

ATMOSPHERIC NOISE MITIGATION FOR LORAN

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

C.O. Lee Boyce Jr.

June 2007

© Copyright 2007

by

C.O. Lee Boyce Jr.

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.

(J. David Powell) 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.

(Per K. Enge)

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.

(Stephen M. Rock)

Approved for the University Committee on Graduate Studies.

Abstract

While aircraft users increasingly rely on the Global Positioning System (GPS) for navigation, the received signal strength from these distant satellites is low and hence may easily be jammed. Of the alternative systems available to provide aircraft with a backup to GPS for non-precision approach (NPA), Loran promises to be an exceptional candidate since its signal is present over the entire continental United States (CONUS).

Loran, a low-frequency hyperbolic radio-navigation system, provides positioning with quarter nautical mile accuracy, 95% of the time. Since it was not originally designed for aircraft navigation, the system requires upgrades to both the transmitters and user receivers to meet NPA requirements. Furthermore, to be viable, Loran's coverage area, defined as where it meets the NPA requirements, must be across most of CONUS. Impeding its success is Loran's susceptibility to lightning and other atmospheric noise.

Determining Loran's coverage requires the estimation of atmospheric noise and the processing gain or credit for noise reduction due to non-linear processing techniques within a typical receiver. The traditional method uses the standard atmospheric noise model from the International Telecommunications Union (ITU), and estimates the processing gain as a constant. While easy to implement, such an analysis produces overly conservative results and, therefore, a poor estimate of coverage area.

This dissertation presents a new method for predicting coverage by extending the ITU

model and by demonstrating empirically that a lower bound on atmospheric noise impulsivity is proportional to the noise strength. Since it is well known that the gain in signal-power to noise-power ratio (SNR) in a receiver produced by non-linear signal processing is proportional to the noise impulsivity, then the correlation of noise impulsivity to noise strength implies that the processing gain is also proportional to noise strength and is not simply constant as previously predicted.

The new method results in an increase of 8 dB to the SNR over those predicted by the traditional method. The result is an increase in system availability across most of CONUS, thereby increasing the likelihood that a modernized Loran can become certified for NPA.

Acknowledgments

It was in the previous millennium that I started my long journey through graduate school. Along my path, I have been blessed and graced with delightful, memorable, and dear company. Like many travels, when the trip got hard, often it was the folks around me that kept me going. So, I would like to thank all of those who have contributed to make this journey one of the most wonderful of my life. Please excuse the length of the list, it is commensurate with the extent of the journey, and my gratitude to so many.

To start off: Thank you to the one who makes all things possible. While, it may not be fashionable in the age of reason, I have been a devout theist and Christian all of my life, and am grateful for the comfort this belief has given me.

Thanks Mom, your encouragement, love, faith and support have helped make it possible for me to make it to this point today. It has been because of your guidance and that of Baba, Dzedo, Linda, and Roy that I have been encouraged all of my life to seek out education and to better myself. I also thank Grandma, Dad, Ellen, Molly, Dennis, Nancy, Hallie, Elise, Lambert, Tim, and Judy for coming out to graduation, taking me out, and being supportive of their relative who has been lost at school these many years.

Thanks to my California family: April, Ron, Caden, Graham, Heather, Luke, Elizabeth, Ted, Sierra, Lucia, Crystal, Lee Walker, Rita, Paula, Skip, and Teddy. The touch of family has been one of my chief joys out here in the West.

There are those who have helped in the direction and management of my graduate career. Their efforts have made my journey through the GPS lab a smooth one. First, thanks to Dave Powell, my advisor; you have helped me to enjoy the journey as much as the destination. Your philosophy on a balanced life is one that I have tried to maintain by keeping work, family, and fun in harmony. I also am grateful that you never gave up on me over the years. Also, thanks to Per Enge, my co-advisor; you have provided me with technical direction and funding throughout my career in Loran, and have kept my spirits up along the way with your bright outlook on life. Steve Rock, thanks for being the third member of my committee; you have generously guided me throughout the years with humor and patience. Thank you all for your insight and comments which have improved the writing of this dissertation. Also, thank you Len Tyler for your generosity in chairing my Defense.

As with any great lab, there must be great leadership. Sherman Lo, as the post-doctorate in charge of Stanford's Loran effort, you have provided me with that leadership. Thank you. By being both a good manager and a good friend, you have selflessly given your time to help me throughout my tenure in the GPS lab. Also, thanks to Todd and Sam for setting fine examples as research associates and helping me whenever I took the time to ask.

Gratefully, I thank Mitch Narins of the FAA, for all of the funding which has supported my Loran work. Your encouragement and passion for Loran is an inspiration and makes me proud to know such fine people as you work for this country. Also, thank you Ben Peterson, Tom Gunther, and all of the members of the LORIPP for your continued support and mentorship in Loran.

For cheerfully helping me through all of the bureaucratic hurdles associated with any academic and government undertaking, and for providing a large supply of chocolate, thank

you Sherann and Dana. Thank you Doug, not only for your countless hours keeping the machines running smoothly and always having the right part under your desk when I needed it, but for defending this country so I can sit and enjoy the life of a graduate student. Encore, mille grazie Aldo, for the last minute machining and for threats of whippings with a wet noodle if I didn't hurry up and finish.

Two individuals within the department provided me with great support during the loss of my first advisor: George Springer and Sally Gressens. Thank you both for helping me through that hard time when I needed it most.

Also, I thank the rest of the Aero/Astro Staff, Lynn, Robin, Carolyn, and William. You have helped me with paperwork as well as provided me office space, supplies, and great conversation for these many years.

A lot of work went into the data collection efforts for this dissertation. For that, I also thank Nick Alexeev and Di Qiu for the late nights helping to build components and getting me packed up for my long expeditions. Thanks Mike S. Koenig for your drawings, support, and rapier wit. I am grateful for the editing efforts of Fiona Walter who has been instrumental in making this dissertation read as if English really is my first language.

The data collection effort would have not been possible without the help of Demoz Gebre-Egziabher, Bill Winn, Bill Beasley, Steve Hunyady, Sandy Kieft, and Graydon Aulich. Thank you all for your contributions of space and time and in putting up with the uncertain schedules.

Tim, Chris, and Chris, I thank you for putting me back together after all of the accidents which accompany living a balanced lifestyle and teaching me the value of ice.

Before wrapping up the business side of things, I thank Ben, Harris, Vaughn, Roger, Matt, Mark, Stuart, and other Silicon Valley folks for supporting my consulting habit. You

helped me take this trip in style.

Now, I turn to those good friends whom it has been my privilege to know. These are those folks whose laughter, conversation, and thoughtfulness have made me richer in mind, body, and spirit.

First on the list are Guttorm, Konstantin, and Gabe. We go as far back as to our Master's years. You are as close to me as brothers, and will always be a part of my life.

Next, there are those brethren from the subterranean lab: Brian, Žsolt, Jim, Caleb, Jim, Eric, and all of the rest of the SSL team. It's not just the radiation which warms my thoughts of our time together. Also thanks Professor How for getting me started; I hope you are now in a better place.

Mike, Eric, Sasha, Chad, Wendy, Keith, Ming, Jenny, Dennis, Frank, Grace, and all of the rest of the GPS lab, I thank you. Also, Edgar, Pascal, Logi, and all of my Aero./Astro and EE academic siblings. I also am filled with gratitude for Greg, Dan, and Nick for being most excellent roommates, Santiago for being a most excellent officemate, and Pete, for simply being most excellent.

There are those from my Maryland years who also have stuck out this ride: thank you John, Jonathan, and Chris, for being both friends and touchstones to my past. For those from my Master's and JPL days, I thank Erin, Özgür, Gary, Mike, and Jolie who also have kept in contact despite these long "four" years of school.

Finally, I save the one for whom I am most thankful: Julie. Despite all of the years, the hardships of being apart, and the struggles, you have continued to love me and to be supportive of me and my pursuit of an education. From the bottom of my heart, I thank you for all of the beauty and joy you have brought into my life and for every moment we have shared together.

Contents

Abstract	iv
Acknowledgments	vi
1 Introduction	1
1.1 Motivation	3
1.1.1 National Airspace System	5
1.1.2 The Global Positioning System	6
1.1.3 Need for a Backup System	7
1.1.4 NAS Costs	10
1.1.5 Aircraft Approach Requirements	12
1.2 Description of RNP 0.3	14
1.3 Prior Loran Research	17
1.3.1 Loran for Navigation	17
1.3.2 Traditional Coverage Model	18
1.3.3 Atmospheric Noise	20
1.4 Contributions	22
1.5 Dissertation Outline	23

2	Loran System Background	25
2.1	History	25
2.2	Loran System Configuration	26
2.3	Loran Receiver Basics	32
2.3.1	Receiver Design	32
2.3.2	Envelope Processing	33
2.4	Loran Signal Vulnerabilities	37
2.4.1	Transmission and Propagation Vulnerabilities	37
2.4.2	Reception Vulnerabilities	39
2.4.3	Signal Interference and Noise	40
2.5	Loran Coverage Tool	41
3	Atmospheric Noise and Signal Processing Review	43
3.1	Atmospheric Physics	44
3.1.1	Cloud-to-Ground Lightning	45
3.1.2	Return Stroke Electric Fields	48
3.2	ITU Atmospheric Noise Model	53
3.2.1	ITU History	54
3.2.2	CCIR/ITU Atmospheric Noise Model	55
3.2.3	ITU Measurement System Details	58
3.2.4	Noise Factor Statistics	60
3.2.5	Instantaneous Noise Envelope Measurements	67
3.2.6	ITU Model Issues	71
3.3	Signal Processing	74

4	Problem Statement	79
4.1	Problem Summary	79
4.2	Traditional Model Parameters	80
4.3	SNR Calculations	81
4.4	Number of Towers Calculation	82
4.5	Traditional Coverage results	84
4.6	Research Questions	84
5	Signal Processing Development	90
5.1	Envelope Ratio Distributions	91
5.1.1	Overview	91
5.1.2	Assumptions	91
5.1.3	Noise-only Case	92
5.1.4	Low SNR case	94
5.1.5	High SNR Case	96
5.2	Probability of Wrong Cycle Selection	99
5.3	Signal Processing with APDs	103
5.3.1	APD Review and Interpretation	103
5.3.2	Reading the APD	104
5.3.3	Signal Suppression	106
5.3.4	Processing gain	109
5.3.5	Total Non-linear Processing Gain	109
5.4	Noise and the Limitations of Average Power	112
5.4.1	Noise	112
5.4.2	Problems with Non-Gaussian Noise	114

5.4.3	Problems with Average Power	114
6	Data Collection System	118
6.1	Brief History	118
6.2	Location Requirements	119
6.3	2005 Data Collection Campaign Locations & System Configurations	121
7	Data Comparison to Atmospheric Physics	127
7.1	Time Domain Comparison	128
7.2	Frequency Domain Comparison	130
7.3	NLDN Comparison	133
7.4	Verification of the Stanford Loran Receiver	134
7.5	Extension to Power Spectral Density	136
7.5.1	Return Stroke Time Interval	137
7.5.2	Return Stroke Power Spectral Density	138
7.6	Effect of Return Strokes on Loran	141
7.7	Time Variation of Power Level	143
8	Improved Loran Coverage Model	147
8.1	Verifying the ITU Model	147
8.2	Estimates of RMS Electric Field Strength	148
8.3	Instantaneous Noise Envelope APDs	154
8.4	Correlation of Voltage Deviation and RMS Noise Field Strength	157
8.5	Improved Coverage Results	159

9	Conclusion and Future Work	165
9.1	ITU Model	165
9.2	Signal Processing	167
9.3	Loran Coverage	167
9.4	Implications	168
9.5	Future Work	169
9.5.1	Improving Loran Coverage	169
9.5.2	Further Testing	170
9.5.3	Refining the Model	173
9.5.4	Time Domain Model	174
9.5.5	Inertial Aiding	175
9.6	Final Thoughts	176
A	Fourier Analysis	177
A.1	Fourier Series	177
A.2	Continuous Fourier Transform	178
A.3	Discrete Fourier Transform	179
A.4	Physical Interpretation	184
A.5	Power Spectral Density	187
B	Antenna Noise Calculation	189
C	Loran Receiver Calibration	194
C.1	Calibration using Loran Signal Models	194
C.2	Calibration Based on Lightning Data	197

D SNR Definition	204
E Acronyms and Symbols	208
Bibliography	215

List of Tables

1.1	Required Navigation Performance 0.3 NM Specifications	14
2.1	Loran Station Phase Codes	29
6.1	Test Location and Lightning Characteristics	120
C.1	Site Location	195
C.2	Return Stroke Data	200

List of Figures

1.1	Comparison of traditional and improved coverage prediction charts.	4
1.2	Earth-bound user receiving GPS signals from satellites.	6
1.3	200 mW GPS jammer available on the web for \$700 in 2006.	9
1.4	Depiction of VOR (bearding) and DME (range) systems.	11
1.5	Dissertation chapter organization. Background and development sections are in blue. Thesis topics and results in red.	24
2.1	Loran towers across the US and Canada.	27
2.2	Loran GRI timing diagram. The Master station is denoted by M and the secondary towers by W, X, & Y.	28
2.3	Idealized Loran pulse (red) and corresponding envelope (blue)	30
2.4	Loran hyperbolas or locus of constant time difference, between Master and Secondary X, TD_{MX} (blue), and between Master and Secondary Y, TD_{MY} (red)	31
2.5	Loran front end and complex envelope generation.	34
2.6	Loran envelope and ratio test.	35
2.7	Loran vulnerabilities, the LORIPP members working on them, and the technologies used to overcome them.	38

2.8	Absolute electric field envelope values in the Loran band. Note the return stroke interfering with the third Loran envelope pulse.	41
3.1	Component time history of negative cloud-to-ground (C2G) lightning flash from [42].	47
3.2	Typical electric fields for first (solid lines) and subsequent (dashed lines) return strokes at distances of 1, 2, 5, 10, 15, 50, and 200 km. Reproduction from Preta [45].	50
3.3	Frequency spectrums of electric fields for first return strokes. Reproduction from Preta [45].	52
3.4	Placement of the original CCIR/ITU testing sites.	57
3.5	Median summer noise Values, F_{am} , in dB above KT_0B at 2000 h	63
3.6	Log-normal F_a distribution with adjustments due to uncertainties σ_m and σ_D . 64	
3.7	Median 99.9% rms envelope noise voltage, E_{rms} , with 50% confidence taken over all hours and seasons.	66
3.8	Reproduction of the APD curves parameterized by V_d from ITU P.372-7 [29]. 69	
3.9	Linear and non-linear processing functions	77
4.1	SNR and corresponding reception area of Boise City tower corresponding to a -24 dB threshold under different noise conditions.	83
4.2	Number of towers available at 50% and 99.9% noise field strength.	85
4.3	Loran availability using only simple constant credit for non-linear processing. 86	
4.4	Current pessimistic Loran coverage prediction based on the outcomes of this research.	89
5.1	Autocorrelation of I samples.	92

5.2	Joint probability distribution of two Gaussian variables with zero-mean and equal variance. Indicative of noise-only case ($\text{SNR} = -\infty$ dB).	93
5.3	Joint probability distribution of one zero-mean and one finite-mean Gaussian variable with equal variance. Indicative of low SNR case ($\text{SNR} = 0$ dB). . .	95
5.4	Joint probability distribution of one zero-mean and one finite-mean Gaussian variables with equal variance. Indicative of high SNR case ($\text{SNR} = 20$ dB). . .	96
5.5	I-Q plane view of noise added to an envelope voltage.	97
5.6	Distributions of the $15\ \mu\text{s}$ ratio point for high and low SNR. Comparison of the exact CDF and APD to the Gaussian approximations for the distribution are shown.	98
5.7	Probability of wrong cycle selection given the post-processing SNR under Gaussian noise.	102
5.8	Amplitude probability density (APD) function of Rayleigh noise ($V_d = 1.05$) and impulsive noise ($V_d = 10$)	105
5.9	APD showing the division of noise as being impulsive or Rayleigh in distribution.	106
5.10	Loss due to signal suppression.	108
5.11	APD depicting the rms of the Rayleigh level as below that of the overall distribution.	110
5.12	Difference between the rms value of the Rayleigh noise and the overall rms value of the distribution.	111
5.13	Total non-linear processing gain as a function of V_d	112
5.14	Noise categories.	113
5.15	Example of the effects of quadrupling the time interval on rms power. . . .	115

5.16	Example showing how the same rms voltage value can take on different shapes in a finite interval.	116
5.17	Example showing how the same rms voltage value can take on different averages and voltage deviations.	117
6.1	Test site locations throughout the 2004-2006 data collection campaigns. . .	120
6.2	Sarkey's Energy Center at the University of Oklahoma, Norman, OK. . . .	121
6.3	View from atop the Sarkey's Center with Locus Loran and TrueTime GPS Antennna.	122
6.4	Circuit board of Stanford Loran Receiver Revision A.	123
6.5	200 kHz bandwidth flat-plate antenna provided by Professor Beasley of the University of Oklahoma. The antenna measures 1 meter on each side and is held down with cinder blocks.	124
6.6	University of Oklahoma 2005 Loran band data collection system. The Langmuir 2005 system is similar, but with only the first two data channels. .	126
7.1	Reprint of Fig. 3.2 for comparison [45].	128
7.2	Electric fields for return strokes at various distances.	129
7.3	Lightning electric field return stroke captured by flat-plate antenna.	133
7.4	Electric field frequency spectrum of a single return stroke pictured in Fig. 7.3 using Preta's method.	134
7.5	Return stroke electric field frequency spectrum at 100 kHz versus distance using Preta's method.	135
7.6	Time domain traces of a return stroke and a return stroke filtered through a Loran bandpass filter (35 kHz wide bandpass filter centered at 100 kHz). . .	136
7.7	Power spectral density (PSD) of return stroke from Figure 7.3.	138

7.8	Power spectral density coefficients for the flat-plate antenna at 100 kHz versus return strokes range.	139
7.9	Electric field strength in dB($\mu\text{V}/\text{m}$) of return strokes within the Stanford Loran Receiver versus range.	140
7.10	Time history of lightning strikes and Loran. Each figure represents differ- ent amounts of enlargement on the time scale from [0, 5] sec.	142
7.11	Logarithm of Loran band electric field during a storm. Loran signals are visible around 1 second.	143
7.12	Range of 20 μs electric fields samples and rms of electric field voltage over 5 s and 15 min intervals during a storm on 13 June 2006 UTC.	145
8.1	Distribution of envelope noise measurements collected in Norman, Okla- homa over the summer of 2005. The blue and red curves show the log- normal distribution predicted by the ITU model above and below the me- dian value.	149
8.2	Comparison of ITU and long term RMS data taken from the 2005 data collection campaign. Time blocks are given in local time.	150
8.3	Radar plot of storm across Oklahoma on 13 June 2005 at 0205 UTC. Cour- tesy of Plymouth State Weather Center.	153
8.4	APD of instantaneous noise envelope voltages for a single five second data record during a storm. Data drops off for low amplitudes due to the dy- namic range of the receiver.	155
8.5	APD of instantaneous noise envelope voltages for a single five second data record during a quiet day. Data drops off for low amplitudes due to the dynamic range of the receiver.	156

8.6	Distribution of rms electric field strength (E_{rms}) and the voltage deviation (V_d) from collected data. 5 s interval data in blue and 15 min interval data in green.	158
8.7	V_d derived from noise estimates.	161
8.8	Resulting noise after non-linear processing.	162
8.9	Number of towers available for a position solution.	163
8.10	Comparison of traditional and improved coverage prediction charts.	164
9.1	APD of instantaneous noise envelope voltages for a single one minute data record during a storm with new Loran Receiver.	172
A.1	Visualization of the effect of interval on DFT and rms value.	186
C.1	Millington's Curves for Loran-C Field Intensity	198
C.2	Time history of electric field (E-field) strength of a nearby lightning return stroke.	199
C.3	Preta predicted electric field spectrum dB($\mu\text{V}/\text{m}/\text{Hz}$) versus distance with unadjusted return stroke data.	199
C.4	Preta predicted electric field spectrum dB($\mu\text{V}/\text{m}/\text{Hz}$) versus distance with return stroke data normalized to 35 kA current.	201
C.5	Preta predicted electric field spectrum dB($\mu\text{V}/\text{m}/\text{Hz}$) versus distance with return stroke data normalized to 35 kA current with 13 dB offset removed.	202
C.6	Measured envelope data (blue) and simulated response with filted flat plate data (red).	203

Chapter 1

Introduction

Always mindful of tight budget constraints, the Federal Aviation Administration (FAA) periodically investigates where it can reduce its overhead by eliminating obsolete or redundant systems. With an increase in aircraft users' reliance on the Global Positioning System (GPS), the FAA has considered reducing some of its older ground-based radio-navigation aids.

GPS is a system of satellites which provides navigation accuracies of approximately 10 meters. However, GPS satellites are distant, so their received signals are low in power and vulnerable to jamming. Therefore, a backup system to GPS is still required; in particular, the system should be capable of assisting an aircraft in making a non-precision approach (NPA) into an airport. Of the alternative systems available to provide a backup to GPS for NPA, Loran, a low-frequency radionavigation system, promises to be an exceptional candidate since its signal is present over the continental United States (CONUS).

Loran, in its current configuration, cannot meet the NPA requirements. However, the FAA is seeking the advice of experts in the Loran community to determine if, with reasonable upgrades, Loran can be made to meet the NPA requirements. If Loran cannot meet

the requirements, then the system will likely be decommissioned. If, however, Loran can meet the requirements, then it can provide a useful and cost-effective backup navigation system for aircraft. Moreover, Loran could replace other navigation aids and save the FAA tens of millions of dollars per year in maintenance costs while ensuring the continued safe navigation of aircraft, even in the event of a GPS outage.

While Loran has a number of weaknesses preventing it from meeting the NPA requirements, the impact of atmospheric noise on a Loran receiver is the most significant. Atmospheric noise is wide-band radio interference generated by discharges in the atmosphere such as lightning. Because atmospheric noise falls across the Loran band, such noise has the potential to produce blunders within a Loran receiver and induce position errors of 3,000 m or more. Such an error is unacceptable since NPA requires a position estimate with errors no greater than 556 m.

Therefore, this dissertation focuses on both improving the way atmospheric noise and mitigation techniques are modeled in a Loran receiver and accurately estimating the probabilities of such large position errors. Moreover, this dissertation is a crucial piece of a larger research program aimed at improving the Loran system and proving Loran capable of assisting aircraft in making non-precision approaches. With the improved noise model developed in this dissertation for a Loran receiver, a more precise coverage model is developed depicting the areas of the country where Loran can support NPA. This improved model shows an increase in Loran coverage across the continental United States (CONUS) which is more capable of meeting the NPA requirements than previously determined through traditional models.

To that end, this dissertation leads to three important conclusions:

1. Previous noise models were overly simplistic and therefore, inaccurate as to their

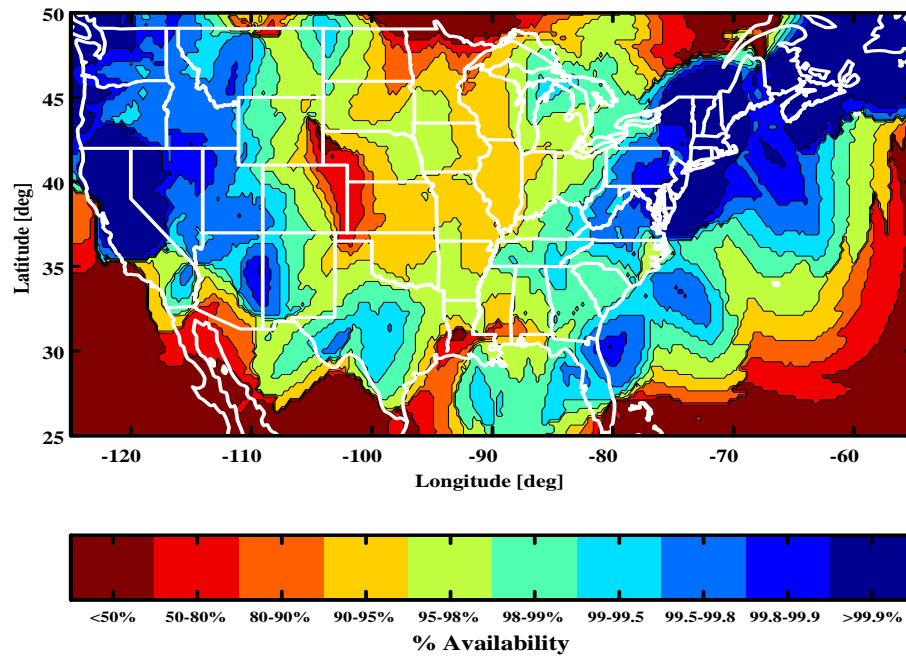
determination of the performance of Loran.

2. Traditional non-linear signal processing methods, such as clipping and hole-punching, effectively mitigated atmospheric noise in a Loran receiver.
3. Combining the new noise model with other proposed upgrades of the Loran system results in a system which is more likely to be certifiable for NPA.

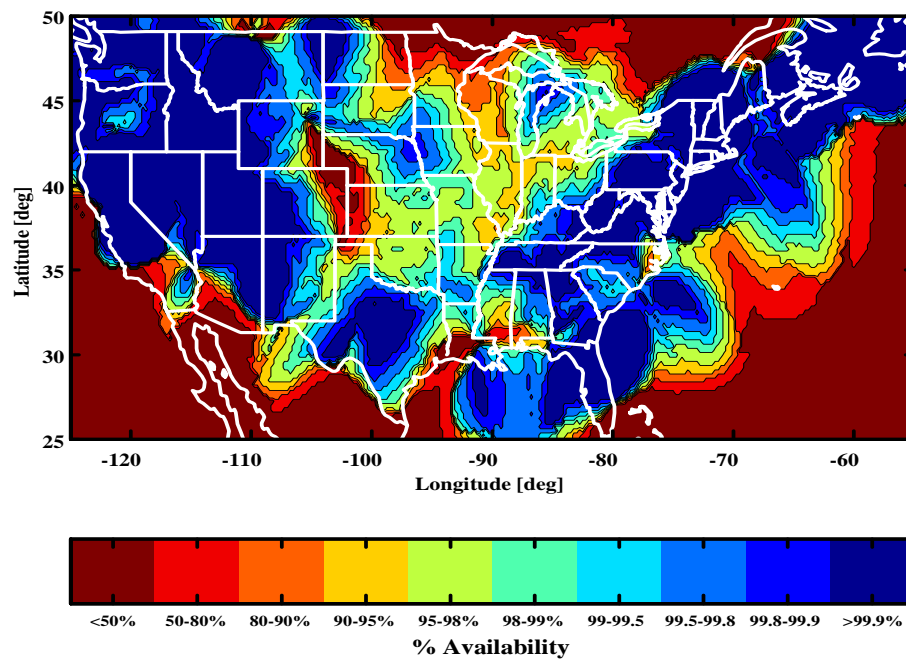
Loran availability is defined as the percentage of time which Loran can support an NPA into an airport at a given location. A view of the final conclusion is summed up by Figures 1.1 (a) and 1.1 (b). The first figure shows the anticipated Loran coverage using techniques prior to this work. While the details of these models will be more fully explained in the upcoming chapters, the dark blue areas indicate where Loran will be available for use more than 99.9% of the time during the worst-case storm conditions, a level the FAA deems the minimum acceptable for NPA. The red area indicates where the system is unavailable more than 50% of the time. The second figure shows that by applying the improved model described in this dissertation, and by giving a more accurate accounting of the noise and the processing techniques available, the coverage is increased by 153%. While the gap of coverage in the middle of the continent is unacceptable for NPA, it can be addressed by the placement of one to two additional Loran towers which may become part of proposed upgrades for Loran. A more thorough explanation of these diagrams is provided in Chapters 4 and 8.

1.1 Motivation

Because of the complexity involved in understanding the policy decisions which motivate this dissertation, the next few sections serve to provide background and an overview of the



(a) Traditional Model: Constant 12 dB Non-linear Processing Credit



(b) New Model: Pessimistic Non-linear Processing Credit

Figure 1.1: Comparison of traditional and improved coverage prediction charts.

key issues. These sections describe the current state of the nation's infrastructure for aircraft navigation and the issues surrounding the vulnerability of GPS. With the overriding issue being safety-of-life concerns, alternative navigation systems to GPS are briefly discussed and compared by cost. The final two sections cover the details required of a backup navigation system which enables aircraft to make non-precision approaches into airports.

1.1.1 National Airspace System

In an ideal world, one unjammable, ever operational, and exact navigation system would satisfy all positioning and timing needs. Unfortunately, such a system is not available today. To provide aircraft with accurate and safe navigation services, the Federal Aviation Administration (FAA) and the Department of Homeland Security use several radio-navigation aids working simultaneously in a complementary manner to form the National Airspace System (NAS) [1]. Radio-navigation aids of the NAS provide pilots with radio-frequency (RF) signals useful in determining aircraft position.

Six such navigation systems currently in use by aircraft are the Global Positioning System (GPS), VHF Omnidirectional Range (VOR) Beacon, Distance Measuring Equipment (DME), Tactical Air Navigation (TACAN), Non-Directional Beacon (NDB), and Loran. Each of these navigation aids range in cost both to the user and to the government for operations and maintenance. They have different levels of precision, allowing them to be used in different phases of flight, such as take-off, en route, approach, and landing.

Due to its precision and low cost, GPS has become one of the most popular means of navigating small aircraft, especially in general aviation for the private pilot. An understanding of the design, history, and rapid growth of GPS will motivate the rationale behind the study of this vital system.

1.1.2 The Global Positioning System

The Global Positioning System consists of 26 satellites or more orbiting Earth at an altitude of more than 20,200 km. These distant satellites, designed by the United States (US) military, carry a free civilian signal available almost anywhere in the world. As pictured in Figure 1.2, by receiving signals from four or more satellites overhead and using trilateration, a civilian receiver can determine its three dimensional position accurate to 10 m, 95% of the time, with precise timing down to 100 ns.

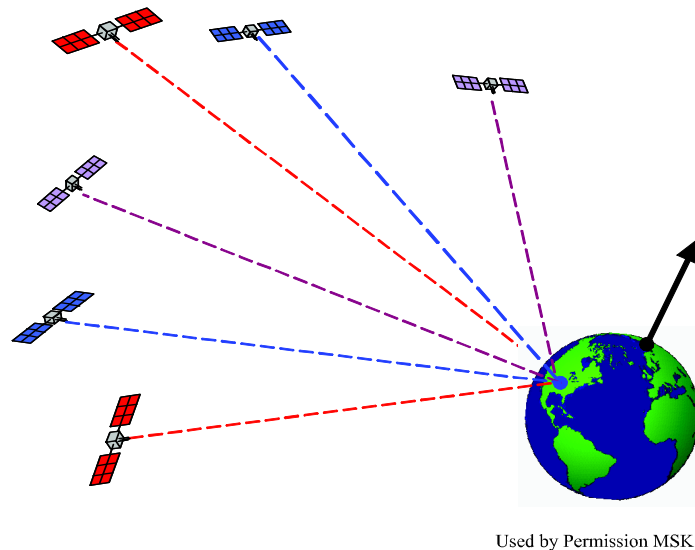


Figure 1.2: Earth-bound user receiving GPS signals from satellites.

The first GPS satellites launched in 1978. By 1992, the Federal Aviation Administration approved aviation use of GPS for en route travel. With accurate and low-cost receivers, GPS swept through the aviation market. Annual sales of aviation grade receivers from Garmin, one of the top manufacturers of GPS avionics, have gone from \$30.5 million in 1996 [2] to \$229.2 million in 2005 [3] reflecting pilots' comfort with the system. While

GPS has made navigation easier, legislators and policy makers have been forced to confront the consequences of becoming too dependent on any one system.

1.1.3 Need for a Backup System

Policy decisions concerning aviation and safety-of-life systems are complex. Providing an in-depth understanding of all of the facets that go into the decision is not the goal of this section, instead the more crucial aspects are highlighted.

The reliance on GPS for aviation is only part of a larger picture of the United States' increasing reliance on GPS for position and timing services. In 1998, then President Clinton issued Presidential Decision Directive 63 [4] noting that the United States' military advantage and economic power are intertwined and inexorably tied to numerous facets of the nation's infrastructure and information systems. He noted that areas such as banking, energy, finance, telecommunications, and transportation had increasingly become dependent on common technologies. Seeing that the economic and military power of the US relied on increasingly few systems, President Clinton urged both public and private agencies to

"...undertake a thorough evaluation of the vulnerability of the national transportation infrastructure that relies on the Global Positioning System."

The Directive goes on to state that,

"The Federal Aviation Administration shall develop and implement a comprehensive National Airspace System Security Program to protect the modernized NAS from information-based and other disruptions and attacks."

With such widespread dependence, President Clinton realized that a new threat had emerged since disruption of the common technologies would have a devastating impact

across multiple areas. Therefore, he called upon the various private and government agencies involved to review possible vulnerabilities from human error, natural disaster, and attack, and asked each agency to address these concerns and secure the United States' future military and economic power.

This directive spawned a number of investigative bodies that undertook the evaluation of risks to the nation's reliance on GPS. Two agencies whose reports epitomize the differing view points on GPS and its vulnerability are Johns Hopkins University Applied Physics Laboratory (JHU/APL) and the Volpe Center.

The first report by JHU/APL, prescribed making GPS the "sole-means" of navigation [5]. This report evaluated the risks to GPS from typical environmental forces, such as ionosphere and satellite geometry, and from malicious or intentional jamming sources. JHU/APL felt that all of the risks may be acceptably mitigated and that the likelihood of hostile jamming was small, especially if GPS systems are coupled with inertial measurement units or nulling antennas. Such a report inspired the notion of the removal of all other radio-navigation aids in favor of GPS in order to reduce NAS costs.

During its early growth period, GPS looked to be so promising that this sentiment transcended all modes of transportation. As a logical conclusion, the 2001 Federal Radio-navigation Plan suggested that the US could dismantle the other radio-navigation aids since they would be redundant over the next two decades and become solely reliant on GPS and instrument landing systems (ILSs) [6].

The second report, representing a contrasting viewpoint, was the Volpe Center's report on GPS Vulnerability [7]. The Volpe Center insisted that the economic and safety-of-life importance of aviation was much too important to entrust to one system alone and that some radio-navigation aids must be kept for backup.

The report stated that GPS's shortcoming stems from the low power nature of its signals. As mentioned previously, GPS satellites are more than 20,200 km above Earth's surface; by the time the signals from the satellites reach Earth, they are quite weak. Volpe determined that a 1 W jammer could deny a large metropolitan area use of civilian GPS.

Currently, jammers of comparable power are readily available. A quick search on the Internet will show that anyone with a credit card and a web browser can buy a 200 mW jammer, as shown in Figure 1.3, for about \$700.



Figure 1.3: 200 mW GPS jammer available on the web for \$700 in 2006.

Volpe published its report warning about the dangers of this over-reliance on a single system on September 10, 2001. While the foresight of this realization remained unrecognized for several years, their concerns were supported by President Bush's policy published in 2004, tasking the Secretaries of Homeland Security and Transportation with maintaining backup systems in the event of a GPS outage.

In an effort to address this question, the FAA established working groups to examine GPS backup systems capable of supporting non-precision aircraft approaches. The decision as to which navigation systems to keep is a hotly debated topic with a wide range of opinions. For example, Demoz Gebre-Egziabher in his thesis, proposed keeping only the

DME beacons [8]. The Aircraft Owners and Pilots Association (AOPA) have emphasized the need to keep the VOR and DME systems since these systems are prevalent in their constituents' aircraft. However, the Loran Integrity Performance Panel (LORIPP), a body composed of government, academic, and industrial researchers, propose making Loran the backup system to GPS. Note that the NDB system, which does not have a strong advocacy group, is being phased out as alternate landing approaches are developed and accepted throughout the country.

The next section compares the system costs for various navigation systems. Due to the ubiquitous nature of the Loran system's signals, Loran outperforms VORs and DMEs.

1.1.4 NAS Costs

Ranging in weight from one ton to 600 tons, aircraft vary as much in weight as they do in the tasks they perform. Similarly, with their variety of tasks, aircraft require different amounts of performance from their navigation systems. Navigation system requirements vary depending on whether the aircraft's cargo is passenger or freight, and on the type of weather conditions in which the aircraft is required to fly.

At the higher end of the weight spectrum are commercial passenger planes. Under such high demands and with a very high cost in the event of an accident, commercial airlines demand the most out of their navigation systems. In addition to using radio-navigation systems, these aircraft often install an inertial navigation system (INS). By measuring only the aircraft's acceleration and rotation rates, an INS can provide an aircraft's position accurate to one nautical mile after one hour of flight. These precision instruments cost over \$200,000 per aircraft. At that price, only large commercial aircraft and high end business jets can afford to carry such an instrument. Smaller regional jets, transport and cargo aircraft, and

general aviation rely on other radio-navigation systems which are less expensive.

Examples of more affordable radio-navigation systems for smaller aircraft are VOR and DME. For the small aircraft user, DME gives range information to the receiver, while a VOR provides bearing information. As shown in Figure 1.4, by using range and bearing, a pilot can position the aircraft horizontally. If coupled with a barometer, the pilot can obtain vertical positioning, and know the aircraft's position accurately enough to make a non-precision approach.




Figure 1.4: Depiction of VOR (bearing) and DME (range) systems.

These systems, however, all require direct line-of-site, so an aircraft must be in view of the radio-navigation aid in order to use it. Because of this line of site requirement, the National Airspace System contains many of these radio-navigation aids. In 2005, the FAA operated and maintained 973 VOR and 953 DME or TACAN sites. The cost in operations and maintenance alone were \$44.3M for VOR and \$24.7M for DME and TACAN [9]. This large amount of infrastructure makes these systems a target for cost-cutting measures.

With few towers, Loran is another terrestrial radio-navigation system which may be a reasonable alternative to VORs and DMEs. As shown in Chapter 2, Loran is a low-frequency system whose signals propagate along the ground. These groundwaves can travel a thousand nautical miles so only two dozen towers are needed to cover the CONUS and Alaska. Cost estimates for Loran range from \$20M-\$40M/year [10, 11], making it less costly than the combination of VOR and DME/TACAN system. However, Loran is currently only approved for en route travel between airports, not for non-precision approach. A result of this dissertation is an improved estimation of the impact of atmospheric noise on Loran when various noise mitigation techniques are used. This analysis will assist the LORIPP to validate Loran as a system capable of aiding aircraft in making non-precision approaches.

1.1.5 Aircraft Approach Requirements

This section reviews the classifications of approaches and the next section details the technical specifications for a non-precision approach. These sections serve mainly as background to reveal some of the difficulty in specifying requirements and the need for rigorous mathematical treatments to assure a high level of safety.

A recreational pilot has the luxury of waiting for good weather before flying, a commercial pilot does not. Navigation techniques vary based on weather conditions, resulting visibility, and intent of operations. In good weather and prior to GPS, a pilot relied on pilotage, the skill of determining position by spotting landmarks and features on the ground to navigate. Along with pilotage, the pilot used dead reckoning, the process of using wind, air speed, and time to estimate the aircraft's position [12]. With the advent of GPS, most pilots enter the coordinates of their destination and have a continuous navigation solution,

so a view out the window is not as critical as in the era of pilotage. However, when weather conditions cloud the view, pilots are forced to rely on a number of instruments mounted in their aircraft to aid them in navigating within a prescribed route in order to avoid air traffic. The required accuracy of these routes dictate the required accuracy for the instrument used for a particular stage of flight.

En route travel between airports requires navigation instruments to be accurate to one nautical mile since the geographic areas are large and the density of air traffic is low. As an aircraft approaches an airport and begins its terminal procedures, the position accuracy requirements stiffen depending on the type of approach the pilot is planning to execute.

The FAA divides approaches into two broad categories, precision and non-precision approaches [13]. A precision approach has glide slope information available. Precision systems include ILSs and the Wide Area Augmentation System-capable GPS receivers which can provide measurements of the aircraft's vertical deviation from a desired vertical profile or glide slope. A non-precision approach is an approach without vertical guidance; the pilot relies on the baro-altimeter for vertical positioning [13]. For each airport, the FAA publishes a separate non-precision approach procedure written for each navigation aid available at the airport. Common non-precision approach procedures are written for NDB, VOR, GPS and DME/DME.

In order to consolidate the number of charts required, the FAA has been producing new approach procedures based on the performance guaranteed by the flight management system instead of an approach based on a particular navigation aid. The FAA designates each flight management system by its overall accuracy.

A flight management system combines various navigation aids and provides the user with area navigation capabilities. Area navigation (RNAV) systems allow the user to fly

Required Navigation Performance 0.3 NM Specifications	
Monitor Limit	0.3 nautical miles (556 m)
Integrity	99.99999% / flight hr
Availability (minimum)	99.9%
Availability (target)	99.99%
Accuracy (95%)	307 m
Continuity (minimum)	99.9% over 150 seconds
Continuity (target)	99.99% over 150 seconds
Time to Alarm	10 seconds

Table 1.1: Required Navigation Performance 0.3 NM Specifications

between arbitrary waypoints denoted by longitude and latitude, rather than points defined by a land-based navigation aid. This additional flexibility leads to better use of the airspace.

One particular categorization for flight management systems is termed required navigation performance of 0.3 nautical miles (RNP 0.3). Any flight management system that can guarantee position solutions with sufficient confidence to 0.3 nautical miles may use the same RNP 0.3 charts for approach that the FAA has already developed for other navigation aids, including GPS. If Loran can be made to meet this requirement, pilots could use Loran to make RNP 0.3 non-precision approaches, thereby making Loran a more viable alternative as a backup navigation system to GPS.

1.2 Description of RNP 0.3

Currently, the International Civil Aviation Organization (ICAO) is working on consolidating all of the current standards to make an international definition for RNP 0.3. Quoting the definitions from [14–16], the FAA and the LORIPP define the demands of an RNP 0.3 system as meeting the criteria listed in Table 1.1.

These criteria vary from obvious to subtle in their description and interpretation. The

Monitor Limit will be explained after defining the other terms since it encompasses a number of these terms. Coverage, though not explicitly mentioned, also has a precise meaning which is described below. The International Civil Aviation Organization defines these terms as follows [directly quoted] [14]:

Accuracy – The degree of conformance between the estimated or measured position and/or the velocity of a platform at a given time and its true position or velocity. Radio-navigation performance accuracy is usually presented as a statistical measure of system error and is specified as:

1. Predictable. The accuracy of a position in relation to the geographic or geodetic coordinates of the earth.
2. Repeatable. The accuracy with which a user can return to a position whose coordinates have been measured at a previous time with the same navigation system.
3. Relative. The accuracy with which a user can determine one position relative to another position regardless of any error in their true positions.

Availability – An indication of the ability of the system to provide usable service within the specified coverage area and is defined as the portion of the time during which the system can be used for navigation during which reliable navigation information is presented to the flight crew, autopilot, or other system managing the flight of the aircraft.

Continuity (of function) – An assurance that, through a combination of sensors or equipment, guidance information permitting navigation to the appropriate level of RNP will continue to be provided for an acceptable period of time

after the loss of a sensor.

Coverage – The coverage provided by a radio-navigation system is that surface area or space volume in which the signals are adequate to permit the user to determine position to a specified level of accuracy. Coverage is influenced by system geometry, signal power levels, receiver sensitivity, atmospheric noise conditions and other factors that affect signal availability.

Integrity – The fraction of time that a system can provide timely warnings to users when the system should not be used for navigation.

Time to alarm – The maximum allowable elapsed time from the start of system failure (i.e., alarm limit) until the time that the integrity alarm is annunciated.

Containment Limit – A region about an aircraft's desired position, as determined by the airborne navigation system, which contains the true position of the aircraft to a probability of 99.999 percent [end of quoted section].

The Monitor Limit is one of the more nuanced terms and is an attempt to define RNP 0.3 in a way that would comply with [15] and be an extension of the Containment Limit. The Monitor Limit ties to the receiver's estimate of its own position errors. In using mathematical models, the receiver determines its Horizontal Protection Level, which is defined as follows:

Horizontal Protection Level (HPL) - The radius of a circle in the horizontal plane (the plane tangent to the WGS-84 ellipsoid), with its center being at the true position, which describes the region that is assured to contain the indicated horizontal position [17].

Part of the definition of the system's integrity is that the probability of true error being within HPL is typically 99.99999% per flight hour. For RNP 0.3, the Monitor Limit, or the maximum tolerable HPL is 0.3 NM or 556 m. Should the receiver determine that its errors could exceed the Monitor Limit, the receiver must notify the user that it is unavailable at that time. The Monitor Limit implies that while the receiver is on, the true errors can exceed the receiver's estimated errors only 1 time in 10,000,000 hours. This stringent requirement demands formal, methodical, and rigorous mathematical treatments of system vulnerabilities that lead to errors.

1.3 Prior Loran Research

The prior art in Loran pertaining to atmospheric noise or to the use of Loran as a navigation system falls short in its mathematical rigor and applicability in addressing the FAA's requirements for assisting aircraft in performing non-precision approaches. In broad terms, the prior art presented here covers the following areas: the use of Loran as a navigation system, the development and use of the *traditional* model to assess the performance of Loran in atmospheric noise, and the evaluation of atmospheric noise. A review of this work provides context for the placement of this dissertation in the on-going progress of Loran coverage estimation and prefaces the transition from the traditional model for atmospheric noise to the model presented in this dissertation.

1.3.1 Loran for Navigation

References [18–22] represent some of the earlier works in Loran which focused on the use of Loran as an independent or sole-means aircraft navigation system. In [20, 21], the

authors proposed the replacement of existing radio-navigation aids with Loran as a means of reducing the overall architecture of the NAS.

Rather than Loran as a sole-means system, other references, such as [23,24], looked to Loran as being used in conjunction with GPS as a more robust navigation solution than that of GPS or Loran alone. However, in these early years, the nation had yet to realize the need for a GPS backup system. Without strong incentive directing this effort, this prior research did not have the mathematical rigor to meet the requirements that the FAA eventually developed for a backup system. It was not until the Loran Integrity Performance Panel's report [10] that a comprehensive set of requirements were developed, enumerated, and addressed. Sections 1.1.5 and 1.2 describe the requirements that currently constitute a backup navigation system for non-precision approach, termed Required Navigation Performance of 0.3 nautical mile (RNP 0.3).

During the late 1980s and early 1990s, Loran as a sole-means system was becoming less of a possibility. Instead, Loran was seen as supportive of GPS or aided by GPS as in References [23,25–27]. The reason for this change is discussed in [25,28], where the authors state that atmospheric noise is the leading cause of Loran non-availability. It is Reference [25], which questions the accuracy and applicability of the ITU model and calls for further research into atmospheric noise in the Loran band. This dissertation is an answer to that call.

1.3.2 Traditional Coverage Model

From the early 1970s to the early 1990s, the method for estimating Loran coverage remained largely unchanged. The basis of the analysis, which defines the *traditional* model, relied on the use of the ITU (CCIR) model of atmospheric noise [29]. With a noise level

between 90% and 99.9%, corresponding to the desired level of availability, and a constant credit received for non-linear processing gain, references [18, 25, 30–33] describe the application of the traditional model and the resulting coverage. Some researchers, notably Feldman in [34] and Carroll and Weitzen in [28, 35, 36], sought to develop an alternative model which reflected the short-term statistical nature of atmospheric noise.

In addition to the issue of selecting an appropriate noise model, a number of early works also set standards as to the expected performance of a typical Loran receiver. In particular, many references quoted that a Loran receiver may acquire signals down to a signal-power to noise-power ratio (SNR) of 1:3, or approximately, -10 dB [18, 25, 26, 31, 37]. None of them justified the number, but simply stated this as fact. Some researchers felt more conservative values of -4 or -6 dB to be more appropriate [28, 30], mainly due to the poor nature of the definition of SNR. The appropriateness of SNR as a measure of noise is addressed extensively in this dissertation in Sections 5.4.2 and 7.7 as well as Appendix A. Some insight as to why -10 dB is an appropriate number can be found in Section 5.2, where this value is tied to the probability of wrong cycle select.

The traditional model developed in these references also assumed that some form of signal processing will exist in a Loran receiver to improve the SNR. Researchers have mainly offered constant credit for non-linear processing of highly impulsive atmospheric noise. This credit usually is in the range of 8-12 dB [25, 30]. The conversion from a constant credit to a more dynamic one based on noise is presented in this dissertation.

Some researchers have worked at identifying the performance of a non-linear processing technique called *hole-punching*, where data which exceed a given threshold are discarded from the estimation solution. This technique is also referred to as *censoring* or *editing* in the literature. Proponents of hole-punching claim performance gains of 10-40

dB in SNR, however, they caution its use since it is dependent on the threshold level chosen [18, 28, 32, 38].

An alternative to hole-punching is clipping. With clipping, incoming samples are limited to a given value, rather than discarded. Clipping proves to be almost as effective as hole-punching and less sensitive to the threshold setting. Gains of 12-13 dB are found in [34, 39] for frontal weather conditions.

While these gains are good, the researchers have shown that the gains are applicable only to the tracking statistics of a Loran receiver. References [25, 26], however, note that it is not the tracking which limits the performance of Loran, but the acquisition of the signal within the receiver that is the difficulty. While Fehlner in [18] also comments on the acquisition process, he only offers experimental data showing the improvement in acquisition. He does not support it with any in-depth analysis. Since the acquisition problem is what limits Loran, it is the focus of this dissertation rather than tracking.

1.3.3 Atmospheric Noise

It is the International Telecommunication Union's atmospheric noise model which formed the basis of the traditional model of Loran coverage; a thorough description of the model is in Section 3.2. The results of the initial coverage analysis based on the ITU model formulate this dissertation's problem statement as described in Chapter 4. However, ITU's model is not the only atmospheric noise model. Other models exist, particularly in the study of lightning and its spectrum rests in the field of atmospheric physics [40–47]. This work will become an alternative basis for comparing the results from the ITU model. The description of the atmospheric physics model may be found in Section 3.1. The perceived discrepancies between the atmospheric physics model and the ITU model will be developed

in Section 5.4.2 and Appendix A, while Section 7.7 reconciles these discrepancies as a matter of time scale of the observations.

Some researchers have also worked in developing their own models by taking atmospheric noise data in the Loran band. The seminal work performed by Feldman in [34] began to address issues of LF radio noise. While he did not publish his results of 100 kHz noise in his dissertation, he did show good correlation in three other bands at 14, 65, and 85 kHz. Carroll and Weitzen followed up in their research with some tests at various airports [27]. These efforts were on the time scales of a few weeks or a summer. Since none of the researchers attempted to refer their work back to a standard model, such as the ITU atmospheric noise model, they inspired the author to examine various noise models and determine if the ITU model and other atmospheric noise models developed by physicists can mutually support one another.

Other than the work presented in this dissertation, only Cutright et al. at Ohio University have pursued such an extensive campaign of data collection which spans from 2003 to 2006 [48–51]. Cutright’s work, however, focuses more on the effects of precipitation static build up on the aircraft and how this charge build-up affects the receiver. Additionally, his work concentrates on flight tests of receivers rather than capturing the extremes in weather conditions as in this dissertation.

By validating the ITU model for high amplitude noise and then extending the model to show that a relationship between noise impulsivity and noise amplitude exists, this dissertation more appropriately credits non-linear processing techniques. With gains proportional to the noise power, as shown by the curves of Chapter 8, impressive improvements in SNR are achieved.

1.4 Contributions

The prior art techniques of using the ITU model and constant processing credit produced the traditional Loran coverage model, which is more fully described in Chapter 4. As shown there, when applying what was considered standard practices, Loran's ability to provide RNP 0.3 coverage across the continental United States (CONUS) was poor. This poor performance was mostly based on the application of the standard ITU model to extreme power values while only providing a constant clipping credit.

Early in its outset, based on the standard practices of the time, the Loran Integrity Performance Panel (LORIPP) developed the algorithms for evaluating the performance of the traditional model for atmospheric noise and the estimation of Loran coverage. The results showed that Loran's ability to provide RNP 0.3 coverage across the CONUS was poor. As stated, this poor performance was mostly based on the extrapolation of the standard ITU model to extreme power values while only providing a constant clipping credit for processing gain. The results of this model were deemed incredulous, even though the technique had been used extensively in the past. Based on these initial results, this dissertation is written with the intent to better understand atmospheric noise and its respective models, and to understand the impact of noise on a Loran receiver. To this end, this dissertation makes contributions in the following areas:

- Improved mathematically modeled evaluation of Loran receiver acquisition errors (i.e., wrong cycle selection) accounting for envelope noise distribution.
- Validation and extension of an existing atmospheric noise model (ITU) for use by Loran for extreme values.

- Analysis of non-linear signal processing techniques in Loran receiver acquisition using the extended noise model.
- First generation of availability maps for Loran to support NPA landings (RNP 0.3) that accurately account for lightning noise.

These contributions are not just a stand-alone effort. They represent an important part of an overall project by the FAA and on behalf of the LORIPP, to show that Loran has the potential for being an adequate backup to GPS. As will be shown in Chapter 8, through more accurate accounting of noise and the gains due to signal processing, Loran's availability has improved across most of CONUS. Based on these results, Loran may become not only a backup system to GPS for small aircraft, but also a backup position and timing solution to a variety of users and modes of transportation.

1.5 Dissertation Outline

The outline of this dissertation is given graphically in Figure 1.5. To make navigating this work easier, review and development chapters are given in blue while the heart of the dissertation is in red. Chapter 2 reviews the history, design, vulnerabilities and limitations of Loran. Chapter 3 covers other major topic areas required for the study: Atmospheric Noise and Signal Processing. With the foundation laid in the background fields, Chapter 4 describes the specifics of how atmospheric noise affects Loran availability coverage and the shortcomings of applying the standard coverage techniques. Chapter 5 covers the development of the necessary mathematics to understand how non-linear signal processing techniques work in a Loran receiver. The design of the data collection equipment is given in Chapter 6. The background in the previous chapters helps to explain the contributions

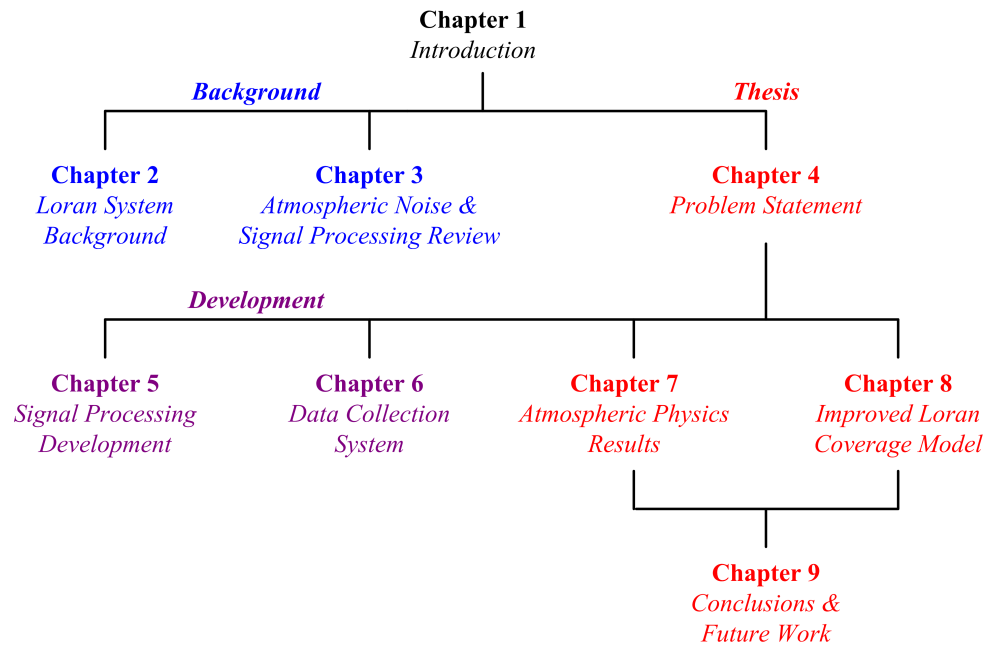


Figure 1.5: Dissertation chapter organization. Background and development sections are in blue. Thesis topics and results in red.

to atmospheric physics in Chapter 7 and the final description and results of the improved Loran coverage model of Chapter 8. The impact of this work is discussed in Chapter 9, as are recommendations for future work.

Throughout much of the background discussion, various specialty topics and areas of knowledge are required. Within the Appendices, additional information on Fourier transforms, non-linear signal processing, Loran receiver design, and antenna noise calculations are provided. Additionally, the calibration technique for the data collection equipment is described. A list of acronyms and symbols is provided in Appendix E for reference.

Chapter 2

Loran System Background

This chapter reviews some of the relevant history of Loran as well as details of the Loran system and typical receiver design. To understand the limitations on a receiver's ability to acquire and track Loran signals, an in-depth understanding of how the Loran signals are processed within a Loran receiver design is required.

2.1 History

A good history of Loran and other radio-navigation aids may be found in [52–54]. For completeness, this section covers some of that history and then covers the Loran system details which are prerequisites to understanding Loran receiver design.

In 1943, the United States military needed to develop an all-weather, 24 hour, navigation system capable of fleet deployment and bombing runs to support the World War II (WWII) effort. The US military with the help of Massachusetts Institute of Technology developed the first LOnG-RAnge Navigation system, Loran-A [52].

The system operated in the medium frequency band of 1850-1950 kHz and had a range of 600 NM and an accuracy of approximately 1 nautical mile [31]. The system proved itself in guiding planes and ships throughout the Pacific Theater in WWII [53].

Loran-A towers operated in pairs, and each tower in the pair emitted pulses. Upon reception of the pulses, a receiver would difference their times-of-arrival to produce a hyperbolic line of position. In later years, the military improved the timing of the signals and allowed groups of towers to work together in chains; this system was Loran-B.

In 1957, the military made further improvements, resulting in Loran-C, by changing the transmission frequency to 90-110 kHz, thus increasing the range of transmissions to 1,600-2,000 NM [52]. This is the current version of Loran supported by the Coast Guard for the public.

2.2 Loran System Configuration

The details of the signal structure of the Loran signal are important since they highlight the difference between Loran and many other forms of navigation and communication signals. The primary difference comes from Loran being a pulsed system rather than a continuous wave system. In a continuous wave system, data is a continuous function of time. Interference falling at any random time will be guaranteed to fall across a data symbol. With a pulsed data system like Loran, navigation data are only provided by the towers over short periods of time, so there is a finite probability that an interference event will not fall across navigation data.


Loran's signal properties and structure provide insight into how atmospheric noise affects the signal. An overview of the system will aid in this discussion. Beginning with the

overall system description, this section explores down to the pulse level details in order to reveal some of this insight.

Due to Earth's propagation characteristics, the low-frequency Loran waves travel long distances along the nap of the Earth and are called *groundwaves*. With groundwaves traveling over 1,000 NM, fewer towers are needed to provide coverage over CONUS than with other higher-frequency line-of-site navigation systems. VORs and DMEs are two line-of-site (LOS) systems which number nearly a thousand towers and yet still provide less coverage than Loran.

As shown in Figure 2.1, the current Loran-C system uses 24 towers spread across CONUS and Alaska and an additional six towers in Canada. These towers are capable of 350 kW to 1.4 MW of transmit power.

The towers are grouped into 11 chains across the US; a 12th chain is shared with Russia. A chain consists of a single Master tower and two to five Secondary towers. Figure 2.1 highlights the US West Coast chain whose Master (M), is located in Fallon, NV, and whose Secondary towers are in George, WA (W), Middletown, CA (X), and Searchlight, NV (Y).



*Courtesy Mitch Narins' SOIT Briefing

Figure 2.1: Loran towers across the US and Canada.

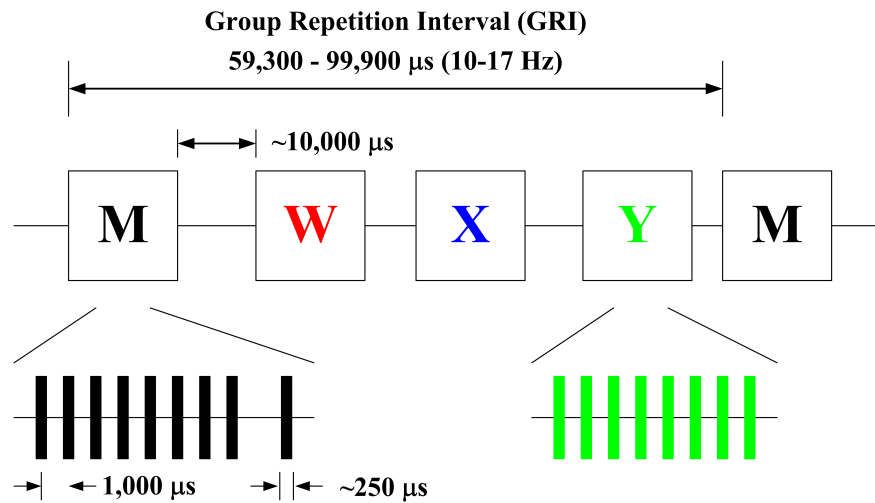


Figure 2.2: Loran GRI timing diagram. The Master station is denoted by M and the secondary towers by W, X, & Y.

Each tower uses three cesium clocks to form timing commands that drive the transmitters to emit a sequence of pulses at precisely timed intervals. Figure 2.2, illustrates the timing of both Master and Secondary stations. The sequence begins with the Master emitting a group of nine Loran pulses. The first eight pulses are spaced exactly one millisecond apart. The ninth pulse, used to identify the Master, is transmitted two milliseconds after the eighth pulse.

With the planned implementation of time-of-transmission control tying all transmitters to UTC, each Secondary will emit its group of eight pulses a fixed amount of time after the Master pulse. This fixed amount of time is called the *emission delay*. The emission delay is long enough such that a secondary group will fall in the same time slot as the master group or another secondary group for a receiver within the anticipated coverage region. Each of the Secondary towers transmits pulses in turn and then the sequence repeats itself starting again with the Master tower.

The group repetition interval (GRI), the time between sets of Master tower transmissions, uniquely identifies each chain. Designation of a chain is expressed as a multiple of ten microseconds (e.g., GRI 9960 = 99,600 μ s). Typical GRIs are between 59,300-99,900 μ s. In order to differentiate the Master pulses from the Secondaries, a low cross-correlation pulse phase encoding is placed on the pulses. The phase code repeats over two GRIs, A and B. The Loran Signal Specification [55] dictates the sequence of the signs of the transmitted pulses for the phase code. This sequence is shown in Table 2.1.

Group	Station															
	Master								Secondary							
A	+	+	-	-	+	-	+	-	+	+	+	+	+	-	-	+
B	+	-	-	+	+	+	+	+	-	+	-	+	-	+	+	-

Table 2.1: Loran Station Phase Codes

Getting deeper into the details of the signal, Figure 2.3 illustrates an ideal Loran pulse: a 100 kHz carrier wave amplitude modulated by a teardrop shaped envelope. The formula for the envelope, E , is given by:

$$E(t) = (t - \tau)^2 \exp(-2(t - \tau)/65)$$

where,

t is time in microseconds, and

τ is the envelope-to-cycle difference (ECD).

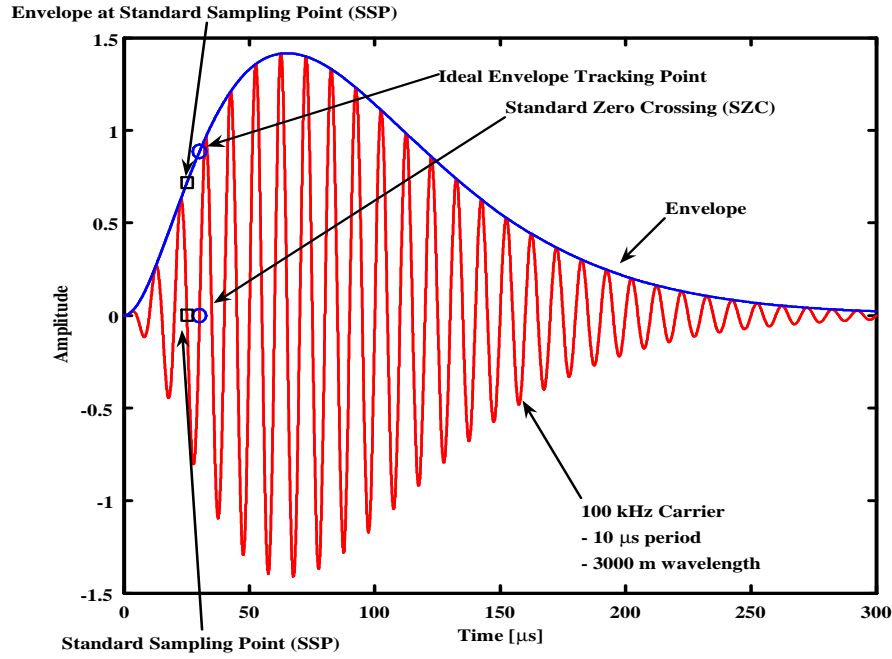


Figure 2.3: Idealized Loran pulse (red) and corresponding envelope (blue)

Typical values for the ECD range from $-5 \leq \tau \leq 5 \mu\text{s}$. The 100 kHz carrier wave corresponds to a $10 \mu\text{s}$ cycle period and the peak of the Loran pulse occurs at the $65 \mu\text{s}$ point.

While the peak of the pulse may seem to be the ideal tracking point, a receiver actually uses an earlier point because of skywave interference. When a Loran tower emits a pulse, it generates two types of signals. The first is the groundwave which propagates along the nap of the Earth. The tower also produces a direct line-of-site wave called a *skywave*. Due to the reflective properties of the ionosphere at low-frequencies, the ionosphere can reflect the direct Loran pulses back down to Earth. The skywave may be stronger in amplitude than the groundwave even though it travels a farther distance since the direct groundwave tends to have a higher path loss. The result is that the received skywave signal may interfere with the groundwave signal, but with a delayed effect. Skywave detection and mitigation efforts are an area of on going research in the Loran community [56].




Figure 2.4: Loran hyperbolas or locus of constant time difference, between Master and Secondary X, TD_{MX} (blue), and between Master and Secondary Y, TD_{MY} (red)

Due to the possible interference of skywaves, the time-of-arrival (TOA) of the Loran pulse is not measured at its peak, but by its third positive zero crossing. Called the standard zero crossing (SZC), it occurs $30 \mu\text{s}$ from the start of the pulse. By differencing the times-of-arrival of two towers, the receiver, in effect, creates a line of constant time difference (TD) that defines a hyperbolic line-of-position on which the receiver must lie. By using the TD between another pair of towers, the receiver calculates a second hyperbola, and the intersection of these two lines-of-position determines the receiver's location as seen in Figure 2.4. The calculation is solved iteratively using a linearized difference equation, in the manner of a 2D GPS solution [57].

Regardless of the solution method, Loran has an accuracy of 0.25 NM (460 m), 95% (2 drms) of the time [25, 31, 37, 58]. This accuracy is not quite absolute accuracy, since the Loran positions do not tie directly to a global reference system like WGS-84. In the design of the Loran system, monitor stations which track the signal's qualities are in place near harbors and channels where Loran accuracy is critical. In these areas, the timing

of Loran stations is adjusted by tower operators to keep the accuracy near the monitor to about 50 m [31, 58]. The 0.25 NM accuracy attributed to Loran reflects the stability of charts created with Loran readings and represents the maximum deviation in position as the user reaches the fringes of coverage. Most of the error is due to uncertainties in path propagation modeling which include natural time-varying effects and terrain. The diurnal and seasonal changes in propagation characteristics vary slowly over the course of a year, giving Loran its repeatable accuracy of 18-120 m over the course of several days. So that a user may return to the same spot within this distance over the course of a few days.

Summarizing this section: Loran is a low-frequency, long ranging, ground propagating, pulsed navigation system and Loran solutions yield accuracies on the order of 0.25 NM, with repeatable accuracies that can be much better. Section 3.1.2 describes how lightning occurs as pulsed, wide-spectrum noise. In order to see how lightning affects the receiver, the following section covers a basic Loran receiver design.

2.3 Loran Receiver Basics

2.3.1 Receiver Design

GPS receivers have been well described in the literature and text books [57, 59], and so receiver models are readily available. Due to a number of factors, primarily the desire to keep their design algorithms proprietary, companies and individuals have not widely published the intricacies of their Loran receivers. Most of the existing Loran receiver design papers center around the tracking loops such as [60–62]. While Reference [63] discusses a novel Master and Secondary tower identification technique. A few papers describe envelope tracking such as [64–68]. However, these techniques center around half-cycle peak

ratio tests and delay-and-add derivers to approximate the first derivative. This dissertation uses an envelope ratio test based on Peterson's Digital Down-Converting Software Receiver [69] built for the Coast Guard Academy. Since he has made his design available in the open literature, it makes an ideal baseline for the Loran community to use when comparing processing techniques.

While all of the details of the design are not necessary for comprehension of this dissertation, some of the basic design principles involved in processing the Loran envelope are useful. They reveal how lightning can introduce errors in the receiver's cycle selection algorithm resulting in 3,000 m position errors.

2.3.2 Envelope Processing

In general, a Loran receiver performs three distinct functions in regard to Loran signals: acquisition, tracking, and position calculation. Before processing a Loran signal, a receiver must first acquire it and it is on this function that the interference due to lightning has its strongest impact. Therefore, this dissertation focuses on acquisition rather than tracking or the position calculation.

Acquisition of a Loran signal begins with the receiver's analog front end. Figure 2.5 shows the schematic of a front end which is based on [69]. The first red pulse is a transmitted Loran pulse exciting the receiver's antenna. Next, the incoming signal is filtered, sampled, and mixed down to the in-phase and quadrature components which are used to form a complex envelope.

Thermal noise, which is Gaussian in distribution, can enter at the antenna and at the mixing process. This noise adds to each of the in-phase and quadrature components. In

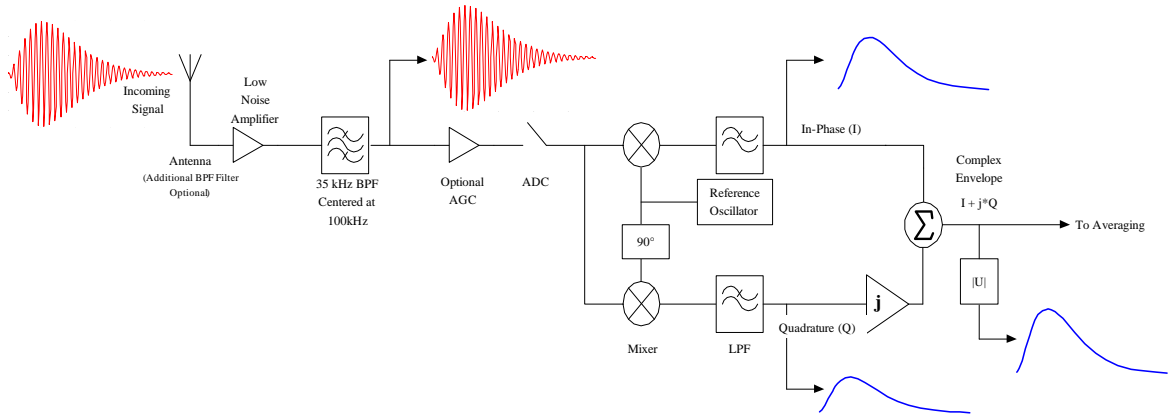


Figure 2.5: Loran front end and complex envelope generation.

taking the magnitude of the complex envelope, the envelope noise becomes a Rayleigh distribution. That is, a Rayleigh distribution is the root-sum-square of two Gaussian variates. The Rayleigh distribution is used in the APDs of Section 5.3.3 to denote “typical” thermal noise and all other noise distributions are measured relative to it.

Since noise interferes with the signal, a receiver averages the complex envelope data to increase the signal strength relative to the noise. Averaging improves the signal power to noise power ratio or signal-to-noise ratio (SNR).

After averaging the envelope, the receiver begins to coarsely estimate the time of arrival of the pulse by correlating the incoming signal with a reference pattern of pulses whose magnitudes were given previously in Table 2.1. Upon finding the approximate start of the sequence, all remaining pulses are averaged to one single pulse to maximize the SNR.

This coarse acquisition algorithm produces timing estimates on the order of $10 \mu\text{s}$. To transition to a finer estimation of its timing measurement, the receiver will need to use the phases of the samples in order to track Loran’s carrier. As mentioned in Section 2.2, the standard zero crossing (SZC), the third zero-crossing of the carrier, denotes the arrival time




Figure 2.6: Loran envelope and ratio test.

of the pulse. However, since the receiver has the envelope data and not the raw carrier, it will determine the correct zero crossing by the means of a ratio test.

To perform the ratio test, the receiver takes the ratio of two sample points $15 \mu\text{s}$ apart along the envelope. The ratio of the delayed sample over the prompt sample is shown in green in Figure 2.6. From the mathematical model of an ideal Loran pulse, an envelope ratio of 0.4 indicates the SZC. The receiver then uses the envelope samples to calculate the ratio along the entire envelope and pick the samples that bound the ideal ratio value of 0.4. Using the phase of the neighboring samples to linearly interpolate, the receiver can refine its estimate of the correct zero crossing, and hence the TOA of the pulse. The red dot at $30 \mu\text{s}$ in Figure 2.6 denotes the SZC.

Noise, however, can distort the envelope and skew the ratio result. This skewing may cause the receiver to select the wrong zero crossing as denoted by black crosses in Figure 2.6. A wrong cycle selection results in the range estimate being off by one or more Loran

cycles, each of which is $10\ \mu\text{s}$. This time offset results in a range error which is an integer multiple of 3,000 m. With such a blunder induced in the receiver by atmospheric noise, Loran will fail the accuracy requirement of 0.3 nautical miles (556 m).

Since the results of an incorrect cycle selection are serious, an understanding of the probability of such a blunder occurring is of great interest. Therefore, this dissertation develops a mathematical description of the impact of noise on the ratio test. Section 5.2 discusses the development of a mathematical model of wrong cycle selection for Gaussian noise as a baseline for evaluating the impulsive noise of lightning.

Traditionally, dealing with Loran envelope errors centered around envelope-to-cycle differences (ECD) [67, 68, 70, 71]. ECD is the difference between the phase of the Loran carrier and the time origin of the envelope waveform [31]. Zero ECD is defined when the $30\ \mu\text{s}$ point on the envelope is coincident with the third zero crossing. Changes in ECD occur when the Loran signal travels across a heterogeneous medium, such as land, whose propagation speed varies with frequency. Over such a medium, the Loran 100 kHz carrier wave propagates slower than the 20 kHz envelope increasing the ECD. A wrong cycle select can occur in the receiver if the ECD is greater than $5\ \mu\text{s}$. Similarly, if the envelope is distorted by noise such that the results of the envelope ratio test are skewed by more than $5\ \mu\text{s}$, then a cycle error can also occur. A $5\ \mu\text{s}$ variation in the TOA at acquisition will be the basis for the occurrence of a wrong cycle selection; the statistics of this event are in Section 5.2.

Before continuing to the next section, it is important to distinguish wrong cycle selection from other cycle related errors such as phase errors and cycle-slip. Wrong cycle select occurs when the receiver is in the processes of acquiring the signal. After acquisition, the receiver will begin tracking the Loran pulses and will do so by generating a replica of the

Loran carrier phase using a phase-locked loop. Should the receiver begin tracking a wrong cycle, there is no mechanism which would adjust the phase-locked loop. Instead, the loop would simply continue tracking the wrong cycle. This is what makes the wrong cycle selection such an insidious error.

However, should the correct cycle be identified, then the tracking loop will continue to track the correct cycle. Noise present during tracking will distort the phase measurements and cause the calculated position to jitter around some value. Should the noise get large enough that the jitter exceeds $5\ \mu\text{s}$, then a cycle error can occur. This would be called a cycle-slip since the receiver slipped off of the correct cycle and onto a wrong one. While a cycle-slip is a serious error, a receiver can detect and correct such an error since the position would tend to jump by 3,000 m. Since the error is correctable, cycle-slip is not as great a concern as wrong cycle selection.

2.4 Loran Signal Vulnerabilities

Since the system relies on the transmission and reception of radio waves, Loran is vulnerable to transmission, propagation, reception, and interference problems. Figure 2.7 depicts several different error mechanisms that affect a receiver's ability to determine its location. This dissertation focuses on atmospheric noise, one the primary sources of error in Loran. Other Loran vulnerabilities are studied by the various members of the LORIPP; the respective teams are listed in the figure.

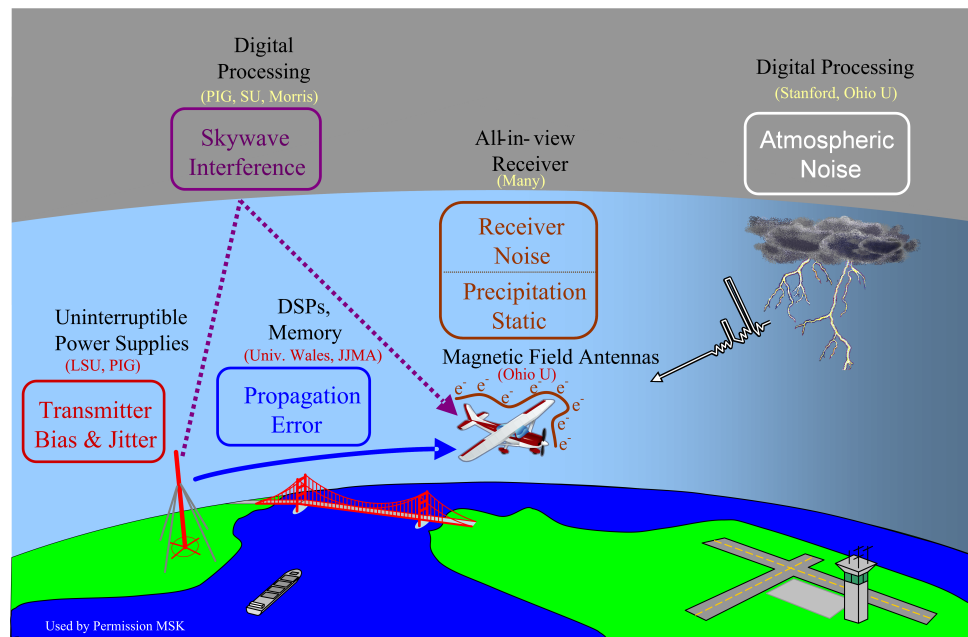


Figure 2.7: Loran vulnerabilities, the LORIPP members working on them, and the technologies used to overcome them.

2.4.1 Transmission and Propagation Vulnerabilities

The accuracy of estimating the position solutions is fundamentally tied to the accuracy of the time-of-arrival estimates of the Loran pulses. The original Loran towers used analog tube transmitters. These transmitters experienced biases and jitter which tower operators manually corrected for throughout the day. Such timing variations, undetectable by the receiver, produce range errors which manifest themselves as position estimation errors. By installing solid-state electronics which are less prone to these errors, the Loran Support Unit (LSU) and Peterson Integrated Geopositioning (PIG) have upgraded the towers' transmission system and eliminated most of the timing variations associated with the older analog tube technology.

During local thunderstorms, power outages would take the tower off air until the backup

generators came on line. To keep power active at the onset of an outage and before the generators turn on, LSU has upgraded most of the towers in the system to use battery powered uninterruptible power supplies.

TOA information is converted to range information through propagation models within the receiver. Daily and seasonal variations in soil conductivity brought on by material, temperature, moisture, and terrain affect the receiver's range estimate and induce range errors. These errors typically generate Loran position inaccuracies of less than 100 m. Through continued experimentation and by using modern processors, the University of Wales and John J. McMullen Associates (JJMA) are developing better models of Earth's propagation characteristics which will reduce Loran's range errors.

2.4.2 Reception Vulnerabilities

As an aircraft travels through clouds, dust particles in the air can cause charge to build up on the plane. Once the charge builds up to a high enough level, arcing between the aircraft and the air takes place all along the aircraft. The build up of charge, called precipitation static or P-static, produces these short duration, high-voltage sparks which generate wide band electrical interference at the antenna. Such interference can prevent reception of the Loran signal in receivers with an electrical field (E-field) antenna, while a magnetic (H-field) antenna is less susceptible to such interference [48–51]. E-Field antennas require the use of static wicks, which are small metal brushes, to be installed on the trailing edges of the airplane. These wicks help reduce charge build-up on the plane, but they have a tendency to deteriorate due to both weather and use over time. The unintentional loss of these wicks have resulted in much anecdotal evidence accumulated by pilots against Loran.

Ohio University leads the effort in designing H-field antennas as an alternative to using

E-field antennas. Through experimentation, the researchers at Ohio University have found that H-field antennas are less sensitive to P-static, do not require the addition of static wicks, and are therefore better for aircraft use. H-field antennas have the added advantage that they employ two loops in their design so the receiver gets additional gain over an E-field antenna by beam steering.

When a stationary E-field Loran antenna is grounded, corona discharges can occur when a strong storm is overhead. The discharge is similar in its effect to P-static, producing bursts of noise at the antenna. Corona was observed in several data records during some intense lightning storms.

2.4.3 Signal Interference and Noise

As described in Section 2.1, Loran towers pulse at different rates. Cross-rate interference occurs when pulses of two Loran towers arrive simultaneously at the receiver and interfere with one another. By using modern digital signal processors, the LORIPP team is developing cross-rate cancellation algorithms and all-in-view receivers capable of accommodating the difference in tower pulse rates.

The vulnerability that this dissertation addresses is Loran's susceptibility to noise generated by severe lightning storms and other atmospheric discharge. Characterized as wide-band noise, atmospheric discharges, such as lightning, interfere with Loran signals. In examining the impact of atmospheric noise on Loran coverage, models show that if not properly mitigated, atmospheric noise severely limits Loran's effectiveness as a viable NPA alternative. Only through understanding the noise process, its statistical properties, and its effects on a Loran receiver, can the problem with the mathematical confidence required to meet RNP 0.3 be mitigated. Figure 2.8 shows the time history of the absolute value of the




Figure 2.8: Absolute electric field envelope values in the Loran band. Note the return stroke interfering with the third Loran envelope pulse.

envelope electric fields within the Loran band. Note the lightning return stroke interfering with the third pulse in the sequence. This return stroke noise is generated from a stroke several hundred kilometers away. A nearby lightning return stroke may produce in-band noise hundreds of times stronger than even a nearby Loran station.

2.5 Loran Coverage Tool

All areas of vulnerability come into play when it comes to evaluating Loran. As part of the LORIPP's effort to determine Loran's coverage, Ben Peterson, Sherman Lo and the author have developed the Loran Coverage Tool to account for numerous error mechanisms encountered by Loran and therefore better determine Loran coverage across the US. This tool calculates how different errors affect Loran coverage. While beyond the scope of this

work, Lo et al. specify all of the algorithms in [72]. The traditional model, as described in Section 1.3.2, which used the ITU noise model 3.2 and a constant gain of 12 dB accounted for non-linear signal processing. Such an analysis followed the established practice of hole-punching in [18, 28, 32, 34, 38] and clipping or hard-limiting in [34, 39, 73].

Even though this model represented the state-of-the-art, it yielded a mediocre estimation of Loran's ability to provide RNP 0.3 coverage. Using this simple coverage tool, most of CONUS was found to have less than 90% availability, while the FAA wanted 99.9%. Availability was driven by the probability of the receiver tracking the wrong cycle. Atmospheric noise limits a receiver's ability to acquire the correct Loran cycle.

In reviewing the coverage estimation of RNP 0.3 of the traditional model, those in the LORIPP most familiar with Loran found the results too poor and inconsistent with their experiences. The conflict between predicted performance and experience motivated this dissertation and the exploration of the impact of atmospheric noise on Loran. As a result, the extension of the atmospheric noise model developed by the author for this dissertation, and the credit for non-linear processing became the model used in the current coverage tool.

Chapter 3

Atmospheric Noise and Signal Processing Review

Since atmospheric noise is the primary issue governing this dissertation, an introduction to the field of atmospheric physics is warranted. This chapter begins by describing the various types of discharges present in clouds and then focuses on the return stroke, the powerful current pulse seen in cloud-to-ground discharges. The electric and magnetic fields developed by the large current of the return stroke produce the noise prevalent at low frequencies. A description of the method used by atmospheric physicists to analyze the noise is presented since it will serve as a basis for the ITU model verification in Chapter 8.

The ITU atmospheric noise model was developed at a time when processing power was limited when compared with today. In particular, details pertaining to the method in which the ITU data were taken by the original researchers raise some questions as to the applicability of the model to the Loran band. This chapter will discuss the methodology used by ITU researchers to collect data and to generate the long and short-term noise statistics.

The ITU model is complex and the documentation has been reduced over the years,

so in this chapter the ITU model will be described and an example given on the model's use. The short-term statistical ITU model is also described since it is the basis of the signal processing algorithm presented in this dissertation.

The chapter concludes with an explanation of signal processing techniques. In particular, non-linear techniques such as clipping, hole-punching, and hard limiting are shown.

3.1 Atmospheric Physics

In 1749, Benjamin Franklin made his first scientific measurements of lightning. Since then, physicists have been attempting to quantify lightning and its effects on the world around us. References [40–42] give vivid descriptions and lengthy discussions of the colorful history behind mankind's pursuit of understanding this ephemeral subject. While a thorough treatment of atmospheric physics and history is beyond the scope of this work, a brief introduction motivates the discussion of lightning and its behavior.

Most people are familiar with lightning as the streak of light that they see flowing across the sky during a summertime thunderstorm. The flash seen comes from the flow of charge between the ground and the cloud. This current, which ionizes and heats the air, can exceed 80 kA, can take place over periods as short as 100 microseconds, and can move tens of Coulombs of charge between the cloud and the ground. Such violent discharges generate wide-band noise which may interfere with radio signals, and in particular with Loran.

While lightning strikes are the most visible of atmospheric discharges, they are not the only type of discharge. Atmospheric discharges are categorized based on the bodies in which charge is dissipated.

The first category pertains to the aforementioned flashes, they are known as cloud-to-ground (C2G) or ground-to-cloud (G2C) lightning, and account for 40% of all atmospheric

discharges. More than half of all discharges are within the cloud itself and hence are called intracloud (IC) lightning. The remaining 10% of discharges are made up of cloud-to-cloud (C2C) and cloud-to-air (C2A) discharges.

All of these atmospheric discharges generate broadband electrical noise through their rapid neutralization of charge. While the frequency content of each category varies, they all generate noise that ranges from 1 kHz to 20 MHz with a peak spectral component around 10 kHz. Since Loran is within the band, it is heavily affected by lightning. The details of these effects are found in Section 7.6 which discusses how lightning affects the Loran envelope during receiver acquisition.

3.1.1 Cloud-to-Ground Lightning

Despite intracloud discharges being more frequent than other types of discharges, researchers have produced the most work on cloud-to-ground lightning since it is the most readily seen, easiest to measure, and has the most immediate impact on safety and property. This dissertation leverages this work in developing an improved model of atmospheric noise. By understanding both the physical and the electrical aspects of lightning the statistical nature of the discharges and the power of the noise becomes apparent. Figure 3.1 reproduced from [42] depicts the time scales and the different processes involved in a negative cloud-to-ground lightning flash.

Currently, understanding the physical mechanisms for the build up of charge and the initiation of lightning discharges within a cloud is still an active area of research [74, 75]. Some researchers suspect charge accumulation begins with a complicated interplay between water, ice, and dust particles flowing in updrafts resulting in accumulate charged

layers within a cloud. Tens of Coulombs of negatively charged particles stratify to the bottom half of the cloud while a complement of positive charge remains at the top of the cloud. Fields on the order of 100 to 150 kV/m develop between these layers. Though these are strong fields, they are an order of magnitude less than what is required for dielectric breakdown of air.

There are two prevailing hypotheses which would explain the local increase in electric field which could initiate lighting [74]. The first is the hydrometeor mediated mechanism, where the hydrometeors, such as ice particles and water droplets, locally augment the field in excess of what is required for electrical breakdown.

The second hypothesis requires that electrons within oxygen and nitrogen molecules in the upper atmosphere are bumped into high energy states by cosmic rays that continuously bombard the Earth. These energetic electrons are accelerated by the fields established within the clouds and can reach relativistic speeds. Collision of these electrons with other electrons within the cloud can produce an avalanche effect.

Under either hypothesis, an initial pulse of moving charge called a *step leader*, marks the beginning process of the lightning stroke. The step leader is a moving packet of charge, which may exceed tens of Coulombs, that is lowered to the oppositely charged Earth in a series of erratic bursts. Depending on whether the leader has positive or negative charge, the categorization of cloud-to-ground lighting can be further refined, with the most common step leader being negative cloud-to-ground lightning.

The erratic bursts of the step leader move towards the ground in short $1\ \mu\text{s}$ pulses with approximately $50\ \mu\text{s}$ of dwell between bursts. The entire process takes roughly $19\ \text{ms}$. As the leader moves downward, it creates an ionized channel of air between the cloud and the ground allowing for easy charge flow later in the process.

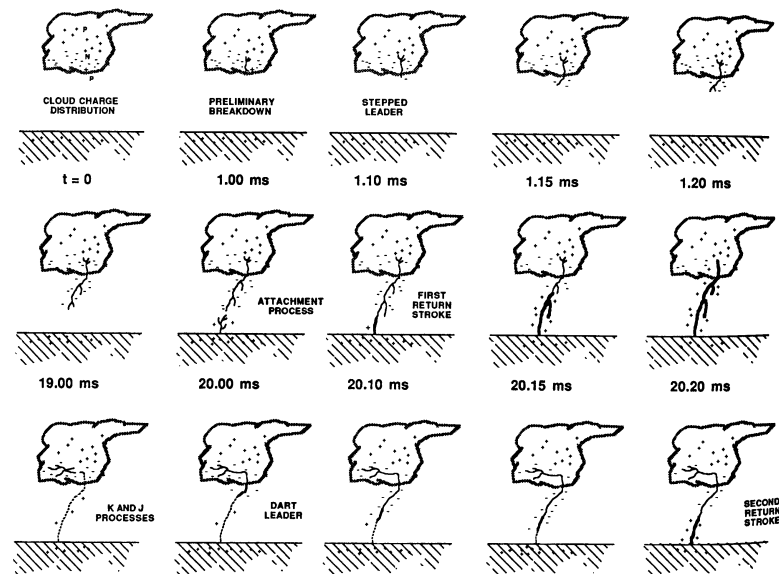


Figure 3.1: Component time history of negative cloud-to-ground (C2G) lightning flash from [42].

As the negatively charged leader approaches the ground, streams of positive charge called, *streamers*, come up from the ground. Because of the large amount of charge involved in the leader, any object connected to the ground can develop a streamer. Pointed objects like poles and metal objects may quickly develop streamers while blunt objects and non-conductors, such as trees, are slower to develop them.

Often multiple sources can generate streamers simultaneously. The first upward traveling streamer which contacts the downward traveling leader forms the attachment point and closes the circuit for lightning discharge; the resulting current bypasses all of the other streamers. Once the attachment is made, the first return stroke, or surge of current from the ground, travels up the ionized air channel left by the stepped leader. The return stroke current typically is on the order of 30 kA and peaks a few microseconds after initiation.

After the first return stroke has reached the cloud, there is a period of time where charge

within the cloud is rearranged. After the charge repositions itself within the cloud, it is possible to start a second leader down the original path of the stepped leader. Since the path is already ionized, the subsequent leaders quickly flow back down the path to ground and are called *dart leaders* which in turn induce additional return strokes. The subsequent return strokes are lower in magnitude than the first one. This repeated process of leaders and return strokes which make up the *lightning flash* give an appearance of flickering since the rate at which they occur is near the 30 Hz rate that human eyes can detect.

In examining the physical process of a return stroke, a pattern emerges. The lightning process proceeds in discrete steps and on a variety of time scales. Return strokes occur over hundreds of microseconds, step leaders occur at approximately one millisecond intervals, and the entire stroke may be spaced by 40 ms. Bursts of noise at various intervals are non-Gaussian in distribution. Because there are quiet periods between the bursts, a Loran signal may come through undisturbed during this quiescent interval. The bursty nature of lightning leads to time varying statistics and make the evaluation of power a challenge. The evaluation of different statistical models is the subject of Section 3.2 while the evaluation of the power of return strokes is developed in Section 7.7.

3.1.2 Return Stroke Electric Fields

Delving further into the details, this section examines the return stroke electric and magnetic fields. The high current and short duration of return strokes generate electric and magnetic fields which produce wide-band noise. The time history of these fields is a characteristic of lightning, and the time history of a lightning return stroke will be used in Section 7.1 to verify the calibration of the data collection system.

The electric field may be divided into three components that vary differently with r ,

the distance to the return stroke from the observer. The first two electric field components are the static and inductive fields that vary as $1/r^3$ and $1/r^2$, respectively. The third component, the radiative field, varies as $1/r$. So while the electric static and inductive fields are more dominant than the radiative field close to the return stroke, it is the radiative field that propagates the farthest. The magnetic fields are similar to the electric fields, with an inductive and radiative magnetic field which vary as $1/r^2$ and $1/r$, respectively; there are no static magnetic fields. At large distances, the electric field and magnetic fields are identical in shape and differ only in magnitude by the impedance of air. Figure 3.2, reproduced from [45], shows the electric field produced by return strokes at varying distances. After 50 km, the radiative field dominates the shape of the return stroke.

Numerous researchers have examined cloud-to-ground lightning [40–45, 47]. Uman in [42] summarizes much of the work done in the field of lightning and provides tables of characteristics such as number of strokes within a flash, interstroke arrival time, and peak return stroke current distributions.

Return Stroke Spectrum

In addition to time domain characteristics, Uman coauthored two papers on analyzing the spectrum of return strokes. The first paper, whose lead author was Serhan [43], measured the spectra of both the first and subsequent return strokes. The improper use of windowing the time domain data skewed Serhan's frequency results, resulting in a follow up paper [45] by Preta with corrections.

Using strip charts of wide-band electric field with frequency content from 1 kHz to 1.5 MHz, Preta digitized the data. He divided the data into return stroke records of $180\ \mu\text{s}$ in length. He then used a definition of the Discrete Fourier Transform (DFT) to define his

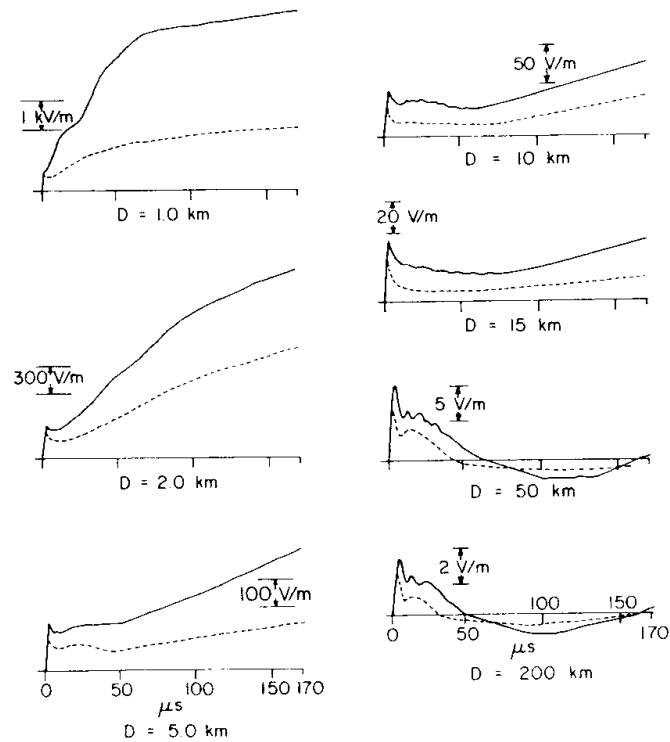


Figure 3.2: Typical electric fields for first (solid lines) and subsequent (dashed lines) return strokes at distances of 1, 2, 5, 10, 15, 50, and 200 km. Reproduction from Preta [45].

form of the electric field frequency spectrum of the return stroke. Appendix A describes the physical interpretation of the mathematics behind Preta's DFT definition. He defines the DFT as

$$E_{\text{preta}}(mf_{bin}) = T_{step} \sum_{n=0}^{N-1} \mathcal{E}(nT_{step}) e^{-2\pi mf_{bin} nT_{step}} \quad (3.1)$$

where,

$E_{\text{preta}}(mf_{bin})$ is Preta's Electric Field Frequency Spectrum [V / m / Hz] or [V s / m] at frequency mf_{bin} ,

f_{bin} is the Fourier frequency bin width [Hz],

T_{step} is the sample time step [s],

$\mathcal{E}(nT_{step})$, is the electric field value [V / m] at time nT_{step} ,

$m \in [0...N]$, is the frequency bin index, and

$n \in [0...N]$, is the sample index.

Figure 3.3 is a reproduction from Preta's paper showing the variation of electric field frequency spectrum of the first lightning return stroke versus frequency and parameterized by various return stroke differences. In this figure, the electric field is converted to dB by

$$E_{\text{preta dB}}(mf_{bin}) = 20 \log_{10}(E(mf_{bin}))$$

While Preta's data is in a format that appeals to atmospheric physicists, it is not as useful in obtaining a power spectrum of a return stroke. The difficulties in obtaining the average power from these data are reserved for Chapter 7 and Appendix A. Those sections

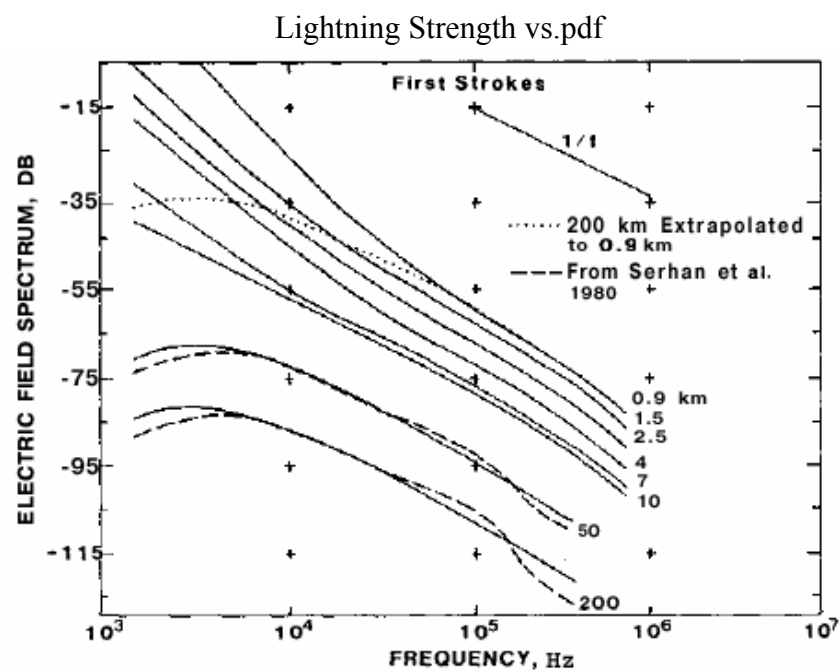


Figure 3.3: Frequency spectrums of electric fields for first return strokes. Reproduction from Preta [45].

discuss the details, extension, and mathematics of Preta's results and how they apply to determining a receiver's performance.

The difficulty in Preta's data format inspired one of the first contributions of this thesis which is to extend atmospheric noise data to the field of signal processing by transforming Preta's data to a form which will lend itself more readily to a power interpretation. This extension is developed in Section 7.5.

While this section pertained to the models and data obtained by atmospheric physicists, the next section discusses a model developed by radio engineers. To understand the difference in the manner in which physicists and engineers cast their data, the standard noise model is reviewed and in so doing, motivate this dissertation's second contribution, the bridging of these two seemingly disparate models.

3.2 ITU Atmospheric Noise Model

The International Telecommunications Union's model of atmospheric noise is a fundamental basis for this dissertation. The following sections cover the history of ITU summarized from [76] and the atmospheric noise model before concentrating on the ITU model's long and short-term statistics. In reviewing the history of how the measurements were made originally, issues are raised as to the applicability of the model, which this dissertation later addresses in Section 3.2.6. Since the current published version of the model is complex and with few examples of the model's use, an example is given in Section 3.2.4 along with the description of the model's statistics for the curious reader. A description of the model's use of the amplitude probability distribution is given since it becomes the basis for the non-linear processing gain described in Section 5.3.

3.2.1 ITU History

Since 1837, when Samuel Morse introduced the electric telegraph to the world, countries have struggled with international communications. Over the twenty years following Morse's invention, telegraphy spread throughout Europe with each country defining its own standard, resulting in cumbersome delays for international telegrams. To streamline and ease international communications, twenty European countries founded the International Telegraph Union (ITU) in Paris on May 17, 1865 whose charter was to draft the first international telegraphy standards.

Over the next sixty years, radio and wireless communication emerged. By 1927, new advances in communication prompted ITU to divide their oversight of the different communication technologies into three smaller committees. The Comité Consultatif International Des Radiocommunications (CCIR), or International Radio Consultative Committee, became responsible for the coordinating of technical studies, tests and measurements of radio communications, as well as for drawing up international standards for the use of the radio spectrum. Similarly, the other two committees, the International Telephone Consultative Committee (CCIF), and the International Telegraph Consultative Committee (CCIT), governed the standardization of telephone and telegraph technologies. In an effort to reflect the continued evolution of communication technologies, ITU changed its name from the International Telegraph Union to the International Telecommunication Union in 1932.

As telephony surpassed the use of telegraphy over the next thirty years, more organizational changes took place. By 1956 the CCIF and CCIT evolved into the International Telephone and Telegraph Consultative Committee (CCITT) to better direct the development of these technologies.

In the past fifteen years, cellular, satellite, and wireless technologies have continued

to grow as part of the world's communication infrastructure. To maintain its flexibility in addressing the ever changing needs of the telecommunication market, ITU was drastically overhauled in 1992. At that time ITU merged CCIR back into itself thus forming the group ITU-R, while ITU also absorbed CCITT. The change also brought about two new ITU branches, ITU-T to continue with telecommunications standards, and ITU-D for telecommunications development. Today, these three branches are the organizational structure of the ITU.

With such changes throughout ITU's history, tracking the development of projects that span many years is confusing. Of the subgroups mentioned, the original CCIR (and now ITU-R) has developed atmospheric noise models favored by statistical signal processing engineers. The details of the development of this model follow in the next section.

3.2.2 CCIR/ITU Atmospheric Noise Model

During its existence, CCIR took on the task of understanding the effects of lightning and atmospheric noise on low-frequency radio communications. To understand the impact of atmospheric noise and possible solar effects, CCIR carried out a measurement campaign from 1956-1970, approximately the length of one solar cycle, with the hopes of capturing any noise which may be related to the cycle. They measured the noise in a frequency band from 10 kHz to 20 MHz.

Data collection efforts began when CCIR established 16 measurement stations around the globe, and staffed them with researchers from national laboratories, universities, and Army research laboratories. Figure 3.4 shows the locations of these 16 sites.

The Committee documented their results in CCIR-322, Characteristics and Applications of Atmospheric Radio Noise Data [77]. Over the next two decades, CCIR continued to update and revise their model, and by 1986 they released the third revision of the document incorporating extra data and corrections; this updated model became known as CCIR-322-3, and was the last version edited under their auspices.

With the restructuring of ITU in 1992, came another revision of CCIR-322. Stripping out most of the historical content of the document, adding extra information on man-made noise, and revising the bandwidth performance conversion graphs, ITU re-released CCIR-322 as Recommendation ITU-R P.372-7, Radio Noise, the current incarnation of the atmospheric noise model.

The details of the system's construction and operation have been omitted from the current published version of the model. Yet, this history has been recovered and recorded here by reviewing the original CCIR-322-3 documents and by interviewing a researcher named Robert Matheson, who collected data at the Boulder site beginning in the 1950's. The history of the details of the measurement system aid in understanding what are potential short-comings and pitfalls are encountered in using the ITU atmospheric noise model.

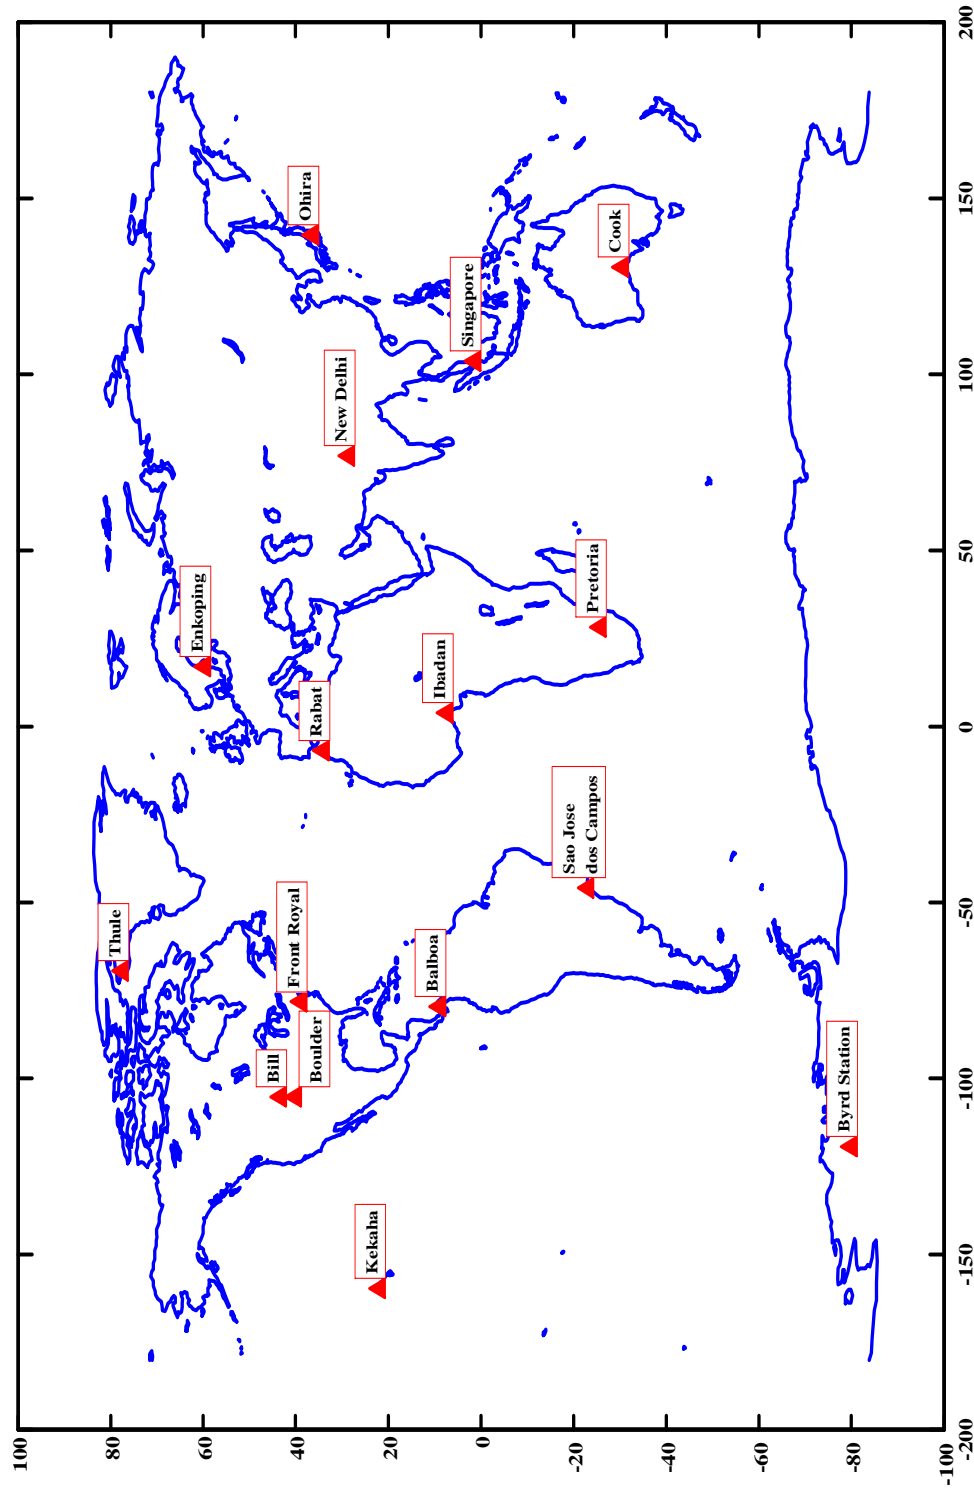


Figure 3.4: Placement of the original CCIR/ITU testing sites.

3.2.3 ITU Measurement System Details

The details of the construction and operation of ITU's noise measurement apparatus provide insight as to the interpretation of the atmospheric noise model. It is this insight which guided the development of this dissertation and led to the extension of the ITU the model. In particular, by seeing the method by which researchers at ITU collected, suggests that improvements could be made.

For the United States, the National Telecommunications and Information Administration (NTIA) in Boulder, CO, under the auspices of the late Arthur Spaulding, collected the noise data for CCIR. While Spaulding led the project, Robert Matheson, a high school student who later became an electrical engineer and employee of NTIA, assisted Spaulding in the data collection efforts. The author interviewed Matheson who described the Atmospheric Radio Noise Model 2 receiver (ARN-2) and the methodology used in collecting the noise data.

As for the physical size of the unit, the ARN-2 used a 3 m high antenna mounted 10 ft off the ground with 50-60 ft of cables creating a circular ground plane. Composed of vacuum tube technology, the receiver measured noise at the RF frequencies of 13 kHz, 110 kHz, 250 kHz, 500 kHz, 2.5 MHz, 5 MHz, 10 MHz, and 20 MHz. Above 13 kHz, the receiver first mixed down the desired frequency band to an intermediate frequency (IF) of 13 kHz before passing the subsequent signal through a 200 Hz bandpass filter.

With the bandwidth of the noise data being only 200 Hz, CCIR and later ITU, provided charts to convert from the 200 Hz noise data into wider bandwidths. The Stanford Loran Receiver (SLR), described in Section 6.3, was used to collect most of the data for this dissertation. The SLR used a 35 kHz wide bandpass filter centered at 100 kHz, thus, the bandwidth ratio between the ITU system and the SLR is 175. The bandwidth ratio was used

to scale the electric field noise as well as the voltage deviation for the Stanford receiver, however, Section 3.2.6, discusses some of the possible problems involved with applying the ITU data to the receiver design.

The scientists who designed the data collection system aimed to obtain the root-mean-square (rms) noise envelope voltages, or the noise power, by measuring the square of the voltage (V^2) for the antenna. The large amplitude variations associated with lightning would saturate any equipment they had if they attempted to measure the square of the voltage directly. Instead of the direct method, the scientists measured the voltage from the second harmonic of the mixer since it correlated well to V^2 which is proportional to power. By measuring the mixer's second harmonic through an attenuator with 2 dB steps, they achieved 100 dB of dynamic range before the 200 Hz bandpass filter. After all of the filtering, the final system had 80 dB of dynamic range, with a ceiling of 1500 V which could capture the wide-ranging rms noise envelope voltages.

The ITU scientists chose to measure the rms noise envelope voltage rather than the instantaneous noise voltages since the analog-to-digital converters and storage technology at the time limited the direct measurement of instantaneous voltage to only short periods of time. In order to smooth their noise envelope measurements, the scientists averaged the receiver's output in an analog fashion to calculate the rms voltage. Since the averaging filter had a time constant that was greater than one minute, they devised a method to measure the rms voltage, or equivalently, power, over a 15 minute time interval.

Using five minute segments, the scientists divided the time interval into three separate phases, each being defined by the treatment of the attenuator in the circuit. During the first five minutes of the interval, the attenuator would automatically adjust the rms output of the filter's second harmonic to match a reference. The scientists would then measure

the attenuator value to determine the noise level. During this phase, the attenuator would be allowed to increase or decrease in value to keep the filtered rms level matched to the reference. Over the next five minute phase, the attenuator would be held constant which kept the gain to the input of the filter constant. Finally, during the last 5 minute phase of the 15 minute interval, the attenuator was only allowed to go down, i.e., attenuate less. This final value was used as the power of the noise over the 15 minute interval.

Through a trial and error process, the ITU scientists deemed this technique the best at averaging the rms noise envelope voltages. At the end of the 15 minutes, they recorded the attenuator setting in 2 dB increments. Paper charts tracked the position of the attenuator over the entire fifteen minute interval, but these records are no longer available.

With two independent ARN-2s, each taking 15 minutes to calculate noise power, the scientists measured the “hourly” values of noise at each of the eight given frequencies from 13 kHz to 20 MHz. Since atmospheric noise is a non-stationary process, the scientists grouped the hourly noise statistics into four 90-day seasons and six 4-hour *time blocks*. In this manner, each of the 24 time blocks had its own set of statistics to describe noise and accounted for the noise intensity variation that occurs with season and local time of day.

3.2.4 Noise Factor Statistics

After collecting years worth of rms noise envelope voltages, the ITU scientists calculated the statistics of the external antenna noise factor, F_a , the power received through a loss-free antenna for each time block and each season. F_a is bandwidth independent, since increasing a receiver’s bandwidth, increases the noise entering the receiver, keeping the ratio of the two constant. By interpolating values between the respective sites, the scientists developed

global maps of F_a . Figure 3.5 shows the F_a values for the summer season and the 2000-2400 h Time Block.

Through their research, the scientists found that the distribution of F_a followed a split log-normal distribution. Log-normal means, if F_a was given in decibels, its distribution would be approximately normal, with the distribution centered at the mean noise factor, F_{am} . However, they also found that values lower than F_{am} , would have one variance, while those above the median would have a separate variance. Due to the nature of normal distributions this introduces a discontinuity or “split” of the distribution at the median value. The solid blue and red lines of Figure 3.6 show an example of such a split log-normal distribution.

Since their data spanned several years, the scientists could give variances upon the different estimated values, such as $\sigma_{F_{am}}$, the variance on F_{am} ; D_l and D_u , the lower and upper deciles of F_a ; and σ_{D_l} and σ_{D_u} , the variance on D_l and D_u , respectively, which accounted for the yearly variations.

These random variables and their respective variances are used to calculate the noise power for different percentages of noise. To illustrate how these variances on top of variances are used, the following design example is given. Suppose a receiver is required to work with 99% confidence when the noise factor achieves 95% of its expected strength. This also implies that the receiver is available 95% of the time with 99% confidence.

Since the ITU model gives the noise factors as a median value, an adjustment must be made to attain the 95% level. ITU provides, not the variance of the distribution, but the 90% noise factor with 50% confidence, D_u . By adding D_u to F_{am} , as shown in. Figure 3.6, F_{am} becomes the median 90% F_a value with 50% confidence. To convert F_{am} to an arbitrary percentage, P , the following formula is used:

$$D = D_u \cdot \text{norminv}(P/100, 0, 1) / \text{norminv}(90/100, 0, 1)$$

$$F_{a,P} = F_{am} + D$$

where,

P is the desired noise level above 50%

$\text{norminv}(X, 0, 1)$ is the inverse of a zero-mean, unit variance Gaussian cumulative distribution function for fraction X , and

$F_{a,P}$ is the noise factor at the desired percentage.

For this example, the 95% F_a value with 50% confidence is desired. Therefore, D is calculated from above for $P = 95\%$ and added to F_{am} resulting in the median 95% F_a value with 50% confidence, $F_{a\ 95\% w/50\%}$.

If the receiver circuit just met the desired performance requirements at a median 95% value, or for median 95% availability, then this circuit would fail half of the time for any given year's 95% noise level. This is due to the definition of median, that half of any year's 95% noise level will exceed the median 95% value half of the time. Therefore, to increase confidence in the circuit working over many years at the yearly median 95% noise level, the noise estimate must be increased to account the yearly variation.

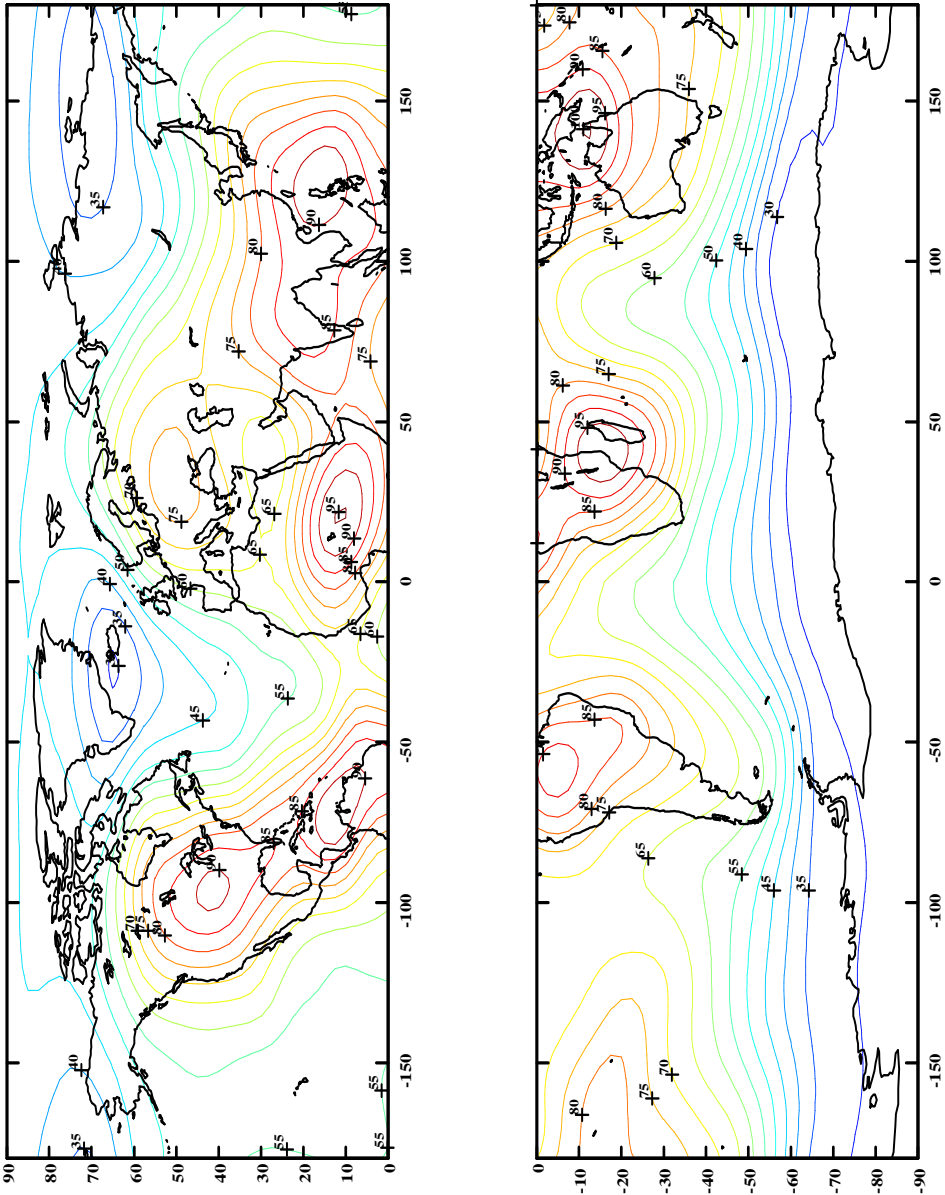


Figure 3.5: Median summer noise values, F_{am} , in dB above KT_0B at 2000 h




Figure 3.6: Log-normal F_a distribution with adjustments due to uncertainties σ_m and σ_D .

K , a factor which is multiplied by the root-sum-square (rss) of the standard deviation of each uncertainty is introduced to account for the yearly variation. K comes into the estimation of the median noise strength and is chosen to reflect the desired confidence level, under the assumption that the errors are normally distributed. The rss of the values are used since the noise uncertainties are assumed to be independent of one another.

Figure 3.6 shows how the uncertainties in the median noise value, and uncertainty in its variance, affect the noise distribution. The solid red and blue curves show the upper and lower log-normal distributions of F_a , respectively. Since F_{am} is a random variable, its variance, $\sigma_{F_{am}}$, shifts the distribution. Adding $\sigma_{F_{am}}$ to the median produces the corresponding dashed curves which would account for an 84% confidence level in the prediction of the median value of F_a , but with only a 50% confidence in D_u .

D_u , the adjustment to F_a to get the 90% value, is also a random variable, so its variance, σ_{D_u} , affects the spread of the overall noise distribution. By adding σ_{D_u} to the original variance, the dotted red line, a wider distribution from the former distribution is found. This may be interpreted as the distribution with a 50% confidence level in the median F_a .

value, but with an 84% confidence in D_u .

To improve the confidence bound of the receiver to 95%, then accounting for both uncertainties in D_u and F_{am} must occur. Since only, $\sigma_{F_{am}}$ and σ_{D_u} are considered,

$$K\sqrt{(\sigma_{F_{am}}^2 + \sigma_{D_u}^2)}$$

can be added to the 90% F_a value to increase the confidence. For a 99% confidence level, $K = 2.33$, and if typical sigma values are $\sigma_m = 3$ and $\sigma_D = 4$, then

$$K\sqrt{(\sigma_{F_{am}}^2 + \sigma_{D_u}^2)} = 2.33 \cdot 5 = 8.22 \text{ dB}$$

must be added to F_a to get the 99% confidence value.

By using these statistics, F_a for various percentages and with various levels of confidence can be calculated. Given the derivation in Appendix B, F_a , the power received through a loss-free antenna, is converted to E_{rms} , the anticipated rms noise envelope voltage received for a monopole antenna above an ideal ground plane. Thus these cumbersome steps enable the calculation of the rms atmospheric noise strength for a given percentage of time during any season and time of day for the data collection system or any arbitrary antenna. Figure 3.7 shows the application of the ITU data by estimating the worst-case median 99.9% noise level with 50% confidence for a Loran antenna across CONUS; the worst-case being over all time blocks and seasons.

F_a is a long-term rms noise value since it corresponds to a period of 15 minutes to 1 hour. While long-term statistics are important, the distribution of the instantaneous noise dictates the processing gain made available by non-linear processing of the received signal. The next section addresses how to model the short-term instantaneous noise envelope

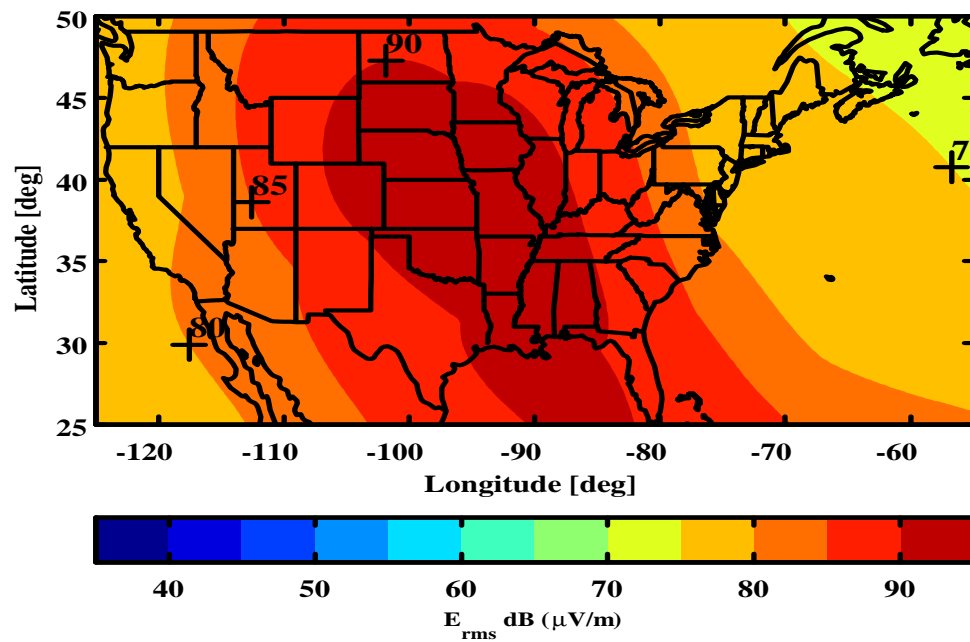


Figure 3.7: Median 99.9% rms envelope noise voltage, E_{rms} , with 50% confidence taken over all hours and seasons.

values.

3.2.5 Instantaneous Noise Envelope Measurements

Amplitude Probability Distribution (APD)

Section 3.2.3, described how ITU scientists took their measurements of long-term noise data. In their analog averaging over three five-minute periods, there appear to be some possible short-comings. Due to their method of limiting the attenuator to only increase over the first five-minute period, a lightning stroke that comes in during the last 10 minutes of the 15 minute measurement interval would be completely ignored. Additionally, the energy of a stroke that arrives in the first five-minutes would have at least 10 minutes of decay time before the final value was recorded.

Clearly, this method smears the power of individual strokes when calculating the rms value of the 15 minute interval. To get more detailed statistics for the interval, the scientists also took some “high-speed” data at 400 Hz. This sample rate was sufficient to fulfill the Nyquist criteria since the ARN-2 mixed all frequencies to an IF of 13 kHz and passed the signal through a 200 Hz bandpass filter.

With these high-speed data samples, the scientists produced amplitude probability distributions (APDs) of the instantaneous noise envelope voltages for numerous 15 minute intervals. The APD relates to the more commonly recognized cumulative distribution function (CDF) by

$$APD = 1 - CDF$$

However, there are differences in how the APD is typically drawn compared to the CDF.

First, the APD is normalized to the rms value of the interval. Also, the axes of the APD are flipped from that of the CDF so that the x-axis is the percentage of time for which the y-axis value, A_0 , will be exceeded relative to the interval's rms value.

Furthermore, the x-axis is scaled by $\frac{1}{2} \log_{10}[-\ln(P[E_{rms} > A_0])]$ [78], where $P[E_{rms} > A_0]$ is the probability that the envelope voltage, E_{rms} , will exceed A_0 . This scaling makes Rayleigh distributed data appear as a straight negatively sloped line, and therefore indicates the extent that the sampled noise envelope voltage follows a Rayleigh distribution. A more impulsive distribution appears as a steeper negative slope when using this scaling. Figure 3.8 is an excerpt from [29] showing APDs by their voltage deviation, V_d .

Voltage Deviation (V_d)

In taking a large number of APDs, correlations emerged from ITU's data collection effort. First, the researchers found that the data set's average, rms, and logarithmic mean governed the shape of the APD. They also found that the logarithmic and arithmetic mean correlated well to one other. Therefore, since only the rms and average values were independent, the researchers could parameterize the APD distributions by the voltage deviation.

The voltage deviation, V_d , is the ratio of the rms envelope voltage to the mean envelope voltage. Since the envelope voltages are all positive values and therefore the log of this ratio exists, V_d is defined as

$$V_d = 20 \log_{10} \left(\frac{\text{the root-mean-square noise envelope voltage}}{\text{the mean noise envelope voltage}} \right)$$

V_d gives a measure of the impulsivity of noise, where noise which is more impulsive has a




Figure 3.8: Reproduction of the APD curves parameterized by V_d from ITU P.372-7 [29].

higher V_d . To illustrate this, take the following two sequences of numbers as an example.

$$\{1, 1, 1, 1, 1, 1, 1, 1, 1, 1\}$$

$$\{1, 1, 1, 1, 1, 1, 1, 1, 1, 10\}$$

The first sequence has a mean and rms of 1, while the second has a mean of 1.9 and an rms of 3.3. The single large value of the second sequence will more than triple its rms voltage but not quite double its average voltage, thereby increasing V_d and demonstrating that the more impulsive noise has a larger V_d .

Since the rms voltage increases quickly given a few large values, noise which on average may be of tolerable power, can have an interval rms voltage that is too high if some large, but infrequent impulsive noise comes through. Section 5.4.3 provides examples illustrating how voltage deviation relates to interval power.

ITU Amplitude Probability Distributions

Engineers typically model noise as Gaussian, since the mathematics pertaining to Gaussian models are tractable. As shown in Section 2.3.2, Gaussian instantaneous noise voltages have a Rayleigh distributed envelope voltage. Since a Rayleigh distribution's mean and variance are analytically calculable, the Rayleigh distribution is found to have a $V_d = 1.049$. Because of the scaling of the APD, the Rayleigh distribution is a straight line as shown in Figure 3.8. The other curves show how more impulsive distributions, with larger V_d , diverge from the Rayleigh noise case and are shown to be more steep.

As mentioned previously, the APD is normalized relative to the rms noise envelope voltage of the interval, which shows how often the rms value is actually achieved. For example, for a data record with a V_d of 4, the corresponding distribution shows that the rms

noise envelope voltage is only exceeded about 16% of the time, while for 60% of the time, the noise is more than 10 dB below the rms value. Therefore, describing impulsive noise by its rms voltage is of limited value since the rms voltage is only achieved a small percentage of the time.

Furthermore, due to the Rayleigh scaling of the x-axis, an even more important observation regarding the behavior of the distribution can be made. Again, using the APD associated with $V_d = 4$, the right end point of the curve at 99% up until approximately the 60% mark is parallel to the Rayleigh curve. This implies that 40% of the time the noise follows a Rayleigh distribution, while 60% of the time the noise is more impulsive than Rayleigh. This powerful observation will be the basis of the signal processing algorithm of Section 5.3 because it leads to logically dividing the noise into two categories: noise which is Rayleigh in distribution and noise which is more impulsive or non-Rayleigh in distribution.

3.2.6 ITU Model Issues

In reviewing the ITU model's background, possible short-comings and questions regarding the applicability of the ITU model on a Loran receiver are raised. The key finding of the ITU model is that it accurately predicts both the long term rms envelope voltage values as well as the short term distributions for ITU's 200 Hz narrow-band system. However, in reviewing their data collection methodology, the ITU researchers gave no assurances that the same data are applicable to a wide-band system such as Loran which is 35 kHz wide and centered at 100 kHz.

Furthermore, questions are raised by the documents themselves. ITU P.372-7 states that some of the data, especially the bandwidth conversion section which allows for scaling

ITU's 200 Hz bandwidth data to other bandwidths, are,

“...strictly valid only at MF and HF frequencies, so care should be exercised when applying these results to lower frequencies (i.e., LF, VLF, ELF).”

Loran is centered at 100 kHz, and therefore falls in the Low Frequency or LF band, so caution is implied. The earlier revision of the document, CCIR-322-3, also warns,

“Measurements have indicated that Figure 26 gives the proper bandwidth conversion of V_d only for bandwidth ratios on the order of 20 or less and generally predicts too high a value of V_d for larger bandwidth ratios. Therefore, Figure 26 should be used with caution for large bandwidth ratios.”

and

“The validity of the idealized APD curves, in representing the actual distributions, and the way in which they vary with V_d and with bandwidth, have so far only been checked against limited data, and further verification is required. For the moment, therefore, the curves and bandwidth conversion factors should be used with caution.”

Since the Stanford Loran Receiver is 35 kHz wide, and the ITU bandwidth is 200 Hz wide, the bandwidth ratio is 175. Again the text itself implores caution in interpreting the ITU model results.

The text also raises issues concerning the model's applicability during a thunderstorm. CCIR-322-3 states,

“The estimates for atmospheric noise levels given in the Report are for the average background noise level due to lightning in the absence of other signals, whether intentionally or unintentionally radiated. In addition, the noise due to local thunderstorms has not been included. In some areas of the world, the noise from local thunderstorms can be important for a significant percentage of the time.”

Matheson confirmed this during his interview, saying that they turned off the data collection system when local storms were present since there was corona discharging at the antenna. The text also goes on to say,

“Atmospheric radio noise is characterized by large, rapid fluctuations, but if the noise power is averaged over a period of several minutes, the average values are found to be nearly constant during a given hour, variations rarely exceeding ± 2 dB except near sunrise or sunset, or when there are local thunderstorms.”

Since the rms values corresponding to nearby thunderstorms are the primary concern of this dissertation, all of these issues have implications on this research. It is precisely during thunderstorms that a Loran receiver must be shown to work. So while the ITU models the rms values well on average, it is necessary to confirm that the values are applicable during a thunderstorm. Even if the ITU models are accurate, there still is an arbitrary selection for the correct percentage envelope noise level to indicate a nearby storm.

The results of Section 8.2 show that the ITU model conservatively predicts the rms atmospheric noise strength up to the 99.9% level, which corresponds to a nearby storm. That section further confirms that the rms values properly correspond to data intervals that are 15 minutes in length, and that the noise’s short term statistics confirm the APDs

predicted by the ITU model. Furthermore, Section 3.2.5 demonstrates that the rms voltage of the envelope noise, E_{rms} , does provide a minimum bound on V_d .

3.3 Signal Processing

Signal processing is the mathematical manipulation, analysis, or interpretation of signals or waveforms. Two applications of signal processing are estimation and detection. Estimation involves trying to recover properties of a signal, such as amplitude or phase, in the presence of noise. Detection involves determining which signal out of a finite set of possibilities has been sent by a transmitter, such as determining if the transmitter sent a one or a zero. Determining a Loran signal's time-of-arrival is an estimation problem. In this section, typical signal processing techniques are discussed which are used to reduce the noise in a Loran receiver.

Another method of generalizing signal processing is by dividing the field up by the type of mathematics used as either linear or non-linear signal processing. The term *linear* refers to the filters or mathematical functions used in the signal processing algorithm. A function is linear if superposition holds. Mathematically, linear functions are those which have the following properties:

$$f(x + y) = f(x) + f(y) \quad (3.2)$$

$$f(\alpha x) = \alpha f(x) \quad (3.3)$$

A Loran receiver algorithm begins with averaging the raw incoming data to form an estimate of the Loran signal. Averaging is an example of a linear filter.

Kalman filtering, where observations are combined optimally taking into account the statistics of both the signal and noise to form an estimate of the signal, is another example of linear estimation or linear signal processing. For a linear dynamical system generating a signal in the presence of Gaussian noise, Kalman filtering is not only the best linear unbiased estimator, it is the optimal estimator, linear or otherwise [79]. In the case of a constant signal corrupted by independent identically distributed Gaussian noise, the Kalman filter is the equivalent of an average of all previous data.

However, in the presence of non-Gaussian noise, Kalman filtering alone is only the best *linear* estimator, not the over-all optimal estimator. When dealing with non-Gaussian noise, a *non-linear* filters can be introduced which do not have the properties of Equations (3.2) and (3.3) that can outperform Kalman filtering.

Take, for example, a signal that is occasionally corrupted by a constant large value. In this instance, rather than simply averaging all of the samples directly, which would be skewed by the large valued outliers, the estimate can be formed by averaging all of the samples that are below some threshold value, but limit or clip the few outliers above the threshold, $x_{threshold}$, to some finite value and then form the average on the resulting data.

This clipping algorithm is defined as

$$y = \begin{cases} x, & \text{if } x < x_{threshold} \\ x_{threshold}, & \text{otherwise.} \end{cases}$$

In this manner, a non-linear estimate is formed since this function does not have the properties of Equations (3.2) and (3.3). A simple counter example shows the function is non-linear.

Let,

$$x = x_{threshold} + x_{threshold}$$

then

$$\begin{aligned} y &= f(x) \\ &= f(x_{threshold} + x_{threshold}) \\ &= x_{threshold} \\ &\neq f(x_{threshold}) + f(x_{threshold}) = 2x_{threshold} \text{ for } x \neq 0 \end{aligned}$$

While any function which does not conform to Equations (3.2) and (3.3) is non-linear, there are a few which are particularly useful in signal processing. In addition to clipping, receivers often employ hole-punching and hard-limiting in their signal processing. Hole-punching is a procedure where the data which exceed a threshold is digitally removed. By discarding the data entirely, a reduction in the quantity of data is exchanged for an improvement in quality.

Hard-limiting employs the signum function, where only the sign of the incoming signal is used. In this manner, given the input, the filter outputs only a -1 or 1. Hard-limiting is an effective method for detection problems, since it helps to distinguish between two possible finite states where the actual amplitude of the signal is not as important. However hard-limiting alone is not as useful in estimation, since a signal may take on an infinite number of possible values over an interval. To reintroduce the range of the measurement, averaging must also be performed in conjunction with hard limiting.

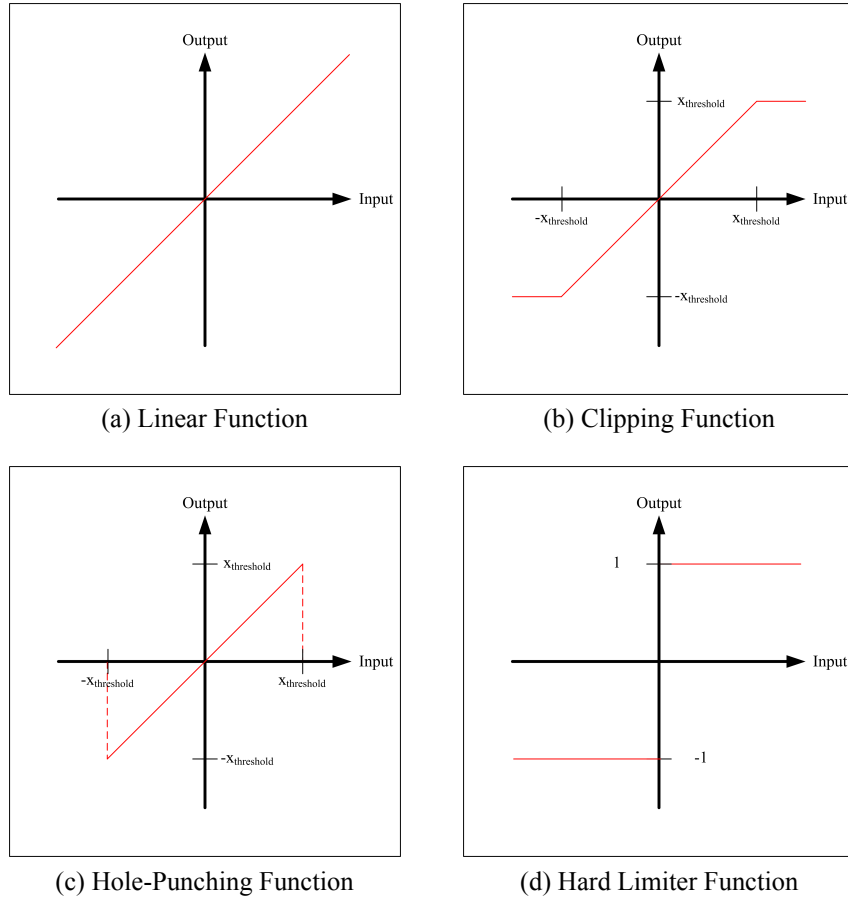


Figure 3.9: Linear and non-linear processing functions

Figures 3.9 (a) through 3.9 (d) are graphs relating the input to the filter's output for the linear, clipping, hole-punching, and hard-limiting cases. Since hard-limiting is mainly used in detection and Loran is inherently an estimation problem because a receiver need to determine envelope ratios and ultimately, a signal's time-of-arrival, only clipping and hole-punching are evaluated in this dissertation.

Many researchers have investigated the use of non-linear signal processing and its application to detection of signals in non-Gaussian noise. In particular, Middleton and Spaulding in [80, 81] have done extensive work in this area, mostly in the field of communications.

Likewise, Enge and Sawate [82] have shown the effects of non-linear signal processing for a spread-spectrum communication signal.

However, most of the literature focuses on communication rather than navigation. In particular, communication receivers employ detection techniques since the receiver needs to determine which discrete symbol a transmitter broadcast. In contrast, Loran receivers are for navigation and use estimation algorithms to determine the signal's time-of-arrival from an infinite range of possibilities. So, communication and navigation reflect the dual nature of detection and estimation.

Despite these differences, the gains produced by non-linear signal processing are similar for Loran and communication algorithms. In particular, Section 5.3.5, shows that the calculated non-linear processing gain for Loran compares well to the work of Enge and Sawate in [82].

Chapter 4

Problem Statement

4.1 Problem Summary

Modeling complex systems requires sufficient detail to accurately estimate the nuances of reality and yet be simple enough to implement. Predicting Loran coverage requires the bringing together of many different models to form a complex system, since coverage depends on a number of factors such as transmission power, terrain, tower geometry, and transmission and receiving antenna patterns. Often in Loran coverage analysis, simplifying assumptions are made to reduce the complexity of the model. However, making too many simplifying assumptions or using overly conservative estimates results in models producing poor or misleading predictions of coverage.

An initial estimation of Loran coverage using the traditional model, produced the results shown in Figure 1.1. The traditional model, more fully described in Section 1.3, used the ITU atmospheric noise model and applied a fixed credit for non-linear signal processing. The coverage diagram depicting the anticipated availability of Loran meeting RNP 0.3

during the worst-case time and season, achieved the FAA's desired level of 99.9% availability only in isolated areas of the East and West coast; these results are both contrary to the expectations and to the experiences of senior researchers in Loran. The two underlying causes for poor coverage were the severity of noise predicted by the ITU atmospheric noise model and the underestimation for the benefits of non-linear signal processing. To illustrate these causes and demonstrate their impact, the following sections review the high level workings of the model to show a more detailed view of Loran coverage. Based on these results, the final section highlights the questions addressed by this dissertation.

4.2 Traditional Model Parameters

While there are numerous parameters to the Loran coverage model to accommodate a range of test conditions, four parameters drove the model's predictions for this dissertation. The first is the estimation of the gain received by non-linear signal processing techniques, such as hole-punching and clipping. For the traditional model, this is taken as a constant gain of 12 dB to the signal-to-noise ratio (SNR), independent of position or noise strength, which is a method similar to References [18, 28, 32, 34, 38, 39, 73].

The second parameter is the minimum tower SNR threshold to accept in determining a position solution. Based on the literature, previous Loran receivers were expected to fail proper acquisition of a signal if the SNR was below -10 dB in Gaussian noise [18, 25, 26, 28, 30, 31, 37], or -22 dB before the 12 dB credit is applied for signal processing. Since the current model coverage software employs receiver integrity monitoring techniques [83] where strong signals are leveraged to aid weaker ones, the minimum threshold SNR before clipping was set to -24 dB.

The final two parameters were the integration time required to acquire the signal and

the integration time required for tracking. Integrating and averaging a signal increases its SNR. Should the SNR be too low after the integration period, the probability of wrong cycle selection may be too high to pass the integrity check, thereby leading to a loss of availability. For acquisition, the integration time was 20 seconds, while for tracking it was 2 seconds.

4.3 SNR Calculations

To illustrate the impact of noise on coverage in more concrete terms, the following example will show how the reception area of the Boise City Loran tower changes as a function of availability. Recall that the percentage availability corresponds to the percentage of noise. Reception area will be defined by the area which the signal-to-noise (SNR) of the tower exceeds a pre-determined threshold. For this example, that threshold is -24 dB.

The Loran Signal Specification [55] defines SNR as the ratio of the rms amplitude of the Loran pulse at the Standard Sampling Point (SSP) to the rms amplitude of the noise present. The SSP is the envelope value $25 \mu\text{s}$ into the Loran pulse. For an ideal Loran pulse, the SSP is 0.506 times the peak amplitude.

The Boise City, OK Loran tower has a transmission power of 900 kW. Using a model of Earth's conductivity and by measuring the distance from each point on the map to the transmitter, the anticipated signal strength of the tower across CONUS is calculated.

After the signal strength is calculated, the next step is to determine the noise. Each percentage of availability corresponds to an identical percentage noise level. Section 3.2.4 describes how the rms noise field strength at different levels of availability is predicted by the ITU model for the worse-case hour and season.


Two availability levels of interest are the 50% or median noise level, corresponding to a “typical” day and the 99.9% median noise level corresponding to the RNP 0.3 requirement. Since the expected values of these levels are used, there is an implicit confidence level of 50%.

Because both the signal strength and noise plots are in $\text{dB}(\mu\text{V}/\text{m})$, subtracting the noise at a given availability from the signal strength plot yields the SNR for that given availability level in the RF bandwidth. Figures 4.1 (a) and 4.1 (b) show the SNR given the median 50% noise level and the 99.9% noise field strengths.

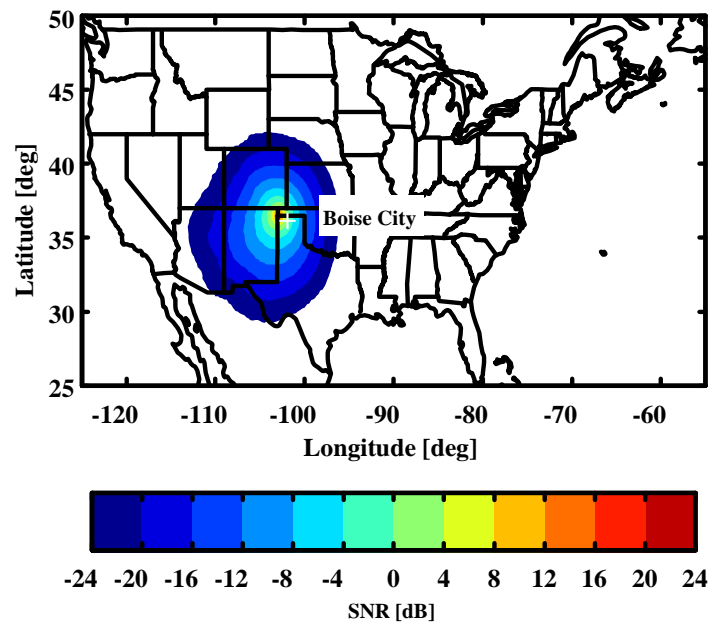
These plots show the reception area for two different levels of availability. If an SNR threshold of -24 dB denotes the minimum SNR that a receiver will use in order to acquire a tower, then the outer dark blue ring would denote the reception area of the Boise City tower. The change from the 50% to the 99.9% noise field strength reduces the reception area by a factor of six.

4.4 Number of Towers Calculation

The Loran coverage software calculated the reception area for each Loran tower. By checking at points across the country, the software determined how many towers were available for a position calculation or fix. For each of the two noise conditions, the number of towers available for a position fix is given in Figures 4.2 (a) and 4.2 (b). A minimum of three towers is required to calculate a position fix, since there are three unknowns, the latitude, longitude, and time offset of the user. When the noise is at the median 50% noise field strength, the coverage software predicts that a Loran receiver can see plenty of towers across CONUS for a solution. When the noise is increased to the 99.9% noise field



(a) Boise City SNR in dB during the Median 50% Noise Field Strength



(b) [Boise City SNR in dB during the Median 99.9% Noise Field Strength

Figure 4.1: SNR and corresponding reception area of Boise City tower corresponding to a -24 dB threshold under different noise conditions.

strength, corresponding to 99.9% availability during the worst-case season and time block, Loran position fixes were unachievable across most of the Plains States.

4.5 Traditional Coverage results

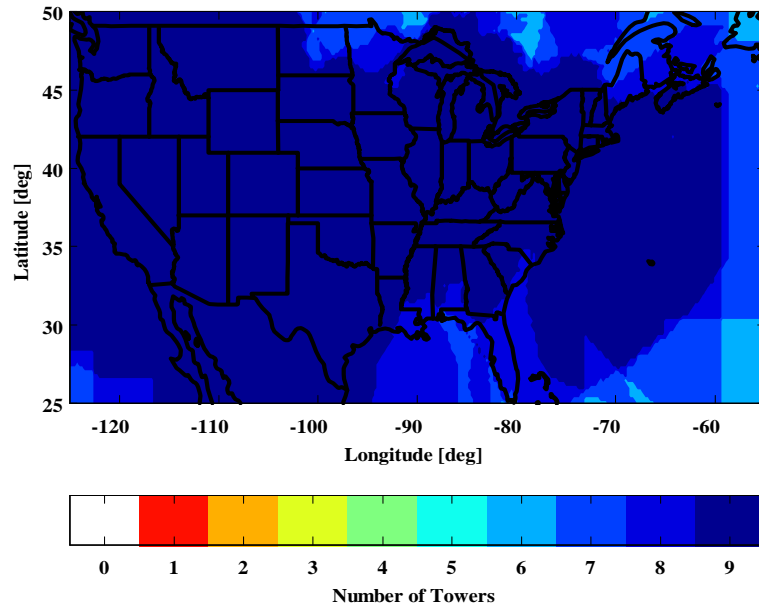
Once the number of towers has been established for each location at a given availability level, the Loran coverage model software calculates whether that location can achieve the requirements for RNP 0.3 by meeting both the monitor limit and the integrity requirements [72]. Figure 4.3 shows the resulting availability for RNP 0.3 coverage across CONUS.

The dark blue regions depict when Loran is available at the 99.9% level, the desired availability for meeting RNP 0.3. Ideally for Loran to be used as a backup navigation system, the FAA would want to see that the 99.9% level is met across all of CONUS, however, the predictions from the traditional model shows this level of coverage for only portions of the East and West Coast.

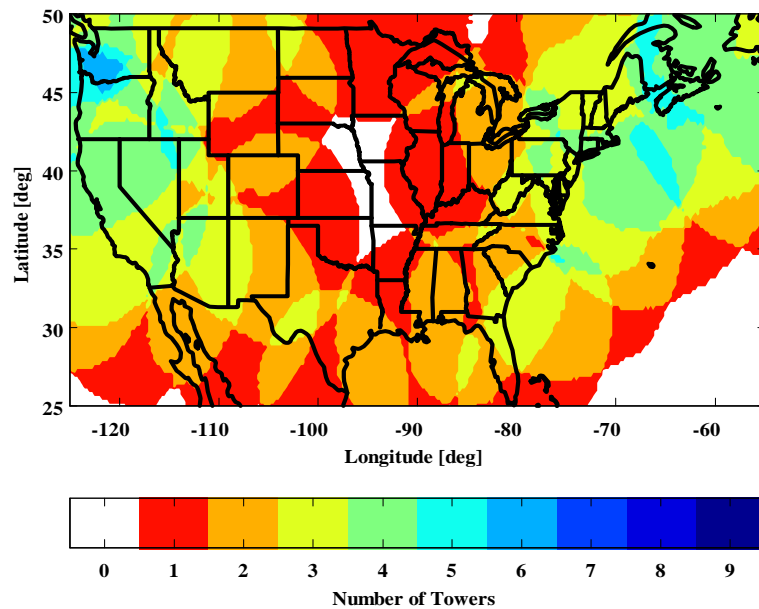
While most of the country shows Loran could meet the RNP 0.3 specifications 90% of the time, there exists a region around Colorado, Kansas, and Nebraska where availability is anticipated to be less than 80% during the worst case season and hour. Additionally, there are regions in the northwest corner of Nebraska and in central Louisiana where the availability is less than 50%. These outages are a result of having too few towers available in those areas for position fixes.

4.6 Research Questions

While the Loran availability for the traditional model is disappointing, it does not accurately reflect the experiences of senior Loran researchers. The conflicting expectations between



(a) Number of Towers for 50% Noise Field Strength



(b) Number of Towers for 99.9% Noise Field Strength

Figure 4.2: Number of towers available at 50% and 99.9% noise field strength.

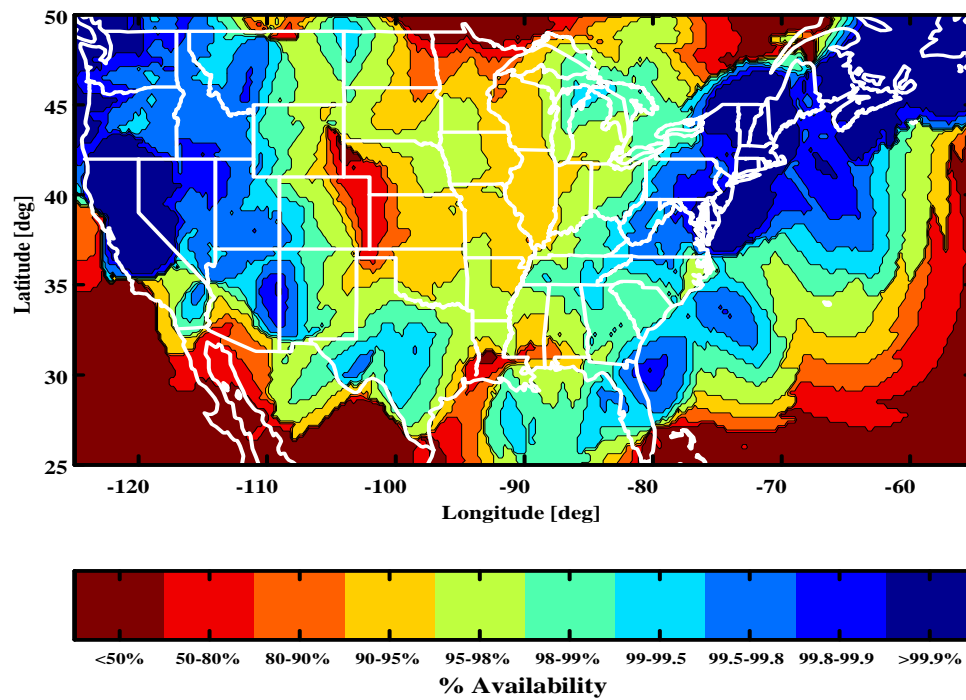


Figure 4.3: Loran availability using only simple constant credit for non-linear processing.

the models and the experts forced the following questions:

1. Is the ITU atmospheric noise model flawed?
2. How does the ITU model compare to other models?
3. Is the coverage model appropriately crediting signal processing performance gains?
4. Are better receiver algorithms required?

This dissertation develops an improved understanding of both the ITU atmospheric noise model and the effectiveness of signal processing on mitigating atmospheric noise in a Loran receiver. Through a more thorough understanding of atmospheric noise and its influence on Loran receivers, this dissertation demonstrates the following:

1. The ITU atmospheric noise model accurately predicts both the rms noise envelopes for 15 minute data intervals over the seasons and the instantaneous noise envelope values for high atmospheric noise.
2. The ITU atmospheric noise model agrees with models made of individual lightning strokes if the appropriate transformations are made.
3. While the rms noise values are correctly predicted by the ITU model, the interpretation of the noise requires careful consideration when it comes to a Loran receiver. The non-linear processing gain, which was previously modeled as constant, is not an accurate assumption, and actually varies with the noise strength.
4. The minimum voltage deviation of the noise and the rms noise envelopes for 15 minute data intervals are correlated. This is used to derive a processing gain proportional to the predicted noise level.

5. Past receiver designs used automatic gain control when processing the noise to more efficiently use the dynamic range of the analog-to-digital converter. This clipped out the majority of the noise and resulted in much improved performance. While any amount of clipping can improve the receiver, a more systematic approach of adjusting the gain relative to the distribution of the noise level can result in further improvements in coverage.

The findings of this dissertation result in an improvement to the coverage diagram. Figure 4.4 shows a lower bound on the expectation of Loran coverage for meeting RNP 0.3 given atmospheric noise. Overall, the availability of Loran for NPA is improved across CONUS. While still not achieving the 99.9% availability level everywhere, there is a 157% increase in the coverage area which meets the FAA's requirements for NPA.

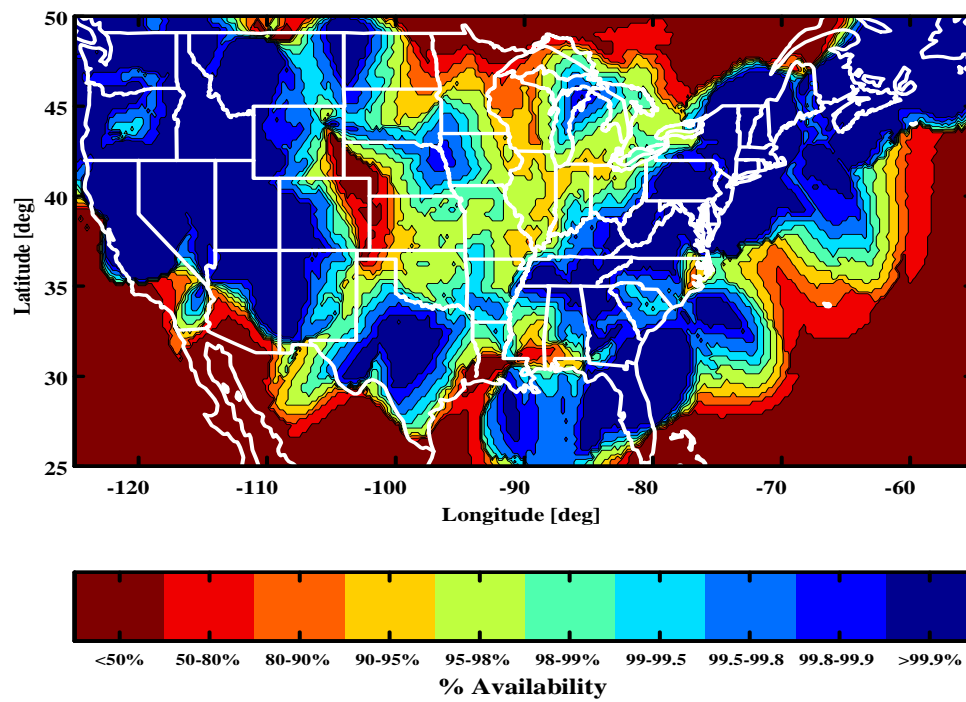


Figure 4.4: Current pessimistic Loran coverage prediction based on the outcomes of this research.

Chapter 5

Signal Processing Development

Due to the interdisciplinary nature of this dissertation, a number of mathematical tools across a variety of fields are employed to properly assess the effectiveness of non-linear processing elements in a Loran receiver. The techniques of the next three sections describe the development of such tools that underpin the results of Chapter 8.

This chapter's first section derives the distribution of envelope voltages found in Loran receiver for varying levels of SNR. With the envelope distribution, the probability of wrong cycle selection is determined. The following section describes the non-linear processing gain calculation derived from the ITU atmospheric noise model's amplitude probability density functions. By recasting the performance of non-linear signal processing gain in terms of noise impulsivity, it sets the stage for this dissertation's contribution of correlating the processing gain to the anticipated noise strength. The final section addresses the shortcomings of using average power as a measure of noise. This discussion leads to another contribution of this dissertation, the relating of the ITU atmospheric noise model to the data from atmospheric physics.

5.1 Envelope Ratio Distributions

5.1.1 Overview

As discussed in Section 2.3, the receiver uses a ratio test to determine which zero crossing is the standard zero crossing (SZC). If the noise is strong enough, it can distort the envelope causing the receiver to track the wrong cycle and induce a 3,000 m range error.

By ignoring propagation errors and considering only thermal noise, then this noise is modeled as Gaussian and the probability of a cycle slip can be derived by integrating analytical models. This section develops $P[\text{Wrong Cycle}]$, the probability of the Loran receiver selecting the wrong cycle based on the ratio test in the presence of Gaussian noise. From this analysis, a reference curve for comparison purposes when evaluating more impulsive noise is generated.

Earlier work in wrong cycle selection was performed by Peterson et al. in [71]. Peterson based most of the analysis on the fact that the envelope in the receiver could be modeled as Gaussian noise. The derivation in this section shows why this is a reasonable assumption when the Loran signal has high SNR, but a poor one for low SNR values.

5.1.2 Assumptions

The derivation starts with the in-phase (I) and quadrature (Q) samples resulting from the injection of Gaussian noise at the receiver's front end. Figure 2.5, from Section 2.3.2, showed the topology of the receiver.

Gaussian noise at the front end results in I and Q noise samples that are also Gaussian. By performing an autocorrelation, I and Q noise samples are found to be weakly correlated in time. A plot of the results for the I channel is shown in Figure 5.1. The Q channel is




Figure 5.1: Autocorrelation of I samples.

similar. For $t = 20 \mu\text{s}$, corresponding to the 50 kHz sample rate, the autocorrelation is below 20%, and approaches zero further out towards $80 \mu\text{s}$. Therefore, the rest of the derivation will ignore this weak correlation, and approximate the I and Q noise samples as independent identically distributed Gaussian variables when sampled at 50 kHz.

While the I and Q noise samples are Gaussian, the envelope is made up of the root-sum-square of these variables. Depending on the amount of Loran signal present on the I channel, a range of distributions for the envelope are formed.

5.1.3 Noise-only Case

If no Loran signal is present, then the receiver measures the noise samples on both the I and Q channels. If I and Q are both zero-mean, equal variance, Gaussian variables, then the joint probability density function (pdf), f_{IQ} , of the variables can be formed. This

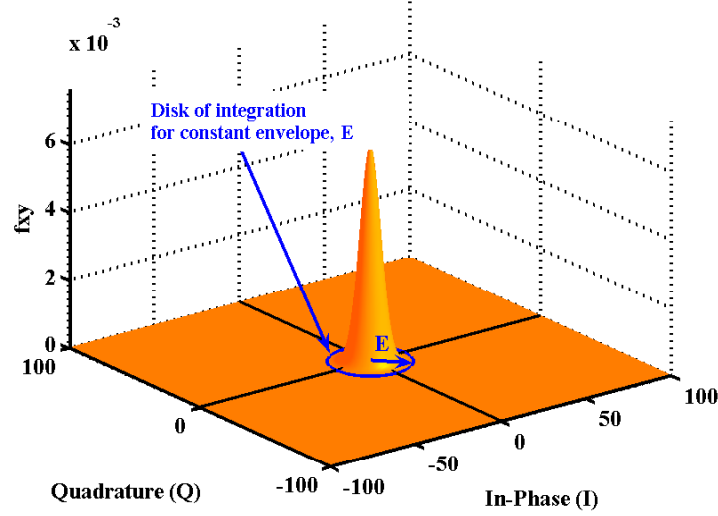


Figure 5.2: Joint probability distribution of two Gaussian variables with zero-mean and equal variance. Indicative of noise-only case ($\text{SNR} = -\infty$ dB).

distribution is shown in Figure 5.2, where the peak of the distribution is located at the origin.

However, the receiver uses envelope values during acquisition, not just the I and Q values. Since the envelope, E , is given by the root-sum-square of the I and Q samples, i.e.,

$$E = \sqrt{I^2 + Q^2}$$

then the envelope defines the radius of a circle centered at the origin of the I - Q plane. To determine the probability of the envelope falling within a particular value, Z , then the blue circle in Figure 5.2 outlines a disk which f_{IQ} must be integrated over to form $P[E = I^2 + Q^2 < Z]$. Integrating the joint pdf over a range of values for Z results in the cumulative distribution function (CDF) of the envelope. Differentiating the CDF then produces the pdf of the envelope. This distribution may be analytically calculated and is called a Rayleigh

distribution [84], with a pdf given by,

$$f_{Rayleigh}(x) = xe^{-\frac{x^2}{2}}.$$

5.1.4 Low SNR case

A Loran signal was not present in the previous case. The next two cases demonstrate what happens to the envelope distribution as the SNR increases. Recall from Section 4.3, the SNR in the RF bandwidth is defined as the Loran rms signal power of the Standard Sampling Point (SSP) divided by the noise power, that is

$$\begin{aligned} SNR &\equiv \frac{SSP_{rms}}{Noise_{rms}} \\ &= \frac{SSP}{\sqrt{2}\sigma_{Noise}}. \end{aligned}$$

In dB this becomes,

$$\begin{aligned} SNR_{dB} &= 20 \log_{10} [SNR] \\ &= SSP_{dB} - 3 \text{ dB} - \sigma_{Noise_{dB}}. \end{aligned}$$

As SNR increases, the peak of the distribution is displaced. Without loss of generality, a receiver would steer the Loran signal towards the in-phase direction, therefore, the I component is modeled now as a Gaussian, with finite mean, ν , corresponding to the SNR, and with a given variance, σ . The Q component will still be zero-mean Gaussian, with the

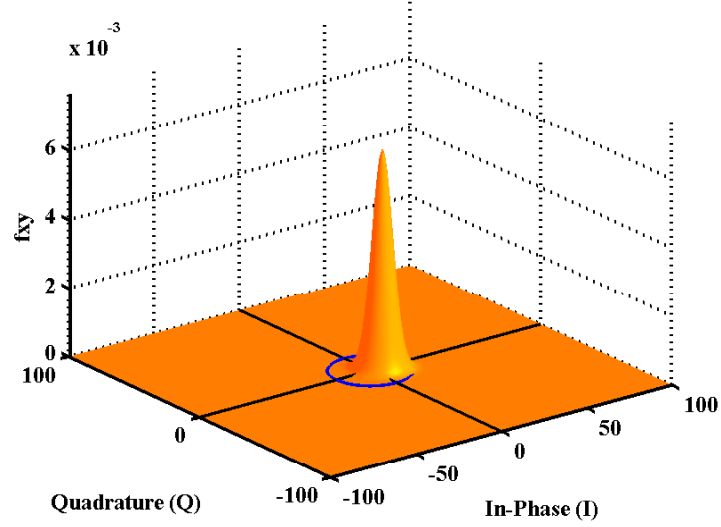


Figure 5.3: Joint probability distribution of one zero-mean and one finite-mean Gaussian variable with equal variance. Indicative of low SNR case (SNR = 0 dB).

same variance as I . As SNR increases, the peak of the joint pdf moves farther along the I channel as seen in Figure 5.3. The result is known as a Rice Distribution [85], whose pdf is

$$f_Z(z|\nu, \sigma) = \frac{z}{\sigma^2} \exp\left(-\left(\frac{z^2 + \nu^2}{2\sigma^2}\right)\right) I_0\left(\frac{z\nu}{\sigma^2}\right) \quad (5.1)$$

where $I_0(x)$ is the modified Bessel function of the first kind.

To determine the cumulative distribution function of the envelope, the Rice distribution is integrated. This requires numerical methods. Note that as the Loran signal goes to zero, the SNR, typically measured in dB, goes to negative infinity, and the distribution reverts back to Rayleigh.




Figure 5.4: Joint probability distribution of one zero-mean and one finite-mean Gaussian variables with equal variance. Indicative of high SNR case (SNR = 20 dB).

5.1.5 High SNR Case

As SNR continues to increase, Figure 5.4 shows that the peak of the joint pdf moves farther out along the I axis. As a result, for SNR > 20 dB, the Q channel has little influence on the magnitude of the envelope, and only affects the phase of the measurement. Therefore, the I channel noise dominates the envelope which takes on the same statistics of the I channel, i.e., Gaussian.

Another manner of viewing the high SNR case is shown in Figure 5.5 which is a perspective from within the I - Q plane. Here, the true envelope value is in blue, while additive noise in both I and Q directions is in red. The measured envelope, shown in green, is the addition of the true envelope and the noise, and the peak of the joint pdf function, f_{IQ} , is superimposed in orange. As the SNR increases, the true envelope will increase relative

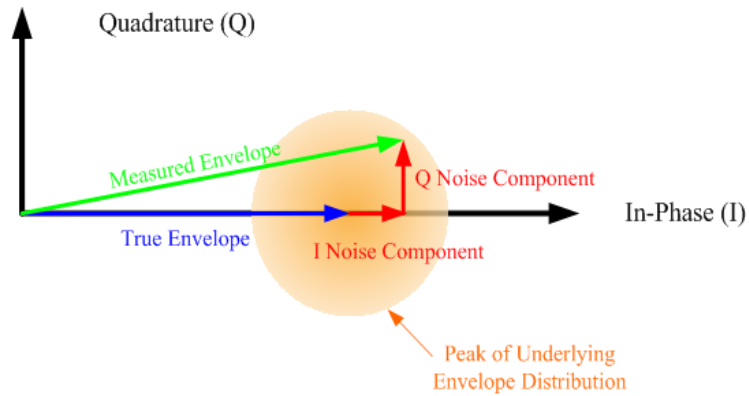


Figure 5.5: I-Q plane view of noise added to an envelope voltage.

to the noise components and eventually only the in-phase noise influences the magnitude of the envelope. The quadrature component does not contribute to the magnitude of the envelope; it will only contribute to its phase error.

The reduction of the effect of the Q component on the magnitude of the envelope for the high SNR case simplifies the modeling of the envelope. Now, the envelope noise model is the same as the I noise component which is Gaussian.

To demonstrate how this effect varies with SNR, the CDF and APD of the $15\ \mu\text{s}$ Envelope Point for two different SNR values are calculated. Recall from Figure 2.6, that the $15\ \mu\text{s}$ Envelope Point and the $30\ \mu\text{s}$ Envelope Point are the two samples whose ratio indicates the correct zero-crossing.

If a Gaussian approximation is used for the envelope noise distribution in the low SNR case, rather than the exact Rice distribution, then Figure 5.6 (a) shows that the Gaussian over-estimates the lower end of the distribution, and therefore, in Figure 5.6 (b) the Gaussian also under-estimates the higher values of the distribution. At higher SNRs, like the ones shown in Figures 5.6 (c) and 5.6 (d), the Gaussian approximates the true distribution better.

Therefore, for SNRs greater than 20 dB, the Gaussian approximation is valid. A typical




Figure 5.6: Distributions of the $15 \mu s$ ratio point for high and low SNR. Comparison of the exact CDF and APD to the Gaussian approximations for the distribution are shown.

Loran receiver integrates its incoming signals to achieve a large SNR. Therefore, when analyzing the statistics of the $15 \mu s$ Envelope Point in a receiver, the Gaussian approximation to the distributions, which has been the case historically [70] is accurate. However, the more exact Rice distribution becomes a better estimate than the Gaussian approximation as the SNR drops below 20 dB.

5.2 Probability of Wrong Cycle Selection

Now that the distribution of any single envelope value has been shown to be Rician, with limiting cases of either Rayleigh for no signal present or Gaussian for high SNR, the next step is to determine the distribution of a ratio of two envelope values. Reference [84] states that the probability density function of the ratio, q , of two variables, z_1 and z_2 , is

$$q = \frac{z_2}{z_1}$$

$$f_Q(q) = \int_{x=0}^{\infty} z_1 f_{Z_1}(z_1) f_{Z_2}(qz_1) dz_1 \quad (5.2)$$

Taking the ratio of two, zero-mean, finite variance Gaussian variables, forms a Cauchy distribution [84]. Since Loran will always be present, to ensure complete generality and exactness, and to cover a variety of SNRs, it is more accurate to take the ratio of two Rice distributions since they account for finite-mean Gaussian variables.

With two envelope values following Rice distributions, Equations (5.1) and (5.2) are combined to form the pdf of the envelope ratio as

$$f_{\text{Ratio}}(q) = \int_{-\infty}^q \xi \frac{\xi}{\sigma^2} \exp\left(-\left(\frac{\xi^2 + \nu^2}{2\sigma^2}\right)\right) I_0\left(\frac{\xi\nu}{\sigma^2}\right) \frac{q\xi}{\sigma^2} \cdot \exp\left(-\left(\frac{q^2\xi^2 + \nu^2}{2\sigma^2}\right)\right) I_0\left(\frac{q\xi\nu}{\sigma^2}\right) d\xi \quad (5.3)$$

$$= \int_{-\infty}^q \frac{q\xi^3}{\sigma^4} \exp\left(-\left(\frac{\xi^2 + q^2\xi^2 + 2\nu^2}{2\sigma^2}\right)\right) I_0\left(\frac{\xi\nu}{\sigma^2}\right) I_0\left(\frac{q\xi\nu}{\sigma^2}\right) d\xi. \quad (5.4)$$

Numerically integrating Equation (5.4) forms the CDF, thereby determining the probability of achieving a particular ratio value. In practice, as the SNR increases, the distribution gets narrower, making numerical integration difficult due to grid size. For $\text{SNR} > 25$, the Gaussian approximation integrates easier and gets around these numerical integration issues, and thus Gaussian distributions were used for higher SNRs.

With Equation (5.4), the ratio of envelope values used in the ratio test within a Loran receiver may now be determined. From Section 2.3.2, the receiver ratio test algorithm requires the ratio of two envelope points $15 \mu\text{s}$ apart. Since the receiver is seeking the ratio test statistics at $30 \mu\text{s}$ to denote the SZC, it is useful to define the function $\text{Ratio}(\tau)$ as

$$\text{Ratio}(\tau) \equiv \text{Envelope}(\tau - 15 \text{ s}) / \text{Envelope}(\tau)$$

where,

$\text{Envelope}(\tau)$ is the Loran envelope amplitude of the $\tau \mu\text{s}$ point, and

$\text{Ratio}(30)$ is bounded by

$$\text{Ratio}(25) \leq \text{Ratio}(30) \leq \text{Ratio}(35)$$

So long as the receiver's ratios calculated from the envelope values falls within this bound, the receiver will select the correct zero crossing as the SZC. However, computed ratio values less than Ratio(25) and greater than Ratio(35) are of interest. Should a receiver's estimate of Ratio(30) exceed one of these threshold values, then the receiver would choose the wrong cycle zero crossing, since the wrong zero crossings will be closer in time than the SZC.

Therefore, given different SNRs and integrating from negative infinity to Ratio(25) gives the probability of the receiver choosing the earlier cycle rather than the correct one. Similarly, integrating from Ratio(35) to infinity, yields the probability of the receiver choosing the later cycle incorrectly. Adding these two probabilities together gives $P[\text{Wrong Cycle}]$, or

$$P[\text{Wrong Cycle}] = P[\text{ratio} < \text{Ratio}(25)] + P[\text{ratio} > \text{Ratio}(35)]$$

Figure 5.7 shows $P[\text{Wrong Cycle}]$ vs. SNR, in black, the probability of choosing the wrong cycle in the presence of Gaussian noise. To verify that the equations are indeed correct, five seconds of Loran signals of various SNRs were simulated and passed through a simulated Loran receiver with averaging as the only processing. Five seconds of averaging improves the SNR by +26 dB so the x-axis is plotted displaying the SNR post averaging. The blue trace shows that the measured performance matches the theoretical predictions well.

In [70], Peterson estimated the envelope errors as Gaussian. He predicted the phase error of the received signal as

$$\sigma_{\text{Phase Noise}} = \frac{28 \mu s}{\sqrt{2 N SNR}}$$




Figure 5.7: Probability of wrong cycle selection given the post-processing SNR under Gaussian noise.

Given that the envelope's time of arrival is given by this variance, integration of a Gaussian of the same variance determines when the timing is off by more than 5 microseconds. Such an offset is another method of viewing the wrong cycle selection blunder. Performing that integration, generated the red curve in Figure 5.7, which shows the Gaussian approximation provided a conservative over-bound of the true distribution.

Since averaging for five seconds added 26 dB to the SNR, a received signal with 0 dB of instantaneous SNR in the RF bandwidth, can be acquired with a 10^{-8} chance of being on the wrong cycle. As SNR decreases, the averaging times must increase to make up this difference, or rely on integrity monitoring algorithms that could work with TOA measurements that have a higher probability of being on the wrong cycle. For example, extending the averaging time to 50 seconds, which is an upper limit due to clock stability, would allow signals as low as -10 dB SNR to be acquired.

The Austron curves of Figure 5.7 are a baseline performance of a Loran receiver for various SNRs. These baseline curves for Gaussian noise provide an easy method of comparing receiver performance when impulsive noise of equivalent power passes through a receiver.

5.3 Signal Processing with APDs

5.3.1 APD Review and Interpretation

Section 3.2.5 discussed how the ITU atmospheric noise model characterizes the distribution of noise using Amplitude Probability Distributions (APDs) which were generated through years of data collection. The ITU model is not the only model. Other researchers have attempted to predict APDs based on various mathematical models of atmospheric noise.

Some researchers make parametric models that fit the noise data, as Hall does in [86] and Feldman in [34], while others approach the matter in more of a physical model approach as in Middleton and Spaulding in [87–90]. Regardless of the manner in which it is generated, the APD gives two important characteristics of the noise that will be used to determine the processing gain for non-linear signal processing: how often the noise is impulsive, and how much lower the Rayleigh portion of the noise is from the rms value. The next few sections describe these steps in detail.

5.3.2 Reading the APD

Recall from Section 3.2.5, that the APD is equivalent to unity minus the cumulative distribution function (CDF). The APD also differs from the CDF in its appearance which makes the APD more useful for noise analysis. As an example to show the utility of the APD, start with the APD curve corresponding to a V_d of 10, taken from the ITU model. The APD is used to find the probability that the instantaneous envelope voltage, A , will exceed a given value relative to the rms value, A_0 . For example, as shown in Figure 5.8, the APD shows that the instantaneous envelope voltage will exceed the rms value by more than 10 dB about 1% of the time. Alternatively, the APD shows that the noise is below the rms value 99% of the time. This is quite different than Rayleigh noise, with a $V_d = 1.05$, where the noise exceeds the rms value 36% of the time.

The choice of x-axis scaling and the red Rayleigh reference line, provides another useful tool. Recall, that Rayleigh noise is the root-sum-square of two zero-mean equal variance Gaussian variates. As mentioned in Section 5.1, the noise on either the I or Q channels when no signal is present is Gaussian. The root-sum-squaring of these variates, that is, taking the magnitude of the phasor made by I and Q , produces the noise envelope which




Figure 5.8: Amplitude probability density (APD) function of Rayleigh noise ($V_d = 1.05$) and impulsive noise ($V_d = 10$)




Figure 5.9: APD showing the division of noise as being impulsive or Rayleigh in distribution.

will be Rayleigh in distribution. This distribution is so prevalent in signal processing that typically a red reference line depicting a Rayleigh distribution is superimposed on the APD. With this reference line, the amount that distribution diverges from Rayleigh is quickly seen.

5.3.3 Signal Suppression

The extent to which a distribution remains parallel to the Rayleigh reference line indicates how long the noise underlying that distribution is Rayleigh. As the noise becomes more

impulsive and less Rayleigh like, the distribution's slope becomes steeper. By examining the point at which the APD deviates from a Rayleigh distribution, the percentage of time the noise is impulsive or the complementary amount of time that the noise is Rayleigh in distribution can be determined.

For this example, with a V_d of 10, the point at which the APD deviates from the straight line portion by more than 3 dB is the criteria for when the noise is non-Rayleigh. As shown in Figure 5.9, the 3 dB point is obtained by drawing a construction line tangent to the Rayleigh portion of the APD, and then by finding the point at which the APD exceeds a Rayleigh distribution by 3 dB. For a V_d of 10, the noise is Rayleigh approximately 52% of the time and it is impulsive the remaining 48%. The result is that both the noise and desired signal would be punched out or suppressed 48% of the time.

Non-linear signal processing techniques such as hole-punching can be used to remove large voltages from the data. As mentioned in Section 3.3, hole-punching and clipping are non-linear processing elements which adjust the signal when the incoming signal surpasses a set threshold level.

If a threshold is set at the 48% level indicated by the APD, which from the y-axis is 18 dB below the rms level, then both signal and noise are removed when the combination of the two exceeds the threshold, resulting in keeping 52% of the signal and throwing out the remainder. Since both the signal as well as the noise is suppressed 48% of the time, this non-linear processing method cost 2.8 dB loss due to signal suppression.

In a similar fashion, the APDs listed in the ITU noise model can be used to determine approximately the amount of time the noise is Rayleigh, or "available." Furthermore, the amount of loss due to signal suppression can be calculated.

In Figure 5.10, the calculated signal suppression values for each of the APD curves up




Figure 5.10: Loss due to signal suppression.

to a V_d of 16 is plotted. The APD curves have been drawn based on limited experimental data, and no analytical expression is given. Therefore, the suppression values are calculated by hand.

With a best-fit line through the points, a linear relationship, between V_d and suppression of the signal is given by

$$\text{Loss} = -0.3 V_d - 0.2 \text{ dB} \quad (5.5)$$

Thus, for any V_d value, the loss due to signal suppression when the noise is non-Gaussian is readily obtained.

5.3.4 Processing gain

While some amount of loss is incurred by clipping, hole punching or hard limiting, this section will now explore the benefits it brings. Figure 5.11 again shows the construction line parallel to the Rayleigh reference line. Here, the Rayleigh portion of the APD is 20 dB lower in power than its rms. Hence, the noise after hole-punching 52% of the time would be 20 dB below the original rms noise level. Thus, the effective noise level drops by 20 dB at the cost of 3.2 dB of signal suppression.

Applying this analysis to the APDs listed in the ITU model, produces Figure 5.12, which shows the amount that the Rayleigh portion of the noise is lower in power than its rms value. The best line fit through the points results in a reduction of the noise level that follows

$$\text{Gain} = 2.1 V_d - 0.9 \text{ dB} \quad (5.6)$$

5.3.5 Total Non-linear Processing Gain

The previous two sections, have shown that the non-linear processing credit is made up of two parts, a loss from the suppression of signal through the implementation of the non-linear element, and a gain from the resulting noise being of lower value than the original distribution's rms value. By separating the noise into Rayleigh and impulsive components and eliminating the impulsive part using non-linear elements, Equations (5.5) and (5.6) can be combined, to find the total gain due to non-linear processing:

$$\text{Non-linear Processing Gain} = 1.8 V_d - 1.1 \text{ dB} \quad (5.7)$$




Figure 5.11: APD depicting the rms of the Rayleigh level as below that of the overall distribution.




Figure 5.12: Difference between the rms value of the Rayleigh noise and the overall rms value of the distribution.

The equation is plotted in red in Figure 5.13. Also plotted are the analytical results derived by Enge in Reference [82] for clipping on a Spread Spectrum Multiple Access (SSMA) system in black. SSMA is a communication method using spreading codes to transmit symbols of information. The analysis of an SSMA system would be similar to Loran since both systems employ bursts of pulses rather than continuous wave modulation. However, a SSMA receiver is attempting to solve a detection problem since digital communication concerns itself with trying to distinguish which symbol was sent, while a Loran receiver is working on a estimation problem since the receiver is trying to determine when a signal arrives. Despite the duality between detection and estimation was discussed in Section 3.3, this dissertation's results are comparable to those obtained by Enge.

While the focus of this section has been on the hole-punching element, clipping implemented in a Loran receiver would perform comparably [18,28,32,34,38,39,73]. There have




Figure 5.13: Total non-linear processing gain as a function of V_d .

been differences found as stated by [34] and [81], but these differences are small compared to the overall gains.

With this analysis, the ITU model APDs parameterized by V_d is tied directly to the performance of non-linear signal processing. In Section 8.4, Equation (5.7) will show that the gains generated by non-linear signal processing based on the anticipated noise level can be bounded.

5.4 Noise and the Limitations of Average Power

5.4.1 Noise

Noise is any man-made or natural process which corrupts the desired signal, in this case Loran. To describe noise, the field can be divided as shown in Figure 5.14. Noise may be

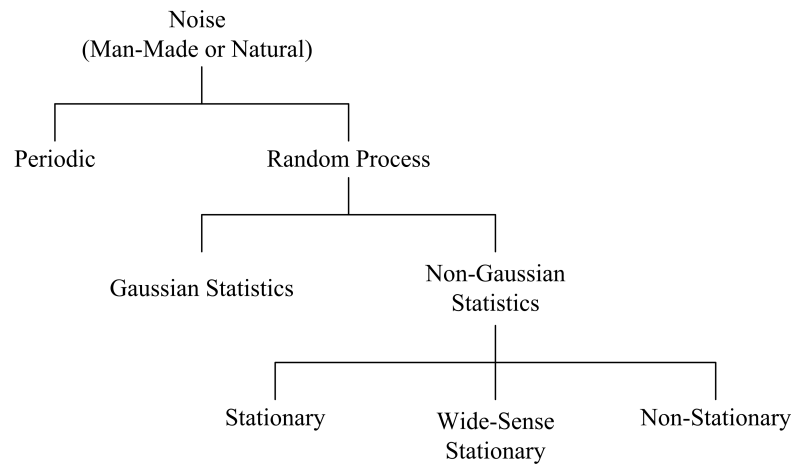


Figure 5.14: Noise categories.

periodic or it may be random. If noise is random in nature, then it can be further divided into two more sub-categories based on its statistical distribution being either Gaussian or non-Gaussian. Impulsive noise falls under the category of non-Gaussian. Additional divisions can also be made to reflect the consistency in time of the statistics. If the statistics are constant over all time then they are called stationary, or if only the first order statistics are constant over time, then they are wide-sense stationary. However, if all of the statistics change over time then the process is non-stationary.

In Gaussian noise, also called a Gaussian random process, each time sample of the process follows a Gaussian distribution. In addition, the mean and the variance of an ensemble of various realizations of the process also follow a Gaussian distribution. A special case of noise is when the samples are independent of one another, but identically distributed (IID). An IID Gaussian process is completely described by its first order statistics, the mean and variance, or power, of any one sample. Because of the mathematical ease of describing Gaussian noise, it is the most readily analyzed and is most familiar to engineers.

5.4.2 Problems with Non-Gaussian Noise

Lightning occurs in bursts, and is neither periodic, nor stationary, nor Gaussian. Therefore, while the use of average noise power works well for periodic or stationary Gaussian random processes, it is not sufficient for atmospheric noise processes.

For this latter case, additional information beyond first order statistics is required; in particular, a time interval over which the process occurs to handle the non-stationary nature of the noise needs to be defined. As discussed in Appendix A, this time interval should be at least as long as the decorrelation time of the receiver's front end. The need to define a time interval makes working with non-Gaussian noise processes a challenge. To emphasize the need for explicitness in defining an interval over which to average power or to calculate the root-mean-square (rms), the following example is provided.

Figure 5.15, shows a sample waveform defined on two different intervals; the first being from $[0,4]$ and the second from $[0, 16]$. The rms values over the two different waveforms, shown in red, are 2.5 V and 1.25 V, respectively. By quadrupling the time interval, the amount of power is halved.

Whereas the instantaneous voltage, and hence instantaneous power, are identical for the two waveforms, the “zero-padding” of the second time interval greatly reduces the average power value for such non-periodic signals. Thus, the need to explicitly state the time period for a proper interpretation of average power is demonstrated.

5.4.3 Problems with Average Power

Average power as a measurement of noise is fundamentally limited. By definition, in averaging power of a non-stationary process, the details of the data on the interval are reduced to a single number, and so information and detail is traded off for simplicity. For Gaussian




Figure 5.15: Example of the effects of quadrupling the time interval on rms power.




Figure 5.16: Example showing how the same rms voltage value can take on different shapes in a finite interval.

noise this is sufficient since only the first order statistics are required to characterize the noise, but as shown in the previous example, for non-Gaussian noise, the length of the time interval can affect the interpretation of the data.

To further emphasize the inherent weakness in only using average power to characterize noise, Figure 5.16 depicts another example where three different noise waveforms in green occur over the time interval from $[0,4]$ seconds. Each of these three waveforms has the same root-mean square value over the interval. Since the square of the rms value is proportional to power, each of these waveforms is of identical power.

If the signal of interest is shown in magenta, then even though the three noise waveforms have the same power, each will affect the signal in very different ways. In the first case, the pulse of interest is completely free of noise, while in the second case, the leading edge is corrupted by noise. In the final case, the entire pulse is corrupted. Again the rms value or average power alone is insufficient to describe the process well; the time period over which the signal is acting or more ideally, the distribution of power as a function of time is required.

As described in Section 3.2, ITU found that the distributions of the noise is tied to




Figure 5.17: Example showing how the same rms voltage value can take on different averages and voltage deviations.

voltage deviation or V_d which is a measure of the impulsivity of the noise. Figure 5.17 shows the same waveforms and includes their mean or average voltage. While they all have the same rms voltage, their means are distinct. Additionally, V_d is shown for each waveform over the interval $[0,4]$. As the noise becomes more impulsive or “spiky,” V_d increases. The role of V_d and signal processing is reserved for Chapter 8, however, this example serves to illustrate that for non-Gaussian noise, rms or average power alone is an abstraction which glosses over some important properties of the noise process.

Despite its insufficiency to completely describe non-periodic and non-Gaussian noise, average power is a useful starting point to begin the discussion of determining the effectiveness of navigation or communication systems. Section 8.4 shows that the statistics for V_d , rms noise electric field strength, and the instantaneous noise electric field strength distributions are all related. Additionally, suggested time intervals are given to capture the short term impulsive power which are on the order of the time of a lightning return stroke, as well as longer time periods which reflect the quasi-stationary power of a storm.

Chapter 6

Data Collection System

Over the course of collecting noise data in the Loran band, the research for this dissertation became focused on validating the ITU atmospheric noise model. The ITU model predicts long and short term atmospheric noise envelope values that varied across the country and over different seasons. It was necessary, especially in light of the concerns mentioned in Section 3.2.6, to verify the large rms values predicted by the ITU model since those excursions drove the Loran system's availability. The data collection system evolved over the years to reflect the refining of the research goals and the need to capture the large variations in atmospheric noise. This chapter gives a summary of the equipment used from 2004 to the present, but focuses on the details of the 2005 system which generated the data for this dissertation.

6.1 Brief History

From 2004 to 2006 the atmospheric noise data collection system evolved. In 2004 and for part of 2005 the system consisted of a Locus receiver, a commercial-off-the-shelf (COTS)

Loran receiver, the Locus LRS III D. Due to its high gain, the COTS receiver clipped the received atmospheric noise generated by nearby lightning. The desire to capture storm data led the author to design the Stanford Loran Receiver (SLR) Revision A in 2005. This new receiver, managed to capture the high energy signals from nearby lightning strokes, but due to its limited dynamic range, the receiver could not capture the lower energy Loran signals. In an attempt to emulate the ITU equipment, this receiver also came with a 200 Hz bandwidth channel. Since capturing both high and low energy signals was desirable, in 2006, the SLR was upgraded to use two Loran band channels, one with high gain, to capture Loran signals, and a second at low gain, to capture the nearby lightning noise. This receiver, designed by the author and Nicolai Alexeev, is the Stanford Loran Receiver Revision B. Currently, this receiver is in the field and is being evaluated. More on this receiver will be discussed in Section 9.5.

6.2 Location Requirements

The choice of test location had to meet a few constraints. Since lightning storms vary in their intensity and frequency of occurrence, locations were desired which would regularly provide highly energetic storms. The data collection system need to run remotely, so Internet access was mandatory.

While that left a large geographic region available for testing, the hospitality extended by the professionals working in the various facilities narrowed the choices considerably. In particular, the generosity of the various laboratory and university managers have been most appreciated and without whom, testing would have been substantially more difficult. Figure 6.1 shows a map of the locations while Table 6.1 lists the seasons tested and qualitative comments on the types of lightning.

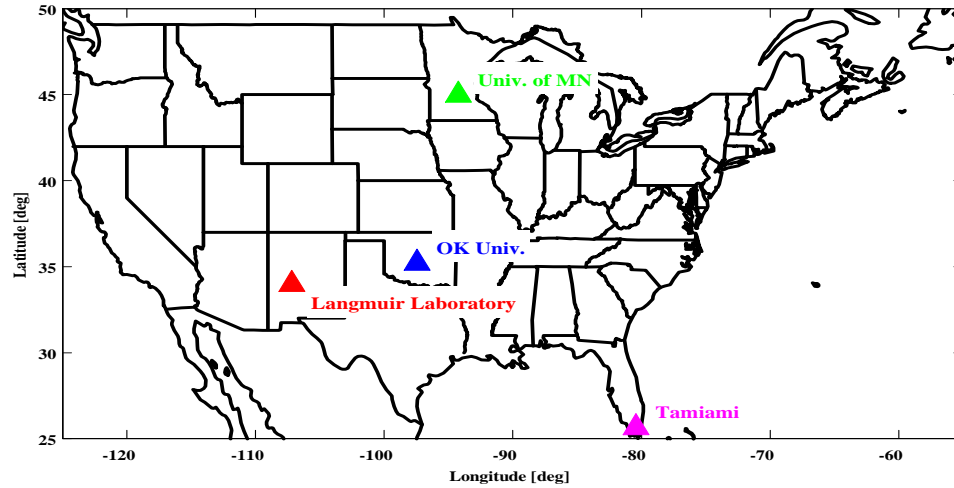


Figure 6.1: Test site locations throughout the 2004-2006 data collection campaigns.

Site	Oklahoma University	Langmuir Lab	Univ. of MN/TC	Tamiami Airport
Location	Norman, OK	Socorro, NM	Minneapolis, MN	Tamiami, FL
Years	2005 - 2007	2004 & 2005	2004	2005 & 2006
Season	Spring – Summer	Summer	Summer – Fall	Summer
Lightning Level	Extremely High	Moderate to High	Moderate to High	Light
Storm Frequency	Common	Daily	Common	Daily
Contact	Dr. Bill Beasley	Dr. Bill Winn	Dr. Demoz Gebre-Egziabher	Dr. Dave Diggle

Table 6.1: Test Location and Lightning Characteristics

University of Oklahoma
Norman, OK



Figure 6.2: Sarkey's Energy Center at the University of Oklahoma, Norman, OK.

6.3 2005 Data Collection Campaign Locations & System Configurations

With the hopes of encountering severe storms, in 2005 a site atop the Sarkey's Center at the University of Oklahoma was established. Figure 6.2 shows the fifteen story tall Center. It is the tallest structure until downtown Oklahoma City, approximately 17 miles away. The top floor of the building houses the Meteorology Department which made roof access feasible and provided an unobstructed view of the sky as seen in Figure 6.3.

After collecting and analyzing the data from the previous year's campaign, it was found that both the COTS antenna and receiver clipped the peaks of the large amplitude lightning strokes. Since the COTS receiver and antenna were designed to capture Loran signals, their dynamic range was not adequate to capture the full power of nearby lightning. The clipping of the data by the receiver reduced the overall rms level of the interval thus rendering them



Figure 6.3: View from atop the Sarkey's Center with Locus Loran and TrueTime GPS Antenna.

useless for validating the ITU atmospheric noise model.

The need for a receiver which could capture the large amplitude fields of lightning led to the development of a new RF front-end. Figure 6.4 depicts the SLR, Revision A. This new Loran Receiver design had two different bandwidth channels. While References [25, 26, 39, 61] give receiver bandwidths from 20 to 30 kHz, the first channel was chosen to be similar to the Locus receiver: a 35 kHz wide band-pass filter centered at 100 kHz. By going with a bandwidth that encompassed all other receiver designs, it was hoped to make the data the most general to all receivers. This would allow for digital filtering of the data to smaller bandwidths in post-processing if necessary for a given receiver design. In addition to a 35 kHz wide bandwidth, a second RF channel was added to make the comparison to ITU unambiguous. This second channel used a 200 Hz wide band-pass filter centered at 100 kHz.

Simultaneous with the raw RF output, an envelope detector converted the raw RF signal to log-envelope values and provided an additional data source on each channel. The




Figure 6.4: Circuit board of Stanford Loran Receiver Revision A.

receiver, therefore, had a total of four analog signal outputs.

As for an antenna, Locus Inc. provided an antenna that had only 6 dB of gain rather than the 26 dB of their standard antenna. The combination of the Stanford receiver and the new antenna could capture the high-amplitude signals of nearby lightning.

In addition to Stanford's equipment, additional hardware was provided by Oklahoma University's Meteorology Department. Professor Beasley, a veteran of atmospheric research at the department, not only helped to secure the site, but provided both an electric field mill, to measure the steady-state component of the electric field and a wide-bandwidth flat-plate antenna. Figure 6.5 pictures the flat-plate antenna on the roof of the Sarkey's Center. The flat-plate antenna had a low-end cut off near 10 kHz and a high-end 3 dB point at 250 kHz. The wide bandwidth of the flat-plate antenna could capture most of the energy of a lightning stroke.

Many atmospheric scientists have used wide-band antennas and much data have been




Figure 6.5: 200 kHz bandwidth flat-plate antenna provided by Professor Beasley of the University of Oklahoma. The antenna measures 1 meter on each side and is held down with cinder blocks.

gathered to describe the expected waveforms seen during lightning return strokes. The flat-plate antenna provided a means of comparing the Loran band data to a large existing body of existing literature in atmospheric physics. In addition, the flat-plate antenna was used to calibrate the system as described in Appendix C.

The data collection system included an ICS-652 card, a 14 bit analog-to-digital converter (ADC). A daughter card mounted on the ADC card provided a means of performing the digital down conversion to 50 kHz I and Q in hardware rather than in software. Data for the Loran band were recorded in five second intervals.

The wide-band flat-plate antenna was connected to the system for a few weeks of data collection in lieu of the log-envelope detector channels. By adjusting the properties on one of the ICS-652 cards, flat-plate antenna data was recorded at 1.6 MHz for 0.3125 seconds within the five seconds of data interval. After a couple of weeks, the log-envelope channels were reinstated and the flat-plate antenna channel removed.

To ensure a stable time base, a function generator with a 10 ppb oven-controlled crystal oscillator (OCXO) provided the sampling clock to the system. A TrueTime GPS receiver was added for long-term stability where the 1 pulse-per-second from the GPS trained the OCXO. By reading the serial port on the GPS every interval, the data collection software incorporated the UTC time in the data file name.

In order to access, modify, and control the system from California, the data collection computer ran under Windows XP. Using Remote Desktop, and a high-speed Internet connection, data could be download and upgrades in the software could be made.

Summary data of both the 200 Hz and 35 kHz bandwidth channels were recorded which included maximum, mean, rms, and voltage deviation values for the interval. Additionally, the log-envelope outputs of both channels made two additional data records. Due to a shortage of available ICS-652 cards, these channels were dropped from the system in order to build a second unit to take to Langmuir Laboratory. Figure 6.6 shows the final configuration for Oklahoma; the Langmuir 2005 system is similar except there was only the first two data channels.

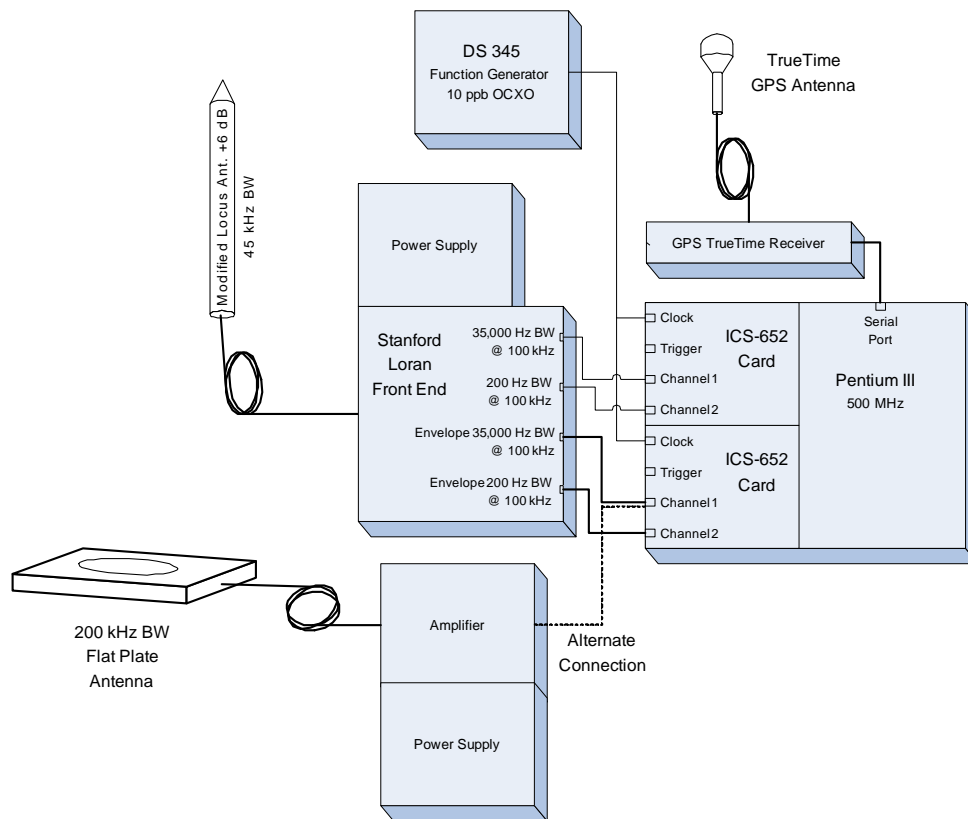


Figure 6.6: University of Oklahoma 2005 Loran band data collection system. The Langmuir 2005 system is similar, but with only the first two data channels.

Chapter 7

Data Comparison to Atmospheric Physics

A contribution of this dissertation is the linking of the ITU atmospheric noise data with that of atmospheric physics. This linking will be accomplished by comparing measurements of return stroke field strengths with the work performed by atmospheric physicists such as Serhan [43], Weidman [44], and Preta [45].

The comparisons begin with the wide-band measurements from the flat-plate antenna. Once these measurements are shown to be accurate, the measurements made by the 35 kHz SLR can be verified. These devices are described in Chapter 6 and their calibration procedure is described in Appendix C.

As a final check, return stroke lightning data was gathered from the National Lightning Detection Network (NLDN). The NLDN provides estimates of return stroke current strengths and locations across CONUS. A comparison of the Stanford data with the NLDN data provides an additional level of cross-checking through an independent means to show that the Stanford Receiver accurately records the absolute magnitudes of the electric fields.

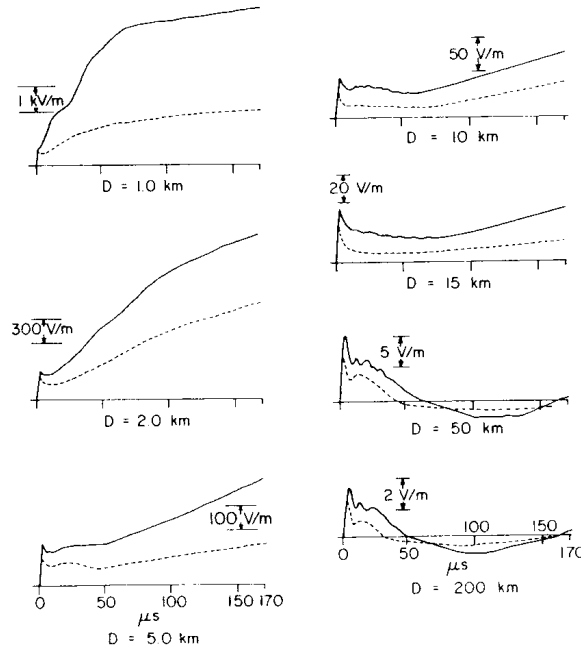


Figure 7.1: Reprint of Fig. 3.2 for comparison [45].

7.1 Time Domain Comparison

Recall that the flat-plate antenna measures the electric fields from 10 kHz to 250 kHz. Atmospheric physicists, such as Preta in [45], have traditionally used such wide-band antennas for return stroke electric field measurements, so a comparison may be readily made between the Stanford data and Preta's.

Figure 7.1, a reproduction of Figure 3.2, depicts Preta's measured electric fields for individual return strokes at various distances, while Figure 7.2 shows the electric field values obtained with the flat-plate antenna. The general shape of Stanford's time domain waveform corresponds well to Preta's data for the respective strike distances. However, the absolute magnitude is off for near-by return strokes by more than an order of magnitude.




Figure 7.2: Electric fields for return strokes at various distances.

Recall that Preta's data collection system had a bandwidth from 1 kHz to 1.5 MHz, and therefore captured more low and high frequency content of the return stroke than the flat-plate antenna which had a bandwidth from 10 kHz to 250 kHz. The difference in low frequency content accounts for the difference in magnitude since the low frequency signals influence the peak amplitude of the return stroke the most for nearby lightning.

7.2 Frequency Domain Comparison

Fourier Transforms convert data from the time domain into the frequency domain and are used to convert the return stroke time traces to spectral data. This transformation may be achieved by a variety of mathematical definitions. The only requirement of a given definition is that the transforms to and from the time and frequency domains be reversible. With such flexibility, researchers have proposed a variety of methods, each lending itself more useful in some situations than in others. However, with such variability arises potential confusion. Appendix A examines the differences in mathematics and physical interpretations between two such definitions of the Discrete Fourier Transform (DFT): those of Preta, used predominantly in the field of Atmospheric Physics, and those of Bracewell, whose definition is well known in the field of Signal Processing.

Bracewell's definition [91], shown here in Equation (7.1), lends itself well to periodic signals or stationary random noise since it averages data over the data interval.

$$E_{\text{Bracewell}}(mf_{bin}) = \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{E}(nT_{step}) e^{-2\pi mf_{bin}nT_{step}} \quad (7.1)$$

where,

$E_{\text{Bracewell}}$ is Bracewell's Fourier Transform of the electric field, \mathcal{E} ,

T_{step} is the interval between data samples,

$f_{bin} = \frac{1}{T_{total}} = \frac{1}{NT_{step}}$ is the frequency bin spacing or the reciprocal of the total record duration T_{total} ,

N is the number of data samples in the record,

$m \in [0...N]$ is the frequency bin index, and

$n \in [0...N]$ is the sample index.

Preta's definition, repeated here from Equation (3.1) for reference, is better at capturing impulsive events since it simply integrates the waveform over the entire interval of the data record.

$$E_{\text{preta}}(mf_{bin}) = T_{step} \sum_{n=0}^{N-1} \mathcal{E}(nT_{step}) e^{-2\pi mf_{bin}nT_{step}}$$

The differences between the two definitions are akin to the differences between power and energy. Average power is defined as the average energy over a given period and, in the limit, as the period becomes infinitesimal, instantaneous power becomes the time rate of change of energy. Conversion between the two Fourier definitions is accomplished by changing the term leading the summation in Equation (3.1) from T_{step} to $1/N$, making it identical to Equation (7.1). Throughout this dissertation, the phrase “average power” requires some explanation. In general, when working with noise voltages, the rms value of the noise voltage is squared to get the average noise power. In the case of atmospheric noise, the electric field strength rather than a true voltage is used, and the corresponding power term will be the square of the electric field strength. Therefore, the noise rms electric field strength and average noise power are interchangeable, since they are proportional to one another and related through the effective antenna length.

The distinction between the two transform methods is necessary since the units for Preta's definition may be misinterpreted by the signal processing community who are more familiar with Bracewell's definition. The possible misinterpretation stems from Preta's choice of units for the frequency spectrum of $V/m/Hz$. Unfortunately, the rms electric field value is not simply calculated by integrating the frequency spectrum over the bandwidth of the receiver, which would be the inclination for those engineers familiar with power spectral density (PSD) curves. This procedure requires more care, and Appendix A details the steps necessary to perform the conversion from Preta's data into a PSD. As implied by Section 5.4.3, a key step in the conversion is to account for the duration of the original data record.

Because his definition is traditional for atmospheric physicists, Preta's definition for the Fourier Transform will be used to transform the Stanford data. A return stroke from 14.7 km away is used here as an example for comparison. Figure 7.3 shows 400 μs of the return stroke electric field captured by the flat-plate antenna.

While Preta used 180 μs data records, the Stanford data record will be extended to 400 μs since the electric-field returns to nearly zero by that time. Without this allowance, additional high-frequency content to the signal would be introduced. This effect stems from the replication of the signal that occurs when the DFT of the data is taken, as described in Appendix A, and the resulting step that would occur if the beginning and ending values of the sequence were unequal. Since Preta's definition is independent of the interval length and only depends on the total amount of energy in the data, and since the values roll off slowly after 180 μs , this extension of the data record introduces a little extra low-frequency energy, but almost nothing in the Loran band.

After applying the scale factor obtained in Appendix C, Preta's DFT is applied to the




Figure 7.3: Lightning electric field return stroke captured by flat-plate antenna.

return stroke data and produces the spectrum shown in Figure 7.4. Preta found that return stroke spectrum varied inversely with frequency, f . Superimposed in red is the $1/f$ relationship referenced to Preta's 7 km data value.

In the figure, the Stanford data falls a little below Preta's prediction, but in general matches to within 5 dB. Additionally, the reduced bandwidth of the flat-plate antenna relative to that of Preta's incurs quicker roll off at both the low and high ends.

7.3 NLDN Comparison

Only a few lightning return stroke electric fields were captured with the flat-plate antenna that could be correlated to National Lightning Detection Network (NLDN) data. By using the NLDN data, range information of the cloud-to-ground return strokes was obtained and provided a means of examining how spectral coefficients vary with distance. Plotting the




Figure 7.4: Electric field frequency spectrum of a single return stroke pictured in Fig. 7.3 using Preta’s method.

100 kHz coefficients from Preta’s curves versus range, r , shows that at a distance greater than about 3 wavelengths, radiative electric fields vary as $1/r$. Figure 7.5 shows that the Stanford data follow this trend after applying the calibration procedure of Appendix C.

7.4 Verification of the Stanford Loran Receiver

Appendix C, shows the calibration of the Stanford Loran Receiver using the anticipated signal strength from the Boise City tower and the calibration of the flat-plate antenna data through the use of the NLDN and Preta’s data. Since these are two independent methods of calibration, they can be compared to check their validity.

The return stroke of Figure 7.3 was captured simultaneously using both the Stanford Loran Receiver and the flat-plate antenna. Digitally passing the return stroke’s wide-band




Figure 7.5: Return stroke electric field frequency spectrum at 100 kHz versus distance using Preta's method.

signal from the flat-plate antenna through a simulated 35 kHz bandpass filter at 100 kHz produces a reference waveform. Theoretically, this reference waveform represents the output from an ideal Loran receiver due to the return stroke. The top graph in Figure 7.6 shows the calibrated, wide-band return stroke waveform captured by the flat-plate antenna in blue and the simulated reference waveform to the return stroke in red. The bottom graph in Figure 7.6 repeats the simulated reference waveform in red, and then superimposes the actual envelope measurements obtained by the Stanford Receiver.

From the plot, the actual Loran envelope closely approximates the envelope of the filtered flat-plate antenna reference waveform. This comparison illustrates that scaling the electric field based on both the predicted signal strength of the Boise City tower and on the propagation models corresponds well to the scaling derived through flat-plate antenna measurements, NLDN, and Preta's data. The independent confirmation shows that the field




Figure 7.6: Time domain traces of a return stroke and a return stroke filtered through a Loran bandpass filter (35 kHz wide bandpass filter centered at 100 kHz).

strength measurements made by either system are accurate and reflect absolute electric field values. With absolute measurements of field strength and power rather than relative measurements, the data from the SLR may be compared to the ITU data.

7.5 Extension to Power Spectral Density

This section covers a minor contribution of this dissertation by showing the conversion from Preta's data to power spectral density. While the details are reserved for Appendix A, applying Preta's Fourier Transform to return stroke electric field data does not directly produce the power spectral density (PSD) of the field. Instead of the power of the pulse, his transform produces something more akin to the total energy of the pulse, making it less familiar to signal processing engineers. Signal processing engineers prefer the use of

noise power, since the signal power to noise power ratio (SNR) is a metric prevalent in the performance analysis of communication and navigation systems.

7.5.1 Return Stroke Time Interval

From Section 5.4.3, average power is defined as the average of instantaneous power over a given interval of time. It may also be defined as the total energy of an interval divided by the duration of that interval, or as the square of the rms value over the interval.

All of these methods require a given interval over which the calculation will be performed. In the case of periodic signals, that interval is typically chosen to be one cycle long. However, for return strokes, which are impulsive, there is no distinctive period since the strokes occur randomly spaced in time. Because there is no clearly defined interval, expressing the power requires the definition of a time interval over which the power measurement is made.

The choice of interval is arbitrary. However, several intervals would be natural: the duration of a lightning stroke, the time constant of a receiver's bandwidth, a longer period of time where the noise may take on stationary or quasi-stationary behavior, or a common agreed upon time scale. Of these possibilities, a time scale of $180\ \mu\text{s}$ has been chosen for several reasons. First, in reviewing Preta's return strokes in Figure 3.2 and the flat-plate antenna's return strokes in Figure 7.2, most of the energy of the stroke resides in the first $180\ \mu\text{s}$ of the stroke, so the chosen interval also is on the order of a lightning stroke. Second, $180\ \mu\text{s}$ is roughly the response time of a receiver's bandpass filter as confirmed in Figure 7.6, where the measured envelope from the Loran receiver is almost zero $200\ \mu\text{s}$ after the return stroke. Finally, $180\ \mu\text{s}$ is identical to the data record lengths recorded by Preta [45]. Therefore, when using his data, all derived power measurements must take into account the




Figure 7.7: Power spectral density (PSD) of return stroke from Figure 7.3.

data record length of $180 \mu\text{s}$.

While this dissertation will use an interval of $180 \mu\text{s}$ for the calculation of individual return stroke power, references like ITU P372-7 [29] use intervals from 15 minutes to 1 hour for atmospheric noise power. Such long intervals will have the effect of averaging the energy of the return stroke over the interval. Section 7.7 discusses this effect and shows how the noise power measurement is reduced as the interval of measurement increases.

7.5.2 Return Stroke Power Spectral Density

As shown in Appendix A, given that the interval for the power measurement is $180 \mu\text{s}$, adding 37.4 dB (Hz) to Preta's values produces the PSD. Since the return stroke pictured in Figure 7.3 generated the spectrum of Figure 7.4, then adjusting the latter figure by 37.4 dB (Hz) results in the PSD shown in Figure 7.7 with units of $\text{dB ((V/m)}^2/\text{Hz})$.




Figure 7.8: Power spectral density coefficients for the flat-plate antenna at 100 kHz versus return strokes range.

From the return stroke time history, the period of the sharp rise at the start of the return stroke is quite brief, on the order of a few microseconds. Such a sharp rise is similar to a step function from -1 to 1, whose Fourier Transform has a magnitude proportional to $1/f$ [91]. Figure 7.7 demonstrates this relationship as seen by the red line with a $1/f$ slope overlaid with the spectrum.

Similarly, Figure 7.5 can be adjusted to yield the PSD of the 100 kHz return stroke coefficients for the flat-plate antenna versus range. Converting Figure 7.5 produces Figure 7.8 which now has units of $\text{dB} ((\text{V}/\text{m})^2/\text{Hz})$, consistent with a power spectral density.

From Figure 7.8 lightning strikes within a few kilometers have an rms electric-field exceeding $-30 \text{ dB} ((\text{V}/\text{m})^2/\text{Hz})$ or equivalently, $90 \text{ dB} [(\mu\text{V}/\text{m})^2/\text{Hz}]$. Since the Stanford Loran receiver has a bandwidth of 35,000 Hz, 45.5 dB can be added to the PSD to account for this bandwidth. The resulting field strength of a return stroke would be 135




Figure 7.9: Electric field strength in $\text{dB}(\mu\text{V}/\text{m})$ of return strokes within the Stanford Loran Receiver versus range.

$\text{dB}(\mu\text{V}/\text{m})$ and the peak electric field within that interval will be even stronger. Even as an average over the interval, a field strength this large will wipe out any underlying Loran signals, since a strong Loran signal, from a tower 500 km away, is on the order of $70 \text{ dB}(\mu\text{V}/\text{m})$. Making this last adjustment for bandwidth converts Figure 7.8 to Figure 7.9, which shows the anticipated electric field strength from a return stroke as measured in the Stanford Loran Receiver as a function of distance.

Figure 7.9 is a key result which will make the results of Section 8.2 more clear. The plot can be used to relate the power generated over any interval to an equivalent amount of lightning continuously occurring at a range of distances. For example, field strengths of $130 \text{ dB}(\mu\text{V}/\text{m})$ are equivalent to lightning discharges occurring within four kilometers of the transmitter.

7.6 Effect of Return Strokes on Loran

To gain an appreciation for the impact a return stroke has on a Loran signal, the following series of plots is presented. Figure 7.10(a) shows five seconds of electric field noise envelope values collected during a storm on June 13, 2005, 01:52:42 UTC, at the University of Oklahoma with the 35 kHz channel of the SLR. Note that the instantaneous electric field values approach 3 V/m for the Loran band. As a comparison, the following series, Figures 7.10(b)-(d) are close-ups of the data around the 1 second time mark. In Figure 7.10(c), Loran pulses are present as repeating signals with a magnitude of 5 mV/m . These pulses are coming from the Boise City tower which is approximately 513 km away. Finally, in Figure 7.10(d), pulses from Boise City are visible at 0.83 s, and both rates of the Grangeville Towers are present at 0.88 s and 0.91 s. The return stroke of Boise City, the largest Loran signal, is 500 times, or 54 dB, weaker than the nearby return stroke. The Loran signals fell within the first three bits of the ADC used on the SLR, since the SLR had a limited dynamic range and was tuned to capture the high energy return strokes. The result is extensive quantization noise on the Loran signal as seen in the final figure.

Loran signals will typically measure between $100\text{-}10,000 \mu\text{V/m}$, yet lighting strokes measure from $1\text{-}100 \text{ V/m}$. Such overwhelming power from a close return stroke will wipe out a concurrent Loran signal. While it may be possible to estimate the shape of the waveform of the return stroke and use the estimate to cancel it out of the signal, this would require more dynamic range which a practical receiver does not have. Therefore, this dissertation considers a Loran signal effectively destroyed should it occur simultaneously with such a large return stroke.

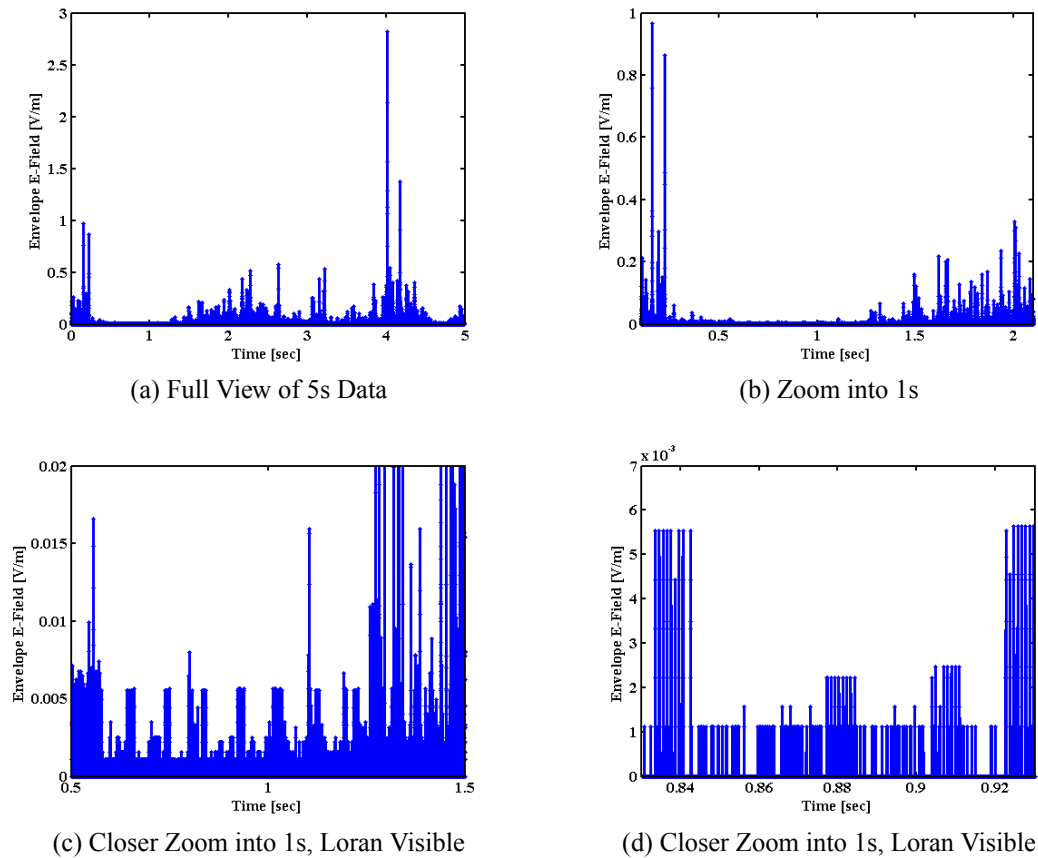


Figure 7.10: Time history of lightning strikes and Loran. Each figure represents different amounts of enlargement on the time scale from [0, 5] sec.




Figure 7.11: Logarithm of Loran band electric field during a storm. Loran signals are visible around 1 second.

7.7 Time Variation of Power Level

Electric fields from Loran are much weaker than return strokes. Changing the units of Figure 7.10 (a) from V/m to $\mu V/m$ and converting the scale from a linear scale of electric field to a logarithmic one in decibels produces Figure 7.11. With this scaling, the instantaneous electric field values for a five second interval where both the instantaneous return stroke fields (approaching $130 \text{ dB}(\mu V/m)$) and the Boise City Loran tower pulses ($75 \text{ dB}(\mu V/m)$) are clearly seen.

The calculated rms of electric field strength over the five second interval is $83 \text{ dB}(\mu V/m)$, almost 50 dB lower than the peak instantaneous electric field strength. Using the models of Section 3.2, ITU predicts a median 95% noise value of $82 \text{ dB}(\mu V/m)$. If either of the rms noise values is used to determine the SNR of Boise City then the SNR is approximately -9 dB.

Section 5.2 showed how Gaussian noise increases the probability of the receiver selecting the wrong cycle. For the typical receiver integration times of 20 s, the lowest SNR for practical acquisition is approximately -10 dB given no other improvements. Applying this rule to the noise pictured in the Figure 7.11, the conclusion drawn would be that the receiver must be on the edge of not being able to track any Loran towers over this interval. However, around the 1 second point the Loran signals are clearly visible, and at that time the reception would be quite good. This example reiterates the fact that non-Gaussian noise cannot be adequately described by rms or average power alone.

Due to the non-stationary nature of the noise, the rms electric field of the noise can vary by 50 dB depending on which interval the average is performed. Such differences between the instantaneous rms electric field strength and the rms taken over a five second interval are due to the statistics of the noise which are governed by the physics of the lightning process and the rate at which atmospheric discharges occur.

The sensitivity on the choice of interval length points to the implications of choosing longer intervals, ones which are much longer than a lightning return stroke. As the interval for averaging the power is increased, the effect of any one stroke on power becomes lessened, and the long-term nature of the processes begin to dominate. If the time interval for determining power is broadened to 15 minutes, the power is seen to be reduced. To illustrate this fact, Figure 7.12 depicts 12 hours of rms electric field values of a storm that occurred around 0200 h June 13th, 2005 UTC. The Loran band data was sampled at 50 kHz and saved in five second data records. By concatenating the five second data records, 15 minute records were created which correspond to the interval used in the ITU model's measurements.

At 50 kHz, each data sample lasts 20 μ s. The averaging inherent in the sampling process




Figure 7.12: Range of $20\ \mu\text{s}$ electric fields samples and rms of electric field voltage over 5 s and 15 min intervals during a storm on 13 June 2006 UTC.

is the rms value of the sample over that period. The green, blue, and red traces in Figure 7.12 show the rms electric field values taken over the intervals of $20\ \mu\text{s}$, 5 s, and 15 minutes, respectively. Given the finite thickness of the lines, if all of the data were plotted, the individual values of the $20\ \mu\text{s}$ data would be indiscernible. To reduce the data overhead, these data were reduced to only the maximum and minimum sample values in each five second interval, thus demonstrating the dynamic range of the $20\ \mu\text{s}$ data samples.

Upon inspection, three characteristics stand out for the data set. First, the noise is clearly non-stationary over this twelve hour period. All three interval durations show the waxing and waning of the noise rms electric field strength as the storm approaches around 0200 h and then departs. Second, by increasing the time interval, the rms electric field strength is reduced. Third, averaging over 15 minutes results in rms field strengths that are relatively insensitive to any single return stroke but have become a measure of the intensity

of the storm.

This graph helps to explain the commonality between the ITU model and the atmospheric physics model of individual lightning strokes. This comparison is a major contribution of this dissertation. From Figure 7.9, Preta's model is shown to predict lightning return strokes greater than 130 dB $\mu\text{V}/\text{m}$, depending on the range of the observer. However, ITU models predict median 95% values that are on the order of 85 dB($\mu\text{V}/\text{m}$). Both are correct, provided the interval over which the observation takes place is understood, thus reconciling the models.

Since the requirement of defining an interval for the calculation of power seems so critical, it implies that the SNR calculation must be tempered by understanding the interval over which it is calculated. From the analysis of Loran cycle selection performance in Section 2.1, the most useful period to take the rms measurement was the time over which the pulses were averaged. However, evaluating the credit for non-linear processing requires working within the bandwidth that of processing, which is the RF bandwidth.

Chapter 8

Improved Loran Coverage Model

This chapter begins by reviewing the collected data and comparing it against the ITU model. After verification by the collected data, the ITU model is extended, showing that noise impulsivity is correlated to the rms electric field value. Leveraging this data relationship, the non-linear processing gain is shown to also be proportional to the anticipated rms electric field strength, rather than a constant as the traditional model would suggest. This change to the model improves the SNR by 8 dB over the traditional model at the 99.9% availability level thereby increasing the RNP 0.3 coverage area by 153%.

8.1 Verifying the ITU Model

Section 3.2.6 raised a number of concerns regarding the applicability of the ITU data to a Loran receiver. Those concerns are reiterated here:

1. The receiver bandwidth is 175 times wider than the ITU receiver, which the model cautions may lead to erroneous estimates of noise and impulsivity;

2. Loran frequencies are in the LF region, which again the ITU model cautions may be too low for accurate estimates;
3. The ITU researchers turned off their equipment during nearby lightning storms, so there is uncertainty as to the validity of their data during a storm; and
4. The ITU model had questions regarding the validity of the APDs because of their limited data.

Each of these concerns will be addressed in the following sections.

8.2 Estimates of RMS Electric Field Strength

As described in Section 6.3, the Stanford Loran Receiver collected noise envelope electric field strength data, E_{rms} , in Norman, Oklahoma almost continuously from the end of May until mid-August 2005. The collection system recorded five second intervals of rms field strengths and the corresponding time tags.

Section 5.4.3 states that when giving an rms value for noise, it is necessary to specify the interval over which the measurement is taken. Though not explicitly stated by ITU, from their measurement technique described in Section 3.2.3, the researchers appeared to have taken the rms noise strength values over a 15 minute interval. Therefore, to be comparable with ITU, 180 five-second data records were concatenated before calculating the rms noise strength over their entirety. Figure 8.1 shows these data accompanied by the log-normal distributions predicted by the ITU model.

To be consistent with ITU, the rms data were partitioned into each of the six time blocks, and the distribution of the data for each block was calculated. According to the ITU model




Figure 8.1: Distribution of envelope noise measurements collected in Norman, Oklahoma over the summer of 2005. The blue and red curves show the log-normal distribution predicted by the ITU model above and below the median value.

and as shown in Figure 3.6, for each time block the rms envelope noise field strength should follow a log-normal distribution. Figure 8.1 shows the ITU model's predicted distribution compared to the collected data for that time block. Comparing the received field strength measurements with the ITU model requires using absolute power measurements, not just relative ones; this need drove the system calibration procedure described in Appendix C.

Plotting the probability density function on a linear scale demonstrates how well the data fit the model for the middle values. However, it is more instructive to convert the data to a CDF and to replot them as shown in Figure 8.2 to see the tail of the distribution and to see how the rms noise strength varies as a percentage of time for extreme values. The figure's y-axis is not just interpreted as a probability, but is the percentage of time that the noise strength will be at or below the corresponding x-axis value. The x-axis then relates




Figure 8.2: Comparison of ITU and long term RMS data taken from the 2005 data collection campaign. Time blocks are given in local time.

to the availability and the 99.9% value would correspond to the RNP 0.3 requirement.

Due to the limitations imposed by the SLR's design, which will be described shortly, only a small portion of the blue curve is visible near the 50 dB($\mu\text{V}/\text{m}$) point. Extreme noise values drive the coverage predictions (i.e., these are values greater than the median 50% expected value) and are denoted by the red curve which represents the upper median values of the worst case time block for Norman, Oklahoma over all seasons. The black dashed and dotted lines vary this prediction by $\sigma_{F_{am}}$ in both directions. The blue traces with the various symbols identify each of the different time blocks with the times given in local time. The data pertaining to the afternoon through late evening has some of the strongest values corresponding to local storm patterns, while midday data can be more than 10 dB lower.

The lowest noise strength recorded is around 56 dB($\mu\text{V}/\text{m}$). Accumulation of data

between 56 and 65 dB($\mu\text{V}/\text{m}$) stems from the details of the SLR's design and the presence of the Loran towers in the measurements. Since the SLR was designed to capture high energy pulses, its sensitivity was biased accordingly, compromising the receiver's ability to measure low level noise. The SLR's sensitivity, which varied from 863 to 1106 $\mu\text{V}/\text{m}$ per count, made quantization noise an issue. Due to quantization, the first bit of the ADC incorporated noise field strengths up to a range of 58.7 to 60.9 dB($\mu\text{V}/\text{m}$). This range of data stems from variations in the experimental setup over time. When the quantization error of the ADC couples with the noise internal to the ADC and the total strength is greater than half of a least significant bit, the first bit will flip. Thus, the receiver cannot measure noise field strengths below a range of 52.7 to 54.9 dB($\mu\text{V}/\text{m}$).

An additional factor limiting the SLR's ability to measure low value noise, and the accumulation of data near 60 dB($\mu\text{V}/\text{m}$), arises from the Loran Towers themselves. Since the SLR receives the Loran Towers pulses, the tower's output contributes to the total "noise" power. The Boise City tower strength measured 75 dB($\mu\text{V}/\text{m}$) with a 2 dB variation. The variation is presumed to be something or someone disturbing the setup that changed the grounding slightly. These variations manifest themselves as jumps in the floor of the rms data values on June 11 and July 28.

Since the duty cycle of a Loran tower is roughly 2%, the Boise City tower alone contributes 57 dB($\mu\text{V}/\text{m}$) to the noise measurement. Even on quiet afternoons, when only the Loran towers and low level background noise from distant storms were present, the towers and background noise accounted for 58.5-60.6 dB($\mu\text{V}/\text{m}$) of the overall rms field strength.

The combination of the Loran power present in the band and the low sensitivity of the receiver places a lower limit on the minimum field strengths the SLR can measure. Since

the receiver cannot effectively measure noise much lower than the median 50% value, all data points which would be below this accumulate in the low-end range of 55-61 dB. Only noise occurring over the interval that is stronger than the Loran signals present is accurate in representing atmospheric noise since the towers only contribute a small amount to the overall noise. The presence of Loran signals in the measurement band limits the SLR's ability to collect low level data. Future efforts may include blanking out the Loran signal or measuring noise just off of band to compensate.

While the low-end is affected, after 70 dB($\mu\text{V}/\text{m}$) the noise is well bounded by the ITU model. From the y-axis of Figure 8.2, noise exceeding 70 dB($\mu\text{V}/\text{m}$) should only happen about 7% of the time. The tracking of the CCIR curve continues up until about 99.5% of the time. After this point, the data distribution of all time blocks begins to roll-off. This phenomenon is likely to be physical in nature.

Data above 99.5% is only generated by very close thunderstorms. The lightning rate is limited by the physical mechanism of charge re-distribution within a cloud. Since the power measurement intervals are 15 minutes in duration, there is a limit as to how many close strikes can occur within a nearby region. So, the rate at which strokes occur, as well as the limit on the density of clouds capable of such strokes close to the receiver, will limit the amount of power received.

During these 15 minute intervals, the SLR recorded lightning stroke electric fields exceeding 135 dB($\mu\text{V}/\text{m}$). Figure 7.12 shows how the instantaneous, the five second, and 15 minute rms values vary with respect to each other. Therefore, the roll-off is more indicative of the rate at which lightning occurs rather than the maximum value the receiver can record.

The collected data corresponding well to the ITU model addresses Concerns 1, 2, and

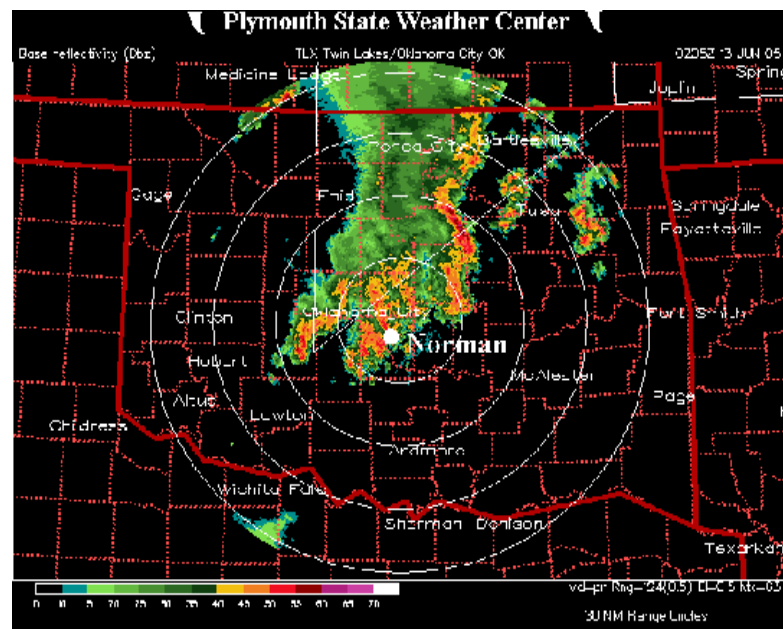


Figure 8.3: Radar plot of storm across Oklahoma on 13 June 2005 at 0205 UTC. Courtesy of Plymouth State Weather Center.

3 from the start of the chapter. Despite the limited data available, the ITU model does predict the data well up to about the 99.5% noise level. Above the 99.9% noise level, the ITU model is conservative since the lightning process seems to be limited, perhaps to some physical phenomena limiting the rate of return strokes in a cloud. In addition, based on the storms, the 99.5% noise level represents strong nearby lightning storms, with 99.9% as a conservative over bound providing 5 dB of margin.

One of the worst storms during the 2005 campaign occurred on June 13th and managed to achieve a 15 minute rms noise field strength of 84 dB($\mu\text{V}/\text{m}$). This corresponds to the 99.5% level of Figure 8.2. Figure 8.3 shows the radar plot from Twin Lakes, OK during the time period when the frontal system ranged over 100 NM in length and swept over the data collection station. While this may not be the largest storm ever, it shows a storm large enough to make flying difficult if not impossible.

The 99.9% predicted rms noise field strength of $89 \text{ dB}(\mu\text{V}/\text{m})$ seems to be a sufficient over bound of a 15 minute interval. While such a high field strength relative to the closest Loran tower ($75 \text{ dB}(\mu\text{V}/\text{m})$) would cast doubt on the efficacy of a receiver in this environment, the rms noise values alone are merely a starting point for evaluating the performance of a receiver. As the example of Section 5.4.3 shows, the distribution of the noise as well as the rms value is important.

8.3 Instantaneous Noise Envelope APDs

The previous section addressed Concerns 1 through 3, by assuring the applicability of the ITU model to the Loran band. This section addresses the final concern, the ability of the ITU model to accurately describe instantaneous envelope noise field strength distributions.

Recall from Section 3.2.5, that the APD provides the distribution of the instantaneous envelope noise voltages or field strengths for the interval. Furthermore, the curves are normalized to the rms value of the interval and are parameterized by the voltage deviation, V_d , of the interval.

By digitizing the data provided by ITU, in between values of V_d can be interpolated and the APDs can be derived for arbitrary V_d . Figure 8.4 shows in blue dots the APD of a five second data file taken on June 13th during an intense thunderstorm. Given the data record, the V_d of the interval was calculated to be 18.2 while having rms value of $84.4 \text{ dB}(\mu\text{V}/\text{m})$. Plotting the interpolated curves generated by the ITU data, the predicted APD for the given V_d is shown in black. Two observations are apparent. First, the data points in blue fall remarkably close to the theoretical curve; and second, the data fall off at the 30% level due to sensitivity of the receiver.




Figure 8.4: APD of instantaneous noise envelope voltages for a single five second data record during a storm. Data drops off for low amplitudes due to the dynamic range of the receiver.




Figure 8.5: APD of instantaneous noise envelope voltages for a single five second data record during a quiet day. Data drops off for low amplitudes due to the dynamic range of the receiver.

This result is typical for high strength noise; the high amplitude values follow the predicted curves well. However, for low strength noise, the Loran signals present in the band play more of an influence, as in Figure 8.5, where V_d is 10.3 for an rms of 64.6 dB ($\mu\text{V}/\text{m}$). In this figure, the Loran towers are skewing the data around the 0.1 to 5 percent levels. The low noise data may not match as well to the theoretical APD. However, this is not of much concern since at high SNR Loran signal acquisition is reliable.

The APDs predicted by CCIR for a given V_d do indeed seem accurate for large noise values, and so Concern 4 has been partially answered. The remainder of the concern is how well does the ITU model predict the range of V_d in the first place?

8.4 Correlation of Voltage Deviation and RMS Noise Field Strength

The previous sections confirmed that the ITU model does well in capturing both the rms and instantaneous noise envelope electric field strengths in the Loran band for high strength noise. This section extends the results of the ITU model and in doing so achieves another contribution of this dissertation.

ITU claims that there is a weak correlation between the voltage deviation, V_d , and noise envelope electric field measurements, E_{rms} . However, they did not comment further. Here, not only will the correlation be shown, but also that a lower bound of V_d may be predicted for a given value of E_{rms} .

Section 5.3 showed that the shape of the distribution determines the effectiveness of non-linear signal processing elements such as hole-punching or clipping since most of the energy of the interval is tied up in noise spikes. The more impulsive the noise, the larger its V_d .

Using the 2005 Norman, Oklahoma data collected over the summer, the five second and fifteen minute rms envelope noise field strength measurements versus the voltage deviation in Figure 8.6 are plotted. The five second rms electric field measurements are shown in blue, while those of the concatenated 15 minute long records are shown in green. For any given value of electric field strength, E_{rms} , the variation of V_d can span more than 10 dB, but there still is a clear trend that as E_{rms} increases, so does V_d .

The red dashed line drawn through the center of the green distribution is an estimate of how V_d varies with E_{rms} by




Figure 8.6: Distribution of rms electric field strength (E_{rms}) and the voltage deviation (V_d) from collected data. 5 s interval data in blue and 15 min interval data in green.

$$V_d = 0.3 E_{rms} - 8 \text{ dB.} \quad (8.1)$$

This is the “Median V_d Estimate,” since for a given value of E_{rms} , approximately half the values will be lower bounded by this curve. Similarly, the solid red line is defined by

$$V_d = 0.3 E_{rms} - 15 \text{ dB.} \quad (8.2)$$

This provides a lower bound for V_d , given E_{rms} values above 70 dB($\mu\text{V}/\text{m}$), and is the “Pessimistic V_d Bound.” It represents a bound on the lowest value of V_d for a given E_{rms} .

Even with a lower bound on V_d , the predicted value of E_{rms} is high, corresponding to a high amount of impulsive noise. Using the signal processing techniques described in Section 5.3.5, a lower bound on processing gain may be predicted that is proportional to

the anticipated noise level, and is given by

$$\text{Non-linear Processing Gain} = 0.54 E_{rms} - 28.1 \text{ dB}.$$

8.5 Improved Coverage Results

Now all of the pieces are in place to apply these results to the Loran coverage tool. The processing gain works as follows:

1. Given that the ITU model has been shown to be valid, it can be used to accurately predict the rms envelope noise field strength, E_{rms} .
2. From the contribution linking V_d and E_{rms} , a lower bound on V_d given E_{rms} across the country is predicted.
3. Given that the APDs for a given V_d have also been validated, the APDs can be used to calculate a non-linear processing gain.
4. The rms envelope noise field strength may be reduced by the non-linear processing gain, yielding a lower post-processed noise
5. With the reduced noise, the SNR is calculated and used in the coverage model to determine RNP 0.3 availability.

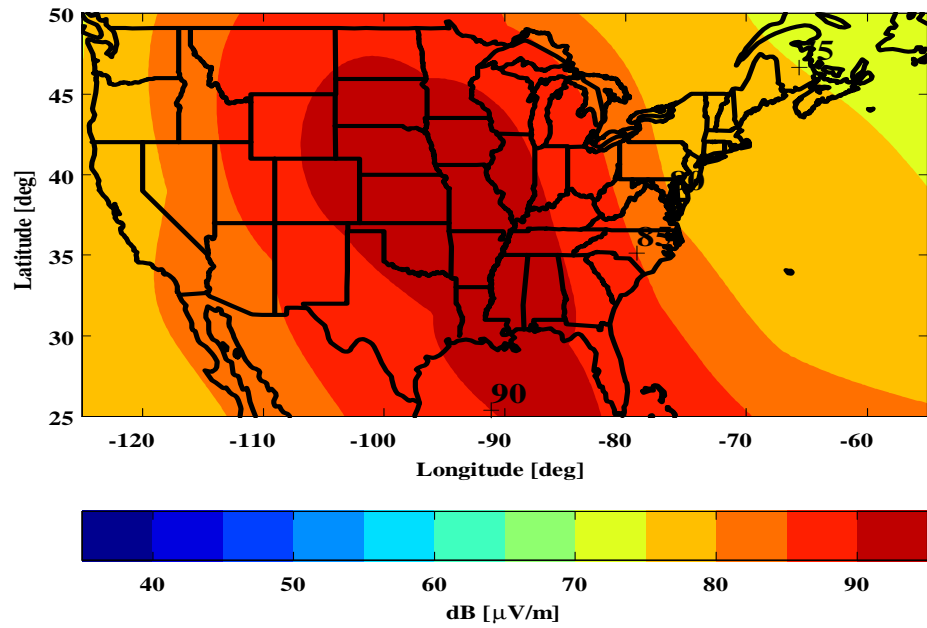
In following these steps, the ITU predictions for the worst-case season and time block at the 99.9% rms noise envelope field strength are used to determine the noise corresponding to the 99.9% availability percentage. With the ITU noise estimates Equation (8.2) is used, which predicts the lower bound on V_d given E_{rms} at this availability level. Figure 8.7(a)

is a repeat of Figure 3.7 while Figure 8.7 (b) shows the lower-bound estimates of V_d across the country. Note here that because of the development of Section 8.4, V_d can now be modeled as varying spatially since it is proportional to the anticipated noise level. This spatial variation is in contrast to the constant value that the ITU model gives for each season's time block.

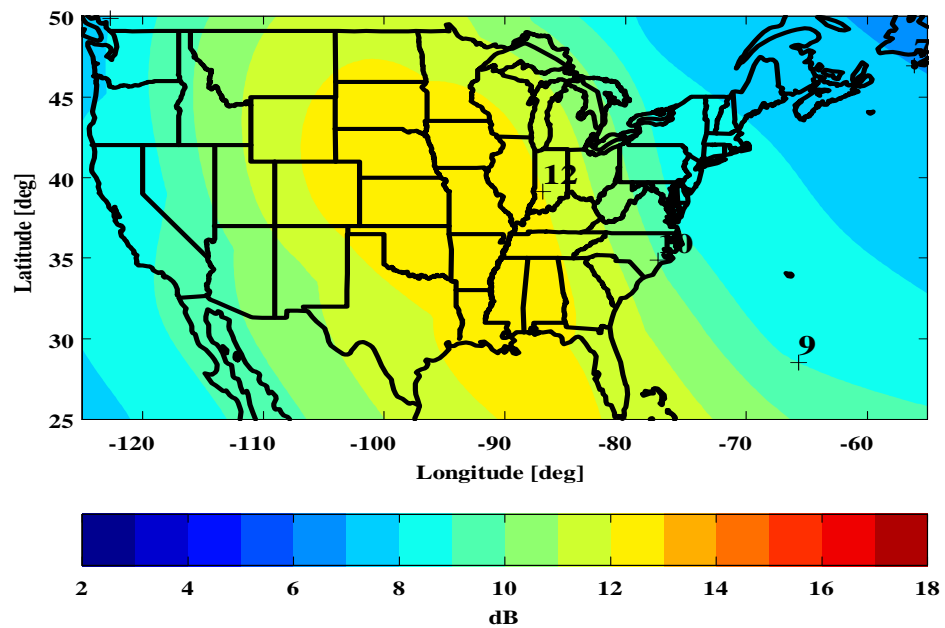
With the lower-bound V_d estimate, the non-linear processing gain is estimated from Equation (5.7). The 99.9% rms noise envelope field strength is then reduced by the processing gain to produce Figure 8.8. Even using the lower-bound on processing gain, improvements of the unprocessed SNR of more than 20 dB are achieved.

With the reduced noise levels across the country, the number of towers available for a position solution is calculated. Figure 4.2 (b) is repeated in Figure 8.9 (a) for comparison, while Figure 8.9 (b) shows the number of towers available under the same noise conditions but after non-linear signal processing of the noise. The improvements in the number of towers goes from few or no towers available across most of the Plains States, to somewhere between three and nine towers.

When applying these results to the coverage tool, the predicted coverage goes from Figure 8.10 (a) to that of Figure 8.10 (b) using the "pessimistic" estimate of V_d . Except for areas along some of the Plains States, Great Lakes, Mississippi, and the most southern tip of Florida, the coverage has been improved from 80% to 99.5%. Still, areas within the center of the country remain at the 95% availability level. Comparing the areas of 99.9% coverage between the two figures shows a 153% increase from the traditional model to the new coverage model.



(a) ITU Estimated Worst-Case 99.9% Noise Strength

(b) Lower-bounds of V_d **Figure 8.7:** V_d derived from noise estimates.

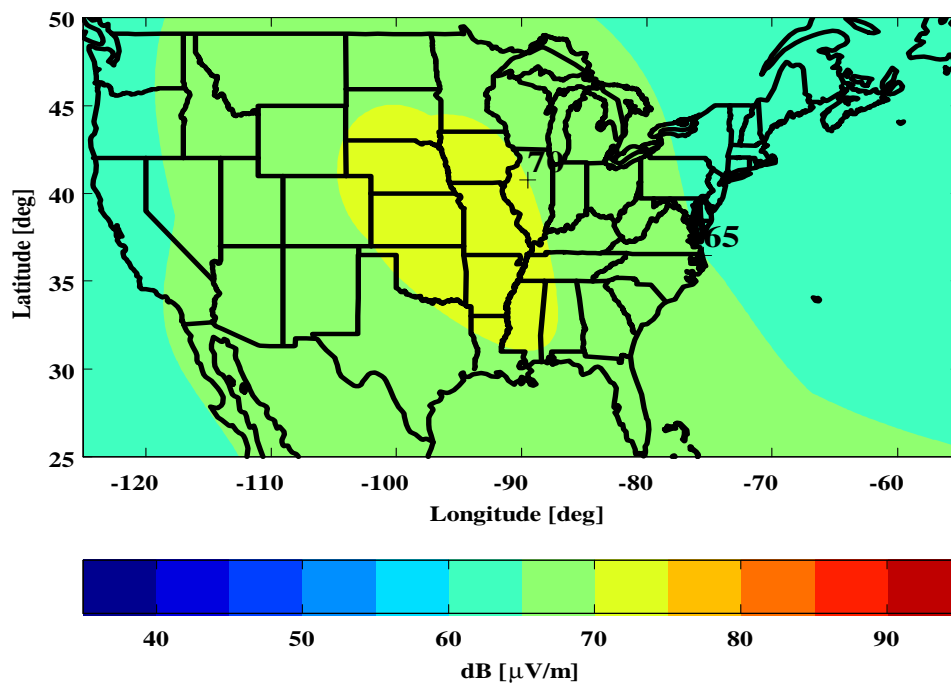
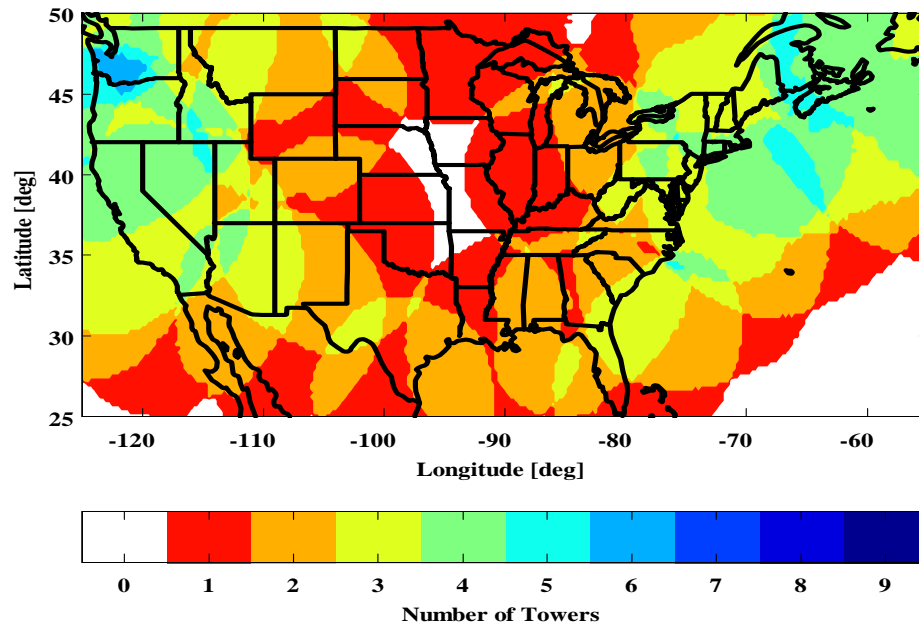



Figure 8.8: Resulting noise after non-linear processing.



(a) Constant 12 dB Clipping Credit



(b) New Clipping Credit

Figure 8.9: Number of towers available for a position solution.

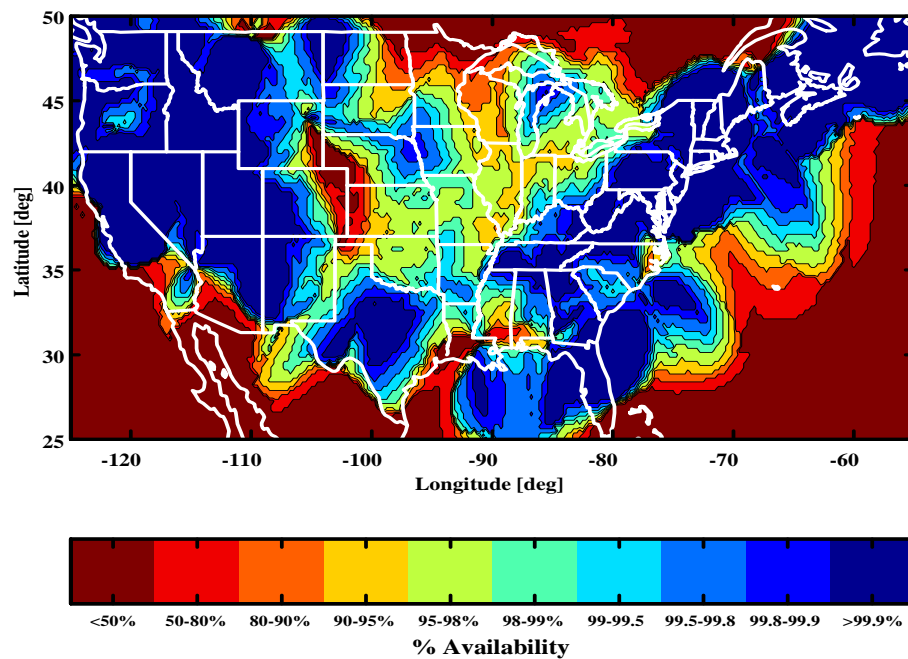
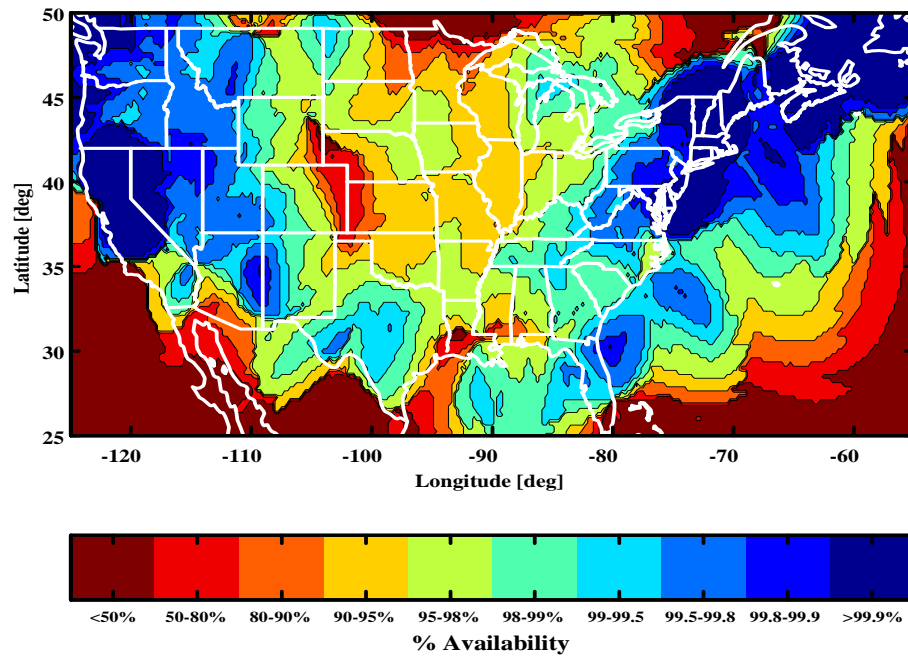


Figure 8.10: Comparison of traditional and improved coverage prediction charts.

Chapter 9

Conclusion and Future Work

Based on the results of the previous chapter, a number of conclusions can be drawn regarding the validity of the ITU model, the effects of signal processing in a non-Gaussian environment, and the improvement of Loran coverage through the use of the new model developed in this dissertation. After reviewing the implications of the results, suggestions for future research will be made.

9.1 ITU Model

Section 5.4.3 described how power calculations for non-Gaussian or impulsive noise required that the interval for the calculation be specified. Having defined the interval as $180\ \mu\text{s}$, Section 7.7 led to the first contribution of this dissertation, and showed that the models developed by atmospheric physicists can be transformed to a similar format as that of the ITU model through the correct understanding of power. Comparison confirmed the agreement of these two models on both short and long term time scales.

When considering the manner in which ITU collected their data (described in Section 3.2), while not explicitly stated, the interval for CCIR should be taken as 15 minutes. Given a 15 minute interval, the rms noise envelope distributions of the Stanford data agree with the ITU model for amplitudes corresponding to 70% through 99% probability amplitudes. By the time the noise level reaches 99.9%, the physics of the atmospheric noise process involved, specifically the minimum spacing between lightning strokes, curtails the noise such that the ITU model is about 5 dB conservative. Therefore, the validity of the ITU models for the long-term noise measurements for the Loran band has been shown to be adequate to the 99.9%. Beyond the 99.9% level, the predictions by the ITU model are shown to be overly conservative and so requiring Loran coverage and availability predictions to use the ITU model is excessively conservative.

By breaking up the collected data into five-second long records, the ITU APDs were found to have accurately described the instantaneous noise envelopes for high amplitude signals even though the bandwidth of the SLR was much larger than that viewed by the ITU model as acceptable. This confirmation of the ITU model for short term noise distributions is another contribution of this dissertation.

In order to verify the applicability of the ITU model to Loran despite the caveats mentioned in Section 8.1, this dissertation extends the ITU model by showing the correlation between the rms noise field strength, E_{rms} , and the Voltage Deviation, V_d . These two quantities were related using a simple line, given in Equation (8.2), and drawn below the lowest data point whose field strength was greater than 70 dB($\mu\text{V}/\text{m}$). This lower bound should produce conservative estimates of V_d and hence conservative estimates of the gains due to non-linear signal processing.

9.2 Signal Processing

Using the ITU APDs of the instantaneous noise envelope, the noise was separated into a high-amplitude impulsive component and a low-amplitude Rayleigh component. By using non-linear signal processing, duty cycle was traded for noise reduction and improvements to the SNR of a Loran receiver were made. This dissertation determined that the non-linear processing gain was proportional to the rms noise field strength, E_{rms} , and not just a simple constant as indicated by the traditional models. This finding is crucial to the accurate modeling of Loran coverage since the lower bound on the non-linear signal processing gain is on the order of 20 dB for higher noise levels rather than a flat gain of 12 dB as described in the prior art.

Such processing gain is inherent in many current receiver designs. Clipping comes about through the limited range of the ADC or by the saturation of the receiver components themselves. Therefore, standard Loran receivers using clipping should be capable of realizing these gains.

9.3 Loran Coverage

When accounting for non-linear processing gains through the use of ITU models, the SNR improved by 8 dB and coverage increased from 90% across CONUS to 95% across most of CONUS. This is the most significant contribution of this dissertation. Additionally, the 99.9% availability coverage area has been increase by 153% from the traditional model. While the increase is impressive, the coverage still falls short of the desired FAA requirements of 99.9% across all of CONUS.

In developing the Loran coverage model, the lack of towers and the probability of wrong

cycle selection was found to limit the availability of Loran. As seen from Figure 8.9 (b), at the 99.9% availability level, there are just barely enough towers available at the center of the country to provide a position fix, let alone support RNP 0.3. The lack of towers in that region has been called the "mid-continent gap." This analysis implies that an additional tower or two may be required to fill in the gap. Only a few towers are likely to be necessary to increase the number of strong stations to three in the troubled area.

Also, since the 99.9% level predicted by ITU is still 5 dB higher than indicated by the data, there may be room for improvement in the coverage prediction if the noise at this level is reduced by this offset. Since the probability of wrong cycle selection is proportional to $1/\sqrt{SNR}$, any reduction of noise or gains in signals would be of benefit.

9.4 Implications

Improvements to Loran's availability coverage were made possible by this dissertation's contributions, thus pushing Loran closer to validation for non-precision approach. With the combination of this work and the work of other LORIPP members, Loran's future is still hopeful.

Furthermore, the gains from signal processing are achievable using standard Loran receivers, so the benefit is realizable today. Since the coverage does not meet the current FAA requirements for RNP, at least two options present themselves: change the requirements and/or add new towers.

Improvements can be made in coverage if the noise is reduced from the 99.9% level. Reducing the RNP 0.3 requirement down to the 99% availability level, increases coverage 344% from the original area. A reduction of the noise strength may also be due to the limitations of lightning stroke spacing, but further data is required for this justification.

In addition, the severity of storm conditions present at the 99.9% level might exceed the operational boundaries for safe aviation. That is, planes would not fly in such storm conditions voluntarily and should they end up in such conditions, the difficulties of keeping the plane in the air may surpass the worries about accurate navigation. If planes are not expected to operate in such conditions, a lessening of the 99.9% availability requirement can be justified.

While a costly fix, this study implies the need for additional towers. One to two additional towers in the Midwest would improve coverage by increasing the number of towers available in the low coverage areas. It is likely that only a couple of towers are needed since there are Loran signals present in the areas where availability is low, just not enough strong signals to pass the integrity requirements at the 99.9% noise level.

9.5 Future Work

In researching and testing the work for this thesis, opportunities for future research have presented themselves. In particular, the short-comings in the current Loran coverage model pose interesting questions as to the improvement of Loran and the improvement of the coverage model. In this section, some of these issues are discussed.

9.5.1 Improving Loran Coverage

Two areas that could lead to improved Loran coverage are changing to H-field antennas and the addition of more towers. Both of these areas require further research to evaluate their potential effect.

The change from an E-field to an H-field antenna stems from reviewing the data records. While large storms were overhead, low amplitude periodic pulses were present in the data records. The speculation is that corona discharge either at the antenna or from one of the antennas on the roof top contributed this noise to the signal.

An H-field antenna may perform better in this environment as it will be less likely to generate corona due to its more rounded shape. Further testing with a field mill and combined H-field and E-field testing would help to confirm this. As mentioned in Section 1.3, the work done by Cutright [48–51] showing that the H-field antenna is less susceptible to P-static helps to support this speculation.

Further work is also required in examining the impact of additional towers in the center of the country. A tower in Omaha, Nebraska, for instance, would help provide the extra tower necessary in the mid-continent gap. Loran requires three towers to make a position fix. However, more towers enable the use of integrity algorithms which can leverage weaker signals reliably. From Figure 8.9 (b), at the 99.9% noise level only three towers are available in some locations. Therefore, it is likely that the addition of one or two more towers would ensure the integrity of the position solution. Further study is required to determine if multiple low-powered towers or just a single high powered one would fill in the mid-continent gap.

9.5.2 Further Testing

Looking back on the experiment from the perspective that history provides gives a chance to examine some of the strengths and weakness of this dissertation. The greatest strength lies in the quantity of collected data. There were three months of data from Norman, Oklahoma in 2005, which allowed runs of more than 40 hours of storm data to analyze Loran receiver

processing algorithms.

While the quantity of data was large, the limited dynamic range of the receivers plagued the quality of the data. In earlier campaigns, only low level signals were captured, while in 2005 only large amplitude ones were captured.

The Stanford Loran Receiver Revision B has been built to fix this problem. The new receiver has two Loran band channels 35 kHz wide centered at 100 kHz . The new design provides a low-gain channel for recording the high amplitude lightning data and has a separate high-gain channel for capturing the weaker Loran signals. The two data channels are combined in software to improve the dynamic range of the data system to 122 dB. This is accomplished by overlapping the ranges of the two channels, where the range of the high gain channel lies within the first 50 counts of the low gain one. From the receiver's calibration, the minimum signal observable would be approximately $5 \mu\text{V/m}$, or $14 \text{ dB}(\mu\text{V/m})$. The maximum field strength observable would be $136 \text{ dB}(\mu\text{V/m})$, $16 \text{ dB}(\text{V/m})$, or 6.3 V/m .

Furthermore, this modification will enable tracking algorithms to be run on the Loran signal data while simultaneously recording the interval's true rms value. Since the 200 Hz data of the SLR Revision A was found to be of limited use, those data channels have been replaced by two channels for a pair of magnetic field (H-field) antennas from Megapulse.

The new receiver design with an improved dynamic range does show promise. Figure 9.1 is an APD of a one minute data record during a thunderstorm. Overall, the APD reasonably follows the theoretical curves predicted by the ITU model for the high amplitude noise and does become Rayleigh for lower amplitude noise, but tends to be higher in amplitude than predicted. The bump in the curve around 1% is generated by the Loran signals present within the band.




Figure 9.1: APD of instantaneous noise envelope voltages for a single one minute data record during a storm with new Loran Receiver.

9.5.3 Refining the Model

A number of refinements can be made to the coverage model to improve its accuracy. Using the new 2006 storm data, real Loran signals can be processed in real atmospheric noise. Preliminary results show that the theoretical gains predicted by this thesis are overly optimistic for large signals, though adequate for weaker ones.

This is because the non-linear processing threshold level cannot be brought low enough for real Loran signals to take out all of the impulsive noise. Referring back to Figure 9.1, the Boise City tower produces the bump in the APD around 1% and at approximately the rms value of the data. In order to punch or clip the impulsive noise such that only the Rayleigh portion is left, the threshold would need to be set to 20 dB below the rms value which would mean processing about 20% of the time. However, the threshold cannot be set this low without processing out the Boise City tower. Instead, the threshold must be set at a level just above the Boise City tower and process out only 1% of the impulsive noise.

Fortunately, the threshold can vary for each of the towers. Since, the time-of-arrival is approximately known to the receiver once one strong tower is received, weaker signals can have a lower threshold applied to them resulting in a higher percentage of processing and hence, more processing gain.

Therefore, based on the 2006 data, the next generation of non-linear processing gain algorithms should incorporate the signal strength of the incoming tower. Knowing the tower strength relative to the rms of the anticipated noise will result in the maximum percentage of time that the receiver may clip or punch out the noise. After removing the noise, a new rms value may be calculated for the remaining noise. This greatly increases the complexity of the Loran coverage tool, since the signal strength of all towers must be taken into

account at each geographic location in order to determine the processing gain. While computationally intensive, this is achievable with modern computer technology and will be the basis of future Loran work.

Although the low amplitude data does follow a Rayleigh distribution, it is offset from the predicted curve specified by V_d . However, this offset is influenced by the presence of Loran in the band which skews the distribution. Since the Rayleigh portion of the curve is almost 10 dB higher than predicted, the processing gain will not be as large as anticipated. Fortunately, as seen by the bump in the distribution at 1%, the Loran signal is shown to be above the rms value. Therefore, the tower actually has a positive SNR most of the time. Thus, not much gain is required from the non-linear signal processing circuit since the tower has such a strong SNR. Additional data will be required to determine if this is indicative of all of the noise data.

Another area of model refinement is tightening or improving the bound between the rms noise field strength, E_{rms} , and the Voltage Deviation, V_d . If the median rather than the pessimistic bound is used, then coverage is improved since more towers are available for a position solution and the quality of the towers is improved. The choice of bounds was made to be conservative. A policy decision from within the FAA is required to determine the appropriate level to take with this bound. However, noise at 99.9% and higher is indicative of severe storms which would inhibit safe flying. Therefore, a less restrictive bound more in line with reasonable flight conditions seems warranted.

9.5.4 Time Domain Model

The quantifying of noise for this dissertation mainly concerned itself with noise power. However, there is a need for a better time-domain model of lightning noise that can be used

for the testing of Loran receiver designs.

In [42], Uman gives statistics on the rate at which lightning flashes occur, the number of strokes within a flash, and the return stroke's inter-arrival time. Additionally, he provides data on return stroke event statistics such as the duration of the leaders, peak return stroke current, and other parts of the lightning process. These statistics may prove useful for additional statistical lightning models.

Middleton and Spaulding also propose a number of models for lightning noise in [87–90]. Extending the models of Middleton and Spaulding by combining their statistical models with the NLDN data to produce new lightning noise models that are the superposition of smaller storm cells and are spatially varying would be a better mix of the physics and storm statistics than the currently available models.

9.5.5 Inertial Aiding

Improvements to the time varying nature of lightning may assist another avenue of improving Loran coverage: inertial aiding. Since lightning does occur in bursts, inertial instruments may help the receiver coast through periods of high noise. Inertial instruments, which measure the aircraft's accelerations and rotations, can be used to aid in the position solution. Improvements of time-domain lightning models, coupled with the statistics from the APDs of lightning, could lead to a better understanding of the requirements for blending inertial instruments with Loran.

9.6 Final Thoughts

Loran has the potential to be a useful and cost-effective backup navigation system for aircraft. While, improvements are still necessary to reach the requirements set by the FAA, much progress has been made over the past few years. This dissertation represents some of that progress which shows Loran coverage can be improved through increased understanding of the physics and the requirements. The final solution for Loran may demand a combination of improved antennas, more towers, and inertial aiding. However, none of these are beyond the capabilities of today's technology. With its ubiquitous signal, simple infrastructure, and frequency separation from GPS, the investment in improving Loran is an investment well spent to provide a robust backup navigation system to GPS.

Appendix A

Fourier Analysis

A.1 Fourier Series

Fourier stated that any arbitrary *periodic* function, $E(t)$, with period T , may be expressed by the sum of an infinite series of sines and cosines

$$E(t) = \frac{1}{2}A_o + \sum_{n=1}^{\infty} [A_n \cos(2\pi n f_0 t) + B_n \sin(2\pi n f_0 t)] \quad (\text{A.1})$$

where,

$f_0 = 1/T$ is called the fundamental frequency of the signal and is defined by the period of the signal,

A_n are the even Fourier coefficients for the even cosine basis functions, and

B_n are the odd Fourier coefficients of the odd sine basis functions.

The orthogonal property of sines and cosines can be used to obtain the Fourier coefficients of the function. If two bases are orthogonal then their inner product is zero. Sine and

cosine are orthogonal when the integral for the inner product is taken over the interval for which the function repeats itself.

$$\int_0^T \cos(2\pi n f) \sin(2\pi m f) = 0 \quad \forall n, m$$

$$\int_0^T \cos(2\pi n f) \cos(2\pi m f) = \begin{cases} 0 & n \neq m \\ T & n = m \end{cases}$$

$$\int_0^T \sin(2\pi n f) \sin(2\pi m f) = \begin{cases} 0 & n \neq m \\ T & n = m \end{cases}$$

This shows that sines and cosines are mutually orthogonal. Additionally, either function is orthogonal with itself for different frequencies.

A.2 Continuous Fourier Transform

For continuous signals, the Fourier Transform is defined as

$$\mathcal{E}(f) = \mathcal{F}(E(t)) \equiv \int_{-\infty}^{\infty} E(t) e^{-j2\pi f t} dt. \quad (\text{A.2})$$

The Fourier transform makes use of this orthogonality since the kernel for the transform can be expanded using Euler's rule,

$$e^{j2\pi f t} = \cos(2\pi f t) + j \sin(2\pi f t). \quad (\text{A.3})$$

If $E(t)$ is periodic, then Equation (A.2) can be expanded using Equations (A.1) and (A.3) resulting in

$$\begin{aligned}
\mathcal{E}(f) &= \int_{-\infty}^{\infty} E(t) e^{-j2\pi ft} dt \\
&= \int_{-\infty}^{\infty} \left[\frac{1}{2} A_o + \sum_{n=0}^{\infty} A_n \cos(2\pi n f_0 t) + \sum_{n=0}^{\infty} B_n \sin(2\pi n f_0 t) \right] [\cos(2\pi ft) + j \sin(2\pi ft)] dt.
\end{aligned}$$

Since sine and cosine are orthogonal, if $f = n f_0$ then,

$$\mathcal{E}(n f_0) = A_n + j B_n.$$

Performing the Fourier transform at a given frequency produces the Fourier coefficients for that frequency which are the amplitudes of sine and cosine required to reconstruct that function. The absolute magnitude of the coefficients can be plotted versus frequency thus producing the spectrum for the signal.

A.3 Discrete Fourier Transform

There are several forms of the Discrete Fourier Transform (DFT). Each form has its own peculiarities when understanding the physical interpretation of the math. This section will compare three different definitions, those of Bracewell, Matlab, and Preta, to see how the three relate to one another both mathematically as well as physically.

Reference [91] gives Bracewell's definition of DFT. By substituting m for ν and n for τ , his definition of the DFT, $E_B(m)$, of a discrete time function, $E(n)$, can be recast as

$$\mathcal{E}_B(m) = \frac{1}{N} \sum_{n=0}^{N-1} E(n) e^{-j2\pi(m/N)n} \quad (\text{A.4})$$

where,

m/N is the frequency component given in cycles/sample,

E is the discrete valued function of interest,

\mathcal{E}_B is the discrete Fourier coefficient of E averaged over the N samples in the units of E ,

$n \in \{0, \dots, N - 1\}$ is the sample number index, and

N is the number of samples taken.

Note that unlike the continuous case, the units of \mathcal{E}_B are the same as E , rather than that of E per Hz. Also, the term m is related to the frequency, f , of the continuous case, but is not exactly the same. For example, for the frequency component at the Nyquist rate, $m = N/2$. Alternatively, this can be interpreted as there being $N/2$ cycles in the total N sample time period.

The following terms are defined to ease the comparison to other DFT definitions. Let,

f_s be the sample frequency,

$T_{step} = \frac{1}{f_s}$ be the sample period or interval,

$T_{total} = NT_{step}$ be the total time interval over which the N samples are taken,

and

$f_{bin} = \frac{1}{T_{total}} = \frac{1}{NT_{step}} = \frac{f_s}{N}$ be a frequency bin-width of the Fourier coefficients.

Rather than using Bracewell's index notation where $\mathcal{E}(n)$ is the n^{th} time sample, $E(nT_{step})$ will be used to explicitly indicate the time, $t = nT_{step}$, at which the continuous function E is sampled. Likewise, m will be converted to an explicit frequency, $f = mf_{bin}$.

In using the above definitions, m/N can be rewritten as

$$\begin{aligned}\frac{m}{N} &= \frac{m}{f_s} \frac{f_s}{N} \\ &= mT_{step}f_{bin}.\end{aligned}$$

Through substitution, Equation (A.4) can be rewritten as

$$\begin{aligned}\mathcal{E}_B(mf_{bin}) &= \frac{1}{N} \sum_{n=0}^{N-1} E(nT_{step}) e^{-j2\pi(mT_{step}f_{bin})n} \\ &= \frac{1}{N} \sum_{n=0}^{N-1} E(nT_{step}) e^{-j2\pi mf_{bin}nT_{step}}.\end{aligned}\tag{A.5}$$

Bracewell's DFT can now be readily compared to alternative definitions. Matlab defines the DFT similarly to Bracewell's, though the indices get a bit confusing. Matlab's DFT definition is

$$\mathcal{E}_M(k) = \sum_{p=1}^N E(p) e^{-j2\pi(p-1)\frac{(k-1)}{N}}\tag{A.6}$$

Where the indices $k, p \in \{1, \dots, N\}$ since Matlab traditionally has not supported zero or negative indices. Letting

$$m = k - 1$$

and

$$n = p - 1$$

and if

$$N = \frac{1}{f_{bin} T_{step}}$$

then Equation (A.6) can be recast in explicit time and frequency to become

$$\begin{aligned} \mathcal{E}_M(m f_{bin}) &= \sum_{n=0}^{N-1} E(n T_{step}) e^{-j 2 \pi m \frac{n}{N}} \\ &= \sum_{n=0}^{N-1} E(n T_{step}) e^{-j 2 \pi m f_{bin} n T_{step}} \end{aligned}$$

thereby relating Bracewell's and Matlab's definition of the DFT by

$$\mathcal{E}_B(m f_{bin}) = \frac{1}{N} \mathcal{E}_M(m f_{bin}).$$

The factor of N^{-1} appears here, since Matlab places it on the inverse definition of the DFT rather than the forward definition as Bracewell has done. Results from the FFT in Matlab, therefore, must be scaled by N^{-1} to get the magnitudes correct.

These machinations set the stage for the third alternative definition for the DFT, that of Preta, which was encountered in this dissertation. In [45], Preta defines the DFT as

$$\mathcal{E}_P(mf_{bin}) = T_{step} \sum_{n=0}^{N-1} E(nT_{step}) e^{-j2\pi mf_{bin} nT_{step}} \quad (\text{A.7})$$

$$\begin{aligned} &= T_{step} \mathcal{E}_M(mf_{bin}) \\ &= NT_{step} \mathcal{E}_B(mf_{bin}) \\ &= T_{total} \mathcal{E}_B(mf_{bin}) \end{aligned} \quad (\text{A.8})$$

$$= \frac{1}{f_{bin}} \mathcal{E}_B(mf_{bin}) \quad (\text{A.9})$$

The multiplier of T_{total} in Equation (A.8) makes it impossible to convert between the two forms without having a value for the total time required to take the data. The determination of this time interval is the subject of Section 7.5. The alternate view of Equation (A.9) shows how confusion may arise in attempting to interpret Preta's results. The units of Preta's DFT are the units of the original function E divided by Hz. This would seem to imply a density function, but it is not.

Preta's definition likely comes from the continuous definition of the Fourier Transform, given by Bracewell as,

$$\mathcal{E}(f) = \int_{-\infty}^{\infty} E(t) e^{-j2\pi ft} dt \quad (\text{A.10})$$

Discretizing the equation, with the integral going to a summation, and

$$\begin{aligned} dt &\rightarrow T_{step} \\ t &\rightarrow nT_{step} \\ f &\rightarrow mf_{bin} \end{aligned}$$

would result in a form similar to Equation (A.7).

In Reference [91], Bracewell defines the energy (or power, depending on the units of E) spectrum of a function as $|\mathcal{E}_B(f)|^2$ which has units of energy density per Hz. Since the units of E are V/m, and in the discrete case, \mathcal{E}_B has the same units, then $|\mathcal{E}_B(f)|^2$ would have units of V^2/m^2 . Likewise, power is proportional to V^2 , so $|\mathcal{E}_B(f)|^2$ can be interpreted as a power density. Equation (A.8) is used to convert $\mathcal{E}_B(f)$ to $\mathcal{E}_P(f)$ by multiplying by the total time of the data record, T_{total} . Thus, if the physical interpretation of Bracewell's definition of the DFT relates to power, then Preta's definition relates to energy. This is the key difference between the two mathematical realizations of the DFT.

A.4 Physical Interpretation

While not always explicitly stated, periodicity of the signal in discrete Fourier techniques is assumed. Therefore, the validity of the periodicity of the signal complicates the physical interpretation of each mathematical definition of the DFT.

Bracewell's definition of the DFT utilizes the periodic property of the Discrete Fourier Transform by averaging the Fourier coefficients over the interval of the data record. His DFT "sees" the data as a single period of an infinitely repeating signal rather than just a finite event. Therefore, replicas of the signal are generated in both directions in time and the spacing of those replicas is dependent upon the duration of the data.




To illustrate this phenomenon, Figure A.1 (a) shows the waveform of a signal which is sampled for four seconds, is finite over the interval $[0, 2]$, and is zero at all other times. If the data existed only on the interval $[0, 2]$, then the DFT would view the waveform of Figure A.1 (a) as Figure A.1 (b), where the signal replicas will be placed two seconds apart.

Taking the magnitude of the DFT, gives the rms of the waveform. Under this interval, the signal has an rms voltage of 3.54 V.

However, if the data is said to exist on the interval $[0,4]$, then the sampled data properties of the DFT result in the waveform shown in Figure A.1 (c). Here, the signal replicas are placed 4 seconds apart and result in an rms voltage of 2.5 V. Thus the duration of the data record used for the DFT determines the spacing of the replication of the signal and thereby influences the root-mean square voltage or power of the perceived signal.

In using Bracewell's definition of the DFT to calculate power or rms voltage, the result is an average power since the Fourier coefficients are obtained from averaging the data over the duration of the data record. This averaging of power only makes sense, however, when dealing with periodic signals. Lightning is an aperiodic phenomena, and therefore care is required when interpreting power since the average power will drop as longer data records are used. This is demonstrated in the example of Section 5.4.2.

Because his formulation of the Fourier coefficients is independent of the duration of the data record of the return stroke, Preta's definition aids in capturing and describing impulsive signals. With his formulation, if the magnitude of the DFT is used to calculate a spectrum, then this spectrum is more energy-like. This is a result of the DFT coefficients being calculated from the total projection of the function across the sinusoidal basis functions rather than just taking an average value. Preta's definition makes comparisons of aperiodic signal spectra between researchers easier since it is independent of the data's time interval. However, it is not equivalent to Bracewell's definition and needs further manipulation to be turned into a power spectrum, as the next section shows.

(a) Single pulse, defined on $[0, 2]$.(b) DFT view of single pulse defined on $[0, 2]$.(c) DFT view of single pulse defined on $[0, 4]$.**Figure A.1:** Visualization of the effect of interval on DFT and rms value.

A.5 Power Spectral Density

The Power Spectral Density (PSD) of a signal from data processed by Bracewell's DFT is a straight forward process. Starting with the double sided power spectrum of the signal, P_{DS} , defined as

$$P_{DS} = |\mathcal{E}_B(f)|^2$$

to produce the PSD, P_{DS} must be divided by the frequency bin width of the data and by the noise power bandwidth ($NPBW$) of any windowing used. If no windowing is used, the $NPBW = 1$. Therefore, the PSD is defined as

$$\begin{aligned} PSD &= \frac{P_{DS}}{f_{bin} NPBW} \\ &= \frac{|\mathcal{E}_B(f)|^2}{f_{bin} NPBW} \end{aligned} \quad (A.11)$$

To convert Preta's data into a PSD, Equation (A.8) must be substituted into Equation (A.11) and the reciprocal nature of T_{total} and f_{bin} is used to find

$$\begin{aligned} PSD &= \frac{|\mathcal{E}_P(f)|^2}{T_{total}^2 f_{bin} NPBW} \\ &= \frac{|\mathcal{E}_P(f)|^2}{T_{total} NPBW} \end{aligned}$$

Taking the logarithm of the above function gives,

$$\begin{aligned}
PSD_{\text{dB}} &= 10 \log_{10} \left(\frac{|\mathcal{E}_P(f)|^2}{T_{\text{total}} NPBW} \right) \\
&= 20 \log_{10} (|\mathcal{E}_P(f)|) - 10 \log_{10} (T_{\text{total}}) - 10 \log_{10} (NPBW)
\end{aligned}$$

The first term is the quantity Preta used for his plots which must then be adjusted by the next two terms to convert his data into a PSD. If windowing was not performed in calculating the DFT, then the last term is ignored. From [45], $T_{\text{total}} = 180 \mu\text{s}$, therefore the conversion from Preta's double-sided data to a double-sided PSD is simply done by adding A , defined as,

$$\begin{aligned}
A &= -10 \log_{10} (T_{\text{total}}) \\
&= -10 \log_{10} (180e^{-6} \text{ s}) \\
&= -10 \log_{10} (180e^{-6}) \text{ dB Hz} \\
&= 37.4 \text{ dB Hz}
\end{aligned}$$

to Preta's values in order to convert them to a PSD. This is shown in Section 7.5, where the return stroke data of Preta is converted into a PSD in order to make a comparison between the return stroke average power and that predicted by the ITU/CCIR model.

Appendix B

Antenna Noise Calculation

As mentioned in Section 3.2.4, the ITU model provides measurements of F_a , the median noise factor. To be useful in this study, F_a must be converted into an electric field strength. The conversion requires assumptions about the antenna. Since Loran is centered at 100 kHz, its wavelength is 3,000 m. With such a long wavelength, a meter long antenna will satisfy the short-dipole approximation. The conversion begins with the antenna's maximum effective aperture, A_M , for an ideal dipole

$$A_D = \frac{3}{8\pi} \lambda^2 \text{ m}^2.$$

Since the Loran antennas used in this dissertation are monopoles, each antenna will be modeled as a monopole above an ideal ground plane. The effective area of a monopole is twice the effective area of a dipole due to the mirroring effect of the ground plane.

$$A_M = \frac{3}{4\pi} \lambda^2 \text{ m}^2. \tag{B.1}$$

The power density, or Poynting Vector, S , in W/m^2 , for an incoming electric field, E , is

$$S = \frac{1}{2} \frac{|E|^2}{\eta} \quad (\text{B.2})$$

where,

E is the rms electric field strength aligned with the dipole in V/m , and

η is the ratio of the electric and magnetic field components (for free-space this is $120\pi \Omega$).

The power received at an antenna is a function of the electric density, S , and the effective area of the antenna, A_e , which is the area of the monopole, A_m .

$$\begin{aligned} P &= S \cdot A_e \\ &= S \cdot A_m. \end{aligned} \quad (\text{B.3})$$

Given that the external noise factor, f_a is defined by

$$f_a = \frac{p_n}{kT_o b}$$

where,

p_n is the available noise power from an equivalent lossless antenna W ,

k is the Boltzmann's constant $= 1.38 \times 10^{-23} \text{ J/K}$,

T_o is the reference temperature (K) taken as 290 K, and

b is the noise power bandwidth of the receiving system (Hz).

Then by taking $10 \log_{10}$ on both sides and representing this shift by capital letters, would result in,

$$\begin{aligned} F_a &= P_n - B - 10 \log_{10}(kT_o) \\ &= P_n - B - 10 \log_{10}(4.002 \times 10^{-21}) \\ &= P_n - B + 204 \text{ dB} \end{aligned}$$

or,

$$P_n = F_a + B - 204 \text{ dB}.$$

For a short vertical monopole ($h \ll \lambda$) above a perfect ground plane, by Equations (B.1), (B.2) and (B.3),

$$\begin{aligned}
P &= S \cdot A_m \\
&= \frac{1}{2} \frac{E^2}{\eta} \cdot \frac{3}{8\pi} \lambda^2 \\
&= \frac{3}{16} \frac{E^2}{\eta \pi} \frac{c^2}{f^2} \\
&= \frac{3E^2}{1920\pi^2} \frac{\bar{c}^2}{f_{MHz}^2} \\
&= \frac{3E^2}{1920\pi^2} \frac{300^2}{f_{MHz}^2} \\
&= 14.248 \frac{E^2}{f_{MHz}^2}.
\end{aligned}$$

Taking $10 \log_{10}$ on both sides

$$P = 20 \log_{10} E - 20 \log_{10} f_{MHz} + 11.538 \text{ dB W}$$

$$P = E - 20 \log_{10} f_{MHz} + 11.538 \text{ dB W}.$$

Combining the equation with the definition of F_a produces

$$P = F_a + B - 204 \text{ dB W} = E - 20 \log_{10} f_{MHz} + 11.538 \text{ dB W}$$

$$E = F_a + 20 \log_{10} f_{MHz} + B - 215.5 \text{ dB V/m}$$

but, E is measured in V/m. To convert to \bar{E} in $\mu\text{V/m}$,

$$E = \bar{E} \times 10^{-6}$$

so

$$\begin{aligned} 20 \log_{10} E &= 20 \log_{10} \bar{E} + 20 \log_{10} 10^{-6} \\ E &= \bar{E} - 120. \end{aligned}$$

Therefore,

$$\begin{aligned} \bar{E}_n &= E + 120 \\ &= F_a + 20 \log_{10} f_{MHz} + B - 95.5 \text{ dB}(\mu\text{V/m}). \end{aligned}$$

Thus, \bar{E}_n represents the absolute vertical component of the rms electric field strength for a short vertical monopole over an infinite ground plane.

Appendix C

Loran Receiver Calibration

Two independent methods were used to calibrate the Loran receiver system. The first method employs a simple propagation model to estimate the electric field strength of Loran signals. It then develops a scale factor based on the received signals and those predicted by the model. The second method uses models of lightning return strokes and compares the responses of the electric field received simultaneously on the flat-plate and the Loran antennas to develop the scale factor. This calibration enables the receiver to measure the absolute field strength of the noise. Absolute measurements allow the comparison of data collected by the SLR to that of the ITU database.

C.1 Calibration using Loran Signal Models

This method estimates the peak electric field value from the Loran tower at Boise City, OK as it would be received in Norman, OK. The location of these two points are listed in Table C.1. The range between these two points is 276.6 NM or 512.2 km. An ADC value of 0.0008 V was typical for the Boise City signal. Given the Boise City tower uses

Site	Latitude [deg min sec]	Latitude [deg]	Longitude [deg min sec]	Longitude [deg]
Boise City, OK	35° 12' 38.1" N	36.5058° N	102° 53' 59.487" W	102.8999° W
Norman, OK	36° 30' 20.783" N	35.2106° N	97° 26' 25.4" W	97.4404° W

Table C.1: Site Location

a 900 kW transmitter, and by using propagation models which contain the 1985 Conductivity maps, the Loran coverage software estimates the peak Boise City signal strength to be 75.1 dB($\mu\text{V}/\text{m}$). An alternate version of the software propagation model estimates the signal to be 76.5 dB($\mu\text{V}/\text{m}$). Given that these are rough models of conductivity, discrepancies of ± 2 dB are not unexpected.

Both of these software packages use Millington's curves, which are shown in Figure C.1, to estimate the propagation of the signal over a 1/4 degree grid. Given that the soil is cultivated throughout Oklahoma, and the distance is 277 NM, the electric field value can be determined directly as 66 dB($\mu\text{V}/\text{m}$) for a 400 kW transmitter at the standard sampling point (SSP). Adding 3.52 dB to account for Boise City being a 900 kW transmitter and adding an additional 5.91 dB to adjust the field strength from the SSP to the peak value produces a third estimate of 75.4 dB($\mu\text{V}/\text{m}$) for the Boise City signal strength. Since Millington's method is in between the other two, 75.4 dB($\mu\text{V}/\text{m}$) which is 5,888 $\mu\text{V}/\text{m}$ will be used as the expected signal strength from Boise City, OK.

Therefore, the calibrated scale factor, SF , of the Stanford Loran Receiver is

$$\begin{aligned}
 SF &= \frac{5,888 \mu\text{V}/\text{m}}{0.0008 \text{ V}} \\
 &= 7,360,000 \mu\text{V}/\text{m}/\text{V} \\
 &= 7.36 \text{ V}/\text{m}/\text{V}.
 \end{aligned}
 \tag{C.1}$$

With SF , the ADC readings can be converted into absolute electric field measurements. This will facilitate comparison between the data collected by the SLR and that of the ITU database.

C.2 Calibration Based on Lightning Data

The second method for calibrating the receiver's scale factor relies on the consistency of the lightning return strokes mechanism. Wide-band electric field strength waveforms of lightning return strokes follow a characteristic time history and atmospheric scientists have documented a modified spectrum of lightning return strokes in [41, 43–47]. Using the wide-band flat-plate antenna as described in Section 6.3, return strokes were captured such as those shown in Figure C.2.

In [45], Preta gives values for what is termed the Electric Field Frequency Spectrum. The values for a lightning return stroke frequency spectrum and a straight line approximation are shown in Figure C.3. Section 7.5 describes the significance of Preta's frequency spectrum and the calculation of the power of a lightning return stroke.

To be able to correlate return strokes captured by the wide-band flat-plate antenna to that of Preta, the distance to the return stroke from the receiver as well as the return stroke's magnitude needs to be known. This information is obtained from the National Lightning Detection Network (NLDN) database. The NLDN is comprised of numerous sensors scattered across the country which use lightning models and time-of-arrival measurements to estimate lightning locations and stroke current magnitudes. Table C.2 lists the values for several captured return strokes that could be correlated to NLDN data.

In order to compare the collected return stroke data with that of Preta, several steps must be followed. First, the Fourier Transformation given by Equation (A.7) must be performed on the return stroke signal captured by the flat-plate antenna. From the spectrum values, the maximum value between 99-101 kHz is taken as the 100 kHz value. Seven of these spectrum values as well as the Preta curve for 100 kHz are shown in Figure C.3.




Figure 2 LORAN-C Field Strength vs. Distance

Figure C.1: Millington's Curves for Loran-C Field Intensity




Figure C.2: Time history of electric field (E-field) strength of a nearby lightning return stroke.




Figure C.3: Preta predicted electric field spectrum $\text{dB}(\mu\text{V}/\text{m}/\text{Hz})$ versus distance with unadjusted return stroke data.

File Name	Time [sec into file]	Distance [km]	Current [kA]	Raw Spectrum [dB V/m/Hz]
Norman20050613T020017.mat	4.032	4.3	-22.0	-65.5
Norman20050613T015122.mat	4.249	7.0	-30.0	-66.2
Norman20050613T015902.mat	4.190	3.5	-10.0	-66
Norman20050613T015132.mat	4.302	8.4	-34.0	-65.4
Norman20050613T015427.mat	4.206	7.1	-22.0	-63.7
Norman20050613T015242.mat	4.020	14.7	-42.0	-71.9
Norman20050613T015532.mat	4.115	21.2	-14.0	-81.1

Table C.2: Return Stroke Data




Figure C.4: Preta predicted electric field spectrum $\text{dB}(\mu\text{V}/\text{m}/\text{Hz})$ versus distance with return stroke data normalized to 35 kA current.

Since both the return stroke current and the distance to the return stroke can vary for each stroke, researchers normalize their data in order to facilitate comparisons. They normalize the electric field proportional to a return stroke current of 35 kA and inversely proportional to a distance of 50 km. The collected data has been normalized in this fashion, resulting in the data varying more linearly as shown in Figure C.4. At this point, the points seem to follow the trend predicted by Preta but offset by a constant gain.

Assuming the difference stems from a calibration error or possibly the gain induced by the building since the antenna is above the ground plane, a constant gain offset of -13 dB is introduced to bring the sample more in line with the values obtained by Preta as shown in Figure C.5.

To show that this 13 dB adjustment is not just arbitrary, but reflects a real gain offset, an additional comparison with the scale factor calculated in the previous section is made.




Figure C.5: Preta predicted electric field spectrum $\text{dB}(\mu\text{V}/\text{m}/\text{Hz})$ versus distance with return stroke data normalized to 35 kA current with 13 dB offset removed.

With the newly calibrated flat-plate data, the return stroke data are filtered through a 35 kHz bandpass filter centered at 100 kHz. The resulting filtered waveform is the component of the stroke that exists within the Loran band. Overlaying the resulting waveform to the actual envelope data obtained directly from the receiver attached to the ICS-652 board highlights the time shifting of the data due to the filter delay. If the delay is compensated for, then the filtered return stroke data aligns well with that of the actual envelope data obtained from the receiver as depicted in Figure C.6. Since two different methods were used to calibrate the receiver and the flat-plate antenna to absolute electric field values, the closeness of this comparison shows the accuracy of the calibration technique. Therefore, the scale factor given in Equation (C.1) has been validated and now the collected data which is in terms of absolute field strength, may be compared directly to the ITU database.




Figure C.6: Measured envelope data (blue) and simulated response with filtered flat plate data (red).

Appendix D

SNR Definition

From the Loran Signal Specification [55], the RF band SNR is defined as the ratio

$$SNR = \frac{SSP_{RMS}}{Noise_{RMS}}$$

Since these are voltages, to convert to dB requires taking the log of SNR,

$$SNR_{dB} = 20 \log_{10} (SNR)$$

If the power in one pulse is proportional to the rms voltage as

$$Power_{Pulse} = SSP_{RMS}^2$$

then N pulses will have N times the power

$$Power_{N \text{ Pulses}} = N \cdot SSP_{RMS}^2$$

To express this in dB, take $10 \log_{10}$ of the above expression

$$\begin{aligned}
Power_{N \text{ Pulses}} \text{ dB} &= 10 \log_{10} (N \cdot SSP_{RMS}^2) \\
&= 10 \log_{10} (N) + 10 \log_{10} (SSP_{RMS}^2) \\
&= 10 \log_{10} (N) + 20 \log_{10} (SSP_{RMS}) .
\end{aligned}$$

Only $10 \log_{10}$ is used for N while $20 \log_{10}$ is used for the voltage. This first term is the processing gain due to averaging or $Gain_{Avg}$.

For N pulses, the SNR cannot be expressed as just a simple fraction in a straight forward manner. It's easier to express this in dB

$$\begin{aligned}
SNR_{N \text{ Pulses}} \text{ dB} &= Power_{N \text{ Pulses}} \text{ dB} - Power_{Noise} \text{ dB} \\
&= 10 \log_{10} (N) + 20 \log_{10} (SSP_{RMS}) - 20 \log_{10} (Noise_{RMS}) \\
&= 10 \log_{10} (N) + 20 \log_{10} \left(\frac{SSP_{RMS}}{Noise_{RMS}} \right) \\
&= 10 \log_{10} (N) + 20 \log_{10} (SNR) \\
&= Gain_{Avg} + SNR_{\text{dB}} .
\end{aligned}$$

As an example, by using a Loran signal strength propagation model and the ITU model, the SNR of a tower is found to be 0 dB. Using two GRIs to form one PCI worth of data, then averaging for 5 seconds with the GRI of the tower 9610 (or 96,100 μs in length), results in a total of p PCIs, given as,

$$\begin{aligned}
 p &= \text{floor}(5 \text{ s} / 96, 100e^{-6} / 2) \\
 &= 26 \text{ PCI.}
 \end{aligned}$$

Since a Master station has nine pulses per GRI, and 18 pulses per PCI, then a Master station will have, N_{Master} pulses to average,

$$\begin{aligned}
 N_{Master} &= 18 p \\
 &= 468 \text{ pulses.}
 \end{aligned}$$

Therefore, the processing gain due to averaging is

$$\begin{aligned}
 Gain_{Avg} &= 10 \log_{10} (N_{Master}) \\
 &= 26.7 \text{ dB.}
 \end{aligned}$$

Making the post averaging SNR,

$$\begin{aligned}
 SNR_{N \text{ Pulses dB}} &= Gain_{Avg} + SNR_{\text{dB}} \\
 &= 26.7 \text{ dB} + (0 \text{ dB}) \\
 &= 26.7 \text{ dB.}
 \end{aligned}$$

Thus, averaging a tower whose SNR is -10 dB for 5 seconds, results in an SNR post averaging of 16.7 dB. With such an SNR, from Figure 5.7 the receiver should select the standard zero crossing with a 9×10^{-7} probability of error.

Appendix E

Acronyms and Symbols

2drms Twice the Root-Mean-Square of Distance.

ADC Analog-to-Digital Converter

AGC Automatic Gain Control

AGL Above Ground Level

AIAA American Institute of Aeronautics and Astronautics

AIM Aeronautical Information Manual

AOA Angle of Arrival

APD Amplitude Probability Distribution

ARN-2 Atmospheric Radio Noise Model-2 Receiver

AWGN Additive White Gaussian Noise

bps Bits per Second

BPF Band-pass Filter

C2A Cloud-to-Air

C2C Cloud-to-Cloud

C2G Cloud-to-Ground

CCIR Comité Consultatif International Des Radiocommunications (International Radio Consultative Committee)

CDF Cumulative Distribution Function

CDMA Code Division Multiple Access

CONUS Continental United States. All of the US except Alaska and Hawaii.

COTS Commercial off the Shelf

CRB Cramer-Rao Lower Bound

CW Continuous Wave

dBW Decibels Relative to One Watt

dB Decibels (logarithmic measurement of power or gain ratios)

DFT Discrete Fourier Transform

deg Degree

DH Decision Height

DLL Delay Lock Loop

DME Distance Measuring Equipment

DoD Department of Defense

DSP Digital Signal Processing

ECD Envelope-to-Cycle Difference

E-field Electric Field

EMI Electromagnetic Interference

FAA Federal Aviation Administration

FCC Federal Communications Commission

FFT Fast Fourier Transform

G2C Ground-to-Cloud

GDOP Geometric Dilution of Precision

GHz Gigahertz (billions of cycles per second)

GIDL Generalized Interference Detection and Localization

GNSS Global Navigation Satellite System

GPS Global Positioning System

GRI Group Repetition Interval

H-field Magnetic field

HMI Hazardously Misleading Information

I In-phase

IC Intracloud

ICAO International Civil Aviation Organization

IDFT Inverse Discrete Fourier Transform

IF Intermediate Frequency

IFR Instrument Flight Rules

ILS Instrument Landing System

ION Institute of Navigation

INS Inertial Navigation System

ITU International Telegraph (or Telecommunication) Union

JHU/APL Johns Hopkins University/Applied Physics Laboratory

kHz Kilohertz (thousands of cycles per second)

kS/s Kilo-Samples/second

L-band All frequencies between 1 and 2 GHz

L1 first L-band GPS frequency: 1575.42 MHz (154×10.23 MHz)

LNA Low Noise Amplifier

Loran LOnge-RAnge Navigation

LOS Line-of-Sight

LO Local Oscillator

LPF Low Pass Filter

m Meters

Mbps Megabits per second

MHz Megahertz (millions of cycles per second)

ML Maximum Likelihood

MSL Mean Sea Level

MS/s Mega-Samples/seconds

mW Milliwatt

NAS National Airspace System

NDB Non-Directional Beacon

NM Nautical Miles

NTIA National Telecommunications and Information Administration

OCXO Oven-Controlled Crystal Oscillator

OU Oklahoma University

PDF Probability Density Function

ppb Part-per-Billion

PSD Power Spectral Density

Q Quadrature

RFI Radio Frequency Interference

RF Radio Frequency

rms Root-Mean-Square

RNAV Area Navigation Systems

RNP 0.3 Required Navigation Performance of 0.3 Nautical Miles

rss Root-Sum-Square

RVR Runway Visual Range

SLR Stanford Loran Receiver

SNR Signal-to-Noise Ratio

SSP Standard Sampling Point

SZC Standard Zero Crossing

TCXO Temperature Controlled Crystal Oscillator

TD Time Difference

TDOA Time Difference of Arrival

TOA Time of Arrival

UHF Ultra High Frequency (all frequencies between 300 MHz and 3 GHz)

US United States

UTC Universal Coordinated Time

V Volts

V_d Voltage Deviation

VOR Very-High-Frequency Omnidirectional Range

VPL Vertical Protection Limit

WAAS Wide Area Augmentation System

WWII World War II

Bibliography

- [1] “2005 Federal Radionavigation Plan,” U.S. Department of Transportation, Washington, DC, Tech. Rep. DOT-VNTSC-RITA-05-12, 2005.
- [2] “Garmin 1996 Annual Report,” Garmin, Caymen Islands, Tech. Rep., 1996.
- [3] “Garmin 2005 Annual Report,” Garmin, Caymen Islands, Tech. Rep., 2005.
- [4] W. J. Clinton, “Presidential decision directive,” May 1998.
- [5] T. M. et. al., “Gps risk assessment study-final report,” Johns Hopkins University-Applied Physics Laboratory, Tech. Rep., January 1999.
- [6] “2001 Federal Radionavigation Plan,” U.S. Department of Transportation, Washington, DC, Tech. Rep. DOT-VNTSC-RITA-01-3, 2001.
- [7] “Vulnerability Assessment of the Transportation Infrastructure Relying on Global Positioning System,” John A. Volpe National Transportation System Center, Boston, Tech. Rep., 2001.
- [8] D. Gebre-Egziabher, “Design and Performance Analysis of a Low-Cost Aided Dead Reckoning Navigator,” Ph.D. dissertation, Stanford University, Stanford, CA, February 2004.

- [9] P. Leman, "Spreadsheet of vor/dme/dmer and tacr operation costs," March 2006.
- [10] "Loran's Capability to Mitigate the Impact of a GPS Outage on GPS Position, Navigation, and Time Applications," Federal Aviation Administration Vice President for Technical Operation Navigation Services Directorate, Boston, Tech. Rep., March 2004.
- [11] "An Assessment of the Proposed Phase out of the Loran-C Navigation System," Booz-Allen Hamilton, Tech. Rep., July 1998.
- [12] "<http://www.navfltsm.addr.com>."
- [13] "2006 Aeronautical Information Manual: Official Guide to Basic Flight Information and ATC Procedures," Federal Aviation Administration, Washington, DC, Tech. Rep., 2006.
- [14] "Manual on Required Navigation Performance (RNP)," International Civil Aviation Organization, Montreal, Canada, Tech. Rep. Doc 9613-AN/937, 1999.
- [15] "Minimum operational performance standards for airborne supplemental navigation equipment using global positioning system (gps)," RTCA, Tech. Rep., July 1991.
- [16] "Loran's capability to mitigate the impact of a gps outage on gps position, navigation, and time applications," Federal Aviation Administration, Tech. Rep., March 2004.
- [17] "<http://www.nstb.tc.faa.gov/terms.html>," October 2006.
- [18] L. F. Fehlner and T. A. McCarty, "How to harvest the full potential of loran-c," in *Proceedings of the Second Annual Convention of the WGA*. Wild Goose Association, October 1973.

- [19] R. H. Erikson, "A review of loran-c at the faa technical center," Wild Goose Association. Rockville, MD: Proceedings of the Sixteenth Annual Technical Symposium, October 1987.
- [20] T. O. P. Team, "Loran-c nonprecision approaches - the ohio program," in *Proceedings of the Sixteenth Annual Technical Symposium*. Wild Goose Association, October 1987.
- [21] R. W. Lilley, "Redesignation of loran-c as a permanent, sole-means navigation system," in *Proceedings of the Sixteenth Annual Technical Symposium*. Rockville, MD: Wild Goose Association, October 1987.
- [22] "Loran for aviation update," in *Proceedings of the Seventeenth Annual Technical Symposium*. Wild Goose Association, October 1988.
- [23] F. van Graas, "Sole Means Navigation and Integrity through Hybridized Loran-C and NAVSTAR GPS," in *Proceedings of the Sixteenth Annual Technical Symposium*. Rockville, MD: Wild Goose Association, October 1987.
- [24] M. Morgenthaler and G. Steiner, "Experimental Results using Differential Loran/GPS for Non-Precision Approaches," in *Proceedings of the Eighteenth Annual Technical Symposium*. Hyannis, Massachusetts: Wild Goose Association, October 1989.
- [25] "Study results of gps/loran interoperability for air navigation (executive summary)," Navcom Systems, Incorporated, Manassas, VA, Tech. Rep. Contract Number DTRS-57-88-C-00088, April 1989.

- [26] D.H. Amos, J.D. Catlin, W.B. Mohin, B.W. Heldt, R.B. Goddard, F. Vicksell, P. Enge, and R.J. Gauthier, "Coordination and Synchronization of GPS and Loran-C Transmitters to Universal Time," U.S. Department of Transportation, USCG, Syntetics, 540 Edgewater Drive, Wakefield, MA 01880, Tech. Rep. DTCG23-86-A-20022/CGA001, August 1989.
- [27] J. V. Carroll and J. A. Weitzen, "Loran-C Noise Model Development for RTCA SC-176, SC-159, and TSO-C60b," in *Proceedings of the Twenty-second Annual Convention and Technical Symposium*. Wild Goose Association, October 1993.
- [28] J. V. Carroll, "Notes on the Computation and Use of Atmospheric Signal-to-Noise Ratio in Loran Receivers," in *Proceedings of the Twentieth Annual Technical Symposium*. Williamsburg, VA: Wild Goose Association, October 1991.
- [29] "Recommendation International Telecommunications Union Radio Noise," ITU, Geneva, Tech. Rep. ITU Recommendation ITU-R P.372-7, 2001.
- [30] "Proposed Radio Aid to Navigation Position Fixing System for the Republic of South Africa: Natal Demonstration," Megapulse, Inc., Tech. Rep., March 1985.
- [31] "Loran-c user handbook, Tech. Rep. COMDTPUB P16562.6, 1992.
- [32] R. A. Reilly, "Loran processing advances," in *Wild Goose Association Proceedings of the First Technical Symposium*, Boston, MA, November 1972.
- [33] D. Last, "Atmospheric Noise, EM Propagation, Skywaves, Coverage, GDOP and Interference - A Tutorial Paper," in *Proceedings of the Twenty-first Annual Technical Symposium*. Birmingham, England: Wild Goose Association, August 1992.

- [34] D. A. Feldman, “An Atmospheric Noise Model with Applications to Low Frequency Navigation Systems,” Ph.D. dissertation, Massachusetts Institute of Technology, Boston, MA, June 1972.
- [35] J. C. J.A. Weitzen and B. Dao, “Use of Simulated Atmospheric Noise in the Calibration and Characterization of Loran-C Receivers,” in *Proceedings of the Twenty-first Annual Technical Symposium*. Birmingham, England: Wild Goose Association, August 1992.
- [36] —, “Real-Time Simulation of VLF Atmospheric Noise for Use in the Calibration and Characterization of LORAN-C Receivers for Aircraft Navigation,” *Simulation*, vol. 60, pp. 357–366, June 1993.
- [37] G. Lachapelle, “Fundamentals of loran-c,” Navtech Seminars, Inc., September 1991.
- [38] W. J. Pelgrum, “Noise from a receiver perspective,” in *Proceedings of the Thirty-fourth Annual Convention and Technical Symposium*, Santa Barbara, CA, October 2005.
- [39] L. N. Postema, “Performance optimization of linear loran receivers operating in a non-gaussian atmospheric noise environment,” in *Proceedings of the Third Annual Convention of the WGA*. Great Gorge, NJ: Wild Goose Association, October 1974.
- [40] Martin A. Uman, *All About Lightning*. Mineola, New York: Dover, 1987.
- [41] —, *Lightning*. Mineola, New York: Dover, 1984.
- [42] —, *The Lightning Discharge*. Mineola, New York: Dover, 2001.

- [43] G.I. Serhan, M. A. Uman, D. G. Childers, and Y. T. Lin, "The RF Spectra of First and Subsequent Lightning Return Strokes in the 1- to 200-km range," *Radio Science*, vol. 15, no. 6, pp. 1089–1094, November-December 1980.
- [44] C. D. Weidman and E. P. Krider, "Lightning Amplitude Spectra in the Interval from 100 kHz to 20 MHz," *Geophysical Research Letters*, vol. 8, no. 8, pp. 931–934, August 1981.
- [45] Joseph Preta Jr., M. A. Uman, and D. G. Childers, "Comment on "The RF Spectra of First and Subsequent Lightning Return Strokes in the 1- to 200-km range" by Serhan et al." *Radio Science*, vol. 20, no. 1, pp. 143–145, January-February 1985.
- [46] Upul Sonnadara, Vernon Cooray, and Mahendra Fernando, "The Lightning Radiation Field Spectra of Cloud Flashes in the Interval From 20 kHz to 20 MHz," *IEEE Transactions on Electromagnetic Compatibility*, vol. 48, no. 1, pp. 234–239, February 2006.
- [47] Joseph E. Nanevich, E.F. Vance, William Radasky, M.A. Uman, Gordon K. Soper, and Joan Ma Pierre, "The Lightning Radiation Field Spectra of Cloud Flashes in the Interval From 20 kHz to 20 MHz," *IEEE Transactions on Electromagnetic Compatibility*, vol. 30, no. 4, pp. 463–172, November 1988.
- [48] M. L. Curtis Cutright and F. van Graas, "Loran-C Band Data Collection Efforts at Ohio," in *Proceedings of the Thirty-second Annual Convention and Technical Symposium*. Boulder, CO: International Loran Association, November 2003.

- [49] J. S. Curtis Cutright and F. van Graas, "Effects of Atmospheric Noise on Loran-C," in *Proceedings of the Thirty-third Annual Convention and Technical Symposium*. Tokyo: International Loran Association, October.
- [50] ———, "Analysis of the Effects of Atmospheric Noise on Loran-C," in *Proceedings of the Thirty-fourth Annual Convention and Technical Symposium*. Santa Barbara, CA: International Loran Association, October 2005.
- [51] F. v. G. Curtis A. Cutright, Janet M. Blazyk and D. W. Diggle, "Loran-C Atmospheric-Noise Flight-Test Results," in *Proceedings of the Thirty-fifth Annual Convention and Technical Symposium*. Groton, Connecticut: International Loran Association, October 2006.
- [52] "http://www.jproc.ca/hyperbolic/loran_a.html."
- [53] L. Melton, *The Complete Loran-C Handbook*. Camden, Maine: International Marine Publishing Company, 1986.
- [54] Myron Kayton, Walter R. Fried, *Avionics Navigation Systems, 2nd Ed.* John Wiley and Sons, Inc., 1997.
- [55] "Specification of the Transmitted Loran C Signal," U.S. Department of Transportation, United States Coast Guard, Washington, DC, Tech. Rep. COMDTINST M16562.4A, May 1994.
- [56] Sherman Lo, Peter Morris, and Per Enge, "Early Skywave Detection Network: Preliminary Design and Analysis," in *Proceedings of the International Loran Association 34th Annual Meeting*, Santa Barbara, CA, October 2005.

- [57] P. Misra and P. K. Enge, *Global Positioning System: Signals, Measurements and Performance*. Ganga-Jumuna Press, 2001.
- [58] R. J. Wenzel and D. Slagle, "Loran-c signal stability study: Neus/seus," U.S. Department of Transportation, United States Coast Guard, Office of Research and Development, Washington, DC 20593, Tech. Rep. Report No. CG-D-28-83, August 1983.
- [59] Bradford W. Parkinson and James J. Spilker Jr., *Global Positioning System: Theory and Applications, volume 1 and 2*. Washington, D.C.: AIAA, 1996.
- [60] D. Marshall, "Digital averaging loran receiver," in *Proceedings of the Sixteenth Annual Technical Symposium*. Rockville, MD: Wild Goose Association, October 1987.
- [61] R. D. van Nee, "An Advanced Loran-C Receiver Structure," in *Proceedings of the Twenty-first Annual Technical Symposium*. Birmingham, England: Wild Goose Association, August 1992.
- [62] R. D. van Nee and H. J. Andersen, "Optimum Loran-C Signal Processing: First Experimental Results," in *Proceedings of the Twenty-second Annual Technical Symposium*. Wild Goose Association, October 1993.
- [63] M. N. Kiyohiko Tatebayashi and T. Yamada, "Stochastic investigation on the search process of a loran-c receiver," in *Proceedings of the Ninth Annual Technical Symposium*. Boston, Massachusetts: Wild Goose Association, October 1980.
- [64] A. Nieuwland, "An Improved Cycle Identification Algorithm," in *Proceedings of the Twenty-first Annual Technical Symposium*. Birmingham, England: Wild Goose Association, August 1992.

- [65] W. N. Dean and R. L. Frank, "Loran-C Receiver System Technology and Phase Coding," in *Proceedings of the Twenty-first Annual Technical Symposium*. Wild Goose Association, August 1992.
- [66] M. Beckmann, "Synchronous interference to loran-c and its influence on cycle identification," in *Proceedings of the 19th Annual Technical Symposium Wild Goose Association*. Wild Goose Association, October 1990.
- [67] Y. Bian and D. Last, "Relative sensitivity of loran-c phase tracking and cycle selection to cwi," *IEEE Electronic Letters*, vol. 29, no. 4, February 1993.
- [68] M. S. David Last and J. McCulloch, "Calculation of ecd variations over inhomogeneous terrain using millington's method," in *Proceedings of the Twenty-second Annual Convention and Technical Symposium*. Wild Goose Association, October 1993.
- [69] Benjamin Peterson, "LORAN C H-Field DDC Receiver," United States Coast Guard Academy Department of Engineering, New London, CT, Tech. Rep. Version 011221, 2001.
- [70] Benjamin Peterson, Thomas Thomson and Jonathon Riffe, "Measurement of LORAN-C Envelope to Cycle Difference in the Far Field," in *Proceedings of the 20th Annual WGA Technical Symposium*, Williamsburg, VA, October 1991.
- [71] Benjamin Peterson, et al., "Hazardously Misleading Information Analysis for Loran LNAV," in *2nd International Symposium on Integration of LORAN-C/Eurofix and EGNOS/Galileo*, June 2002.

- [72] Sherman Lo, et al., “Loran Availability and Continuity Analysis for Required Navigation Performance 0.3,” in *Proceedings of GNSS 2004 Ū The European Navigation Conference*, Rotterdam, The Netherlands, May 2004.
- [73] A. C. Service, “Airborne Area Navigation Equipment Using Loran C Inputs,” Department of Transportation Federal Aviation Administration, Washington, D.C., Tech. Rep. TSO-C60b, May 1988.
- [74] Vladimir A. Rakov and Martin A. Uman, *Lightning: Physics and Effects*. Cambridge University Press, 2003.
- [75] W. H. Beasley, “Correspondence,” October 2006.
- [76] “<http://www.itu.int/aboutitu/overview/history.html>.”
- [77] “International Radio Consultative Committee Characteristics and Applications of Atmospheric Radio Noise,” CCIR, Geneva, Tech. Rep. CCIR Recommendation 322-3, 1988.
- [78] R. J. A. et al., “Man-Made Noise in the 136 to 138-MHz VHF Meteorological Satellite Band,” National Telecommunications and Information Administration, Boulder, CO, Tech. Rep. NTIA Report 98-355, 1998.
- [79] Gene F. Franklin, J. David Powell, Michael Workman, *Digital Control of Dynamic Systems, 3rd Edition*. Addison Wesley, 1997.
- [80] A. D. Spaulding and D. Middleton, “Optimum Reception in an Impulsive Interference Environment,” National Telecommunications and Information Administration, U.S. Department of Commerce, Tech. Rep.

- [81] D. Middleton and A. D. Spaulding, "Optimum Reception in Non-Gaussian Electromagnetic Interference Environments: II Optimum and Suboptimum Threshold Signal Detection in Class A and B Noise," National Telecommunications and Information Administration, U.S. Department of Commerce, Tech. Rep. NTIA Report 83-120, May 1983.
- [82] Per K. Enge and D. V. Sawate, "Spectrum Multiple-Access Performance of Orthogonal Codes: Impulsive Noise," *IEEE Transactions on Communications*, vol. 36, no. 1, January 1988.
- [83] P. E. Sherman Lo, "Demonstrating Integrity for Loran Cycle Selection Using Weighted Sum Square Error (WSSE) Statistic," in *Proceedings of the Thirty-fifth Annual Convention and Technical Symposium*. Groton, Connecticut: International Loran Association, October 2006.
- [84] Alberto Leon-Garcia, *Probability and Random Processes for Electrical Engineers 2nd Ed.* Redding, Massachusetts: Addison-Wesley Publishing Company, May 1994.
- [85] "<http://mathworld.wolfram.com/ricedistribution.html>," October 2006.
- [86] H. Hall, "A New Model for Impulsive Phenomena: Application to Atmospheric-Noise Communications Channels," Stanford Electronics Laboratories (SEL), Stanford University, Stanford, CA, Tech. Rep. Report SEL-66-052, 1966.
- [87] D. Middleton, "Statistical-Physical Models of Man-Made Radio Noise Part I," Office of Telecommunications, U.S. Department of Commerce, Tech. Rep. OT Report 74-36, April 1974.

- [88] ———, “Statistical-Physical Models of Man-Made Radio Noise Part II: First Order Probability Models of the Envelope and Phase,” Office of Telecommunications, U.S. Department of Commerce, Tech. Rep. OT Report 76-86, April 1976.
- [89] ———, “Statistical-Physical Models of Man-Made Radio Noise Part III: First-Order Probability Models of the Instantaneous Amplitude of Class B Interference,” National Telecommunications and Information Administration, U.S. Department of Commerce, Tech. Rep. NTIA-CR-78-1, June 1978.
- [90] ———, “Statistical-Physical Models of Man-Made Radio Noise Part IV: Determination of the First-Order Parameters of Class A and B Interference,” National Telecommunications and Information Administration, U.S. Department of Commerce, Tech. Rep. NTIA-CR-78-2, September 1978.
- [91] Ronald N. Bracewell, *The Fourier Transform and Its Applications*. McGraw-Hill Science/Engineering/Math, June 1999.