

SPREAD-SPECTRUM VLF REMOTE SENSING OF THE IONOSPHERE

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
David Christopher Shafer
July, 1994

© Copyright by David Christopher Shafer 1994
All Rights Reserved

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.



Umran S. Inan

(principal advisor)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.



Robert A. Helliwell

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.



Arögyaswami Paulraj

Approved for the University Committee on Graduate Studies:

To my mother and father
William H. and Polly A. Shafer
and my fiancée
Kara

Abstract

A novel method for the characterization of transient, localized disturbances to the lower ionosphere, making use of the modulated Very Low Frequency (VLF) transmissions used for communication and navigation purposes, is presented in this work. These disturbances include those produced by energetic electron precipitation from the magnetosphere, which have been intensively studied for several years, and those produced by other effects, such as localized heating above intense lightning discharges. Radio signals transmitted for other purposes, such as the stations in the 20 to 30 KHZ range operated by the United States Navy, provide useful diagnostics of these disturbances. These signals propagate in the earth-ionosphere waveguide to long distances from the transmitter, and are influenced by both ionospheric and ground conditions along the path. A transient disturbance of the ionospheric conductivity within a few hundred kilometers of the signal path thus appears as a transient perturbation of the received signal. The disturbance can have a different effect on each of the frequencies which make up the modulated signal, particularly when the receiver is near a point along the signal path where several waveguide modes interfere destructively. Near these points, the predicted steady-state signal strength can vary by several DB between frequencies only 200 HZ apart, and the predicted perturbation due to changing ionospheric conditions can vary by as much as 1 DB across the same range. The measurement method exploits the finite bandwidth of these signals of opportunity (several hundred HZ in the case of most transmitters), employing a descision-directed technique to infer the probing signal, and using input/output analysis to estimate the transfer function of the propagation channel. With the addition of frequency as an independent variable, a direct comparison between experimental results and theoretical models can be made. This previously unavailable frequency dimension is explored by applying this new spread-spectrum analysis method to a data set collected in the Antarctic; the results suggest that several effects previously neglected in modeling lightning-associated ionospheric disturbances must now be taken into account, including coupling of waveguide modes within the disturbance, and the structure of the received field at the antenna. In addition, the steady-state propagation conditions often show variations with frequency which are indicative of the existence of several significant waveguide modes, which should further constrain ionospheric models.

Acknowledgements

As I write this, it's hard to believe that this long, sometimes frustrating, sometimes distracted, and oftentimes confusing process is coming to a close. Left to my own devices, I would have abandoned this path a long time ago; the fact that this document exists is largely the result of the work, support, encouragement, and love of many other people. I would like to thank my advisor, Professor Umran Inan, for his belief in me and in this work; it is largely his efforts which have kept this work on track over the years. The same is true for Professor Robert Helliwell, whose enthusiasm for the field, and formation of a world-class Antarctic scientific effort, brought me to Stanford in the first place, and for which I will be ever grateful. I was also very fortunate to have Professor Arogyaswami Paulraj on my reading committee, and I would like to thank him for taking the time to discuss this work, and to evaluate this dissertation. Much of the research history of these fields, and the relationships among the subjects that form the foundation of this work, are the result of the efforts of these three men. I would also like to thank Professor Abbas El-Gamal, for agreeing to chair my oral defense committee in the middle of a hot summer, when it would have been easy to bow out, and turn to more pressing matters.

My education, both specific and general, has benefited greatly from discussions with many of the research faculty, staff, and fellow students in the Space, Telecommunications and Radioscience Laboratory. In particular, I would like to thank Don Carpenter, Ron Bracewell, Tim Bell, Tony Fraser-Smith, and Martin Walt, for access to the wealth of information contained in their files and their heads; Vikas Sonwalker, for sharing his knowledge about everything from Jovian lightning to where to find good Dosas; Bill Trabucco and Jerry Yarbrough, for all the fine lunch conversations, and always knowing how everything worked and where everything was in the lab.

My fellow students at Stanford have also been supportive and helpful, as well as being the best spotters of a specious argument around. When I arrived at Stanford, I had the good fortune to have Tom Wolf as a mentor and friend, who showed me the way around so well. Tom and Bill Burgess did much of the field work at Palmer Station, and I would like to thank them both for doing such a fine job of setting up and calibrating the VLF equipment there, so that the results of Chapter 5 actually worked out. Also, I had a lot of fun with Bill on returning from Palmer in 1990, and many fine Chilean meals. Though

I never had the pleasure of working in the field with him, Juan Rodriguez has been a good friend, and very calming as we prepared for our orals. Much of the field work in this research group has been carried out by students, and I have had the pleasure of working with Steve Reising, Paul McGill, Alan Brown, and Dave Lauben in particular. As my background is in experimental work, I was most fortunate to have access to Lee Poulsen's waveguide propagation model, and to Steve Cummer, Sean Lev-Tov, and the aforementioned Juan, to explain how it worked, and what to do when it didn't. In addition, I have enjoyed the company and wisdom of many of my fellow students; this work is the better for the contributions of Tom Mielke, Jane Oh, Wai-Yeung Yip, Sasha Draganov, Yuri Taranenko, Jasna Djurovic, Victor Pasko, Curt Carlson, Hoc Ngo and Luis Sa, and I thank them for their discussions, comments, and work.

My working style could only be described (charitably) as peripatetic, and I would like to thank our strange crew, including Ken Britten, Inka Mimms, Dave Ellison, Linda Fletcher, Kim Fletcher, Aleksey Novicov, Steve Schoenberg, Nancy Serpiello, Alisa Koreneff, Stewart Carl, Nina Demediuk, Sarah Kelly, Karen Bollinger, Bob Smith, Linda Gass, Jamey Graham, Noam Eisen, Andy Butler, and Sue and Ed Smith, for hanging in there. This is doubly true for those friends who are also now clients, and have been waiting for this work to be done with bated breath.

Finally, special thanks are due to: my parents, William H. and Polly A. Shafer, for everything, and above all teaching me not to fear the truth; John Katsufakis, the 1983 Siple Station crew, Bill Trabucco, and Ev Paschal, for showing me the good results from hard work; Bill Scott, for sharing all these times, both good and bad, and being there; and Kara, for the future, and what we are becoming.

DAVID C. SHAFER

Stanford, California

July 25, 1994

The collection of the field data used in this dissertation, and the subsequent analysis and development which led to the bulk of the findings reported here, was made possible through the sponsorship of the National Science Foundation, Polar Programs Division, under grant DPP-9020687. The initial conception of the work was supported by Stanford University, through Seed Research Project 172-C-062 of the Center for Integrated Systems. The author wishes to thank the people of the United States, and the Board of Trustees of Stanford University, for making this work possible.

Contents

Abstract	v
Acknowledgements	vi
Contents	viii
List of Tables	x
List of Illustrations	xi
1. INTRODUCTION.....	1
1.1. MOTIVATION	2
1.2. PROBLEM DESCRIPTION	10
1.3. CHAPTER DESCRIPTIONS AND CONTRIBUTIONS.....	12
2. THEORETICAL BASIS	15
2.1. MEASUREMENT OF AMBIENT ELECTRON DENSITY PROFILES	15
2.2. MEASUREMENT OF ELECTRON DENSITY AND TEMPERATURE PERTURBATIONS	23
2.3. USE OF SPREAD-SPECTRUM SIGNALS AS MULTI-FREQUENCY PROBES	24
3. PROCESSING ALGORITHMS	27
3.1. FSK SIGNALS	29
3.2. MSK SIGNALS	30
3.3. PROCESSING OVERVIEW	35
3.3.1. Filtering and Clipping of Broadband Data.....	36
3.3.2. Demodulation of Broadband Data	39
3.3.3. Detection of Transmitted Bit Sequence.....	42
3.3.4. Reconstruction of the Transmitted Signal.....	44
3.4. COMPARISON WITH RECEIVED SIGNALS	45
3.5. SMOOTHING AND DISPLAYING RESULTS.....	50
4. SIMULATION OF PROCESSING SYSTEMS	57

4.1. PROCESSING OF SIMULATED MSK SIGNAL.....	58
4.1.1. Construction of the Test MSK Signal	59
4.1.2. Simulation of the Slowly-Varying Propagation System	61
4.1.3. Use of Spread-Spectrum Processing.....	67
4.2. PERFORMANCE OF EXISTING AMPLITUDE AND PHASE RECEIVERS	68
4.2.1. Amplitude Receivers.....	69
4.2.2. Phase Receivers	70
5. APPLICATION TO DATA.....	75
5.1. SIGNAL PATHS STUDIED	75
5.2. RECEIVING, RECORDING AND PROCESSING OF DATA	80
5.2.1. Antenna Efficiency.....	82
5.2.2. Preamplifier Frequency Response and Noise Figure.....	84
5.2.3. A/D Converter and Storage System.....	87
5.2.4. Noise Calibration Measurements and Dynamic Range	88
5.3. PROCESSING SYSTEM.....	89
5.4. MSK SIGNALS ON EXPECTED SINGLE-MODE PATHS.....	90
5.4.1. Time-Averaged Propagation Channel Transfer Functions.....	90
5.4.2. LEP-Induced Perturbations to Transfer Functions.....	96
6. SUMMARY AND SUGGESTIONS FOR FUTURE WORK.....	103
6.1. SUMMARY OF RESULTS AND CONTRIBUTIONS	103
6.2. SUGGESTIONS FOR FUTURE WORK	104
6.2.1. Impulsive Noise Removal for Broadband Recordings	105
6.2.2. Implementation of Real-Time Processing	105
6.2.3. Measurements at Separated Stations	108
6.2.4. Studying Shorter, Overland Paths	108
6.2.5. Use of Other Transmitters	108
BIBLIOGRAPHY.....	111

List of Tables

1. Transmitter and Arrival Characteristics	76
2. Processing Rate Requirements for Field Receiver	107

List of Illustrations

1.	Example of Lightning-induced Electron Precipitation (LEP) events.....	3
2.	Scattering of energetic electrons by a whistler-mode wave launched by lightning	6
3.	Perturbation of subionospherically-propagating VLF or LF signals by energetic electrons	7
4.	Decomposing a wave into a sum of normal modes	11
5.	System model used in this work	13
6.	Nighttime ambient profiles of electron density vs. altitude	16
7.	Predicted signal amplitude as a function of distance along the path, for a fixed ambient electron density.....	17
8.	Predicted signal amplitude as a function of distance along the path, for a fixed frequency (26.00 KHZ), and two ambient electron density profiles.	18
9.	Predicted signal amplitude as a function of frequency, for the NSU-to-Stanford path	19
10.	Predicted signal phase as a function of frequency, for the NSU-to-Stanford path	20
11.	Predicted measured signal phase as a function of frequency, for the NSU-to-Stanford path	21
12.	Comparison of measured propagation channel transfer function with the prediction, for the NPM-to-Palmer path using electron density profile II	22
13.	Predicted amplitude of the propagation channel transfer function perturbation, as a function of frequency, on the NSU-to-Stanford path near 26.0 KHZ	23
14.	Common VLF modulation methods	28
15.	Power Spectral Density of MSK-modulated signal.....	31
16.	Operation of an MSK modulator, implemented by phase-shift keying and amplitude modulation of two coherent carriers.....	32
17.	MSK waveform, illustrated for a specific bit sequence	33

18. Phase of the MSK-modulated signal, relative to an unmodulated carrier at the same center frequency	34
19. Overview of the processing steps used in this work, for MSK signals.....	35
20. Effect of initial processing steps on the broadband data	36
21. Amplitude and phase response of the high-pass filter used as the RF filter.....	37
22. Setting of the clipping threshold	39
23. Amplitude and phase of baseband filter used in the processing	40
24. Amplitude and phase output from baseband stage.	41
25. Reconstruction of the transmitted phase from the observed phase of the decimator output	43
26. Comparison of output to input fourier transforms to yield transfer function estimate	46
27. Output of the propagation channel transfer function estimate, for a sample data set	51
28. Smoothed sample data, based on the same data set as Figure 27	52
29. Estimate of the average transfer function of the propagation channel for a 4 minute data set, over ± 300 HZ in frequency	53
30. Spectrogram of the propagation channel transfer function residual, after removing the average transfer function estimate from the data	54
31. An expansion of two overlapping transfer function perturbations from Figure 30	55
32. Processing flow for the simulation.....	58
33. Autocorrelation of a sample received bit sequence.....	60
34. Location of simulated VLF transmitter ("NSU"), and the simulated ionospheric perturbation on the great-circle propagation path to a receiver at Stanford	61
35. Electron density profiles used in the simulation.....	62
36. Mode structure of the received signal at Stanford, from the fictitious 26.0 KHZ transmitter in Nebraska	63
37. Amplitude and phase of transfer function for propagation at 26 KHZ from "NSU" to Stanford, under the ambient conditions	64
38. Amplitude of propagation from NSU to SU as a function of frequency	65
39. Phase of propagation from NSU to SU as a function of frequency	66
40. Spectrum of the broadband VLF, with the synthetic NSU signal added at 26.0 KHZ	67

41. Spectrogram resulting from processing the simulated NSU signal	68
42. Typical block diagram of a frequency-doubling MSK phase receiver, with amplitude output.....	71
43. Waveform of MSK signal and an FSK signal synthesized from it by a square-law frequency doubler	72
44. Spectra of MSK signal and frequency-doubled FSK signal generated from it.....	73
45. Great-circle propagation paths for the signals studied in this section	76
46. Mode structure of the arriving VLF signals at Palmer Station	77
47. Predicted propagation from NPM in Hawaii to Palmer Station	78
48. Predicted propagation from NSS in Maryland to Palmer Station.....	78
49. Predicted propagation from NAA in Maine to Palmer Station	79
50. Predicted propagation from NLK in Washington to Palmer Station.....	79
51. Diagram of the receiving and recording equipment used to acquire the data analyzed in this chapter.....	81
52. Amplitude response of the Palmer antenna and preamp	85
53. Phase response of the Palmer antenna and preamp	86
54. Sensitivity of the Palmer antenna and preamp.....	87
55. Data recording system noise and signal detectability.....	89
56. Average propagation channel transfer function and 95% confidence limits for the NPM-to-Palmer path	91
57. Average propagation channel transfer function and 95% confidence limits for the NSS-to-Palmer path.....	91
58. Average propagation channel transfer function and 95% confidence limits for the NAA-to-Palmer path.....	92
59. Average propagation channel transfer function and 95% confidence limits for the NLK-to-Palmer path	92
60. Effect of mode structure of incoming wave on received signal, for a small loop antenna	94
61. Average transfer function of the propagation channel from NPM to Palmer, as observed on the E/W antenna.....	96
62. Two transfer function perturbations on the NPM-to-Palmer path,	97
63. An expansion of the time axis in Figure 62, showing a single LEP event.....	98
64. Four transfer function perturbations on the NSS-to-Palmer path	99
65. An expansion of the time axis in Figure 64, showing two LEP events	100

- 66. Four transfer function perturbations on the NAA-to-Palmer path,
coincident with those on NSS101
- 67. An expansion of the time axis in Figure 66, showing two LEP events.....102
- 68. Possible implementation of the processing described in this dissertation,
in a form suitable for use in the field107

1.

INTRODUCTION

Our direct experience shows that the earth's atmosphere is far from a steady environment; near the surface, weather conditions can change dramatically and rapidly. These changes are driven by thermal differences between parts of the Earth, and depend on the presence of a substantial atmosphere. As one ascends, however, the atmosphere thins out, and the thermal gradients which drive the weather near the surface become smaller. Weather phenomena familiar to surface observers rarely reach altitudes greater than 10 km or so. At the highest altitudes attainable with aircraft and balloons, about 40 km, "weather", as we know it on the ground, is non-existent, and one might expect that the atmosphere would simply continue to thin out, and blend into the vacuum of space.

In fact, this does not happen at all, because ionizing radiation from space, primarily in the form of the ultraviolet light from the Sun and cosmic rays from deep space, strips electrons from their atoms, and renders the upper reaches of the atmosphere electrically conductive. The ionosphere and magnetosphere, as these regions are known, exhibit a different type of "weather", in which electromagnetic forces play a central role. The ionosphere conventionally begins at about 60 km, and extends up to about 500 km. At ionospheric altitudes, particularly the lower regions which are the concern of this work, the rate of collisions for the free electrons is high, and chemical processes are important. At higher altitudes, the collision rate drops along with the density, and the medium begins to be dominated by the motion of the charged particles in the Earth's magnetic field; this latter region is known as the magnetosphere.

In the upper ionosphere and magnetosphere, it is possible to make *in situ* measurements from satellites, because the atmospheric density is low enough to allow a satellite to orbit for a reasonable amount of time. Below about 200 km, however, the atmospheric drag is high enough that satellite orbits decay rapidly, so that the only way to make *in situ* measurements of this region is with the use of sounding rockets. These

measurements can be important, but repetitive measurements are expensive, and it is difficult to distinguish between spatial and temporal variations.

Fortunately, we can make use of the electrical conductivity of the lower ionosphere to make indirect measurements, by utilizing the fact that the medium reflects low-frequency radio waves. The frequencies of interest here are in the VLF and LF range of the electromagnetic spectrum, covering the ranges of 3–30 KHZ (VLF), and 30–300 KHZ (LF). Radio waves at these frequencies are used for communication and navigation purposes, and propagate over long distances, being efficiently guided in the earth-ionosphere waveguide. Thus, they are highly sensitive to the ionospheric conditions along their paths, and form the focus of this work.

1.1. MOTIVATION

Interest in radio propagation in the VLF and LF bands dates back many years, and several studies were made to characterize the propagation channel [Davies, 1990, pp. 380–399; Field and Lewinstein, 1978; Turtle *et. al.*, 1989]. These studies were undertaken to support communication and navigation uses of the VLF and LF bands, and thus generally concern the measurement or modeling of the amplitude and phase of propagation, or the atmospheric noise, on coarse frequency and time scales. This type of approach has been sufficient for the intended uses – the design and evaluation of VLF and LF communication systems – but is not well suited to the study of localized and transient effects, which are of interest here.

The earth-ionosphere waveguide constitutes a dynamic channel for the propagation of VLF and LF radio waves, since the waves are sensitive to the boundary conditions at both the earth's surface and the ionosphere, and the electrical characteristics of the ionosphere vary widely. The electron density and temperature in the ionosphere result from a balance between the ionization of neutral atoms and molecules by energy input from a number of sources, and the removal of the free electrons by attachment and recombination processes, resulting in substantial variations over many time and distance scales. The main variation in energy input comes as a result of the change in the flux of ionizing radiation from the sun, which causes strong diurnal variations in the lower ionospheric layers. The solar UV energy during the daytime increases the equilibrium ionization levels, and, particularly at the lower elevations where the collision rate is high, the ionization levels drop soon after sunset. This diurnal




Figure 1. An example of Lightning-induced Electron Precipitation (LEP) events, observed on the phase of a subionospherically-propagating VLF signal originating in Hawaii, and received at Palmer Station, Antarctica. The top two panels show the “carrier phase” of the VLF signal at 23.3–23.5 KHz at different time scales, while the bottom panel shows the spectrogram of the broadband VLF received at the same time, displayed at the same time scale as the middle panel. The $\sim 5^\circ$ advances in phase, lasting ~ 1 second, are LEP events, and are seen to be associated with the arrivals of magnetospheric whistlers. The $\sim 3^\circ$ dips in the phase visible in the expanded scale are artifacts due to bit errors made in pairs by the MSK decoder.

variation is important for propagation of VLF and LF radio waves, since these frequencies are more sensitive to changes in the ionosphere at lower altitudes than are higher frequencies.

At night, when the background ionization at the lower ionospheric altitudes is low, other energy sources can have a strong effect on the electron density and temperature at these altitudes. Many of these sources are transient in time, and localized in space, and are results of complex processes occurring in near-earth space. An important example of a transient propagation effect, which has recently emerged as a significant process in the D region and for its coupling to the magnetosphere [e.g., Burgess and Inan, 1993], was first observed on VLF signals propagating to the Antarctic [Helliwell et al., 1973], and is illustrated in Figure 1. Such transient VLF perturbation events, originally named “Trimpi Events” after the Stanford researcher who first observed the coincidence between these

VLF signal perturbations and magnetospheric whistlers, are also referred to as LEP (Lightning-induced Electron Precipitation) events. These transient perturbations are believed to be due to secondary ionization enhancements produced by bursts of precipitating energetic electrons. The electrons are precipitated as a result of interactions in the magnetosphere with whistler waves originating in lightning discharges (Figure 2). LEP events have received much attention in recent years, both in the form of experimental work [e.g., *Wolf and Inan, 1990; Inan et. al., 1988b; Dowden and Adams, 1989; Smith and Cotton, 1990*], and theoretical studies [e.g., *Poulsen et. al., 1993a,b; Cotton and Smith, 1991*]. Although all the VLF perturbations analyzed in this work are of this type, there are several other known causes for sudden changes in VLF signal amplitude and phase, including excess ionization in the D region due to intense heating from lightning flashes [*Inan et. al., 1991; 1993; Taranenko et. al., 1992*], direct heating of the D region by powerful VLF, LF and HF transmitters [*Inan, 1990; Rodriguez et. al., 1992*], and heating and excess ionization following high-altitude nuclear explosions [*Allcock et. al, 1963; Helliwell and Carpenter, 1963*].

An LEP event occurs when high-energy electrons trapped in the magnetosphere interact with a whistler-mode wave launched by a lightning flash, are thereby scattered out of their stably-trapped orbits, and 'precipitate' into the lower ionosphere. There, they collide with the neutral constituents, and deposit their energy in the form of heat, x-rays, ionization, and light. The source of this whistler-mode wave most commonly is the radio wave radiated from the impulsive current contained in a lightning flash [*Uman, 1987, p. 114*], but other sources are possible as well. Much of this energy remains confined to the waveguide formed by the conducting surfaces of the Earth and the ionosphere, but some of the energy in the VLF frequencies can couple through the ionosphere, and into whistler-mode waves in the magnetosphere [*Helliwell, 1965*].

In the presence of enhancements to the ionization in the magnetosphere, aligned with the magnetic field lines, the whistler-mode wave can be guided along the field lines, and can again couple through the ionosphere to the ground in the geomagnetically conjugate hemisphere [*Helliwell, 1965*]. The whistler-mode wave is circularly polarized, and the index of refraction of the magnetized plasma for the whistler mode is large, so the velocity of propagation of the wave is relatively slow (typically 0.01–0.1 c , where c is the speed of light), so that it may take several seconds for the wave to propagate back to the Earth's surface. Furthermore, the medium is dispersive, with the lowest frequencies having the lowest propagation velocities, and the impulse which entered the magnetosphere is observed in the conjugate region as a falling tone spread over one

second or so; this characteristic accounts for the name of the propagation mode, *i.e.*, the *whistler*.

While in the magnetosphere, the slowly-moving wave encounters a population of trapped electrons, which execute a helical motion in the magnetic field. The transverse component of this motion occurs at the *gyrofrequency*, ω_H , given by:

$$\omega_H = \frac{eB}{m} \quad (1.1)$$

where e is the electronic charge, B is the local magnetic field, and m is the electron mass [Ramo *et al.*, 1984, p. 705]. The motion consists of two velocity components: v_{\parallel} , which is parallel to the magnetic field, and v_{\perp} , which is perpendicular to the external field. The sense of the electrons' motion is such that v_{\parallel} is opposite to the direction of propagation of the whistler-mode wave, while v_{\perp} is in the same direction of rotation as the perpendicular components of the wave electric and magnetic field. Thus, the electrons can experience a resonant interaction with the wave, called cyclotron resonance [Dungey, 1963], when:

$$\omega_H \cong \omega - k_{\parallel} \cdot v_{\parallel} \quad (1.2)$$

where:

- ω is the frequency of the whistler wave component,
- k_{\parallel} is the whistler wave number component parallel to the local magnetic field, and
- v_{\parallel} is the parallel component of the electron velocity.

Near the geomagnetic equator, where the local magnetic field exhibits a minimum, electrons can remain in resonance with a particular wave frequency for an extended time, and the wave-particle interaction is most efficient for fixed-frequency signals [Helliwell, 1970]. For whistler signals with varying frequency, the strongest interaction occurs at points along the field line which are as much as 10–12° away (in latitude) from the equator [Inan *et al.*, 1989]. A common result of the interaction is *pitch angle scattering* of the electrons, where the angular momentum of the particles are redirected due to their interaction with the wave, and the angle between the electrons' velocity vectors and the magnetic field are changed. While individual electrons may have their pitch angles either increased or decreased, there is a minimum pitch angle below which the electrons' mirror points will be in the lower, denser, regions of the ionosphere. Those electrons which have their pitch angles increased will remain trapped in the magnetosphere, but some of those which experience a pitch angle decrease will be scattered into the *loss cone*, which is the set

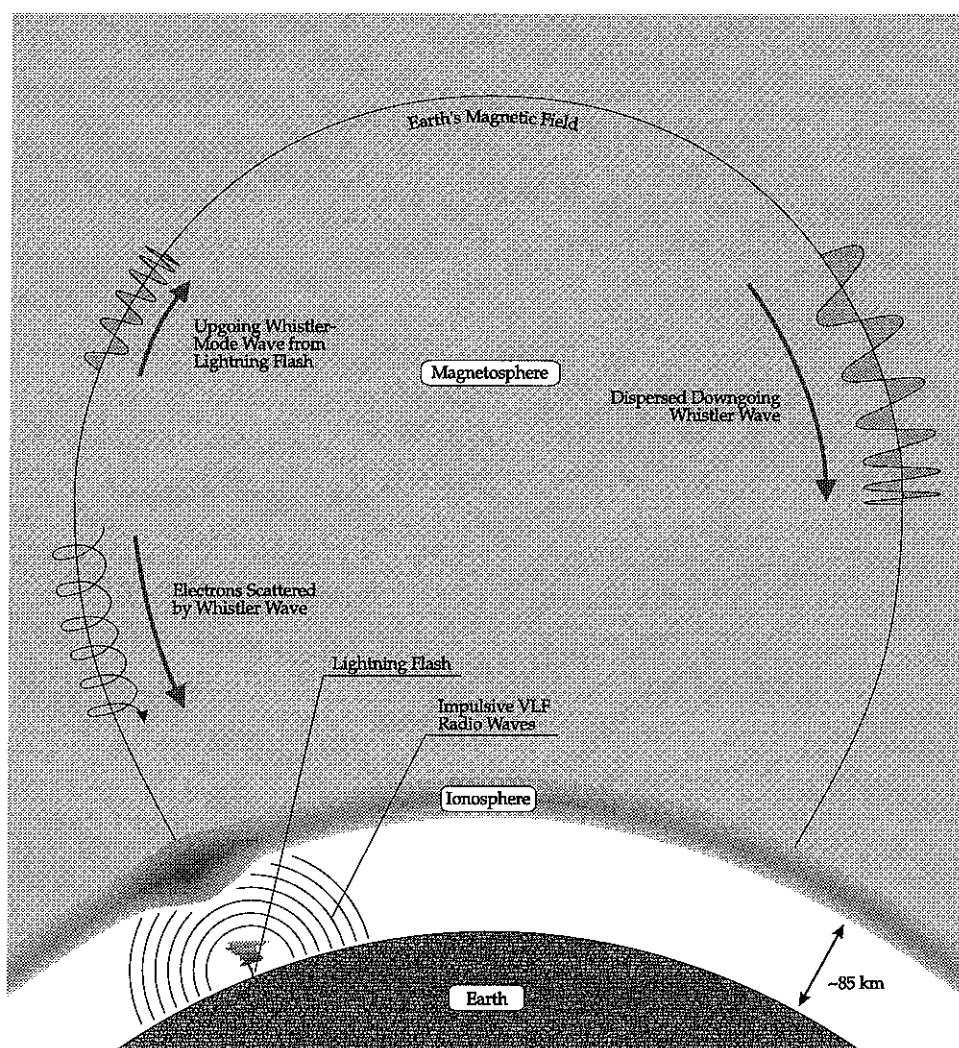


Figure 2. Scattering of energetic electrons by a whistler-mode wave launched by lightning. The lightning flash on the Earth's surface releases an intense impulse of electromagnetic energy, some of which is in the VLF region of the spectrum and can penetrate the ionosphere and travel along the magnetic field in the whistler mode. This wave can interact via cyclotron resonance with energetic electrons trapped by the magnetic field, which can amplify the wave, and lower the mirror points of the electrons' orbits. In order to show both the ionospheric and magnetospheric events in this figure, it is not to scale.

of pitch angles that correspond to mirror points in the lower ionosphere. As these scattered electrons subsequently travel down the field line, they are set to mirror at a lower altitude, and have a greatly increased probability of collision with neutrals in the ionosphere [Inan *et. al.*, 1985], as illustrated in Figure 2.




Figure 3. Perturbation of subionospherically-propagating VLF or LF signals by energetic electrons. When electrons' mirroring points are sufficiently lowered (for example, by the mechanism in Figure 2), collisions with neutral ionospheric constituents cause the electrons' energy to be distributed into the ionosphere, and create a region of secondary ionization. This appears to the VLF or LF wave as a change in conductivity of the earth-ionosphere waveguide boundary, which changes the propagation characteristics through the perturbed region. This change in propagation characteristics causes the signal observed by a suitably-located fixed receiver to undergo an amplitude and phase change; illustrated here as a phase advance and amplitude decrease by the thickness and position of the wave crests.

The incoming >100 KEV electrons carry a much larger energy than is required to ionize neutral constituents of the ionosphere, and can create a large patch of increased ionization at altitudes from 40–90 KM [Rees, 1963]. This excess ionization changes the conductivity of the affected ionospheric regions, and alters the propagation characteristics of VLF and LF waves in regions under or near the disturbance [Galejs, 1972;

Poulsen, 1991], as illustrated in Figure 3. It is these changes in amplitude and phase, observed with an appropriate receiver, that contain information about the ionospheric disturbance, and which are examined in this work.

Following the termination of the burst of precipitating electrons, the disturbed patch of the ionosphere is no longer in equilibrium, and chemical processes (recombination and attachment) cause the excess ionization to disappear over a period of time [*Glukhov et. al., 1992*]. This loss of the excess free electrons over time leads to a recovery of the perturbed signal levels over a period of 10–100 seconds, and accounts for the characteristic form of the disturbance, as seen in Figure 1. The details of this recovery process depend on the chemistry in the D region [*Glukhov et. al., 1992*], and the resulting signatures are expected to be an important source of information about the D region [*Pasko and Inan, 1994*].

Previous work on LEP events and associated VLF perturbations have concentrated on establishing the association of VLF signatures with whistlers [*Carpenter and LaBelle, 1982; Inan and Carpenter, 1986*] and lightning discharges [*Inan et. al., 1988a; Yip et. al., 1991*], on the determination of their occurrence characteristics [*Carpenter and Inan, 1987*], geographic and geomagnetic distribution [*Inan et. al., 1988b; Freidel and Hughes, 1993*], and geomagnetic conjugacy [*Burgess and Inan, 1990*]. Especially in the light of a new realization of the possible importance of this phenomenon for radiation belt loss processes on a global scale [*Burgess and Inan, 1993*], there is a need to extract quantitative information from the observed VLF signatures of LEP events. This task requires the use of an accurate model of the propagation of VLF waves in the earth-ionosphere waveguide, and their scattering from localized disturbances.

In the most general case, the propagation of radio waves, and their interactions with the surfaces of the Earth and the ionosphere, is accurately described by Maxwell's equations. However, it is impossible to solve the equations in closed form for realistic boundary conditions, and we must turn to approximate, computer-aided methods. One approach for the formulation of the propagation problem is to consider all the possible paths from the transmitter to the receiver, which satisfy the boundary conditions at all points. This leads to an infinite sum of 'ray' paths, though the higher-order rays contribute less to the total solution, because they are attenuated more at each ionospheric reflection, and travel over longer paths [*Morfitt and Shellman, 1976*]. This ray approach is most workable for shorter paths, because the angles of incidence for the rays rapidly become small for short separations between transmitter and receiver, with correspondingly higher attenuation, so that the solution may be written in terms of only

a few rays. However, for longer paths, many significant rays exist, and because the path geometry must be recalculated for each transmitter-receiver separation, the ray approach becomes cumbersome.

An alternative approach involves the expansion of the propagating wave in an infinite sum of normal 'modes' [Budden, 1985]. The modes are the eigenfunctions of Maxwell's equations with the appropriate boundary conditions applied, so that the propagation of a single mode (*i.e.*, the evolution of the amplitude and phase of the mode as it propagates along the path), in a waveguide with homogeneous boundary conditions, can be expressed in terms of a single propagation constant. Because Maxwell's equations are linear, an arbitrary input excitation can be expanded as an infinite sum of such modes, which can then be propagated independently to the receiver, as long as any changes to the boundary are slowly-varying, and satisfy the WKB approximation [Budden, 1985]. This process is illustrated in Figure 4, which shows the variation with height of a field component at the input to the system, and its decomposition into a sum of two waveguide modes. The modes propagate independently to the receiver, where they can be summed to give the structure of the output wave. Since, in general, the attenuation (*i.e.*, the amplitude variation) and propagation speed (*i.e.*, the phase variation) is different for different waveguide modes, the complete output wave need not look like the input in structure. Coupling between modes can occur at discontinuities (*i.e.*, changes occurring over distances of under a wavelength) in either the ionosphere or ground conductivity, and can be formulated by imposing the electromagnetic boundary conditions at such interfaces [Morfitt and Shellman, 1976].

For long (≥ 6000 km) propagation paths over seawater and at VLF frequencies, the effect on the VLF signal of a disturbance in electron density, N_e , can be modeled using a single dominant waveguide mode [Inan and Carpenter, 1987]. In such cases, the received signal perturbations are predicted to be in the form of an amplitude decrease and a phase advance, which qualitatively compares well with the experimental results for such paths [Wolf and Inan, 1990]. For paths where more than one waveguide mode is significant, the situation is more complex, and we must turn to computer models for predictions. Two-dimensional models, such as that of Tolstoy and Rosenberg [1983], can reproduce some of the observed features, but for studies of the effects of ionization patches of finite size and occurring in general on or off the great-circle propagation path, a full three-dimensional model must be employed. Such a three-dimensional model of VLF earth-ionosphere waveguide propagation was recently developed, and is utilized in this dissertation [Poulsen *et al.*, 1993a, 1993b]. The model is based on the two-dimensional VLF

earth-ionosphere waveguide propagation model known as the Long Wave Propagation Capability (LWPC) [Ferguson *et. al.*, 1989].

The three-dimensional model, and the underlying LWPC code, is in turn based on two approximations to the full solution of the modal propagation. First, the propagation of the waves from the transmitter to the receiver, from the transmitter to the scattering center, and from the scatterer to the receiver, are all assumed to satisfy the WKB approximation along a finite set of path segments. That is, the paths are broken into a finite number of "slabs", and coupling between propagating modes is ignored, except at the boundaries between slabs. The slab boundaries allow variations in ground and ionospheric conductivities (due to changes in electron density, temperature, and magnetic field) to be taken into account. Second, the change to the refractive index in the perturbed region is assumed small, so that the Born approximation would be applicable. This involves the use of the unperturbed field inside the disturbance, rather than a partly-scattered field, to determine the total field by summing the contributions of fields scattered by each small area of the disturbed region. These assumptions are justified for the disturbances of geophysical interest in our work, as well as for the ambient ionospheric conditions considered here [Poulsen *et. al.*, 1993b].

1.2. PROBLEM DESCRIPTION

To determine the electrical properties of the ionosphere, and their variations over time, we have at our disposal several relevant tools. First, the 3-D waveguide propagation model code described above can accurately predict, for a given set of ionospheric parameters, the received signal amplitude and phase at any frequency. In addition, there is available a network of VLF transmitting stations, operating in the 10-50 KHZ range, with radiated powers up to 1,200 KW. These VLF transmitters exist for other purposes, typically navigation and communication, and the strongest signals are useful for these purposes worldwide. While the transmitters employ a variety of modulation schemes, the three types that will most concern us here are the communication transmitters employing FSK, or, more commonly, MSK modulation, and the Omega navigation transmitters using a pulsed-cw format, with each station transmitting pulses at several frequencies sequentially.




Figure 4. Illustration of the process of decomposing a wave into a sum of normal modes. The input wave and output wave are shown schematically at the top panel. The variation of field strength with altitude is complex, and changes as the wave propagates. In a modal decomposition, the input wave is written as a sum of a (possibly infinite) set of normal modes, as illustrated in the bottom panel. The normal modes propagate with only an amplitude and phase shift if the medium is homogeneous, and linearity allows the output wave to be written as the sum of the modified normal modes.

The measurement of the temporal variations in amplitude and phase of waves which originate at a transmitter and arrive at the receiver over a given propagation path, can be viewed as a system identification problem, with the input consisting of a modulated signal covering several hundred Hz of bandwidth, and a center frequency in the 20-30 kHz range. The approach taken here involves treating the radiation, propagation, and reception of the modulated signal as a "gray box", and, using the

modulated VLF signals as a probe, attempting to identify parameters of the resulting system model which can be directly compared with the propagation model. Assuming the source excitation remains steady, the measured temporal variations are a direct result of variations in the characteristics of the propagation path. The system model used for this is illustrated in Figure 5.

The receiving system is well-characterized, and is both linear and time-invariant, so its transfer function can be transposed with the noise component additions, which reduces the system to a known input $x(t)$, followed by a linear, but slowly time-varying, system. The system output $y(t)$ is formed by adding to this the noise components, transformed by the action of the receiver. The effects of the propagation can be isolated, because the time scales over which significant variations occur in the ionosphere is quite different from that of the other variations in the system. Variations in the transfer functions of the transmitter and antenna, and of the receiver, are much slower than variations caused by the ionospheric disturbances of interest. The time scales for these variations is measured in hours or more, while significant transient changes in the propagation path are seen on time scales ranging from less than a second to many minutes.

By using the VLF probe signals, the transfer function from the modulator output to the receiver output may be measured. The propagation model referred to earlier gives theoretical information about the propagation effects at selected frequencies in the range of interest. These two may be directly compared, to test and refine estimates of the parameters which govern the VLF wave propagation, and which are of geophysical interest.

1.3. CHAPTER DESCRIPTIONS AND CONTRIBUTIONS

This work is organized into six chapters, which describe the background of the work, the techniques developed for analysis, and the results of the application of these techniques to the data set. Below are descriptions of the remaining five chapters, and a summary of the contributions they describe.

Chapter 2 describes the theoretical background for this work, and presents results of simulations of VLF propagation in the earth-ionosphere waveguide. The model calculations show that, under realistic ionospheric conditions, the propagation of the VLF waves is sensitively dependent on the electron density profile in the ionosphere, and that

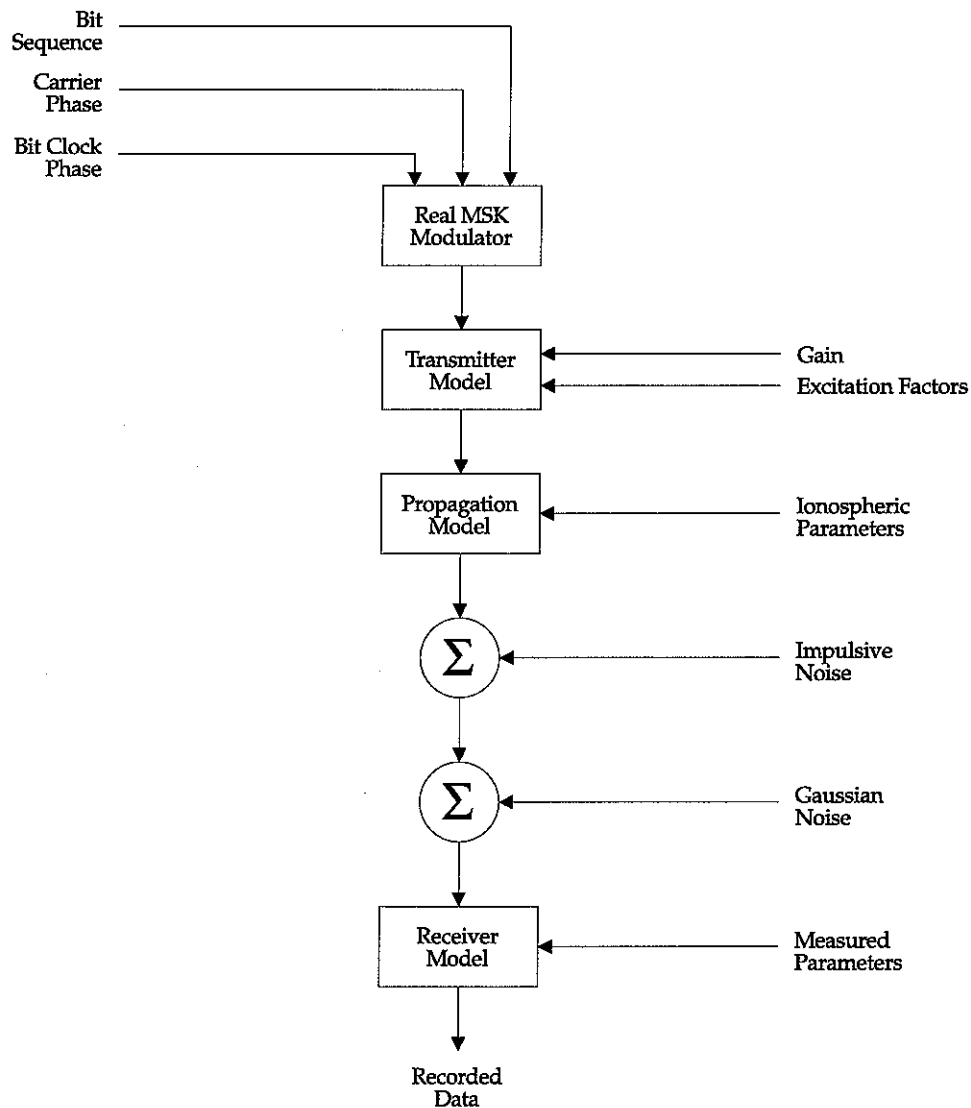


Figure 5. System model used in this work. The transmitter is represented by the top two blocks, with the unknown factors, which must be estimated from the data or elsewhere, listed as inputs. The propagation model is the 3-D propagation code, and estimating the ionospheric parameters listed as its inputs is the ultimate goal of this type of work. The noise sources represent the additive atmospheric noise, and the receiver model accounts for effects in the receiver and processing, which can be measured in advance.

this sensitivity may manifest itself as a strong dependence on the wave frequency of the signal amplitude and phase at a fixed receiver location.

Chapter 3 describes a means of using the modulated signals from VLF and LF transmitters to probe this frequency dependence, by decoding the message information contained in the modulated transmission, and feeding this information back to estimate the transmitted signal. This estimate is then compared in the frequency domain to the

received signal, resulting in a measurement of the transfer function of the propagation channel from transmitter to receiver, over a finite frequency range, with high time resolution. The contributions of the work described in Chapter 3 include:

- Development of a new VLF remote sensing technique, which fully utilizes the characteristics of modulated signals to determine substantially more information about propagation conditions than techniques currently in use.
- Introduction of frequency as an independent variable in VLF propagation measurements and calculations, providing additional experimental tests for theoretical models of geophysical effects.

Chapter 4 presents the results of a simulation of this processing, as applied to a short, overland propagation path. The results of this simulation show that:

- The technique described in Chapter 3, in the presence of a realistic atmospheric noise environment, can detect multi-mode propagation and mode coupling when these effects are important.

Chapter 5 uses the new technique to analyze a data set obtained from Palmer Station, Antarctica, which was recorded for this purpose digitally with high bandwidth and dynamic range. A number of LEP perturbations were recorded and analyzed, with the result that:

- Currently available propagation models give results consistent with the largely single-mode propagation studied.

The final Chapter of this work, Chapter 6, summarizes the above results. In addition, since the method used to undertake this study, involving high-resolution, broadband recording and post-analysis, is difficult to apply where large amounts of field data are desired, the final contribution of this work is:

- Consideration of the design of a field receiver implementing this technique, which allows reduction in interference and sensitivity to transmitter conditions compared with current designs.

2.

THEORETICAL BASIS

In order for a frequency-dependent radio diagnostic technique to yield useful information on geophysical processes, the propagation paths and/or the scattering of energy must exhibit discernible differences over the band of frequencies illuminated by the transmitted signals. These differences may be due to many conditions, including the ambient electron density profile, the ambient electron collision frequency profile, or the perturbations from the ambient levels which occur due to electron precipitation or heating. In this section, we investigate the frequency dependence of the propagation characteristics of VLF signals over the roughly 300 Hz bandwidth illuminated by the available transmitters. For this purpose, we utilize the multiple-mode, three-dimensional model of VLF propagation in the earth-ionosphere waveguide in the presence of localized D-region disturbances [Poulsen *et. al.*, 1993a], hereafter referred to as the 3-D waveguide propagation model.

2.1. MEASUREMENT OF AMBIENT ELECTRON DENSITY PROFILES

An important input to the 3-D waveguide model is the altitude profile of electron density and collision frequency, which must be specified for each segment along the propagation path. Typical values of electron density, observed at night at mid-to-low latitudes, are encompassed by the three different profiles in Figure 6; these representative profiles were also used in other recent work [Inan *et. al.*, 1992]. With the electrical conductivity of the ionosphere, which is a function of the electron density (N_e) and collision time (τ), as defined, and using a tabulated data base for the conductivity of the ground over the globe, the 3-D waveguide propagation model can be used to determine the amplitude and phase as a function of frequency and distance along any great circle path between a transmitter and receiver. The path chosen for our analysis is the same as

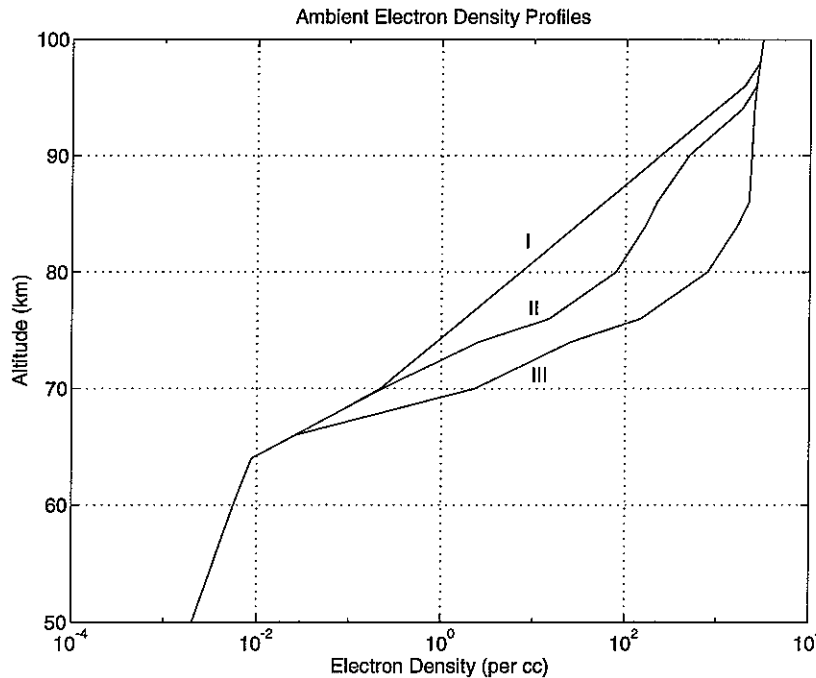


Figure 6. Three nighttime ambient profiles of electron density vs. altitude, which have been used in recent past work [Inan *et. al.*, 1992], and which encompass a 1:10:100 range over the important altitude range of 75–100 km. Profile I represents a tenuous ionosphere, Profile II represents a normal density, and Profile III represents an increase in density at the reflection altitude of a factor of about ten.

that used later in Chapter 4, and runs from a fictitious transmitter at 41.5° N, 97.6° W, to a receiver at Stanford (37.4° N, 122.2° W), operating at a center frequency of 26.0 KHZ.

Figure 8 shows the signal strength at the receiver, as computed using the 3-D waveguide propagation model, for the case where the density profile II shown in Figure 6 was assumed to be in effect all along the propagation path. One feature of the propagation, which is well-illustrated in Figure 8, is the occurrence of sharp nulls in the amplitude of the received signal, as a function of distance from the transmitter along a great-circle path. These nulls are due to the fact that the wave energy propagates in several waveguide modes simultaneously, and each of these modes has a different propagation constant. If no conversion of energy from one mode to another occurred, and the modes all had the same attenuation, the nulls would repeat every time destructive interference occurs between two or more modes with similar amplitudes, due to differences in the propagation constants. For VLF waves propagating in the earth-ionosphere waveguide, the attenuation rates of different modes differ [Poulsen *et. al.*,

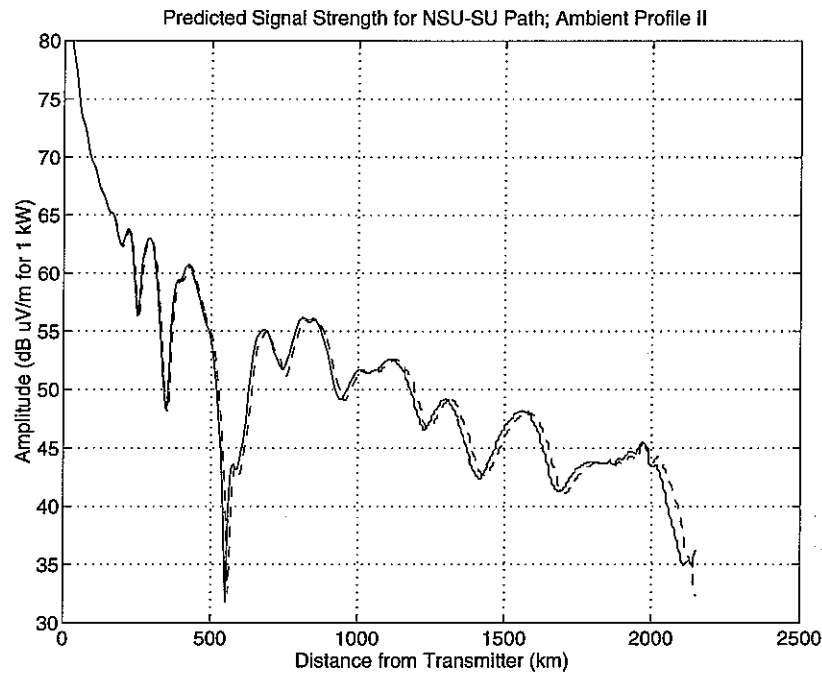


Figure 7. Predicted signal amplitude as a function of distance along the path, for a fixed ambient electron density. The solid line is for a frequency of 25.85 kHz, and the dashed line for 26.15 kHz.

1993a], and the nulls do not appear in a simple, repetitive pattern. Even though the attenuation of the waveguide modes differ, there are still places where two modes have nearly the same amplitude, since certain modes having higher attenuation rates are excited more efficiently by typical transmitting antennas. This can cause sharp nulls near the transmitter, with a transition to a smoother characteristic as many of the modes are attenuated with distance. The path illustrated in Figure 8 is relatively short, at about 2,100 km, and shows nulls of various depths all along the propagation path.

Although the difference between the predicted amplitude at 150 Hz below and above the center frequency is only one dB or so, in many locations along this particular path, this difference is easily measurable over time scales of a few seconds, as will be shown in Chapter 4. In addition, the modal interference produces numerous places along the path where there is much greater sensitivity to frequency; for example, note the shift in the position ($\Delta d \approx 20$ km) of the null near 550 km, and the large shift near the end of the path at 2,100 km. Differences in phase as a function of frequency are similar, with deviations of several degrees over a 300 Hz frequency difference, which are also easily measurable.




Figure 8. Predicted signal amplitude as a function of distance along the path, for a fixed frequency (26.00 KHZ), and two ambient electron density profiles. The solid line is for Profile II, and the dashed line for Profile I.

Although these differences are interesting, it is the sensitivity of these frequency differences to the selected electron density profile which is of interest here. These ambient density profiles, and the associated collision frequency profiles, are the result of a sum of many complex physical processes. Some of these processes create free electrons in the ionosphere (*e.g.*, solar uv and x-ray energy, cosmic rays, and precipitation of electrons from the magnetosphere), while others remove them (*e.g.*, chemical changes affecting attachment and recombination processes), and measurement of the actual ambient profiles under a variety of geophysical conditions can be expected to give information on the importance of these processes and the parameters governing them. In addition, the ambient profiles themselves enter into the calculation of the importance of other ionospheric effects, such as the heating of the ionosphere over powerful VLF transmitters. Finally, knowledge of the ambient profiles is important for predicting those aspects of VLF propagation of interest in communication and navigation applications. As shown in Figure 8, for example, two different density profiles lead to quite different predictions of the field strength as a function of distance along the path, with frequency held fixed at 26.00 KHZ; differences of as much as 15 DB are seen at some locations, with at least 5 DB

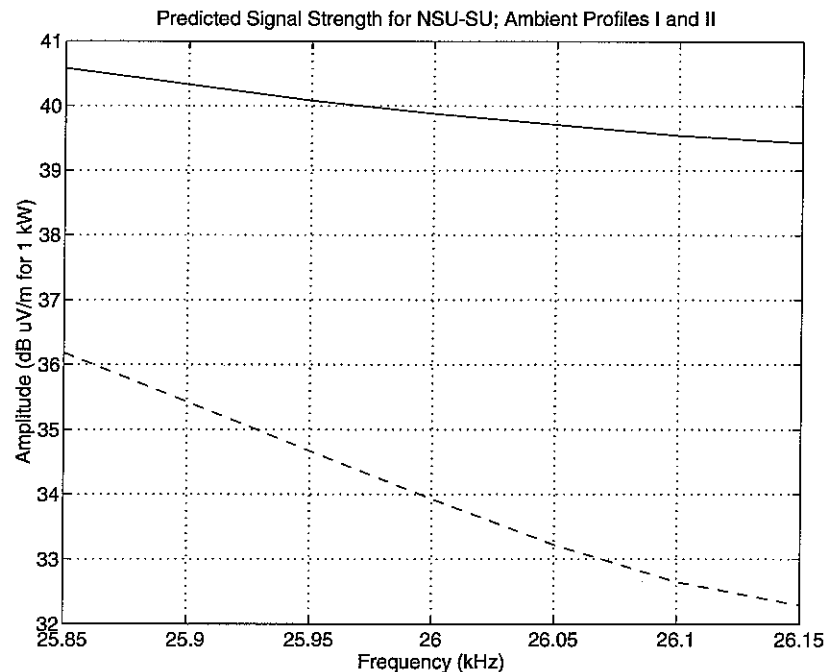


Figure 9. Predicted signal amplitude as a function of frequency, for the NSU-to-Stanford path. Two profiles are included; the solid line is for Profile I, and the dashed line for Profile II.

being common at most locations along the path. While these differences as a function of distance along the path provide another way to constrain the ambient profiles, direct utilization of this method requires simultaneous measurements at multiple points. The method used in this dissertation provides a means for the determination of the same constraints from a single-point measurement.

When frequency is introduced as an independent variable, a plot of predicted signal amplitude vs. frequency, for a set of candidate profiles, may be produced. Such a plot is shown in Figure 9, where the predicted signal amplitude is shown over a ± 150 Hz range around the nominal center frequency of 26.00 kHz, for two electron density profiles. It is clear from Figure 9 that an additional experimental test might now be applied to the ambient profile, since not only does the overall amplitude differ substantially between the two profiles, but there is an additional measurable difference in the dependence of the amplitude on frequency over the 300 Hz range.

A similar result may be obtained for the phase of the received signal, as shown in Figure 10. Here, phase differences of 15° – 40° are predicted, but, as will be discussed in Chapter 3, absolute phase delay and absolute group delay cannot be measured without

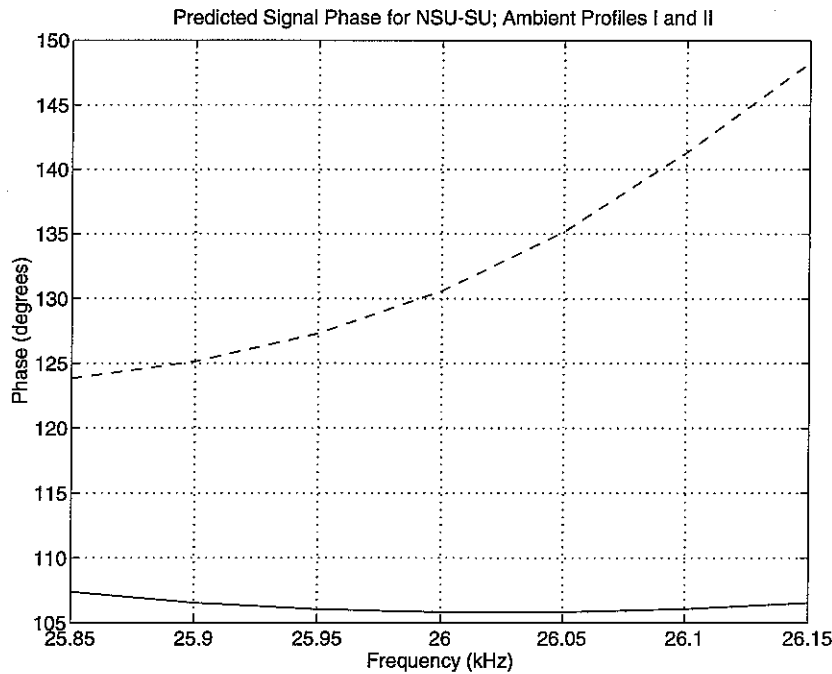


Figure 10. Predicted signal phase as a function of frequency, for the NSU-to-Stanford path. Two profiles are included; the solid line is for Profile I, and the dashed line for Profile II.

information about the transmitted carrier phase and bit transition times. An absolute phase delay results in a simple phase offset, and an absolute group delay in a linear trend of phase *vs.* frequency, so these must be removed from the predictions to match what is available in the data. Since recordings were not made at the transmitter in this study, the phase data are replotted in Figure 11, with the linear trends removed.

It is clear from Figure 11 that, even without absolute timing, the phase *vs.* frequency plots are sufficiently different across the 300 Hz band of frequency to discriminate between the two profiles.

The 3-D waveguide model makes predictions of the received *E*- and *H*-fields, but normalizes the result to 1 kW of radiated power, independent of frequency, and models the transmitting and receiving antennas as short vertical electric monopoles, using the techniques of *Ferguson and Snyder [1980]* to determine the initial excitation factors. Referring back to Figure 5, it is apparent that the technique developed in this work measures the transfer function from the MSK modulator output to the recorded signal from the receiver, and therefore measures the product of three possibly significant transfer function components: the transmitter, the propagation channel, and the receiver.

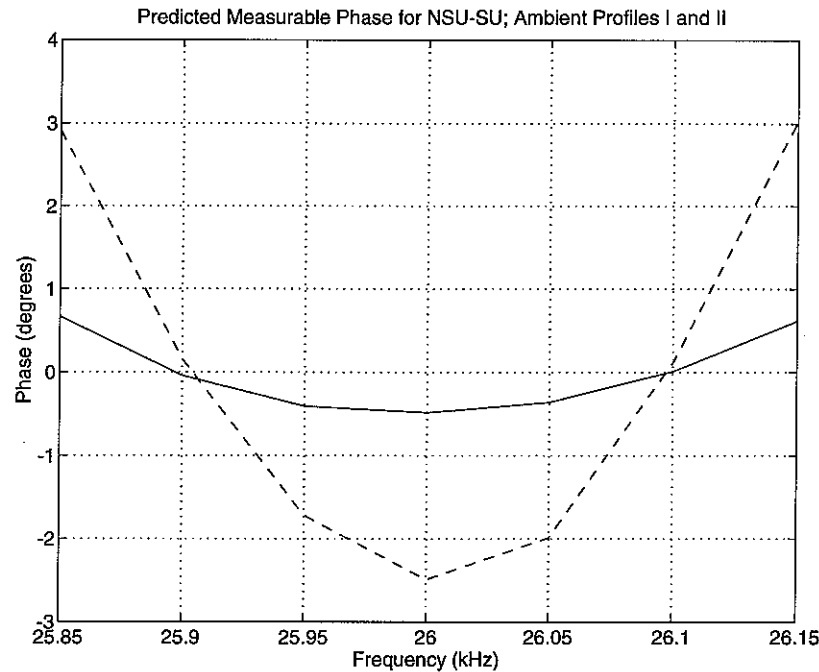


Figure 11. Predicted measured signal phase as a function of frequency, for the NSU-to-Stanford path. Two profiles are included; the solid line is for Profile I, and the dashed line for Profile II. This differs from Figure 10 due to suppression of constant- and linear-phase terms, which cannot be measured without knowledge of the transmitter's absolute phase and bit transition times.

The contribution of the receiver transfer function is easy to estimate, since the receiver is accessible for direct measurement, and its characteristics do not change over the time scales of this experiment.

The contribution from the transmitter requires more analysis to estimate, and, in the general case, would need to be measured. However, for the present study, we can make use of the regular properties of single-mode propagation over seawater to check the validity of modeling the transmitter as a vertical electric monopole over a ground plane, as is done in the 3-D waveguide model. To do this, we can use available data (see Chapter 5) from the NPM-to-Palmer path, which is a long path (12,349 km), entirely over seawater. The propagation along this path is well-understood, and fits the predictions of a single-mode theory quite well [Inan and Carpenter, 1987]. Because there is very little modal interference, and since the propagation constant does not change radically with frequency, the predicted amplitude and phase as a function of frequency are quite smooth, so that any variations which show up must be due to antenna effects.

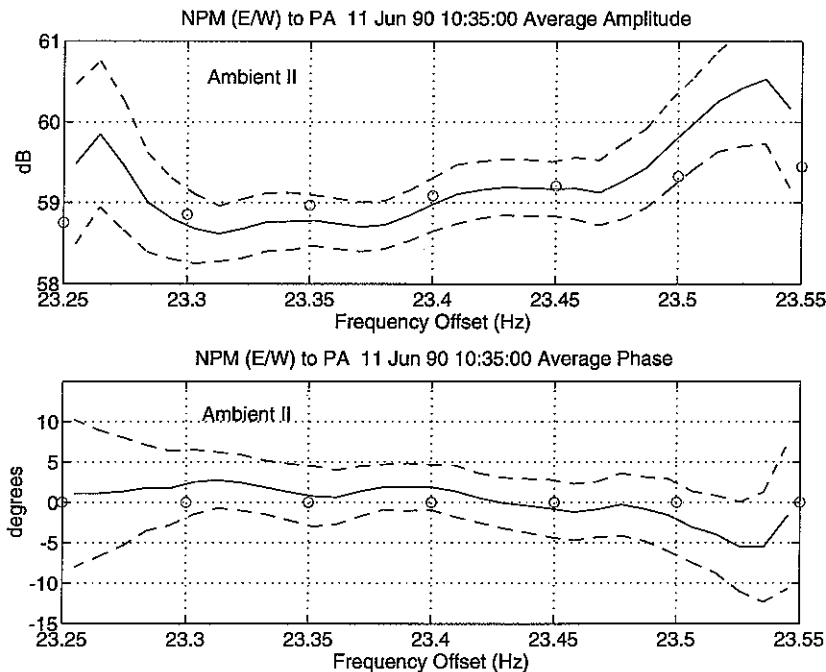


Figure 12. Comparison of measured propagation channel transfer function with the prediction, for the NPM-to-Palmer path using electron density profile II. The predictions (shown as circles) have been normalized so the total predicted amplitude is in agreement with the total received power, and to remove any constant phase or group delay, as required by the measurement technique. The measured transfer function was computed from data taken on the magnetic east-west (E/W) antenna, and is displayed as a solid line, with the 95% confidence error range indicated by dashed lines.

The results of this comparison, in Figure 12, show that the predicted received signal amplitude and phase match that actually observed to within the measurement error. The nature of the measurement and its error will be discussed at length in Chapter 3, but it is useful to note here that the measurement agrees with the prediction (for the chosen ambient electron density profile), to within ± 0.2 dB and $\pm 2^\circ$. The assumption outlined above, which treats both the transmitting and receiving antennas as short vertical electric monopoles, was used in formulating these predictions. The good agreement between data and model predictions indicates that this assumption is reasonable. For the remainder of this work, we will use this assumption, unless otherwise stated.

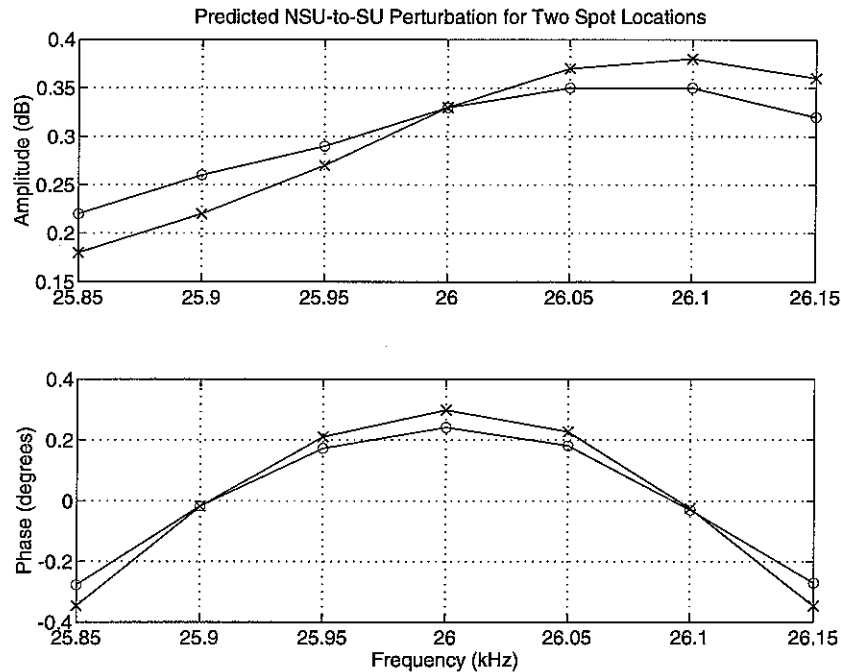


Figure 13. Predicted amplitude of the propagation channel transfer function perturbation, as a function of frequency, on the NSU-to-Stanford path near 26.0 KHZ. The circles represent the prediction for the perturbation located 500 KM from the receiver, and the crosses the prediction for the perturbation being 1000 KM distant from the receiver.

2.2. MEASUREMENT OF ELECTRON DENSITY AND TEMPERATURE PERTURBATIONS

Although the measurement of the ambient profile is of scientific interest, it is made more difficult by the important role played by the transmitter antenna excitation factors, which are difficult to estimate in the general case. However, *changes* to the ambient electron density and collision frequency profiles can be measured without this uncertainty, provided that the time scale of the changes to the density or collision frequency profiles are sufficiently different from any time scales associated with the transmitting antenna, and that the perturbations are not occurring too close to the transmitting or receiving antennas. In this case, the amplitudes and phases of the waveguide modes launched by the transmitter depend only on conditions which do not vary quickly with time. The receiving antenna, which converts incoming waveguide modes into a signal at the

receiver, acts similarly, and the contribution of the antenna excitation factors, although unknown, may be assumed to be fixed, and removed from the data record.

The predictions illustrated in Figure 13 were obtained by utilizing the 3-D waveguide model for two different density perturbations, on the 26.0 KHZ Nebraska-to-Stanford path considered earlier.

2.3. USE OF SPREAD-SPECTRUM SIGNALS AS MULTI-FREQUENCY PROBES

From the above discussion, it is apparent that the existing theory predicts significant differences in received amplitude and phase as a function of frequency, as either the ambient electron density profile, or certain characteristics of the perturbation to that profile (such as its location along the path), are varied. The performance of existing receivers, which treat the MSK signal as if it were a CW signal in many respects, are discussed in Chapter 4. However, the real advantage of the techniques presented here lie in the ability to make measurements across a band of frequencies, and the comparison of those measurements with the existing models.

There have been a few attempts to extend the existing single-frequency technology to make measurements at two or more frequencies. All of these efforts have used a two-frequency approach, which essentially uses two receivers, one tuned to each of the two main sidebands generated when an MSK signal is run through a frequency-doubling operation to generate an FSK signal (see Chapter 4). In this manner, it is possible to use MSK signals to measure perturbations in group, as well as phase, delay, provided that a model consisting solely of phase and group delay is adequate for representing the propagation (see Chapter 4 for a discussion of this approach).

The first use of the two-frequency technique in the literature appears to be in 1981, when it was applied to magnetospheric studies [Thompson, 1981]. In this study, the subionospheric VLF signal was assumed constant, and removed from the analysis, in order to focus on the (much weaker) long-delayed signals resulting from propagation in the ducted whistler mode. The MSK coding was demodulated only to examine the residual whistler-mode signal, and frequency spread of the MSK signal was not directly addressed. Nevertheless, this technique allowed the accurate determination of several important properties of whistler ducts.

Dowden and Adams [1989, 1992] have applied the two-frequency technique to the measurement of group delays on subionospheric signal perturbations. Although not noted in the paper, their method only applies to cases where significant inter-symbol interference (ISI) is not present. In such cases, the two-frequency method does indeed give an accurate measurement of the phase- and group-delay perturbation. The Dowden and Adams papers do not indicate whether the amount of ISI was monitored, although this information is in principle available from the MSK discriminator. Hence, it is not clear whether the data presented are truly a measurement of the best-fit phase and group delays.

Finally, the Omega navigation system has been used to measure the perturbations to the electron collision frequency profile, with low time resolution [*Barr et. al.*, 1985].

3.

PROCESSING ALGORITHMS

Since the primary purpose of VLF transmitters is to transmit data, the effects of this modulation on the received signal must be considered, and used in the analysis, in order to measure the transfer function of the VLF propagation channel. To do this, we must know the type of modulation employed, any parameters employed in the modulation, such as bit rate and carrier phase, and the data being transmitted. In analyzing the characteristics (*e.g.*, usable bandwidth) of various modulation schemes, some of the modulation parameters are often taken as unknowns, and treated as random variables, and the data transmitted are treated as samples from a random process. It is important to note, however, that in the case of any real measurement, the transmitted signal constitutes a single realization or sample function of this random process.

By treating the transmitted data bits as samples from a random process with known statistics, we can calculate the power spectral density (PSD) of the modulated signal. For this calculation, the data bits are assumed to be equiprobable, independent, and identically distributed, so that the resulting PSD is determined solely by the characteristics of the modulation method. Of the many digital bandpass modulation methods, three are commonly used in the VLF range: continuous wave (CW), frequency-shift keying (FSK), and minimum-shift keying (MSK). CW, in the ideal case, transmits no data at all, and the CW signal consists simply of a single sine wave, extending at constant amplitude and phase from $t = -\infty$ to $t = +\infty$. In practice, CW is a good approximation to those transmissions which have very small bandwidths (*i.e.*, constant-frequency pulse with durations much longer than one period), such as those of the Omega navigation system. The PSD of a CW process is simply a pair of delta functions at $\pm f_c$, where f_c is the carrier frequency, as shown (for positive frequencies only) in the left-hand panels of Figure 14. If we take the squared magnitude of the transform of any sample of such a process, we obtain a so-called *periodogram*, which converges to the PSD as the sample length tends toward infinity, since no randomness is involved.

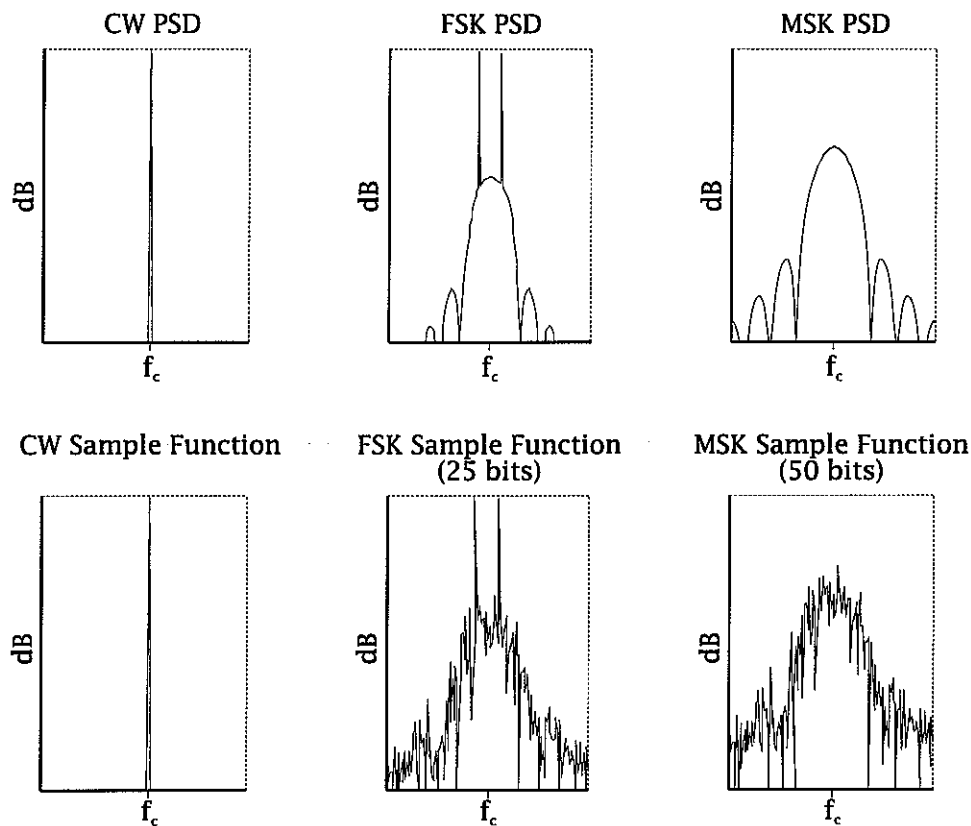


Figure 14. Common VLF modulation methods. The top panels show the power spectral density (PSD) of the three common modulation methods, while the bottom panels illustrate the finite-time transform of a particular realization of the underlying random process. The frequency axes are linearly scaled and zero-suppressed, and would apply, with appropriate scaling, to any combination of bit rate and carrier frequency. These plots cover ± 500 Hz from the carrier, and were generated with a carrier frequency of 10 kHz, and a frequency deviation of ± 50 Hz (bit rate of 100 bps for FSK, and 200 bps for MSK).

For any signal modulated in a non-trivial manner, however, the periodogram does not converge to the PSD as the length of the data examined increases. This is illustrated in the center- and right-hand panels of Figure 14, where the difference between the PSD and the finite-time spectrum is apparent. If longer sequences are used in computing the periodogram, the frequency resolution becomes finer, but the variance does not decrease. The longer periodograms can be smoothed [Oppenheim & Schaefer, pp. 548–549], to yield a better estimate of the PSD, but of interest here is the fact that all the “noise” apparent in the lower panels of Figure 14 is a consequence of the particular bit sequence transmitted. If the modulation method, the bit sequence, and the other relevant parameters such as carrier phase, are all known, then the transmitted waveform, and hence its spectrum, is completely known.

3.1. FSK SIGNALS

When the modulation employed is coherent binary FSK, the PSD may be derived by writing the transmitted signal as an infinite sum of pulses at two frequencies, at $\omega_c \pm \omega_\Delta$:

$$x(t) = A \sum_{k=-\infty}^{\infty} \cos(\omega_c t + \phi_c + \omega_\Delta a_k t) p(t - k / r_b) \quad (3.1)$$

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,
$\omega_\Delta = 2\pi f_\Delta = \pi r_b$	is the frequency offset from the carrier,
a_k	are the message bits (± 1),
r_b	is the bit rate, and
$p(t) = u(t) - u(t - r_b^{-1})$	is a one-bit pulse of unit amplitude.

An expansion of (3.1) into in-phase and quadrature (I and Q) components yields:

$$x(t) = \cos(\omega_c t + \phi_c) \cdot A \cos(\pi r_b t) + \sin(\omega_c t + \phi_c) \cdot A \sum_k (-1)^k a_k \sin[\pi r_b (t - k / r_b)] p(t - k / r_b) \quad (3.2)$$

The I component is simply a coherent cosine function, which has zero mean, and the Q component is a series of half-sine pulses, whose signs depend on the message bits a_k . As long as the a_k are equiprobable ± 1 , the Q component will also have zero mean, and the I and Q components will be independent, since the sine and cosine functions are orthogonal over complete half-cycles. Under this assumption, the I component produces a pair of impulses in the lowpass-equivalent spectrum, offset by $\pm r_b / 2$, and the Q component provides all the message-modulated component [Carlson, 1986, p. 520]:

$$\begin{aligned} G_Q(f) &= \frac{1}{4r_b} \left[\text{sinc} \frac{f - (r_b/2)}{r_b} + \text{sinc} \frac{f + (r_b/2)}{r_b} \right]^2 \\ &= \frac{4}{\pi^2 r_b} \left[\frac{\cos(\pi f / r_b)}{(2f / r_b)^2 - 1} \right]^2 \end{aligned} \quad (3.3)$$

This PSD is illustrated in the top-center panel in Figure 14, while the panel below it gives the transform of an actual bit sequence generated with this coherent binary FSK modulation. A receiver attempting to measure the amplitude and phase of such a signal can do so by measuring the amplitude and phase of either or both of the impulse components, since these are always transmitted with coherent FSK, and do not depend on the message bits being sent. Since the phase may be measured at two frequencies, the group delay, as well as the phase delay, of the propagation channel may be measured.

In practice, such a measurement suffers from two drawbacks. First, as a practical matter, it is impossible to suppress the information-carrying, modulated, portion of the FSK signal, without also affecting the measurement of the amplitude and phase of the carrier component, and any tradeoff which narrows the measurement bandwidth to reduce the message modulation reduces the time resolution in direct proportion. Second, when the propagation channel affects the received signal in a way which differs from simple phase and group delays, this information cannot be extracted from the received signal, since the spread-out Q component is removed in the processing. Therefore, even in the case of FSK, where coherent signal components exist, a measurement which identifies the bit sequence being transmitted and feeds back the result to estimate the propagation characteristics can be expected to perform better than simple filtering of the narrowband components.

3.2. MSK SIGNALS

For the frequency-dependent remote sensing needs as outlined in Chapter 2, the VLF stations transmitting MSK modulated data are nearly ideal. First, they provide relatively high signal levels, with several at or near 1 MW input power. Second, the strong MSK stations are geographically well situated, so that the propagation paths to receivers such as those at Palmer Station and Stanford cover a variety of terrain and expected ionospheric conditions. Finally, the MSK signal covers about 300 Hz of bandwidth with significant energy, and does so on a fairly continuous basis, provided only that the bit sequence being transmitted does not contain long strings of ones or zeros. Although this bandwidth is small by normal spread-spectrum standards, it is much wider than that explicitly used in previous remote-sensing work, which generally treats all VLF transmitters, no matter what the modulation, as if they were CW signals, by attempting to ignore the modulated components.




Figure 15. Power Spectral Density of MSK-modulated signal, illustrated here for a bit rate of 200 bits/sec, which is a rate commonly used by the transmitters in this study. The spectrum covers at least 200 Hz with significant energy.

MSK modulation, however, differs significantly from the CW and FSK discussed above, in that the MSK signal carries information on both the I and Q components, and there are no discrete spectral impulses present. The MSK carrier can be written in a form similar to Equation 3.1 above, except that the bit rate is twice that for the FSK, and in order to preserve phase continuity at the bit boundaries, the carrier phase must depend on the previous transmitted bits:

$$x(t) = A \sum_{k=-\infty}^{\infty} \cos(\omega_c t + \phi_k + \omega_{\Delta} a_k t) p(t - k/r_b) \quad (3.4)$$

where:

$$\phi_k = \frac{\pi}{2} \sum_{j=0}^{k-1} a_j \quad \text{is the carrier phase,}$$

$$\omega_{\Delta} = 2\pi f_{\Delta} = \pi r_b / 2 \quad \text{is the frequency offset from the carrier}$$

Since the carrier phase is now a function of the transmitted bits, the I and Q components become:

$$\begin{aligned} x_i(t) &= \sum_{k=0}^{\infty} \cos \left[\phi_k + a_k \frac{\pi r_b}{2} (t - k/r_b) \right] p(t - k/r_b) \\ x_q(t) &= \sum_{k=0}^{\infty} \sin \left[\phi_k + a_k \frac{\pi r_b}{2} (t - k/r_b) \right] p(t - k/r_b) \end{aligned} \quad (3.5)$$

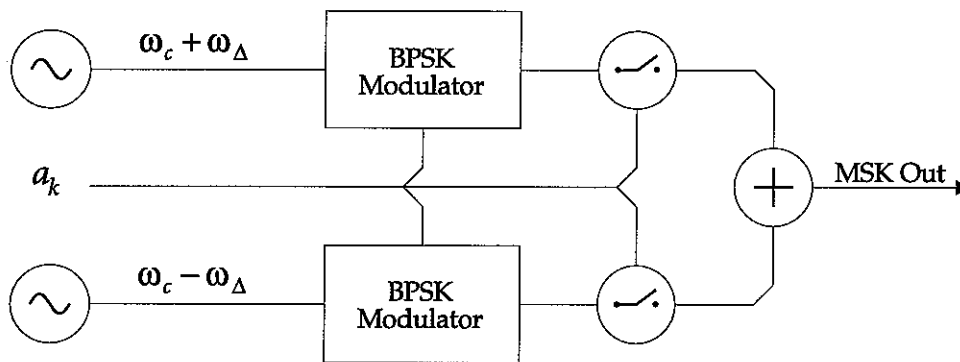


Figure 16. Operation of an MSK modulator, implemented by phase-shift keying and amplitude modulation of two coherent carriers. The BPSK modulators invert the phase of each carrier based on the number of message bits which have been sent on the other carrier. This preserves phase continuity at the transitions between bits, and destroys the coherence of the carriers.

and the I component no longer contributes fixed impulses to the spectrum, as was the case for FSK. Instead, both the I and Q components contain message information, which doubles the bit rate for the same bandwidth, thus making more effective use of available transmitter power and bandwidth. It can be shown [Carlson, 1986, pp. 522-525] that, for the case where the message bits are independent and identically distributed (I.I.D.), the lowpass equivalent spectrum of the MSK signal is:

$$G_p(f) = \frac{16}{\pi^2 T_b} \left[\frac{\cos(2\pi f / T_b)}{(4f / T_b)^2 - 1} \right]^2 \quad (3.6)$$

which is illustrated in Figure 15.

An alternate view of the MSK modulation operation, which is useful for the analysis of existing phase receivers [Wolf, 1990], is illustrated in Figure 16, and may be obtained by considering the two carrier frequencies separately. In order to satisfy the requirement of phase continuity at the bit boundaries, the phase of the “mark” and “space” frequencies (at $\omega_c \pm \omega_\Delta$) must be inverted whenever an odd number of bits have been sent on the other frequency. This can be done by passing each of the carriers through a binary phase-shift keying (BPSK) modulator, prior to on/off modulation which selects one carrier or the other for transmission. Because of the BPSK modulation, there is no coherent component in the output spectrum (for I.I.D. message bits), even though the amplitude modulation is always positive. This BPSK modulation spreads out the carrier energy over

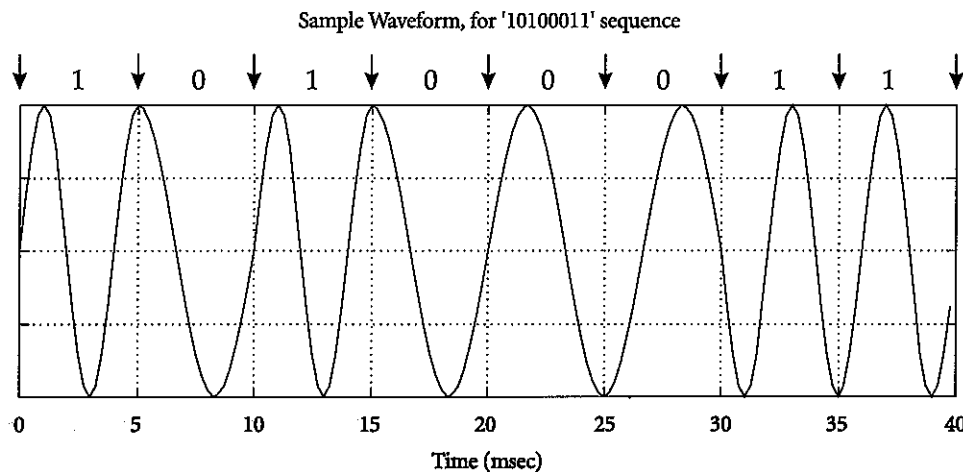


Figure 17. MSK waveform, illustrated for a specific bit sequence of '10100011', a bit rate of 200 bits per second, on a carrier of 200 Hz. For a real MSK wave, the carrier would be higher in frequency by two orders of magnitude, but this illustrates the relationship between the transmitted waveforms for each bit. Note, for example, that the 5-msec waveform corresponding to the second '1' is out of phase with the continuation of that for the first '1'. These inversions, which depend on the message bits, destroy the coherent carrier waves present in the FSK case.

a frequency range on the order of the inverse of the bit time, while ensuring that the two (BPSK-modulated) carriers are at the same phase at the bit boundaries.

While the above analyses lead to an understanding of the power spectrum of the MSK signal, the processing system described in this thesis works initially in the time domain to decode the bit sequence. Thus, an understanding of the MSK modulation in the time domain is needed. In MSK, at each bit interval, either the mark or space frequency is transmitted, but the signal may be inverted as needed to maintain phase continuity. This is illustrated in Figure 17 for a particular bit sequence, transmitted on a carrier which is at an artificially low frequency to illustrate better the modulation. In each bit time, the transmitted signal either gains or loses 90° of phase ('1' or '0' bit transmitted, respectively), with respect to an unmodulated sinusoid at the carrier (center) frequency. Because of this, both phases of each frequency are needed to generate the MSK-modulated signal, and, in general, the coherence of the carrier impulses present in the FSK spectrum is broken up.

One of the first steps in the signal processing described below is to mix the incoming broadband signal with a complex exponential at the carrier frequency to shift the desired transmitter down to baseband, where a suitably lowpass filter can extract the transmitter signal of interest. The results of this operation are a time-varying amplitude




Figure 18. Plot of the phase of the MSK-modulated signal, relative to an unmodulated carrier at the same center frequency, for the same '10100011' sequence used in Figure 17. The possible phase paths are shown as light lines, and all transitions occur on the bit boundaries, which here are 5 msec apart. The dark line is the phase path taken by the sample bit sequence. This phase gains 90° when a '1' is transmitted, and loses 90° when a '0' is transmitted.

and phase, relative to an unmodulated carrier exactly at the center frequency. For the MSK signals, the phase of this complex baseband signal is analyzed to demodulate the MSK signal, and thus determine the bit sequence transmitted. Since the MSK signal either gains or loses 90° of phase relative to the center frequency in each bit time, a plot of baseband signal phase vs. time can be described as a walk on a lattice, with the steps at each vertex (corresponding to the bit transitions) determined at the starting edge of the message bits. Such a plot is illustrated in Figure 18, and is used in the working of the demodulator algorithm described later in Section 3.3.3.

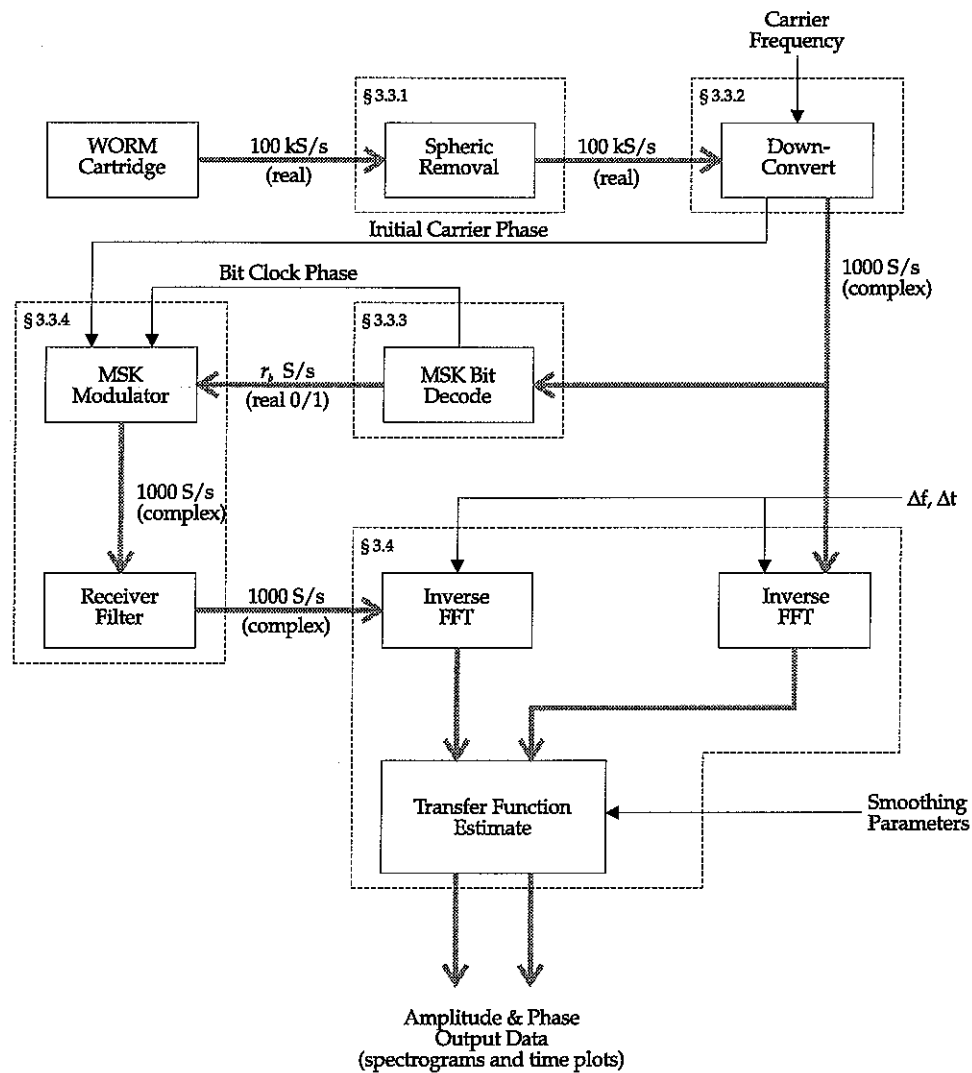


Figure 19. Overview of the processing steps used in this work, for MSK signals.

3.3. PROCESSING OVERVIEW

The data of interest for the development and application of the spread-spectrum remote sensing method put forth in this thesis concerns the measurements of transmissions from MSK-modulated VLF stations, operating in the 20–30 KHZ range. The processing flow for these data is shown in Figure 19. The input data, which consist of 16-bit samples of broadband VLF amplitude taken at a rate of 100 KHZ, are pre-processed to remove as much interfering energy as possible, and then downconverted directly to a




Figure 20. Effect of initial processing steps on the broadband data. The top panels display the time series and spectrum of the raw data, sampled at 100 KHZ, which shows the strong contamination from the Omega transmitter in Argentina, and a pair of sferics near 2 msec. The middle panels show the same data after the RF filter; the Omega interference is gone, but the sferic remains at high amplitude. The lower panels show the clipped output, which is used in all subsequent processing.

complex baseband signal sampled at 1 KHZ. The phase of the baseband signal is then processed to determine the transmitter parameters needed to simulate the transmitted signal, so that the simulated signal can be used to compare with the received signal in the frequency domain. Differences between the simulated transmitted and the received signals reflect propagation effects, which is what we aim to measure. The following sections describe the processing steps in more detail.

3.3.1. Filtering and Clipping of Broadband Data

As described in Chapter 5, the raw data arrive in the form of 16-bit samples of broadband signal amplitude taken at 100 KHZ, which are contaminated by various noise sources. The main noise source in this frequency band is impulsive noise due to lightning discharges, known as “sferics,” but there are also other interfering signals in the received data. These

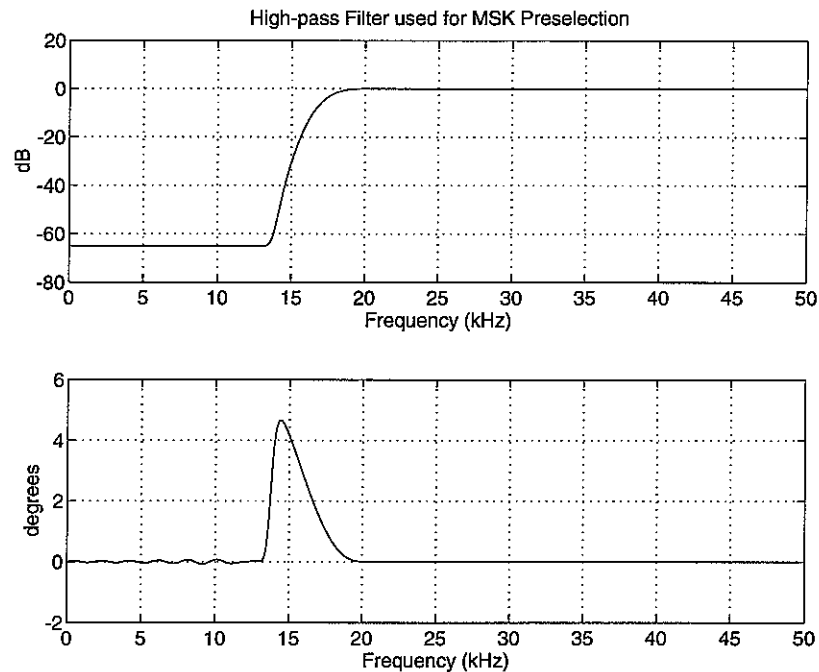


Figure 21. Amplitude and phase response of the high-pass filter used as the RF filter. The response is tailored to allow accurate measurement above 20 kHz, while rejecting the Omega stations and magnetospheric signals below 15 kHz.

are due to other transmitters, magnetospherically-generated signals such as auroral hiss (observed typically at high latitudes), and thermal noise. A representative segment of the raw data is shown in the top panel of Figure 20. For data recorded at Palmer Station, Antarctica, which constitutes the data set studied here (Chapter 5), contamination from a nearby VLF navigation transmitter (Omega Argentina) is particularly strong, with a signal intensity at least 10 dB above that of any of the MSK signals to be measured. In Figure 20, the 12.9 kHz Omega signal manifests itself as a dominant oscillation with a period of 77 μ sec. Also seen in Figure 20 is a sferic at about 2 msec, which has a long, low-frequency tail. During the initial sferic burst, both the Omega signal and the desired MSK signals are overwhelmed, but they can be recovered in between the individual bursts of impulsive noise, and during the slow tail.

To remove most of the magnetospheric signals, and some of the interfering transmitters, the first operation on the raw data is to pass it through a band selection filter. This filter is illustrated in Figure 21, and, since this filter is applied to the broadband data before resampling, it is relatively short (101 points). The results of applying this filter to the previously-illustrated broadband data are shown in the second panel of Figure 20.

The amplitude of the interfering Omega signal is considerably reduced, and the desired MSK signals are visible as a non-uniform signal envelope in the figure. This envelope is due to the fact that there are several MSK signals at comparable amplitudes, which interfere constructively and destructively over a time equal to the reciprocal of their frequency difference. It should be noted here that this difference is several KHZ, and is unrelated to the bit time of 5 msec.

Once the data have been filtered, considerable energy may remain in the sferics, which, because they are impulsive, have spectra that extend into the band of interest. The sferics produce two effects: they mask the desired signals, by causing the receiver to saturate, and they cause ringing in the narrowband stages to follow, particularly in the baseband filter. Since the larger sferics are strong enough to cause the receiver to saturate, there is no information available about the behavior of the desired MSK signals while the sferic is present (typically for 0.5–3 msec), and the only effect we can mitigate is the ringing in the baseband filter. To this end, the filtered output is clipped, with the threshold of clipping set on the basis of the distribution of the filtered signal.

To set the clipping threshold, a histogram is first built up of the amplitudes reached during the entire data run. This distribution of signal levels in the data consists of two parts. The main component is due to the sum of many MSK transmitters in the band, and, as the sum of a large number of independent sinusoids, is approximately gaussian. The other component is a high-amplitude tail, caused by the rare, but high-amplitude, impulsive sferics. Therefore, a simple analysis of the amplitude histogram can be made, and an appropriate cutoff found, above which the large majority of the excursions are due to interference from the sferics. This procedure is illustrated in Figure 22, which determines a clipping level of 0.30. The results of clipping the filtered data to this level are shown in the bottom panel of Figure 20.

In the course of this work, attempts were made to go beyond simply clipping out the sferics, by performing the threshold procedure above, and then setting the contaminated data to a "no data" value, rather than clipping it. However, such algorithms did not produce noticeable gains, relative to the dynamic clipping discussed above. In addition, suppression of data brings about substantial additional complexity in the subsequent processing, since the "no data" values must actively be handled at each processing step. Therefore, all the data presented in this work have been processed by the filtering and clipping methods discussed in this section.

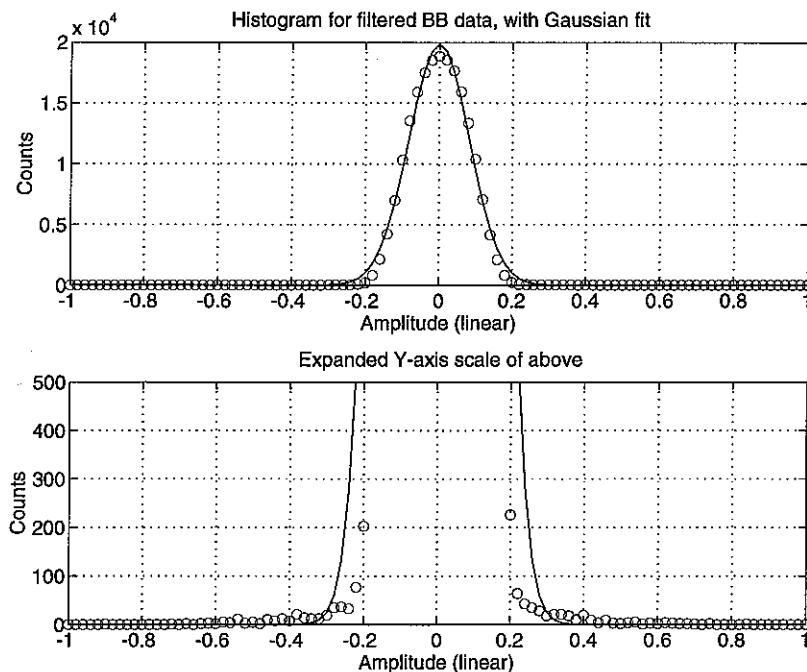


Figure 22. Setting of the clipping threshold. The algorithm fits the amplitude histogram data (open circles), to a gaussian, and uses the fit to adjust the clip levels. The bottom panel, on an expanded scale, shows the “tails” of the distribution, which are largely due to sferic energy.

3.3.2. Demodulation of Broadband Data

The first step in processing the filtered and clipped data is to bring it down to baseband, where the subsequent processing takes place. For this work, the baseband bandwidth needed is about ± 300 Hz, and the baseband is chosen to be sampled at 1 kHz. Since the processing is digital, frequency translation is implemented by mixing the broadband signal with a complex exponential, passing the result through a lowpass filter, and decimating the (complex) result down to the baseband, at zero frequency, and a sample rate of 1 kHz. The lowpass filter used in this work is shown in Figure 23, and is flat enough in both amplitude and phase so that the errors it introduces are not measurable compared to the transmitter and receiver unknowns. This filter is a very long FIR filter (3001 points), but, since it is a decimation filter, it only needs to be run once for every

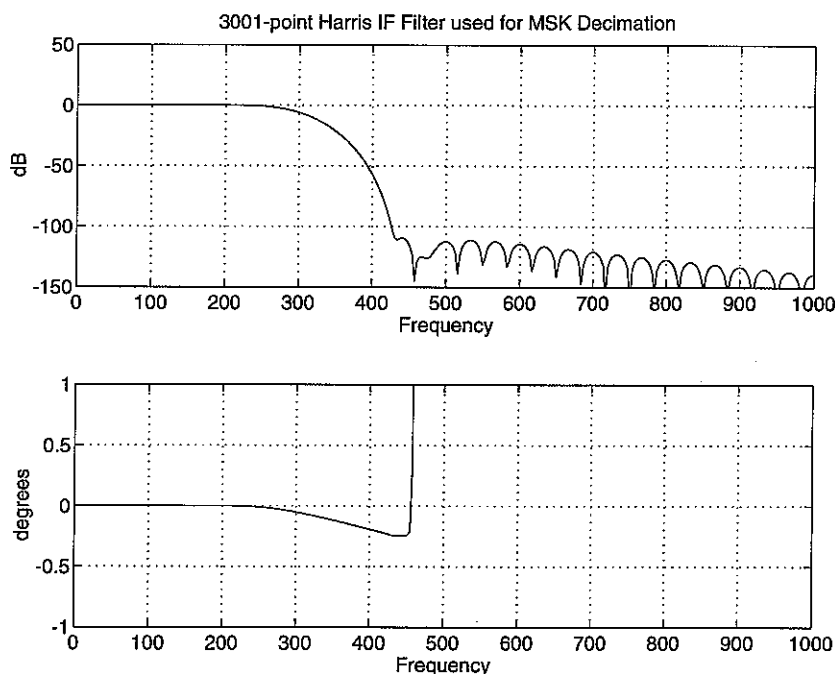


Figure 23. Amplitude and phase of baseband filter used in the processing. The filter illustrated is suitable for processing MSK signals, where the first spectral nulls are at ± 150 Hz from the center frequency. For this frequency range, the filter is flat to within 0.001 dB, and 0.001°. The contribution from a similar interfering transmitter 600 Hz away from the center frequency is down by 100 dB.

output point. All the data processed through this filter uses a decimation ratio of 100, which results in a substantial computation savings.

The baseband signal is a complex (I and Q) signal. This is convenient, since all the processing done after this step can be done at the baseband, and since the bandwidth of the VLF stations is only a few hundred hertz, the sample rate needed to capture this information is low. The use of a complex signal at DC, rather than a real signal centered about a subfrequency, allows the various filters in the system to be (possibly complex-valued) lowpass filters. For the case where the filter is symmetric about zero frequency, which is true of most of the filters used in this work, the filter coefficients are real, and the processing load is reduced by nearly a factor of two.

In addition, the amplitude of the baseband signal resulting from this process is simply the integrated amplitude of the MSK signal at the receiver, as would be read by a conventional heterodyne receiver, and the phase of the baseband contains the message modulation, in the form illustrated in Figure 18. The amplitude and phase of the

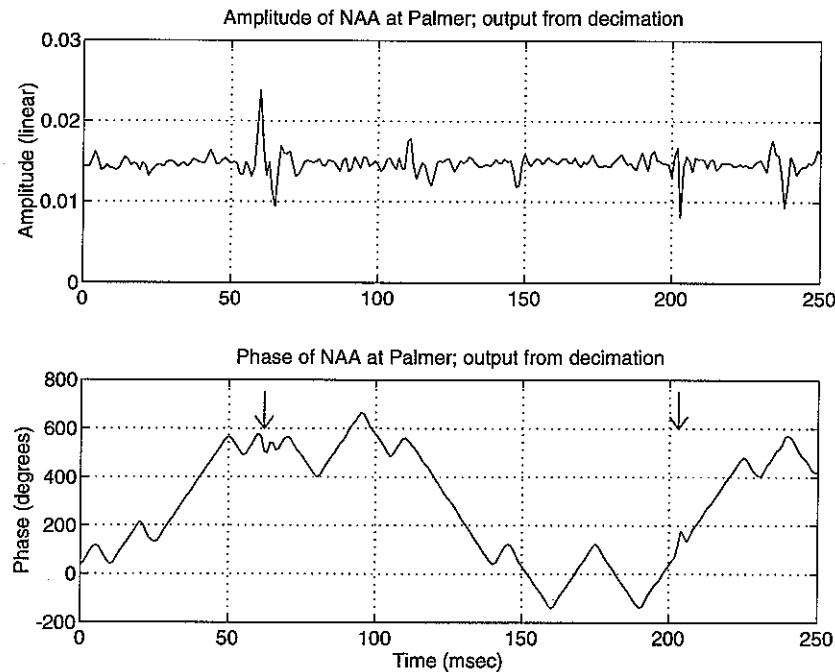


Figure 24. Amplitude and phase output from baseband stage. The top panel shows the amplitude of the output of the baseband filter, and the bottom panel the phase. The MSK modulation is clearly visible in the phase.

baseband signal for a particular example are shown in Figure 24, where the MSK modulation can clearly be seen in the lower panel. The positive and negative phase ramps from the modulation differ from the theoretical value of 90° per bit time, due to the additive noise (marked by arrows) and filtering. In addition to small deviations, there are three significant spurs during this 250-msec interval: one at 60 msec, one at 200 msec, and one at 240 msec. The effect of the baseband filter on the data is seen in the rounding of the transitions between positive- and negative-going phase ramps (*i.e.*, between '1' and '0' bits). This filter suppresses the more distant sidebands of the MSK signal, and thus rounds the transition edges, but, since the filter is zero-phase, this rounding is symmetric, and the location of the bit transition times can be accurately determined from the filtered phase.

3.3.3. Detection of Transmitted Bit Sequence

In order to estimate the transfer function of the propagation channel, the input signal must be known. For an MSK transmitter, five quantities must be known in advance, or must be reliably estimated from the data: f_c , the modulator's center frequency, a_k , the sequence of transmitted bits, Δf , the frequency deviation (or, alternatively, the baud rate), ϕ_c , the carrier phase, and T_g , the time of the bit transitions. For the VLF signals used here, the center frequency is precisely known, and controlled to a high degree of precision [ITU, 1979] by a frequency standard. Similarly, the baud rate is also known, from an inspection of the received data, since in practice only a few baud rates are used, with 200 bits/sec being by far the most common. Hence, only three parameters must be estimated from the data: a_k , ϕ_c , and T_g .

The time of the bit transitions, T_g , is estimated by considering the entire data record, and looking for the times at which the phase changes direction. Referring to Figure 25, those times ideally occur at bit boundaries. The program estimates the time of the transition, modulo the bit time, to a small fraction of the bit time, by comparing the phases before and after each stationary point on the phase plot. By this procedure, it is possible to get T_g to better than 1% of a bit time, while only observing the phase at a few points in each bit time. For each transition from '0' to '1' bits, or *vice versa*, the phase plot has an inflection point. There are spurious inflection points due to noise and interference, but these are randomly distributed in time, while those which define the bit time boundaries all occur nearly at multiples of the bit time. The estimation algorithm for T_g first builds a histogram of the inflection point time estimates (modulo the bit time), and then determines where the best-fit peak of this histogram falls, which results in a single value of T_g which is then used in the subsequent stages for the entire data run. The data runs used in this work are short enough (5 minutes maximum), so that this procedure works well, though a dynamic estimate of T_g could be implemented to allow the processing of large data sets.

Once the bit transition times are known, there are two possible hypotheses for the data received in any bit time: either a '1' was transmitted, or a '0' was transmitted. By examining the data points for a bit interval, the goodness-of-fit with each possibility is computed using a quadratic cost function. A provisional identification of the transmitted bit is made based on best fit at this stage.

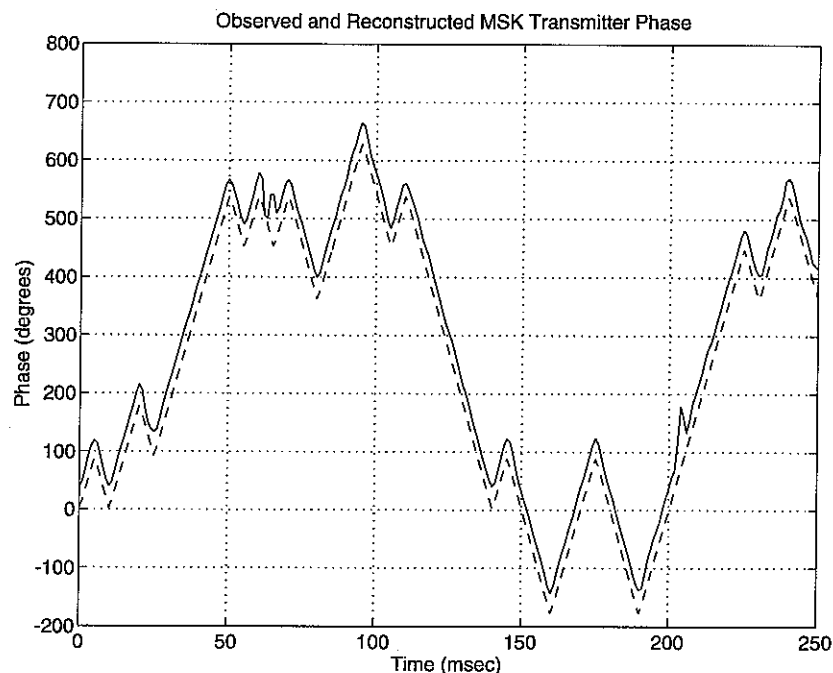


Figure 25. Reconstruction of the transmitted phase, and hence demodulation of the MSK signal, from the observed phase of the demodulator output. The solid line in the above graph is the received phase, as observed at the output of the baseband. There is noise, and the effects of two spherics can be seen in the data, as well as the rounded corners at the bit transition points, which is caused by the finite bandwidth of the baseband filter. The dashed line is the reconstruction of the transmitter phase from this data.

While the above scheme would be sufficient in the absence of noise, there is a good chance that some of the bit decisions will be incorrect. Referring to the observed data in Figure 25, the spheric which arrives at about 60 msec disrupts the phase completely for nearly 5 msec, and could easily lead to a bit error. A single bit error will result in the reconstructed phase advancing by 90° when it should have retarded, or *vice versa*, and, after the disturbance has vanished, the reconstructed phase will differ from the transmitted phase by a full 180° . This is a particularly egregious form of error for our purposes, because when the transfer function is estimated with this data, the presence of a 180° flip in the reconstruction causes the magnitude of the estimate to drop to zero, reverse phase, and rise again. This appears as sharp drops, all the way to zero, in the amplitude, and coincident 180° shifts in the phase, and makes the resulting data very hard to interpret.

In order to prevent such occurrences, the MSK decoder does not accept the provisional bit decisions without looking ahead; instead, it uses the goodness-of-fit values

to correct decisions which are likely to be in error. After a provisional bit decision is made on two bits, if the phase during the second bit is close to 180° off, and the goodness-of-fit is better for the second bit than the first, the decoder assumes that the first bit is in error, and reverses its earlier decision. Even though this scheme is not perfect at correcting bit errors, it forces the decoder to always make the errors in pairs, and the reconstructed phase never gets too far from the received phase. Since the expected perturbations dealt with in this work are always small, over the time spans comparable to the bit rates, this correction mechanism is not likely to be activated by a propagation or LEP perturbation. The reconstructed phase, as determined by the algorithm described above, is shown as the dashed line in Figure 25.

One additional possibility should be mentioned, and that is that use could be made of any redundancy present in the transmitted message to correct bit errors and determine the input signal to the propagation medium to a high degree of confidence. For many messages transmitted over such a channel, some form of forward error-correcting code is most likely applied, and knowledge of the basic code structure would provide an ideal way to determine the transmitted signal, extending the usefulness of these algorithms to lower signal-to-noise ratio regimes. If the system has been partitioned into a source encoder, which might contain a cryptographic component, and a channel coder, which serves only to correct channel errors, only knowledge of the latter would be needed for measurements of propagation characteristics. The algorithms investigated in this thesis need only obtain a good estimate of the transmitted signal; they do not need to interpret the message sent in any other way.

3.3.4. Reconstruction of the Transmitted Signal

In order to estimate the transfer function of the propagation path, which, as discussed in Section 1.2, gives information about the electron density profiles and perturbations to these profiles, the signal parameters measured above (a_k , ϕ_c , and T_g) must be fed back and used to reconstruct the signal launched from the transmitter into the earth-ionosphere waveguide. In addition, the output signal from the waveguide must be measured. Both of these processes involve antennas and transmitting or receiving electronics, and the contribution of these to the system model (see Figure 5) must be accounted for.

On the transmitting side, the signal parameters measured above can directly be used to reconstruct the modulated VLF signal. Referring to Figure 19, this function is performed by the box labeled *MSK modulator*. The inputs to this stage are a_k , ϕ_c , and T_g , and the output consists of an MSK-modulated signal at the system baseband, which, as discussed above, is a complex signal centered at zero frequency. If the estimated parameters are correct, the output will be an accurate representation of the signal which was presented to the power amplifier of the VLF transmitter. The issue of the relation between the transmitter input signal and the EM wave launched into the earth-ionosphere waveguide was discussed in Chapter 2, and the excitation coefficients for the various modes are included in the 3-D waveguide propagation model, so the MSK-modulated signal constructed here should be suitable for comparison with the model results, at least over the time scales of interest in this work.

Representation of the receiver is somewhat simpler. First, the receiver is available for analysis, and, as the characterization data presented in Chapter 5 show, the receiving hardware can be well-represented by a combination of a group delay, a phase delay and a frequency offset. The data are processed as 32-bit floating-point numbers, which have sufficient dynamic range for the contribution of the processing software to be largely due to amplitude and phase distortions in the baseband filter. The first two of these, the phase and group delays, are absorbed into the unknown transmitter phase ϕ_c and bit transition time T_g . The frequency offset is small enough to be adjusted for in the data display, which leaves only the response of the baseband filter to be accounted for. This is done in Figure 19 by the step labeled *receiver filter*, and is important only when bandwidths wider than about ± 300 Hz are examined, or a narrower than normal baseband filter is used. For these cases, the reconstructed MSK signal is resampled at the higher frequency used in the original data sampling (100 kHz), the baseband filter applied, and the result decimated back down to the baseband sample rate. Thus, if the estimates for the signal parameters are correct, and the transmitter is well-modeled by the 3-D waveguide propagation model, the differences between the signal at the output of this stage and the received signal are solely due to noise and propagation effects.

3.4. COMPARISON WITH RECEIVED SIGNALS

To perform the comparison between the received signal and the reconstruction of the transmitted signal, in a form which allows direct comparison with the model results, both

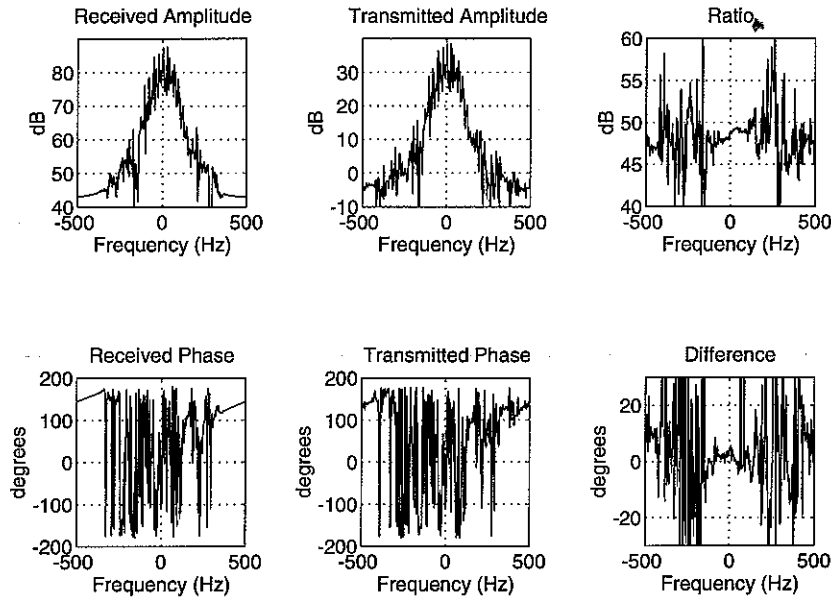


Figure 26. Comparison of output to input Fourier transforms to yield transfer function estimate. The left-hand panels show the transform of the received signal, over a time interval of 256 samples (about 50 bits), and the middle panels show the same for the reconstructed transmitted signal. The right-hand panels show the ratio of the two transforms, which is the raw transfer function estimate.

signals are converted into the frequency domain, and compared as spectra. This is illustrated in Figure 26, for a segment of 256 samples. The transform of both the transmitted and received signals contains a large amount of message variation, but, if the transfer function from transmitter to receiver is nearly constant, most of these variations will be common to both signals, and will drop out of the comparison. This removes the effects of the message, leaving only the transfer function estimate, as a function of frequency, plus noise. The noise level is roughly constant across the band of frequencies processed here, and so the accuracy of the transfer function estimate drops off at frequencies where the input signal power is less.

While Figure 26 illustrates the general procedure employed to estimate the transfer function, there are several refinements which must be introduced in order to obtain useful estimates in the presence of noise. First, the propagation is modeled as a simple linear system, with a known input and output, an unknown transfer function, and a stochastic noise term which is independent of the input:

$$y(t) = g_0(t) * u(t) + v(t) \quad (3.7)$$

where:

$g_0(t)$	is the unknown channel impulse response,
$u(t)$	is the input signal (inferred from data),
$y(t)$	is the output signal (measured), and
$v(t)$	is the noise.

If the additive noise $v(t)$ is uncorrelated with the input $u(t)$, and has zero mean, then the cross-spectrum of the output and input signals will be given by:

$$\Phi_{yu}(\omega) = G(e^{j\omega})\Phi_u(\omega) \quad (3.8)$$

where:

$G_0(e^{j\omega})$	is the unknown transfer function,
$\Phi_u(\omega)$	is the spectrum of the input signal,
$\Phi_{yu}(\omega)$	is the cross-spectrum of output and input signals.

This representation suggests forming a simple transfer function estimate as the ratio of the finite-length transform of the output signal to that of the input signal. This “raw” transfer function estimate is performed simply by taking the ratio of the Fourier transforms of the input and output functions, neglecting the additive noise term, and is referred to as the empirical transfer function estimate [Ljung, 1987, p. 146]:

$$\hat{G}_N(e^{j\omega}) = \frac{Y_N(\omega)}{U_N(\omega)} \quad (3.9)$$

where:

$Y_N(\omega) = \frac{1}{\sqrt{N}} \sum_{t=1}^N y(t)e^{-j\omega t}$	is the FFT of the output signal,
$U_N(\omega) = \frac{1}{\sqrt{N}} \sum_{t=1}^N u(t)e^{-j\omega t}$	is the FFT of the input signal, and
$\hat{G}_N(e^{j\omega})$	is the raw estimate at frequency ω .

$\hat{G}_N(e^{j\omega})$ is an asymptotically unbiased estimate of the transfer function $G_0(e^{j\omega})$, provided that the additive noise has zero mean, and that the input, $U_N(\omega)$, is not periodic over the length of the sample window N . This estimate is very general, in that it assumes only linearity of the propagation channel, but it is also very noisy, because it

calculates an independent estimate at each frequency point ω , and it is possible to show [Ljung, 1987, pp. 146–151] that the variance of the empirical transfer function estimate does not decrease with N . For this problem, as with many practical systems, we can obtain a better estimate by smoothing the raw estimates, based on the assumption that, for this propagation channel, the transfer function is not independently determined at each frequency. If we introduce the assumption that the true transfer function, $G_0(e^{j\omega})$, varies slowly compared to the frequency difference between each of the N output points, then each of these output points is an independent estimate of the same transfer function point, and has a variance that depends on the observed output signal spectrum at that point:

$$\text{Var} \left[\hat{G}_N(e^{2\pi jk/N}) \right] = \frac{\Phi_v(2\pi jk/N)}{|U_N(2\pi jk/N)|^2} \quad (3.10)$$

where:

Φ_v is the spectrum of the noise, and
 U_N is the observed transform of the input.

Therefore, a better estimate of $G_0(e^{j\omega})$ is:

$$\hat{G}_N(e^{j\omega_0}) = \frac{\int_{-\pi}^{\pi} W(\xi - \omega_0) \frac{|U_N(\xi)|^2}{\Phi_v(\xi)} \hat{G}_N(e^{j\xi}) d\xi}{\int_{-\pi}^{\pi} W(\xi - \omega_0) \frac{|U_N(\xi)|^2}{\Phi_v(\xi)} d\xi} \quad (3.11)$$

where:

$W(\xi - \omega_0)$ is a weighting function centered about ω_0 .

Now, we assume that the noise spectrum Φ_v is also smooth, and does not change much over the interval where the weighting function is non-zero. Then, this term in the integrals is a constant, and cancels out, since it appears in both the numerator and denominator. Thus, we can simplify (3.10) to:

$$\begin{aligned}
\hat{G}_N(e^{j\omega_0}) &= \frac{\int_{-\pi}^{\pi} W(\xi - \omega_0) |U_N(\xi)|^2 \hat{G}_N(e^{j\xi}) d\xi}{\int_{-\pi}^{\pi} W(\xi - \omega_0) |U_N(\xi)|^2 d\xi} \\
&= \frac{\int_{-\pi}^{\pi} W(\xi - \omega_0) |U_N(\xi)|^2 \frac{Y_N(\xi)}{U_N(\xi)} d\xi}{\int_{-\pi}^{\pi} W(\xi - \omega_0) |U_N(\xi)|^2 d\xi} \\
&= \frac{\int_{-\pi}^{\pi} W(\xi - \omega_0) U_N^*(\xi) Y_N(\xi) d\xi}{\int_{-\pi}^{\pi} W(\xi - \omega_0) U_N^*(\xi) U_N(\xi) d\xi} \tag{3.12}
\end{aligned}$$

which is the form of the estimator used in this work. In our implementation, a rectangular weighting function was generally used, and both the width of the weighting function and the length of the transform (N) used to compute U_N and Y_N can be chosen depending on the presentation desired.

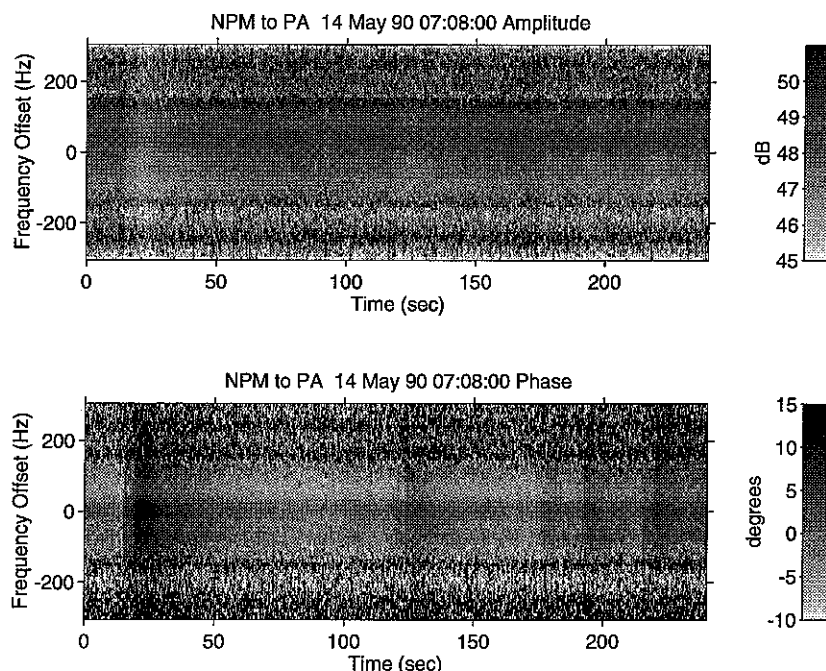


Figure 27. Output of the propagation channel transfer function estimate, for a sample data set. The top panel gives the amplitude, and the bottom the phase, of the transfer function, over a ± 300 HZ frequency range. Note that the falloff in the input signal power past about 120 HZ causes the transfer function estimate to be very noisy at these frequencies.

3.5. SMOOTHING AND DISPLAYING RESULTS

The results from this computation, for a typical data set, are shown in Figure 27. For this Figure, the estimator described above was used, with the FFT size (N) set to 256, with 100 points between FFT segments, and with a rectangular weighting function covering 10 HZ. Thus, there are 10 output slices per second, with 61 separate frequencies displayed, covering ± 300 HZ from the center frequency. The computed transfer function is displayed as a spectrogram, in two panels, the top panel being the amplitude (in dB) of the transfer function, and the bottom panel being the phase. This format is used from here on throughout this dissertation. The phase data have been processed so as to make best use of the spectrogram's limited gray-scale range, by setting the zero point of the phase and eliminating any drift between the local oscillator and the average of the received data. The darker portions of the display indicate higher signal amplitudes and

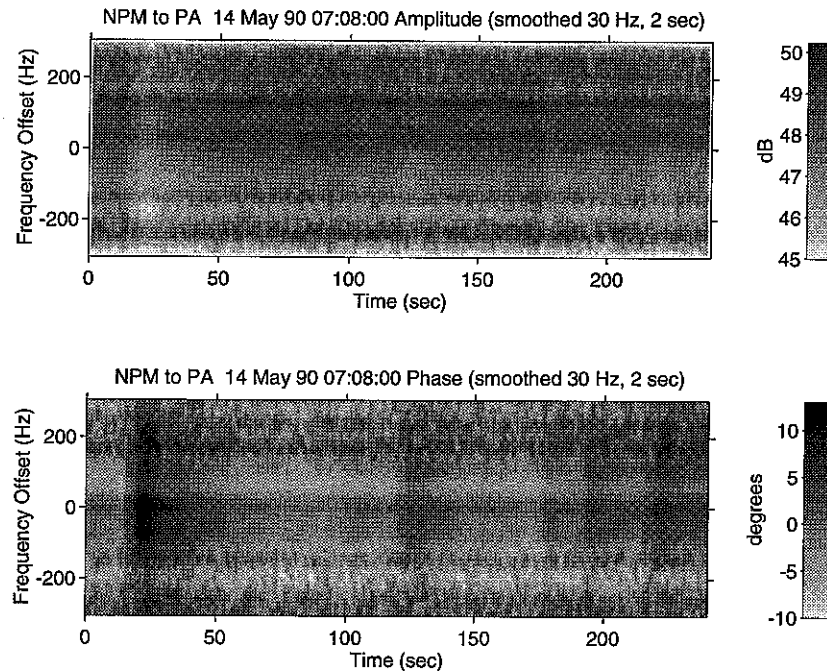


Figure 28. Smoothed sample data, based on the same data set as Figure 27. Note the changed z-axis scales.

phase advances¹. Visible in Figure 27 are several LEP-associated perturbations, as well as variations in the background propagation channel characteristics. The rapid falloff in the MSK signal spectrum beyond about 120 HZ causes the transfer function estimate at these frequencies to become dominated by the noise, and sets the limit on the frequency range over which the transfer function of the propagation channel may be determined by this technique.

The data presented may be further smoothed, depending on the intended use of the display, with some further reduction in the apparent noise, at the expense of time and frequency resolution. For example, the data in Figure 28 are averaged over 2 seconds, and 30 HZ. The sideband structure of the MSK-modulated signal is easily visible, since the noise which adds to the input becomes more significant as the available signal power goes down. However, in this example, as with most of the data, only the central ± 150 HZ band will probably be useful for most applications.

¹ This choice is appropriate since most LEP event signatures shown in this dissertation are phase advances.

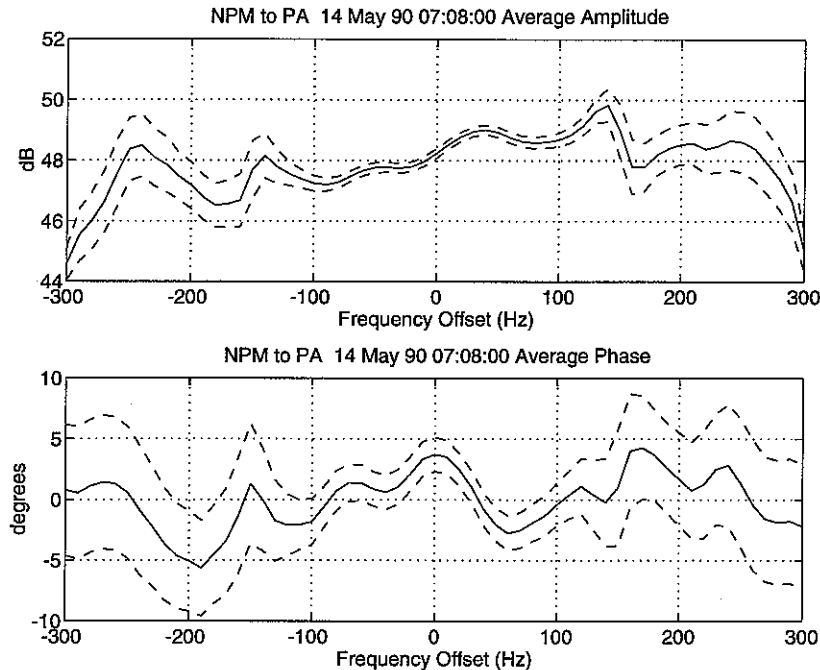


Figure 29. Estimate (solid lines) of the average transfer function of the propagation channel for the 4 minute data set, over ± 300 Hz in frequency, with the 95% confidence limits (dashed lines). The falloff in the power of the input msK signal beyond ± 120 Hz is visible here as an increase in the confidence limits.

These spectrograms contain all the information available about the propagation channel transfer function, for a given choice of FFT and smoothing parameters; however, for some purposes, further processing is useful to highlight certain features. For example, the time-averaged transfer function estimate may be derived by averaging the FFT “slices,” and compared directly against results of the 3-D waveguide propagation model, as shown in Figure 29. The high signal-to-noise ratio near the center frequency results in very good estimates of the transfer function in this range, though beyond about ± 150 Hz of the center frequency, the input signal is below the noise, and the transfer function estimate is biased and noisy. For this reason, the data presentations in the rest of this dissertation are generally limited to the useful range of ± 150 Hz for the msK-modulated signals.

When the background propagation channel transfer function is removed from the data, the residual highlights transient changes to the transfer function, such as those caused by the LEP events discussed in Chapter 2. Figure 30 displays this residual, as the deviation from the average transfer function (Figure 29). In this figure, the full 4 minutes and ± 300 Hz range of the data are displayed.

For detailed examination of an individual perturbation to the propagation channel transfer function, the time and frequency scales may be expanded, as shown in Figure 31. These transfer function perturbations are typical of the LEP-induced events analyzed in Chapter 5. They show a sharp onset (~ 1 second risetime), and a slow decay (~ 20 seconds fall time), and are uniform across the ± 150 HZ frequency band.

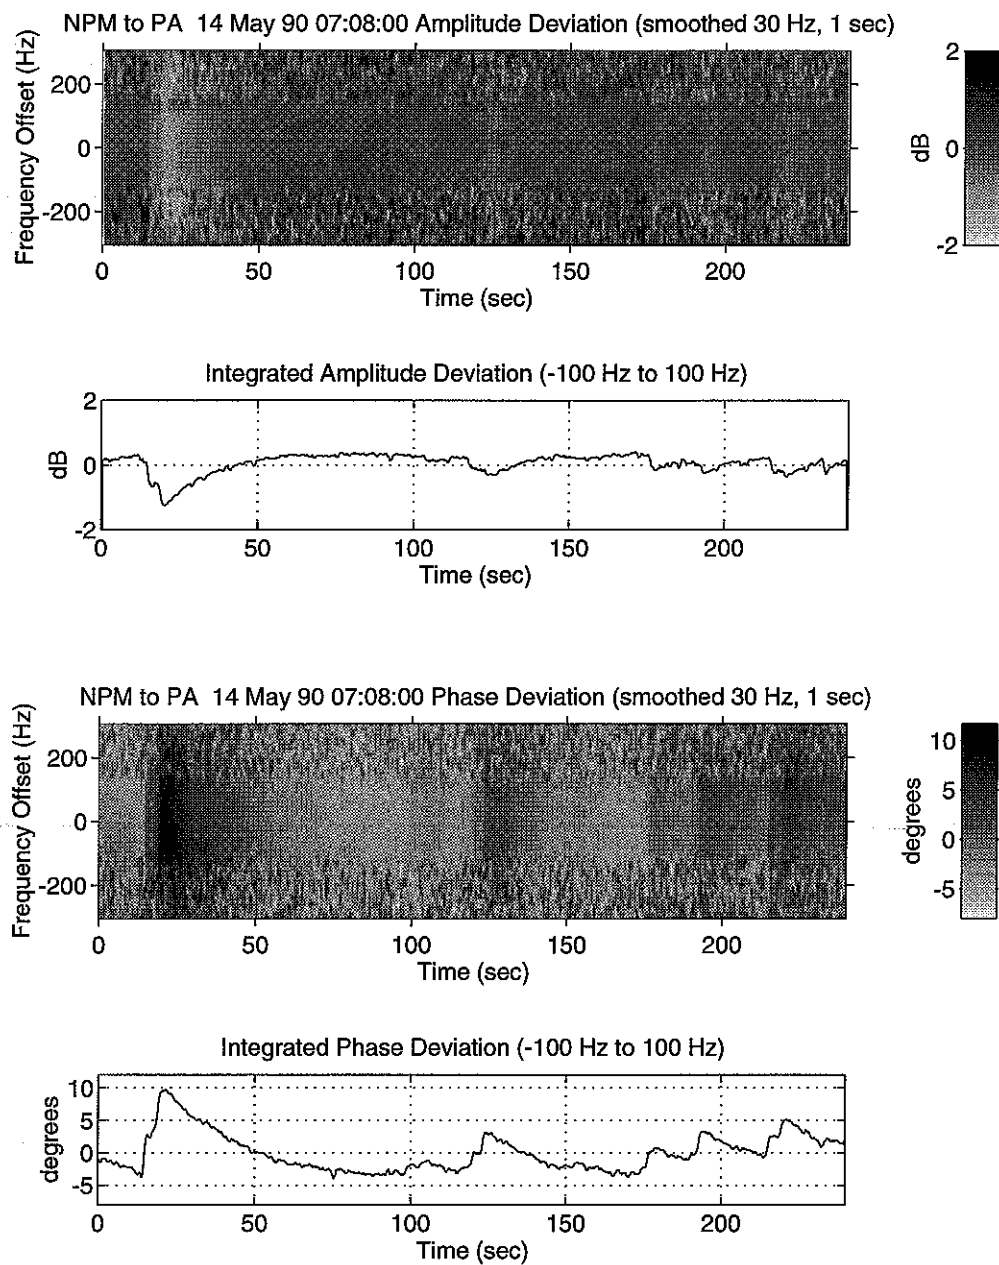


Figure 30. Spectrogram of the propagation channel transfer function residual, after removing the average transfer function estimate from the data. This display highlights the transient disturbances to wave propagation, of which at least six events characteristic of LBP are visible. All these perturbations are amplitude decreases and phase advances, and do not show significant variation with frequency.

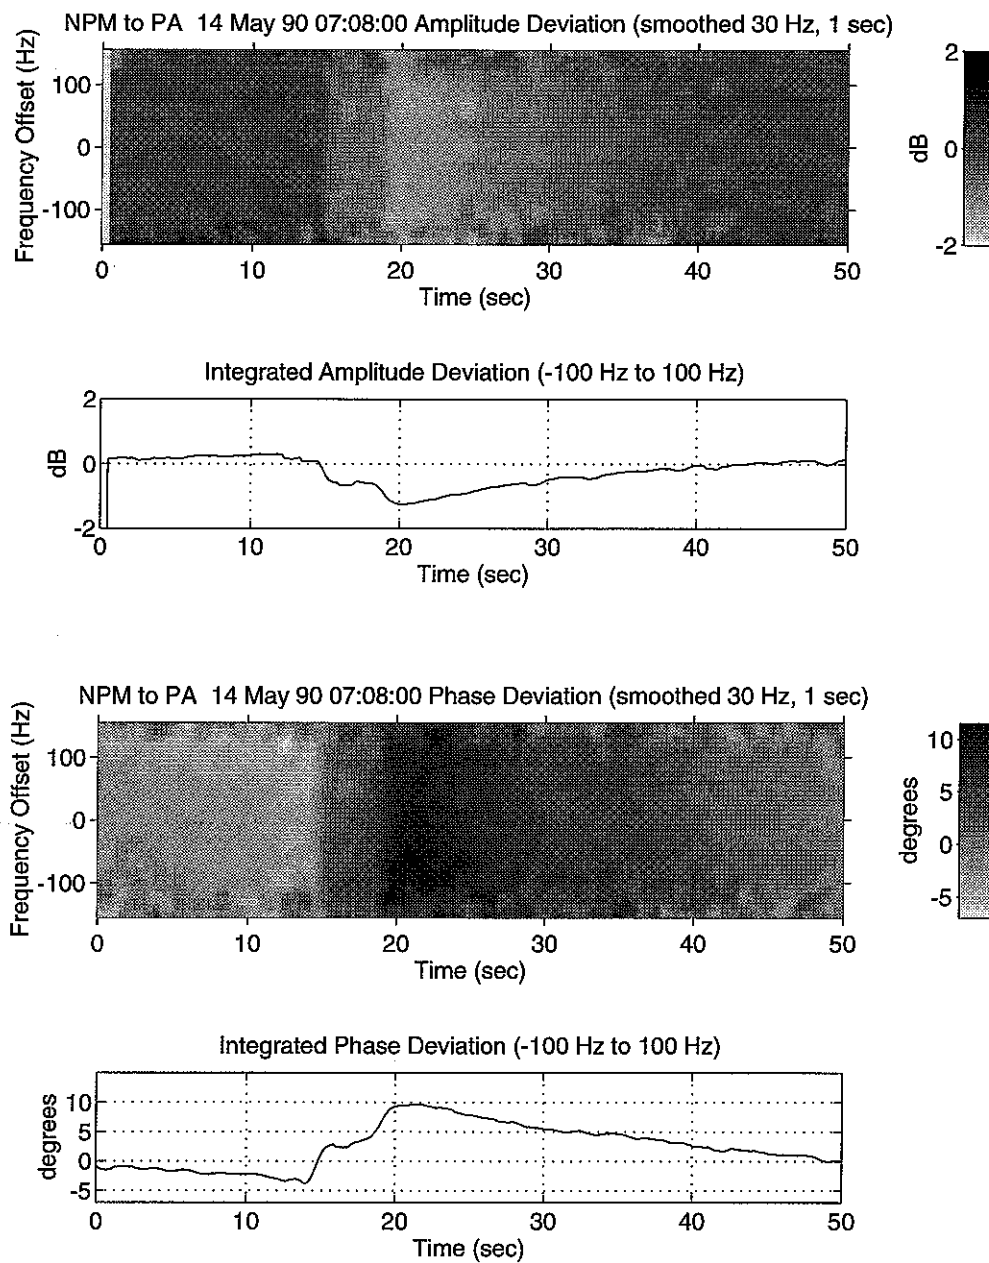


Figure 31. An expansion of two overlapping transfer function perturbations from Figure 30, with an expanded frequency scale showing only frequencies within ± 150 Hz of the carrier, where the input signal is strong enough to give good estimates of the transfer function. The first perturbation has a sharp onset at about 16 seconds, and is followed by a larger event at about 20 seconds. Both perturbations are amplitude decreases and phase advances which are smooth in frequency.

4.

SIMULATION OF PROCESSING SYSTEMS

Since many physical processes influence the propagation of VLF signals, and the net effect of these processes is difficult to predict, it is not always possible or cost-effective to assess the performance of a new measurement technique using only data obtained in the field. While the theoretical basis for this technique has been described above, it remains to be shown that it is capable of resolving geophysically significant propagation differences in the presence of ambient interfering signals and atmospheric noise. To demonstrate this, we turn to simulation, and generate a VLF signal at a frequency not utilized by any of the existing transmitters, alter the signal by passing it through filters which simulate the characteristics of a propagation path of interest, superpose the modified signal on to wideband VLF data recorded in the field, and analyze the result using the techniques described in Chapter 3. The objective of this simulation is to determine whether the characteristics of the propagation path (as artificially simulated by a filter) can be assessed via measurements of signal amplitude and phase as a function of frequency. The degree to which the results of the analysis agree with the filters used in simulating the propagation channel reflects the utility of the measurement technique on real, noisy, data.

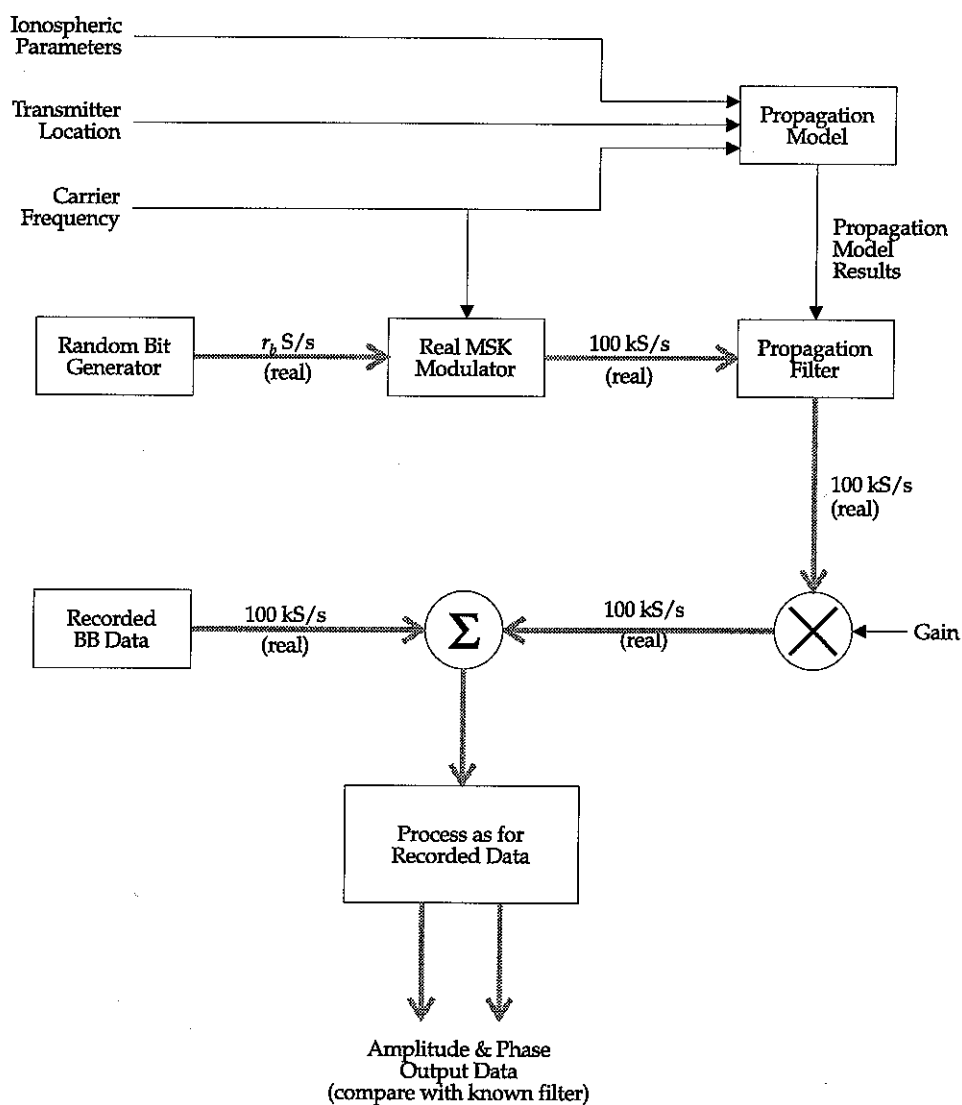


Figure 32. Processing flow for the simulation. By simulating a chosen propagation path with the 3-D waveguide model, and using the results to generate a channel transfer function filter, a signal can be generated which can be processed as described in Chapter 3 to yield propagation amplitude and phase information. These results can be compared with the channel transfer function filter to test the analysis method.

4.1. PROCESSING OF SIMULATED MSK SIGNAL

Figure 32 diagrams the simulation process. An assumed set of values of the ionospheric parameters chosen to match those used in recent work [Inan et. al., 1992],

combined with an arbitrary transmitter and receiver location and carrier frequency, are used in the 3-D waveguide propagation model [Poulsen *et. al.*, 1993a] to calculate the propagation channel characteristics at several closely-spaced frequencies. The model results are then used to design a time-varying filter which simulates both the phase and amplitude changes due to the static propagation path, and a time-varying perturbation, representing a localized ionospheric disturbance somewhere along the path. The filter was designed using the window method [Oppenheim & Schaefer, pp. 239–250], modified to accept an arbitrary amplitude and phase in the desired transfer function. A 4-sample Blackman-Harris window [Harris, 1978] was used. This filter may contain multiple stages: a first stage which is time-invariant and which accounts for the effects of the ambient earth-ionosphere waveguide; and a second stage, which is slowly time-varying and which simulates the effects of an LEP-induced ionospheric density change on the propagation of the VLF wave. The output of the filter is then mixed in with recorded data, with the gain adjusted to give a signal amplitude similar to those observed on VLF transmitter signals in an actual field-recorded data set. This data is then analyzed using the measurement methods described in Chapter 3, and the resulting time-varying propagation channel spectrum is compared with the characteristics of the propagation channel filter to determine the degree to which the measurement method can determine the characteristics of the filter, over time scales of geophysical interest, and in the presence of a realistic noise environment.

4.1.1. Construction of the Test MSK Signal

As discussed in Chapter 3, most of the high-power VLF transmitters currently in operation employ MSK modulation, since the bandwidth of these systems is limited, and the higher data rates obtainable with MSK, as opposed to FSK, modulation justify the slightly greater complexity in the modulator and receiver. All of the transmitters studied in this work utilize MSK, which is the modulation chosen for our simulated signal.

There are six variables which fully determine the MSK-modulated signal (See Chapter 3), at the input to the transmitter: the amplitude (A), the carrier frequency (ω_c), the carrier phase (ϕ_c), the bit rate (r_b), the time offset (T_b), and the bit sequence being transmitted (a_k). Of these, A is normalized to unity, since it is simply a multiplier which determines the signal level (the relative amplitudes of the antenna excitation factors for




Figure 33. Autocorrelation of a sample received bit sequence. The value at zero lag is normalized to unity, and is off the top of the graph. No simple repeating structures are seen, out to a lag of 5 seconds (1000 bits).

each of the waveguide modes are computed by the 3-D waveguide propagation model), and ω_c and r_b are known. The simulation code must assign values to the other three variables. For the performance tests described in this Chapter, ϕ_c and T_b are both set to zero, since the processing described in Chapter 3 cannot estimate these without absolute time and phase information, which is not present in the field data. The code which generates the simulated MSK waveform can set these variables to arbitrary values, however, and this capability was used in developing and testing the analysis code. Note that in real data acquired without absolute phase and time information, these quantities would similarly be unknown, and would be suppressed in the data analysis, as discussed in Chapter 3.

With the above assumptions, the MSK-modulated waveform is determined solely by the bit sequence being transmitted, a_k . Although the actual data being transmitted has structure, since it presumably is the output of some channel-coding process at the transmitter, an examination of the second-order statistics of actual received data (in this case, the NPM signal observed at Palmer Station) failed to reveal any significant departure from independence between the bits (a sample autocorrelation function is shown in

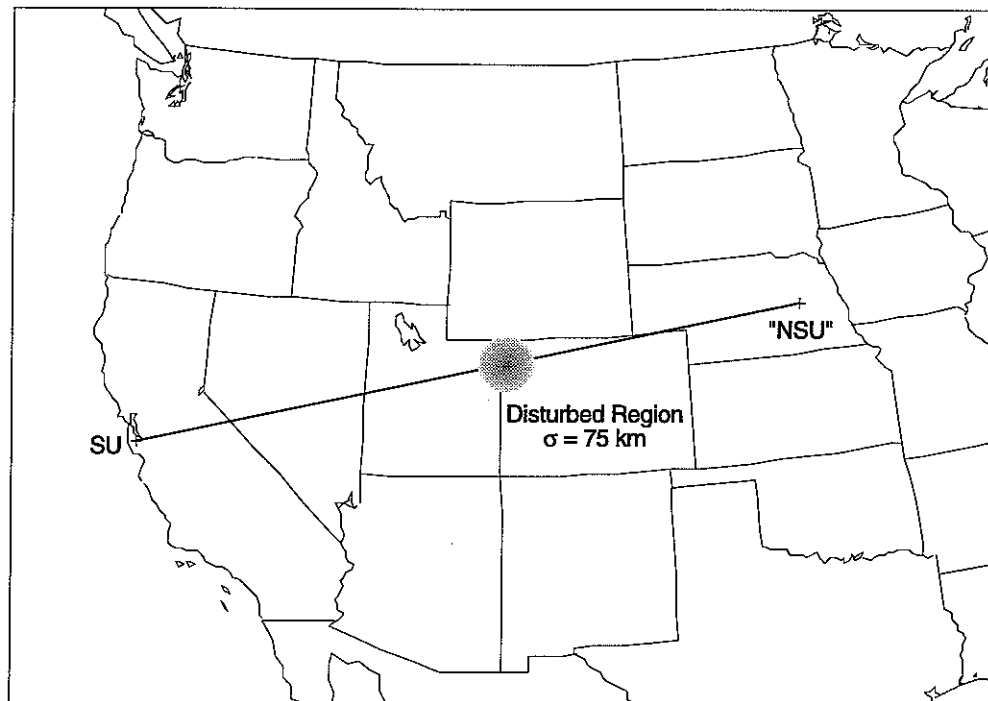


Figure 34. Location of simulated VLF transmitter ("NSU"), and the simulated ionospheric perturbation on the great-circle propagation path to a receiver at Stanford. The disturbance to the ionospheric density profile is assumed gaussian in shape, with $\sigma=75 \text{ KM}$.

Figure 33). Therefore, in the simulation, the data bits were simulated as independent and identically distributed (IID) random variables. For such a case, it is a simple matter to construct a MSK waveform which is representative of the signal at the input to the power amplifier stages of a real VLF transmitter. Once this test waveform has been constructed, the bit sequence used in the modulation, as well as the carrier phase and time origin, are not used in further processing, as they are generally not available to commonly used receivers in the field.

4.1.2. Simulation of the Slowly-Varying Propagation System

The novel feature of the propagation channel measurable by the technique described here is the variation of the amplitude and phase over the $\sim 300 \text{ Hz}$ frequency band illuminated by the signal. In order to test the method, therefore, we need to find a set of conditions which are not only of geophysical interest, but which can be expected to lead to variations in the signal amplitude and phase over the 300 Hz bandwidth occupied by a typical VLF




Figure 35. Electron density profiles used in the simulation. The ambient profile is "Profile II" from Figure 6, while the perturbation is related to a whistler described in the text.

msk-modulated signal, as a result of propagation through the channel. As discussed in Chapter 2, these variations are expected to be most significant when the received signal consists of many waveguide modes with similar amplitudes, which occurs for the case of short, land-based propagation paths.

The path chosen for our simulation is illustrated in Figure 34, from a hypothetical transmitter at the location of the commonly-observed 48.5 KHZ Air Force transmitter in Silver Creek, Nebraska ($41^{\circ} 30' \text{ N}$, $97^{\circ} 36' \text{ W}$), but operating at 26.0 KHZ, to a receiver located at Stanford ($37^{\circ} 25' \text{ N}$, $122^{\circ} 10' \text{ W}$). The 3-D waveguide propagation model was utilized, for frequencies from 25.80 KHZ to 26.20 KHZ, in 50 HZ steps, and for an ambient nighttime D-region electron density profile in effect along the entire path, except under the disturbance, where the altitude profile of electron density was taken to be one representative of conditions following an LEP event. The ambient profile used corresponds to Profile II of Figure 6, while the disturbed profile is that which would be expected for a lightning-induced electron precipitation burst of peak flux




Figure 36. Mode structure of the received signal at Stanford, from the fictitious 26.0 KHZ transmitter in Nebraska, as computed by the 3-D waveguide propagation model. Many modes contribute significantly to the total field at the receiver, so small perturbations to individual modes may affect the total field strength disproportionately.

$2 \times 10^{-2} \text{ erg/cm}^2/\text{sec}$, with a duration of 0.2 seconds, as would be produced by a 0.2–6.0 KHZ whistler, propagating at $L=2.5$ [Inan *et. al.*, 1988b]. These two density profiles are illustrated in Figure 35.

When the 3-D waveguide propagation model is used for the ambient conditions, there are at least six propagating modes which contribute significantly to the received signal (*i.e.*, within 6 DB of the total). These modes are shown in a polar display in Figure 36. From this figure, it is apparent that the total predicted field at the receiver is less than the amplitude of the single largest mode, due to partial cancellation by other, slightly weaker modes. Poulsen [1993b] has shown that the strength of the field scattered by a typical LEP disturbance depends strongly on the waveguide mode; in particular, the quasi-transverse magnetic (QTM) modes are scattered more effectively than quasi-transverse electric (QTE) modes, and the higher-order modes are scattered more effectively than the lower-order modes. When the total received signal consists of multiple significant modes, the change in the total field at the receiver is very sensitive to the phase changes to the different modes. Since the scattering may also be a function of




Figure 37. Amplitude and phase of transfer function for propagation at 26 KHZ from "NSU" to Stanford, under the ambient conditions. The predicted amplitude varies by more than 3 DB over the 300 HZ band of frequency; the predicted phase varies by nearly 15° over the same interval.

frequency, small differences in scattered fields at different frequencies making up a single modulated signal may combine to yield a perturbation which is a strong function of frequency.

The extent of these differential amplitude and phase perturbations may be predicted by using the 3-D waveguide propagation model for the above conditions, for a closely-spaced set of frequencies covering the range over which the simulated MSK signal has significant energy. For this simulation, model calculations were carried out for nine frequencies spaced 50 HZ apart, ranging from 25.8 KHZ to 26.2 KHZ. The predicted variation of the received signal amplitude and phase, for the chosen ambient electron density, are shown in Figure 37. This result illustrates that, for this short, multi-mode path, the transfer function of the propagation channel exhibits significant variations over the frequency band of interest.

When the localized perturbation to the electron density is introduced, the amplitude and phase of the signal at the receiver changes, with the difference due to the perturbation also being a function of frequency. As discussed in Chapter 3, these




Figure 38. Amplitude of propagation from NSU to SU as a function of frequency, as predicted by the 3-D waveguide propagation model, for both ambient and perturbed conditions (open circles in the top panel). A 4001-point FIR filter is fit to each of these predicted propagation characteristics (solid lines in the top panel). The difference between the two characteristics and filters is shown in the bottom panel, and is the predicted amplitude of the LEP-induced perturbation to the electron density, as a function of frequency. The 3-D waveguide propagation model was run only out to ± 200 Hz; the FIR filter was designed to have constant amplitude beyond these points, since signal components beyond ± 200 Hz are heavily attenuated by the baseband filter during processing.

differential perturbations are more easily measured, since they are not sensitive to assumptions about the transmitter and antenna system. The detectability of differential changes in signal amplitude and phase, due to typical LEP-induced electron density changes, is of most interest in this Chapter. These propagation effects can be simulated by using a slowly time-varying filter, which normally has an impulse response appropriate to the ambient propagation prediction, but changes to an impulse response appropriate to the perturbed condition, over a time scale of several seconds, which is 10–100 times the length of the impulse response of the filter. These specific propagation channel transfer functions, and filters which represent the channel transfer function predictions, are shown in Figures 38 and 39.




Figure 39. Phase of propagation from NSU to SU as a function of frequency, as predicted by the 3-D waveguide propagation model, with the simulation filters (top panel). The difference in the bottom panel is the predicted phase of the LEP-induced phase perturbation. As with the amplitude response, the phase of the FIR filter was selected to be constant beyond the ± 200 Hz range covered by the 3-D waveguide propagation model runs.

To simulate the effects of an LEP-induced ionospheric disturbance, the generated MSK signal is passed through a time-varying filter, as indicated in Figure 32. This filter varies from the "ambient" filter to the "disturbed" filter in one filter length (40 msec), and then relaxes to the "ambient" filter with a first-order time constant of 10 seconds. Such a time variation approximates the typical onset and decay times seen in LEP-induced ionospheric disturbances generated by mid-latitude whistlers [Inan and Carpenter, 1987].

Finally, the output from this filter, which simulates the propagation over the short, land-based path indicated in Figure 34, is combined with VLF data recorded in the field. The amplitude of the synthetic signal is set to a value at the low end of the range observed for real VLF transmitters in the Antarctic, which results in a spectrum similar to that shown in Figure 40. This composite signal is then used as the input to the analysis procedure, without reference to the filters used to generate the perturbation.




Figure 40. Spectrum of the broadband VLF, with the synthetic NSU signal added at 26.0 kHz. The NSU signal is mixed in at a low enough amplitude to simulate a relatively weak VLF transmitter.

4.1.3. Use of Spread-Spectrum Processing

Using the simulated signal as input, the processing steps detailed in Chapter 3 were performed, with the result displayed as a spectrogram of the transfer function in Figure 41. The two perturbations, at 60 and 180 seconds, are clearly visible, and show some variation over frequency. The amplitude perturbation is greatest above the center frequency (peaks at about +100 Hz); the phase perturbation is positive below the center, and negative above it. Both observations agree with the characteristics of the filter used to simulate the propagation, indicating that frequency-dependent amplitude and phase variations imposed during the propagation of the signal can be recovered upon analysis.

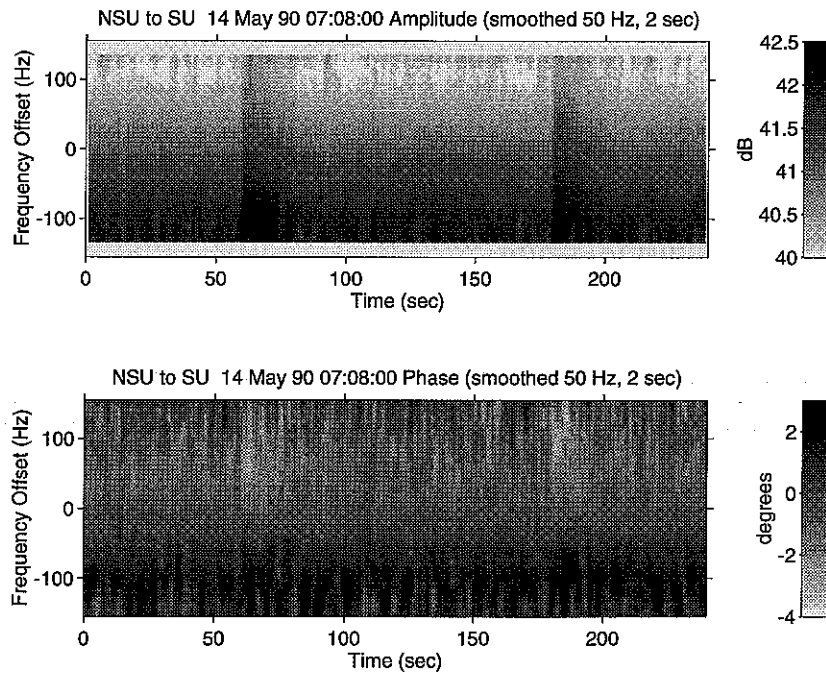


Figure 41. Spectrogram resulting from processing the simulated nsu signal. Two simulated perturbations were introduced, at 60 and 180 seconds from the start of the record. The nature of the perturbation is clearly visible in the amplitude record, and, to a lesser extent, in the phase as well. In addition, the asymmetry of the perturbation over frequency is visible, particularly in the phase.

4.2. PERFORMANCE OF EXISTING AMPLITUDE AND PHASE RECEIVERS

Most of the receivers used in previous studies in this field measure the integrated amplitude of the VLF signal. When phase receivers are used, they track the phase of a single coherent component of the modulated VLF signal using phase-locked loop (PLL) techniques [Wolf, 1990]. The operation of these receivers, and the limitations they impose on the data analysis, are described in this section.

4.2.1. Amplitude Receivers

Amplitude receivers which measure the integrated signal amplitude are easy to build, and also easy to analyze, since they are, ideally, linear receivers except for the envelope detector. The forms of modulation commonly used in the VLF band are constant-amplitude, and so the resulting measurement, whether on MSK, FSK, or CW signals, is a true measurement of the overall attenuation of the signal as a result of the propagation. Since the signal amplitude is a function of frequency, as are the propagation effects and the receiver filters, the detector in such a receiver ideally measures the cumulative result of three components: the transmitted signal, the propagation, and the receiver filter shape. Specifically, if $X(f)$ is the spectrum of the transmitted signal (for one particular bit sequence), $P(f)$ is the propagation effect to be measured, and $R(f)$ is the cumulative transfer function of all the filters used in the receiver, then the measurement output from an ideal amplitude detector is:

$$Y = \int_0^{\infty} X(f)P(f)R(f)df \quad (4.1)$$

With careful design, $R(f)$ may be made nearly flat over the bandwidth where $X(f)$ is significant. If we also assume that the effect of the propagation is also flat in amplitude and linear in phase over the occupied bandwidth, then:

$$\begin{aligned} y(t) &= \int_{-\infty}^{\infty} Y(f)e^{j2\pi ft}df \\ &= \int_{-\infty}^{\infty} X(f)A_P e^{-j(2\pi fT_P + \phi_P)} A_R e^{-j(2\pi fT_R + \phi_R)} e^{j2\pi ft}df \\ &= A_P e^{-j\phi_P} A_R e^{-j\phi_R} \int_{-\infty}^{\infty} X(f)e^{j2\pi f(t-T_P-T_R)}df \\ &= A_P e^{-j\phi_P} A_R e^{-j\phi_R} x(t) * \delta(t-T_P-T_R) \\ &= A_P e^{-j\phi_P} A_R e^{-j\phi_R} x(t-T_P-T_R) \end{aligned} \quad (4.2)$$

so the signal into the detector is simply a time-delayed replica of the input, multiplied by a complex amplitude factor. The detector will ideally strip all the phase information, and thus will simply measure the propagation as:

$$\begin{aligned}\hat{P}(t) &= |A_P e^{-j\phi_P} A_R e^{-j\phi_R} x(t - T_P - T_R)| \\ &= A_P A_R\end{aligned}\quad (4.3)$$

since the detector bandwidth is much smaller than the carrier frequency of $x(t)$, and therefore $|x(t)|$ is equal to unity. Although the above analysis is strictly valid only when the receiver and propagation effects are linear and time-independent, it can be applied to the common case in the VLF bands where variations on the propagation are much slower than the frequencies contained in the transmitted signal. Finally, since the receiver filter characteristics are (ideally) stable, they can be measured, and the only time-varying effect at the output is the propagation amplitude A_P .

If the amplitude and phase of the propagation varies significantly over the bandwidth occupied by the transmitted signal, then an integrated-amplitude receiver will measure the integral of the propagation constant, weighted by the spectrum of the transmitted signal:

$$\hat{P}(t) = A_R \left| \int_{f_c - B/2}^{f_c + B/2} A_P(f) X(f) df \right| \quad (4.4)$$

where B is the bandwidth occupied by $x(t)$.

Thus, this technique cannot resolve differences in propagation as a function of frequency, and only produces a weighted integral of the propagation amplitude. While this is sufficient for many useful measurements, it introduces a practical problem as well: if the receiver filters are not perfectly flat, the deviations from flatness will multiply the short-time transmitted signal spectrum. The output will exhibit variations which depend on the short-time spectrum of the transmitted signal, and thus on the message bits being sent. This message modulation gets worse as the integration time of the detector output is reduced, because the integral of the short-time spectrum of the product varies more as the time interval is reduced.

4.2.2. Phase Receivers

For CW and FSK-modulated signals, an ideal phase receiver can be described which is even simpler than the integrated-amplitude receiver, since CW and FSK-modulated signals both contain unmodulated signal components. A phase receiver for CW and FSK signals can

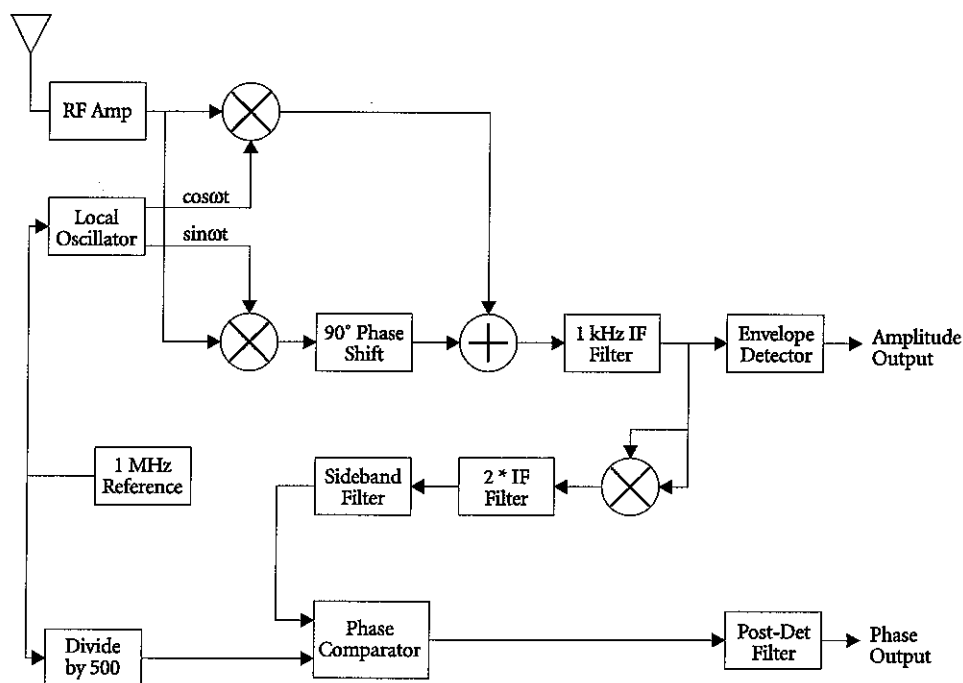


Figure 42. Typical block diagram of a frequency-doubling MSK phase receiver, with amplitude output, after Wolf [1990].

simply measure the phase of one or more of these unmodulated components (see Figure 14). In a practical FSK receiver, the finite bandwidth of the phase-locked loop (PLL) which tracks the phase of one of the unmodulated carrier components would admit some of the modulated signal as well, and therefore the output would exhibit some amount of message modulation. Further, there would be a tradeoff between measurement time and the variations due to message modulation; a shorter time constant in the PLL accepting more of the modulated signal, and resulting in greater message modulation in the output.

A related disadvantage of this type of phase receiver is that, even with FSK signals, measurements may be made at only two frequencies, which correspond to the unmodulated carrier components. Even though there is significant energy at other frequencies, it cannot be used to measure phase at those frequencies. The main advantage of this type of receiver is that it is relatively simple to build, and needs no information about the message being sent, other than the spacing of the mark/space frequencies from the center frequency.

For MSK signals, which have no unmodulated components, the coherent-sideband type of receiver can still be used, if the MSK signal is first put through a frequency-doubler circuit. This serves to construct an artificial FSK signal from the MSK input. In the MSK




Figure 43. Waveform of MSK signal and an FSK signal synthesized from it by a square-law frequency doubler. The center frequency of the MSK signal has been lowered to 400 Hz to exaggerate the effects, and offset in amplitude for clarity. The lower trace is the original MSK signal, and the upper the frequency-doubled result. The message bits being sent here are '01011'.

input, both carrier frequencies are spread out because the requirement for phase continuity forces each carrier to be present in two phases, 180° apart. Where FSK may be constructed from two coherent carriers, by switching between them at zero-crossings, the MSK, with double the bit rate, can only maintain phase continuity by biphase-modulating the carriers appropriately, and this modulation spreads the carriers out over a bandwidth comparable to the signaling rate. Figure 42 is a block diagram of such a receiver, currently used by Stanford at a number of field sites for continuous measurements [Wolf, 1990].

Figure 43 illustrates the frequency-doubling operation in the time domain. The original MSK signal contains no coherent carrier components, as can be seen by examination of the phases of the carrier for the '0' bits. The first and second '0' of the transmitted message are separated by an odd number of '1' bits, and so the carrier is switched in phase by 180° between the two bits. However, in the FSK signal synthesized from the MSK, the frequency-doubling operation has mapped both 0° and 180° phases of the carrier into the same, 0° , phase. With reference to Figure 16, this operation negates the




Figure 44. Spectra of msk signal and a frequency-doubled fsk signal generated from it. The carrier frequency of the msk signal is set to 1 kHz, and the spectrum computed for a random sequence of data bits (2000 total bits).

effect of the BPSK modulators on the “mark” and “space” frequency carriers, and the resulting modulation is FSK, not MSK. This restores the spectral peaks at the mark and space frequencies, which can be seen in Figure 44.

Receivers which use these techniques exhibit the same limitation described above for FSK signals: as the tracking loop bandwidth is made wider to obtain higher time resolution, more of the modulated message energy is passed to the output, and the amount of message modulation increases. In addition, there is a more fundamental problem involved with the frequency doubling: since it is a non-linear operation, it performs frequency conversion on the components of the transmitted MSK signal. The carrier components which appear in the FSK signal at the output of the frequency doubler were not transmitted, and, in the general case where the propagation is a sensitive function of frequency, the FSK carriers may be synthesized by the doubling operation from MSK components which were attenuated and phase-shifted by differing amounts in propagation. Thus, measurement of the amplitude and phase of the synthesized carriers does not, in general, give measurements of the amplitude and phase of the channel transfer function at the carrier frequencies. This deficiency may be reconciled with the

time-domain analysis by noting that large variations in the amplitude or phase of the channel transfer function over the frequency range occupied by the MSK signal imply that the impulse response of that transfer function is significant over a time comparable to the duration of one transmitted bit. Such a circumstance would be a case of large amounts of Inter-Symbol Interference (ISI), and the waveform during one bit interval will depend on those before it. As discussed in Section 4.1.2, cases where this is expected to occur include short, land-based propagation paths, where many waveguide modes are present at comparable amplitudes.

5.

APPLICATION TO DATA

While the simulation results described above in Chapter 4 demonstrate that the frequency-dependent analysis techniques used in the present work are suitable for accurate measurements of the transfer function of the VLF radio propagation channel, it is the application of these techniques to field-recorded data which is of scientific interest. In this chapter, we describe a data set digitally recorded for this purpose at Palmer Station, Antarctica, and give examples of the application of the new analysis techniques to these data.

5.1. SIGNAL PATHS STUDIED

For the purposes of this study, four signal paths were studied, all terminating at Palmer Station, Antarctica (64°W , 65°S). Table 5.1 shows the parameters for the signals observed. All of the signals measured arrive at Palmer over long paths, two of them (NPM, NLK) almost entirely over seawater, from transmitters in the northern hemisphere. All the listed signals are modulated with 200 bits/sec MSK, and are strong enough at Palmer to allow accurate decoding of the bit sequence. The characteristics of these signals as observed at Palmer have been studied extensively in previous work [*Carpenter and LaBelle*, 1982; *Tolstoy et al.*, 1986; *Inan and Carpenter*, 1986; *Wolf and Inan*, 1990].




Figure 45. Great-circle propagation paths for the signals studied in this section. All originate in the northern hemisphere, and terminate at Palmer Station, Antarctica.

Table 1. Transmitter and Arrival Characteristics

Call Sign	Frequency (KHZ)	Power (KW)	Latitude (degrees)	Longitude (degrees)	Mag Bearing (degrees)	Path Length (km)
NPM	23.4	502	21° 25' N	158° 09' W	275° 48'	12,349
NSS	21.4	213	38° 59' N	76° 27' W	350° 00'	11,606
NAA	24.0	1,200	44° 39' N	67° 17' W	357° 30'	12,187
NLK	24.8	245	48° 12' N	121° 55' W	318° 24'	13,526

The amplitudes and relative phases of the significant waveguide modes that constitute the received signals at Palmer Station, as predicted by the 3-D waveguide propagation model, are shown in Figure 46. For all four signals, the received mode structure is dominated by a single mode, although the dominant mode is different for the different paths. According to *Poulsen* [1991, p. 46], under typical conditions, the lowest attenuation rates are obtained for the low-order QTE modes (QTE₁ AND QTE₂), and for the lowest-order QTM mode (QTM₁); however, the QTM₂ mode is most efficiently excited by the vertical monopole transmitting antennas, and can remain dominant out to quite large distances, in spite of its somewhat higher attenuation rate.

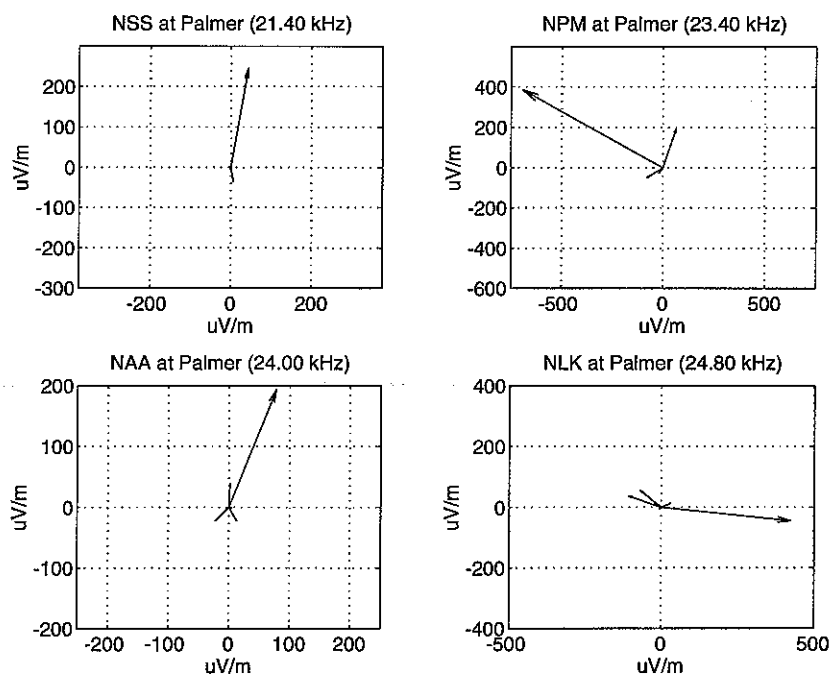


Figure 46. Mode structure of the arriving VLF signals at Palmer Station, as predicted by the 3-D waveguide propagation model, for the center frequency of each transmitter. In all four cases, a single mode is dominant (> 10 dB above the other mode strengths).

With a single dominant mode, the propagation characteristics of the VLF wave should be a relatively smooth function of frequency, since the ionospheric and ground conductivities vary smoothly with frequency so that rapid changes with frequency are most likely caused by the interference between modes of nearly equal amplitude. Figures 47–50 show the amplitude and phase at the receiver as a function of frequency in a ± 200 Hz range around the center frequency, as predicted by the 3-D waveguide propagation model for the indicated paths. The ambient electron density profile used was Profile II from Figure 6 for all these calculations. In comparison with Figure 37, which shows the predicted propagation channel transfer function for a short, land-based, path with several significant modes, all the transfer functions in Figures 47–50 vary smoothly with frequency, with maximum amplitude variations across the ± 200 Hz bands shown of about 1 dB, and phase variations of about 0.3° . In addition, there are no sharp features in the spectrum; sharp features would only be expected in cases where several significant waveguide modes coexisted, in which case destructive interference can occur at some frequencies.




Figure 47. Predicted propagation from NPM in Hawaii to Palmer Station. The predicted phase has been adjusted to remove constant phase- and group-delay terms, as described in Chapter 3, since these cannot be measured from the field data without absolute timing information at both the transmitter and receiver.




Figure 48. Predicted propagation from NSS in Maryland to Palmer Station. The predicted phase has been adjusted to remove constant phase- and group-delay terms, as described in Chapter 3, since these cannot be measured from the field data without absolute timing information at both the transmitter and receiver.




Figure 49. Predicted propagation from NAA in Maine to Palmer Station. The predicted phase has been adjusted to remove constant phase- and group-delay terms, as described in Chapter 3, since these cannot be measured from the field data without absolute timing information at both the transmitter and receiver.




Figure 50. Predicted propagation from NLK in Washington to Palmer Station. The predicted phase has been adjusted to remove constant phase- and group-delay terms, as described in Chapter 3, since these cannot be measured from the field data without absolute timing information at both the transmitter and receiver.

Although the above results illustrate signal characteristics under ambient (unperturbed) ionospheric conditions, the perturbation to the signal due to the appearance of localized disturbances along its propagation path (*i.e.*, due to lightning-induced electron precipitation) is also expected to be a smooth function of frequency, when a single, low-order, mode is dominant. For a single waveguide mode, the phase perturbation ($\Delta\phi$) induced by a change in ionospheric reflecting height (Δh) is given by [Inan and Carpenter, 1987]:

$$\frac{\Delta\phi}{d\Delta h} \equiv -\frac{2\pi f}{hc} \left[\frac{h}{2R_e} + C_n^2 \right] \quad (5.1)$$

where:

$$C_n = \frac{(2n-1)\lambda}{4h}$$

n	is the order of the mode,
d	is the length of the path under the perturbation,
λ	is the wavelength,
R_e	is the radius of the Earth,
f	is the wave frequency, and
c	is the speed of light.

For low-order modes in the 20–30 kHz range, and a typical reflecting height of 90 km, the first term in the brackets dominates the C_n^2 term, and the phase perturbation is simply proportional to frequency, representing essentially a group delay. In such a model, the amplitude perturbation is simply due to increased absorption at the lower altitudes, and is varies insignificantly with frequency over the range occupied by these signals.

5.2. RECEIVING, RECORDING AND PROCESSING OF DATA

The receiving equipment for this experiment consisted of a set of 1GX-style crossed-loop antennas, which are single loops of #6 cable, supported in an isosceles triangle 9 meters tall and 18 meters across at the base [Paschal, 1978], a broadband VLF preamp and line receiver, and a digital data recording system, as is illustrated in Figure 51. All parts of this system, except for the digital recorder, have been in use since 1988 at Palmer station for a variety of ionospheric and magnetospheric studies [Wolf, 1990]. In addition to the digital

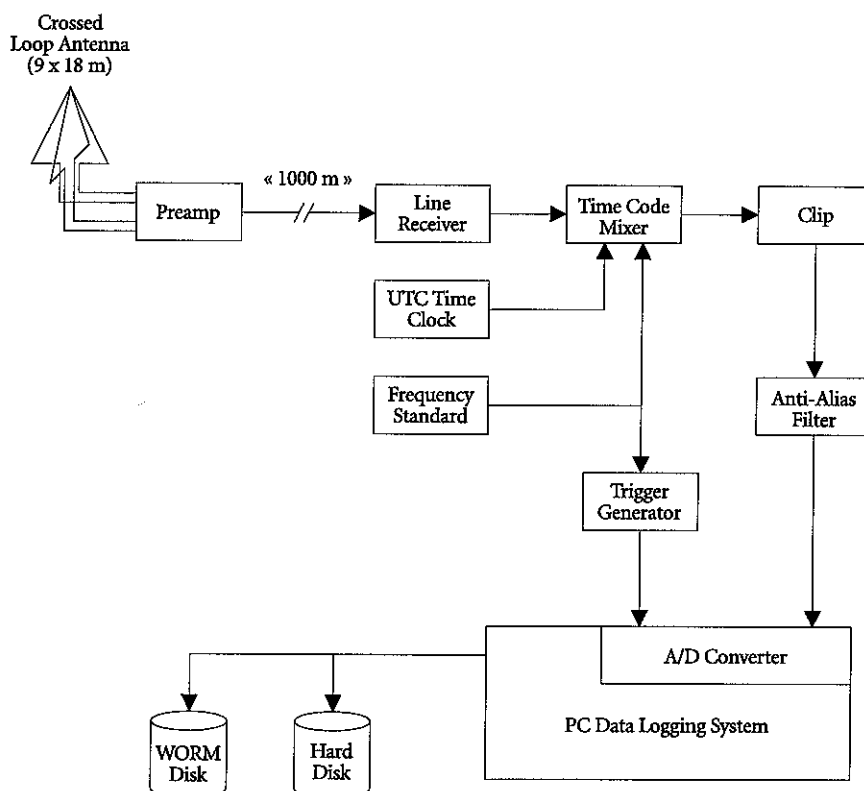


Figure 51. Diagram of the receiving and recording equipment used to acquire the data analyzed in this chapter.

wideband recordings made for this experiment, the existing narrowband data system was used to monitor activity on the signal paths, in order to determine when to record broadband data. In this way, many examples of lightning-associated VLF signal perturbations were obtained, despite the limited recording capacity of the digital wideband recorder (5 minutes per event, and about 5 hours total).

To reduce interference from the generators and other equipment at Palmer Station, the loop antennas are situated on a glacier in back of the station, about 1 km from the main buildings. The signals received by the loops were amplified by the preamps, and driven differentially on lines back to the recording hut. There, a line receiver unit supplies gain, filtering, and buffering, and allows the injection of calibration signals. The conditioned analog signal is then passed through an anti-alias filter, digitized at a rate of 100 kHz, and stored on optical disk.

5.2.1. Antenna Efficiency

Since the antennas used are electrically small loops, compared to the wavelengths of signals in this frequency range, they are sensitive to the B-field of the incoming waves, and are most sensitive to plane waves arriving in the plane of the antenna loop; thus, the N/s loop is most sensitive to waves arriving from the North or South, and the E/w loop to waves from the East or West. The antennas are aligned to within $\pm 5^\circ$ of geomagnetic N/s and E/w.

Most of the data used in this work were recorded from the N/s loop antenna, since it gives the largest signals from the North American transmitters NSS, NAA, and NLK. Although the N/s antenna orientation nearly nulls the reception of NPM, this is partly compensated by the fact that the NPM signal to Palmer propagates almost entirely over water, with relatively little attenuation. The great-circle propagation paths for the other frequencies are over land for significant parts of their length, and the signals arriving at Palmer are attenuated. The net result of these effects is that the amplitude of the listed signals in the wideband spectrum observed by the N/s antenna are close to equal, as can be seen in Figure 40.

The ultimate limit to the performance of a VLF receiver using a loop antenna is set by the resistance of the wire used for the antenna, and its relationship to the area enclosed by the loop [Paschal, 1988]. The wire resistance contributes thermal noise to the receiver input, and, even with a noiseless receiver, limits the weakest detectable signals. This noise voltage is given by:

$$E_a = \sqrt{4kTR_a} \quad (5.2)$$

where:

E_a	is the noise voltage density, in $V \cdot Hz^{-1/2}$,
k	is Boltzmann's constant,
T	is the temperature of the wire, and
R_a	is the resistance of the wire.

Since an incoming plane wave, arriving from a direction in the plane of the loop, induces a voltage on the antenna of:

$$V_a = j\omega N A E_w / c \quad (5.3)$$

where:

N	is the number of turns of wire in the antenna,
A	is the area of the antenna bundle,
E_w	is the electric field strength of the wave, and
c	is the speed of light,

the noise introduced by the resistance of the antenna wire is equivalent to an arriving wave of magnitude:

$$S_a = \frac{c\sqrt{4kTR_a}}{\omega NA} \quad (5.4)$$

where:

S_a	is the antenna sensitivity, in $V \cdot m^{-1} \cdot Hz^{-1/2}$.
-------	---

This sensitivity is a spot noise density [Motchenbacher and Connelly, 1993, p. 44], and is a function of frequency. Over the bandwidths of typical VLF signals, however, the sensitivity only varies by ~1 percent, so that we can use the carrier frequency in the above expression, to calculate the minimum detectable signal from the spot noise density and the post-detection bandwidth, B , occupied by the signal:

$$\begin{aligned} E_{min} &\approx S_a \sqrt{B} \\ &= \frac{c\sqrt{4kTR_a B}}{2\pi f_c NA} \end{aligned} \quad (5.5)$$

where E_{min} is the incoming signal strength, in $V \cdot m^{-1}$, which would give 0 dB signal-to-noise ratio at the output of a noiseless receiver connected to the antenna in question. For the type of loop antennas used at Palmer, Paschal [1978], estimates the antenna noise component at about 0.016 $\mu V/m$, over a bandwidth of 500 Hz, at a center frequency of 25 kHz. This is a theoretical limit to the detectable signal with the Palmer antenna, and does not account for the noise in the receiver, nor for atmospheric sources of noise. However, the sensitivity of the IGY-type antenna is typically about 20 dB better than the atmospheric noise levels normally present in the 20–30 kHz frequency range [Fraser-Smith and Helliwell, 1985].

A loop antenna responds to the rate of change of the magnetic field passing through the enclosed loop area, so that the voltage induced on the loop increases linearly with frequency. However, the loop impedance contains an inductive component, and, above the frequency where the loop's inductive reactance equals the input impedance of

the loop and receiver input stage, the increasing antenna reactance offsets the increasing induced voltage, and the current flowing into the receiver input becomes independent of frequency. The input turnover frequency is given by [Paschal, 1988]:

$$f_i = \frac{R_a + R_{in} / m^2}{2\pi L_a} \quad (5.6)$$

where:

R_{in}	is the input resistance of the preamplifier,
m	is the turns ratio of the input transformer, and
L_a	is the inductance of the loop.

The frequency response of the system is limited on the low end by f_i , which for the Palmer system is about 300 Hz¹.

At the higher-frequency end, the magnetic loop antenna performance begins to deviate from that described above when the size of the antenna is a significant fraction of a wavelength of the incoming wave, so that, at any instant of time, the magnetic field strength varies across the loop. However, since this occurs at frequencies outside the band of interest by several orders of magnitude, the high-frequency response is limited only by circuit and system considerations discussed below.

5.2.2. Preamp Frequency Response and Noise Figure

Since the antenna is located at some distance from the recording equipment, its output is fed into a preamplifier, and then driven on a balanced differential analog transmission line, down to the recording hut about 1000 meters from the antenna site. The combination of the antenna and preamp determine most of the characteristics of the recording system, including the low-frequency response, flatness of gain, phase delay, and noise level. The only major characteristic controlled by the recording system itself is the high-frequency response, which is sharply rolled off by the anti-alias filter.

¹ The antenna/preamplifier system used at Palmer has the following parameters: $R_{in} = 263 \Omega$, $m = 64$, $L_a = 65 \mu H$, and $R_a = 0.0615 \Omega$.




Figure 52. Amplitude response of the Palmer antenna and preamp. The system is flat to within 1 dB from 1–100 KHZ, as shown in the top panel. The bottom panel is an expansion of the response over the range of interest in this work, showing that the response is flat to within better than 0.01 dB over the typical 500 Hz bandwidth occupied by an MSK-modulated VLF transmitter.

The frequency response of the Palmer antenna and preamp are shown in Figures 52 and 53, and from those figures it is apparent that the gain and phase errors are insignificant over the 20–30 KHZ frequency range occupied by the MSK-modulated VLF transmitters studied in this work.

The system noise is dominated by the noise in the front end of the receiver, which was designed to provide a noise figure of about 4 dB at the lower frequencies, where the antenna sensitivity is lower. At higher frequencies, the noise factor increases with the square of the frequency, which offsets the linearly increasing antenna sensitivity to produce a broad minimum in system sensitivity. The signals studied in this work are in the 20–30 KHZ range, for which the combination of antenna and preamp has a noise density equivalent to an incoming plane wave of about $0.02 \mu\text{V} \cdot \text{m}^{-1} \cdot \text{Hz}^{-1/2}$. Over the entire 30 KHZ bandwidth of the recording system, this translates into a broadband system noise of about $3.5 \mu\text{V/M}$, RMS, which is much lower than the noise from atmospheric sources [Fraser-Smith and Helliwell, 1985].




Figure 53. Phase response of the Palmer antenna and preamp. In the bottom panel, covering the 10–30 KHZ frequency range, a constant phase and group delay have been removed, to show the phase error. Over any 500 HZ range, the phase error is much less than 0.1° .

After transmission to the recording hut, the balanced signal from the line is converted to a single-ended signal, filtered, and amplified to useful levels by the receiver. The receiver consists of two components: the line receiver, which performs most of the processing, and a time code mixer, which mixes in a 10 KHZ modulated time- and phase-reference tone, without altering the signal level. As all this processing is performed with signals which have already been amplified well above the noise floor of the processing equipment, the effect of noise in the later amplification stages on the gain, phase, and noise are minimal. Since the calibration data were taken with this equipment, this noise is included in the above figures, but is insignificant in comparison with the effects of the preamp and antenna.




Figure 54. Sensitivity of the Palmer antenna and preamp, expressed as the electric field of an incoming wave which gives an output equal to the system noise density. The noise figure of the preamp is the difference between the "System" and "Antenna" curves, and approaches 4 dB at frequencies below a few kHz. Above 10 kHz, the preamp noise figure increases with the square of the frequency, and the system becomes less sensitive.

5.2.3. A/D Converter and Storage System

The wideband VLF data are recorded in fully digital form, using commercial equipment and custom software for the filtering, conversion and storage of the broadband data. The output of the receiver first passes through a "brickwall" low-pass filter, with a cutoff of over 100 dB per octave beyond 33 kHz, to remove signal components which would cause aliasing. The data are then sampled at 100 kHz, using a 16-bit A/D converter encapsulated in a metal enclosure to reduce the electromagnetic interference from the host PC, so that true 16-bit performance can be achieved.

Converted data are stored first on the hard disk, in a circular buffer capable of retaining up to five minutes of data. The sampling system can run continuously, saving the most recent data, which allows the operator to trigger saving of the data when an

interesting event is observed. While the system is recording, there is no analysis possible of the incoming data. However, other instruments available at Palmer give good indications of activity, including on-line displays of VLF signal perturbations on the paths studied in this work. These available displays were used to control the capture of data.

Once five minutes of data are stored on the hard disk, they must be copied to the optical (WORM) disk before another recording may be made, due to storage limitations on the hard disk. This process takes nearly an hour, after which the system is again ready to record. For post-processing, the data from the WORM disk are loaded onto the hard disk of the analysis computer, using the algorithms described in Chapter 3.

5.2.4. Noise Calibration Measurements and Dynamic Range

The recording system was checked for noise performance by replacing the antenna with a dummy loop, having the same impedance, but no collecting area, effectively removing the incoming VLF signal from the output. What is left at the output is largely preamp noise, the time mark, and its harmonics. By injecting a low-level signal into the calibration input, we can measure the minimum detectable signal in the presence of the preamp noise.

A spectrum recorded during such a test is shown in Figure 55, which shows that a test tone at 5.5 KHZ, 95 DB below the clipping level, is easily detectable. Also visible are the harmonics of the time mark, and the shape of the anti-alias filter cutoff at 33 KHZ. From Figure 55, it can be seen that the preamp noise is about 10 DB above the noise floor set by the A/D converter, so that the recording system does not limit the detection of low-level signals. The preamp has a very large dynamic range, about 120 DB in a 1 HZ bandwidth [Paschal, 1988], larger than the recording range of the A/D converter. Therefore, the recording level is set so that the sum of all the coherent VLF transmitters in the 0-33 KHZ range is about 10 DB below clipping, so that large-amplitude signals, such as intense spherics, are clipped, and most of the dynamic range is available for recording the signals of interest.




Figure 55. Data recording system noise and signal detectability. The lower line shows the noise floor of the A/D and processing system only, while the upper line shows a recording through the entire system, with the time mark at 10 kHz, and a calibration signal injected at 5.5 kHz, 95 dB below clipping. The calibration signal is easily detectable, even at this low level. Harmonics of the time mark can be seen, as can the rolloff of the anti-alias filter above 33 kHz. The vertical scale is calibrated so that a sine wave input which just reaches the clipping levels is 0 dB.

5.3. PROCESSING SYSTEM

The analysis of the data recorded at Palmer was conducted at Stanford using an IBM PC-compatible computer equipped with a board containing a WE DSP-32C chip and associated memory. Since the data were already recorded digitally, this board contained no signal conversion circuitry, but accelerated the processing of the data by several orders of magnitude. The software to implement the processing algorithms was written in a combination of C and DSP-32C assembler, with the C code on the host PC performing data I/O and high-level tasks, and the DSP-32C code optimized for running the large filters and transforms involved. With minimal attempts at optimization, the software runs in nearly

real time on this hardware, with most of the processing work involved with the downconversion from broadband data to the complex baseband signal.

The results of the processing were displayed with custom software, either on-screen, or printed on charts for comparison with existing data, or converted into Matlab for further processing or image display.

5.4. MSK SIGNALS ON EXPECTED SINGLE-MODE PATHS

As discussed in Section 5.1., the signals recorded and analyzed in this work are all predominately single-mode, and are expected to yield propagation channel transfer function estimates which are smooth functions of frequency. In addition, the single-mode theory predicts that, when LEP-associated perturbations to this transfer function are seen, the perturbation should also be constant across the band of frequencies illuminated by the signal. We now compare representative data from the field with the predictions, both for the time-average propagation channel transfer function, and for the LEP-induced perturbations to this transfer function.

5.4.1. Time-Averaged Propagation Channel Transfer Functions

Following the procedure described in Chapter 3, the time-averaged propagation channel transfer functions were computed for selected data runs; a typical set of results are presented in Figures 56–59. When these measured transfer functions are compared directly with the predictions in Figures 47–50, the following differences are seen:

- 1) The signal received from NPM on the N/s antenna is higher than the prediction by about 7 DB, and has more variation with frequency.
- 2) The NSS measurement has an amplitude slope of about +1 DB over 200 HZ, while the predicted slope is about -1 DB over the same range.
- 3) The NAA measurement shows no amplitude slope, and is also about 5 DB lower than the prediction.
- 4) The NLK measurement is about 5 DB lower than the prediction.

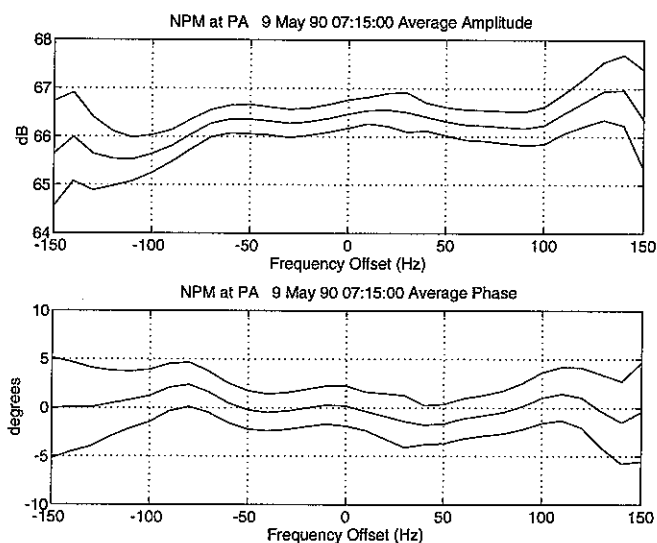


Figure 56. Average propagation channel transfer function and 95% confidence limits for the NPM-to-Palmer path, for a 5-minute data set collected on the N/s loop antenna. This antenna has a null nearly in the direction of NPM; the amplitudes have been corrected for the expected antenna response.




Figure 57. Average propagation channel transfer function and 95% confidence limits for the NSS-to-Palmer path, for a 5-minute data set collected on the N/s loop antenna.




Figure 58. Average propagation channel transfer function and 95% confidence limits for the NAA-to-Palmer path, for a 5-minute data set collected on the N/s loop antenna.




Figure 59. Average propagation channel transfer function and 95% confidence limits for the NLK-to-Palmer path, for a 5-minute data set collected on the N/s loop antenna.

There are many uncertainties which affect the predictions, and, to a lesser extent, the measurements:

- 1) The electron density and collision frequency altitude profiles were assumed constant ("Profile II" from Figure 6) everywhere along the propagation path; in reality, the profiles are not known, and can be expected to vary along the path. For the predominantly single-mode signals studied here, the main effect of increasing the electron density would be to lower the reflection height, and the increased collision frequency at the lower altitudes would be expected to increase the absorption, and reduce the overall amplitude.
- 2) Although the NSS, NAA, and NLK paths are well within the nighttime hemisphere at the time these data were taken, NPM was near the terminator (2043 local time). Since the ionospheric electron density is very different between dark and illuminated portions of the ionosphere, the different propagation characteristics might cause refraction of some of the signal energy, resulting in an enhancement at the receiving site.
- 3) The radiated power from the transmitting antenna is known only approximately [Rodriguez, 1994]. This should only affect the comparison of the overall signal levels.
- 4) The transmitting and receiving antennas are assumed to be electrically small dipoles by the 3-D waveguide propagation model, when computing modal excitation coefficients. This mirrors the actual construction of typical transmitting antennas [Watt, 1967, p. 253], but the receiving antennas are actually small loops (magnetic dipoles). Thus, the comparison of the measurement to the prediction depends on both the ratio of the E- and H-field components of the incoming wave, and the orientation of the receiving antenna.

Of these factors, (2) is an important subset of (1), since the effect of the day/night terminator on VLF radio propagation is due to the very different electron density profiles between daytime and nighttime. The third uncertainty is of no consequence except when comparing absolute amplitudes, and may be measurable during the daytime, when propagation conditions are more stable. Therefore, if we can account for the effects of the last item in the above list, the processing described in this work becomes a tool for the

determination of parameters of geophysical interest: the electron density and collision frequency profiles along the signal propagation path.

Figure 60 illustrates the importance of knowing the modal structure of the incoming signal, particularly when that signal is being received on an antenna which is nearly nulled to the direction of arrival. In the ideal case, where the incoming wave is a




Figure 60. Effect of mode structure of incoming wave on received signal, for a small loop antenna. The small loop is sensitive to the time-varying component of the magnetic field which is perpendicular to the plane of the antenna. In the top panel, the wave has no magnetic field component in the direction of propagation, and the familiar cosine antenna pattern results. In the middle, a mode which has a significant magnetic field in the direction of propagation will produce an additional output. In the bottom panel, when such a wave is received by an antenna which is nulled (perpendicular to the direction of arrival of the signal), the output is due solely to the non-transverse component of the magnetic field.

pure plane wave, or is pure transverse magnetic (TM), the relation of the H_y component to the E_z component is simply the wave impedance, and there is no H_x component. In such a case, the measurements may be compared to the predictions of the 3-D waveguide propagation model simply by allowing for the $\cos\theta$ response of the antenna.

If the incoming wave is not TM, then we must consider the effect of the H_x component on the recorded signal. The non-zero H_x component rotates the total magnetic field vector toward the axis of propagation, and it is the angle between the total H vector and the antenna normal which must be used when calculating the antenna response. If the mode in question has non-zero H_x , then a significant error may be induced by using the simple $\cos\theta$ response, when the antenna is nearly nulled, to infer the total H or total E for comparison with the 3-D waveguide propagation model results. In addition, the ratio of H_x to E_z varies with the waveguide mode number, and so this effect may alter the predicted mode structure as seen at the receiver.

In the data recorded for this work, the only signal which is nearly nulled on the N/S loop antenna is NPM. To evaluate the importance of this effect, data recorded on the E/W antenna were also analyzed, since for this antenna, the direction of arrival of the signal from NPM is nearly in the antenna plane, rather than along a null. The time-averaged propagation channel transfer function from these data is shown in Figure 61. Unfortunately, the recording equipment used at Palmer does not have the capacity to record from both the N/S and E/W antennas simultaneously, so the data in Figure 61 is from a different day. Nevertheless, the variation with frequency is smoother (compared to Figure 56), and the amplitude agrees well with the prediction of the model.




Figure 61. Average transfer function of the propagation channel from NPM to Palmer, as observed on the E/W antenna, which has a maximum response nearly in the direction of the great circle path to NPM.

5.4.2. LEP-Induced Perturbations to Transfer Functions

In the data recorded at Palmer, there are a number of perturbations to the time-averaged propagation channel transfer function which display the sudden onset and slow recovery typical of the LEP-induced perturbations described in Chapter 2. The single-mode theory discussed in Section 5.1 predicts that these disturbances should result in changes to the propagation channel which are only a slowly-varying function of frequency.

Data showing 10 perturbations on three signal paths (some simultaneous on two of the paths) are presented in Figures 62–67. Within the resolution of these displays, the perturbations all are uniform across frequency, as would be expected for a single dominant waveguide mode.

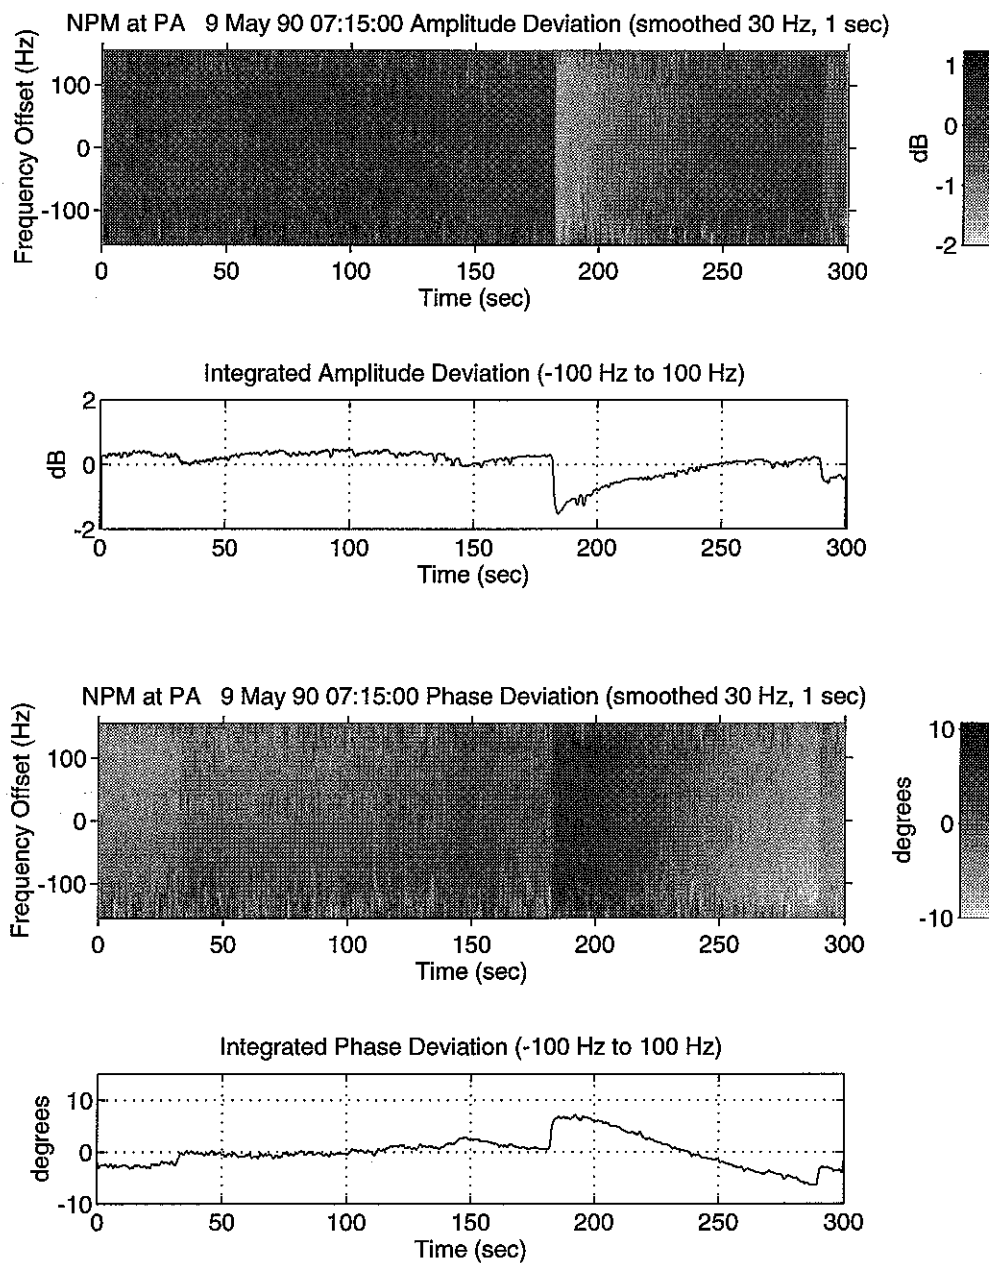


Figure 62. Two transfer function perturbations on the NPM-to-Palmer path, one at ~180 seconds and one at ~290 seconds. The data have been smoothed, and the static portion of the transfer function suppressed.

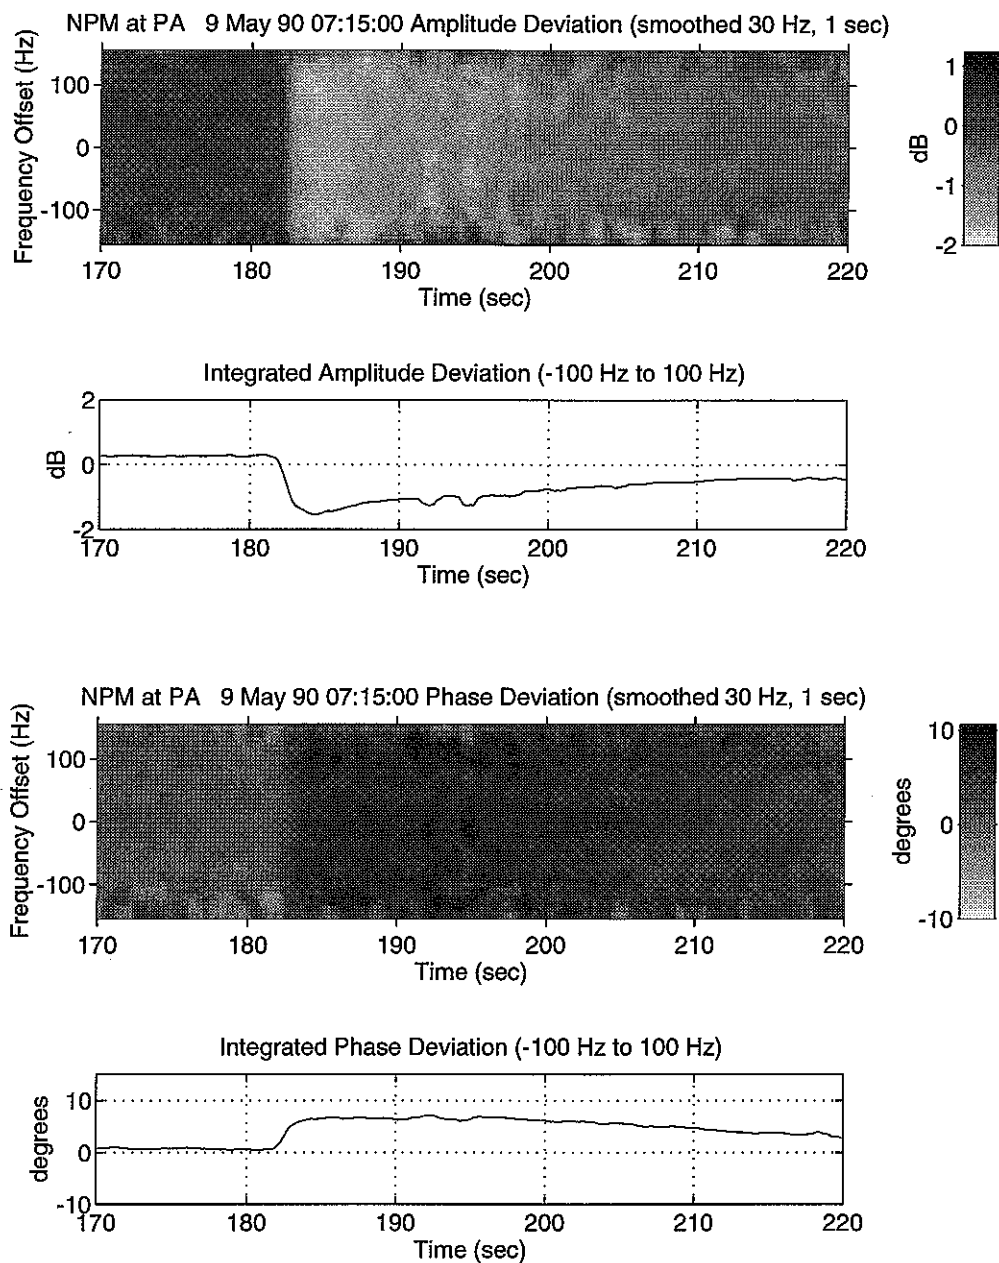


Figure 63. An expansion of the time axis in Figure 62, showing the single LEP event with an onset at about 182 seconds from the start of the record (07:18:02 UT). The perturbation is an amplitude decrease and a phase advance, and is smooth in frequency.

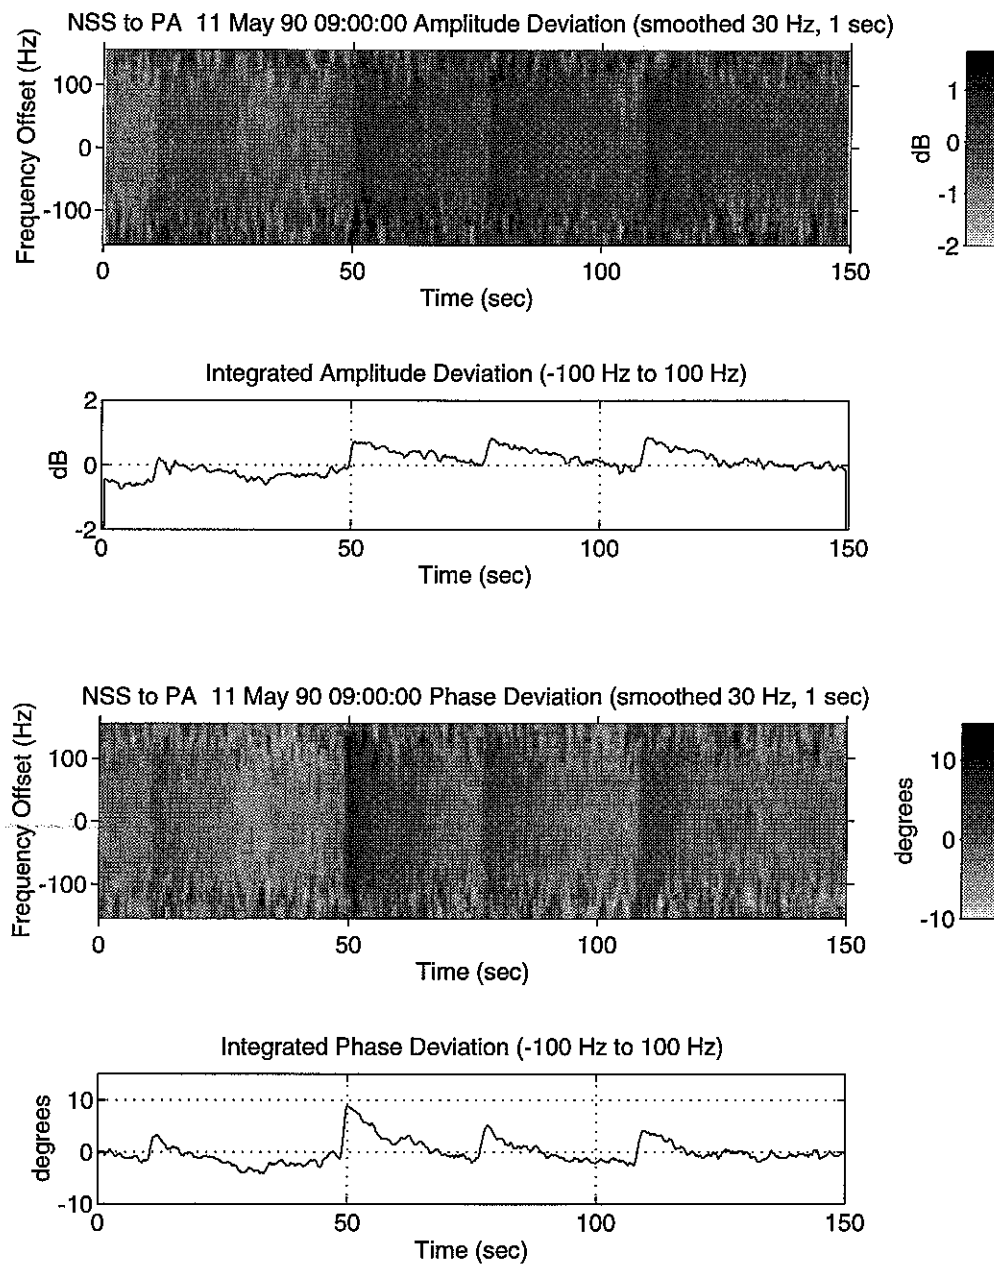


Figure 64. Four transfer function perturbations on the nss-to-Palmer path, at ~10, ~50, ~75, and ~110 seconds. The data have been smoothed, and the static portion of the transfer function suppressed.

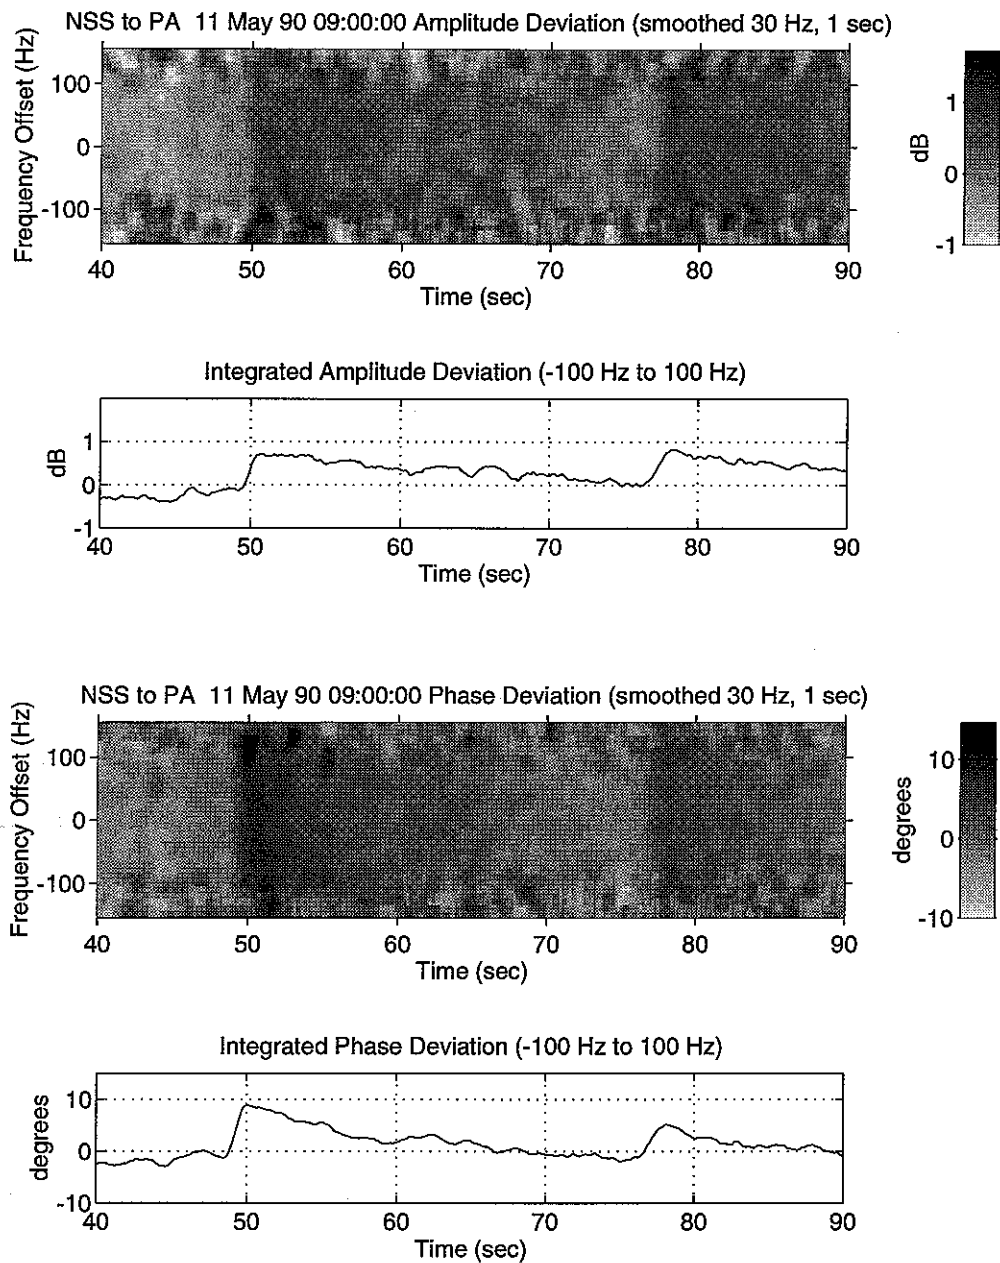


Figure 65. An expansion of the time axis in Figure 64, showing two LEP events with onsets at ~50 and ~77 seconds from the start of the record (09:00:50 and 09:01:17 UT). These perturbations are amplitude increases and phase advances, and are smooth in frequency.

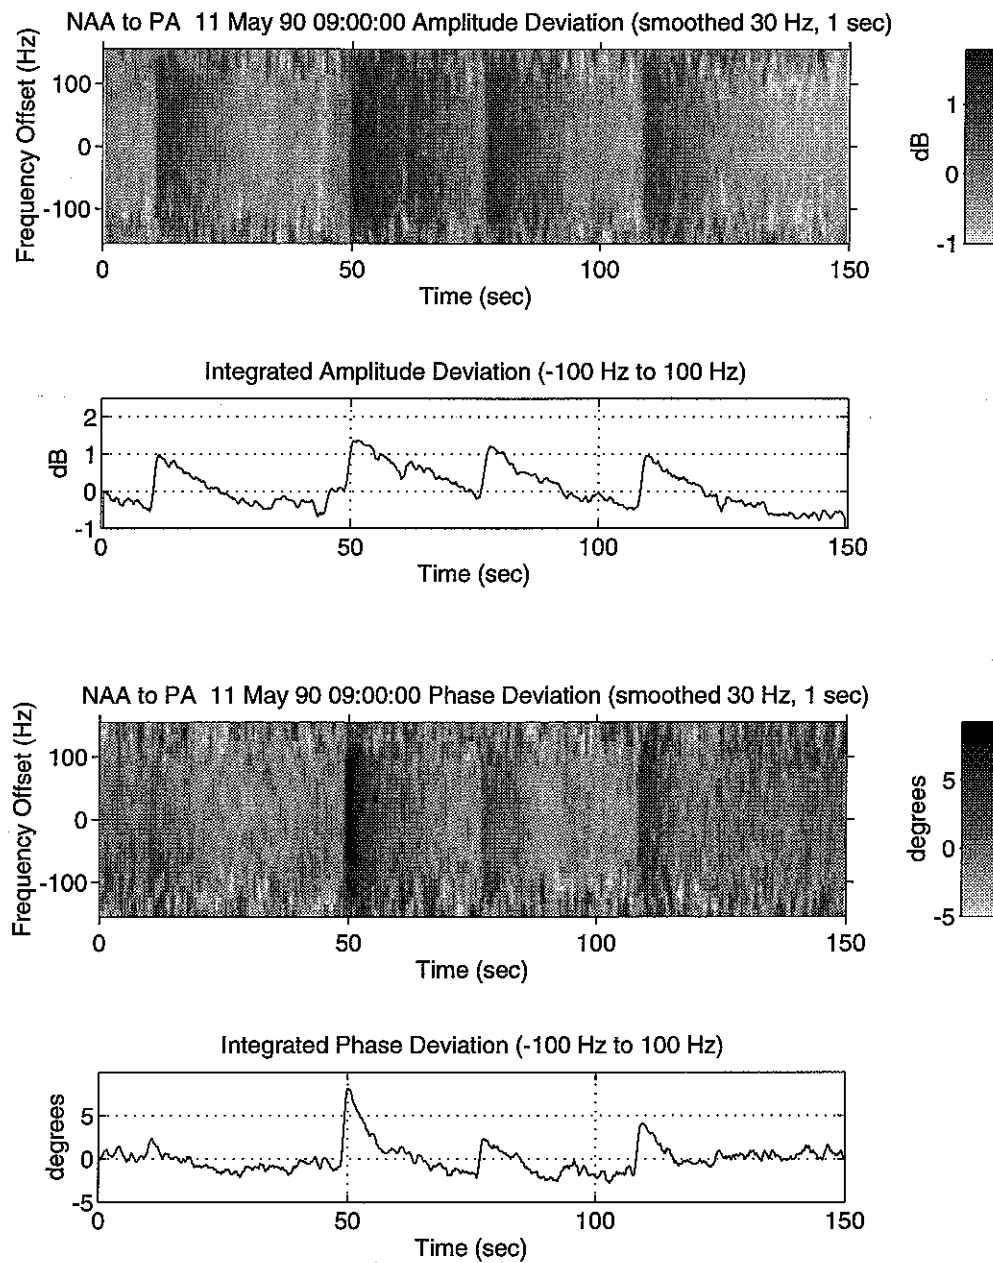


Figure 66. Four transfer function perturbations on the NAA-to-Palmer path, coincident with those on NSS, from Figure 64. The data have been smoothed, and the static portion of the transfer function suppressed.

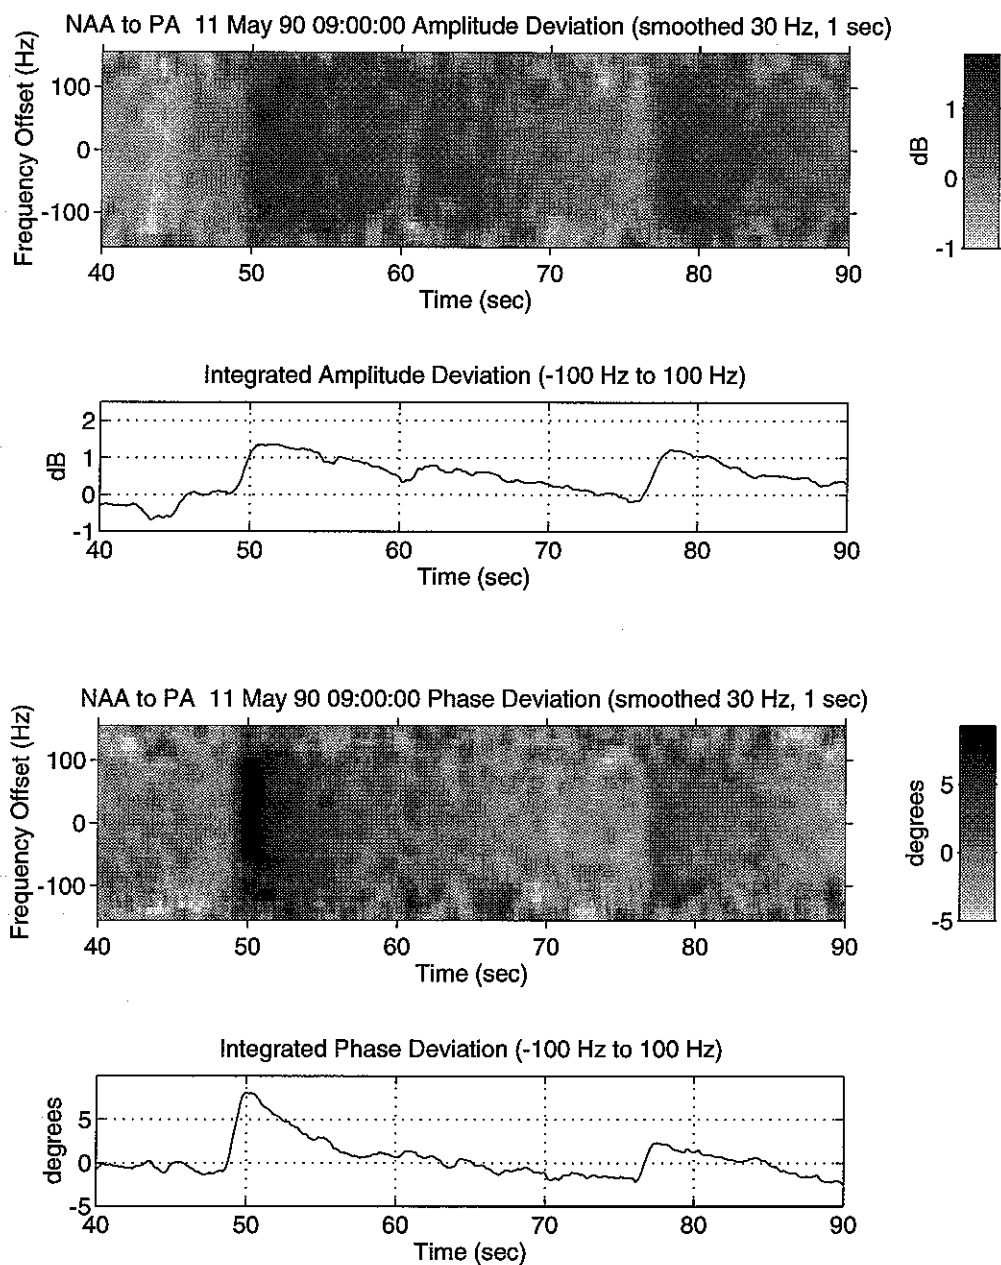


Figure 67. An expansion of the time axis in Figure 66, showing two LEP events with onsets at ~ 50 and ~ 77 seconds from the start of the record (09:00:50 and 09:01:17 UT). These perturbations are amplitude increases and a phase advances, and are smooth in frequency, and are coincident with the perturbations on NSS.

6.

SUMMARY AND SUGGESTIONS FOR FUTURE WORK

6.1. SUMMARY OF RESULTS AND CONTRIBUTIONS

In this dissertation, we have developed a novel technique for the measurement of the propagation characteristics of VLF radio waves in the earth-ionosphere waveguide, which takes advantage of modulated transmissions to characterize the propagation channel over a band of frequencies. The technique for measurement of the transfer function of the propagation channel does not rely on assumptions about the characteristics of that transfer function, as have all previous techniques reported in the literature.

This introduction of frequency as an independent variable in the propagation measurements allows direct comparison with existing propagation theory, and thus constrains the ionospheric parameters (electron density and temperature, as a function of altitude) which control the propagation of the VLF waves. Knowledge of these parameters is important in studies of the physics of the lower ionosphere, and measurements of them are difficult to perform by other means. Comparison of results of model calculations with measured transfer function data will allow the remote sensing of these parameters, on a continuous-time basis, and on a global scale.

In addition, simulation results presented here show that this technique is capable of detecting evidence of multiple-mode propagation and mode coupling, when these effects are significant. It is important to consider the mode structure of the incoming wave at the receiver, since common receivers use small magnetic loop antennas, which are not directly sensitive to the vertical electric field of the wave. This effect is of particular

importance when the direction of arrival of the incoming VLF wave is nearly in the plane of the loop.

Finally, application of this technique to a data set collected at Palmer Station, Antarctica shows that the measured transfer functions of the propagation channel do not exhibit sharp variations in frequency. This result is in general agreement with the predictions of the 3-D waveguide propagation model, and indicates that, on these paths, the mode structure of the received signal is dominated by a single waveguide mode. However, there are quantitative differences between the measurements and the predictions, which indicate that there are significant differences from the assumed electron density and/or temperature profiles.

6.2. SUGGESTIONS FOR FUTURE WORK

Since the focus of this work has been the development of the measurement technique, only a very limited data set has been analyzed. There is much room for the application of the techniques described here to larger data sets, and to propagation paths expected to show effects of more geophysical interest. For example, study of larger data sets from Palmer could determine the diurnal variations in the transfer functions of the four propagation channels studied in this work, as well as variations with geomagnetic activity. These data, in combination with model calculations using assumed ionospheric density and temperature profiles, could be used to test the extent to which the profiles are influenced by other factors, since the frequency-dependent processing yields many points of comparison with the models. In addition, a study of shorter, land-based paths might be expected to capture LEP-induced VLF signal perturbations which show detected differences across the available band of frequency; these would yield direct measurements of the transfer function perturbation, which could be directly compared with model results for the scattering from such perturbations.

Finally, some of the processing techniques investigated as a part of this work may hold promise for other analysis of field data in the future. This is particularly true given the rapid advance in the capability of inexpensive signal processing hardware. Examples of these techniques would include impulsive noise removal from broadband recordings, and the use of the phase of the subionospheric MSK signals as a time reference for phase studies of magnetospheric signals.

6.2.1. Impulsive Noise Removal for Broadband Recordings

One of the processing steps described in Chapter 3 is applied to the broadband data, before filtering and decimation to a lower sampling rate: the suppression of impulsive interference. In this work, we simply followed standard practice, and filtered and clipped the broadband data samples. However, some experimentation was done with more sophisticated techniques of interference suppression, by recognizing the impulses and excising them from the data set. This produced little gain in the narrowband analyses performed here, but may be of value to broadband studies, where spectral representations are used which cover 10 KHZ or so, and in which the masking effect of the impulsive interference from spherics is significant.

6.2.2. Implementation of Real-Time Processing

The data in this study were gathered via broadband recording, followed by post-processing to yield the propagation channel transfer function estimates. While this is useful for development of these techniques, this implementation is not suitable for general use in the field, because limitations on recording rates and capacity do not allow the analysis of the large data sets typically required to reveal novel geophysical phenomena. In addition, the upper frequency is limited by the recording rate of the sampling hardware, and does not allow the analysis of frequencies above about 33 KHZ in the current implementation.

To overcome these limitations, the techniques described in this work must be implemented in a field-equipment receiver, and must operate in real time. Fortunately, the complex processing is largely done after the individual transmitter has been isolated from the broadband data, and resampled at a lower data rate. The isolation of individual transmitters is a much simpler process, but must be implemented efficiently, since this processing must be applied to the entire broadband data stream.

This partition of the processing into a high data rate, low complexity, front end, and a low data rate, high complexity, back end, is not atypical of modern signal processing systems, and may be used to advantage in the design of a field receiver. A block diagram of one possible receiver is shown in Figure 68. In this receiver design, the entire processing path is digital. There are identical channel units, each of which performs the estimation process on one propagation path, and which plug into a front

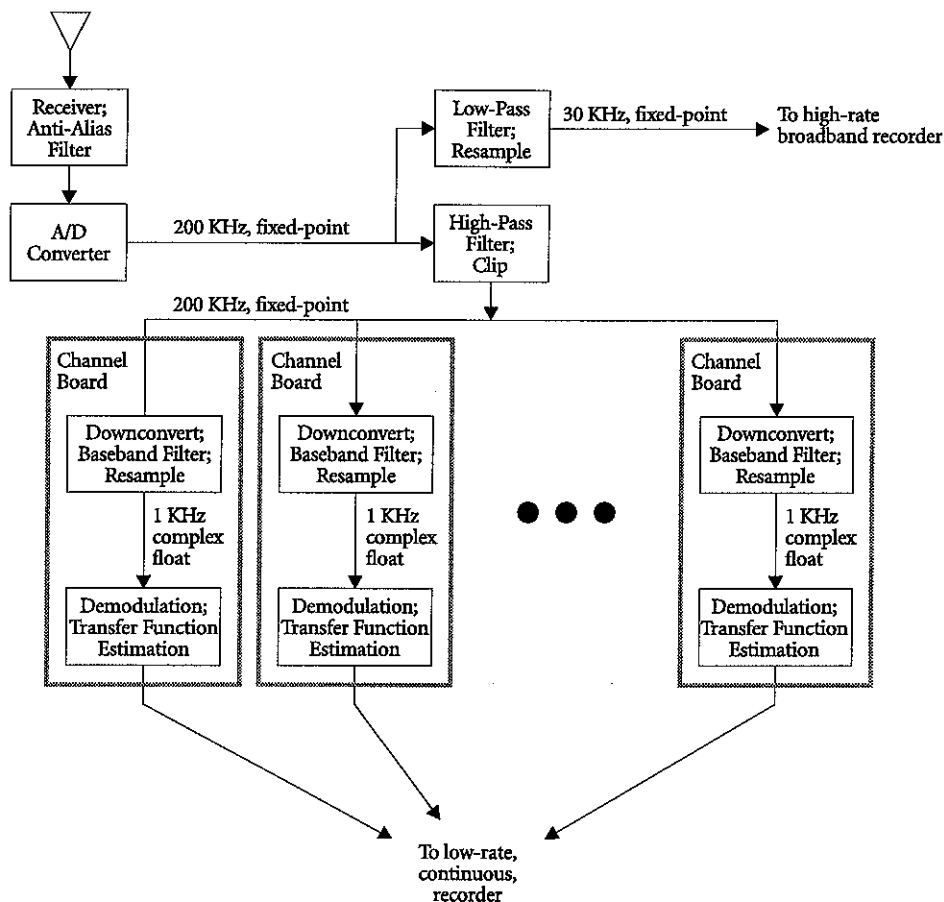


Figure 68. Possible implementation of the processing described in this dissertation, in a form suitable for use in the field. In addition to the propagation channel transfer function estimation, this receiver also provides a broadband output, which contains the whistler data. There is a single A/D converter and front-end processor, and a number of identical channel cards, each of which performs the transfer function estimation for one signal path.

end containing the A/D converter and broadband filtering and clipping functions. Each of the narrowband cards can analyze a single transmitter in the 10–60 KHz range. In addition to the narrowband outputs (one for each channel card), this receiver provides a broadband digital output covering the frequencies below about 10 KHz, which contains the whistlers and other magnetospheric signals of interest.

In this receiver, the processing is divided among DSP elements with different capabilities, which allows us to perform those filtering and conversion functions which occur on the broadband data, and therefore at high rates, using fixed-point processing,

and to perform the more complex estimation procedure using a slower, floating-point, DSP unit. Approximate processing loads for the different types of processors are:

Table 2. Processing Rate Requirements for Field Receiver

Element	Math Type	Rate
Low-Pass Filter, Clip	Fixed, 12 bits	40 MOPS
High-Pass Filter, Clip, Resample	Fixed, 12 bits	40 MOPS
Downconvert, Filter, Resample	Fixed, 12 bits	50 MOPS
Demodulate, Spectral Estimation	Floating-point	2 MFLOPS

Although the processing algorithms used in the final step could be scaled into fixed-point form, the low computing rate required (due to the fact that all this processing is happening at the lowered sample rate of 1 KHZ) makes a floating-point processor attractive, for ease of implementation.

In such a receiver, the channel cards are identical, and the parameters which govern their operation (center frequency, modulation, bit rate, *etc.*) could be downloaded at the start of operation from a host computer, which could also provide the recording system. For example, a receiver designed along these lines could interface to the host computer via a SCSI interface, and appear as a peripheral to the host system. This would allow the entire code run by all the processors in the receiver to be loaded from the host computer, and additional types of processing could be accommodated with software changes.

An additional advantage of such an approach is that channel cards may be programmed to analyze other types of signals. For example, cards to process LORAN and Omega signals could be built, and operate alongside those working on the MSK-modulated signals. Although not shown in the diagram, a data path from the front-end lowpass filter output could be provided to the channel cards, and a channel card could be used to detect interesting events, such as whistler arrivals; these triggers could be used to "freeze" a buffer on the broadband data, and dump it to the recording system for later, off-line, analysis.

6.2.3. Measurements at Separated Stations

One of the possible applications of the receiver described above is to make measurements at separated stations, using clocks synchronized with GPS receivers. Clocks may be synchronized with this method to better than 10 nsec accuracy, which corresponds to 0.1° of phase at 30 KHZ. This interferometry technique may be able to reconstruct the field scattered by a transient disturbance to the earth-ionosphere waveguide, and thus provide information on the source and location of the disturbance.

6.2.4. Studying Shorter, Overland Paths

Although all the data analyzed in this work was collected on long propagation paths, largely over seawater, the simulation results presented in Chapter 4 indicate that the shorter, overland, paths should provide the largest deviations with frequency, and, hence, the largest sensitivity for measurements of the electron density and temperature profiles using these techniques. Such data could be gathered, for example, at Stanford, from NSS (Annapolis, MD), NAA (Cutler, ME), and particularly the Air Force transmitter operating at 48.5 KHZ (Silver Creek, NE). The latter is a path which often experiences transient disturbances, but the higher frequency of this transmitter would require a recording system with larger bandwidth than that used in this work, which only allowed measurements to about 33 KHZ.

6.2.5. Use of Other Transmitters

The VLF and LF regions of the spectrum are crowded with transmitters, in addition to the large MSK transmitters used in this study. The transmitters most useful for use in these studies are the LORAN transmitters, which transmit a coded pulse format centered at 100 KHZ, and the Omega transmitters, which transmit long pulses in the 10–14 KHZ frequency range. Both systems are used for navigation: the Omega system is less accurate, but has global coverage; LORAN covers mainly coastal and populated areas, but provides more accurate position information.

An advantage of the LORAN transmissions is that they consist of accurately-controlled pulse trains, which are transmitted first by a master transmitter, and echoed

after a known delay by a series of slave transmitters. The pulse format is controlled so that the third zero crossing is accurately placed, and the duration of the pulses is short enough so that the spectrum of a LORAN transmitter extends up to 40 KHZ or so away from 100 KHZ. This broad spectrum may prove useful in multi-frequency measurements in the future, though, again, measurements can only be made with hardware built for the purpose. Since the transmissions are pulsed, the estimation technique used in this case may estimate the impulse response of the earth-ionosphere waveguide, rather than performing the estimation in the frequency domain.

The other widely-used navigation system operating in the VLF portion of the spectrum is the Omega system, which consists of eight transmitters scattered around the globe. These transmitters operate using signals which consist of a series of pulses in the 10–14 KHZ range. An advantage of the Omega system, as with LORAN, is that no message information is being sent, and so the demodulator need only estimate the time of arrival of the signal of interest, rather than estimating a bit sequence as well. Also, the Omega system transmissions cover a range of frequencies much wider than that occupied by the MSK transmitters. The main drawback to the Omega transmitters is that, because of the pulsed-cw format, they cover their entire list of frequencies every 10 seconds. This provides sufficient resolution for studies of the ambient electron density and temperature profiles, but is not sufficient to study LEP-induced perturbations and similar events.

One final possibility which should be mentioned in connection with this work is the eventual use of purpose-built transmitters, sending signals coded to optimize the estimation process, rather than to communicate data. Such a modulation would ideally provide a wide bandwidth (limited only by antenna matching considerations), with equal power density in all the frequency bands of interest. This is akin to traditional spread-spectrum techniques, because a spreading code is used to widen the transmitted bandwidth far beyond the baseband spectrum of the data involved.

BIBLIOGRAPHY

- Adams, C. D. D., and R. L. Dowden, VLF group delay of lightning-induced electron precipitation echoes from measurements of phase and amplitude perturbations at two frequencies, *J. Geophys. Res.*, 95, 2457, 1990.
- Allcock, G. McK., R. A. Helliwell, C. K. Branigan, and J. C. Murphy, Whistler and other very low frequency phenomena associated with the high-altitude nuclear explosion on July 9, 1962, *J. Geophys. Res.*, 68, 735, 1963.
- Barr, R., M. T. Reitveld, H. Hopka, and P. Stubbe, Effects of a heated patch of auroral ionosphere on VLF-radio wave propagation, *J. Geophys. Res.*, 90, 2861, 1985.
- Budden, K. G., *The Wave-guide Mode Theory of Wave Propagation*, Prentice-Hall, Englewood Cliffs, New Jersey, 1961.
- Budden, K. G., *The Propagation of Radio Waves*, Cambridge University Press, Cambridge, 1985.
- Burgess, W. C., and U. S. Inan, Simultaneous disturbance of conjugate regions in association with individual lightning flashes, *Geophys. Res. Lett.*, 17, 259, 1990.
- Burgess, W. C., and U. S. Inan, The role of ducted whistlers in the precipitation loss and equilibrium flux of radiation belt electrons, *J. Geophys. Res.*, 98, 15643, 1993.
- Carlson, A. B., *Communication Systems*, McGraw-Hill, New York, 1986.
- Carpenter, D. L., and J. W. LaBelle, A study of whistlers correlated with bursts of electron precipitation near $L=2$, *J. Geophys. Res.*, 87, 4427, 1982.
- Carpenter, D. L., and U. S. Inan, Seasonal, latitudinal, and diurnal distributions of whistler-induced electron precipitation events, *J. Geophys. Res.*, 92, 3429, 1987.
- Cotton, P. D., and A. J. Smith, Signature of burst particle precipitation on VLF signals propagating in the Antarctic earth-ionosphere waveguide, *J. Geophys. Res.*, 96, 19,375, 1991.
- Crombie, D. D., Further observations of sunrise and sunset fading of VLF signals, *Radio Science*, 1, 47, 1966.
- Crombie, D. D., A. H. Allen, and M. Newman, Phase variations of 16 kc/s transmissions from Rugby as received in New Zealand, *Proc. IEE*, 105(B), 301, 1958.

- Davies, K., *Ionospheric Radio*, Peter Peregrinus, Ltd., London, 1990.
- Dowden, R. L., and C. D. D. Adams, Modal effects on amplitude perturbations on subionospheric signals (Trimpis) deduced from two-frequency measurements, *J. Geophys. Res.*, 94, 1515, 1989.
- Dowden, R. L., and C. D. D. Adams, Ionospheric heating by VLF transmitters?, *J. Atmos. Terr. Phys.*, 54, 1051, 1992.
- Dungey, J. W., Loss of Van Allen electrons due to whistlers, *Planet. Space Sci.*, 11, 591, 1963.
- Ferguson, J. A., and Snyder, F. P., Approximate VLF/LF waveguide mode conversion model, *Tech. Doc. 400*, Naval Ocean Systems Center, San Diego, California, 1980.
- Ferguson, J. A., F. P. Snyder, D. G. Morfitt, and C. H. Shellman, Long-wave propagation capability and documentation, *Tech. Doc. 1518*, Naval Ocean Systems Center, San Diego, California, 1989.
- Field, E. C., and Lewinstein, Amplitude-probability distribution model for VLF/ELF atmospheric noise, *IEEE Trans. Commun.*, COM-26, 83, 1978.
- Fraser-Smith, A. C., and R. A. Helliwell, The Stanford University ELF/VLF radiometer project: measurement of the global distribution of ELF/VLF atmospheric noise, *Proc. IEEE Internat. Symp. on Electromag. Computability*, Wakefield, Massachusetts, 1985.
- Freidel, R. H. W., and A. R. W. Hughes, Effect of magnetic field distortions on gyroresonance at low l values, *J. Geophys. Res.*, 98, 1,669, 1993.
- Gallejs, J., *Terrestrial Propagation of Long Electromagnetic Waves*, Pergamon Press, Oxford, 1972.
- Glukhov, V. S., V. P. Pasko, and U. S. Inan, Relaxation of transient lower ionospheric disturbances caused by lightning-whistler induced electron precipitation bursts, *J. Geophys. Res.*, 97, 16971, 1992.
- Harris, F. J., On the use of windows for harmonic analysis with the Discrete Fourier Transform, *Proc. IEEE*, 66, 51, 1978.
- Helliwell, R. A., *Whistlers and Related Ionospheric Phenomena*, Stanford University Press, Stanford, California, 1965.
- Helliwell, R. A., Intensity of discrete VLF emissions, in *Particles and Fields in the Magnetosphere*, B. M. McCormac (ed.), p. 292, 1970.
- Helliwell, R. A., and D. L. Carpenter, Whistlers excited by nuclear explosions, *J. Geophys. Res.*, 68, 4409, 1963.
- Helliwell, R. A., J. P. Katsufakis, and M. L. Trimpis, Whistler-induced amplitude perturbation in VLF propagation, *J. Geophys. Res.*, 78, 4679, 1973.

- Imhof, W. L., E. E. Gaines, H. D. Voss, J. B. Reagan, D. W. Datlowe, J. Mobilia, R. A. Helliwell, U. S. Inan, J. Katsufakis, and R. G. Joiner, The results from the SEEP active space plasma experiment: effects on the ionosphere, *Radio Science*, 20, 511, 1985.
- Inan, U. S., VLF heating of the lower ionosphere, *Geophys. Res. Lett.*, 17, 729, 1990.
- Inan, U. S., and D. L. Carpenter, On the correlation of whistlers and associated subionospheric VLF/LF perturbations, *J. Geophys. Res.*, 91, 3106, 1986.
- Inan, U. S., and D. L. Carpenter, Lightning-induced electron precipitation events observed at L~2.4 as phase and amplitude perturbations on subionospheric VLF signals, *J. Geophys. Res.*, 92, 3293, 1987.
- Inan, U. S., D. L. Carpenter, R. A. Helliwell, and J. P. Katsufakis, Subionospheric VLF/LF phase perturbations produced by lightning-whistler induced particle precipitation, *J. Geophys. Res.*, 90, 7457, 1985.
- Inan, U. S., D. C. Shafer, and W. Y. Yip, Subionospheric VLF signatures of nighttime D-region perturbations in the vicinity of lightning discharges, *J. Geophys. Res.*, 93, 11455, 1988a.
- Inan, U. S., T. G. Wolf, and D. L. Carpenter, Geographic distribution of lightning-induced electron precipitation observed as VLF/LF perturbation events, *J. Geophys. Res.*, 93, 9841, 1988b.
- Inan, U. S., M. Walt, H. D. Voss, and W. L. Imhof, Energy spectra and pitch angle distributions of lightning-induced electron precipitation: analysis of an event observed on the S81-1 (SEEP) satellite, *J. Geophys. Res.*, 94, 1379, 1989.
- Inan, U. S., T. F. Bell, and J. V. Rodriguez, Heating and ionization of the lower ionosphere by lightning, *Geophys. Res. Lett.*, 18, 705, 1991.
- Inan, U. S., J. V. Rodriguez, S. Lev-Tov, and J. Oh, Ionospheric modification with a VLF transmitter, *Geophys. Res. Lett.*, 19, 2071, 1992.
- Inan, U. S., J. V. Rodriguez, and V. P. Idone, VLF signatures of lightning-induced heating and ionization of the nighttime D region, *Geophys. Res. Lett.*, 20, 2355, 1993.
- ITU, *International Frequency List*, Volume 1, International Telecommunications Union, 1979.
- Ljung, Lennart, *System Identification Theory for the User*, Prentice-Hall, Englewood Cliffs, New Jersey, 1987.
- Morfitt, D. G., and C. H. Shellman, MODESRCH: an improved computer program for obtaining ELF/VLF/LF mode constants in an earth-ionosphere waveguide, *Interim*

- Report Number 77T*, Propagation Technology Division, Naval Ocean Systems Center, San Diego, California, 1976.
- Motchenbacher, C. D., and Connelly, J. A., *Low-Noise Electronic System Design*, John Wiley and Sons, New York, 1993.
- Naval Ocean Systems Center, *Fixed VLF/LF Data Base*, Naval Ocean Systems Center, San Diego, California, 18-OCT-89.
- Oppenheim, A. V., and R. W. Schaffer, *Digital Signal Processing*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1975.
- Paschal, E. W., Calibration of very-low-frequency receiving systems for Stanford University (internal report), Stanford Electronics Laboratories, Stanford, California, 1978.
- Paschal, E. W., The design of broad-band VLF receivers with air-core loop antennas (internal report), Stanford Electronics Laboratories, Stanford, California, 1988.
- Pasko, V. P., and U. S. Inan, Recovery signatures of lightning-associated VLF perturbations as a measure of the lower ionosphere, *J. Geophys. Res.* (in press), 1994.
- Poulsen, W. L., Modeling of very low frequency wave propagation and scattering within the earth-ionosphere waveguide in the presence of lower ionospheric disturbances, PH. D. dissertation, Stanford University, Stanford, California, 1991.
- Poulsen, W. L., T. F. Bell, and U. S. Inan, Three-dimensional modeling of subionospheric VLF propagation in the presence of localized D-region perturbations associated with lightning, *J. Geophys. Res.*, 95, 2355, 1990.
- Poulsen, W. L., U. S. Inan, and T. F. Bell, A multiple-mode, three-dimensional model of VLF propagation in the earth-ionosphere waveguide in the presence of localized D-region disturbances, *J. Geophys. Res.*, 98, 1705, 1993a.
- Poulsen, W. L., T. F. Bell, and U. S. Inan, The scattering of VLF waves by localized ionospheric disturbances produced by lightning-induced electron precipitation, *J. Geophys. Res.*, 98, 15553, 1993b.
- Ramo, S., J. R. Whinnery, and T. Van Duzer, *Fields and Waves in Communication Electronics*, John Wiley and Sons, New York, 1984.
- Rees, M. H., Auroral ionization and excitation by incident energetic electrons, *Planet. Space Sci.*, 11, 1209, 1963.
- Rodriguez, J. V., Modification of the earth's ionosphere by very-low-frequency transmitters, PH. D. dissertation, Stanford University, Stanford, California, 1994.
- Rodriguez, J. V., U. S. Inan, and T. F. Bell, D-region disturbances caused by electromagnetic pulses from lightning, *Geophys. Res. Lett.*, 19, 2067, 1992.

- Smith, A. J., and P. D. Cotton, The Trimpf effect in Antarctica: observations and models, *J. Atmos. Terr. Phys.*, 52, 341, 1990.
- Taranenko, Y. N., U. S. Inan, and T. F. Bell, Optical signatures of lightning-induced heating of the D-region, *Geophys. Res. Lett.*, 19, 1815, 1992.
- Thompson, N. R., Whistler mode signals: spectrographic group delays, *J. Geophys. Res.*, 86, 4795, 1981.
- Tolstoy, A., and T. J. Rosenberg, A quasi three-dimensional propagation model for subionospherically propagating VLF radio waves, *Radio Science*, 20, 535, 1985.
- Tolstoy, A., T. J. Rosenberg, U. S. Inan, and D. L. Carpenter, Model predictions of subionospheric VLF signal perturbations resulting from localized, electron-precipitation induced, ionization enhancement regions, *J. Geophys. Res.*, 91, 13473, 1986.
- Turtle, J. P., E. C. Field, Jr., C. R. Warber, and P. R. McGill, Low-frequency transverse electric atmospheric noise: measurement and theory, *Radio Science*, 24, 325, 1989.
- Uman, M. A., *The Lightning Discharge*, Academic Press, Orlando, Florida, 1987.
- Warber, C. R., and E. C. Field, A long wave TE/TM noise prediction model, *Proc. of NATO-AGARD conference on ELF/VLF/LF radio propagation and systems aspects*, Brussels, Belgium, 1993.
- Watt, A. D., *VLF Radio Engineering*, Pergamon Press, Oxford, 1967.
- Wolf, T. G., Remote sensing of ionospheric effects using very low frequency radio signals, PH. D. dissertation, Stanford University, Stanford, California, 1990.
- Wolf, T. G., and U. S. Inan, Path-dependent properties of subionospheric VLF amplitude and phase perturbations associated with lightning, *J. Geophys. Res.*, 95, 20997, 1990.
- Yip, W.-Y., U. S. Inan, and R. E. Orville, On the spatial relationship between lightning discharges and propagation paths of perturbed subionospheric VLF/LF signals, *J. Geophys. Res.*, 96, 249, 1991.