different rates of attenuation, with different phase velocities, and with different group velocities. The signals eventually reach the receiving antenna and excite a voltage that depends on the inductance of the loop, the area of the antenna loop, the number of turns in the loop, and the altitude and orientation of the antenna. This voltage is proportional to the derivative of the sum-of-modes magnetic field component perpendicular to the plane of the antenna loop.

Given a measurement of the MSK-modulated signal, the spread-spectrum method developed herein provides an estimate of the transmitted signal. The transfer function (received/transmit estimate) is calculated and a weighted average is performed as a function of frequency. The output of this process is proportional to the transfer function of the Earth-ionosphere waveguide, weighted by the amplitude of the transmission (i.e., the sum-of-modes solution). In the following sections each of the important steps in the process are detailed and characterized.

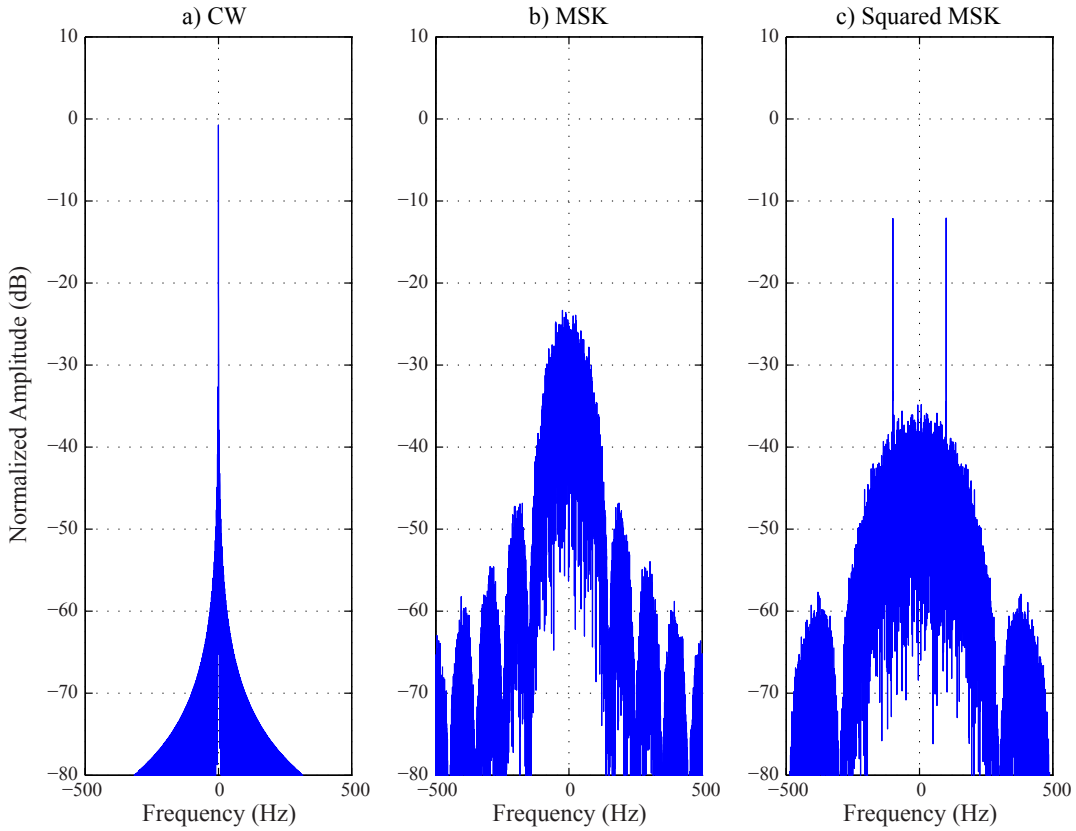

Figure 3-2. The modulation spectra for three different methods: (a) continuous wave (CW), (b) MSK, and (c) squared MSK.

3.2 Initial Signal Conditioning

An overview schematic of the processing methodology is shown in Figure 3-3. The signal incident upon the antenna can be expressed:

$$y_1(t) = \sum_n x(t) * h_{wn}(t) \tag{3–1}$$

where $x(t)$ is the time series of the transmitted waveform, and $h_{wn}(t)$ is the overall impulse response for the n^{th} mode of the Earth-ionosphere waveguide, including waveguide excitation factors, receiver excitation factors, and mode-coupling at waveguide boundaries. The signal is then processed by the receiver electronics and is sampled by the ADC. The resulting signal can be expressed:

$$y_2(t) = (y_1(t) * r(t))\, \text{Ш}(F_s t) F_s = \left(\sum_n x(t) * h_{wn}(t) * r(t) \right) \text{Ш}(F_s t) F_s \tag{3–2}$$

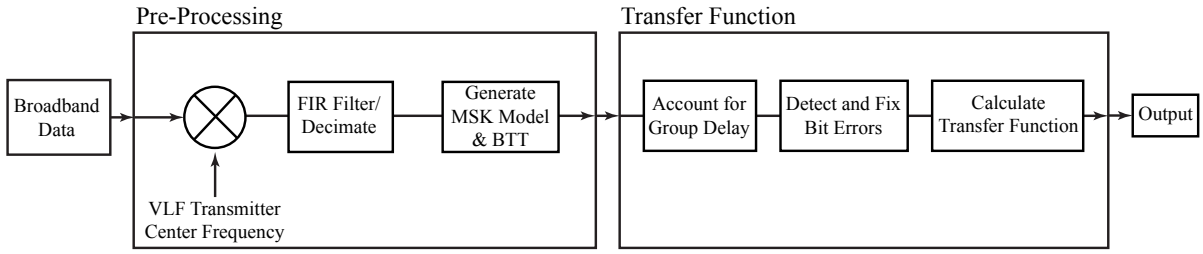

Figure 3-3. A block diagram for the processing method which shows the nuance associated with the pre-processing and transfer function processing blocks.

where $r(t)$ is the impulse response of the receiver, F_s is the sampling frequency of the ADC, and Ш is the sampling function defined as:

$$\text{Ш}(at) = \frac{1}{|a|}\sum_n \delta\left(t - \frac{n}{a}\right) \tag{3–3}$$

The signal is then mixed-down to baseband, low-pass filtered, and decimated by a factor of 100:

$$y_3(t) = \{[y_2(t)\,(\cos(\omega_c t) - j\sin(\omega_c t))] * b(t)\}\,\text{Ш}\left(\frac{F_s t}{100}\right)\frac{F_s}{100} \tag{3–4}$$

where $\omega_c = 2\pi f_c$ is the center frequency of the transmission of interest and $b(t)$ is the impulse response of the low-pass filter. $y_3(t)$ is the complex-valued, baseband time-domain signal prior to MSK signal processing.

The frequency-domain version of this signal can be expressed:

$$Y_3(f) = \left\{\left[\left(\sum_n X(f-f_c)H_{wn}(f-f_c)R(f-f_c)\right) * \text{Ш}\left(\frac{f-f_c}{F_s}\right)\right]B(f)\right\} * \text{Ш}\left(\frac{100f}{F_s}\right) \tag{3–5}$$

Assuming the signal is appropriately band-limited yields the following simplified equation:

$$Y_3(f) \approx \sum_n X(f-f_c)H_{wn}(f-f_c)R(f-f_c)B(f) \tag{3–6}$$

3.3 Construction of the Model Transmission

At this point, the phase of the baseband time-domain signal, $y_3(t)$, is used to estimate the transmitted MSK phase trellis (i.e., the bit sequence). Because the received phase trellis (the phase

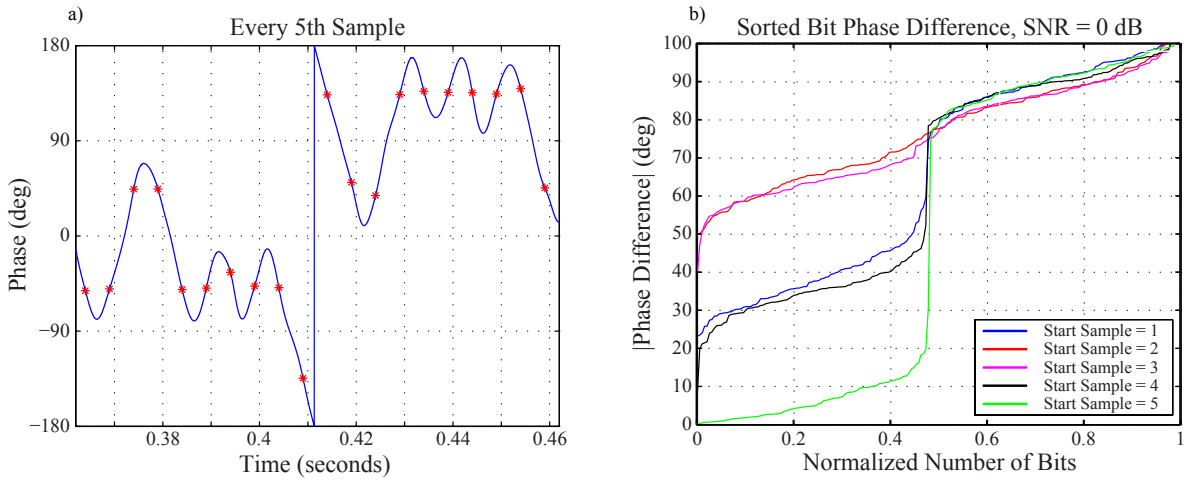

Figure 3-4. Calculating the The received phase with every 4th bit sample plotted in red, (a), and the sorted bit phase difference for different starting samples, (b).

of $y_3(t)$) includes the effective propagation delay of the Earth-ionosphere waveguide, the delay of the receiver electronics, and the delay of the low-pass filter (if any), the estimated transmission signal will also exhibit these same time delays.

These time delays are characterized by the bit transition time (BTT) of the received phase trellis. Calculating the BTT makes use of the statistics of the MSK transmission: the bits have a Bernoulli(1/2) distribution that is independent and identically distributed (iid). With a sampling frequency of 1 kHz, a single bit period is 5 samples (5 ms). To determine the sample (modulo 5) closest to the BTT, every fifth sample is analyzed. Figure 3-4 (a) shows the received phase and every 5th sample in red. It is clear that the highlighted sample is not the desired BTT, but in order to find the correct position, the phase difference between every fifth sample is calculated. The result is plotted in Figure 3-4 (b) where the absolute phase differences for each of the five starting bit positions are sorted and shown.

Because of the nature of the phase trellis as an Bernoulli(1/2) process, approximately 50% of the bit phase differences will be close to 90 degrees while the other 50% will correspond to the phase difference during a bit transition. The start position with the lowest bit phase difference per bit period yields the location of the middle-most sample, and thus identifies the BTT to within

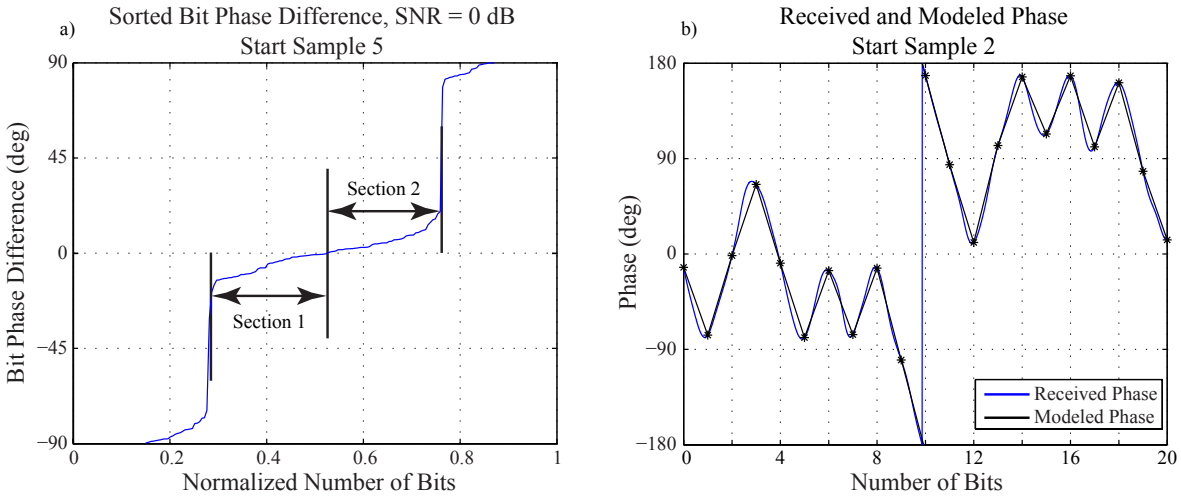

Figure 3-5. The sorted bit phase difference for start sample 5, (a), and the received and model phase when using start sample 2 for the model, (b).

one sample. The second lowest bit phase difference per bit period gives the sign associated with calculating the sub-sample BTT, which will be discussed shortly.

With the integer-valued BTT in hand, the sorted bit phase difference can be created, as shown in Figure 3-5a. Sections are identified based on the zero-crossing and the rapid change in phase near 25% and 75%. The phase difference, $\Delta\Phi$, is calculated as the median of Section 2 minus the median of Section 1. The sub-sample BTT is then calculated as:

$$\mathrm{BTT} = -\left(\frac{|\Delta\Phi|}{90^\circ} \pm 1\right)\left(\pm\frac{1}{2}\right) \tag{3–7}$$

with the $\pm$ sign determined by the second lowest bit phase difference per bit period. Using the resulting bit transition time, the entire phase trellis is reconstructed, evaluating bit values based on the median phase change over each bit period, resulting in a model phase as shown in Figure 3-5b.

Given the calculation of the BTT, one can also calculate the perceived group delay from transmitter to receiver. This value is used later to estimate the original time of transmission. Assuming the first bit transmitted was synchronized to the beginning of a second, the perceived propagation delay is $M(5\text{ msec}) + \mathrm{BTT}$, with M integer valued. Given the GPS coordinates of both the receiver and transmitter (and thus the distance between the two), M may be calculated as the integer which produces a perceived propagation speed that is closest to the speed of light. It is not unusual for

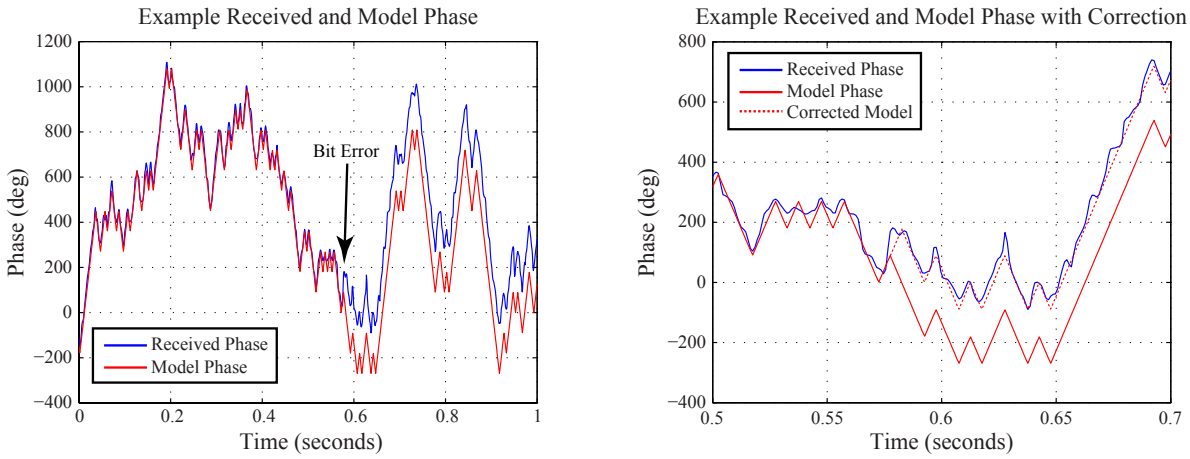

Figure 3-6. An example of a received and model phase trellis with a bit error, (a) and correction, (b).

this perceived speed to be greater than the speed of light. We note that it is only the *perceived* propagation speed, which may be distorted slightly by the interference between waveguide modes. The end result is the total delay, T, between transmission and the first perceived bit transition time.

3.4 Bit Error Correction

Because the MSK phase trellis is forced to be continuous at bit transitions, bit errors can have a profound impact on the estimated transmit spectrum, and thereby the output transfer function. Correcting bit errors is not a simple endeavor, however. While the statistics of an MSK phase trellis helped to determine the BTT and the model phase trellis, the statistics works against the process for correcting bit errors. A given bit can only be determine by the samples associated with a current bit. Previous or future bits cannot be relied upon to correctly determine the current bit. Figure 3-6 is an example of a bit error where the fitted model phase chose the incorrect bit polarity. The bit error results in an offset between the received and model phase of 180° and distorts the model MSK spectrum.

Shafer's method employs a two bit look-ahead scheme that compares the received and model phase and corrects the current bit if the phase exceeds 180 degrees. This is a novel method to correct for bit errors but takes somewhat of a statistical leap of faith. There are several cases, Figure 3-7 (a), (b), and (c), where the the original phase cannot be correctly reproduced by the

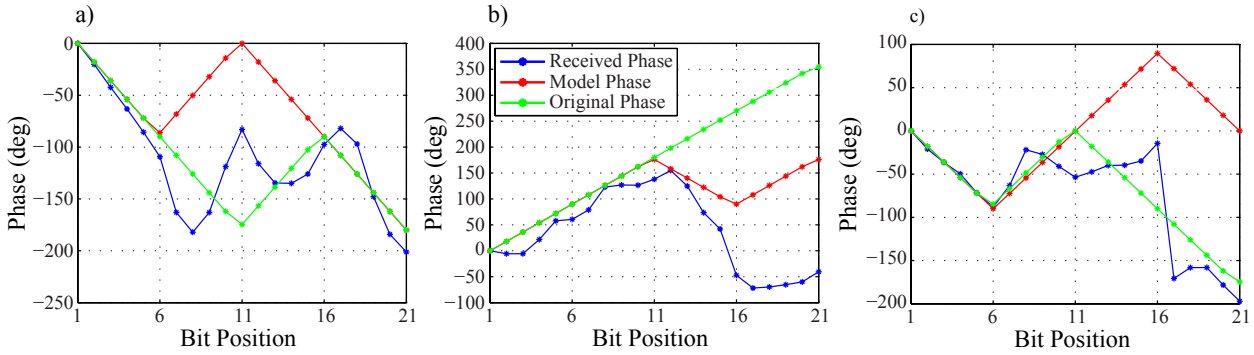

Figure 3-7. Three examples of unrecoverable bit errors found in the phase of MSK signals. With hindsight it is easy to see what the correct bit should be but the nature of the MSK phase trellis and its statistical properties limits the ability to correct bit errors. When choosing a bit there is only knowledge for the current bit and the choice cannot be correctly made based on previous or future bits.

model and received phase. The first case, (a), is an example of two bit errors in a row which effectively cancel out so that the phase spectrum is preserved. The second case, (b), shows an unrecoverable bit error where the received phase is affected by noise and the resulting model phase cannot be corrected. Using Shafer's method, a potentially correct future bit would be incorrectly corrected due to the criteria of a 180 degree phase difference. The final case, (c), is similar to (b) where the incorrect bit cannot be corrected based on the received phase. Both cases result in a future correct bit being erroneously corrected by Shafer's bit error correction scheme. It is impossible to correct every bit, but the bit correct technique used by this spread spectrum method focuses on maintaining the model spectrum without erroneously correcting bits.

The key difference with this method's approach to bit error correction is to find and correct all of the bit errors in a given dataset simultaneously. This solves the problem of having to iterate through the bit sequence one bit at a time. By comparing the phase to a filtered version of itself in conjunction with the modulo command, phase differences in excess of 180 degrees, which correspond to bit errors, can be isolated. The correct bit is then found and corrected accordingly. While this method is similar to Shafer's method it does not suffer the computational penalty associated with iterating through every bit and has the additional benefit of not over correcting through clever

use of the modulo command. Creating the transfer function is then a simple step with the correct group delay and model phase.

3.5 The Output Transfer Function

First, the estimate of the baseband version of the transmitted MSK time series is identified as $\hat{x}(t-T)$, where T is the total delay included in the model phase trellis (calculated in a previous section). $\hat{x}(t-T)$ is a unit amplitude signal with a perfect phase trellis. In order to make a proper comparison with the received signal, this signal must be filtered by $b(t)$. In practice, if $b(t)$ is a causal filter, this process will introduce an additional time delay in the output, and this effect is undesirable. A new filter constructed, $b'(t)$, whose amplitude response is the same as that of $b(t)$ but whose phase response is zero. The time-domain expression for the model signal is now:

$$\hat{x}_1(t-T) = \hat{x}(t-T) * b'(t) \tag{3–8}$$

with a frequency-domain expression:

$$\hat{X}_1(f-f_c)e^{-j(\omega-\omega_c)T} = \hat{X}(f-f_c)e^{-j(\omega-\omega_c)T}B'(f) = \hat{X}(f-f_c)e^{-j(\omega-\omega_c)T}|B(f)| \tag{3–9}$$

The desired transfer function is: $Y_3(f)/\hat{X}_1(f-f_c)$. In the digital domain, however, both time series must be limited to a finite number of points. This will be performed by windowing using a Hanning window. Additionally, the discrete Fourier transform behaves as though the input signal is periodic. The time-domain and frequency-domain expressions including these effects are:

$$y_4(t) = \left[\left(\{[y_2(t)\,(\cos(\omega_c t) - j\sin(\omega_c t))] * b(t)\}\,\text{Ш}\left(\frac{F_s t}{100}\right)\frac{F_s}{100}\right)w(t)\right] * \text{Ш}\left(\frac{t}{N}\right)\frac{1}{N} \tag{3–10}$$

and

$$Y_4(f) \approx \left[\left(\sum_n X(f-f_c)H_{wn}(f-f_c)R(f-f_c)B(f)\right) * W(f)\right]\text{Ш}\,(Nf) \tag{3–11}$$

where $w(t)$ is the windowing function, and N is the number of seconds over which the signal is now assumed to be periodic (this value can be selected using zero-padding, for instance). Applying

the same processes to the model transmission yields a time-domain expression:

$$\hat{x}_2(t-T) = \left[\left(\hat{x}(t-T) * b'(t)\right) w(t)\right] * \text{Ш}\left(\frac{t}{N}\right)\frac{1}{N} \tag{3–12}$$

and a frequency-domain expression:

$$\hat{X}_2(f-f_c)e^{-j(\omega-\omega_c)T} = \left[\left(\hat{X}(f-f_c)e^{-j(\omega-\omega_c)T}|B(f)|\right) * W(f)\right] \text{Ш}\,(Nf) \tag{3–13}$$

Noting that there is only access only to $\hat{X}_2(f-f_c)e^{-j(\omega-\omega_c)T}$, the final resulting transfer function, $Y_4(f)/\hat{X}_2(f-f_c)$ can therefore be expressed:

$$\hat{H}(f-f_c) = \frac{Y_4(f)}{\hat{X}_2(f-f_c)} = \frac{\left[\left(\sum_n X(f-f_c)H_{wn}(f-f_c)R(f-f_c)B(f)\right) * W(f)\right] \text{Ш}\,(Nf)}{\left[\left(\hat{X}(f-f_c)e^{-j(\omega-\omega_c)T}|B(f)|\right) * W(f)\right] \text{Ш}\,(Nf)\, e^{+j(\omega-\omega_c)T}} \tag{3–14}$$

where we note the additional $e^{+j(\omega-\omega_c)T}$ term in the denominator that counteracts the inherent delay (T) of the model transmission. Shafer's method does not account for this factor.

In the digital domain, the Ш function in the numerator and denominator simply represent frequency selection (or sampling), and they may be dropped from the equation above for clarity. The windowing function $W(f)$ essentially sets the minimum frequency resolution. For instance, a wide time-domain Hanning window produces a fairly narrow (several Hz wide) frequency-domain window, and has the effect of smoothing the frequency response. Suppressing these items, and assuming that $\hat{X}(f)$ is a perfect model of $X(f)$, the resulting transfer function is an approximate measure of the transmitted signal magnitude times ionospheric transfer function times the receiver response:

$$\hat{\hat{H}}(f) \approx \sum_n |X(f-f_c)|H_{wn}(f-f_c)R(f-f_c) = |X(f-f_c)|R(f-f_c)\sum_n H_{wn}(f-f_c) \tag{3–15}$$

The receiver response can be removed from the above equation using the amplitude *and* phase receiver calibration described in Chapter 2. The result is the sum-of-modes amplitude and phase, assuming the transmission at each frequency was transmitted with zero phase (as for the CW case). In practice, this transfer function is noisy, and a frequency averaging technique is implemented to smooth the response.

The output transfer function used for the spread spectrum processing method is defined in [*Ljung*, 1987, p. 146] as the Empirical Transfer Function Estimate (ETFE). The derivation of the transfer function is discussed in [*Shafer*, 1994] and an alternate version will be shown here.

The ETFE can be expressed:

$$\hat{G}(f) = \frac{\sum_{BW} |\hat{X}(f)|^2 \hat{\hat{H}}(f) \Delta f}{\sum_{BW} |\hat{X}(f)|^2 \Delta f} \tag{3–16}$$

such that frequency averaging is performed, weighted by the signal-to-noise ratio. The end-user can select the frequencies over which to average. We note once again that $\hat{\hat{H}}(f)$ includes a linear phase term in the denominator that Shafer's method does not account for. One important point, however, is that this ETFE predicts the "best" amplitude and phase under the assumption that the amplitude and phase are constant across the bandwidth BW. Because our transfer function can include large phase slopes with frequency, this particular ETFE is not valid. Instead, the linear phase progression with frequency is temporarily removed from $\hat{\hat{H}}(f)$, the averaging is performed, and then the linear phase progression is re-introduced to $\hat{G}(f)$. For averages centered on f_c, the linear phase progression does not affect $\hat{G}(f)$, but for averages centered on frequencies other than f_c, the phase can be shifted significantly.

In the end, the user may select the output sampling frequency, the frequency spacing, and size of bandwidth averaging, and these choices can produce visually distinctly outputs. The desired output will largely depend on the dataset being analyzed; transient events can vary from seconds to hours while ambient data depends on the length of the dataset or the time of day. Figure 3-8 is an example of ambient data for the NPM transmitter received at Palmer Station, Antarctica on the North/South antenna on April-05 2013 from 0030-2330 UT. The transfer function has been plotted two different ways. The first row, (a) and (b), is the transfer function for the the entire 200 Hz bandwidth with an output sample frequency of 2 Hz and a frequency spacing of 0.5 Hz for amplitude and phase. With the normalized amplitude and phase there appears to be a slight variation with frequency but it is hard to tell given the time scale used. The second row, (c) and (d), is the same transfer function averaged over the entire 200 Hz bandwidth. In the single frequency

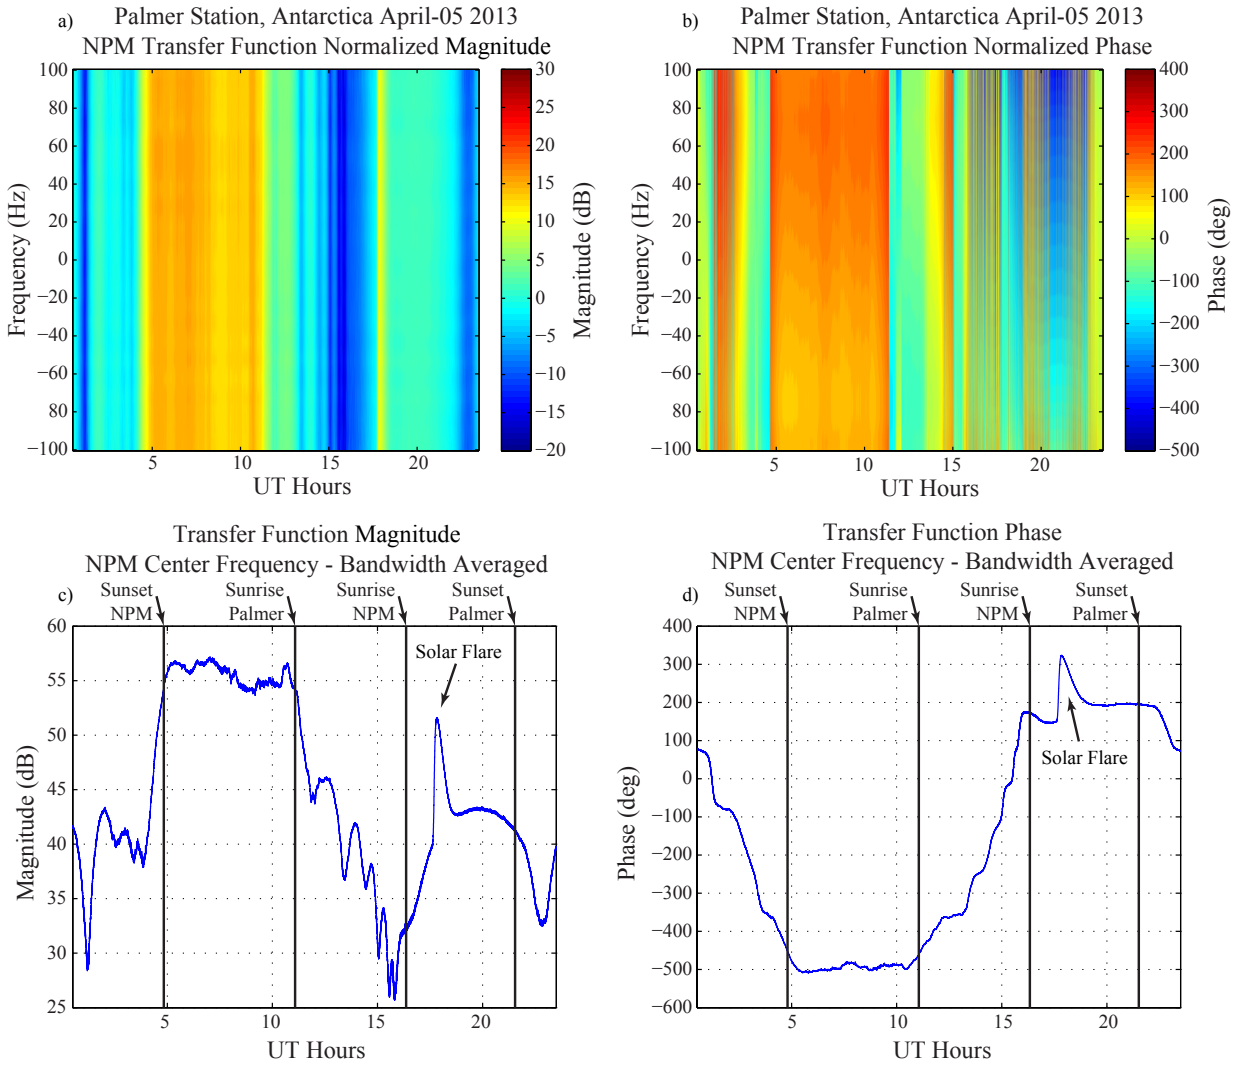

Figure 3-8. An example of the spectrum spectrum transfer function for ambient data over a 23 hour period for the NPM to Palmer Station, Antarctica path on April-05 2013. Normalized amplitude for the entire 200 Hz bandwidth is plotted in (a) while the bandwidth averaged received amplitude is plotted in (c). Similarly, the normalized phase is plotted in (b) and the bandwidth averaged received phase is plotted in (d).

figures it is easy to see that there is an amplitude and phase perturbation, around 1800 UT, which is caused by a solar X-ray flare. Ambient data will serve as an excellent metric to compare different processing methods but a statistical analysis of the method as a function of SNR will show the method's ability to correctly recover the amplitude and phase.

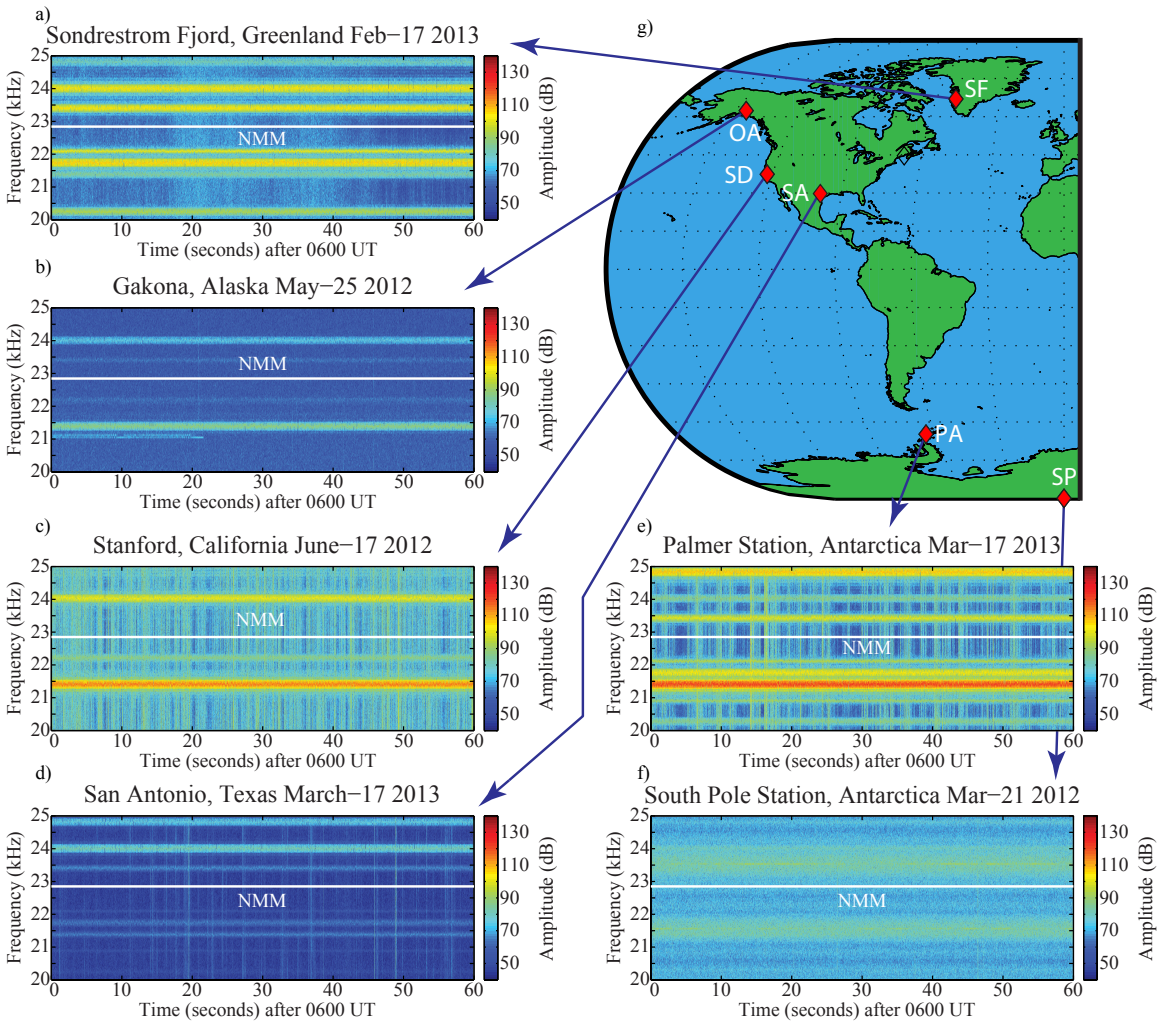

Figure 3-9. A plot of different background noise spectra for a variety of receiver locations from as far North as Greenland to as far south as the South Pole. Each receiver has its own background noise environment as well as the different VLF transmitter bands based on its location. By adding a random MSK signal, NMM at 22.8 kHz, at an unused VLF band and varying the SNR, the ability of the spread spectrum method to recover the correct amplitude and phase can be quantified.

3.6 Method Performance

Given the distributed nature of the network of receivers used to record VLF data there will be distinct differences in the noise background at each receiver. Even extremely remote locations, such as Antarctica, can suffer from spectral pollution due to the proximity to station equipment or noisy power lines despite the best efforts to prevent these problems. Figure 3-9 (a) - (f) shows a variety of background noise environments for the 20-25 kHz band. By adding a random MSK signal to a given noise profile the ability of the spread spectrum method to correctly recover the original amplitude and phase can be quantified as a function of SNR. Choice of the center frequency for the random MSK signal is dictated by the noise backgrounds and other VLF transmitters. In this case, 22.8 kHz is selected because of the relatively clean background spectrum at that frequency and its corresponding 200 Hz bandwidth. This artificial transmitter has been named NMM to distinguish it from other real transmitters.

The simplest case of adding a signal to a noise background and analyzing the resulting output is for a CW signal added to a impulsive noise environment. A CW signal has no modulation and exists solely at one frequency and the results of its recovered amplitude and phase compared to the original signal should serve as the upper boundary for performance of what is achievable by the spread spectrum method. The performance of a CW signal is tested by generating a large number of trials for different SNR levels in order to quantify the ability to correctly recover the initial amplitude and phase in a variety of noise environments. In order for the results to be statistically significant, there must be a larger number of trials which is easily accomplished by utilizing a long period of background noise data. With 30 minutes of background noise at each location, sampled at 100 kHz, it is easy to use one second of data as an individual trial giving 1800 trials. A CW signal with a random phase and a duration of 30 minutes sampled at 100 kHz is scaled to a pre-determined SNR and added, second by second, to 30 minutes background noise data. The resulting combined data set is then processed to extract the CW signal in the presence of noise and compared against the original signal. The results are shown in Figure 3-10 from 0 to 30 dB SNR

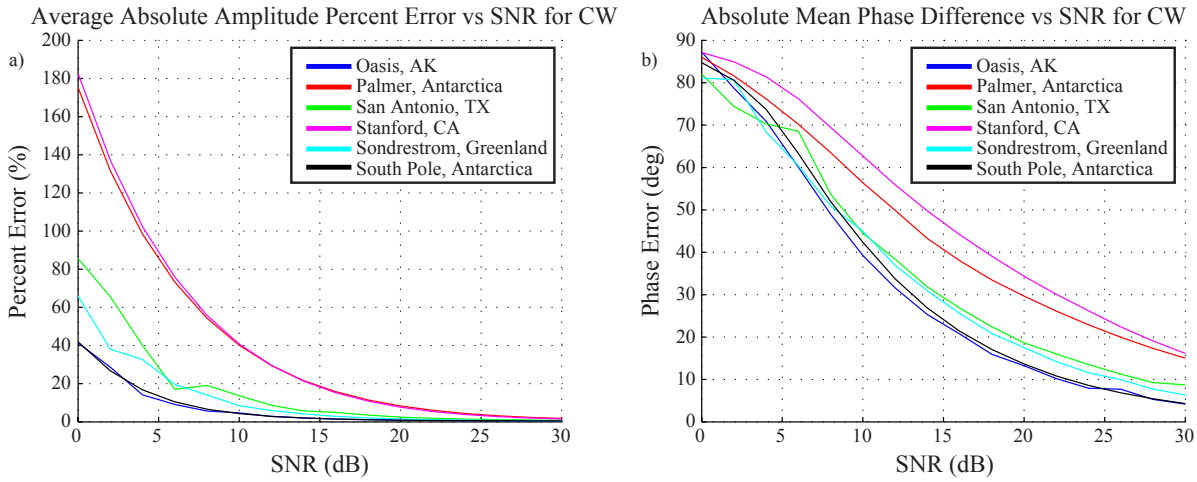

Figure 3-10. The average absolute amplitude percent error , (a), and the absolute mean phase difference, (b), versus different SNR levels for CW signals

for the average absolute percent error of amplitude, (a), and the absolute mean phase difference, (b).

The performance varies with receiver site with the best amplitude performance at Oasis and South Pole, San Antonio and Sondrestrom falling in between, and Stanford and Palmer having the worst performance. There is a similar grouping for phase based on location but the traces are overall better grouped when compared to amplitude. These results give a strong indication that the noise performance at each site will have a strong effect on the ability to correctly recover the amplitude and phase of even the simplest cases, such as CW, and will have a pronounced effect on the spread spectrum's method ability to correctly recover a randomized MSK signal.

Figure 3-11 shows the average percent of noise in the sixth standard deviation for each background noise environment, as a measure of the impulsiveness of the noise environment. The worst performance for the CW cases stems from the higher impulsive noise environments, at Palmer and Stanford Dish, where the problem is further compounded by clipping in this particular data set. The large amount of lightning strikes that occur in a given second in this data set produces an impulsive noise environment at each receiver which is greater or lesser depending on the receiver location relative to the thunderstorms. For the data set considered here, the extremely sensitive receivers at Palmer and Stanford Dish clipping occurs much more frequently than at the other sites,

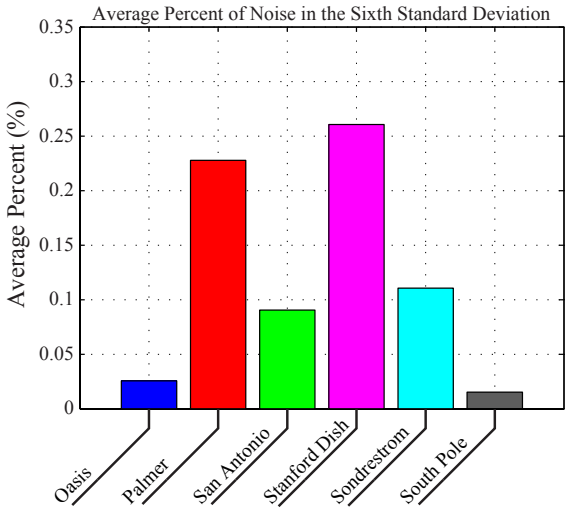

Figure 3-11. The average percent of noise in the sixth standard deviation for different noise backgrounds.

distorting the frequency spectrum and hindering the ability to correctly recover the amplitude and phase of both CW and MSK signals.

The process for adding an MSK signal to the background is similar to the CW case but for an MSK signal the bit trellis, initial phase, and bit transition time must also be randomized. A random MSK signal is created with a set of initial random variables and a duration of 30 minutes sampled at 100 kHz. It is scaled to the appropriate SNR level, and added to the noise background one second at a time, giving 1800 trials for sixteen different SNR levels from 0 to 30 dB. Figure 3-12 shows the results for the spread spectrum method for amplitude and phase.

The average absolute percent error for amplitude as a function of SNR plots shows good performance for Oasis and South Pole with a massive reduction in percent error below 10 dB SNR when comparing the traces with (solid lines) and without (dashed lines) error correction. There is similar performance increase with San Antonio and Sondrestrom, but the difference with and without error correction persists through to a much higher SNR. Unfortunately, all four previous receivers do not converge to a percent error of zero due to the impulsive noise environment. Palmer and Stanford both have extremely poor performance due to the excessive amount of impulsive noise that leads to hard clipping in this data set. This effect additionally results in worse performance

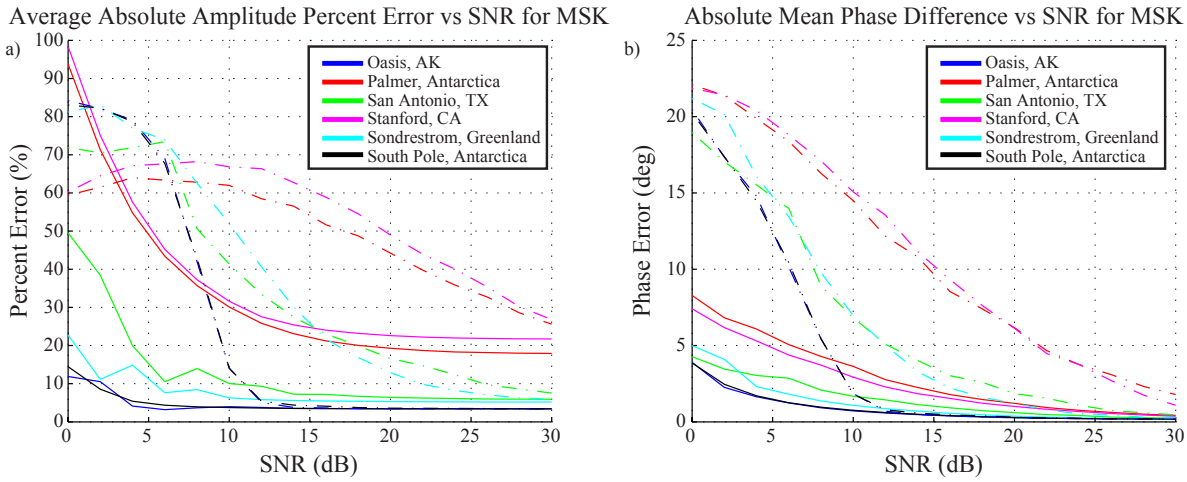

Figure 3-12. The average absolute percent error versus SNR, (a), and the absolute mean phase difference versus SNR, (b), with, solid lines, and without, dashed lines, error correction.

for the corrected versus un-corrected traces at low SNR. The amplitude percent error for these two receivers also converges to a percent error closer to 20%! The results for the percent error of amplitude for Palmer and Stanford illustrate the destructive effects of clipping on the recovered amplitude and should be avoided at all costs in the hardware implementation.

The ability of spread spectrum method to correctly recover the phase is not as adversely affected as the amplitude. There is good performance increases for all four receivers when comparing the traces with and without error correction and a convergence toward a phase difference of zero at higher SNR levels. Once again, the receivers are grouped by their noise environment where Oasis and South Pole have the lowest phase error with and without error correction followed by San Antonio and Sondrestrom, and finally Palmer and Stanford with the worst performance of the six for the data set considered.

The method's performance with bit errors and the the bit transition time calculation has been plotted in Figure 3-13 (a) and (b). As has been previously stated, bit error correction is an important component of the spread spectrum method because it helps to maintain the MSK phase spectrum. In Figure 3-13 (a) there is a clear distinction between the cases with (solid lines) and with out (dashed lines) error correction. While both Palmer and Stanford fair the worst as far as un-corrected performance, the corrected traces are closely grouped with the other receivers making

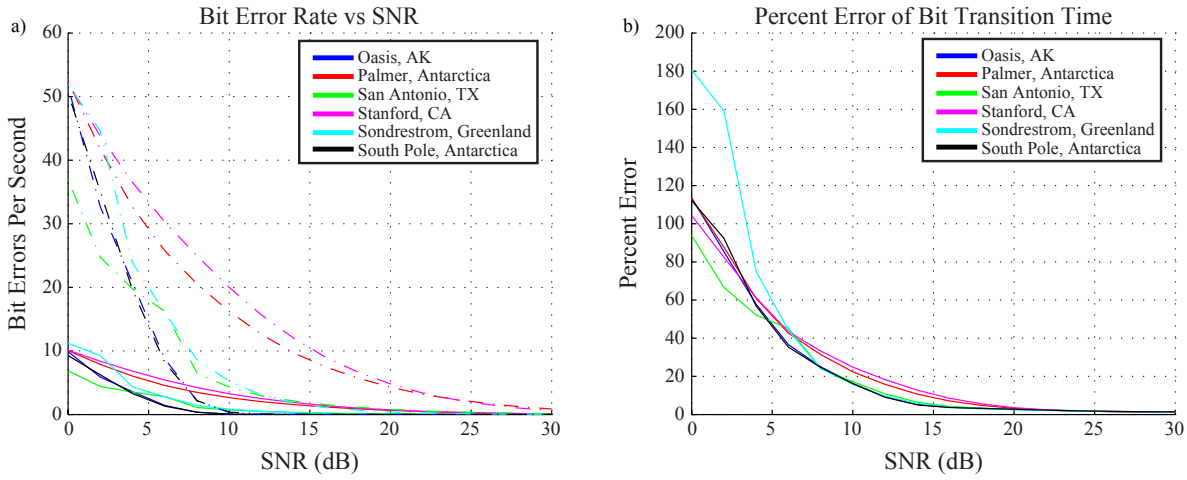

Figure 3-13. The Bit Error Rate, (a), and the Bit Transition Time, (b), plotted as functions of SNR for six different receiver locations. The Bit Error Rate is plotted with, solid, and without, dashed, bit error correction.

the error correction process mostly independent of noise. Nevertheless, the noise prevents rapid convergence to zero bit errors. The two best receivers, Oasis and South Pole, reach a zero bit error rate at nearly the same time, 10 dB SNR, for the corrected and un-corrected traces but the corrected traces have a much lower error rate for lower SNR values. The remaining two receivers fall somewhere in between with better and worse performance at lower SNR but ultimately taking longer to converge to a bit error rate of zero.

The percent error for the bit transition time, Figure 3-13 (b), shows a similar agreement with the bit error rate. All six traces are closely grouped and the performance for each receiver varies predominately by its noise profile. Overall, the bit transition time calculation converges to 1.5% (i.e., better than 75 microsecond accuracy) and is largely immune to the noise background.

3.7 Comparison of Processing Methods: The Ambient Case

With the methodology behind the spread spectrum processing technique established it is now important to focus on comparing different methods. A variety of methods were introduced in Chapter 1 and they will be analyzed in depth and compared to the spread spectrum method by taking a site-to-site transfer function. The site-to-site transfer function is the absolute difference in

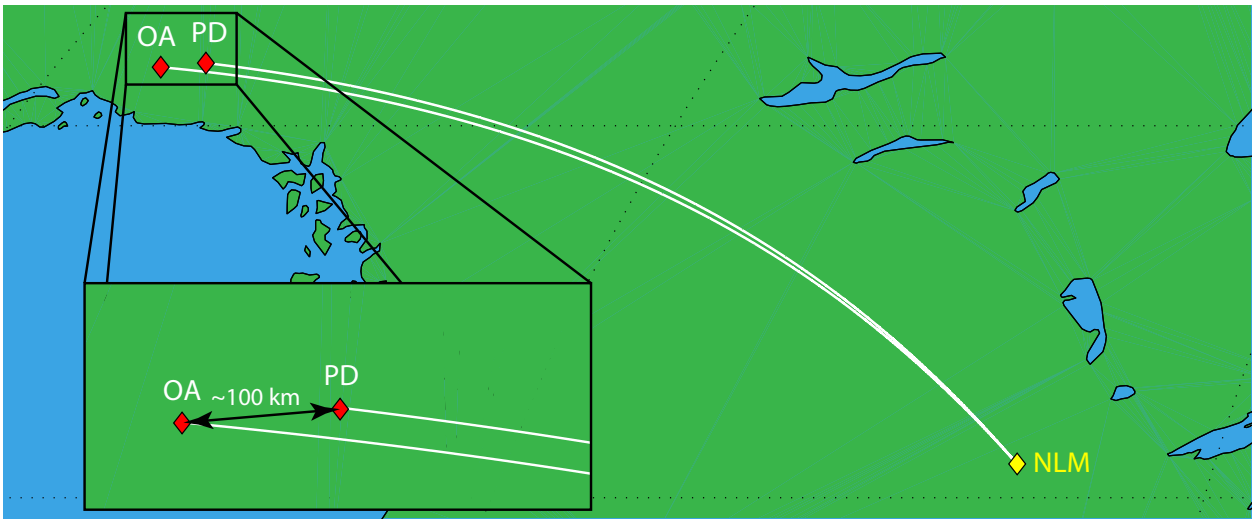

Figure 3-14. The propagation paths from the NLM transmitter to receivers in Alaska, Oasis and Paradise, which are separated by a short distance, about 100 km.

amplitude and phase between receivers and can be compared to the actual transfer function derived from broadband data.

A site-to-site transfer function can be taken between any two receivers. Receivers in Alaska are relatively close together, separated by 100 km or less, and provide a good opportunity to apply a site-to-site transfer function. Figure 3-14 shows the propagation path of the NLM transmitter to the Oasis and Paradise receivers. The distance separating the two differs by about 100 km resulting in different propagation paths but received signals at each receiver should be almost the same but with distinct differences.

The site-to-site transfer function is generated in a similar fashion to the transfer function itself by taking the division of one transfer function by another. In this case, the site-to-site transfer function between Oasis and paradise was taken on August-25 2012 during ambient conditions. Figure 3-15 shows the transfer function for a fixed time as a function of frequency for the University of Florida, Shafer, and broadband methods. There is good agreement for all three cases for the amplitude as a function of frequency. However, only the UF method matches the both the amplitude and phase of the broadband transfer function, i.e., the actual transfer function. The Shafer method phase is nearly flat and does not correctly account for the phase ramp that occurs between sites, making it an imperfect measure of the site-to-site transfer function as a function of

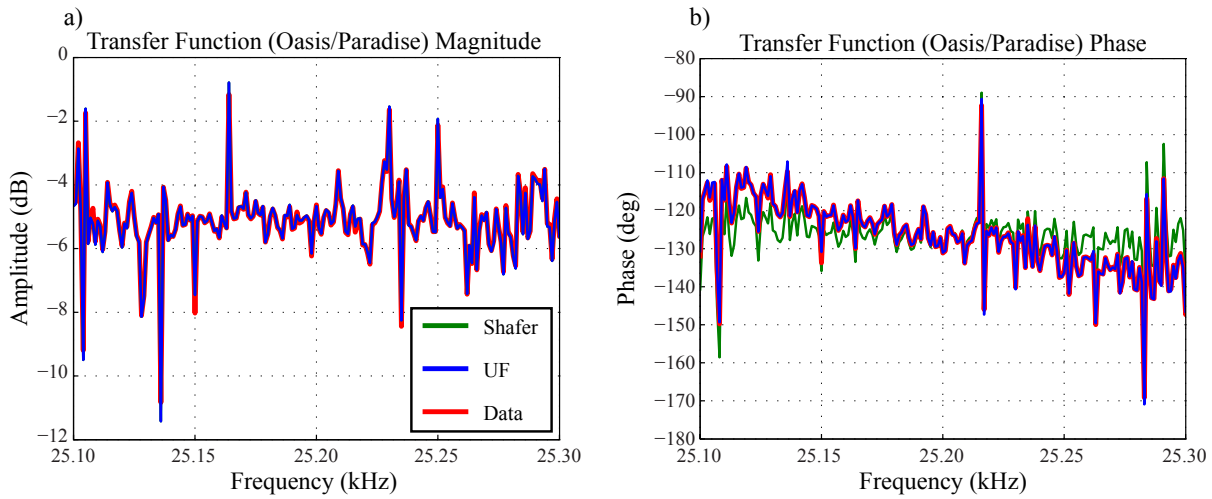

Figure 3-15. The site to site transfer function for the Shafer and UF methods as a function of frequency.

frequency. Considering the relatively large (10°) difference in phase measured with only 100 km distance between sites, it is likely that the Shafer method performs worse than expected, especially for larger distances.

Figure 3-16 compares different methods at the center frequency of transmission as a function of time. Both the UF method and Shafer method match well with the broadband data while the Stanford method is off by a fixed value, about 0.5 dB, for amplitude. For phase, all three methods

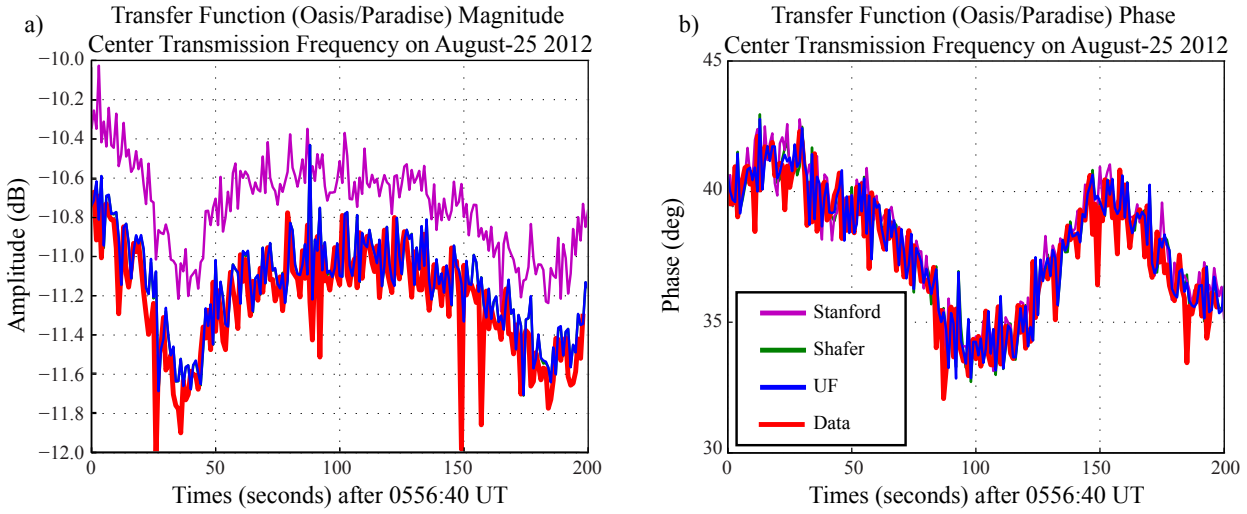

Figure 3-16. The site to site transfer function for several methods at the center frequency, 25.2 kHz, as a function of time for both amplitude, (a), and phase, (b).

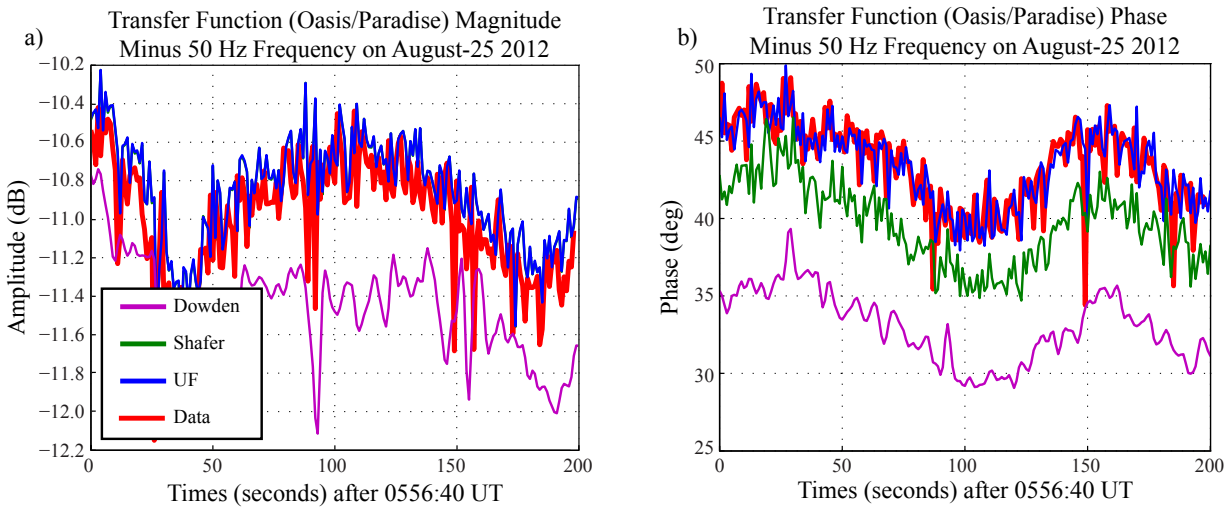

Figure 3-17. The site to site transfer function for several methods at the minus 50 Hz frequency, 25.15 kHz, as a function of time for both (a) amplitude and (b) phase.

match well with the broadband data transfer function. Based on the previous figure, it is expected that Shafer's method would match well at the center frequency. Stanford's method also shows good agreement, however: the amplitude is not exact, but it is close, and the phase is almost exact as well. Because Stanford's method is likely to deviate from the "correct" answer under multi-mode propagation conditions, this analysis will be re-visited later.

The two Dowden/Adams frequencies are $\pm$50 Hz relative to the center transmission frequency or 25.25 and 25.15 kHz. Figure 3-17 shows the UF, Shafer, and Dowden methods at the +50 Hz frequency. Both the UF and Shafer methods are accurate measurements of amplitude while only the UF method is an accurate measurement for the phase. The Dowden method amplitude is considerably different from the broadband data and is not off by a fixed value in the way the Stanford method was at the center frequency. The Dowden phase is similar to the broadband data and seems to be off by a fixed constant but when taking both the amplitude and phase together it illustrates that the two frequency measurement does not have a strong physical meaning. There are similar results for the -50 Hz as shown in Figure 3-18 where the amplitude of the Dowden/Adams method is in a similar range as the broadband data and the phase is almost off by a fixed constant. Again, both the UF and Shafer methods show good agreement with the broadband data for amplitude while the UF method is the only one that also has good phase agreement.

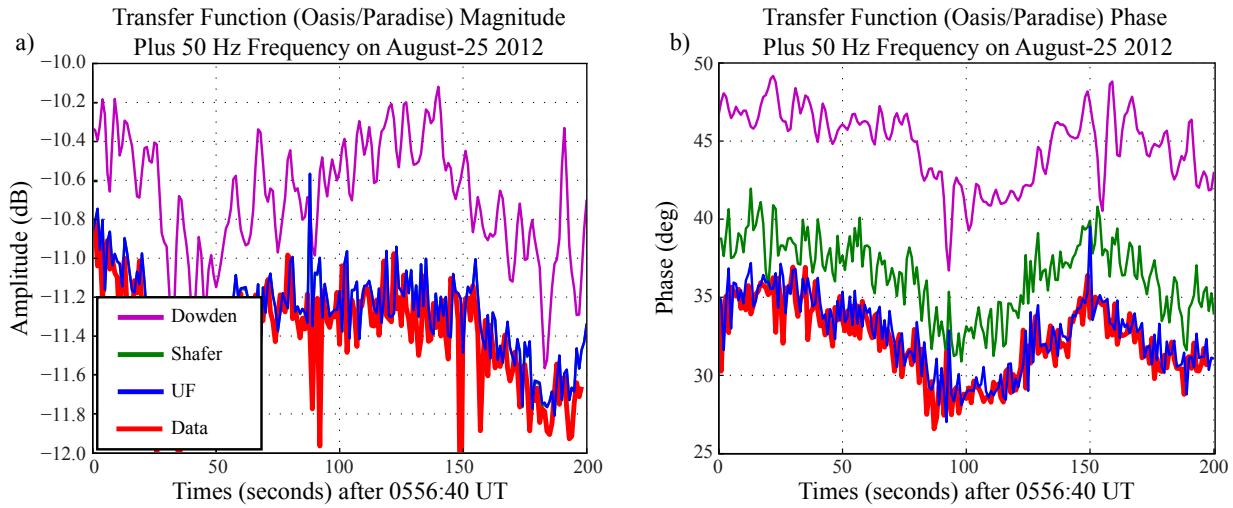

Figure 3-18. The site to site transfer function for several methods at the plus 50 Hz frequency, 25.25 kHz, as a function of time for both (a) amplitude and (b) phase.

In addition to the site-to-site transfer function, the different methods can be compared over long periods of time for ambient data. Figure 3-19 shows the difference between the UF and Stanford method's amplitude over the course of a 23 hours for three different transmitters, NAA, NLM, and NLK. Each path has a different diurnal variation due to the times when the path is sun lit, indicated by the sunrise or sunset at the transmitter and receiver, resulting in a quiet day or a turbulent night time. During the day, there is good agreement between the methods with a fixed difference on the order of 0.1 or 0.2 dB. However, a solar flare occurs during the day time and results in an increase in the difference between methods with the most pronounced event difference for the NLK to San Antonio path. We will return to this transient event in Chapter 4. During the night time, the amplitude shows a much larger variation compared to the daytime. This is due to the additional multi-mode content that propagates at night which the Stanford method does not accurately assess; the day-night terminator exhibits the greatest amount of difference (several dB) between the two methods due to multi-mode content.

The plot of phase difference also shows a a diurnal variation for the difference between methods for ambient data as shown Figure 3-20. All three transmitters show good agreement during daytime with almost zero phase difference, but there is a noticeable change during the solar flare event yielding a phase difference of one or two degrees. As was the case with amplitude, there

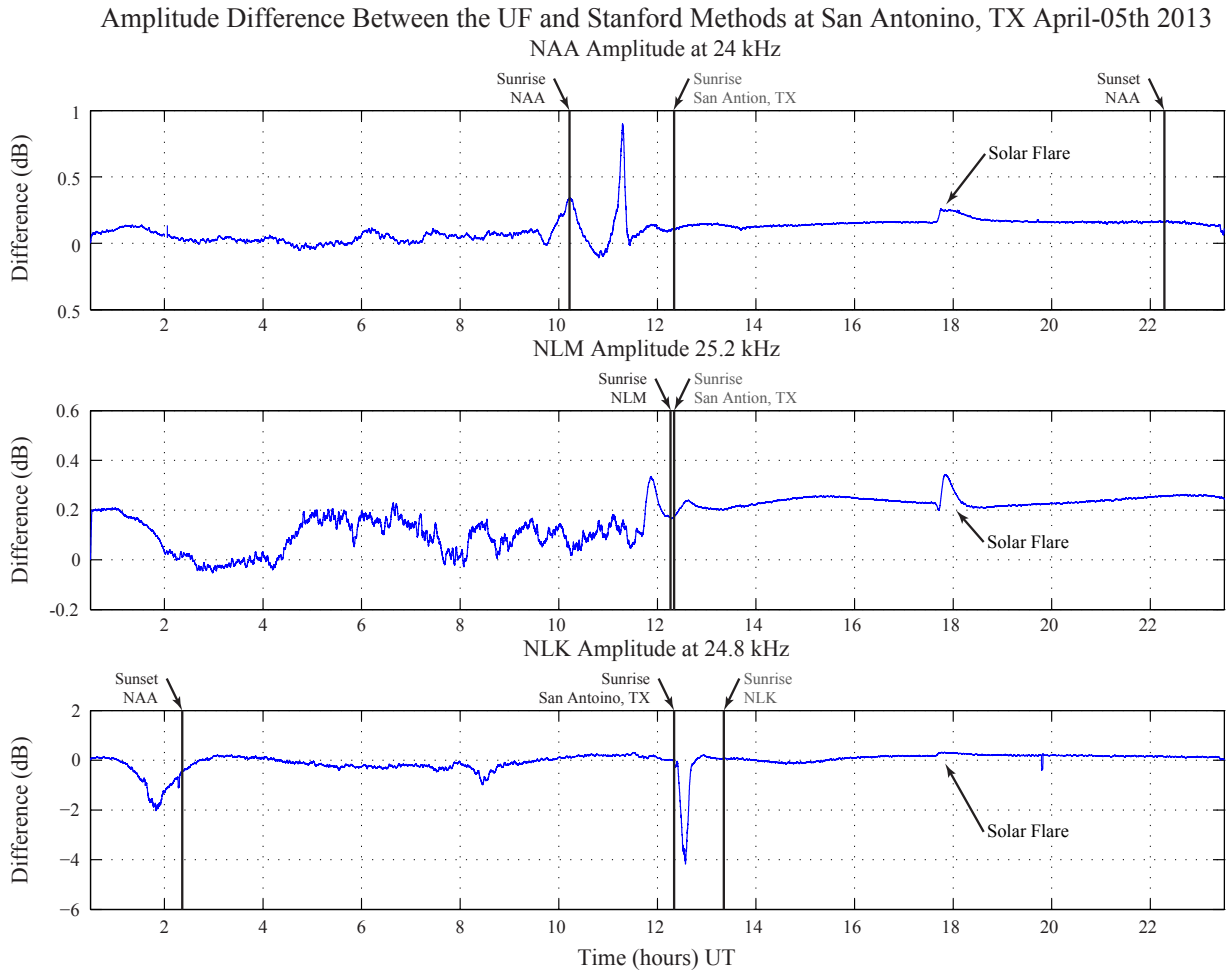

Figure 3-19. An example of the differences between the UF and Stanford method's amplitude over a 23 hour period at San Antonio, TX for multiple transmitters on April-5th 2013.

is a more pronounced difference during night time with the difference reaching as much as three degrees for the NLM path. It is important to note that the difference in phase slowly drifts apart over the course of hours and only returns to a difference close to zero after the rapid change, a difference of several degrees over a few short hours, associated with the day-night terminator. This plot is a perfect example of the limitations of the Stanford method. While the Stanford method serves as a good approximation, it cannot accurately predict the amplitude and phase for a given transmitter to receiver path in the presence of multi-mode content.

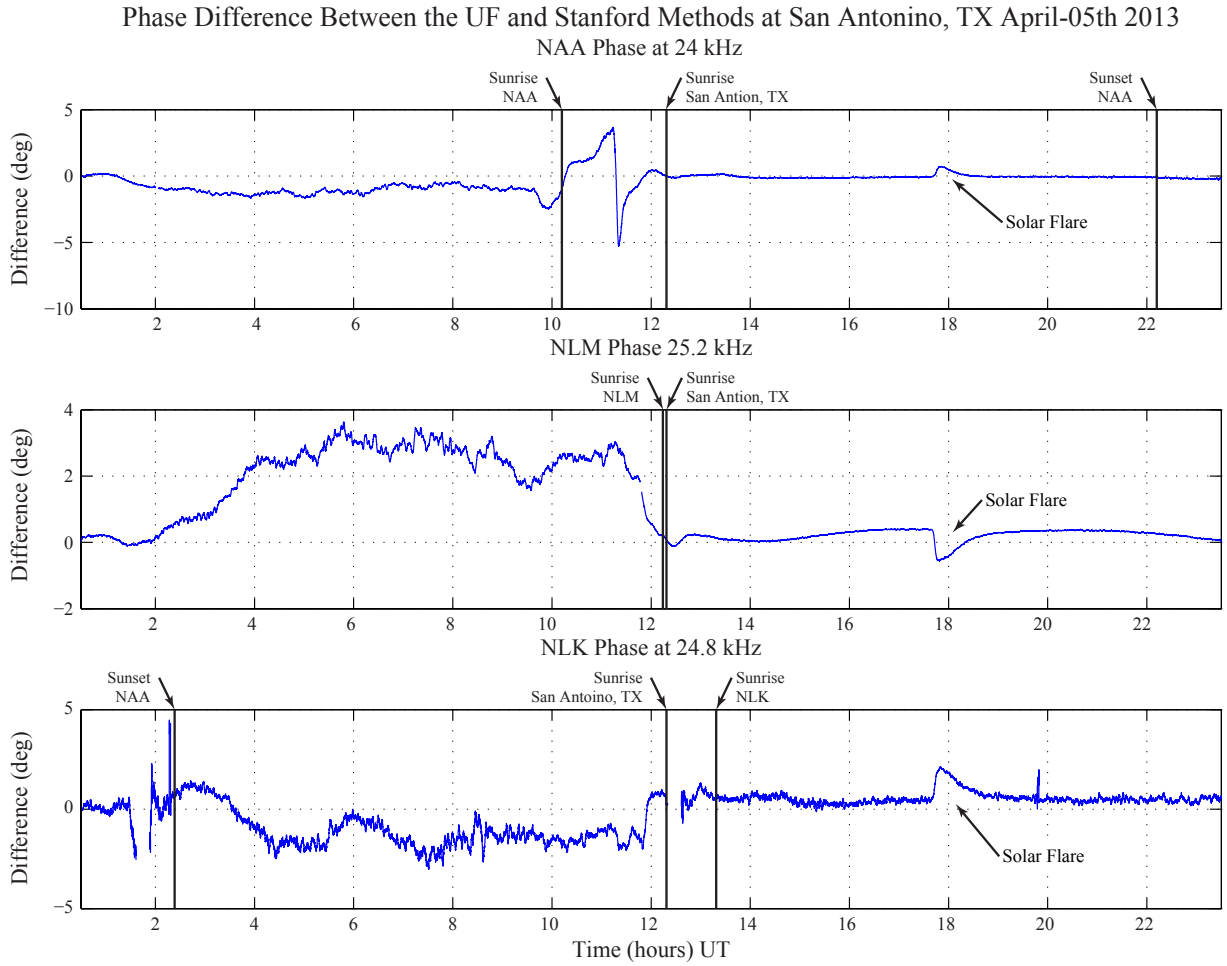

Figure 3-20. An example of the differences between the UF and Stanford method's amplitude over a 23 hour period at San Antonio, TX for multiple transmitters on April-5th 2013.

Nevertheless, the Stanford method is typically within 1 dB and 3 degrees of the UF result (excluding the terminator region), indicating that it is a good, although not perfect, approximation of the sum-of-modes amplitude and phase.

CHAPTER 4
TRANSIENT VLF EVENT ANALYSIS

In this chapter, the spread spectrum processing technique is applied to transient events for two purposes: 1) to determine whether transient events exhibit detectable structure as a function of frequency, and 2) to compare the response of the method to transient signal perturbations with those of other signal processing methods. On a heuristic level, any detected dependence on frequency could be attributed to the physics governing the development of the transient ionospheric disturbance, the geometry of the disturbance in relation to the transmitter/receiver propagation path, and/or the timing of the ionospheric disturbance in relation to other detected transient events. Three distinct event types are considered: lightning induced electron precipitation (LEP) events, early/fast events, and solar X-ray flares. Different aspects of these three events exhibit distinct time scales. Both LEP and early/fast events are associated with lightning and can be identified as sudden increases or decreases in the amplitude and/or phase of the received narrowband VLF signal followed by a $\sim$10–100 second recovery. LEP events exhibit onset durations (the time between the start of the event and the maximum amplitude/phase perturbation) on the order of a second, whereas early/fast events exhibit onset durations on the order of 10's of milliseconds. Due to their association with lighting, LEP events and early/fast events will sometimes occur during the recovery of another event. In contrast, solar X-ray flares events typically occur over much longer time scales and are usually observed as a singular event. These events typically persist for tens of minutes to hours. Scattered field analysis is used together with the spread spectrum method to analyze all of these events as a function of frequency and time.

Before discussing scattered field analysis and proceeding to evaluate the response of the method to transient events, it is important to briefly discuss frequency dependence for the ambient case. Figure 4-1 shows a spectrogram of data collected at Northside exhibiting an ambient, or undisturbed, period from $\sim$0–15 minutes followed by a number of events between $\sim$15 and 60 minutes. The ambient data clearly exhibits frequency dependence in the amplitude of the received signal, and the same can be shown for the ambient phase (not shown). Based on the data shown in

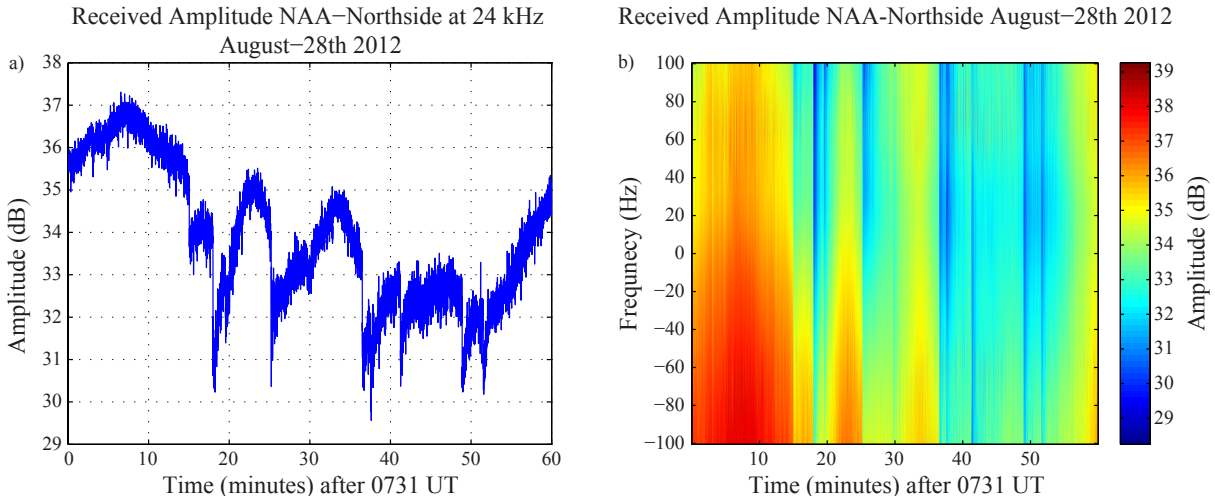

Figure 4-1. The received amplitude for the NAA to Northside path for a single frequency, (a), and as a function of frequency, (b), on August-28th, 2012. while transient events can potentially result in multi-frequency content, it is also possible for the ambient ionosphere to vary as a function of frequency.

this figure (and on countless other examples of ambient data), if the end-user is interested in analyzing the ambient ionosphere, it is clearly imperative to save measurements for the entire 200-Hz frequency band.

On the other hand, Figure 4-1 exhibits a clear frequency dependence in the amplitude variation observed during the events as well. In this case, the dependence of the ambient field on frequency might affect the dependence on frequency observed during the events. For this reason, scattered field analysis is employed to investigate whether transient events exhibit a frequency dependence that cannot be attributed to the ambient field. A discussion of scattered field analysis is now covered in detail.

4.1 Single Transient Events

4.1.1 Scattered Field Analysis

Transient VLF perturbations are often analyzed using scatterd field analysis. Any ionospheric disturbance occurring along or near the great-circle path (GCP) from VLF transmitter to VLF receiver can be interpreted as producing a scattered field [*Dowden et al.*, 1991; *Poulsen et al.*, 1990]. This scattered field adds to the ambient field observed in the absence of the disturbance, as illustrated by the phasor diagram of Figure 4-2. The ambient field propagates along the GCP to

Figure 4-2. An event will create a region that scatters VLF signals. This figure shows the direct path from transmitter to receiver and an additional path from transmitter to the event and then from the event to the receiver. Each path has its own amplitude and phase.

the receiver, and is signified by amplitude A_0 and phase ϕ_0 in the diagram. The scattered field is the VLF signal scattered by the ionospheric disturbance in the direction of the receiver, signified by amplitude A_{SF} and phase ϕ_{SF} in the diagram.

The process for extracting the scattered field requires knowledge of the ambient field. It is common to take the ambient field as the average (over several seconds or minutes) field prior to the event, and assume this ambient field remains valid throughout the remainder of the event. Of course, the true ambient field can change over the duration of the event (particularly for long-duration events), but this possibility will be neglected for the purposes of the analysis included here. The normalized scattered field is calculated by dividing all fields shown in Figure 4-2 by the ambient field, $A_0e^{j\phi_0}$, yielding:

$$\text{Normalized Scattered Field } = \frac{A_{SF}e^{j\phi_{SF}}}{A_0e^{j\phi_0}} = \frac{A_Te^{j\phi_T}}{A_0e^{j\phi_0}} - 1 \tag{4–1}$$

Throughout this chapter, the analysis of transient events will be performed using the normalized scattered field phasor described above.

4.1.2 Lightning Induced Electron Precipitation

Lightning induced electron precipitation was identified by Michael Trimpi, a graduate student at Palmer Station, Antarctica in the 1970's, when he noticed the repeated detection of lighting sferics coincident with VLF signal perturbations [*Helliwell et al.*, 1973]. The electromagnetic wave emanated by lightning couples to the magnetosphere, interacts with energetic electrons in the Earth's radiation belts, and scatters them onto the ionosphere below. *Voss et al.* [1984] presented

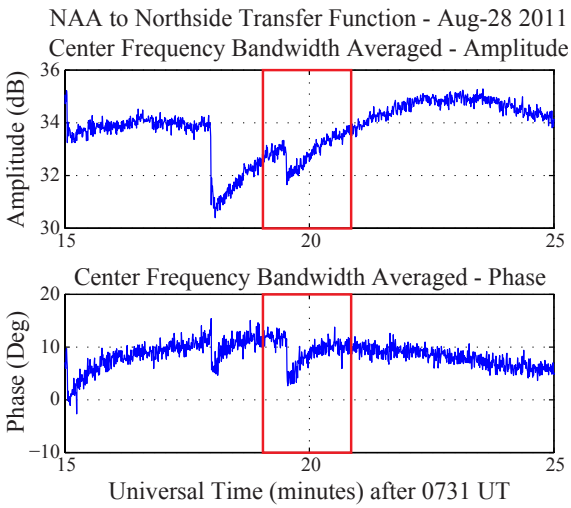

Figure 4-3. An example of an LEP event perturbation which causes a sudden drop in amplitude and phase and recovers over the course of tens of seconds.

the first example of LEP measurements made from space where a one-to-one correlation between ground based measurements of sferics and satellite measurements of electron precipitation was found. Electron precipitation modifies the upper boundary of the Earth-ionosphere waveguide, resulting in a perturbation of the amplitude and phase of the received signal followed by a recovery lasting several 10's of seconds [*Ferguson and Snyder*, 1987; *Inan et al.*, 1988a,b,c; *Poulsen et al.*, 1990; *Shellman*, 1986]. LEP research continues today in an effort to understand the mechanisms governing the dynamics of Earth's radiation belt population.

Figure 4-3 shows an example of an LEP event. The transfer function amplitude and phase are both plotted for the center frequency of NAA, 24.0 kHz, at Northside on 28 August 2011 for the North/South antenna. The data is averaged over the entire 200 Hz bandwidth to produce a single frequency measurement. This LEP event in particular, outlined in red, occurs in close proximity to another event, resulting in the change in amplitude occurring during the recovery of the previous event. This event has an amplitude change of $\sim$1 dB and a phase change of $\sim 10^{\circ}$. In general, LEP events can exhibit different changes in amplitude ranging from small events (< 0.5 dB) to large events (several dB), while phase changes can vary from a few degrees to tens of degrees. Often there will be a number events that occur within a small time frame, such as a few hours.

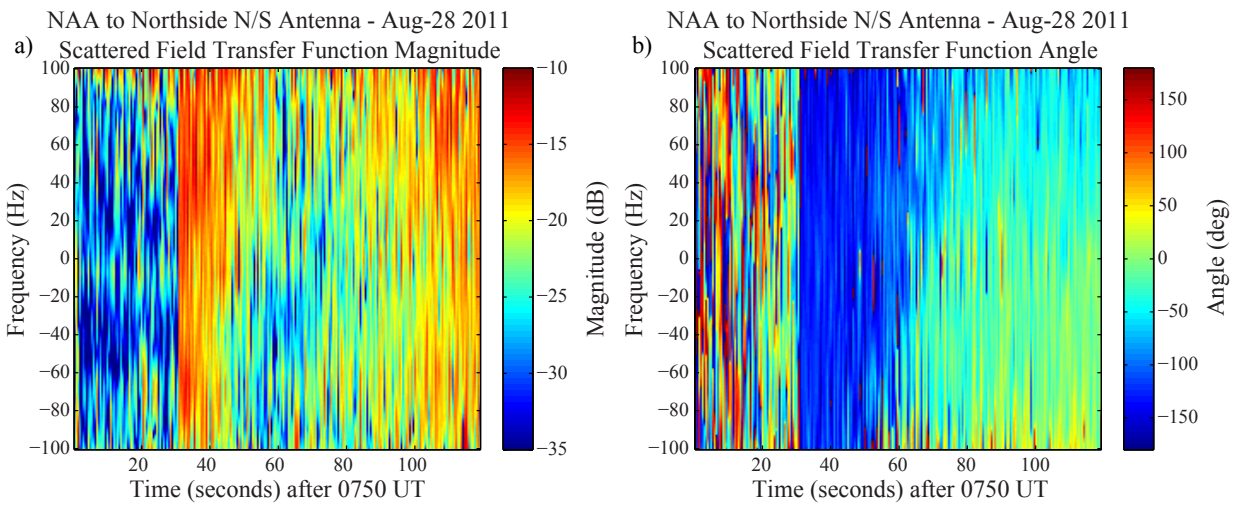

Figure 4-4. The normalized scattered field magnitude, (a), and angle, (b), transfer functions for an LEP event. The transfer function has an output sampling frequency of 2 Hz, frequency spacing of 1 Hz, and bandwidth averaging of 30 Hz.

The normalized scattered field transfer functions for this LEP event have been plotted in Figure 4-4 as a function of time and frequency. The transfer function averaged over a bandwidth of 30 Hz for display purposes.

The normalized scattered field magnitude does not exhibit a significant dependence on frequency over the course of the event. The angle of the normalized scattered field, on the other hand, exhibits a dramatic frequency dependence, taking $\sim$30 seconds to recover at the lowest frequencies and $\sim$40 seconds to recover at the highest frequencies. It is noteworthy, however, that it was extremely difficult to find an example of an LEP event exhibiting frequency dependence in the normalized scattered field. While Figure 4-4 demonstrates that such a frequency dependence can exist during transient events, it is not common. Indeed, it may be the case that this frequency dependence is only observed when the event in question occurs during the recovery stage of a previous event, significantly complicating the analysis. A statistical analysis examining this effect will be presented later in this chapter.

The performance of the method is now analyzed in comparison with other signal processing methods. Figure 4-5 compares the amplitude and phase response for three different methods (it is expected that the Shafer method will produce the same results, except offset by a phase constant

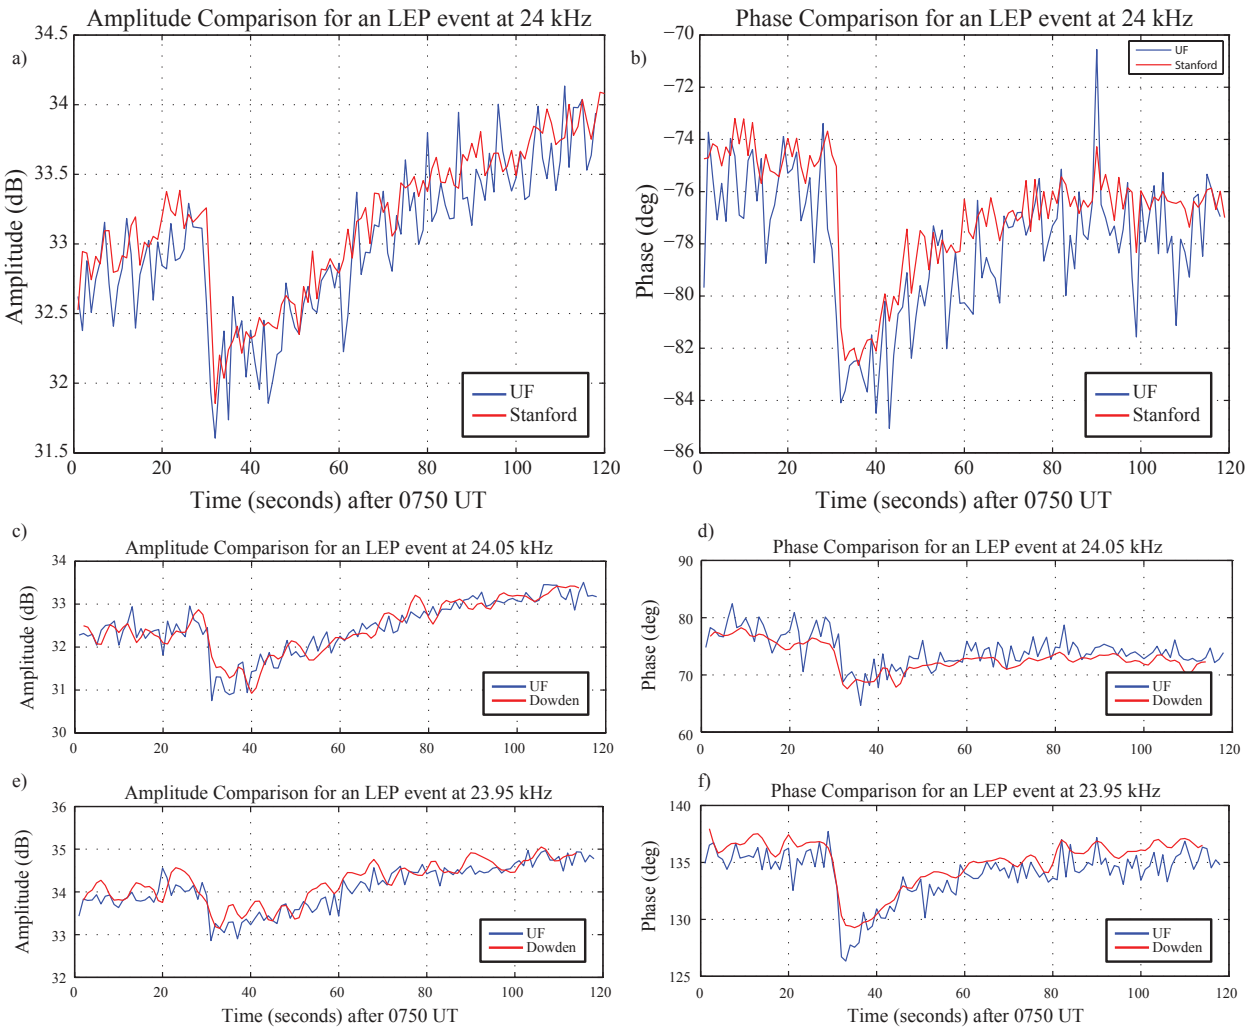

Figure 4-5. A comparison of the amplitude and phase for different methods when analyzing an LEP event. The Stanford amplitude, (a), and phase, (b), are compared to the UF method at the center frequency for NAA at Northside. While the Dowden method is compared at the plus and minus 50 Hz frequencies for amplitude and phase, (c) - (f). Note that in order to make an accurate comparison, the amplitudes and phases have been shifted so that the line up, neglecting any offset associated with other methods.

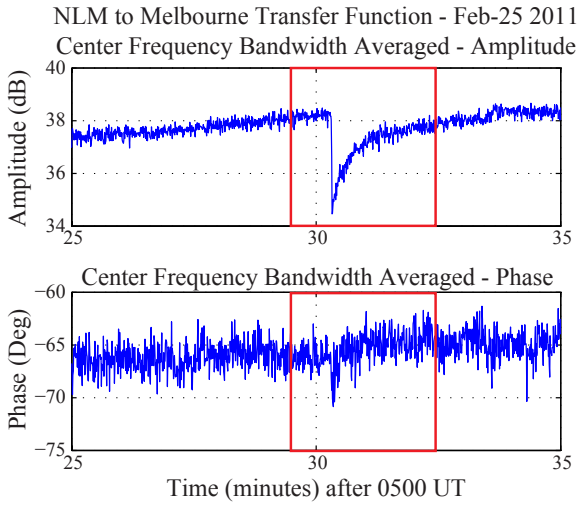

Figure 4-6. An example of an Early/Fast event with a large change in amplitude, 3.5 dB, and a negligible change in phase.

that depends on frequency). During this specific event, the Stanford method and the UF method show good agreement for both amplitude and phase. This is not always the case, as will be shown later. Moderately good agreement is exhibited by the comparison between the Dowden/Adams and the UF methods (at $\pm$50 Hz). The detailed differences are significant, however. For example, the Dowden/Adams method exhibits amplitude changes that are different by $\sim$1 dB and phase changes that are different by $\sim 2^{\circ}$. Depending on the application, these differences can be very important.

4.1.3 Early/Fast Events

Early/fast events are also produced in association with lightning. They were first identified by *Armstrong* [1983], who noted that the time delay between the causative lightning return stroke and the onset of the event was much too short ($<$ 100 msec) for the event to fit the physical mechanism for LEP. Early/fast events have been related to so-called transient luminous events (TLEs), although the specifics of the mechanism(s) involved have yet to be satisfactorily identified [e.g., *Moore et al.*, 2003].

An example early/fast event is shown in Figure 4-6 and boxed in red. The 28 August 2011 amplitude and phase transfer functions from the VLF transmitter NLM to the receiver at Melbourne have been plotted for the East/West antenna. Both amplitude and phase have an output sampling

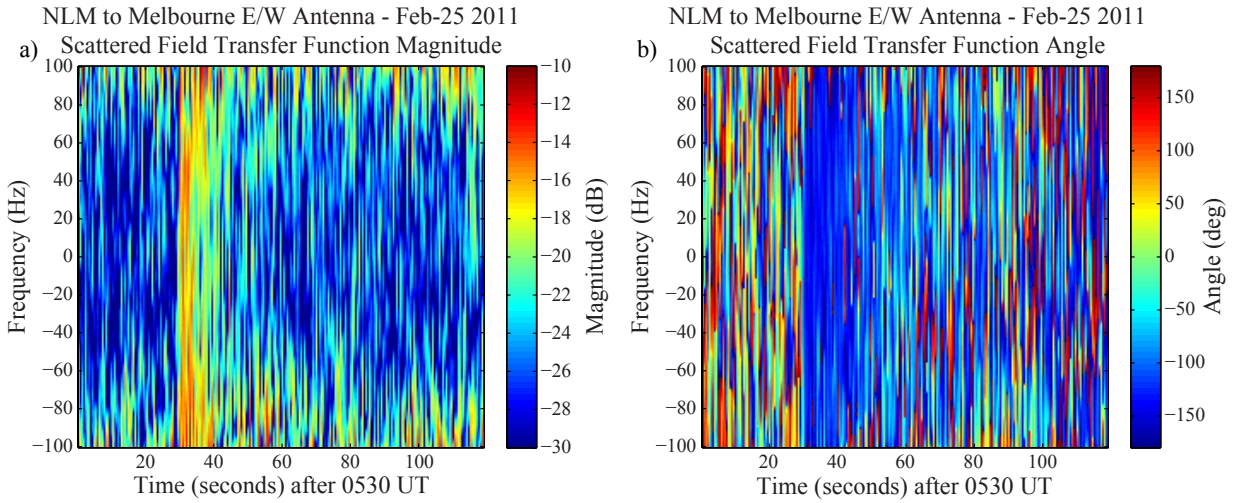

Figure 4-7. The normalized scattered field magnitude, (a), and angle, (b), transfer functions for an Early/Fast event. The transfer function has an output sampling frequency of 2 Hz, frequency spacing of 1 Hz, and a bandwidth average of 30 Hz but is considerably more noisy when when compared to similar plots for LEP.

frequency of 2 Hz and have been averaged over the entire 200 Hz bandwidth, yielding a measurement at the center frequency of NLM, 25.2 kHz. This event has a large amplitude change, 3.5 dB, but a nearly negligible phase change. The recovery lasts about one minute.

Figure 4-7 shows the normalized scattered field magnitude and angle plotted as functions of time and frequency. The same transfer function output settings have been used (2 Hz sampling frequency and 30 Hz bandwidth averaging) as for the LEP case in the previous section. Neither the normalized scattered field magnitude nor the normalized scattered field angle for this early/fast event exhibit any significant frequency dependence during the event. This is found to be the typical case, not only for early/fast events, but also for LEP events.

Figure 4-8 compares the results for different methods associated with this early/fast event. Similar to the LEP case, the Stanford method shows good agreement with the UF method in both amplitude and phase. The amplitudes of the Dowden/Adams method do not compare so well, but the phases are accurate. It is therefore concluded that all three methods produce a reasonable approximation of the event in this case.

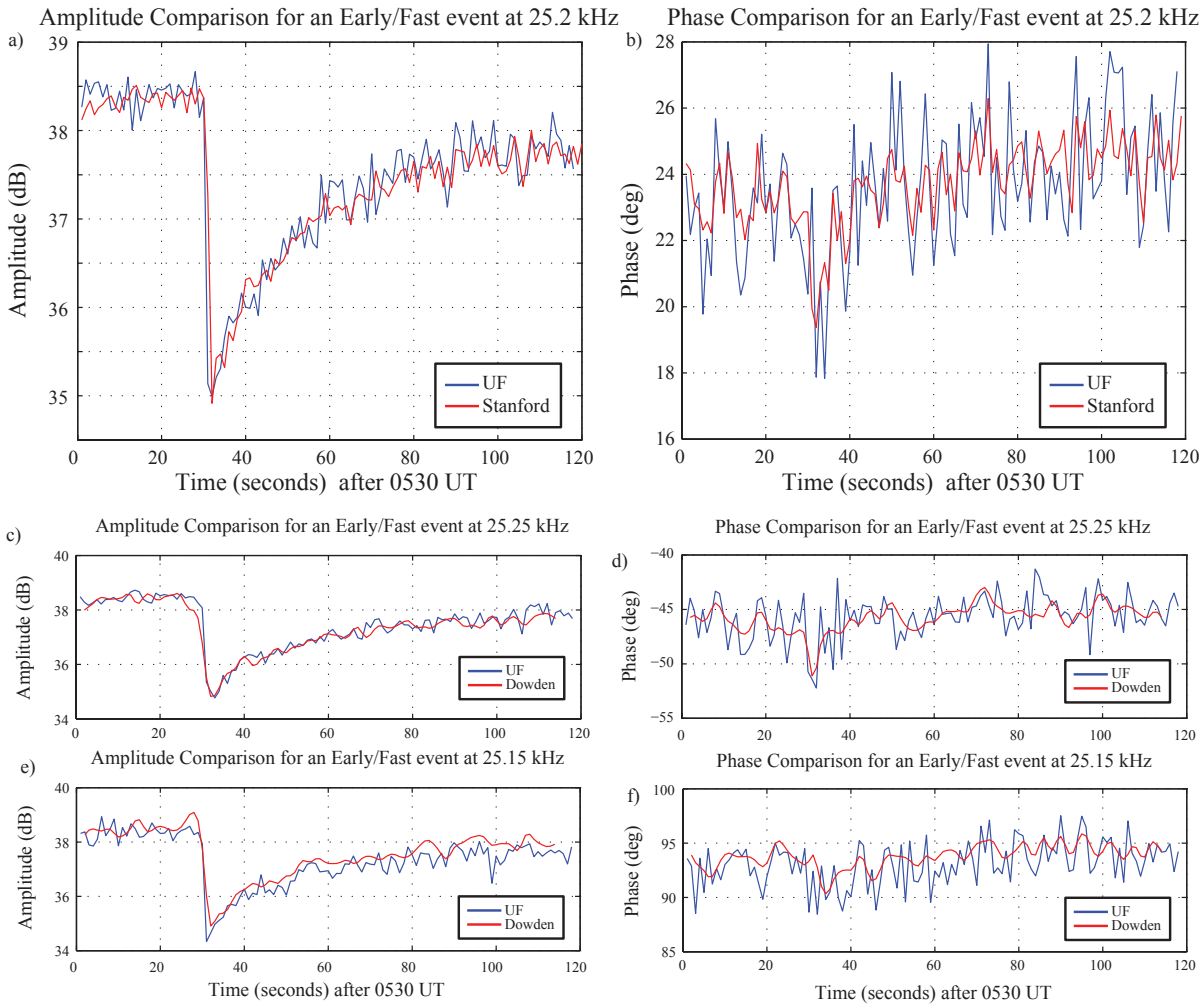

Figure 4-8. A comparison of the amplitude and phase for different methods when analyzing an Early/fast event. The Stanford amplitude, (a), and phase, (b), are compared to the UF method at the center frequency for NLM at Melbourne. While the Dowden method is compared at the plus and minus 50 Hz frequencies for amplitude and phase, (c) - (f). Note that in order to make an accurate comparison, the amplitudes and phases have been shifted so that the line up, neglecting any offset associated with other methods.

4.1.4 X-Ray Solar Flares

One of the major driving forces of the ionosphere is the sun, and X-ray emissions associated with solar flares have an especially pronounced effect. A typical emission will have a wavelength in the 1-8 Angstroms range and will ionize the *D*-region ionosphere resulting in additional radio wave absorption, but also lowering the VLF reflection height [*Davies*, 1990]. Unlike the localized perturbations of an LEP or early/fast event, solar flares bathe half the Earth in X-rays, resulting

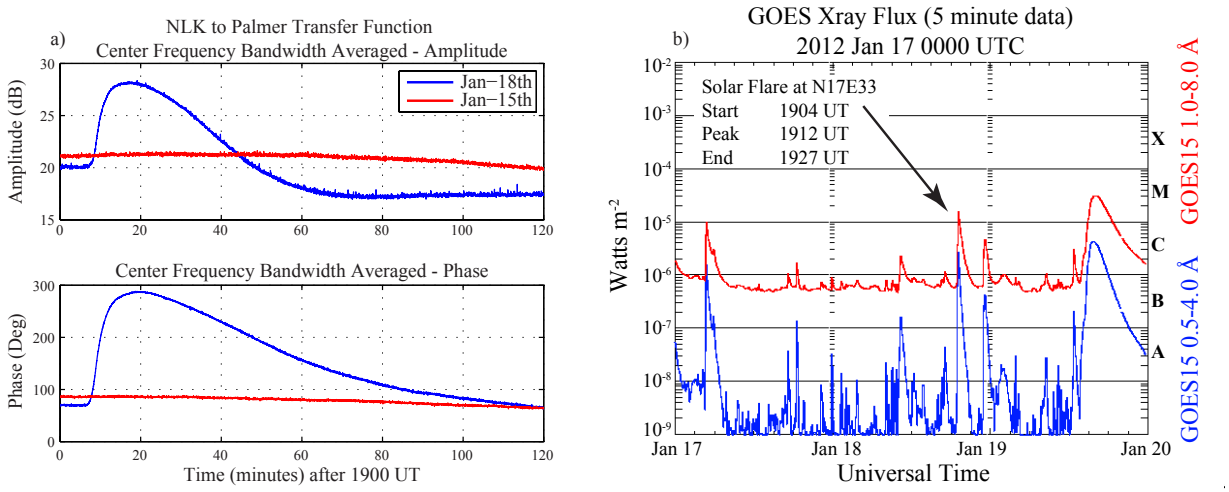

Figure 4-9. An example of an X-ray flare event perturbation in amplitude and phase of a VLF transmitter, (a), and the corresponding GOES X-Ray flux, (b), over a three day period.

in large amplitude and phase perturbations that last from tens of minutes to hours. Much of the research associated with solar flares has focused on the High Frequency (HF) and Very High Frequency (VHF) bands. Initial work focused on sudden frequency deviations of high frequency radio signals [*Davies and Donnelly*, 1966], the relaxation time of the F-region ionosphere [*Baker and Davies*, 1966], and the associated increases in electron density [*Garriott et al.*, 1967]. Early work regarding the effects of solar flares on the VLF band was performed by *Rosenberg et al.* [1971] who used ground-based VLF receivers in conjunction with balloons carrying VLF receivers and scintillation counters to show that X-ray flares were generating VLF perturbations. Later work would refine the associated models and particle interaction with the help of satellites in order to make direct X-ray measurements. The Geostationary Operational Environmental Satellite (GOES) program has maintained weather satellites over the past few decades: GOES-D/I/N/R [*Goodman et al.*, 2012], with the more recent iterations including X-ray, and electron flux measurements.

An example of an X-ray flare event can be seen in Figure 4-9. The amplitude and phase transfer functions are plotted in panel (a) for NAA received at Palmer Station, Antarctica on two different days: 15 and 18 January 2011. The solar X-ray flare occurred on the 18th. Both transfer functions have an output sampling frequency of 1 Hz and are averaged over the entire bandwidth to give a single measurement centered on 24.8 kHz. There is a clear change in amplitude and phase

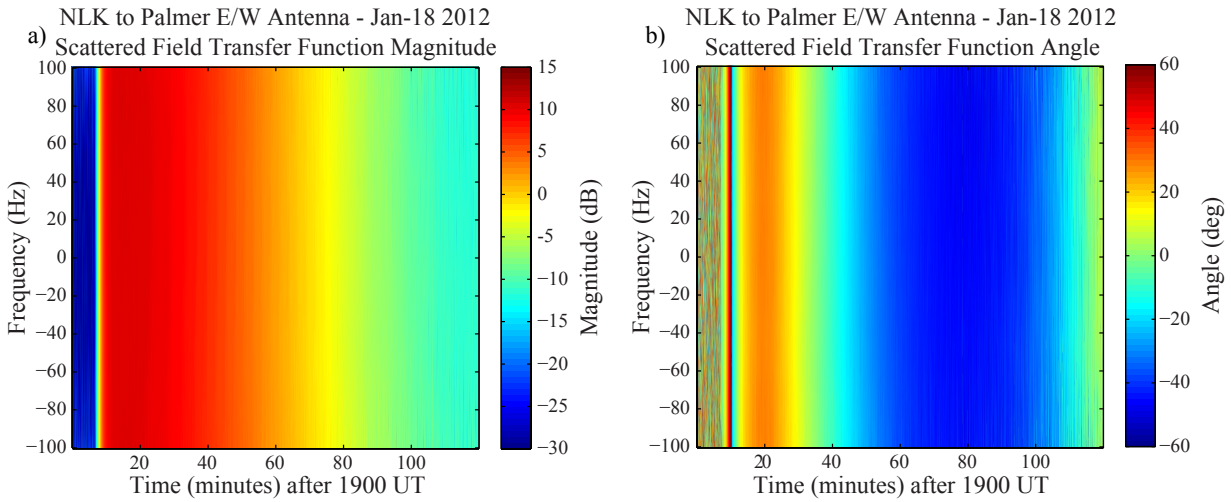

Figure 4-10. The normalized scattered field magnitude, (a), and angle, (b), transfer functions for a solar X-ray flare. The transfer functions use a output sampling frequency of 1 Hz, a frequency spacing of 1 Hz, and a bandwidth average of 30 Hz.

associated with the solar flare with the amplitude increasing by 15 dB and the phase increasing by more than 200°. The event quickly reaches its peak over the course of ten minutes and takes tens of minutes to recover to its ambient level in both amplitude and phase. The corresponding GOES X-ray flux is shown in panel (b), demonstrating that the X-ray flare starts at 1904 UT and reaches its peak at 1912 UT. The event ends at 1927 UT, but the flux takes tens of minutes to decay back to its ambient levels.

Figure 4-10 shows the normalized scattered field magnitude and angle for the solar flare event on January 17th, 2011, plotted as functions of frequency and time. In both cases, there is essentially no variation in frequency across the bandwidth. Once again, it appears that under normal circumstances, the scattered field magnitude and angle do not exhibit significant dependence on frequency within the 200-Hz bandwidth of the narrowband VLF transmission.

Figures 4-11 and 4-12 compare the results for different methods associated with a different solar X-ray flare at San Antonio, TX. Unlike the LEP and early/fast event cases, the Stanford method produces amplitude and phase measurements that are distinctly different from the UF method as shown by the differences in amplitude (a couple dB) and phase (a couple degrees) during the event. As expected, the amplitudes and phases of the Dowden/Adams method (not shown) are

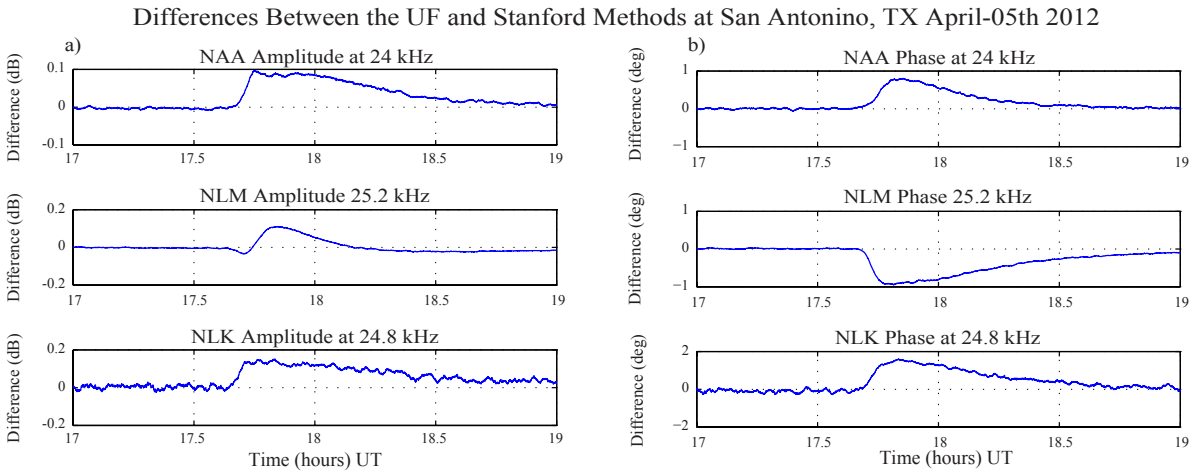

Figure 4-11. A comparison of the amplitude and phase differences for the UF and Stanford methods when analyzing a solar flare event. The Stanford amplitude, (a), and phase, (b), are compared to the UF method at the center frequency for NAA, NLM, and NLK at San Antonio, TX.

also different for both the plus and minus 50 Hz frequencies. It is therefore concluded that while the three methods typically produce similar amplitude and phase responses during transient events, there are cases where the differences are significant and quantifiable.

4.2 Multiple Transient Events

The spread spectrum processing technique has been applied to individual cases of LEP events, early/fast events, and solar X-ray flares events with some of the analysis exhibiting frequency content. This section presents a statistical analysis of events to determine whether frequency dependence is a common or rare occurrence.

For each type of event considered, an event selection criteria has been defined (to be delineated in each section below). Choosing which events to include in the statistical analysis possibly biases the result based on the selection criteria. Events with large amplitude and phase changes have been selected to achieve the highest event signal-to-noise ratio. Events have been selected that can be confirmed using an independent source. Both LEP and Early/Fast events are associated with lightning; events have only been included that can be associated with lightning detected by the National Lightning Detection Network (NLDN), which records intra-cloud lightning, cloud to ground lighting, and the return stroke peak current. In a similar fashion, the National Oceanic

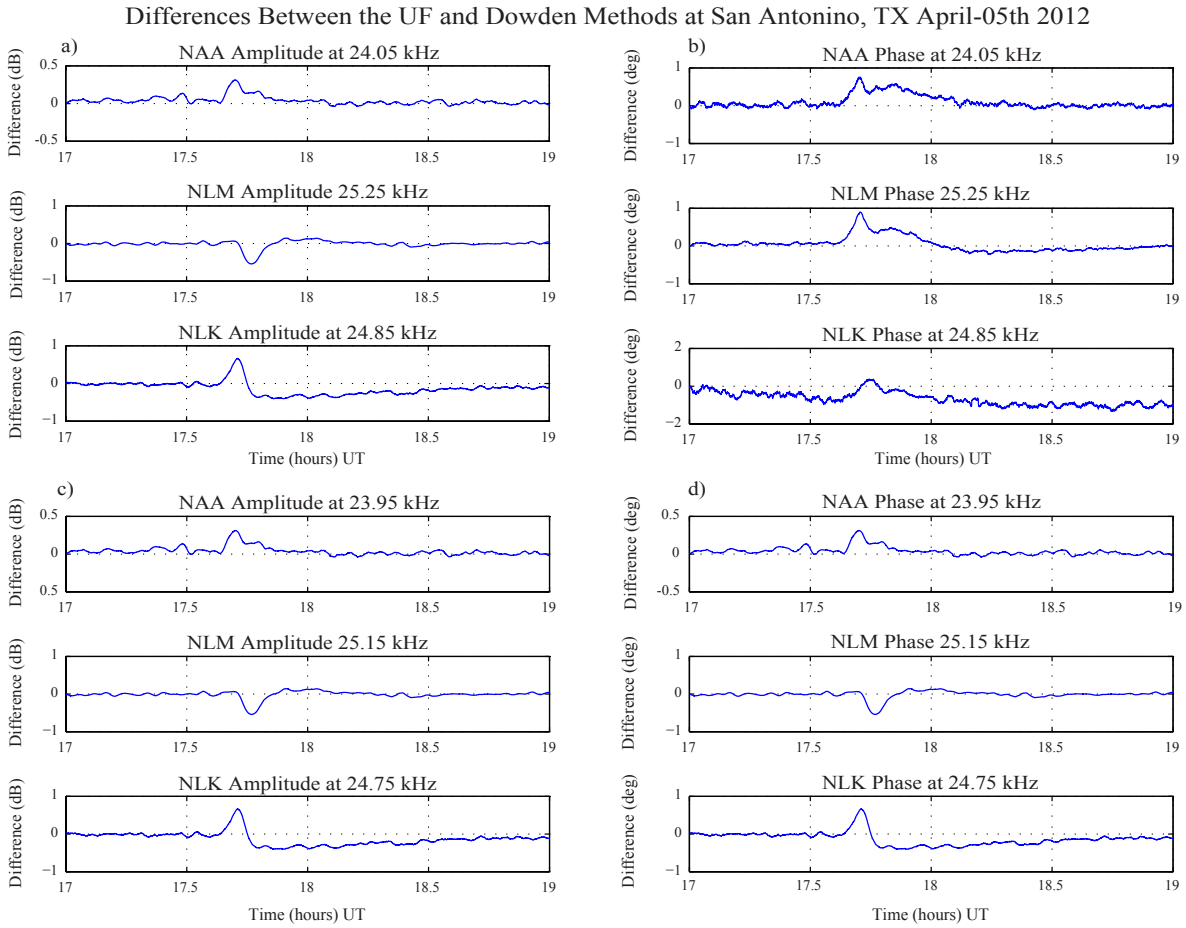

Figure 4-12. A comparison of the amplitude and phase differences for the UF and Dowden methods when analyzing a solar flare event. The Dowden amplitude, (a), and phase, (b), for the +50Hz frequency and the amplitude, (c), and phase, (d), for the -50Hz frequencies are compared to the UF for NAA, NLM, and NLK transmitters at San Antonio, TX.

and Atmospheric Administration (NOAA) maintains the Geostationary Operational Environmental Satellites (GOES) which host X-ray detectors whose data can be used to corroborate perturbations produced by solar X-ray flares.

4.2.1 LEP Events

The single LEP event was a good example of an event with amplitude and phase perturbations that resulted in frequency content in the scattered field. The criteria for perturbations included in this statistical analysis include: amplitude changes of more than 0.6 dB, phase changes of more than 4°, and a recovery time of less than three minutes. The primary mechanism for comparison

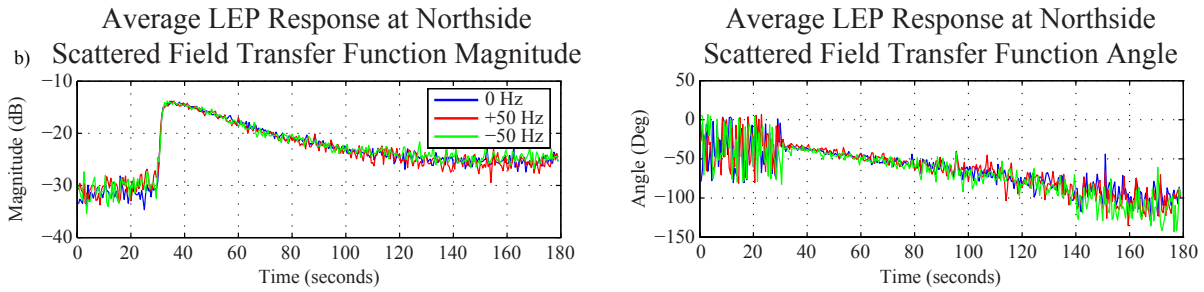

Figure 4-13. The normalized , (a), and the scattered field, (b), transfer functions for amplitude and phase plotted for three frequencies, -50, 0, and 50 Hz relative to the center transmitter frequency of NAA, 24 kHz. Both transfer functions represent the average for a number of LEP events.

will be the normalized scattered field. While smaller events exist with fractional changes in amplitude and phase, they can adversely affect processing by adding noise to the average of other events when the change in amplitude or phase is flat.

Because the focus of comparing events is on the normalized scattered field, it is important that events be chosen which return to their initial ambient levels. Calculating the scattered field requires normalization by the ambient signal level and events that recover to a different amplitude or phase level may falsely introduce frequency content. This criteria is especially problematic for LEP events that perturb the NAA to Northside path because there are many cases where an event occurs within the recovery period of another event. These events have been excluded.

Figure 4-13 shows the average normalized scattered field magnitude and angle. Three different frequencies, $+50$, 0, and -50 Hz, relative to the center frequency of NAA, 24 kHz, are plotted. Each transfer function has the output variables of a 2 Hz sampling frequency, 1 Hz frequency spacing, and a 30 Hz bandwidth averaging. There does not appear to be any frequency dependence exhibited in the average normalized scattered field.

The normalized scattered field is again plotted in Figure 4-14, but as functions of time and frequency. Neither the scattered field magnitude nor the scattered field phase exhibit any frequency dependence during the transient event.

In fact, checking each individual LEP event used in the statistical analysis, no frequency dependence exists in any of the normalized scattered fields. It thus appears that observed frequency

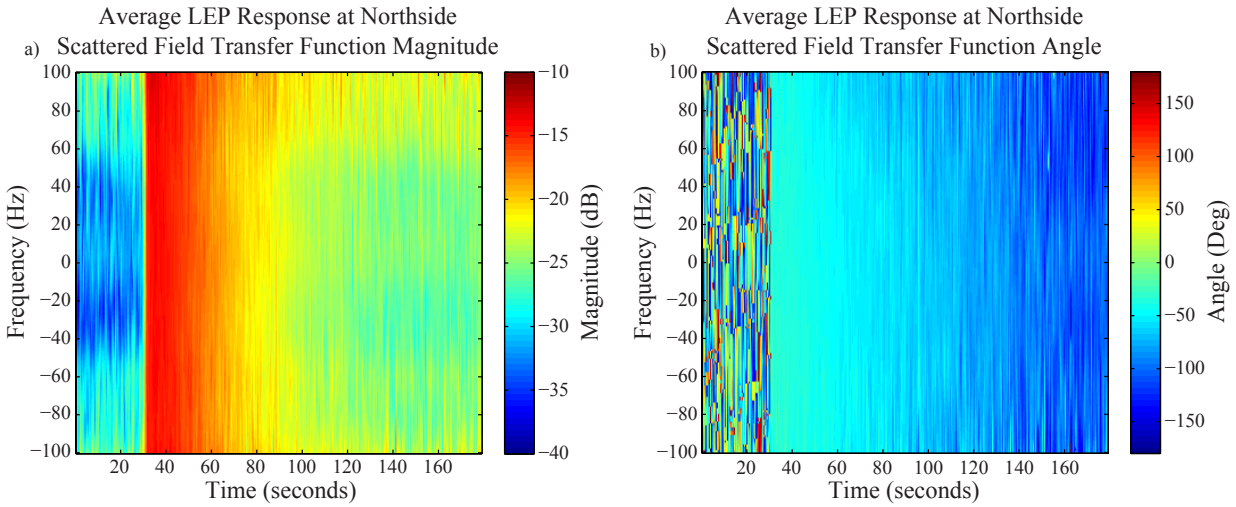

Figure 4-14. The normalized scattered field magnitude, (a), and angle, (b), transfer functions for the average LEP response for the NAA to Northside path.

dependence in the normalized scattered field during LEP events can most likely be attributed to a changing ambient field value that is not accounted for in the scattered field analysis. This effect would occur for LEP events occurring during the recovery of another event, for example, or when the ambient field is not particularly reliable. Although it is beyond the scope of this dissertation, it may be possible to extract any changes in the ambient field value during the transient event by enforcing zero frequency dependence of the normalized scattered field throughout the event.

4.2.2 Early/Fast Events

The criteria for identifying "usable" early/fast events is the same as that for LEP events: amplitude changes of more than 0.6 dB, phase changes of more than 4°, and a recovery time of less than three minutes. Each event is distinct, easily discernible, and does not occur during the recovery period of another event. Lastly, the event fully recovers to the initial ambient field level.

Figure 4-15 shows the normalized scattered field magnitude and angle plotted for three different frequencies as functions of time. Each transfer function uses the same output variables as for the LEP transfer functions in the previous section. In the average case, no frequency dependence exists. While a majority of events may have no frequency dependence, it is still possible to have individual cases where an early/fast event exhibits a variation as a function of frequency.

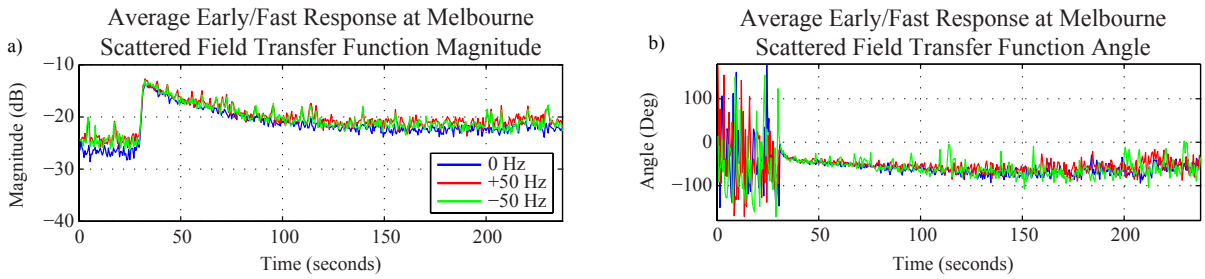

Figure 4-15. The normalized, (a), and the scattered field, (b), transfer functions for amplitude and phase plotted for three frequencies, -50, 0, and 50 Hz relative to the center transmitter frequency of NLM, 25.2 kHz. Both transfer functions represent the average for a number of LEP events.

The complete normalized scattered fields are plotted in Figure 4-16 as functions of both time and frequency, supporting the conclusion that there is no average frequency dependence.

Similar to the LEP case, each individual early/fast event used in this statistical analysis was checked for frequency dependence, and none was found in any of the normalized scattered fields. This result implies that, although analysis of the ambient fields would benefit from high frequency resolution, full bandwidth-averaged narrowband analysis is sufficient to characterize LEP and early/fast events.

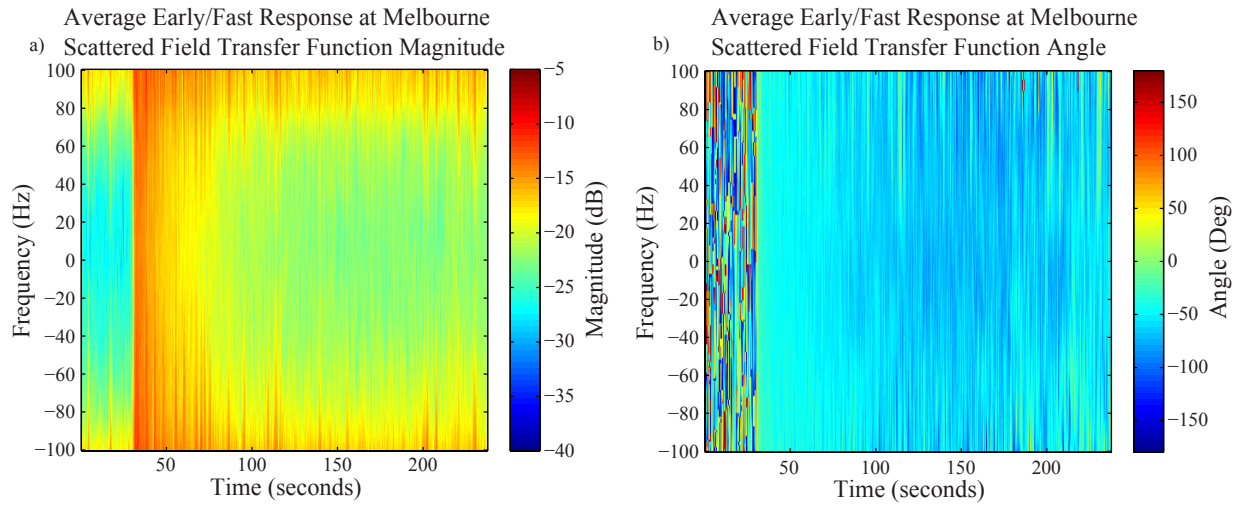

Figure 4-16. The normalized scattered field magnitude, (a), and angle, (b), transfer functions for the average Early/Fast response for the NLM to Melbourne path as functions of frequency and time.

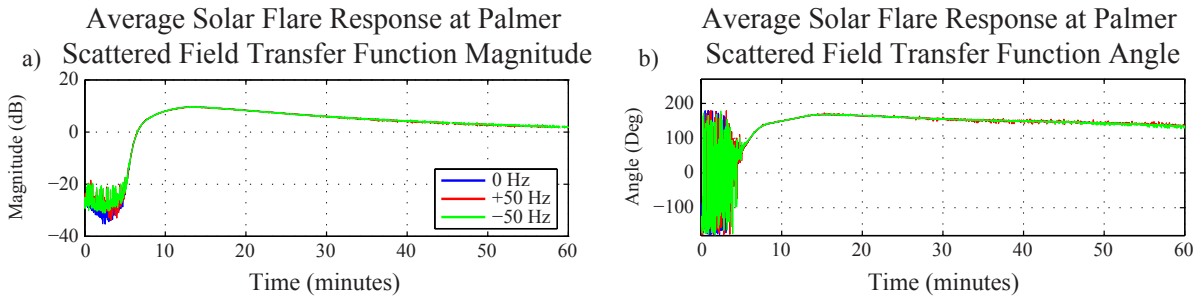

Figure 4-17. The normalized, (a), and scattered field, (b), transfer functions for the averaged solar flare event plotted for three frequencies relative to a center transmitter frequency. Unlike LEP or Early/Fast events, solar flares can occur on different channels and resulting in different paths such as NAA, NLK, NLM, and NPM to Palmer Station.

4.2.3 Solar X-Ray Flare Events

The example solar X-ray flare was an event with an 8 dB amplitude increase and a 205° increase in phase at its peak with a recovery time that lasted about two hours. Frequency content for the event was overall flat across the 200 Hz bandwidth. Choosing the solar flares to compare is a more delicate process than for LEP and early/fast events because the characteristics of solar X-ray flares vary so widely. In some cases, the onset of the event is much longer than other events, and the recovery period may last several hours or several minutes. The perturbation can change in amplitude 1 dB or 10's of dB, while the phase can change between 10's and 100's of degrees. For this statistical analysis, events have been limited to have the following properties: amplitude changes greater than 4 dB, recovery times less than two hours, and a corresponding GOES-detected solar flare that persists for at least 10 minutes.

Figure 4-17 shows the averaged normalized scattered field magnitude and angle. The transfer functions used has an output sampling frequency of 1 Hz, a frequency spacing of 1 Hz and, a bandwidth average of 30 Hz. Each of the plots can be broken up in to four distinct regions: the initial ambient, a rapid increase, a slower increase to the maximum level, and then a slow recovery. The rapid and slow increases are purely a function in the differences between solar flare events with each have its own distinct onset duration. Neither of the plots exhibit a frequency dependence.

The complete normalized scattered field magnitudes and phases are again plotted as functions of time and frequency in Figure 4-18. No strong dependence on frequency is detected. Each of the

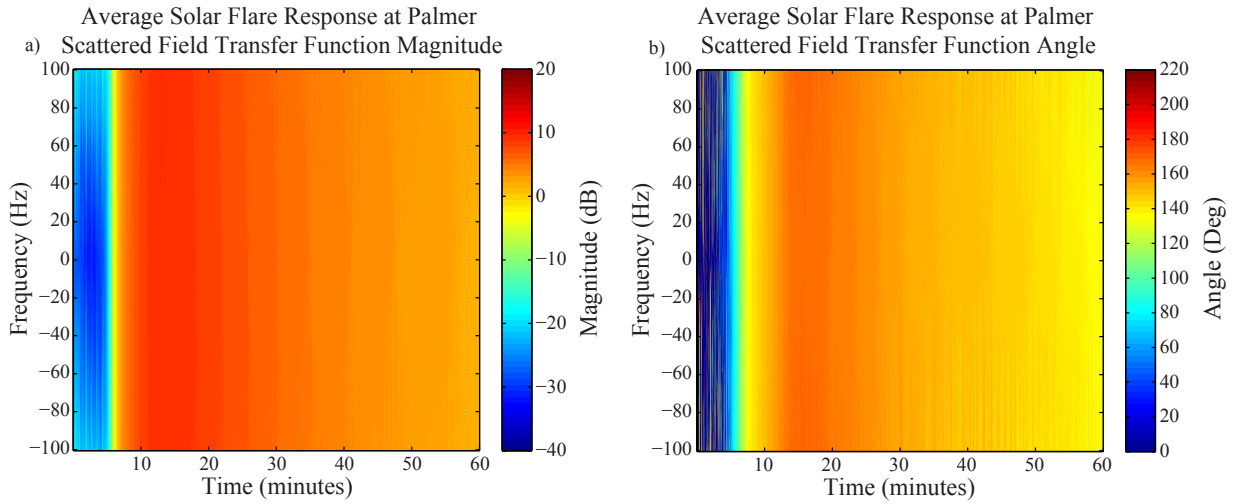

Figure 4-18. The normalized scattered field magnitude, (a), and angle, (b), plotted as functions of time and frequency for the average solar flare event.

events making up the statistical analysis were inspected individually, and no frequency dependence was observed in any of the cases.

Together the analysis of LEP, early/fast, and solar X-ray flare events strongly imply that no significant frequency dependence is observed during transient VLF scattering events. While it is possible that frequency dependence may be observed in the future, it is not present in this data set. Future researchers who may detect a frequency-dependent scattered field should consider the selection criteria outlined in this chapter, especially the possible time-variation of the ambient field over the course of the event, before concluding that the frequency dependence is real.

4.3 Choices in Sampling

Based on the results presented in this chapter, it appears to be the case that almost no frequency dependence is exhibited during the majority of LEP, early/fast, and solar X-ray flare events. This implies that, for transient event analysis, the output of the spread-spectrum method should average over the entire bandwidth in order to achieve the highest signal-to-noise ratio measurement. It is possible that there are other events which are not analyzed in this work that exhibit strong frequency dependence in the scattered field. Researchers interested in exploring this possibility should, of course, preserve measurements across the entire 200-Hz band.

It should also be noted that ambient field observations frequently exhibit strong variations as a function of frequency. Studies focusing on analyzing the ambient field values would strongly benefit from retaining the frequency dependent information.

CHAPTER 5
SUMMARY AND FUTURE WORK

5.1 Summary of Contributions

This dissertation has presented a new spread spectrum processing technique that analyzes the full 200-Hz bandwidth of MSK modulated VLF signals propagating in the Earth-ionosphere waveguide. The theory for the method has been presented, and the results of the implemented method have been evaluated. Previous difficulties with propagation delay, signal filtering, and signal windowing have been successfully addressed in the presented technique. As opposed to all other techniques currently available, the presented method allows for the direct comparison of amplitude and phase measurements observed at different receiver locations.

Experimental observations performed using the new signal processing technique have been presented to demonstrate the existence of frequency-dependent VLF signal propagation under ambient conditions (within the 200-Hz bandwidth of the signal). This observation is important because it has implications for ionospheric characterization, and it cannot be performed using other available techniques.

Analysis of the method performance during transient ionospheric events was also conducted. Observed frequency-dependent scattering was attributed to varying ambient conditions over the course of the event, rather than to the physical mechanism of the event itself or the observation geometry, and this perspective was supported by a statistical analysis of three types of transient VLF events. Additionally, all signal processing methods performed similarly during transient events, producing nearly the same changes in amplitude and phase during the events. The differences observed between the measurements were attributed to the effects of multi-mode signal propagation on the signal processing methods.

New VLF receiver hardware was detailed; laboratory measurements demonstrated the successful implementation of a high ($>$90 dB) linear dynamic range, high ($>$90 dB) spurious-free dynamic range (SFDR), and high ($>$90 dB) signal-to-noise and distortion ratio (SINAD) VLF receiver system. Field measurements were presented, demonstrating the successful deployment and

operation of the instrument. A new scheme for calibrating the amplitude *and* phase of the system in the field was also presented and successfully implemented.

Overall, the material presented in this dissertation has several important geophysical implications: 1) past observations of transient ionospheric events (using other analysis techniques) are likely to have been approximately accurate; 2) observations would benefit from the more accurate analysis technique presented here; 3) analysis of ambient signals at multiple receivers sites could not previously be performed, but now can be performed using the new processing technique; and 4) the impact of lightning-generated sferics on narrowband signal processing has not yet been fully evaluated and directly impacts the error bars on any narrowband signal measurement.

Based on the observations presented in this dissertation, the following VLF recording recommendations are made:

1. For occasional analysis of ambient conditions, and continual coverage of transient VLF events: save relatively short (1-minute duration) snippets of broadband data at the desired cadence for ambient analysis, and save 200-Hz bandwidth averaged narrowband amplitude and phase measurements at the sample frequency of choice. These settings reduce the total amount of memory required to perform the desired observations.

2. For continual coverage of ambient ionospheric conditions, and continual coverage of transient VLF events: retain information for the entire 200-Hz bandwidth of the VLF transmission at the desired sample frequency. Frequency averaging and temporal averaging can be performed in post-processing.

3. For continuing the search for frequency-dependent scattering during transient VLF events (which may occur for near-field scattering observations): save full resolution broadband data continually throughout the day. The memory cost of complete narrowband coverage is expensive, and at equivalent cost, one can retain the entire broadband data set.

5.2 Suggestions for Future Research

1. This dissertation has not presented measurements of the so-called "waveguide polarization" of the narrowband VLF signal. This is the polarization (eccentricity and major axis direction) of the VLF signal measured at the ground (using measurements from two orthogonal antennas). It may be possible that the waveguide polarization exhibits strong frequency dependence within the 200-Hz band of the narrowband transmission, and this possibility has not been evaluated. The eccentricity of the polarization ellipse would be approximately related to the QTE/QTM signal content, while the major axis direction would be related to the perceived "direction-of-arrival" of the VLF signal, which has been shown to vary with

time. It is possible that the waveguide polarization may exhibit strong frequency dependence during day/night transitions, for instance, and possibly during transient events.

2. This new narrowband VLF signal processing method can be implemented in a real-time data acquisition system. It is presently implemented in Matlab, which is slow compared to a real-time application, and the signal processing is completed in significantly-less-than-real-time, implying that implementation in a real-time package in C/C++ is definitely possible. Implementation in a real-time processing package would make the technique more widely accessible and of greater use to the scientific community at-large.

3. Differential ambient VLF signal measurements (made at multiple receiver sites) contain the best potential to quantify the altitude profile of ionospheric conductivity between receiver sites. This application has potentially significant impact for VLF-based global positioning systems, and should be investigated rigorously.

REFERENCES

Appleton, E. V., and F. W. Chapman (1937), On the nature of atmospherics. iv, *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, *158*(893), 1–22, doi:10.1098/rspa.1937.0001.

Armstrong, W. C. (1983), Recent advances from studies of the trimpi effect, *Antarctic Journal*, *18*, 281–283.

Baker, D. M., and K. Davies (1966), Solar flare effects and the relaxation time of the ionosphere, *Journal of Geophysical Research*, *71*(11), 2840–2842, doi:10.1029/JZ071i011p02840.

Bendaoud, S., and G. Marino (2004), Practical techniques to avoid instability due to capacitive loading, *Analog Dialogue*, *38*.

Bracewell, R., K. Budden, J. Ratcliffe, T. Straker, and K. Weekes (1951), The ionospheric propagation of low- and very-low-frequency radio waves over distances less than 1000 km, *Proceedings of the IEE - Part III: Radio and Communication Engineering*, *98*(53), 221–236, doi: 10.1049/pi-3.1951.0043.

Budden, K. (1957), The "waveguide mode" theory of the propagation of very-low-frequency radio waves, *Proceedings of the IRE*, *45*(6), 772–774, doi:10.1109/JRPROC.1957.278471.

Budden, K. G. (1955a), The numerical solution of differential equations governing reflexion of long radio waves from the ionosphere, *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, *227*(1171), pp. 516–537.

Budden, K. G. (1955b), The numerical solution of the differential equations governing the reflexion of long radio waves from the ionosphere. ii, *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, *248*(939), pp. 45–72.

Budden, K. G. (1961), *The Wave-Guide Mode Theory of Wave Propagation*, Logos Press, London.

Carter, B., and T. R. Brown (2001), *Handbook of Operational Amplifier Applications*, Texas Instruments.

Chapman, F. W., and R. C. V. Macario (1956), Propagation of audio-frequency radio waves to great distances, *Nature*, *177*, 930–933, doi:10.1038/177930a0.

Chapman, J., and E. Pierce (1957), Relations between the character of atmospherics and their place of origin, *Proceedings of the IRE*, *45*(6), 804–806, doi:10.1109/JRPROC.1957.278478.

Cohen, M. B., U. S. Inan, and E. W. Paschal (2010), Sensitive broadband elf/vlf radio reception with the awesome instrument, *IEEE Transactions on Geoscience and Remote Sensing*, *48*, 3–17, doi:10.1109/TGRS.2009.2028334.

Crombie, D. D. (1959), On the mode theory of very-low-frequnecy propagation in the presence of a transverse magnetic field, *Journal of Research of the National Bureau of Standards*, *64D*, 265–267.

Crombie, D. D. (1961), Reflection from a sharply bounded ionosphere for vlf propagation perpendicular to the magnetic meridian, *Journal of Research of the National Bureau of Standards*, *65D*, 455–463.

Davies, K. (1990), *Ionospheric Radio*, Peter Peregrinus Ltd., London.

Davies, K., and R. F. Donnelly (1966), An ionospheric phenomenon associated with explosive solar flares, *Journal of Geophysical Research*, *71*(11), 2843–2845, doi:10.1029/JZ071i011p02843.

Dobrott, D., and A. Ishimaru (1961), East-west effect on vlf mode transmission across the earth's magnetic field, *Journal of Research of the National Bureau of Standards*, *65D*, 47–52.

Dowden, R., S. Hardman, C. Rodger, and J. Brundell (1998), Logarithmic decay and doppler shift of plasma associated with sprites, *Journal of Atmospheric and Solar-Terrestrial Physics*, *60*(7âĂŞ9), 741 – 753.

Dowden, R. L., and C. D. D. Adams (1989), Modal effects on amplitude perturbations on subionospheric signals (trimpis) deduced from two-frequency measurements, *Geophysical Research*, *94*(A2), 1515–1519.

Dowden, R. L., C. D. D. Adams, M. T. Rietveld, P. Stubbe, and H. Kopka (1991), Phase and amplitude perturbations on subionospheric signals produced by a moving patch of artificially heated ionosphere, *Journal of Geophysical Research: Space Physics*, *96*(A1), 239–248, doi:10.1029/90JA01783.

Ferguson, J. A., and F. P. Snyder (1987), *The Segmented Waveguide Program for Long Wavelength Propagation Calculations*, Naval Ocean Systems Center, San Diego, CA.

Fraser-Smith, A. C., and R. A. Helliwell (1985), The stanford university elf/vlf radiometer project measurement of the global distribution of elf/vlf electromagnetic noise, in *International Symposium on Electromagnetic Compatibility*, edited by T. C. Rozzell, pp. 305–311.

Füllekrug, M. (2009), Wideband digital low-frequency radio receiver, *Measurement Science and Technology*, *21*(1).

Galejs, J. (1972), *Terrestrial Propagation of Lone Electromagnetic Waves*, Pergamon Press, New York.

Garriott, O. K., A. V. da Rosa, M. J. Davis, and O. G. Villard (1967), Solar flare effects in the ionosphere, *Journal of Geophysical Research*, *72*(23), 6099–6103, doi:10.1029/JZ072i023p06099.

Golomb, S. W. (1967), *Shift Register Sequences*, Holden-Day, San Francisco.

Goodman, S. J., J. Gurka, M. DeMaria, T. J. Schmit, A. Mostek, G. Jedlovec, C. Siewert, W. Feltz, J. Gerth, R. Brummer, S. Miller, B. Reed, and R. R. Reynolds (2012), The goes-r proving ground: Accelerating user readiness for the next-generation geostationary environmental satellite system, *American Meteorological Society*, *93*, 1029–1040.

Gronemeyer, S., and A. Mcbride (1976), Msk and offset qpsk modulation, *Communications, IEEE Transactions on*, *24*(8), 809–820, doi:10.1109/TCOM.1976.1093392.

Helliwell, R. A., A. J. Mallinckrodt, and F. W. Kruse (1951), Fine structure of the lower ionosphere, *Journal of Geophysical Research*, *56*(1), 53–62, doi:10.1029/JZ056i001p00053.

Helliwell, R. A., J. P. Katsufrakis, and M. L. Trimpi (1973), Whistler-induced amplitude perturbation in vlf propagation, *Journal of Geophysical Research*, *78*, 4679–4688, doi:10.1029/JA078i022p04679.

Holzer, R. E., and O. E. Deal (1956), Low audio-frequency electromagnetic signals of natural origin, *Nature*, *177*, 536–537, doi:10.1038/177536a0.

Inan, U. S., D. C. Shafer, W. Y. Yip, and R. E. Orville (1988a), Subionospheric vlf signatures of nighttime d region perturbationsin the vicinity of lightning discharges, *Journal of Geophysical Research*, *93*, 11,455–11,472, doi:10.1029/JA093iA10p11455.

Inan, U. S., T. G. Wolf, and D. L. Carpenter (1988b), Geographic distribution of lightning-induced electron precipitation observed as vlf/lf perturbation events, *Journal of Geophysical Research*, *93*, 9841–9853, doi:10.1029/JA093iA09p09841.

Inan, U. S., W. C. Burgess, T. G. Wolf, D. C. Shater, and R. E. Orville (1988c), Lightning-associated precipitation of mev electrons from the inner radiation belt, *Geophysical Research Letters*, *15*, 172–175, doi:10.1029/GL015i002p00172.

Johnson, H., and M. Graham (1993), *High-Speed Digital Design A Handbook of Black Magic*, Prentice Hall.

Jung, W. (2006), *Op Amp Applications Handbook*, Analog Devices.

Kasper, J., and C. E. Hutchinson (1978), The omega navigation system–an overview, *Communications Society Magazine, IEEE*, *16*(3), 23–35, doi:10.1109/MCOM.1978.1089729.

Kester, W. (2008), *Understanding SINAD, ENOB, SNR, THD, THD + N,, and SFDR so You Don't Get Lost in the Noise Floor*, Analog Devices, rev. a ed., mT-003.

King, G. (1997), Op amps driving capacitive loads, *Analog Dialogue*, *31*.

Lay, E. H., R. H. Holzworth, C. J. Rodger, J. N. Thomas, O. Pinto, and R. L. Dowden (2004), Wwll global lightning detection system: Regional validation study in brazil, *Geophysical Research Letters*, *31*(3).

Ljung, L. (1987), *System Identification Theory for the User*, Prentice-Hall, Englewood Cliffs, New Jersey.

Moore, R., C. P. Barrington-Leigh, U. S. Inan, and T. F. Bell (2003), Early/fast vlf events produced by electron density changes associated with sprite halos, *Journal of Geophysical Research (Space Physics)*, *108*, 1363, doi:10.1029/2002JA009816.

Oppenheim, A. V., and R. W. Schafer (1975), *Ddigital Signal Processing*, Prentice Hall.

Paschal, E. (2013), Variable length comb generator, private correspondance.

Paschal, E. W. (1988), The design of broad-band vlf receivers with air-core loop antennas.

Pasupathy, S. (1979), Minimum shift keying: A spectrally efficient modulation, *Communications Magazine, IEEE*, *17*(4), 14–22, doi:10.1109/MCOM.1979.1089999.

Pierce, J. (1989), Omega, *Aerospace and Electronic Systems Magazine, IEEE*, *4*(7), 4–13, doi:10.1109/62.34731.

Poulsen, W. L., T. F. Bell, and U. S. Inan (1990), Three-dimensional modeling of subionospheric vlf propagation in the presence of localized d region perturbations associated with lightning, *Journal of Geophysical Research*, *95*, 2355–2366, doi:10.1029/JA095iA03p02355.

Rosenberg, T. J., R. A. Helliwell, and J. P. Katsufrakis (1971), Electron precipitation associated with discrete very-low-frequency emissions, *Journal of Geophysical Research*, *76*, 8445–8452, doi:10.1029/JA076i034p08445.

Round, H., T. Eckersley, K. Tremellen, and F. Lunnon (1925), Report on measurements made on signal strength at great distances during 1922 and 1923 by an expedition sent to australia, *Electrical Engineers, Journal of the Institution of*, *63*(346), 933–, doi:10.1049/jiee-1.1925.0116.

Shafer, D. C. (1994), Spread-spectrum vlf remote sensing of the ionosphere, Ph.D. thesis, Stanford University, Stanford, California.

Shellman, C. H. (1986), A new version of modesrch using interpolated values of the magnetoionic reflection coefficients, *Technial report*, Naval Ocean Systems Center, San Diego, CA.

Taylor, W. L. (1960), Vlf attenuation for east-west and west-east daytime propagation using atmospherics, *Journal of Geophysical Research*, *65*(7), 1933–1938, doi:10.1029/JZ065i007p01933.

Voss, H. D., W. L. Imhof, M. Walt, J. Mobilia, E. E. Gaines, J. B. Reagan, U. S. Inan, R. A. Helliwell, D. L. Carpenter, and J. P. Katsufrakis (1984), Lightning-induced electron precipitation, *Nature*, *312*, 740–742, doi:10.1038/312740a0.

Wait, J. (1957), The mode theory of vlf ionospheric propagation for finite ground conductivity, *Proceedings of the IRE*, *45*(6), 760–767, doi:10.1109/JRPROC.1957.278469.

Wait, J. R. (1959), Terrestrial propagation of very-low-frequency radio waves, *Journal of Research of the National Bureau of Standards*, *64D*, 153–203.

Wait, J. R. (1961), *Electromagnetic Waves in Stratified Media*, Pergamon Press, New York.

Wait, J. R., and K. Spies (1960), Influence of earth curvature and the terrestrial magnetic field on vlf propagation, *Journal of Geophysical Research*, *65*, 2325–2331.

Warren, R., P. Morris, R. Gupta, and G. Desrochers (1992), Omega system performance assessment, in *Position Location and Navigation Symposium, 1992. Record. 500 Years After Columbus - Navigation Challenges of Tomorrow. IEEE PLANS '92., IEEE*, pp. 88–95, doi:10.1109/PLANS.1992.185825.

Yabroff, I. (1957), Reflection at a sharply-bounded ionosphere, *Proceedings of the IRE*, *45*(6), 750–753, doi:10.1109/JRPROC.1957.278467.

BIOGRAPHICAL SKETCH

Michael Mitchell earned the B.S. degree in electrical engineering from University of California Davis, Davis, CA, in 2009. He earned the M.S. degree in electrical and computer engineering from the University of Florida, Gainesville, FL, in 2012, and continued to complete Ph.D in 2015 with advisor Dr. R. C. Moore. His research interests include VLF signal propagation in the Earth-ionosphere waveguide, VLF signal scattering from ionospheric disturbances, and the application of those observations toward the modeling of ionospheric processes, such as transient luminous events and lightning-induced electron precipitation. He is a member of AGU and IEEE.