

REMOTE SENSING AND STATISTICAL ANALYSIS OF THE
LIGHTNING-IONOSPHERE INTERACTION

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

Patrick R. Blaes

August 2015

© Copyright by Patrick R. Blaes 2015
All Rights Reserved

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

(Umran S. Inan) Principal Adviser

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

(Robert Marshall)

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.

(John Gill)

Approved for the Stanford University Committee on Graduate Studies

Abstract

Lightning is an extremely frequent phenomenon and an important part of the global electrical circuit with an average of 2000 active storms globally every day and an average flash rate of 40-50 flashes per second. Each lightning stroke radiates an intense electromagnetic pulse (EMP) in the VLF/LF frequency band (300 Hz—300 kHz). This EMP propagates efficiently in the earth-ionosphere waveguide and can be measured long distances away on the ground, in excess of several thousand kilometers. The upwardly directed part of the EMP interacts with and heats plasma in the Earth’s ionosphere. Some of this energy propagates through the ionosphere where it enters the magnetosphere as a whistler-mode wave which may resonantly interact with energetic electrons in the Earth’s radiation belts.

Intense lightning EMPs may produce brief optical flashes known as “elves” in the lower part of the ionosphere (~ 88 km altitude). Elves, first discovered in 1991, appear as rapidly expanding rings of light centered about the causative lightning stroke. They have durations of less than 1 ms and radial extents up to several hundred kilometers.

In this work, we present ground-based observations of elves made using an optical free-running photometer along with VLF/LF observations of the lightning EMP magnetic field. We use these experimental observations to investigate the properties of the lightning return stroke that control the production of optical elve emissions. Two summers of data containing observations of over six hundred elves along with the LF magnetic field of the associated lightning are analyzed. By training a classifier on features of the EMP ground wave, we find that we are able to accurately predict whether or not a stroke produced an elve. Additionally, we show that the geometric features of elves can be used to estimate the current return stroke speed

of lightning. This estimation constitutes the first experimental measurement of the return stroke speed that relates directly to the current propagation rather than its optical manifestation. Finally, we use data from the GLD360 lightning geolocation network to extrapolate the global elve production rate and to estimate the amount of ionospheric heating due to lightning.

Acknowledgements

I would like to begin by taking the opportunity to thank the numerous people whose help made this dissertation possible.

First and foremost, I am immensely grateful to my advisor, Professor Umran Inan, for welcoming me into the VLF group many years ago. Over the years, he has provided me the opportunity to work on several exciting projects and to travel to interesting, far-away places. His wisdom and enthusiasm for the research made it a pleasure to work with him.

I would also like to express my extreme gratitude to Dr. Robert Marshall for mentoring me and serving as my second advisor during the course of this research. He has been an endless source of knowledge and has taught me much of what it means to be a good scientist. Also, his PIPER instrument and lightning EMP model were essential tools used in the completion of this dissertation. It has been a pleasure and a privilege to work with him.

Additionally, I would like to thank Professor John Gill for graciously agreeing to both serve as a reader and on my defense committee, to Professor Sigrid Close for serving on my defense committee, and to Professor Antony Fraser-Smith for serving as my committee chair.

The Stanford VLF group has been a constant source of support and I feel privileged to have been a part of it. I would especially like to thank Dr. Maria Spasojevic for serving as a mentor during my first couple of years, Dr. Nikolai Lehtinen and Dr. Morris Cohen for sharing their expertise during our weekly NIMBUS meetings, and Jefferey Chang for his endless support during field campaigns and for teaching me much of what I know about electronics and instrument design. I would also like to

thank Shaolan Min and Helen Niu for all for all of their help and support over the years. The students in the VLF group have been supportive colleagues and fantastic friends. I would like to thank, in no particular order, Can Liang, Justin Li, Fadi Zoghzoghy, Vijay Harid, Rasoul Kabirzadeh, Drew Compston, Austin Sousa, Chris Young, Forrest Foust, George Jin, Dan Golden, and Kevin Graf.

This dissertation would also not have been possible without the support of colleagues at other universities and field-sites. I would like to thank William Winn, Ken Eack, and Harald Edens of New Mexico Tech for graciously allowing me to install a PIPER photometer at Langmuir Laboratory and for providing support on numerous occasions. I would also like to express gratitude to Thomas Barnes, Coyne Gibson, and David Doss of University of Texas, Austin for hosting me at McDonald Observatory and for providing remote debugging assistance on many occasions when instruments stopped working.

Last, but not least, I would like to thank my family — my parents Brent and Maureen, and my sisters Katie and Mairead.

Patrick Ryan Blaes

Stanford, CA

August 16, 2015

This work was supported by Defense Advanced Research Projects Agency grant agreement HR0011-10-1-0058.

Contents

Abstract	iv
Acknowledgements	vi
1 Introduction	1
1.1 Lightning	2
1.2 Ionosphere	5
1.3 Transient Luminous Events	7
1.3.1 Sprites and Halos	8
1.3.2 Elves	8
1.3.3 Jets	10
1.4 Review of Prior Work	11
1.5 Contributions of this Thesis	13
1.6 Thesis Organization	14
2 The Lightning Return Stroke	15
2.1 Background	15
2.2 Modeling the Return Stroke Current	15
2.2.1 Engineering Models	16
2.2.2 Gas Dynamic Models	19
2.3 Radiated Fields	21
2.4 Earth-Ionosphere Waveguide Propagation	23
2.5 Radio Remote Sensing of the EMP	26
2.5.1 VLF/LF Receiver Design	27

2.6	Summary	29
3	Elves	31
3.1	Elves and the EMP-Ionosphere Interaction	31
3.2	General Features of Elves	38
3.3	Photometric Imaging of Elves	40
3.3.1	Instrumentation	41
3.4	Summer 2013 Observational Campaign	43
3.4.1	Experimental Setup and Observation Techniques	45
3.4.2	Langmuir Laboratory	48
3.4.3	McDonald Observatory	49
3.4.4	Example Observations	50
3.4.5	Occurrence Maps	51
3.5	Summary	53
4	Properties of Elve Producing Lightning	54
4.1	Dataset	54
4.2	General Features of Sferics from Elve Producing Lightning	55
4.3	Elve Production Probability Modeling	61
4.3.1	Features of the Sferic Ground Wave	61
4.3.2	Logistic Regression Model	62
4.3.3	Results	64
4.4	Return Stroke Speed of CG Lightning Using Elve Hole Radii	67
4.4.1	Emission Profile Reconstruction	69
4.4.2	Dataset	71
4.4.3	Return Stroke Speed Estimation	73
4.4.4	Results and Discussion	75
4.5	Conclusions	77
5	Global Effects of the Lightning EMP	79
5.1	Global Elve Production Rates	79
5.1.1	Prior Satellite Observations	80

5.1.2	Production Rate Estimated from Ground Observations	82
5.2	Heating of the D-Region by Lightning	84
5.2.1	Numerical Simulations	84
5.2.2	Global Heating Estimate	87
5.3	Summary	88
6	Conclusions and Suggestions for Future Work	91
6.1	Summary	92
6.2	Suggestions for Future Work	93

List of Tables

2.1	Return stroke engineering model parameters	17
3.1	O ₂ and N ₂ Collisions	34
3.2	Optical Excitation Rates	35
3.3	Optical Quenching Rates	35

List of Figures

1.1	Lightning occurrence map	2
1.2	Lightning schematic	4
1.3	Return stroke schematic	6
1.4	Ionospheric density profile	7
1.5	Sprites	9
1.6	Elves	10
1.7	Example Jets	11
1.8	Previous elve production probability	12
2.1	Return Stroke Base Current	18
2.2	Engineering Models	20
2.3	TDFL current propagation	21
2.4	Analytical EMP	23
2.5	Earth-Ionosphere Waveguide	26
2.6	LF Receiver Block Diagram	28
2.7	LF Receiver Photo	29
2.8	Example VLF/LF Spectrogram	30
3.1	Cartoon showing the EMP-ionosphere interaction	32
3.2	D-region Densities	33
3.3	Optical Rates	36
3.4	N ₂ Emission Spectra	37
3.5	Theoretical elve profiles	39
3.6	Elve photon delay effect	40

3.7	PIPER Schematic	44
3.8	PIPER data explanation	45
3.9	Elve occurrence map	47
3.10	Langmuir Photo	49
3.11	McDonald Photo	50
3.12	Example PIPER data	52
4.1	Sferic Sky-Wave	57
4.2	Distance binned sferic groundwaves	58
4.3	Ground wave rise time vs. distance	59
4.4	Peak current binned ground waves	60
4.5	Groundwave Features	62
4.6	Ground wave feature scatter plot matrix	63
4.7	Logistic Regression Example	65
4.8	Elve Production Probability	67
4.9	Elve Peak Current vs. Distance	68
4.10	Normalized EMP radiation pattern	69
4.11	Example elve emission profile	72
4.12	Histogram of elve hole radii	73
4.13	Histogram of elve hole radii compared to RS speed	74
4.14	Distribution of return stroke parameters estimated from elve hole radii	78
5.1	Elve Production Rate from ISUAL	81
5.2	Elve Production Rate Extrapolated From PIPER Observations	83
5.3	Example heating pattern from EMP	85
5.4	D-region Heating vs. Return Stroke Parameters	86
5.5	D-region Heating vs. Peak Current	87
5.6	Map of D-region heading due to the lightning EMP	89

Chapter 1

Introduction

Lightning is one of the most frequently occurring natural phenomena on Earth and forms an integral component of the global electric circuit. At any given moment, there are on average 2000 active thunderstorms covering about 10 percent of the Earth's surface [*Rakov and Uman, 2003*, p. 10] with a global average flash rate of about 45 lightning discharges per second [*Christian et al., 2003*]. Lightning activity is not uniformly distributed; particular hotspots such as the Caribbean, Central Africa, and Southeast Asia account for much of the global activity, with much of it occurring during summer months. Figure 1.1 shows a map of lightning flash rates as inferred by the satellite borne Optical Transient Detector instrument. Lightning can be very destructive and significantly impacts human society. It is estimated that lightning kills 24,000 people globally and causes and results in \$300 million in insured property damage in the US each year [*Holle, 2008*].

In addition to the direct impact of strikes on the ground, lightning also radiates an intense electromagnetic pulse (EMP). This EMP propagates long distances and may interfere with electronics and radio communications [*Wait and Spies, 1964*]. A portion of the EMP energy propagates upwards where it interacts with plasma in a region of the upper-atmosphere called the ionosphere. This EMP-ionosphere interaction results in energy transfer to free electrons and, subsequently, electron-neutral collisions which may excite optical emissions known as elves [*Inan et al., 1997*]. Elves belong to the broad class known as transient luminous events (TLE)

and are a relatively recent discovery, having only been first scientifically observed in 1990 [*Fukunishi et al.*, 1996; *Inan et al.*, 1991].

Despite the frequency of lightning and the danger to humans posed by ground strikes, there is still a great deal unknown about the phenomenon. In particular, there is uncertainty in the properties of the current propagation known as the *return stroke*. Furthermore, there are outstanding questions regarding the conditions that are necessary for elve production and the associated effects of elves on the ionospheric plasma. In this thesis, we utilize radio remote sensing and photometric imaging techniques to observe elves and to answer questions pertaining to elve production, the lightning return stroke speed, and ionospheric heating associated with elves.

In this chapter, we begin by introducing the basic processes and concepts related to lightning, the ionosphere, and transient luminous events. We then review prior work in this field and present the contributions and organization of this thesis.

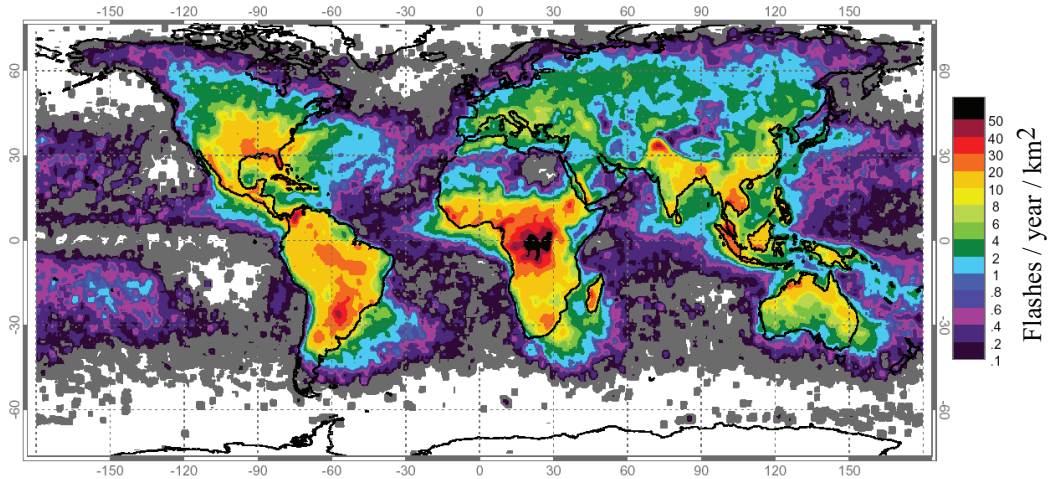


Figure 1.1: A global map of lightning occurrence rates estimated from the Optical Transient Detector instrument [*Christian et al.*, 2003].

1.1 Lightning

Lightning is predominantly produced by cumulonimbus clouds, which typically form at altitudes ranging between approximately 3 and 20 km [*Rakov and Uman*, 2003,

p. 67]. These clouds are formed by convective current systems and are characterized by strong updrafts. Before lightning can occur, it is first necessary that a separation of electric charge is established in the cloud. The most commonly accepted charge separation model is displayed in Figure 1.2, and is characterized by a main positive charge layer at the top of the cloud, a main negative layer in the middle, and a much smaller lower positive layer at the bottom of the cloud [Rakov and Uman, 2003, p. 69]. This tripole charge structure is well supported by in-situ balloon measurements [Byrne *et al.*, 1983] and by ground based electric field-mill measurements [Krehbiel, 1986]. A small negative screening layer has sometimes also been observed at the cloud top. It should be noted that in realistic thunderclouds, charge probably does not separate into ideal flat layers, but likely forms more complicated structures and charge pockets.

The exact mechanism for charge creation and separation is not entirely understood, but the most commonly held theory is that it is due to precipitation and collisional charge transfer between hydrometeor particles [Jayaratne, 1998]. Hydrometeors are the various liquid and frozen water molecules in the cloud, including rain, hail, and graupel. As heavy hydrometeor particles fall downward due to gravity, they collide with lighter particles carried upward by updraft currents. During these collisions, charge may be transferred between the particles, with the heavier precipitation particles tending to accumulate negative charge and the lighter particles gaining positive charge. Subsequently, the heavier, now negatively-charged, particles continue to fall downward to form the main negative charge layer, and the lighter positively-charged particles continue upward to form the main positive layer. At warmer temperatures (greater than -20° C), collisions tend to transfer positive charge to the heavy particles [Jayaratne *et al.*, 1983], which may be a contributing mechanism in the formation of the lower positive charge layer [Jayaratne and Saunders, 1984].

This charge separation in the cloud establishes strong local electric fields. When these fields exceed the atmospheric breakdown threshold, a highly-conductive ionized channel called a leader is formed. Current transferred along the leader serves to neutralize charge pockets in the cloud. There are two main classes of lightning discharges. The first type of discharge, and most common comprising about 75% of




Figure 1.2: Thunderclouds are typically modeled as a tripole charge structure, with a main positive layer on top, a main negative layer in the middle, and a smaller positive layer at the bottom. Cloud-to-ground (CG) discharges remove charge from one of these layers to the ground, while intracloud (IC) discharges neutralize charge between layers.

all lightning activity, is intra-cloud (IC) lightning. As the name implies, the leaders in IC lightning form between, and neutralize, charge pockets inside the cloud. The second type of discharge, cloud-to-ground (CG), connects a charge layer in the cloud to the ground. CG lightning is the more destructive form of lightning, carries much larger currents than ICs, and produces intense radio emissions which may be observed at large distances. Channels that connect the main negative layer and the ground are termed -CG and account for 90% of all CG activity, and channels that connect the main positive layer and ground are called +CG and account for the remaining 10%. CG lightning also has the greatest impact on the upper-atmosphere and ionosphere, and hence is the main type of discharge considered in this thesis.

The steps involved in -CG lightning are shown in Figure 1.3. After breakdown occurs inside the cloud, a leader is formed. This “stepped leader” grows downward toward the ground, ionizing the neutral atmosphere in a tortuous manner with many branches. Eventually, one of the leader branches makes contact with the ground (attachment), forming a highly conductive, ionized channel between the ground and

the negative charge layer in the cloud. Since the ground is a conductor and thus a large charge reservoir, a large current pulse is then injected into the base of the channel. This pulse can carry currents as high as hundreds of kA and is termed the “return stroke.” It propagates upward, decaying along the way due to recombination and channel heating, until it reaches the cloud where it neutralizes charge. The return stroke produces visible light, acoustic thunder due to heating, and radiates an intense electromagnetic pulse (EMP). A portion of the EMP propagates upward where it interacts with the ionosphere, while some of the energy also propagates long distances with little attenuation in the Earth-ionosphere waveguide. Such signals observed at long distances from the causative lightning are called radio atmospherics, or “sferics” [*Wait and Spies, 1964*].

Following the return stroke, in cloud discharges known as J and K processes occur and redistribute charge [*Rakov and Uman, 2003*, p. 112], after which a “dart leader” may begin propagating down the already ionized channel. When the dart leader reaches the ground, a second return stroke propagates up the channel. The first and subsequent return strokes are collectively called a lightning “flash.”

1.2 Ionosphere

This work focuses on the interaction between lightning radiation and the Earth’s ionosphere. In the upper parts of the Earth’s atmosphere, photoionization due to solar radiation produces a partially-ionized plasma region known as the ionosphere. This region has its lower cutoff at around 50 km during the day and at around 80 km during the night with no well-defined upper limit. The ionosphere exhibits substantial vertical structure due to the complicated production and chemical loss processes and is further divided into vertical layers known as the D-region, the E-region, and the F-region. Typical electron density profiles for the daytime and nighttime ionosphere are shown in Figure 1.4. The formation of these layers is due to the fact that different parts of the solar spectrum deposit their energy at various heights depending on the density and composition of the neutral atmosphere and recombination rates also depend on the atmospheric density [*Tascione, 1994*, p. 89].




Figure 1.3: A thundercloud begins with positive and negative charges arranged in layers (1). Eventually the local electric field exceeds the breakdown threshold and a leader is formed (2). This “stepped leader” propagates downward toward the ground (3) where it attaches (4). A large current is injected into the channel, termed the return stroke (5). Following the return stroke, J/K discharge processes occur inside the cloud and redistribute charge (6). A dart leader may follow, depositing charge and in and heating the already existing channel (7). When the dart leader reaches the ground, a second return stroke propagates along the channel (8). Additional dart leaders and subsequent return strokes may follow.

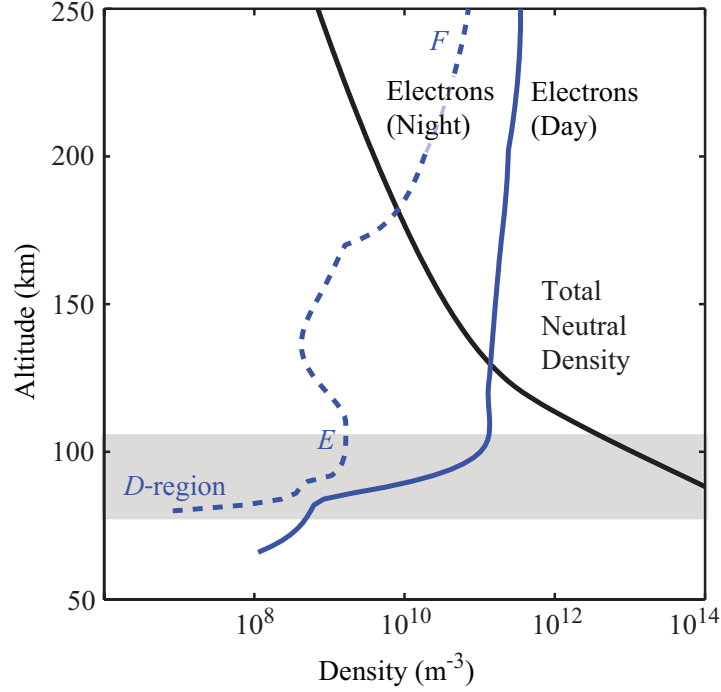


Figure 1.4: Typical electron and neutral density profiles in the ionosphere for daytime and nighttime from the International Reference Ionosphere (IRI) model. Adapted from [Marshall, 2009].

1.3 Transient Luminous Events

Though lightning is such a frequent global phenomenon, it is only in the last 25 years that science has begun to better understand the effects of lightning on the upper atmosphere and ionosphere. It was in 1989 that the first high-altitude optical flash, now known as a “sprite”, was scientifically observed [Franz *et al.*, 1990]. Since then, a number of high altitude optical phenomena collectively known as Transient Luminous Events (TLE), have been observed occurring above thunderstorms. In the following sections, we briefly introduce the main TLE classes in the order of their first discovery: sprites, elves, and jets. The properties of elves and the underlying causative lightning are the main focus of this work.

1.3.1 Sprites and Halos

Sprites are large, brief, flashes of light that occur above a thunderstorm following a CG lightning stroke. They occur at ~ 50 to 90 km altitude [Sentman *et al.*, 1995]. The lower portions of a sprite feature highly structured, tendril-like emissions known as streamers [Gerken *et al.*, 2000]. Often, the upper portion of a sprite features a diffuse red glow known as a “halo” [Barrington-Leigh *et al.*, 2001]. Halos can also occur without an associated sprite. Sprites and halos have been observed tens of kilometers laterally displaced from the causative CG by using triangulation techniques [São Sabbas *et al.*, 2003; Wescott *et al.*, 2001]. It has been theorized that this displacement is due to the effects of the Earth’s background magnetic field [Kabirzadeh *et al.*, 2015].

Sprites and halos are both results of the quasi-electrostatic (QE) field that is established after large CG strokes, which heats and ionizes the upper-atmosphere [Pasko *et al.*, 1997]. The strength of this QE field is related to the charge moment, Qh , of the lightning stroke, where Q is the amount of charge and h is the altitude from which it was removed. Sprites are typically associated with +CG strokes due to their larger charge removal and initiation altitudes [Williams, 1998], but have been occasionally observed following –CG strokes. The quasi-electrostatic field endures on the order of milliseconds to seconds – at higher altitudes, the electron conductivity is high and screening charges quickly redistribute, whereas the relaxation time is much longer at lower altitudes [Pasko *et al.*, 1997].

Several example images of sprites are shown in Figure 1.5.

1.3.2 Elves

Elves are a class of TLE that appear as rapidly expanding rings of light in the lower D-region ionosphere at altitudes of 80-90 km, and are centered above a causative CG lightning stroke, emitting primarily red photons with peak intensities in the tens of megarayleighs [Taranenko *et al.*, 1992; Taranenko and Inan, 1993]. They are brief in duration, 1-2 millisecond, have an apparent expansion rate faster than the speed of light, and have maximum radial extents of several hundred kilometers [Fukunishi *et al.*, 1996; Inan *et al.*, 1996, 1997]. Several images of observed elves are shown in




Figure 1.5: Some example sprite observations. a) The first image of sprites ever recorded, from [Franz *et al.*, 1990]. b) The first color image of sprites, from [Sentman *et al.*, 1995]. c) A color image of a sprite captured from the ISS. The red sprite can be faintly seen above the cloud. Courtesy of NASA.

Figure 1.6. Elves were predicted by Inan *et al.* [1991] and then subsequently first observed from space aboard the space shuttle using an intensified video rate camera [Boeck *et al.*, 1992]. The first ground-based observations were performed using three photometers [Fukunishi *et al.*, 1996], while the radial expansion of elves was verified using the “Fly’s Eye” photometer array [Barrington-Leigh, 2000; Inan *et al.*, 1997].

Elves are due to collisional heating of D-region electrons by the intense lightning-radiated electromagnetic pulse [EMP; Inan *et al.*, 1991, 1997]. As the EMP propagates through the ionospheric plasma, electrons are accelerated by the electric field. This acceleration increases the mean kinetic energy of the electron population (heating them). At D-region altitudes, the neutral atmosphere is still relatively dense – indeed, several orders of magnitude more dense than electrons. When these energetic electrons collide with neutrals, energy is transferred and the neutrals are excited to a higher energy state. As the neutrals relax to their ground state, a photon is emitted. This entire process – heating, collision, and relaxation – occurs on the order of microseconds, essentially instantaneous compared to EMP propagation and QE field relaxation timescales.

Elves are the primary focus of this dissertation. In Chapter 3, we provide a more detailed treatment of elves and the underlying physics of their production. In Chapter 4, we provide new results related to predicting elve production and estimation of the properties of the causative lightning. Finally, in Chapter 5, we estimate elve

production and heating effects on a global scale.




Figure 1.6: Some examples of elves observations. a) An elve captured by the ISUAL instrument onboard the FORMOSAT-2 satellite [Kuo *et al.*, 2007]. b) An elve captured from the space shuttle [Yair *et al.*, 2003]. c) An elve simulated using the model described in Marshall [2012]. d) An elve observed in January 2015 over the Mediterranean (courtesy Oscar van der Velde)

1.3.3 Jets

The final class of TLEs that we consider are known as “jets”. Whereas sprites and elves are direct consequences of electromagnetic fields due to CG lightning, the quasi-electrostatic field and the return stroke-radiated EMP respectively, jets are discharges emanating from the tops of clouds and can be considered a form of upwardly-directed lightning. Jets are divided into two main classes: blue jets and gigantic jets. Blue jets are relatively frequent discharges extending from the cloud-top to around 40 km altitude and, as the name suggests, are primarily blue in color [Wescott *et al.*, 1995]. Gigantic jets, on the other hand, are exceedingly rare, extend from the cloud top to the bottom of the D-region ionosphere (about 80 km altitude) and exhibit streamer-like structure in their upper portions [Pasko *et al.*, 2002]. Some example blue jet and gigantic jet observations are shown in Figure 1.7.




Figure 1.7: Some example images of blue jets and gigantic jets. a) One of the first blue jets observed by [Wescott *et al.*, 1995]. b) A color image of a blue jet observed in Australia (courtesy of Peter Jarver and Earle Williams). c) The first gigantic jet observed by [Pasko *et al.*, 2002].

1.4 Review of Prior Work

Several prior studies have attempted to quantify the production rate and probability of elves as a function of lightning parameters, in particular the National Lightning Detection Network (NLDN)-reported peak current. The first such study was conducted by Barrington-Leigh and Inan [1999] using the “Fly’s Eye” photometer which recorded 34 manually-triggered events that were identified as elves. All of the elves had NLDN peak currents greater than 56 kA, suggesting a peak current threshold for elve production. Measurements from the satellite borne Imager of Sprites and Upper Atmospheric Lightning (ISUAL) instrument, another triggered photometer, suggested that all CGs with peak current magnitudes greater than 60 kA produce elves [Chen *et al.*, 2008]. A later study by [Newsome, 2010; Newsome and Inan, 2010] conducted over multiple summers using the non-triggered Photometric Imaging of Precipitation of Electron Radiation (PIPER) photometer in which ~ 1000 elves were observed suggested, however, that only about 10% of 60 kA CG strokes produce elves. Newsome’s probability of elve production given NLDN peak current is shown in Figure 1.8. Puzzlingly, the results of [Newsome, 2010] show a surprising number of large CG strokes, many greater than 100 kA, for which an elve was not observed.

A global distribution and production rate for elves was also given by [Chen *et al.*, 2008] using satellite data from the ISUAL photometer. This data, however, is biased




Figure 1.8: Elve production probability given NLDN peak current as calculated by [Newsome, 2010]

by the fact that ISUAL used triggered recording and the fact that it was in a polar orbit with a narrow field of view. No prior studies have shown a global elve rate inferred from non-triggered ground based measurements.

The velocity at which current propagates along the lightning channel, termed the return stroke speed, remains a somewhat open question due to the difficulty of making direct measurements. The return stroke speed is an important assumption that must be made when inferring peak currents from remote field measurements. This speed has typically been assumed to lie between $\frac{1}{5}$ and $\frac{2}{3}$ of the speed of light based on optical streak camera and high-speed digital camera measurements [Hubert and Mouget, 1981; Idone and Orville, 1982; Idone et al., 1984; Rakov, 2007; Wang et al., 1999; Mach and Rust, 1989]. However, recent thermodynamic modeling of the return stroke channel [Liang et al., 2014] has suggested, however, that the optical and current wave speed may differ, with the current speed being as fast as $0.8c$. Rakov and Tuni [2003] studied the relationship between return stroke speed and electric field intensity at D-region altitudes, suggesting that a return stroke speed $\geq 0.8c$ is needed for a 30 kA CG to produce an elve.

1.5 Contributions of this Thesis

This dissertation aims to develop a detailed understanding of the conditions under which elves are produced. In particular, we use a dataset that was collected using free-running ground-based photometers and VLF/LF radio receivers to develop a statistical model of elve production. We also develop a model using the geometric features of elve observations, in conjunction with numerical simulations of the EMP-ionosphere interaction, to determine the most likely distribution of lightning return stroke speeds. This dissertation constitutes the first measurement of return stroke speed that relates directly to the current propagation rather than optical measurements. Finally, using our elve production model, we extrapolate a global elve production rate and estimate the contribution of lightning to D-region heating.

To summarize, the main contributions of this thesis are:

1. Determination of the lightning return stroke properties that control elve production using optical elve observations and VLF/LF measurements of lightning-generated sferics. In particular, it is shown that the peak magnetic field of a sferic, and consequently the NLDN peak current estimate, can be used to predict whether a stroke produces an elve with high accuracy.
2. Development of a new method for estimating the lightning current return stroke speed using the geometric features of optical elve observations. It is found that the size of the elve “hole” radius is related to the return stroke speed, with faster return strokes producing elves with smaller holes. A Bayesian inference model is built that uses elve hole radii in conjunction with physical simulations of elves to estimate the probability distribution of return stroke speeds in the dataset.
3. Estimation of the global production rate and geographic distribution of elves using a logistic regression model trained with optical elve observations and NLDN data collected over North America along with peak current estimates from the GLD360 network. Good agreement is found with production rates estimated using the satellite-borne ISUAL instrument.

4. Development of a model that predicts the amount of D-region heating by the lightning EMP as a function of the return stroke peak current using numerical physical simulations of the lightning-ionosphere interaction. The total amount of EMP-driven D-region heating is extrapolated globally using GLD360 events.

1.6 Thesis Organization

- In Chapter 2, we discuss the details of the lightning return stroke. In particular, we describe modeling of the current propagation, the EMP radiation, and remote sensing of the EMP.
- Chapter 3 is concerned with elves. We discuss the detailed physics of the EMP-ionosphere interaction, photon production, photometric remote sensing of elves, and we discuss the experimental campaign performed during 2013.
- Chapter 4 uses experimental observations to explore detailed properties of elve-producing lightning. In particular, we model the probability of elve production as a function of return stroke properties. We also use the geometric features of elve observations to infer the distribution of lightning return stroke speeds.
- Chapter 5 uses the results presented in Chapter 4 to extrapolate elve production to a global scale and also the contribution of lightning to D-region heating.
- Finally, in Chapter 6 we summarize the contributions of this thesis and provide suggestions for future work.

Chapter 2

The Lightning Return Stroke

2.1 Background

The return stroke is the most important and extensively studied aspect of the lightning process. It carries extremely large currents, dissipates large amounts of energy, and is believed to cause the most damage to structures on the ground. The return stroke also produces most of the electromagnetic radiation in the lightning flash, produces the visible optical flash, and contributes to heating the surrounding channel air which produces thunder.

In this chapter, we describe the method in which the return stroke current and its propagation are modeled. We also derive the form of the return stroke electromagnetic pulse (EMP), the propagation of the EMP in the Earth-ionosphere waveguide, as well as how it is recorded using ground-based VLF/LF receivers.

2.2 Modeling the Return Stroke Current

The lightning return stroke is a highly complex physical process, hence it is often necessary to use simple mathematical models that are amenable to computer simulation. Most return stroke models used fall into one of four main categories:

- (i) Gas dynamic models solve for the radial evolution of the lightning channel along

with the gas temperature, pressure, and density along the channel segment. One such model, the Stanford Time-Domain Fractal Lightning (TDFL) model is briefly described in a later section.

- (ii) Electromagnetic models approximate the lightning channel as a lossy, thin wire and numerically solve Maxwell's equations to find the current distribution along the channel at each point in time.
- (iii) Distributed-circuit models treat the channel as a vertical RLC transmission line which is parameterized by an inductance (L), capacitance (C), and resistance (R) per unit length.
- (iv) Engineering models are simple, heuristic models which give the current profile along the channel at each time step and are specified by a very small number of parameters.

In the following sections we discuss the formulation for the various engineering models, as they are by far the most common return stroke models due to their simplicity. We also provide a brief description of gas dynamic modeling.

2.2.1 Engineering Models

Return stroke engineering models are simple mathematical equations that relate the channel base current $I(0, t)$ to the current at altitudes z' along the channel $I(z', t)$. They effectively take a current pulse that is injected into the bottom of the channel and specify the propagation of the pulse upward along the channel. This propagation can be generalized according to the following equation [[Rakov and Uman, 1998](#)]:

$$I(z', t) = u(t - z'/v_{\text{RS}})P(z')I(0, t - z'/v) \quad (2.1)$$

where $u(\cdot)$ is the Heaviside step function, $P(z')$ is an altitude dependent attenuation factor, and v_{RS} is the upward-propagating return stroke speed. A tabulation of these parameters for some of the most commonly used engineering models is given in Table [2.1](#).

Model	$P(z')$	v
TL	1	v_{RS}
MTLL	$1 - z'/\ell_{\text{chan}}$	v_{RS}
MTLE	$\exp(-z'/\lambda)$	v_{RS}
BG	1	∞
TCS	1	$-c$

Table 2.1: Altitude attenuation constants and velocities for five return stroke engineering models. Table adapted from [[Rakov and Uman, 2003](#)].

Before an engineering model can be used, the current pulse injected into the channel at the ground, called the channel base current $I(0, t)$ must first be specified. This pulse is most typically modeled as either a linear rise with exponential decay

$$I(0, t) = \begin{cases} I_p t / \tau_r & t < \tau_r \\ I_p \exp[-(t - \tau_r)/\tau_f] & t \geq \tau_r \end{cases} \quad (2.2)$$

or a linear rise with Gaussian decay

$$I(0, t) = \begin{cases} I_p t / \tau_r & t < \tau_r \\ I_p \exp[-(t - \tau_r)^2 / \tau_f^2] & t \geq \tau_r \end{cases} \quad (2.3)$$

In both of these equations, I_p is the *peak current*, τ_r is the *rise time*, and τ_f is the *fall time*. Optionally, a continuing current may be specified (not shown here), which greatly affects the return stroke charge moment and, hence, sprite and halo production. It is also important to note that these base current formulations result in a non-differentiable high-frequency cusp at $t = \tau_r$, so low-pass filtering the current is often necessary for stability when simulating the return stroke using numerical finite-difference schemes. Example channel base current pulses with linear rise and both Gaussian and exponential decays are shown in [Figure 2.1](#).

Once the channel base current pulse is specified, an engineering model can be used to model the upward propagation along the lightning channel. There are two main classes of engineering models: transmission line-type models in which a current pulse is injected at the base, and traveling current source-type models in which a current source travels upwards at the speed of light and injects current downwards.

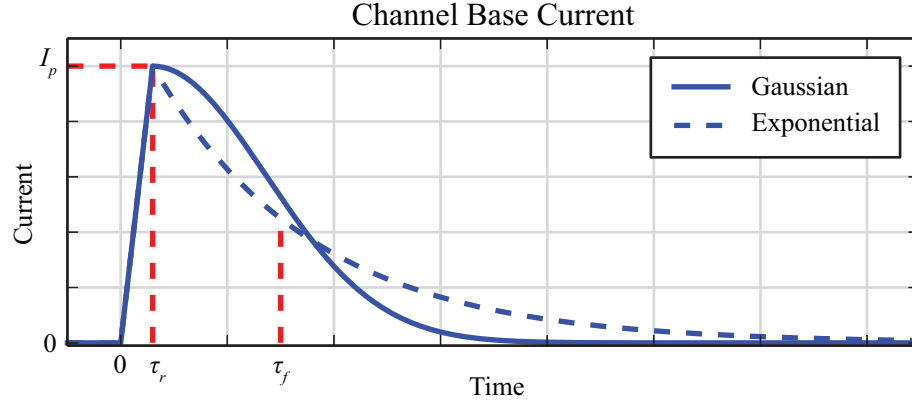


Figure 2.1: Typical channel base current pulses with a linear rise time, τ_r , a peak current of I_p , and both Gaussian and exponential decays with fall time τ_f .

The simplest transmission line model, called TL [[Uman and Mclain, 1969](#)], propagates the pulse upward along the channel at the *return stroke speed* v_{RS} :

$$I(z', t) = I(0, t - z'/v_{RS})u(t - z'/v_{RS}) \quad (2.4)$$

The TL is extremely simplistic and not very physical – we intuitively expect the current to attenuate as it propagates upward due to heating losses and recombination along the channel. The two most popular transmission line models accounting for altitude decay are the *modified transmission line with linear decay* (MTLL) model [[Rakov and Dulzon, 1987](#)]

$$I(z', t) = I(0, t - z'/v_{RS}) \left(1 - \frac{z'}{\ell_{\text{chan}}}\right) u(t - z'/v_{RS}) \quad (2.5)$$

and the *modified transmission line with exponential decay* (MTLE) model [[Nucci et al., 1988](#)]

$$I(z', t) = I(0, t - z'/v_{RS}) \exp(-z'/\lambda) u(t - z'/v_{RS}) \quad (2.6)$$

In the previous two equations, ℓ_{chan} is the total length of the lightning channel and λ is an exponential decay constant. The MTLL and MTLE models are both very popular due to their simplicity and the close match between their simulated

fields and observations. Both models are used extensively in simulations presented in subsequent chapters of this thesis.

Of the traveling current source-type models, the Bruce-Golde (BG) model [[Bruce and Golde, 1941](#)] is the simplest. It imposes that the channel base current $I(0, t)$ propagates uniformly along the channel at a front speed v_f , essentially with an infinite return stroke speed:

$$I(z', t) = I(0, t)u(t - z'/v_f) \quad (2.7)$$

The traveling current source (TCS) model [[Heidler, 1985](#)] models a current source propagating upward at v_f injecting a downward current pulse, itself traveling at the speed of light

$$I(z', t) = I(0, t + z'/c)u(t - z'/v_f) \quad (2.8)$$

Finally, the Diendorfer-Uman (DU) model [[Diendorfer and Uman, 1990](#)] modifies the TCS model to also include an upward current with exponential time decay so as to better match observed fields

$$I(z', t) = I(0, t + z'/c) - \exp\left(-\frac{t - z'/v_f}{\tau_D}\right) I\left(0, z'\frac{1 + v_f/c}{v_f}\right) u(t - z'/v_f) \quad (2.9)$$

where τ_D is a time constant specifying how the amount of charge per unit length deposited along the channel decays with time [[Rakov and Uman, 2003](#), p. 402].

It should be noted that the engineering models described here are only useful insofar as the radiated fields agree with data. An extensive comparison of merits and drawbacks for each of these models is presented in [[Rakov and Uman, 1998](#)].

2.2.2 Gas Dynamic Models

The engineering models presented in the previous section are extremely simplistic and easy to use and parameterize, at the expense of being rather unphysical. The family

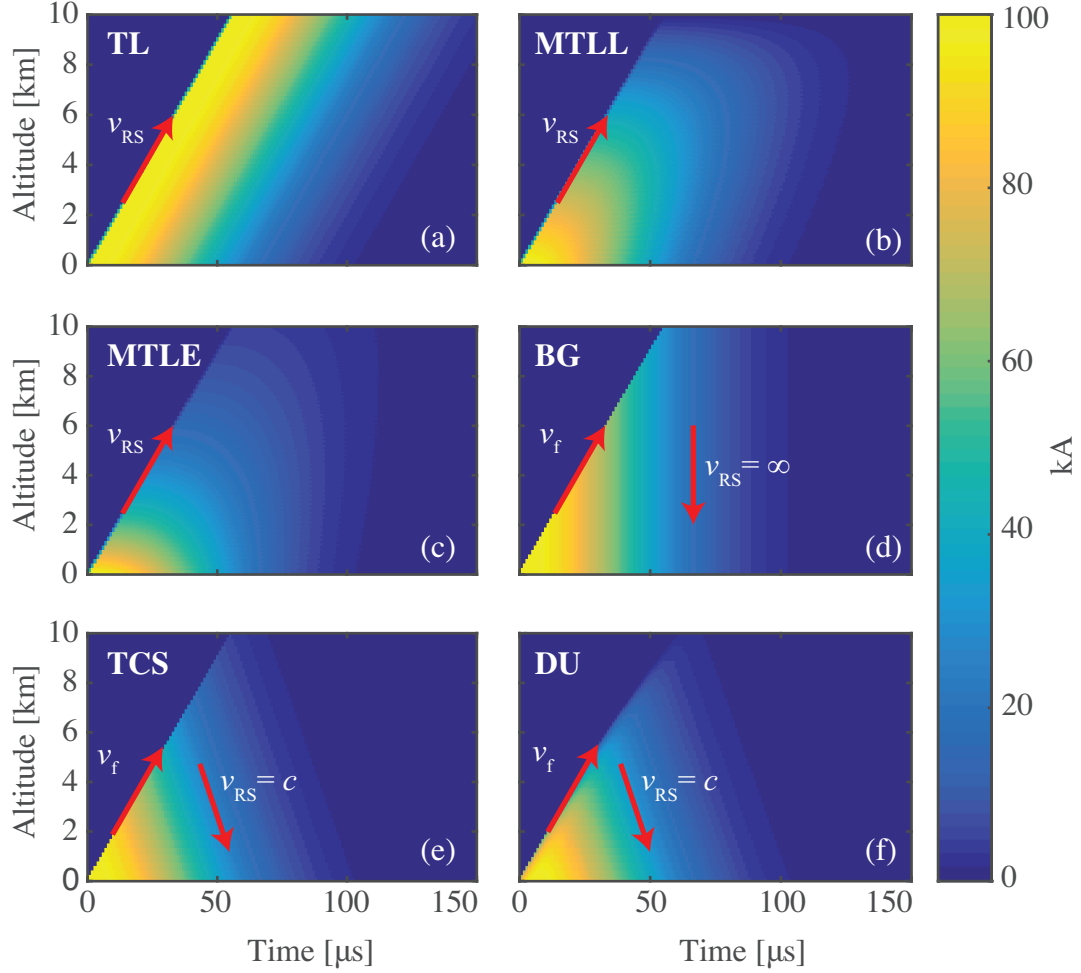


Figure 2.2: Current history plots from the major return stroke engineering models.

of gas dynamic models, in contrast, capture the physics of the return stroke process in much greater detail [e.g. [Dubovoy et al., 1995](#)]. Such models typically involve self-consistent solutions for the temporal and spatial evolution of the channel thermodynamic and electrodynamic properties. The thermodynamic properties solved for include the temperature and pressure of the channel core as well as the electrical and thermal conductivities. Once these parameters are determined at each simulation time step, the electrodynamic update solves the current propagation and associated electromagnetic radiation. The main downside of gas dynamic return stroke models

is their computational complexity as well as their large number of uncertain input parameters. For example, these models typically require initial temperature, pressure, and core radius distributions which may not be well known.

One example of a gas dynamic model is the recently-developed Stanford Time Domain Fractal Lightning (TDFL) model [Liang, 2014; Liang *et al.*, 2014]. This model includes the thermodynamics of multiple fluids (ions, electrons, and neutrals), accounts for the dynamics of the channel sheath, and simulates optical photon production in the channel. This model was used to suggest that the current and optical return stroke speeds may differ, and that the current return stroke speed may be much faster than previously believed – as fast as $0.83c$ [Liang *et al.*, 2014]. An example current history from the TDFL model is shown in Figure 2.3.

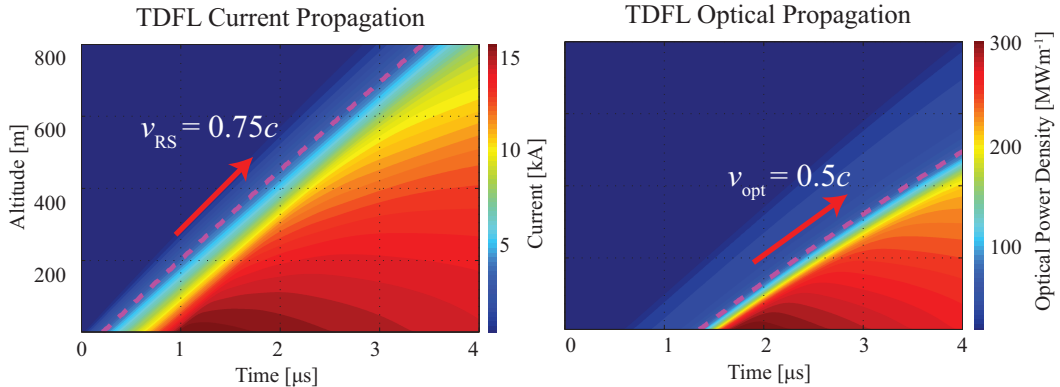


Figure 2.3: An example current history from the TDFL return stroke model simulated with an initial core radius of 4 mm, temperature of 20000 K, and pressure of 1 atm. This simulation resulted in a return stroke with a speed of $0.75c$ and a current history that is very similar to those produced by the MTLE engineering model.

2.3 Radiated Fields

The observed magnetic field from a return stroke EMP can be easily computed by making two simplifying assumptions: that the lightning channel is a vertical straight line and that the ground is flat and perfectly conducting. With these assumptions, the method of image currents, also assuming that current is symmetric about the ground

plane, the magnetic vector potential (in cylindrical coordinates) on the ground a distance r from the source is given by:

$$\mathbf{A}(r, t) = \frac{\mu_0}{4\pi r} \int_{-\ell_{\text{chan}}}^{\ell_{\text{chan}}} I\left(z', t - \frac{r}{c}\right) dz' \hat{z} \quad (2.10)$$

where r is radial distance along the ground, z' is altitude, ℓ_{chan} is the length of the vertical lightning channel, and c is the speed of light. Using the Lorentz Gauge, the magnetic flux density can now be found according to

$$\mathbf{B} = \nabla \times \mathbf{A} \quad (2.11)$$

which yields the azimuthal field

$$\mathbf{B}(r, t) = \underbrace{\frac{\mu_0}{2\pi r^2} \int_0^{\ell_{\text{chan}}} I(z', t') dz' \hat{\phi}}_{\text{Induction Field}} + \underbrace{\frac{\mu_0}{2\pi cr} \frac{\partial}{\partial t'} \int_0^{\ell_{\text{chan}}} I(z', t') dz' \hat{\phi}}_{\text{Radiation Field}} \quad (2.12)$$

where $t' = t - r/c$ is the retarded time. The first term on the right-hand side of Equation 2.12 is the induction field and the second term is the radiated field. Since the induction field decays as $1/r^2$ while the radiation field varies as $1/r$, the radiation field is more important for remote sensing distant lightning. Accordingly, in the far-field (at distances of several hundred meters), the magnetic field is well approximated by

$$\mathbf{B}(r, t) = \frac{\mu_0}{2\pi cr} \frac{\partial}{\partial t} \int_0^{\ell_{\text{chan}}} I\left(z', t - \frac{r}{c}\right) dz' \hat{\phi} \quad (2.13)$$

The integral of current with respect to altitude, $\int_0^{\ell_{\text{chan}}} I(z', t') dz'$, is called the *current moment*. Hence, the EMP magnetic field strength is proportional to the time derivative of the return stroke current moment. This proportionality is illustrated in Figure 2.4, and allows us to build some useful intuition: properties that increase the rate of change (return stroke speed, rise-time) or that increase the current moment magnitude (channel length, peak current) result in stronger EMPs.




Figure 2.4: In the left hand panel, we show the current history along the channel for a 100 kA return stroke. This current propagation is modeled using the MTLL engineering model. The right hand panel shows the resulting EMP groundwave calculated using Equation 2.13 at a distance of $r = 100$ km from the channel.

A more general expression was derived in [Newsome, 2010] for the radiated magnetic field in spherical coordinates which is not restricted to observation on the ground:

$$\mathbf{B}(r, \theta, t) = \frac{\mu_0 \sin \theta}{4\pi cr} \frac{\partial}{\partial t'} \int_{-\ell_{\text{chan}}}^{\ell_{\text{chan}}} I\left(z', t' + \frac{\cos \theta}{c} z'\right) dz' \hat{\phi} \quad (2.14)$$

where θ is the spherical elevation angle. This expression can be used to approximate the radiation patterns from lightning return stroke currents.

2.4 Earth-Ionosphere Waveguide Propagation

As introduced in Section 1.2, the ionosphere is a plasma region in the upper atmosphere, starting at around 50 km during the day and 80 km at night, which is primarily created by photoionization. As the EMP expands radially outward from the lightning source, as derived in the previous section, it propagates through the lower atmosphere, which is well approximated as free space. However, once it reaches the D-region ionosphere, it encounters the dense, collisional, magnetized plasma and the propagation of the VLF wave energy becomes much more complex.

Propagation through the ionosphere is anisotropic, frequency dependent, and subject to variations in electron density and the collision frequency (i.e. neutral density).

Neglecting ion motion, the dispersion relation in the D-region is well approximated by the Appleton-Hartree equation (e.g. [*Inan, Umran S and Gołkowski, 2010*, p. 192]):

$$n^2 = 1 - \frac{X}{U - \frac{Y^2 \sin^2 \theta}{2(U-X)} \pm \sqrt{\frac{Y^4 \sin^4 \theta}{4(U-X)^2} + Y^2 \cos^2 \theta}} \quad (2.15)$$

where

$$U = 1 + i \frac{\nu}{\omega} \quad (2.16)$$

$$Y = \frac{\omega_c}{\omega} \quad (2.17)$$

$$X = \frac{\omega_p^2}{\omega} \quad (2.18)$$

and

$$\omega_p^2 = \frac{q N_e}{m_e \epsilon_0} \quad (2.19)$$

$$\omega_c = \frac{q B_0}{m_e} \quad (2.20)$$

In these equations, ω_p is the electron plasma frequency, ω_c is the electron gyro-frequency, ν is the electron-neutral collision frequency, ω is the frequency of the incident wave, N_e is the electron density, q is electron charge, m_e is electron mass, B_0 is the Earth's background magnetic field, and θ is the angle between B_0 and the wave vector. Ignoring the background magnetic field and assuming $\nu \gg \omega$ (valid for VLF/LF waves in the lower ionosphere), the index of refraction can be simplified as:

$$n^2 = 1 - \frac{X}{U} \quad (2.21)$$

$$= 1 + i \frac{\omega_p^2}{\nu \omega} \quad (2.22)$$

Reflection occurs when the real and imaginary components of Equation 2.22 are equal [Ratcliffe, 1959], which occurs when $\omega = \omega_p^2/\nu$. For VLF waves in a nighttime ionosphere, this reflection occurs at an altitude of approximately 85 km.

The conductive earth surface and the reflective ionosphere essentially form a waveguide in which lightning EMPs, called radioatmospherics or *sferics* [Wait and Spies, 1964], can propagate to long distances with very low attenuation. This waveguide deviates from the ideal case of a perfectly conducting parallel plate waveguide due to earth curvature and the anisotropy of the ionosphere due to B_0 . It supports the so-called quasi-transverse-electric (QTE), quasi-transverse-magnetic (QTM), and quasi-transverse-electromagnetic (QTEM) modes, each with small field components in the propagation direction.

The sferics utilized in this thesis were observed relatively close to the causative CG stroke (<1000 km), close enough so that the LF frequency components were not significantly attenuated. In such cases, it makes sense to model the sferic propagation using a ray-hop approach. Figure 2.5 shows a sketch of the Earth-ionosphere waveguide along with an actual sferic magnetic field waveform received at a station in Kansas 300 km away from the causative CG stroke. The first part of the sferic to be received propagates along the direct line of sight and is referred to as the *ground wave*. Some of the EMP energy propagates upward and reflects from the D-region exactly halfway between the CG and the receiver. This ionospheric reflection, known as the *1st sky wave*, appears at the receiver as an inverted, time-delayed version of the ground wave pulse, but with significant distortion due to ionospheric anisotropy and frequency dependent absorption. Similarly, further sky waves can be observed that undergo multiple reflections between the ionosphere and the ground. The ground wave is unaffected by the ionosphere, and hence contains the most information about the return stroke current, making it particularly useful for inferring return stroke

properties. The sky wave, on the other hand, can be useful for estimating the ionospheric reflection height, and hence electron density, by observing its time delay with respect to the ground wave [e.g., [Lay et al., 2014](#)].

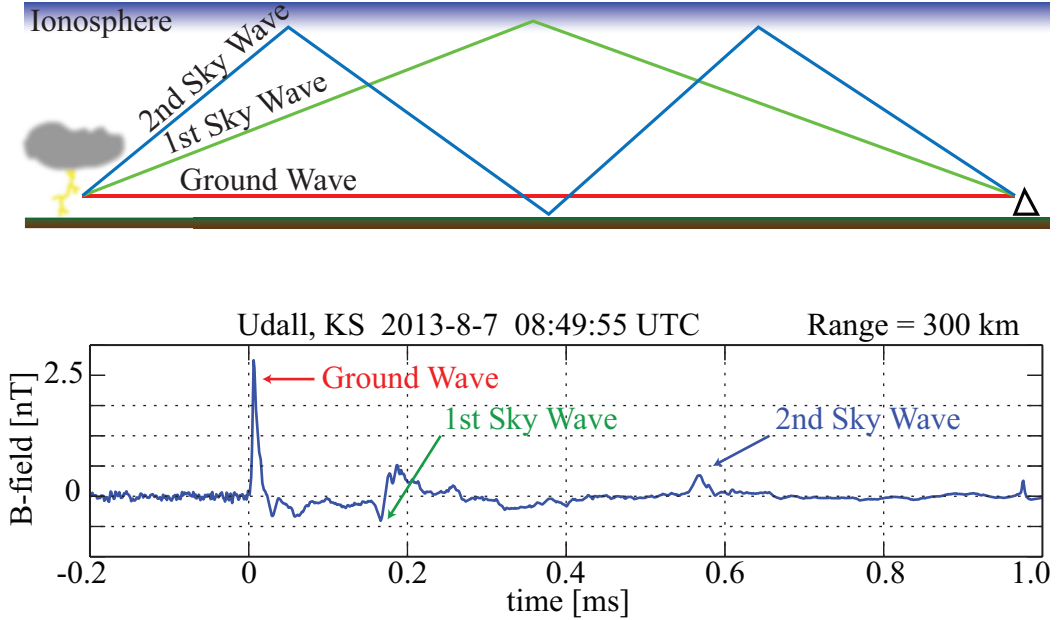


Figure 2.5: Here we show how multiple reflections in the Earth-ionosphere waveguide map to peaks in the received sferic waveform.

2.5 Radio Remote Sensing of the EMP

The sferics radiated by the lightning return stroke propagate efficiently in the earth-ionosphere waveguide, attenuating only 2-3 decibels per megameter [[Wait and Spies, 1964](#)] at VLF frequencies, making it efficient to detect lightning on the ground far from the causative stroke. The analysis in this thesis makes extensive use of individual sferics observed at several ground stations to infer properties of the source lightning.

This long distance propagation also enables accurate geolocation of the lightning stroke. The National Lightning Detection Network (NLDN) is a network of LF receivers in North America. It uses magnetic direction finding (MDF) and time of arrival (TOA) for sferics observed at multiple stations to triangulate the lightning

location [Cummins *et al.*, 1998]. NLDN also provides an estimated peak current by correlating the observed sferic signal strength with the expected radiation from a return stroke with TL current propagation. The Global Lightning Detection Network (GLD360) is a global network of VLF receivers based on a long-range waveform-bank technique [Said *et al.*, 2010]. GLD360 computes the arrival azimuth of observed sferics and correlates against an empirical sferic waveform bank to estimate distance. These results from multiple stations are then sent to a central processor where the lightning location is triangulated. GLD360 also provides peak current estimates. NLDN provides significantly better geolocation accuracy due to the fact that it uses LF sferic observations, but its coverage is limited to North America. GLD360, on the other hand, has global coverage but worse geolocation accuracy since it uses VLF frequencies.

In the next section, we provide a description of the instrumentation used to record the lightning EMP and sferic emissions for this thesis.

2.5.1 VLF/LF Receiver Design

Most of the energy radiated by the lightning return stroke falls in the Very Low Frequency (VLF) and Low Frequency (LF) bands. The VLF band is defined from 300 Hz – 30 kHz and LF from 30 kHz – 300 kHz. To sense waves in this frequency range, we use a magnetic field sensor; in particular an LF variant of the Atmospheric Weather Electromagnetic System for Observation, Modeling, and Education (AWESOME) receiver is used [Cohen *et al.*, 2010]. A high-level block diagram of this system is shown in Figure 2.6.

The magnetic field sensor used is an air-core B -field loop antenna which detects time varying magnetic fields orthogonal to the loop plane. Two orthogonal loops are used so that the arrival azimuth of incoming waves can be determined. A photograph of one of the antenna pairs used is shown in Figure 2.7. The antenna characteristics can be described by four parameters: The wire resistance R , the loop inductance L , the loop area A , and the number of turns N . The sensitivity S , defined as the smallest detectable magnetic field normalized the factor $1/f$, in units of $(\text{T} - \text{Hz}^{1/2})$

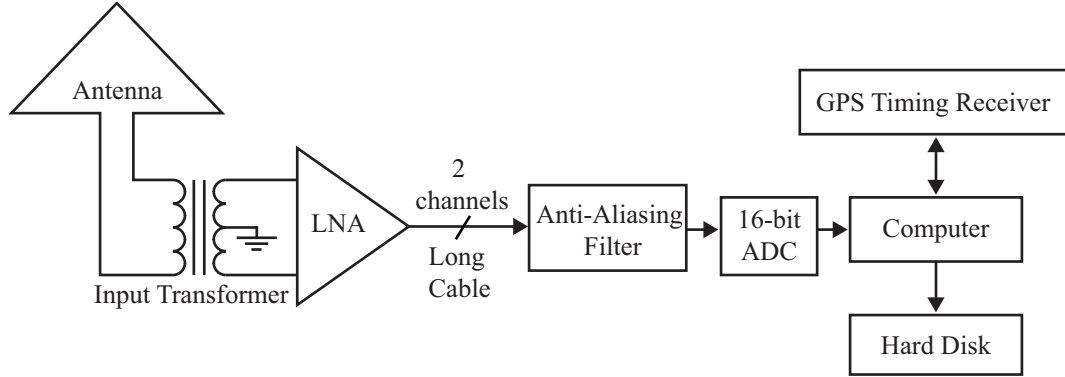


Figure 2.6: Schematic block-diagram of the VLF/LF receiver systems

is limited by the thermal noise of the wire and is given by

$$S = \frac{\sqrt{4kTR}}{2\pi NA} \quad (2.23)$$

The inductance L essentially determines the low-frequency roll-off of the antenna. Hence, an antenna with a low-impedance, $1\ \Omega$ – $1\ \text{mH}$ is used, which gives a lower cutoff of roughly $160\ \text{Hz}$. The output of the antenna is connected to a transformer, which is used to match the low-impedance antenna loop to the high-impedance input of a low-noise amplifier (LNA). A detailed description of this analog front-end is provided in [Harriman *et al.*, 2010].

The output of the LNA is used to drive a long ($\sim 100\ \text{m}$) cable which feeds into the digital portion of the system. The analog front-end is separated from the digital portion of the system in order to prevent emissions from the noisy digital electronics from coupling into the antenna. The signal is then filtered by a $500\ \text{kHz}$ anti-aliasing filter (AAF) and then sampled at $1\ \text{Msample/second}$ by a 16-bit analog-to-digital converter (ADC). The digitized data is stored on a hard drive and accurately timestamped by a GPS timing receiver. The GPS disciplined timing is necessary to correct for clock drift so that data from disparate receivers and optical instruments can be compared, while also enabling correlation of detected sferics with the timestamps of NLDN reported events.

The AWESOME receiver is a broadband receiver, meaning that the signal from



Figure 2.7: A photograph of a VLF/LF receiver loop antenna used in the Oklahoma array.

the antenna is digitized directly, without down conversion or demodulation. This broadband property is important because the lightning sferics span almost the entire VLF/LF range, from hundreds of Hz up to hundreds of kHz. A spectrogram of one second of broadband data is shown in Figure 2.8. Several features stand out: the horizontal lines at constant frequencies are VLF transmitters operated by the US Navy for submarine communications, while the impulsive, vertical lines are lightning generated sferics.

2.6 Summary

The return stroke is the process by which a large current pulse propagates along the cloud-to-ground lightning channel, ultimately neutralizing charge in the cloud. The intense current pulse, up to several hundred kiloamperes, propagates upward along the channel at a significant fraction of the speed of light. Several methods exist for modeling this current propagation, with the MTLL and MTLE engineering models being the most popular due to their simplicity and reasonable match with

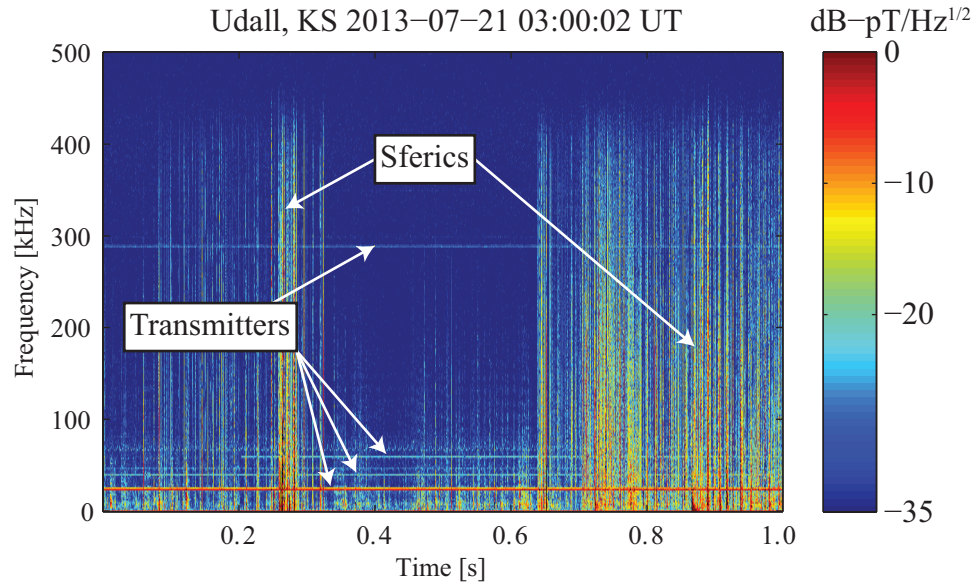


Figure 2.8: An example VLF/LF spectrogram captures at the Udall, KS field site during the summer of 2013.

observations.

The return stroke is responsible for most of the electromagnetic radiation in lightning. Each return stroke radiates an intense EMP which expands radially outward from the lightning source, interacting with the ionosphere and propagating in the Earth-ionosphere waveguide as a “sferic.” Sferics can be detected using broadband VLF/LF receivers far away from the causative stroke, enabling us to geolocate the stroke and infer properties of the return stroke current.

Chapter 3

Elves

In this chapter, we discuss the details of elve production and fast photometry techniques. Finally, we describe an experiment involving the observation of hundreds of simultaneous elve and lightning sferic emissions.

3.1 Elves and the EMP-Ionosphere Interaction

The D-region ionosphere is a plasma comprised of free electrons, ions, and neutral molecules. As the lightning return stroke EMP propagates through this plasma, the electrons are accelerated by the electric field, increasing their mean kinetic energy. Ions are also heated by the EMP, but since they are several orders of magnitude more massive than electrons, ion heating is negligible compared to electron heating. It is these EMP-heated electrons and their subsequent collisions with neutral molecules that produces the optical photon emissions known as elves. A schematic illustration of this effect is shown in Figure 3.1.

At D-region altitudes, the background neutral atmosphere is comprised of molecular nitrogen (N_2), molecular oxygen (O_2), argon (Ar), atomic oxygen (O), helium (He), hydrogen (H), and atomic nitrogen (N). N_2 accounts for about 79% of total neutral content while O_2 accounts for about 20%. The neutral densities as a function of altitude are shown in Figure 3.2, with the densities taken from the MSIS-E-90 atmosphere model [*Hedin*, 1991], along with a nighttime electron density profile, taken




Figure 3.1: Cartoon showing the EMP-ionosphere interaction

from the IRI-2007 ionosphere model [[Bilitza and Reinisch, 2008](#)].

The total neutral density in the D-region is many orders of magnitude higher than the electron density. Hence, the lower ionosphere is a highly collisional plasma – electrons that are accelerated by the lightning EMP very quickly collide with the neutrals and transfer their energy. There are two main types of electron-neutral collision processes: elastic and inelastic collisions. Elastic collisions are those in which the total kinetic energy of the electron-neutral system remains constant. The internal state of the neutral atom or molecule is unchanged. Such collisions have the general effect of returning the electron distribution to thermal equilibrium. During inelastic collisions, on the other hand, some of the electron’s kinetic energy changes the state of the neutral molecule.

There are several types of inelastic electron-neutral collisions. Impact ionization strips an electron from the neutral, enhancing the local free electron density. Electron attachment to the neutral, especially O_2 , decreases the electron density. Various excitations also occur which raise the neutrals to a higher energy internal state, including rotational, vibrational, and electronic excitation states. When the excited

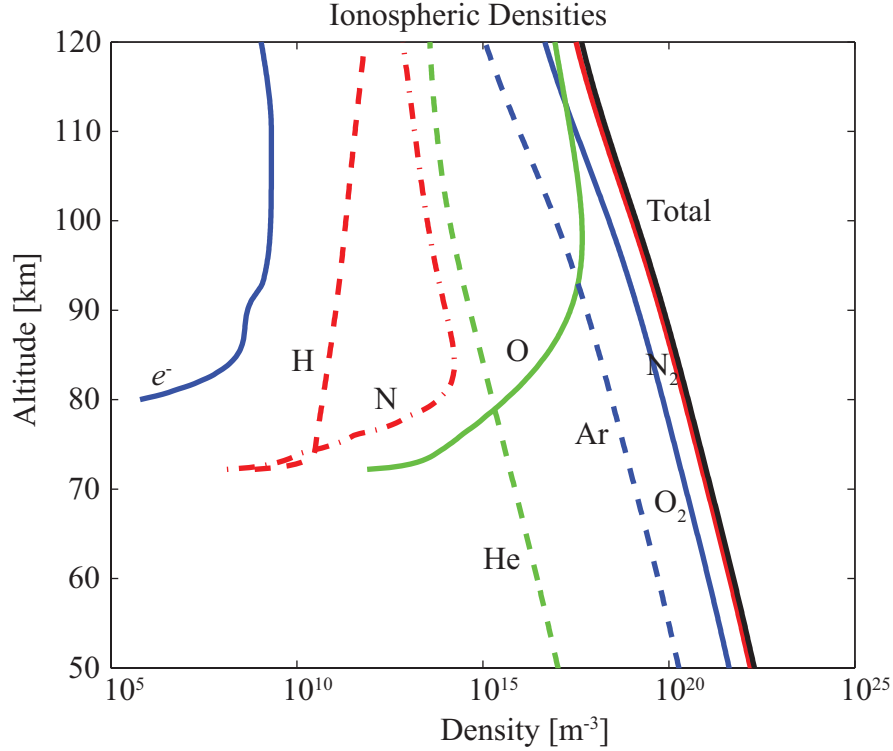


Figure 3.2: Altitude density profile for the major constituents of the D-region ionosphere. The nighttime electron density, provided by the IRI ionosphere model, is shown as a thick blue curve. The neutral atoms and molecules are provided by the MSIS atmosphere model.

neutral relaxes to a lower-energy state, a photon is produced, giving rise to the optical emissions such as elves. The energy thresholds for some of the important collision processes are tabulated in Table 3.1 and are adapted from [Moss *et al.*, 2006].

Once a neutral is excited into a higher energy state, there are two main mechanisms through which the molecule or atom may relax to a lower energy state: photon emission and quenching. Quenching is a transfer of energy through neutral-neutral collisions. Table 3.2 provides the optical excitation rate, A_k , for neutrals in state k , and Table 3.3 provides the quenching rates for N_2 , $\alpha_k^{N_2}$, and O_2 , $\alpha_k^{O_2}$. These rates are calculated by numerically solving the Boltzmann equation with the ELENDIF code and are tabulated in [Moss *et al.*, 2006]. The total radiative lifetime, τ_k , for a neutral in energy state k is then given by:

Collision Process	Reaction	Threshold (eV)
elastic	$e + \text{N}_2 \rightarrow e + \text{N}_2$	–
	$e + \text{O}_2 \rightarrow e + \text{O}_2$	–
ionization	$e + \text{N}_2 \rightarrow 2e + \text{N}_2^+ (\text{X}^2\Sigma_g^+ + \text{A}^2\Pi_u)$	15.60
	$e + \text{O}_2 \rightarrow 2e + \text{O}_2^+ (\text{X}^2\Pi_g)$	12.06
three-body attachment	$e + \text{O}_2 + \text{A} \rightarrow \text{O}_2^- + \text{A}$	–
two-body attachment	$e + \text{O}_2 \rightarrow \text{O}^- + \text{O}$	–
electronic excitation	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{A}^3\Sigma_u^+)$	6.17–7.80
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{B}^3\Pi_g)$	7.35
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{W}^3\Delta_u)$	7.36
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{B}'^3\Sigma_u^-)$	8.16
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{a}'^1\Sigma_u^-)$	8.40
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{a}^1\Pi_g)$	8.55
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{w}^1\Delta_u)$	8.89
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{C}^3\Pi_u)$	11.03
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{E}^3\Sigma_g^+)$	11.88
	$e + \text{N}_2 \rightarrow e + \text{N}_2 (\text{a}''^1\Sigma_g^+)$	12.25
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (\text{a}^1\Delta_g)$	0.977
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (\text{b}^1\Sigma_g^+)$	1.627
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (\text{c}^1\Sigma_u^-)$	4.50
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (^3\text{P}) + \text{O} (^3\text{P})$	6.00
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (^3\text{P}) + \text{O} (^1\text{D})$	8.40
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (^1\text{D}) + \text{O} (^1\text{D})$	10.00
	$e + \text{O}_2 \rightarrow e + \text{O}_2 (^3\text{P}) + \text{O} (^3\text{S}^0)$	14.70

Table 3.1: O_2 and N_2 collisions in the lower-ionosphere. Adapted from [Moss et al., 2006].

$$\tau_k = \frac{1}{A_k + \alpha_k^{\text{N}_2} N_{\text{N}_2} + \alpha_k^{\text{O}_2} N_{\text{O}_2}} \quad (3.1)$$

where N_{N_2} and N_{O_2} are the N_2 and O_2 densities, respectively [Sipler and Biondi, 1972]. For example, at an altitude of 85 km, $N_{\text{N}_2} \simeq 10^{20} \text{ m}^{-3}$ and $N_{\text{O}_2} \simeq 10^{19.5} \text{ m}^{-3}$. Hence, the radiative lifetimes at this altitude for the five optical bands considered in Tables 3.2 and 3.3 are each less than 6 μs . Furthermore, Monte Carlo simulations of the electron density enhancement due to the lightning EMP performed by [Glukhov and Inan, 1996] showed that, at 90 km altitude, the heated electron distribution

function returns to near-isotropy and quasi-stationarity within $\sim 2 \mu\text{s}$. Since this relaxation is faster than the timescale of the applied EMP, we can consider steady-state photon production. We can also assume that the electrons and neutrals relax extremely quickly, and that photon production stops, after the EMP has propagated through the lower D-region.

Reaction	$A_k \text{ (s}^{-1}\text{)}$	Band System
$\text{N}_2 (\text{B}^3\Pi_g) \rightarrow \text{N}_2 (\text{A}^3\Sigma_u^+) + h\nu$	1.7×10^5	N_2 1st Positive
$\text{N}_2 (\text{C}^3\Pi_u) \rightarrow \text{N}_2 (\text{B}^3\Pi_g) + h\nu$	2.0×10^7	N_2 2nd Positive
$\text{N}_2 (\text{A}^2\Pi_u) \rightarrow \text{N}_2 (\text{X}^2\Sigma_g^+) + h\nu$	7.0×10^4	N_2^+ Meinel
$\text{N}_2 (\text{B}^2\Sigma_u^+) \rightarrow \text{N}_2 (\text{X}^2\Sigma_g^+) + h\nu$	1.4×10^7	N_2^+ 1st Negative
$\text{O}_2^+ (\text{b}^4\Sigma_g^-) \rightarrow \text{O}_2^+ (\text{a}^4\Pi_u^+) + h\nu$	8.5×10^5	O_2^+ 1st Negative

Table 3.2: Optical excitation rates in the lower ionosphere. Adapted from [Vallance-Jones, 1974, p. 119].

Band System	$\alpha_k^{\text{N}_2} \text{ (m}^3 \text{ s}^{-1}\text{)}$	$\alpha_k^{\text{O}_2} \text{ (m}^3 \text{ s}^{-1}\text{)}$
N_2 1st Positive	10^{-17}	0
N_2 2nd Positive	0	3×10^{-10}
N_2^+ Meinel	5×10^{-10}	0
N_2^+ 1st Negative	2×10^{-10}	0
O_2^+ 1st Negative	4×10^{-10}	0

Table 3.3: Optical quenching rates in the lower ionosphere. Adapted from [Vallance-Jones, 1974, p. 119].

In the steady state, the photon production rate for state k , γ_k , is well modeled by

$$\gamma_k = \nu_k N_e \quad (3.2)$$

where ν_k is the optical excitation rate that depends upon the applied electric field and N_e is electron density. Figure 3.3 shows the density-normalized values of ν_k (again obtained from ELENDIF simulations presented in Moss *et al.* [2006]) for five band systems as a function of reduced electric field. We can clearly see that the N_2 1P band system is the first emission to “turn-on” and that it has the highest emission rate, except at extremely high applied fields where it is slightly surpassed by the N_2

2P band system. Hence, the N_2 1P system is the dominant optical emission in the lower D-region ionosphere and accounts for most of the photons observed as elves.




Figure 3.3: Optical excitation rates in the D-region ionosphere, from the ELEN-DIF Boltzmann equation solver as presented in [Moss et al. \[2006\]](#). The atmospheric breakdown field, E_k , is shown as a vertical dashed line.

The spectra radiated by the N_2 1P, N_2 2P, N_2^+ 1N, and N_2^+ M band systems are shown in Figure 3.4. These spectra are obtained from the SPECAIR optical excitation code [[Laux, 1993, 2003](#)]. We see that the dominant emissions in the D-region, the N_2 1P band system, radiate in the red and near-infrared parts of the spectrum, between 600-1200 nm. This property gives elves their red color and determines our choice of filters for elve observing instruments.




Figure 3.4: Spectra for the major N_2 emission band systems in the D-region. Obtained from SPECAIR software code.

3.2 General Features of Elves

In the previous section, we showed that the lightning EMP heats the D-region electron distribution and that the subsequent electron-neutral collisions produce photon emissions, giving rise to elves. Hence, the EMP propagation and the geometry of its interaction with the ionosphere determine the shape and temporal evolution of elves, giving rise to their interesting observed features.

With some simplifications, we can gain valuable intuition about the temporal expansion properties of elves. As a first approximation, we assume that the EMP has a spherical wavefront that expands at the speed of light. We also approximate the bottom of the D-region to be perfectly flat, and we assume that optical emissions occur where the EMP “intersects” this surface and that emissions cease several tens of microseconds after the local EMP magnitude has decayed significantly. Then, the location of photon emissions, $R(t)$, relative to the point on the ionosphere directly above the CG lightning stroke, follows a simple hyperbolic trajectory

$$R(t) = \sqrt{c^2 t^2 - h^2} \quad (3.3)$$

where c is the speed of light, t is time, and h is the “height” of the ionosphere above the ground. Several example trajectories for different ionospheric heights are shown in Figure 3.5, giving the location of photon emission relative to the elve center as a function of time, referred to as an “emission profile”. Several important features become apparent from this plot. First, we notice that at their onset, elves expand at an infinite velocity (the slope of the traces in Figure 3.5 are vertical) and that the expansion rate approaches the speed of light after $\sim 150 \mu\text{s}$. It is important to note that, though the elves have an apparent expansion rate faster than c , nothing is actually propagating that fast and no fundamental laws of physics are violated – the effect is simply due to the geometry of the interaction and is reminiscent of a plane wave incident on a flat boundary appearing to have a phase velocity faster than c .

The apparent faster-than-light expansion of elves causes them to appear distorted when observed, an effect which we refer to as the *photon delay effect*. Essentially, the first photons to arrive at the observer are not those that were emitted first in

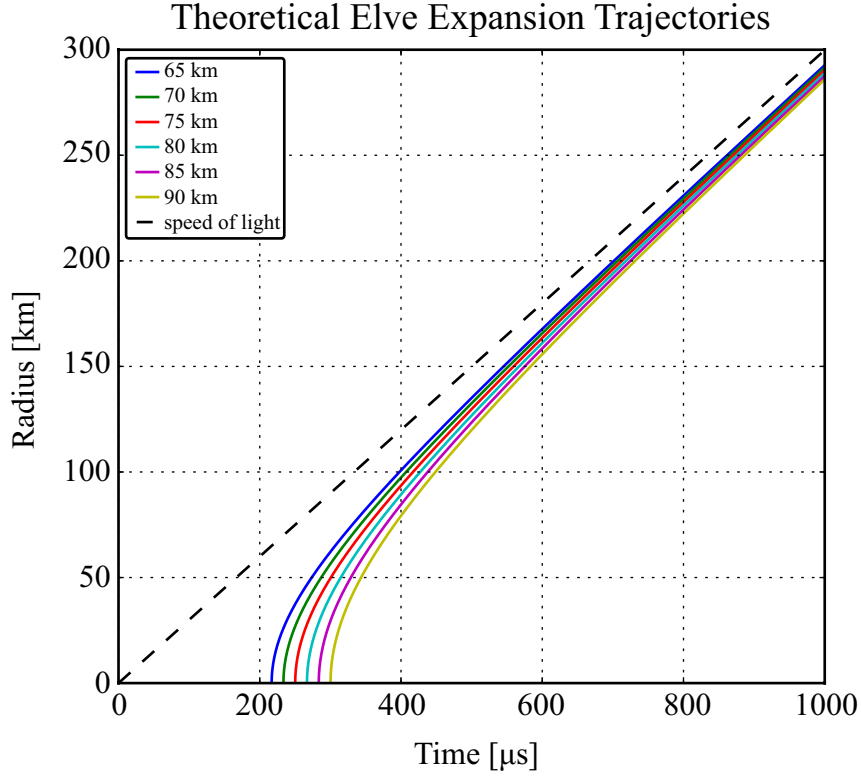


Figure 3.5: The theoretical elve expansion trajectories along the bottom of the D-region for various ionosphere heights. A speed of light expansion rate is shown by the black dashed line.

time, but those emitted at the point along the ionosphere that is halfway between the causative CG and the observer. When observed from the ground, this effect makes the elve appear as a downward descending arc rather than an expanding ring. The photon delay effect is illustrated in Figure 3.6. First consider the points A and B in which photons are first produced at times t_1 and t_2 , respectively. By the time the photon from A propagates from r_1 to r_2 , the elve has expanded and emitted a photon at B. Hence, the observer located a distance r_0 away from the elve center observes the photon B before photon A. Next, consider points C and D. Due to the faster than c elve expansion, the arrival time-difference at the observer for photons C and D is longer than $t_2 - t_1$.

Though the plots in Figure 3.5 show photon emissions at the center of the elve, such is not actually the case – the EMP is not actually spherical. The CG lightning

channel is essentially a small, vertical, dipole antenna. Hence, the return stroke EMP radiation pattern has a null in the vertical direction. This null results in a “hole” in the center of the elve in which no photons are ever produced. This feature was predicted early [Inan *et al.*, 1997] and can be seen clearly in time-integrated elve images captured by video-rate cameras (e.g., Figure 1.6). It turns out that the size of this hole is dependent upon the properties of the lightning return stroke. The use of the observed hole radius to estimate return stroke properties, in particular the return stroke speed, is discussed in detail in Section 4.4.




Figure 3.6: An illustration of the elve photon delay effect.

3.3 Photometric Imaging of Elves

The short timescales and low brightness of elves make them particularly difficult to observe. Intensified video cameras can be used to image elves, but they typically cannot resolve the temporal evolution of the elve expansion – the entire emission is time-integrated into a single frame. This short timescale necessitates the use of sensitive, high-speed photometers which provide much better temporal resolution at the cost of poor spatial resolution. The first ground-based observations of elves were made using three photomultiplier tubes (PMT) [Fukunishi *et al.*, 1996]. Subsequent experiments [Barrington-Leigh, 2000; Inan *et al.*, 1997] verified the radial expansion of

elves using an instrument called the “Fly’s Eye” consisting of 13 PMTs. Photometric imaging of elves has also been conducted from space. The ISUAL instrument on-board the FORMOSAT-2 satellite conducted a three-year observational campaign using an array of PMT sensors in which thousands of elves were observed. In this thesis, we extensively use elve observations collected using the Photometric Imaging of Precipitation of Electron Radiation (PIPER) instrument [Marshall *et al.*, 2008]. The PIPER photometer was originally designed at Stanford in an attempt to observe the optical signatures of lightning-induced electron precipitation (LEP) events, however it was found to be particularly useful for observing elves [Newsome, 2010; Newsome and Inan, 2010]. Due to the high data rates, many previous instruments operated in modes based on triggered recordings of lightning, introducing a possible sampling bias. PIPER, on the other hand, is a free-running instrument and does not suffer from such bias. Data is sampled continuously at a high rate and all data is saved. Event detection and labeling is then performed in post-processing.

3.3.1 Instrumentation

A schematic illustration of the PIPER instrument is shown in Figure 3.7. The optical front-end consist of a clear acrylic window, followed by filters, and camera lenses. The filters used are red, 650 nm long-pass filters which pass most of N₂ 1P band system photons, the dominant photon emissions in elves. These filters are followed by 50 mm f/1.4 Nikon camera lenses. The original PIPER described in Marshall *et al.* [2008] contained four photometers: two with red filters and two with blue filters. However, the blue photometers are omitted from the redesigned PIPER used in this work as they are not necessary for studying the occurrence rates and geometric features of elves. This redesign results in a more compact instrument.

The sensors used in PIPER are Hamamatsu R5900U-L16-20 photometers. These are 16-anode linear PMT arrays. The two photometers are rotated perpendicular to each other as shown in Figure 3.8. This arrangement allows us to photometrically “image” TLE events – one PMT array captures the vertical distribution of incident photons while the other captures the horizontal distribution. The outputs of the two

photometer arrays can then be used to reconstruct elve photon emission profiles, as discussed in Section 4.4.1. The PMTs are square with a 16 mm active area, which in combination with the 50 mm lenses, provide a 18° field of view. The PMTs themselves are powered Hamamatsu C4900 power supplies which take a 15 VDC input and output between -100 and -1200 VDC.

Photomultiplier tubes convert incident photons into electrical current. The current outputs from each of the 32 PMT channels are amplified with transimpedance amplifiers and are then passed into low-pass filters with 12 kHz cutoffs. The filter outputs are then sampled at a rate of 25 kHz by a National Instruments NI-DAQ USB-6259 data acquisition card. This sample rate provides a temporal resolution of $40 \mu\text{s}$, enough to observe the expansion of the elve. Data from the DAQ card is then transferred over the USB bus to a desktop computer and saved on a hard disk where it can be accessed remotely over the internet.

In addition to the core photometric sensor, PIPER interfaces with several supporting pieces of hardware. For scientific studies, it is essential that the data is time-stamped to very high accuracy. We often need to correlate events observed by PIPER with other data sources, such as NLDN events and VLF/LF sferics, for example to match an elve observation with its causative lightning. It may also be desirable to correlate elves observed by multiple PIPER stations, for example to perform triangulation. To address these needs, PIPER uses a GPS antenna to accurately trigger the start of sampling and to timestamp the recorded data. The specific receiver we use is a Motorola M12M Oncore which guarantees 20 ns absolute timing on its one-pulse-per-second (1PPS) clock edge. The 1PPS signal is tied to one of the digital inputs of the NI-DAQ card and is used as the start of sampling trigger. The actual timestamps are sent over USB to the host computer where they are used to tag data files.

The PMTs used in PIPER are extremely sensitive and may actually be damaged if exposed to too much light, for example sunlight. Hence, it is desirable to remove power to the PMTs during the daytime. This protection is achieved by removing the 15V input to the Hamamatsu high voltage power supplies. In particular, we use a USB controlled Phidget InterfaceKit 0/0/4, which consists of four single pole, double

throw (SPDT) mechanical relays to switch this voltage rail on and off. The relays are connected to the host desktop computer which turns the PMT on at dusk and turns them back off at dawn automatically.

Another aspect that must be considered is the control of the pointing direction of the instrument. The location of thunderstorms changes on a daily and even hourly timescale, so it is very desirable to have the ability to set the PIPER pointing direction to arbitrary locations. This requirement was addressed by bolting PIPER to a Pelco pan-tilt mount which is also controlled by a set of relays connected to the computer. This arrangement enabled us to remotely (from Stanford) log into the computer every night, point PIPER in the direction of any thunderstorms (based on real-time observation of NLDN data), and to systematically move the pointing direction to track the storm motion during the night.

The photometer arrays used by PIPER only provide two 16-channel time-series. It is thus difficult to determine, by looking at PMT data alone, whether there is anything interesting happening in the pointing direction. For example, we must know if there is cloud cover on the horizon or if the moon is in the PIPER field of view. To address this requirement, a Wattec low light CCD camera is mounted on top of PIPER with its pointing direction coaxial with PIPER's. A 30-second exposure is captured once every ten minutes, allowing us to track the local weather conditions and to discard PIPER data when such conditions are poor. The CCD camera also occasionally captures sprites. Furthermore, by looking at star-field patterns in the CCD, we are able to determine the exact PIPER pointing direction to better than one degree accuracy. This determination turns out to be essential in the emission profile reconstruction and hole extraction algorithm presented in Section 4.4.1.

3.4 Summer 2013 Observational Campaign

An experimental field campaign was run in the western United States from mid-June until September of 2013 in which we observed 671 elves along with the coincident LF sferics of the causative lightning return strokes. This dataset was used for the bulk of the analysis carried out in this thesis as presented in Chapters 4 and 5.




Figure 3.7: A schematic of the PIPER photometer deployed during the summer 2013 campaign.

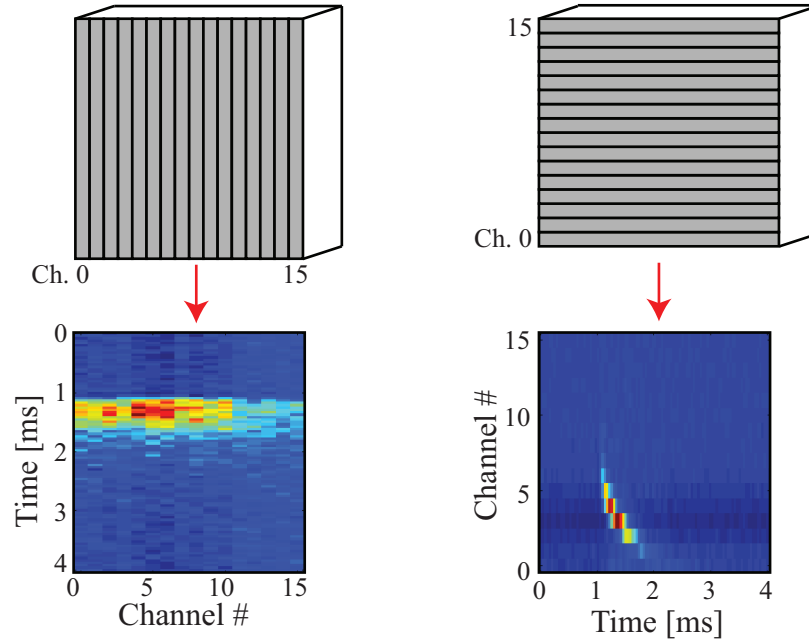


Figure 3.8: Here we show how the two orthogonal photometer orientations capture the horizontal distribution of light (left) and the vertical distribution (right).

3.4.1 Experimental Setup and Observation Techniques

Two PIPER photometers were deployed, one at McDonald Observatory, Texas and one at Langmuir Laboratory, New Mexico. Both were pointed over roughly the same region of northern Texas and western Oklahoma. Both photometers recorded simultaneously and could be pointed independently to track storms. The major advantage of the multi-site photometer setup was that, though both sites pointed over the same region, it dramatically increased the size of our dataset due to the redundancy. The probability that one of the sites would have its viewing obscured due to local weather is quite high, but the probability that both sites have cloud cover on the same night is much lower.

In addition to the photometers, an array of six AWESOME VLF/LF receivers were installed in Oklahoma, Kansas, and Texas, roughly located under the fields-of-view of the two PIPERs. The intended purpose of these receivers is to record the radio emissions from nearby lightning, especially the causative lightning that

produced elves, so as to infer the return stroke properties that are important for elve production.

The geographic layout of the experiment is shown in Figure 3.9. The PIPER photometers are shown as red circles with their fields-of-view in dashed black lines. Since the PIPERs are mounted on remote-controlled pan-tilt mounts, it is possible to rotate these fields-of-view dynamically to track individual storms. A brief description of the two PIPER deployment configurations is provided in the following sections. The locations of the 6 VLF/LF receivers are shown as blue dots in Figure 3.9. The sites were located at a high school in Minneola, KS; a high school in Udall, KS; a bed and breakfast in Terrapin Peaks, OK; a high school in Silo, OK; an elementary school in Norman, OK; and a dude ranch in Clarendon, TX. The locations of all 671 observed elves, inferred from NLDN event locations, are also shown in the figure. The elves are color coded according to the storm in which they were observed.

Conducting a data collection experiment of this scale is very laborious and requires a great deal of coordination. The general steps that were carried out over the course of the summer campaign are outlined below:

1. Each day, just before local dusk, the lightning activity as reported by NLDN was checked. If there was a storm within range of either of the PIPERs (≤ 1000 km), then we logged into the PIPER computer and rotated the instrument to point in the general direction of the storm. We also logged into the VLF/LF field sites to ensure that they were running and recording data.
2. Halfway through the night, we would log back into the PIPER computers and re-point the instruments to account for storm motion during the night.
3. The following morning, after the PIPERs stopped recording, a script would automatically scan the data recorded the prior night for “events” of interest. A simple threshold detector was used to find peaks in any of the 32 channels, where the threshold was usually set to ~ 2 standard deviations. The events were then transferred over the internet to a computer at Stanford.
4. The raw broadband data from the VLF/LF sites would also be transferred back

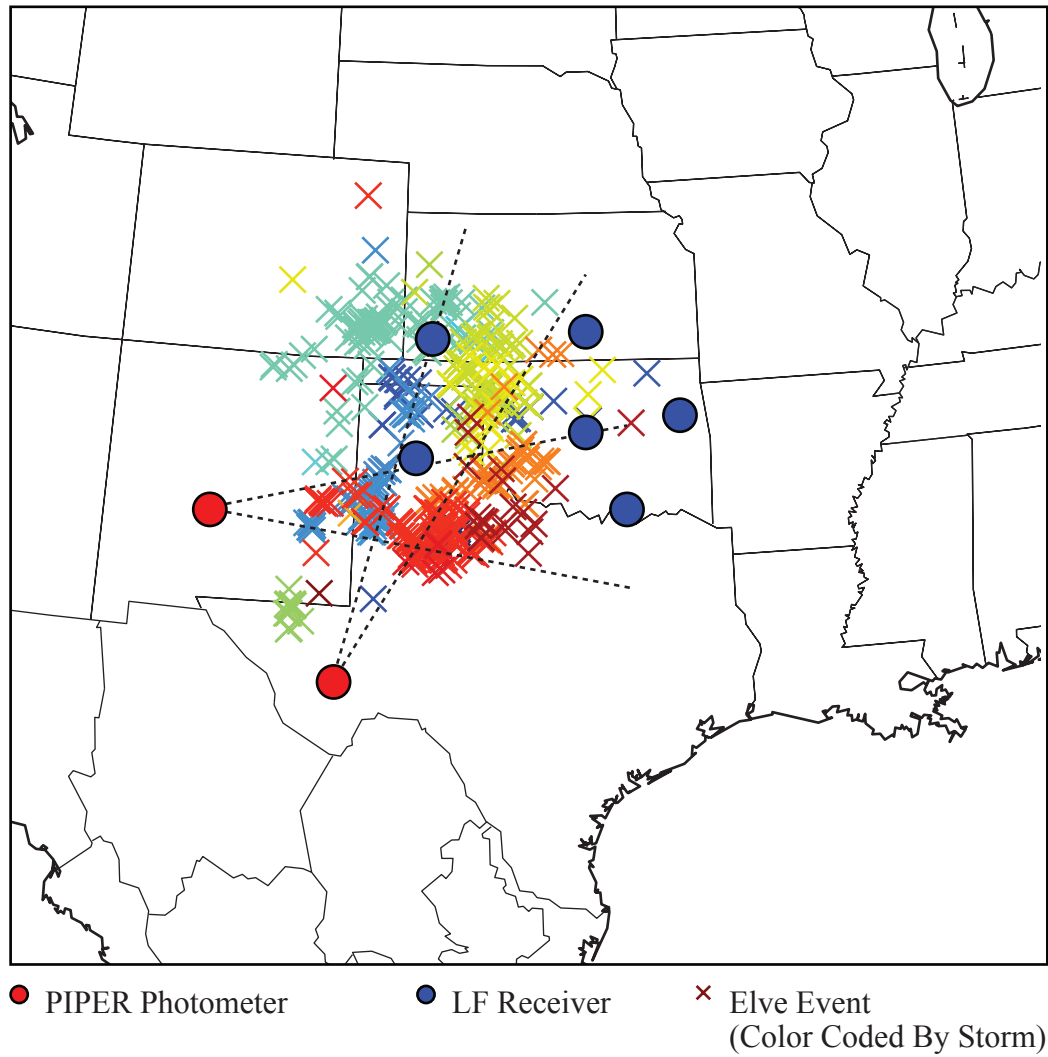


Figure 3.9: Elve occurrence map from Summer 2013 campaign in which 641 elves were observed from sites located at Langmuir Lab, NM and McDonald Observatory, TX.

to Stanford.

5. Each of the PIPER events would be manually labeled as an elve, sprite, halo, doublet, or noise. Note that due to the simplistic threshold detection, the number of false positives, from flashes of light such as airplanes or local lightning, usually far outnumbered the true positives.
6. Each of the PIPER elves would be associated with the temporally closest large NLDN event and matched with the associated lightning sferic at one of the VLF/LF field sites. This information was then inserted into a MongoDB database for easy querying.
7. The CCD images captured would be analysed with an astronomy plate solver [[Lang et al., 2010](#)] to determine the instrument pointing elevation and azimuth. This technique essentially involves matching stars in the image field of view with a database of known stars.

3.4.2 Langmuir Laboratory

Langmuir Laboratory is a site dedicated to atmospheric and lightning research. It is situated on the summit of South Baldy Peak in the Magdalena Mountains of central New Mexico at an altitude of roughly 3255 m. The laboratory is operated by the New Mexico Institute of Technology, located in nearby Socorro.

During the summer of 2013, a PIPER photometer was installed at Langmuir. The instrument itself was mounted on a wooden platform on the railing of the facility's south annex, providing an unobstructed view to the north and the east. The recording computer was located on a workbench inside the building, with the wiring passed through a cable feed. The high altitude location provided excellent viewing conditions, allowing us to detect many elves distant on the horizon. A photograph of the installed system is shown in Figure [3.10](#).



Figure 3.10: A photograph of the PIPER photometer installed at Langmuir Laboratory, New Mexico.

3.4.3 McDonald Observatory

McDonald Observatory is primarily an astronomical observatory, operated by the University of Texas at Austin, located approximately 280 km to the southeast of El Paso, TX. The observatory is situated at a roughly 2070 m altitude (the highest point in Texas, in fact).

A PIPER was installed at McDonald Observatory in late June 2013 and operated the entire summer. It was installed on the northern side of catwalk of the Otto Struve 82-inch telescope. The recording computer was on the floor of a closet inside the telescope dome, and the wiring were fed through an air-duct. This location provided an excellent, unobstructed view to the north into North Texas and Western Oklahoma, though local cloud cover was often a problem.

A photograph of the installed system is shown in Figure 3.11. Note that the hills in the background are quite shallow, rising only a couple degrees off the horizon, and hence do not obstruct any elves in the PIPER field of view.



Figure 3.11: A photograph of the PIPER photometer installed at McDonald Observatory, Texas.

3.4.4 Example Observations

As described in Section 3.3.1, the PIPER photometer contains two 16-anode linear PMT arrays. The two arrays are oriented 90° with respect to each other, allowing the recording of the spatial distribution of incident light in both the vertical and horizontal direction. Each photometer channel integrates the incident photons in its strip of the field of view, and outputs a time-series waveform. Each PMT waveform can then be stacked together, essentially providing a photometric “image”. Several example PIPER observations which were recorded during the 2013 summer campaign are shown in Figure 3.12.

Each row of panels in Figure 3.12 displays a single PIPER event, with the horizontal photometer recording on the left and the vertical photometer on the right. Panels a) and b) show a single elve. Due to the fact that elves are expanding rings of light, elves have a very distinctive shape in the PIPER vertical photometer, appearing as downward descending arc. This feature makes elve emissions very obvious in the PIPER dataset, and in fact enables the manual discrimination of elve and non-elve emissions when tagging events by eye. Elves appear in the horizontal PIPER photometer as a relatively flat line with slight downward concavity, though sometimes the elve hole can be seen in bright, nearby emissions.

Panels c) and d) show a PIPER observation containing an elve, halo, and sprite. The elve appears first, as evidenced by the distinctive elve shape in the vertical photometer. It is followed roughly one millisecond later by a halo. The halo appears as a diffuse glow in several of the PIPER channels and lasts for about one millisecond. Finally, a sprite event immediately follows the halo. The sprite appears as a long enduring light in a single horizontal/vertical PIPER channel. Cases such as this, with three different TLEs, are relatively common in the PIPER dataset, especially following large positive CG lightning. The return stroke that produced these events was a 136 kA +CG stroke.

Another class of event that is found very frequently in the PIPER dataset is shown in panels e) and f). Examining the vertical photometer, we see a bright vertical bar followed immediately by the distinctive elve emission. The vertical bar is brighter in the lower channels (closer to the horizon). This emission preceding the elve is in fact light produced by the lightning stroke itself which is reflected by clouds. These cloud reflections which precede elves are especially common for nearby lightning where the cloud tops are still visible on the horizon, and they render the reconstruction of the elve emission profiles infeasible (discussed in Section 4.4.1).

Finally, panels g) and h) show what is known as an “elve doublet” event. Doublets are two elves immediately following each other, separated by roughly 100 μ s, and are associated with a single lightning stroke. The causative mechanism behind doublets is not well understood, but it has been hypothesized that they result from compact intracloud discharges (CID) and their subsequent ground reflections [[Newsome and Inan, 2010](#)]. Doublets are relatively rare, comprising about 2% of the PIPER dataset.

3.4.5 Occurrence Maps

Figure 3.9 shows the occurrence map for the Summer 2013 campaign. Each cross is the location of an elve observed by PIPER, as inferred from the NLDN reported location for the causative CG stroke. Each elve observed on the same night is displayed as the same color, so we can clearly see how they clustered into storms.

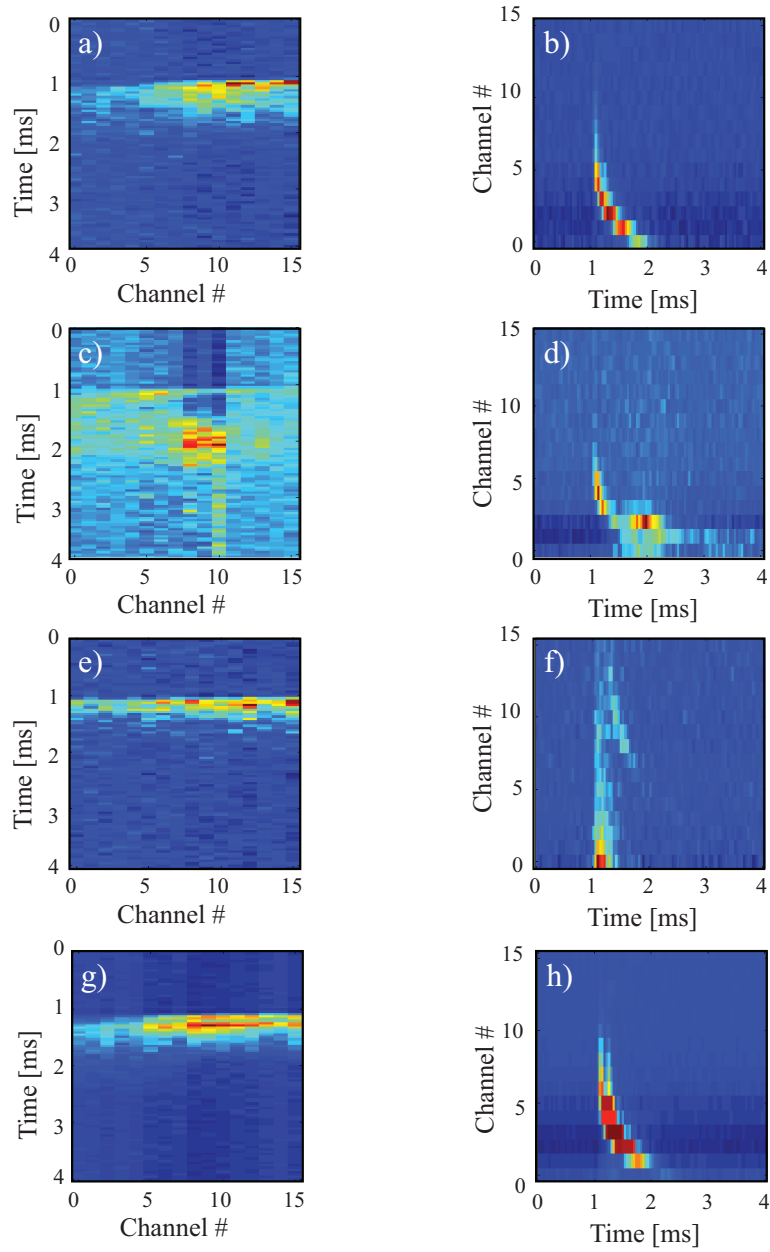


Figure 3.12: Example elves observed by PIPER in 2013. Panels a) and b) show a single elve. c) and d) show an elve followed by a halo and sprite. e) and f) show an elve preceded by a cloud flash. g) and h) show an elve doublet event.

We note that the majority of the events occurred above North Texas and Oklahoma, close enough to the VLF/LF receivers that sferics produced by the causative CGs should contain a good deal of information in the LF portion of the spectrum.

3.5 Summary

In this chapter, we showed that elves are a signature of the interaction between the lightning return stroke radiated EMP and the D-region ionosphere. In particular, the EMP heats the electron distribution in the D-region. When these higher energy electrons collide with neutral molecules, the neutrals may be excited into a higher energy state. As these neutrals relax, they may emit a photon. The dominant photon emissions in elves are from the first positive band system of molecular nitrogen (N_2 1P). Next, we described some general features of elves: they are red in color, appear as expanding rings of light due to the EMP radiation pattern, are significantly distorted due to the photon delay effect, and exhibit a hole in the center.

We also described the imaging of elves photometrically using the PIPER photometer. Finally, we described the experimental field campaign conducted during the summer months of 2013 in which 671 elves were observed over North Texas and Oklahoma along with their coincident VLF/LF sferics.

Chapter 4

Properties of Elve Producing Lightning

In this chapter, we use the dataset of elve observations from PIPER in conjunction with VLF/LF sferics from the Oklahoma LF array to address several important questions about elve-producing lightning. We begin by observing the qualitative features of sferics from elve-producing lightning. We then use the features of the LF sferics to determine the important factors that discriminate between elve-producing and non-elve-producing lightning. Lastly, we use the temporal-spatial features of elve emissions, along with a physics-based numerical model, to determine a likely distribution of current return stroke speeds for elve-producing lightning.

4.1 Dataset

Before sferics from elve causative lightning strokes can be compared with those from non-elve causative strokes, a reliable method of labeling the two classes of sferics must be determined. For the analysis presented in this chapter, each PIPER elve observation is matched with the causative NLDN CG event based on its timestamp and whether or not the discharge occurred within the PIPER photometer field of view. The sferic from the elve causative lightning stroke is then extracted from the broadband VLF/LF recordings based on the timestamp after accounting for the

propagation delay. Determining that a lightning discharge did not produce an elve, however, is more difficult and requires great care. We first create a list of all large (>50 kA) NLDN CG events that occurred during the same storm as a known elve-causative discharge. Based on the latitude and longitude reported by NLDN, all events that occurred outside of the PIPER field of view are discarded. We then use the PIPER Watec CCD images to ensure that there was no cloud cover during these events (i.e. stars are visible in the images). Finally, for the remaining events, we review the PIPER dataset to ensure that there was not actually an elve at that time. After performing these steps, we can be reasonably confident that we have obtained a set of sferics from large lightning discharges that did not produce elves, because if elves did occur, they would have been visible in the PIPER data.

Once sferics have been classified into elve and non-elve causative events, they still cannot be directly compared because they occur at different distances from the receiver. As a sferic propagates a distance r along the Earth's surface, its intensity is attenuated by a factor of $\sim 1/r$. Furthermore, dispersion distorts the sferics due to the finite conductivity of the ground. Due to these facts, the shape of a sferic waveform changes significantly as it propagates, and we cannot reliably compare the features of sferics from two different distances. Hence, we bin sferics based on their distance from the VLF/LF receiver before computing average waveforms and extracting features.

4.2 General Features of Sferics from Elve Producing Lightning

The first sferic feature that we examine is the arrival time of the first ionospheric reflection, also called the first sky-wave. Some authors have reported variations in the nighttime ionospheric density above thunderstorms, with the hypothesis that the height of the D-region is measurably lowered over the course of a storm [[Shao et al., 2012](#); [Lay et al., 2014](#)]. When such lowering occurs, the VLF reflection height is lower in altitude, so the first sky-wave should arrive at the receiver earlier in time due to the shorter propagation path. The question naturally arises as to whether the ionospheric

height may play a role in elve production – perhaps a lower ionosphere makes it more difficult to produce an elve since the neutral density is higher and thus the threshold for optical excitation (E/N) is lower. In Figure 4.1, we show average distanced-binned sferics for the two classes of lightning, with the ± 1 standard deviations shown as shaded regions. A vertical, black dashed line marks the point at which we would expect the sky-wave to start if the ionosphere was a perfectly conducting flat plane at 88 km altitude. The first sky-wave is the negative peak following the ground-wave. Several features stand out in Figure 4.1. The arrival time of the first sky-wave is nearly the same for elve and non-elve causative sferics across all distances, implying that the height of the ionosphere probably does not play a role in elve production. Also, the sky-wave peak lines up nicely with its expected location for a 88 km reflection height.

Next, we turn our attention to the portion of the EMP which propagates along the line of sight from the lightning channel to the receiver on the ground, known as the sferic ground wave. As it does not interact with the ionospheric plasma, this wave is directly related to the return stroke current propagation in the lightning channel. Assuming a flat earth and ignoring ground conductivity, the ground wave magnetic field is given by

$$\mathbf{B}(r, t) = \frac{\mu_0}{2\pi cr} \frac{\partial}{\partial t} \int_0^{\ell_{\text{chan}}} I\left(z', t - \frac{r}{c}\right) dz' \hat{\phi} \quad (4.1)$$

In Figure 4.2, we compare the distance-binned groundwaves from elve and non-elve causative lightning. The waveforms are normalized by dividing by the peak of the groundwave. The most immediately obvious feature that differentiates the two classes of lightning is in the portion of the waveform that precedes the groundwave peak – across all of the distance bins, sferics from elve causative lightning exhibit a slightly longer rise-time. We compute the 10–90% rise time for each sferic in the dataset and plot it as a function of distance from the receiver in Figure 4.3. To test whether the difference in rise-times is due to them being drawn from different distributions or just random coincidence, a two sample Kolmogorov-Smirnov test [Drijard *et al.*, 1971, pp. 269-271] was performed and found $p \leq 0.05$ in all of the distance bins.

In Figure 4.4, we bin the groundwaves by their NLDN reported peak current

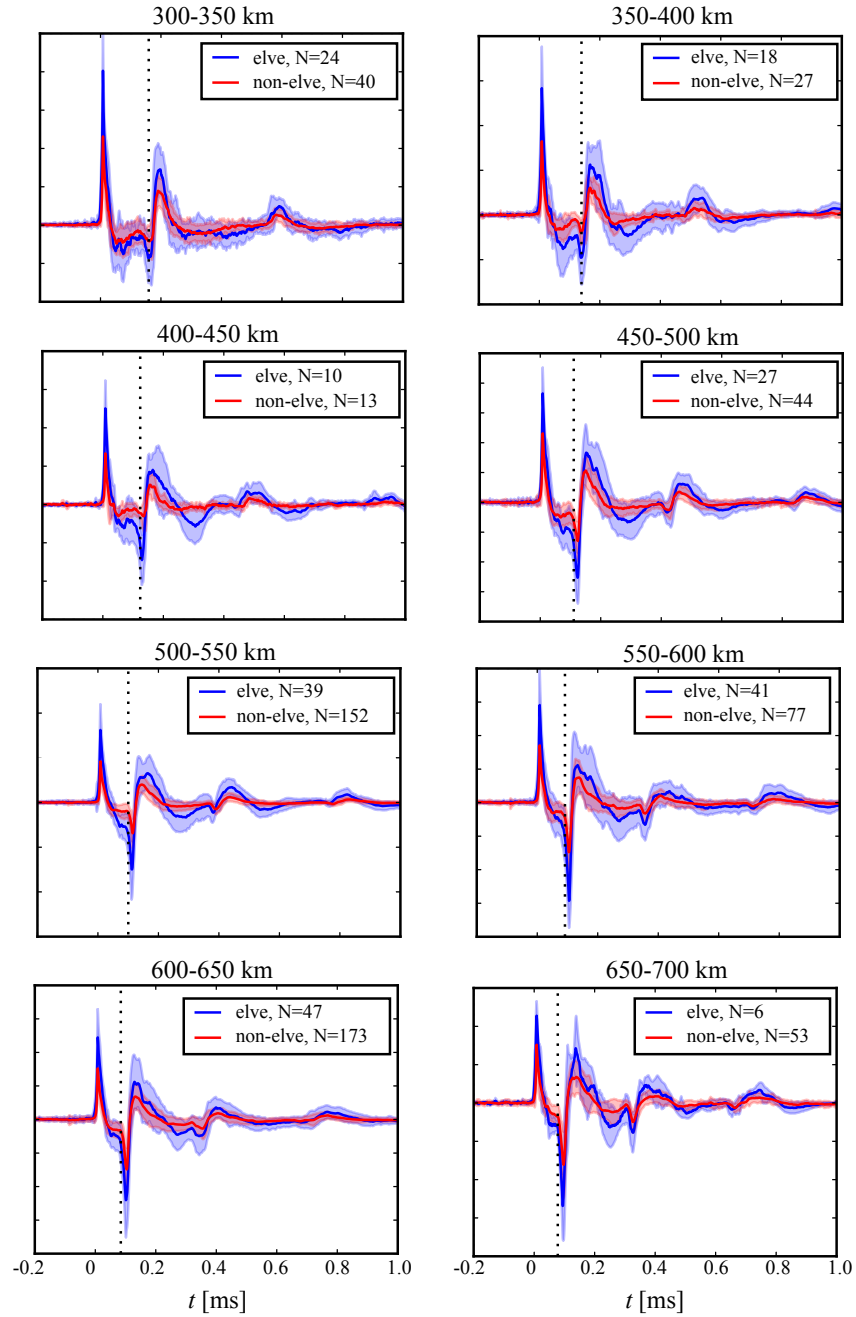


Figure 4.1: Distance binned sferics from elve and non-elve producing lightning. We don't see much variation in the sky-wave arrival times, suggesting that the ionosphere isn't lower when elves are produced. In fact, the sky-wave zero crossing occurs for both classes precisely where it should for a 88 km ionosphere height.

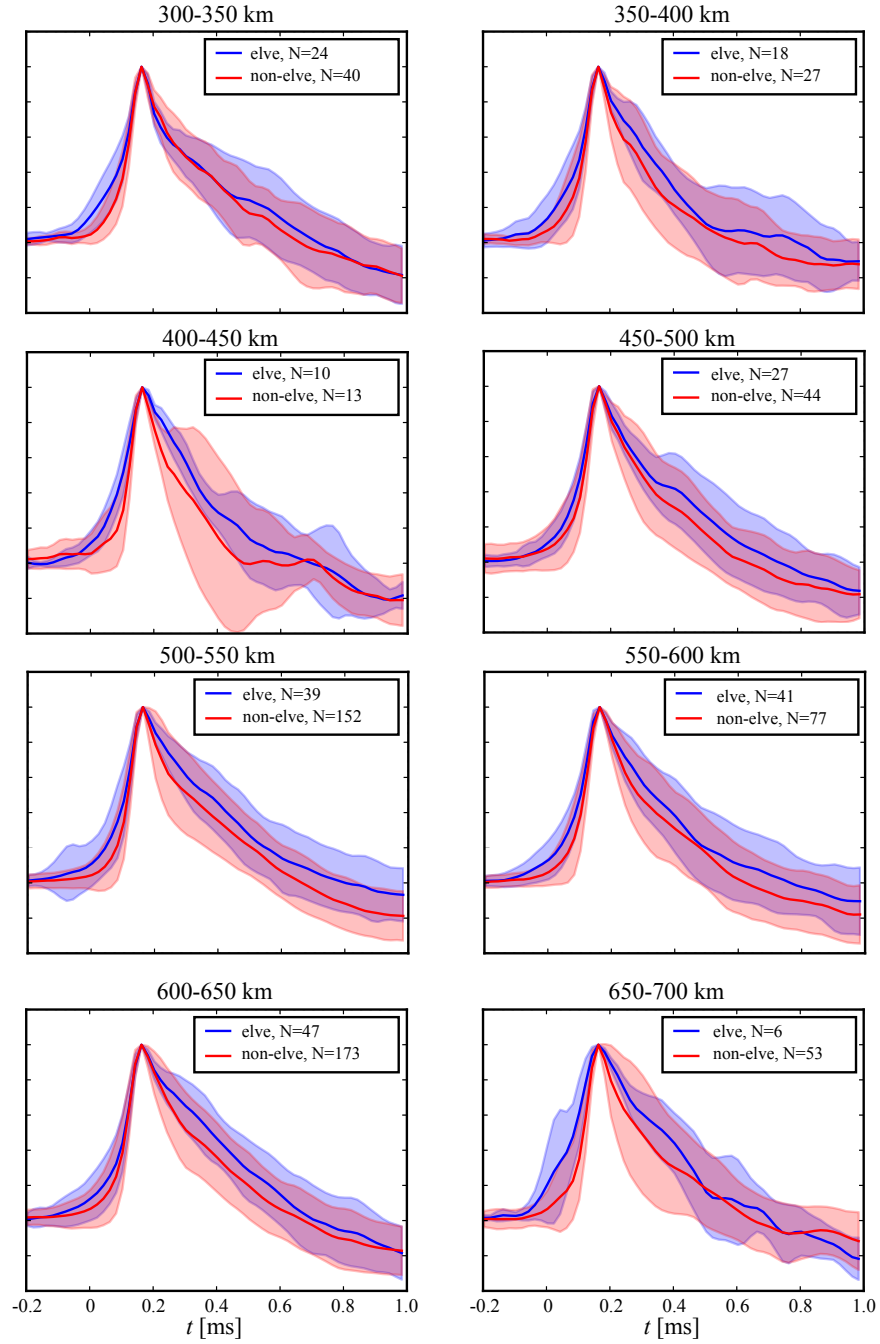


Figure 4.2: Distance binned sferic ground waves from elve and non-elve producing lightning.

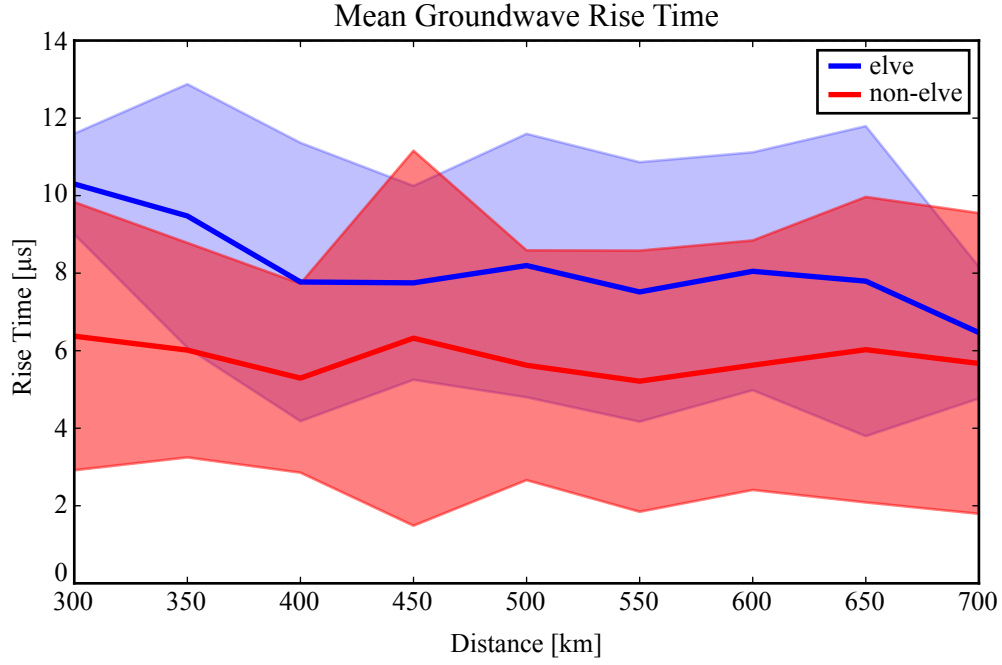


Figure 4.3: Ground wave 10-90% rise times as a function of distance. Elve causative lightning produces sferics with statistically significant longer rise times.

rather than distance from the receiver. Propagation is accounted for by multiplying by r , the distance from the receiver, to effectively cancel out the factor of r^{-1} in Equation 4.1. It should be noted that the r^{-1} variation is only an approximation; dispersion and attenuation due to finite ground conductivity and earth curvature play secondary roles. After multiplication by r , we observe that the difference in ground wave rise-time disappears — the slope and rise-times in Figure 4.4 match almost perfectly between the two classes. This matching suggests that rise-time is linearly correlated with the peak of the radiated field, as NLDN estimates peak current using the range normalized magnetic field with the assumption of a transmission line return stroke current propagation model. This result also hints that the main differentiating factor between elve and non-elve causative strokes is that elve producing causative strokes simply have a larger peak current. In the following section, we develop a more sophisticated statistical model to estimate the probability that an elve was produced based on multiple features of the observed sferic.

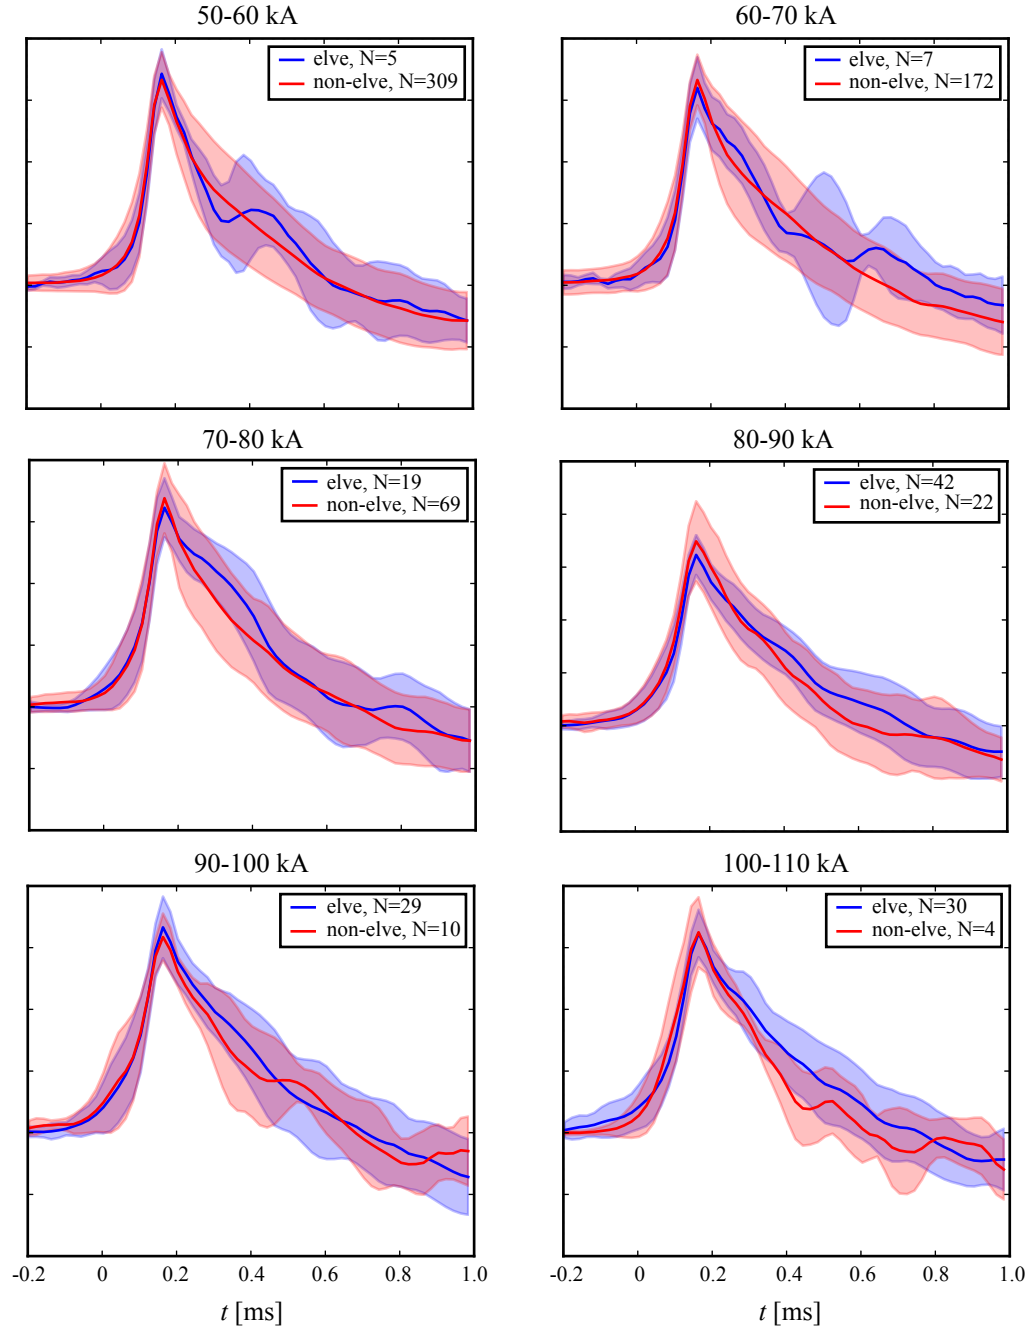


Figure 4.4: Range normalized ground waves binned by NLDN peak current. We find that within each peak current bin, the two classes of lightning have approximately the same peaks and rise times.

4.3 Elve Production Probability Modeling

In the previous section, we compared sferics from large lightning discharges, some of which produced elves and some of which did not. Qualitative differences are evident between the two classes, namely that at a given distance from the lightning stroke, ground waves from elve causative lightning has a significantly longer rise-time. However, when we control for NLDN peak current, the rise time difference for range-normalized ground waves disappears, suggesting that elves have a tendency to be produced by larger peak current return strokes.

In this section, we develop a more rigorous statistical model for analyzing the differences between elve and non-elve causative lightning. In particular, we train a logistic regression classifier model to estimate the probability that an elve was produced for a given observed sferic waveform. We begin by discussing the features used, then provide a derivation of the logistic regression model, and finally discuss the results of the model.

4.3.1 Features of the Sferic Ground Wave

Five hand-picked waveform features are used as features in the logistic regression. These features are illustrated with a typical LF sferic in Figure 4.5.

First, we use the range normalized groundwave peak, also referred to as the range normalized signal strength (RNSS). This feature is measured by first multiplying the waveform by the range from the receiver to the stroke to account for the $1/r$ propagation decay and then searching for the peak. A parabola is fit to the three largest points and the maximum computed.

The second feature used is the 10–90% rise time, τ_r . This risetime is measured by first detecting the ground wave peak, and then computing the time it takes for the magnetic field to change from 10% to 90% of that value in the preceding time segment. The 10–90% fall time is computed in a similar manner on the time segment immediately following the ground wave peak.

The full-width half-maximum (FWHM) is determined by computing the time it takes for the waveform to go from 50% of the peak value, up to the peak, and then

back down to 50%. This determination gives a sense of the “width” of the ground wave.

The ground wave to sky wave ratio (GW/SW) is computed by dividing the range normalized ground wave peak by the value of the first skywave peak. The skywave peak is found by searching for the smallest value in a time window determined by the range from the stroke and an assumed ionosphere height. The motivation for using this feature is to capture the radiation pattern of the return stroke EMP, i.e., by determining the ratio of upward propagating to laterally propagating energy.

Finally, we augment these features extracted from the sferic waveform itself with the peak current estimate provided by NLDN.




Figure 4.5: Example ground wave with the features labeled.

In Figure 4.6, we show a scatter plot matrix of the features extracted from our dataset for both classes of lightning. The off-diagonals show pairwise scatter plots for each pair of features, while the diagonals show the probability distributions of each individual features. The distributions are estimated using a kernel density estimator (KDE) with a Gaussian kernel [Parzen, 1962; Rosenblatt, 1956]. From this figure, we can gain a sense as to which features are important. For example, RNSS, τ_r , and peak current all have some discriminatory power, while the distributions for τ_f , FWHM, and GW/SW are nearly identical for elve and non-elve causative lightning.

4.3.2 Logistic Regression Model

In order to model the production probability of elves given the observed ground wave features, we use a logistic regression. This logistic regression is a classic binary




Figure 4.6: Feature scatter plot matrix. The diagonal plots show the KDE estimated probability distributions for each feature, while the off-diagonal plots show pairwise scatter plots for each pair of features.

classifier, in fact a type of generalized linear model, which maps our feature vector to a class probability. A brief derivation of the model is given below.

We first define the random variable $y_i \sim \text{Bernoulli}(\phi_i)$ where, for each training event i , $y_i = 1$ if an elve occurred and $y_i = 0$ otherwise. We then define the probability

$$\phi_i = p(y_i = 1 \mid x_i, \theta) = \sigma(\theta^T x_i) \quad (4.2)$$

and conversely

$$p(y_i = 0 \mid x_i, \theta) = 1 - \sigma(\theta^T x_i) \quad (4.3)$$

where $x_i \in \mathbf{R}^n$ is the feature vector, $\theta \in \mathbf{R}^n$ are the model weights, and $\sigma(z) = [1 + \exp(-z)]^{-1}$ is the sigmoid, or logistic, function. We can combine Equations 4.2 and 4.3 into a single equation

$$p(y_i \mid x_i, \theta) = \sigma(\theta^T x_i)^{y_i} [1 - \sigma(\theta^T x_i)]^{1-y_i} \quad (4.4)$$

Next, we define the likelihood of all m training examples

$$p(Y \mid X, \theta) = \prod_{i=0}^m \sigma(\theta^T x_i)^{y_i} [1 - \sigma(\theta^T x_i)]^{1-y_i} \quad (4.5)$$

Taking the negative log-likelihood, we get

$$J(\theta) = \sum_{i=1}^m y_i \log \sigma(\theta^T x_i) + (1 - y_i) \log [1 - \sigma(\theta^T x_i)] \quad (4.6)$$

which is known as the cross-entropy loss, a convex objective function. Taking advantage of the convenient fact that $\partial \sigma(z) / \partial z = \sigma(z) [1 - \sigma(z)]$, the gradient of the loss with respect to the model weights is

$$\nabla_{\theta} J(\theta) = \sum_{i=1}^m [\sigma(\theta^T x_i) - y_i] x_i \quad (4.7)$$

We can now use a gradient descent method to minimize the loss and find the optimal model weights. For a more detailed treatment of logistic regression, see for example *Bishop [2006, Chapter 4]*.

A simple illustration of how logistic regression works is shown in Figure 4.7.

4.3.3 Results

Our logistic regression model is trained on the sferic features and the classifier performance is evaluated using a k -fold cross-validation procedure with $k = 3$. This training




Figure 4.7: This plot provides a simple illustration of how logistic regression works. Here, we have two classes of 2D points: red and blue. The logistic regression finds a straight line decision boundary, shown in yellow, that splits the two classes. Distance away from the decision boundary is interpreted as confidence in that class.

involves randomly partitioning the dataset into 3 subsets, training the model on two of the subsets, and evaluating the accuracy on the remaining 1/3. This procedure is performed for each permutation of the subsets and the accuracies from each of the three folds is averaged to provide a single estimate of the classifier accuracy. Using this approach provides a more robust estimate for the manner in which the classifier may generalize to new data.

When trained using the five waveform features discussed in Section 4.3.1 in addition to the NLDN peak current, the logistic regression is able to correctly classify an event as being elve causative or non-elve causative with 91.7% accuracy. Remarkably, when trained on a single feature, NLDN peak current alone, the model is able to classify an event with 90.2% accuracy. The reasons for this are apparent after careful examination of Figure 4.6. Fall-time, full-width half-maximum, and ground wave to sky wave ratio are distributed nearly identically for the two classes of lightning and thus do not provide much discriminatory information. All of the other features — RNSS, rise-time, and peak current — are highly correlated. In fact, NLDN uses a variant of the RNSS, albeit using electric field measurements, along with transmission

line return stroke assumptions to calculate the peak current estimate. Also, as was shown in Figure 4.4, higher peak current strokes tend to have a longer ground wave rise-time.

In Figure 4.8, we show the probability that an elve is produced given the peak current, as modeled by the logistic regression. The red circles show the empirical fraction of events that produced elves in 15 kA wide peak current bins. We see that the modeled probability provides a very reasonable fit with the empirical point estimates. This model predicts that CG return strokes with $|I_p| = 88$ kA have a 50% probability of producing elves, while those with $|I_p| = 106$ kA have a 90% probability. These results differ significantly from those of *Newsome* [2010] shown in Figure 1.8, who found a significantly worse fit at larger peak currents and estimated that elves are produced with 50% probability at approximately 160 kA. The reasons for the differing results are unclear, but we suspect that more care was taken in producing our dataset, in particular by excluding events which occur outside of the PIPER field of view and during cloud cover.

We examine a possible distance bias in Figure 4.9 by plotting the peak current of each event as a function of distance from the PIPER photometer. The mean peak current is shown as a solid curve and the $\pm 1\sigma$ intervals are shown as shaded regions. We observe that the elve causative strokes have statistically significant higher peak current across all distances. This observation discounts the alternative hypothesis that PIPER preferentially detected large peak current events far from the photometer, thus biasing the results.

The conclusion that elves can be reliably predicted using only the NLDN peak current is extremely powerful, as it enables us to classify events when LF sferics are not available. In particular, this result can be applied to the peak current estimate from global lightning detection networks, such as GLD360, to predict elve production outside of NLDN's coverage area. Further validation of this result is discussed in Section 5.1.2, where we find good agreement with satellite based elve observations from ISUAL.




Figure 4.8: Elve production probability as a function of NLDN peak current. In blue, we show the analytical probability provided by the logistic regression, while in red we show the empirical fraction of events which produced elves in several peak current bins.

4.4 Return Stroke Speed of CG Lightning Using Elve Hole Radii

The lightning return stroke is characterized by a brief current pulse propagating upward from the ground towards the cloud. As mentioned in Chapter 2, a simplified “engineering model” is typically used to model this propagation. The transmission line class of models are widely used, including the basic transmission line (TL), modified transmission line with linear altitude decay (MTLL), and modified transmission line with exponential decay (MTLE) [Rakov and Uman, 1998]. These models are parameterized by a channel base current pulse profile, a return stroke speed at which the pulse propagates up the channel, and an altitude decay parameter in the case of MTLL and MTLE. The current return stroke speed, which cannot be measured




Figure 4.9: The NLDN peak currents of all elve and non-elve events in our dataset, plotted as a function of distance. We see that elve causative strokes consistently have higher peak currents across all distances, suggesting that there is not a bias resulting in only observing the most intense events at long distances.

directly, has typically been assumed to be the same as the optical wave speed – between $\frac{1}{5}c$ and $\frac{2}{3}c$ [e.g. [Rakov, 2007](#)]. Recent numerical modeling of the return stroke which accounts for the thermodynamics of the channel, however, has predicted differing speeds for the optical and current pulses with the current pulse propagation being significantly faster [[Liang et al., 2014](#)].

In this section, we show that it is possible to estimate the current return stroke speed from the geometric features of elves, the optical ionospheric signatures of the return stroke EMP. Figure 4.10 shows the normalized radiation pattern of simulated lightning EMPs using an MTL return stroke model with a range of return stroke speeds. We see that as the return stroke speed approaches c , the resulting EMP has a more upwardly-directed radiation pattern. Hence, we expect that the size of the “hole” in the center of an elve emission to be correlated with the return stroke speed




Figure 4.10: Normalized EMP radiation patterns simulated using using the MTLL model with a range of return stroke current speeds. We observe more upwardly directed radiation patterns for return strokes with a high current wave speed.

of the causative stroke, as faster return strokes result in a smaller null in the vertical direction. Using elve observations captured with the PIPER high speed photometer [Marshall *et al.*, 2008] in conjunction with an FDTD model of the lightning-ionosphere interaction [Marshall, 2012], we are able to estimate the return stroke speed of elve causative strokes. Our results provide evidence for a return stroke speed $> \frac{2}{3}c$, in line with the results of [Liang *et al.*, 2014]. Indeed, we find that the return stroke speed may take on a distribution of speeds. The contents of the following sections were previously published [Blaes *et al.*, 2014].

4.4.1 Emission Profile Reconstruction

Elves appear as rapidly expanding rings of light in the lower D-region ionosphere (~ 90 km altitude). They have a duration of less than 1—2 milliseconds and have an apparent expansion rate greater than the speed of light. Hence, the first photons observed by a ground-based instrument are actually emitted at a point on the ionosphere halfway between the observer and the causative CG rather than from the

emissions that occurred first in time. We refer to this as the “photon delay effect” and it results in significantly distorted images of elves in the field-of-view of ground-based instruments.

In light of the photon delay effect, it is very difficult to estimate the geometric features of elves directly from PIPER observations. Thus, we transform the PIPER data into a more natural representation which we call the “photon emission profile”. This technique was first developed by [Newsome \[2010\]](#). An example photon emission profile, along with the PIPER elve observation from which it is calculated, is shown in Figure 4.11. The photon emission profile displays the radial location of photon emissions relative to the center of the elve, assumed to be directly above the causative CG, at each point in time. The overlaid white line in Figure 4.11c shows the theoretical hyperbolic trajectory an elve should follow, assuming a spherical EMP expanding at c which is incident upon a flat ionosphere, similar to Figure 3.5 for an ionosphere height of 88 km.

Several assumptions are made during the emission profile reconstruction. We assume that the elve occurs directly above the causative lightning stroke and that the elve expansion is radially symmetric. We also assume that the viewing geometry is known perfectly, including the range to the lightning, the azimuth, and the elevation angle. [Newsome \[2010, pp. 82-87\]](#) characterized the sensitivity of the algorithm to imperfectly known viewing geometries, finding that the hole radius is most sensitive to unknown elevation angle. However, our photometer included a coaxially aligned CCD camera, which allows us to measure the elevation angle accurately based on star-field patterns.

The emission profiles are determined using a constrained least-squares reconstruction by solving the following convex optimization problem:

$$\begin{aligned}
 & \text{minimize} && ||y_s - A_s x||_2 + \lambda ||D_u x||_2 \\
 & \text{subject to} && x \succeq 0 \\
 & && x^i = 0, \quad i \notin I
 \end{aligned} \tag{4.8}$$

The first term of the objective minimizes reconstruction error. Here, x is the emission profile we are solving for and y_s is the observed elve from PIPER. The matrix A_s

takes into account the known viewing geometry (location of CG along with location, pointing azimuth, and pointing elevation of PIPER) and raytraces photons from the ionosphere back onto the PIPER aperture. The second term of the objective is a regularization term which promotes smoothness in the reconstructed emission profile. The matrix D_u is a first-order finite difference operator which penalizes discontinuous jumps in the elve expansion [Newsome, 2010, pp. 141-142]. The regularization parameter λ controls the degree to which we promote smoothness. The constraint $x \succeq 0$ is used because photons are emitted rather than absorbed, and the constraint $x^i = 0$ is used to reject noise outside of the region of interest, defined by the set of pixels I in which elve photons are present.

Once an emission profile is found, it is trivial to extract the elve hole radius. In the emission profile shown in Figure 4.11, the hole radius is simply the vertical offset of the profile from the radius=0 axis. In practice, for both simulated and observed elves, we take the hole radius to be the radius at which the time-integrated emission profile reached 50% of its peak.

4.4.2 Dataset

Our elve observations were collected using the PIPER high speed photometer [Marshall et al., 2008] at Langmuir Laboratory, New Mexico during the summer of 2013. The PIPER instrument contains two 16-anode photomultiplier tube (PMT) arrays oriented orthogonal to one another, allowing the spatial distribution of transient luminous events (TLEs) to be observed. Each PMT is sampled at 25 kHz, fast enough to temporally resolve the brief (~ 1 ms) elve emissions. An example elve observation is shown in panels (a) and (b) of Figure 4.11.

During the 2013 lightning season, we observed hundreds of elves from Langmuir Laboratory. However, we find that the elve hole can only be observed and reliably measured under a narrow range of conditions. If the elve is too close, then light from the causative CG may be in the field of view and obscure the hole. On the other hand, if it is too far away (> 600 km) then the elve will be low on the horizon and atmospheric scattering and absorption become problematic. In some instances,

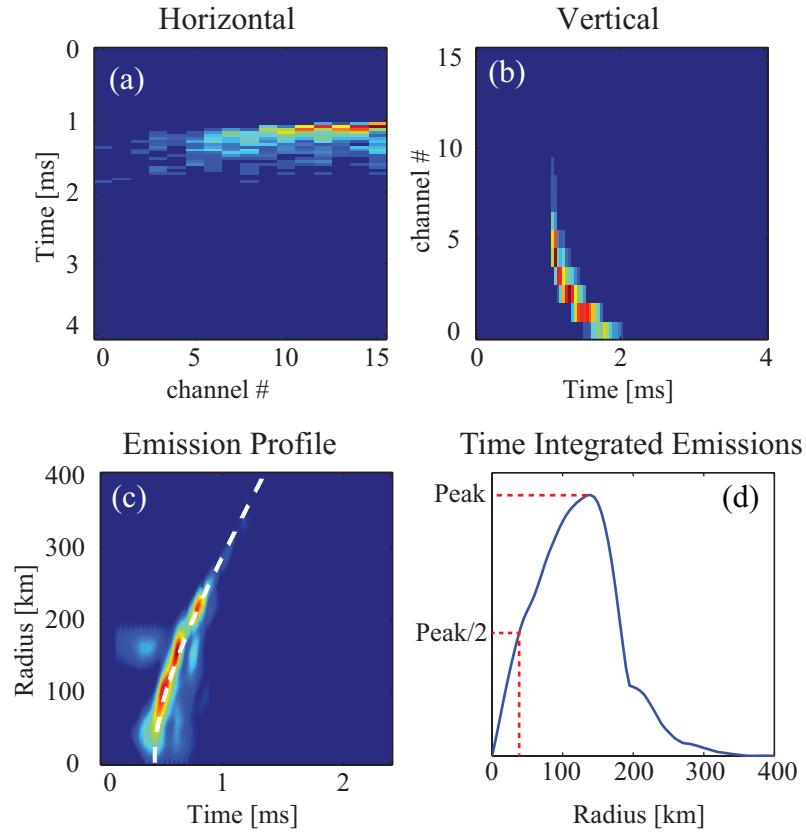


Figure 4.11: Example reconstructed elve emission profile. (a) and (b) show the elve as seen in the horizontal and vertical PIPER photometers respectively. (c) shows the reconstructed emission profile. The overlaid dashed line is the trajectory followed by a spherically expanding EMP incident on a flat ionosphere at 88km altitude. (d) shows the time-integrated emission profile. The elve hole radius is taken to be the radius at which the time-integrated emissions reach $\frac{1}{2}$ of the peak.

especially after large +CG strokes, the elve is followed by sprites and/or halos which obscure the elve hole and make emission profile reconstruction difficult. Finally, since PIPER has a 18° field of view, the center of the elve must be within 9° of the field of view center in order for the hole to be measurable.

Of the hundreds of elves observed in 2013, 55 had the appropriate viewing conditions to make hole radius measurements possible and were not followed by sprites or halos. A histogram of the extracted elve hole radii from these observations is shown in Figure 4.12.




Figure 4.12: A histogram of hole radii extracted from elves observed during the 2013 field campaign at Langmuir Laboratory.

4.4.3 Return Stroke Speed Estimation

We consider three parameters of the lightning return stroke, under a MTLL propagation model, which have a significant effect on the elve hole radius – the return stroke speed, channel length, and the current pulse rise time:

$$\theta = [v_{\text{RS}} \quad \ell_{\text{chan}} \quad \tau_r]^T \quad (4.9)$$

These three parameters are swept in the numerical EMP model over a realistic range of values and the resulting elve hole radii are extracted by computing the radius at which the N_2 1P emissions reach half of their peak brightness. A typical exponential ionosphere from [Wait and Spies, 1964] is used with height $h' = 88\text{km}$ and sharpness $\beta = 0.5$. The results of these simulation runs are shown by the colored curves in Figure 4.13. Each of the curves shows the hole radius as a function of return stroke speed with different choices of channel length and rise time. It is clear from these simulations that return stroke speed is the dominant factor in determining the elve hole radius. The parameter accounting for most of the remaining observed variation is the channel length, which determines the altitude decay rate for MTLL.




Figure 4.13: The simulated elve hole radius as a function of return stroke speed. Other lightning parameters were varied, namely the current wave rise time and the channel length. We see a clear relationship between faster return stroke speeds and smaller elve hole radii. Most of the remaining variation in the hole radius is due to the effective length of the channel, as is shown in the colored line plots.

We perform a quadratic fit on these results to produce an analytical model giving an estimated hole radius for a given set of return stroke parameters, denoted by $r_{\text{model}}(\theta)$.

Assuming Gaussian errors in the radius measurement from data, the likelihood of a given radius measurement is normally distributed according to

$$r_{\text{meas}}|\theta, \sigma \sim \mathcal{N}(r_{\text{meas}} - r_{\text{model}}(\theta), \sigma^2) \quad (4.10)$$

Here, r_{meas} is the measured radius estimate and σ^2 is the variance of the measurement. Now, using this likelihood function in conjunction with “reasonable” prior probabilities for the lightning parameters, we use Bayes’ rule to estimate the joint posterior probability for the three parameters

$$p(\theta, \sigma|r_{\text{meas}}) = \frac{p(r_{\text{meas}}|\theta, \sigma)p(\theta, \sigma)}{\int p(r_{\text{meas}}|\theta, \sigma)p(\theta, \sigma)d\theta d\sigma} \quad (4.11)$$

The marginal distribution for the return stroke speed, given our observed elve hole radii, is then

$$p(v_{\text{RS}}|r_{\text{meas}}) = \iiint p(\theta, \sigma|r_{\text{meas}}) d\ell_{\text{chan}} d\tau_r d\sigma \quad (4.12)$$

Rather than numerically integrate these four-dimensional integrals, we estimate the posterior distribution by performing Markov Chain Monte Carlo (MCMC) sampling [Foreman-Mackey *et al.*, 2013]. This approach uses a large number of Markov chains which take a random walk in model parameter space, providing an estimate of the underlying density.

As little is experimentally known about the actual distributions of the return stroke parameters, we assume uniform prior probabilities over reasonable ranges for the return stroke speed, current rise time, and channel length. This can be thought of as a non-informative choice of priors which does not bias the resulting model posterior probability. We assume that return stroke speed varies between $\frac{1}{3}c$ to c . Current rise-time is assumed to vary between $1\mu\text{s}$ and $7\mu\text{s}$, a reasonable choice for first return strokes [Rakov and Uman, 2003]. Channel length is assumed to vary between 5km and 10km. Finally, σ is assumed to be distributed as a Jeffreys prior, $p(\sigma) \propto 1/\sigma$, a standard non-informative prior for the variance of Normal distributions [Gelman *et al.*, 2013]. While estimating the posterior density through MCMC sampling, we used 200 random walkers, performed 2000 steps, and kept track of the Markov chain locations after a burn-in period of 1000 steps. Our MCMC sampler had an acceptance rate of 0.45 and an autocorrelation time of 61 steps, indicating that it properly converged and was drawing samples from the posterior distribution.

4.4.4 Results and Discussion

Our estimated posterior distribution of elve-causative lightning parameters is shown in Figure 4.14. This “corner plot” displays the parameter marginal distributions on the diagonal and the pairwise joint probabilities on the off-diagonal.

Of particular interest is the panel in the top-left showing the marginal probability distribution of current return stroke speed. The distribution is sharply peaked with a maximum *a posteriori* probability at $0.64c$. This speed is faster than the commonly accepted range of $0.3c$ to $0.5c$ [Rakov and Uman, 1998]. Note that, according to

our model, there is very low probability that an elve is produced by a return stroke with $v_{\text{RS}} < 0.5c$. This makes sense given the histogram of radii in Figure 4.13; a $v_{\text{RS}} < 0.5c$ implies a hole radius greater than 40 km, of which we observe very few. The distribution flattens at about $0.75c$ because the elve hole radius becomes less sensitive to v_{RS} at high speeds. This is apparent in the flattening of the line plots in Figure 4.13 as v_{RS} approaches c . The weak dependence of the elve hole radius upon channel length and current rise time results in flatter posterior distributions with peaks that are not as well-defined.

The uncertainty of the return stroke speed estimation in our model is tied to the width of the marginal posterior distribution. While the distribution has a maximum at $0.64c$, we can say with 95% confidence that the return strokes in our data set had speeds between $0.52 - 0.94c$. Additionally, we observe that the 16—50—84th percentiles of our posterior samples are 0.62 - 0.71 - $0.88c$ respectively. Given that the viewing geometry and distance to the causative stroke used in our hole radius estimation are fairly accurate, we expect that the most significant source of physical uncertainty affecting this return stroke speed estimate is actually the effective lightning channel length (or alternatively the rate at which current decays with altitude), as this has an impact on the current moment and EMP radiation pattern.

The faster return stroke speeds estimated by our model lend credence to the results of [Liang et al., 2014], who predicted a current return stroke speed of $0.84c$, significantly faster than the optical speed. It is important to note that Liang et al. [2014] only simulated return strokes with initial conditions consistent with those of subsequent strokes. All elves we have observed are produced by first return strokes, which may have slower return stroke speeds due to the fact that the initial conditions of the channel include significantly lower temperature and pressure [Rakov and Uman, 2003, pp. 163-164]. We further note that since elves are used to compute these return stroke speeds, these speeds may be particular to elve-producing lightning; lightning that does not produce elves may have a different distribution of return stroke speeds.

4.5 Conclusions

In this chapter, we have presented two studies in which the properties of elve producing lightning are examined. In the first study, PIPER observations are used in conjunction with NLDN events to build up a dataset of known elve and non-elve causative return strokes. Features of the resulting sferic waveforms are compared to determine the salient features that contribute to elve production. A machine learning classifier is trained and is able to discriminate between elve and non-elve causative lightning to a high degree of accuracy. In the second study presented, we examine the geometric features of PIPER elve observations. By comparing these features with numerical simulations of elves, we are able to estimate a distribution of likely return stroke speeds. The key results of these studies include:

1. The NLDN reported peak current is a good predictor for elve production. Strokes with $I_p = 88$ kA have a 50% probability of producing an elve.
2. The elve hole radius is found to have a strong correlation with the return stroke speed of CG lightning. Faster return stroke speeds produce EMPs with a more upwardly directed radiation pattern, resulting in smaller holes in the elve center.
3. A Bayesian inference model with MCMC sampling is used to estimate the probability distribution of return stroke speed, given a dataset of observed elve hole radii.
4. Elve causative lightning was found to have a maximum posterior probability of $v_{RS} = 0.64c$, which is faster than commonly believed and in good agreement with the gas dynamic modeling results of [Liang et al. \[2014\]](#).




Figure 4.14: The posterior distribution of return stroke parameters found through MCMC sampling. This “corner plot” provides a convenient means of visualizing the multi-dimensional parameter space using only 2 dimensions. The diagonal line plots show the marginal distributions for the parameters, while the off-diagonal plots show the joint distribution for each pair of parameters.

Chapter 5

Global Effects of the Lightning EMP

All of the elves and associated lightning radiation observed and utilized in the analysis for this thesis were observed above thunderstorms in the western United States using PIPER photometers and AWESOME VLF/LF receivers in conjunction with NLDN geolocation data. In this chapter, we extrapolate the patterns and relationships learned from our North American dataset to estimate a global production rate for elves. Furthermore, using numerical simulations of the EMP-ionosphere interaction, we provide estimates for the amount of energy deposited into the ionosphere by lightning discharges through collisional heating.

5.1 Global Elve Production Rates

The global production rate of elves, that is the expected geographic distribution of elves and the number produced per unit time, is not well understood. It is of particular importance for determining the amount of energy coupled between the lower troposphere, thunderstorms, into the ionosphere through heating by the return stroke EMP and injection of wave energy into the magnetosphere. In the following sections, we review prior studies on global elve production using satellite observations, and then we introduce our results using ground based measurements.

5.1.1 Prior Satellite Observations

Elves have been observed many times from Earth orbit, including their first observation in 1991 onboard the Space Shuttle [Boeck *et al.*, 1992]. In particular, the Imager of Sprites and Upper Atmospheric Lightning (ISUAL) instrument [Chern *et al.*, 2003], a scientific payload on the Taiwanese FORMOSAT-2 satellite, observed thousands of elves (in addition to sprites and jets) over the course of multiple years. Over a three year period from July 2004 to June 2007, ISUAL observed 5,434 elves, 633 sprites, 657 halos, and 13 gigantic jets, from which global distributions and occurrence rates for the various TLEs were estimated [Chen *et al.*, 2008]. A map of the elve distribution and occurrence rates is shown in Figure 5.1a.

There are, however, a number of issues with the ISUAL dataset which may bias its predicted elve occurrence rate. The FORMOSAT-2 satellite was in a sun-synchronous orbit with an altitude of 890 km, and the ISUAL payload had a relatively narrow field of view with an eastward pointing direction. These limitations can be seen in the coverage and accumulative observation time map shown in Figure 5.1b. Furthermore, ISUAL had to be powered off during its passage through the South Atlantic Anomaly (SAA) to prevent radiation damage to the instruments. These factors introduce a significant observation bias which may affect the global rate estimate.

ISUAL contained three sensors: a CCD camera, a six-channel spectrophotometer, and a 16-anode photometer array (similar to the one used in PIPER) which provided light variation at various altitudes. Due to limited on-board data storage and down-link bandwidth, the instruments could not be operated in a continuous acquisition, free-running manner. Instead, the spectrophotometer was used to trigger events based on emissions from lightning and TLEs. The triggering required an empirically-tuned threshold, so as to correctly save actual TLEs while rejecting false negatives. This triggering threshold inherently introduces a triggering bias into the dataset: there may be many dim TLEs that ISUAL did not register as events, hence lowering the estimated global occurrence rates.

Due to these potential limitations and potential biases, it is desirable to develop a method for predicting global elve production using only continuous ground-based observations. A method for doing this is presented in the following section.

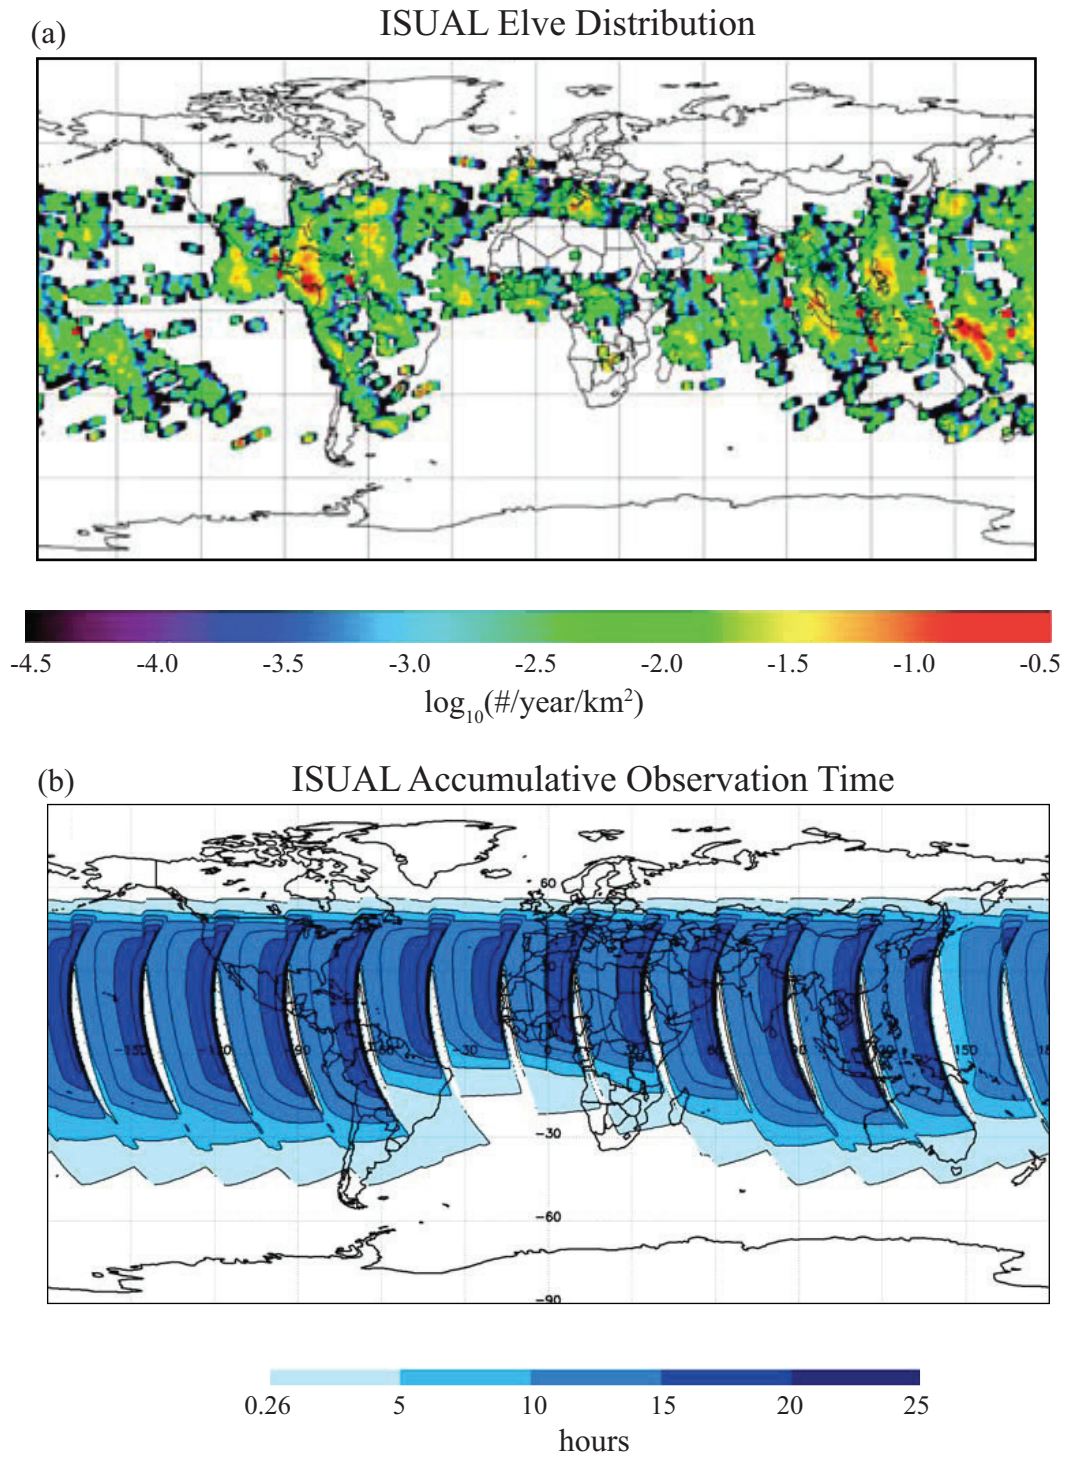


Figure 5.1: (a) The global elve distribution estimated from ISUAL satellite observations. (b) The accumulative observation time of ISUAL. We notice that the polar orbit produces significant gaps in coverage. Reproduced from [Chen et al. \[2008\]](#).

5.1.2 Production Rate Estimated from Ground Observations

In Chapter 4, we compare the properties of elve and non-elve causative lightning and find that NLDN peak current provides an accurate method for estimating the probability that a lightning stroke produced an elve. In particular, we train a logistic model which estimates the probability that an elve was produced given the NLDN reported peak current of the form

$$p(\text{elve} \mid I_p) = \frac{1}{1 + \exp(-\theta_1 |I_p| - \theta_0)} \quad (5.1)$$

where I_p is the peak current and θ_1 and θ_0 are learned model parameters. In this section, we apply this model to the peak current estimates from another lightning geolocation network, GLD360.

GLD360 is a global network of VLF receivers maintained by Vaisala, Inc. which provides accurate global geolocation of CG lightning strokes. The network uses a combination of time of arrival along with correlation with waveform banks of average sferics to perform the geolocation [Said *et al.*, 2010]. In addition to location, GLD360 provides an estimate for the return stroke peak current by calibrating the received VLF peak magnetic fields with NLDN reported peak currents. Using these peak currents in Equation 5.1 enables us to predict whether a stroke produced an elve, even if it occurred outside of the NLDN coverage area, and thus determine the global elve production density.

The procedure we use to estimate the global elve production rate is as follows:

1. Begin with a dataset consisting of one year of GLD360 events.
2. For each event, label it as elve producing with probability $p(\text{elve} \mid I_p)$.
3. Bin each elve causative event into a two dimensional latitude-longitude histogram with $0.5^\circ \times 0.5^\circ$ bin sizes.
4. Normalize each bin by its surface area to obtain units of $\#/\text{km}^2/\text{year}$.

The resulting distribution of elve production rates is shown in Figure 5.2. We

estimate that 4658579 elves were produced in 2013, representing 0.82% of all GLD-recorded lightning events that year. When compared to the empirical elve distribution from ISUAL observations shown in Figure 5.1, we see remarkable similarities. The occurrence rates are roughly on the same order of magnitude — the peak production rates occur at roughly 0.5 per year per km^2 . Additionally, the locations of peak elve production occur in similar locations in Central America, the Caribbean, and Southeast Asia. The fact that our statistical model, trained only on events over North America within NLDN coverage, extrapolates to a global scale and is in such good agreement with satellite observations provides a good deal of confidence about the global validity of the model.




Figure 5.2: The estimated global distribution of elve production estimated using GLD360 peak currents. The probability that a return stroke with a given peak current will produce an elve is estimated using PIPER elve observations and NLDN peak current estimates.

5.2 Heating of the D-Region by Lightning

The upward propagating portion of the lightning EMP interacts with the ionospheric plasma, transferring energy through a collisional heating process, also called Joule heating [Inan *et al.*, 1997]. In this section we perform numerical simulations to estimate the amount of heating due to lightning EMPs and then extrapolate the results to find the global contribution of elve causative lightning to D-region heating.

5.2.1 Numerical Simulations

To investigate the relationship between the return stroke and D-region heating, we use the FDTD EMP model from Marshall [2012] to simulate return strokes with a variety of parameters and to calculate the total heating from each one. The return stroke parameters that were varied include the peak current I_p , the return stroke speed v_{RS} , the channel length ℓ_{chan} , and the current rise time τ_r . These simulations are performed with IRI ionosphere profiles typical for day and night. Additionally, the daytime simulations were run three times with the IRI electron density profiles multiplied by $\frac{1}{3}$, 1, and 3 to account for ionospheric variability.

At each time-step, the model computes the time integrated Joule heating $\mathbf{J} \cdot \mathbf{E}$ in every grid cell. Once the elve has dissipated, we are left with the spatial distribution of the total Joule heating associated with the EMP, $S(r, h)$, which is a function of range r and altitude h . Example heating distributions for both day and night are shown in Figure 5.3. The total energy transferred to the D-region can then be computed by integrating vertically and radially

$$\text{Total Energy} = U = 2\pi \int_{h_{\text{iono}}}^{\infty} \int_0^{r_{\text{max}}} S(r, h) r dr dh \quad (5.2)$$

where h_{iono} is taken to be 60 km during the day and 75 km at night and r_{max} is the maximum radial extent of the elve. We start integrating at these higher altitudes so as to avoid contributions due to the QES field and the lightning channel itself.

Of the return stroke parameters that were varied in these simulations, the peak current and return stroke speed have the greatest impact on the resulting D-region

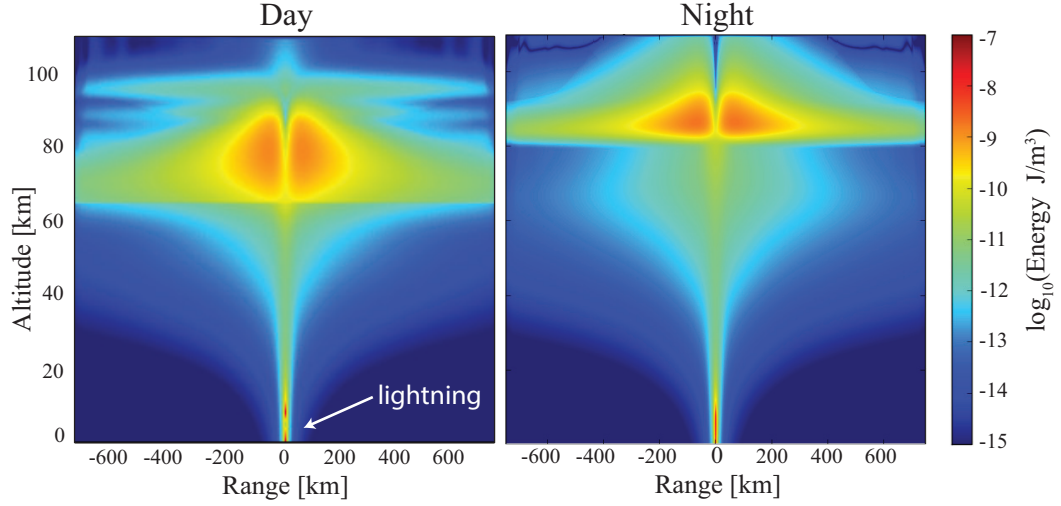


Figure 5.3: The simulated Joule heating of the D-region due to a lightning EMP during the day (left) and at night (right).

heating. The total energy as a function of these two parameters is shown in Figure 5.4. Larger peak currents result in greater heating, as intuitively expected. Return stroke speed plays a significantly smaller role, with faster return stroke speeds resulting in slightly more heating.

Using the result from Section 4.4 in which we found a maximum a posteriori probability estimate for return stroke speed of $0.64c$ for elve causative lightning, we can extract a horizontal slice through the data in Figure 5.4 at $2/3c$ to obtain heating as a function of peak current only. This result is plotted for both day and night in Figure 5.5 along with quadratic best fit lines. The shaded red region shows the range in heating as the daytime electron density is varied from IRI/3 to IRI \times 3, while the solid red line shows the heating for the unmodified IRI profile. Since this ionospheric variation leads to only a small variation in the total heating, using the quadratic fit for the typical IRI daytime profile is sufficient. Using this analytical quadratic fit, we can now estimate the total Joule heating in the D-region due to a CG return stroke using only the return stroke peak current.

The question naturally arises as to what fraction of the total energy dissipated by a lightning stroke gets transferred to the D-region ionosphere through heating. *Quick and Krider [2013]* used calibrated optical measurements of the visible and

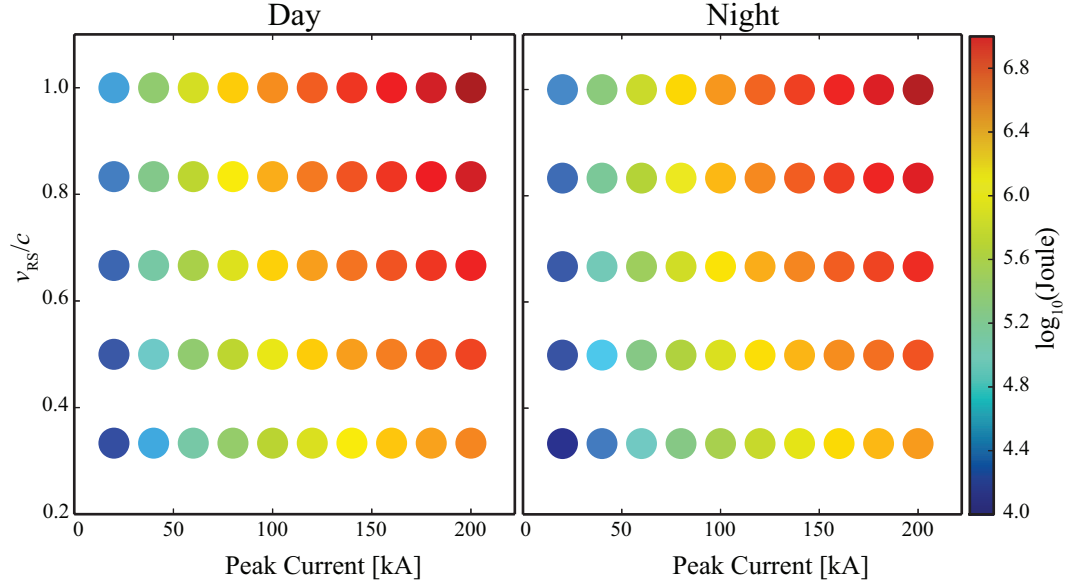


Figure 5.4: Simulated D-region Joule heating as a function of peak current and return stroke speed for day-time and night-time ionospheres.

near-infrared radiation produced by CG lightning to find average optical energies of 3.6 MJ and 1.2 MJ for first strokes and subsequent strokes respectively. [Krider and Guo \[1983\]](#) estimated the peak EMP power by assuming a small vertical dipole channel and then computing the surface integral of the Poynting flux through an enclosing hemisphere, finding

$$P_{\text{peak}} = \int_S (\mathbf{E} \times \mathbf{H}) \cdot \hat{\eta} dS = \frac{4\pi}{3} R^2 \frac{E_p^2}{\mu_0 c} \quad (5.3)$$

where R is range along the ground and $E_p \simeq I_p/3$ is the peak electric field at a range of 100 km. If we approximate the return stroke current as a triangular current pulse, the width of which is the fall-time τ_f , then the total EMP energy is given by

$$U_{\text{EMP}} = \frac{\tau_f P_{\text{peak}}}{2} \quad (5.4)$$

This total EMP energy is computed for a range of peak currents and is plotted as the black line in Figure 5.5. As expected, we find that this curve lies slightly above our D-region heating estimates for both day-time and night-time ionosphere profiles.

The ratio of the heating energy to the total EMP energy is shown in the right panel of Figure 5.5. We find this fraction to be approximately constant across the range of peak currents, with approximately 77% of the EMP energy being dissipated in the ionosphere during the day and 70% at night. This result suggests that the majority of the total lightning-radiated energy is actually transferred to the D-region plasma by the causative lightning stroke.




Figure 5.5: Left: Simulated D-region Joule heating as a function of peak current. We find that this follows a simple quadratic relationship. The points show the simulated heating results for day (red) and night (blue) while the solid lines show the quadratic best fit. The shaded region shows the range in daytime heating as the electron density profile is varied from IRI/3 to IRI×3. Right: The fraction of total EMP energy transferred to the D-region as a function of peak current. We find this to be relatively constant across peak currents with roughly 77% dissipated in the D-region during the day and 70% at night.

5.2.2 Global Heating Estimate

Using an approach similar to the one in Section 5.1.2, we can utilize the peak current estimates from GLD360 to estimate the contribution of lightning EMPs to D-region heating. Using one year’s worth of GLD360 CG events, we label each stroke as occurring in either day or night based on whether the sun was visible at D-region altitudes directly above the lightning. Subsequently, using the peak current and the quadratic fits shown in Figure 5.5, the amount of energy transferred to the D-region

plasma is estimated. These heating results are then integrated into latitude-longitude bins and normalized by area to give units of $\text{J}/\text{km}^2/\text{year}$. In Figure 5.6, we show the geographic distribution of the EMP-driven heating. Notably, the locations of the night-time “hot-spots” are highly correlated with the regions of maximum elve production shown in Figure 5.2, as expected.

Next, we integrate the power density distributions shown in Figure 5.6 over the surface area of the Earth to obtain the total energy deposition due to lightning EMP-driven D-region heating. Doing so, we estimate that CG lightning transfers a total of 61.04 TJ of energy into the D-region ionosphere due to Joule heating each year, with 39.87 TJ transferred during the day and 21.17 TJ at night. This deposition corresponds to an average power of 1.94 MW.

By comparison, this estimate of the energy transfer by the lightning EMP is several orders of magnitude lower than reported auroral energy deposition rates. Observations using the Ultraviolet Imager on-board the Polar spacecraft showed rates of auroral energy deposition by the solar wind as low as 1 GW during the night and varying between 4 and 10 GW during the day, even during times of relatively quiet solar activity ($0 \leq K_p \leq 1$) [Brittnacher *et al.*, 1997].

It should also be noted that the 1.94 MW power estimate only includes the effects of CG lightning, due to the limitations of the two dimensional numerical model used. This model is only capable of simulating vertical lightning channels as a result of the spherical symmetry of the simulation grid. CG lightning has significantly higher peak currents, and thus more intense EMPs, than IC lightning, but only comprises approximately 25% of total lightning activity. Hence, the total lightning EMP driven D-region heating may be higher if the contribution from IC strokes is included.

5.3 Summary

In this chapter, we used statistical models of elve production that were trained using data from North America and extrapolated the results globally. In particular, peak current estimates from the GLD360 geolocation network were used to estimate the

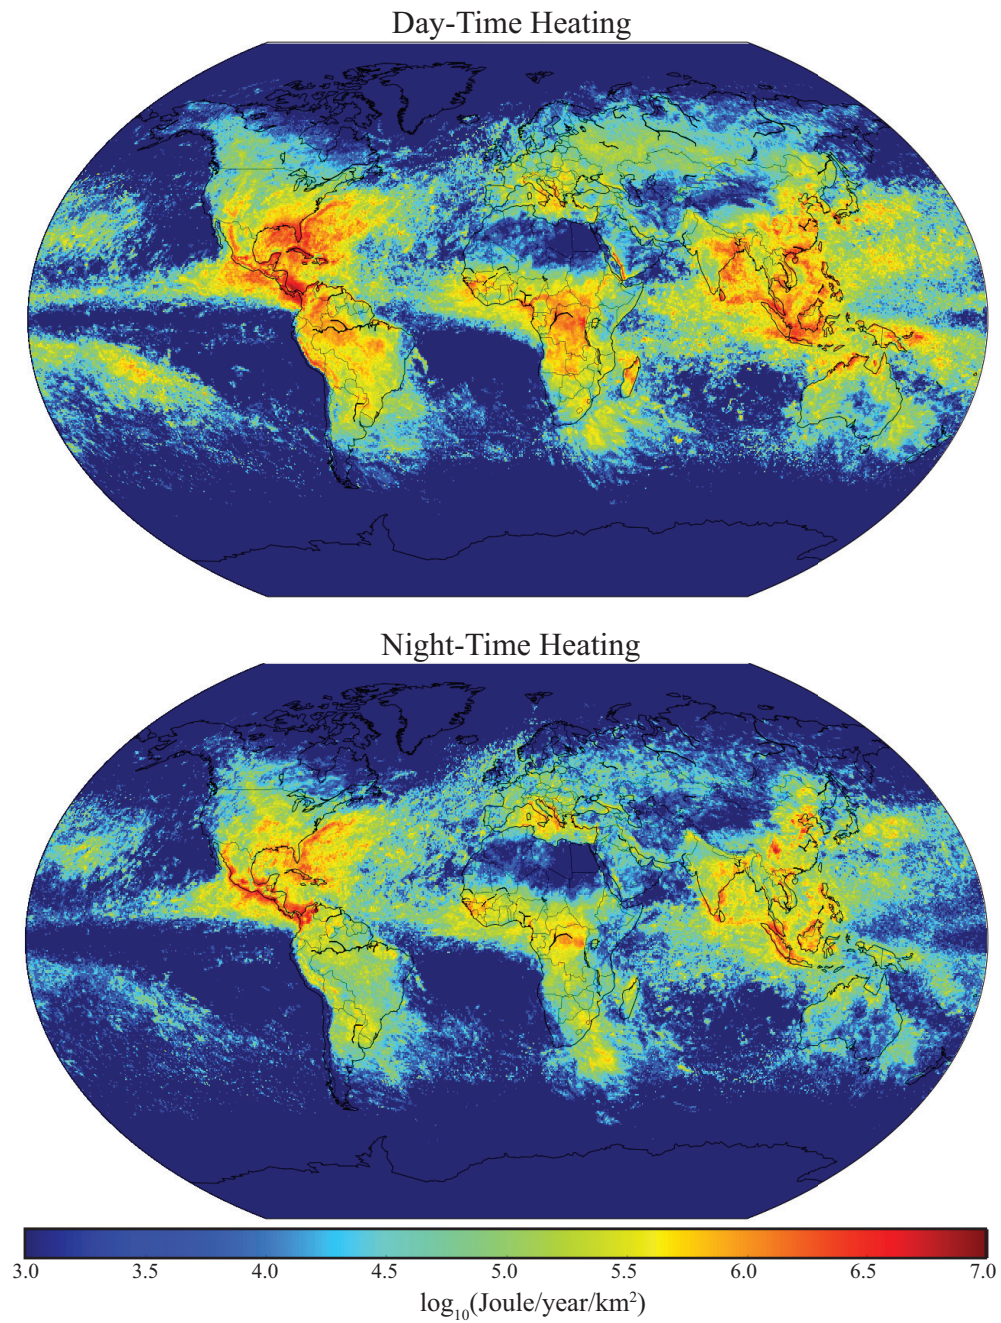


Figure 5.6: Here we show the estimated geographic distribution of one year's worth of D-region heating at local day-time (top) and night-time (bottom).

global elve production rate. Then, using numerical simulations of the lightning-ionosphere interaction, the total contribution of lightning to D-region heating was estimated. The key results of these studies include:

1. A statistical model of elve production was applied to GLD360 data to produce an elve density map which is in good agreement with satellite observations from ISUAL.
2. Numerical simulations of D-region heating due to lightning EMPs was performed. The total heating is found to follow a quadratic relationship with peak current.
3. The analytical heating model was applied to one year's worth of GLD360 events, and found that lightning transfers approximately 61.04 TJ of energy to the D-region every year, which corresponds to an average power of 1.94 MW.

Chapter 6

Conclusions and Suggestions for Future Work

In this work, we present results addressing several outstanding questions pertaining to elves, the lightning return stroke, and the lightning-ionosphere interaction. Using simultaneous optical observations from high speed photometers and VLF/LF magnetic field measurements, the properties of lightning that control elve production were examined and peak current alone was found to explain elve production to a high degree of accuracy. Using the geometric features of elves, namely the radius of the hole in the center, along with numerical simulations of elves, we present a Bayesian model which predicts the probability distribution of the causative return stroke speeds. A maximum probability return stroke speed of $0.64c$ was found. This is the first experimental measurement of the return stroke speed that relates directly to the current propagation itself, rather than the return stroke optical pulse. We extrapolate the rate of elve production to a global scale using GLD360 events, finding good agreement with satellite observations, and use numerical modeling to estimate the contribution of elve causative lightning to D-region heating.

6.1 Summary

In Chapter 1, we provide a introduction to lightning, the ionosphere, TLEs, and a brief overview of prior work in the field. In Chapter 2, we provide a detailed introduction to the lightning return stroke. This includes descriptions of all of the major return stroke engineering models as well as gas dynamic models such as the time-domain fractal lightning (TDFL) model. We provide a derivation that relates the return stroke current propagation to the radiated electromagnetic field, describe ionospheric reflections, and provide a brief discussion of sferic propagation in the Earth-ionosphere waveguide. Finally, we provide information on the VLF/LF reciever instrumentation that is used to observe lightning data.

In Chapter 3, we provide a detailed introduction to the physics of elves. This includes a description of ionospheric and atmospheric chemistry, the various collision processes in the lower ionosphere, and optical excitation rates. Next, we provide a general description of the geometric features of elves and explain how the faster than light “expansion” of elves results in the photon delay effect, distorting them when observed by ground-based instruments. We then give a detailed overview on the design of the PIPER photometer and how it is used to observe elves. Finally, we describe the 2013 summer campaign, the data collection methodology, and show example events and occurrence maps from the dataset.

Chapter 4 uses PIPER and VLF/LF sferic data to examine the detailed properties of elve causative lightning. We first explain how lightning events are labeled as elve causative or non-elve causative. We then compare average waveforms from the two classes of lightning, finding that elve-causative ground waves tend to exhibit a longer 10-90% risetime for a given distance from the receiver. However, when we control for peak current, sthe differences in the risetimes dissappear. Next we describe how features are extracted from the sferics and how a logistic regression model is used to model elve production probability given those features. Doing this, we find that we are able to classify the events in our dataset with 90.2% accuracy using NLDN peak current alone. In the second half of Chapter 4, we observe in simulations that faster return stroke speeds produce elves with smaller hole radii. We then show

that hole radii can be extracted from PIPER elve observations after reconstructing a photon emission profile. Finally, we construct a Bayesian inference model and perform Markov chain Monte Carlo (MCMC) sampling to obtain a likely distribution of return stroke speeds given the empirical hole radii. In particular, we found a MAP return stroke speed estimate of $0.64c$. This is on the faster end of what’s commonly accepted in the literature, but agrees well with recent thermodynamics modeling of the return stroke.

In Chapter 5, we use the logistic regression classifier presented in Chapter 4 to classify a year’s worth of GLD360 events. This allows us to estimate global elve occurrence rates based on our PIPER observations, despite the fact that the PIPER data was collected in North America. Our resulting global occurrence map agrees well with satellite-based measurements from the ISUAL instrument, providing confidence that our elve production probability model using peak current is correct. We find that 0.8% of all lightning located by GLD360 produces elves, amounting to 4.7×10^6 total elves per year. Next, we perform numerical simulations to find the total Joule heating in the D-region due to the lightning EMP as a function of peak current. These results are then applied to the GLD360 dataset to provide a global estimate for D-region heating. We find that lightning EMPs transfer approximately 120 TJ of energy into the D-region through collisional heating.

6.2 Suggestions for Future Work

Finally, we present several suggestions for future work:

1. *Improved PIPER photometer.* A PIPER photometer with a wider field of view, better horizontal resolution, and faster sampling frequency should be developed. In our dataset, many elves were only partially visible because they occurred off-center in the instrument’s field of view. This prevented us from using many of the elves in the hole radius study. Increasing the field of view and improving the spatial resolution, e.g. by using a 32-channel photometer rather than the current 16-channel one, will enable one to capture the full horizontal extent of

the elve emission and more accurately measure the size of the elve hole radius. Also, improving the temporal resolution by increasing the sampling frequency would enable one to better measure the time delay between optical peaks in elve doublet emissions, perhaps allowing one to validate the hypothesis that they are produced by CID lightning.

2. *Better sferic waveform features.* The sferic features used in our analysis in Chapter 4 were easy to compute, but rather crude. It may be worthwhile to investigate the use of unsupervised feature learning and deep learning techniques for analyzing and extracting features from VLF/LF waveforms. Deep convolutional networks have had recent success at automatic speech recognition [[Abdel-hamid et al., 2013](#)] and may also be useful for classifying VLF/LF sferics.
3. *Use of more physical return stroke models.* Most of the analysis in this thesis used return stroke engineering models such as MTLL due to their simplicity and ease of use. However, these engineering models are not very physically realistic. Using a more physical models, such as the TDFL model, may enable one to investigate the impact of the geometry, pressure, and temperature of a channel upon the resulting lightning EMP.

Bibliography

- Abdel-hamid, O., L. Deng, and D. Yu (2013), Exploring Convolutional Neural Network Structures and Optimization Techniques for Speech Recognition, *14th Annual Conference of the International Speech Communication Association*.
- Barrington-Leigh, C. P. (2000), Fast Photometric Imaging of High Altitude Optical Flashes Above Thunderstorms, *Thesis*, pp. 1–119.
- Barrington-Leigh, C. P., and U. S. Inan (1999), Elves triggered by positive and negative lightning discharges Elves Triggered by Negative CG Discharges, *Geophysical Research Letters*, *26*(6), 683–686.
- Barrington-Leigh, C. P., U. S. Inan, and M. Stanley (2001), Identification of sprites and elves with intensified video and broadband array photometry, *Journal of Geophysical Research*, *106*, 1741, doi:10.1029/2000JA000073.
- Bilitza, D., and B. W. Reinisch (2008), International Reference Ionosphere 2007: Improvements and new parameters, *Advances in Space Research*, *42*, 599–609, doi:10.1016/j.asr.2007.07.048.
- Bishop, C. M. (2006), *Pattern recognition and machine learning*, Springer, New York.
- Blaes, P. R., R. A. Marshall, and U. S. Inan (2014), Return stroke speed of cloud-to-ground lightning estimated from elve hole radii, *Geophysical Research Letters*, pp. 9182–9187, doi:10.1002/2014GL062392.
- Boeck, W. L., O. H. Vaughan Jr., R. Blakeslee, B. Vonnegut, and M. Brook (1992),

- Lightning induced brightening in the airglow layer, *Geophys Res Lett*, 19(2), 99–102.
- Brittnacher, M. J., R. Elsen, G. K. Parks, L. Chen, G. a. Germany, and J. F. Spann (1997), A Dayside Auroral Energy Desposition Case Study Using the Polar Ultraviolet Imager, *Geophysical Research Letters*, 24(8), 991–994, doi:10.1029/97GL00257.
- Bruce, C. E. R., and R. H. Golde (1941), The Lightning Discharge, *Journal of the Institute of Electrical Engineers-Part II: Power Engineering*, 88(6), 487–505, doi:10.1259/0007-1285-5-49-85-a.
- Byrne, G. J., A. A. Few, and M. E. Weber (1983), Altitude, Thickness and Charge Concentration of Charged Regions of Four Thunderstorms During Trip 1981 Based Upon In Situ Balloon Electric Field Measurements, *Geophysical Research Letters*, 10(1), 39–42.
- Chen, A. B., C. L. Kuo, Y. J. Lee, and H. T. Su (2008), Global distributions and occurrence rates of transient luminous events, *J Geophysical Res.*
- Chern, J. L., R. R. Hsu, H. T. Su, and S. B. Mende (2003), Global survey of upper atmospheric transient luminous events on the ROCSAT-2 satellite, *Journal of atmospheric . . .*
- Christian, H. J., R. J. Blakeslee, D. J. Boccippio, W. L. Boeck, D. E. Buechler, K. T. Driscoll, S. J. Goodman, J. M. Hall, W. J. Koshak, and D. M. Mach (2003), Global Frequency and Distribution of Lightning as Observed from Space by the Optical Transient Detector, *Journal of Geophysical Research*, 108, doi:10.1029/2002JD002347.
- Cohen, M. B., U. S. Inan, and E. W. Paschal (2010), Sensitive Broadband ELF/VLF Radio Reception With the AWESOME Instrument, *IEEE Transactions on Geoscience and Remote Sensing*, 48(1), 3–17.
- Cummins, K. L., M. J. Murphy, E. a. Bardo, W. L. Hiscox, R. B. Pyle, and A. E. Pifer (1998), A Combined TOA/MDF Technology Upgrade of the U.S. National

- Lightning Detection Network, *Journal of Geophysical Research*, 103(D8), 9035, doi:10.1029/98JD00153.
- Diendorfer, G., and M. A. Uman (1990), An Improved Return Stroke Model with Specified Channel-Base Current, *Journal of Geophysical Research*, 95(D9), 13,621–13,644.
- Drijard, D., W. T. Eadie, F. E. James, M. G. W. Roos, and B. Sadoulet (1971), *Statistical methods in experimental physics*, North Holland, Amsterdam.
- Dubovoy, E. I., M. S. Mikhailov, a. L. Ogonkov, and V. I. Pryazhinsky (1995), Measurement and numerical modeling of radio sounding reflection from a lightning channel, *Journal of Geophysical Research*, 100(D1), 1497, doi:10.1029/94JD00965.
- Foreman-Mackey, D., D. Hogg, D. Lang, and J. Goodman (2013), emcee: The MCMC Hammer, *Publications of the Astronomical Society of the Pacific*, 125, 306–312, doi:10.1086/670067.
- Franz, R. C., R. J. Nemzek, and J. R. Winckler (1990), Television image of a large upward electrical discharge above a thunderstorm system., *Science (New York, N.Y.)*, 249, 48–51, doi:10.1126/science.249.4964.48.
- Fukunishi, H., Y. Takahashi, and M. Kubota (1996), Elves: Lightning-induced transient luminous events in the lower ionosphere, *Geophysical Research Letters*, doi:10.1029/96GL01979.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013), *Bayesian data analysis*, CRC press.
- Gerken, E., U. Inan, and C. Barrington-Leigh (2000), Telescopic imaging of sprites, *Geophys. Res. Lett.*, 27(17), 2637, doi:10.1029/2000GL000035.
- Glukhov, V. S., and U. S. Inan (1996), Particle simulation of the timedependent interaction with the ionosphere of rapidly varying lightning EMP, *Geophysical Research Letters*, 23(16), 2193, doi:10.1029/96GL02185.

- Harriman, S. K., E. W. Paschal, and U. S. Inan (2010), Magnetic Sensor Design for Femtotesla Low-Frequency Signals, *IEEE Transactions on Geoscience and Remote Sensing*, 48(1), 396–402.
- Hedin, a. E. (1991), Extension of the MSIS thermosphere model into the middle and lower atmosphere, *Journal of Geophysical Research*, 96(90), 1159–1172, doi:10.1029/90JA02125.
- Heidler, F. (1985), Traveling current source model for LEMP calculation, in *6th Int. Symp. on Electromagnetic Compatibility, Zurich, Switzerland*, pp. 157–162.
- Holle, R. L. (2008), Annual rates of lightning fatalities by country, *20th International Lightning Detection Conference*, pp. 1–14.
- Hubert, P., and G. Mouget (1981), Return stroke velocity measurements in two triggered lightning flashes, *Journal of Geophysical Research*, 86, 5253, doi:10.1029/JC086iC06p05253.
- Idone, V. P., and R. E. Orville (1982), Lightning return stroke velocities in the thunderstorm research international program (TRIP), *Journal of Geophysical Research*, 87, 4903, doi:10.1029/JC087iC07p04903.
- Idone, V. P., R. E. Orville, P. Hubert, L. Barret, and A. Eybert-Berard (1984), Correlated observations of three triggered lightning flashes, *Journal of Geophysical Research*, 89, 1385–1394, doi:10.1029/JD089iD01p01385.
- Inan, U. S., T. F. Bell, and J. V. Rodriguez (1991), Heating and ionization of the lower ionosphere by lightning, *Geophysical Research Letters*, doi:10.1029/91GL00364.
- Inan, U. S., W. a. Sampson, and Y. N. Taranenko (1996), Spacetime structure of optical flashes and ionization changes produced by lightingEMP, *Geophysical Research Letters*, 23(2), 133, doi:10.1029/95GL03816.
- Inan, U. S., C. P. Barrington-Leigh, S. Hansen, V. S. Glukhov, and T. F. Bell (1997), Rapid lateral expansion of optical luminosity in lightning-induced ionosphere flashes referred to as 'elves', *Geophysical Research Letters*, 24(5), 583–586.

- Inan, Umran S and Golkowski, M. (2010), Principles of plasma physics for engineers and scientists, in *Principles of plasma physics for engineers and scientists*, Cambridge University Press.
- Jayaratne, E. R. (1998), Possible laboratory evidence for multipole electric charge structures in thunderstorms, *Journal of Geophysical Research*, *103*, 1871–1878.
- Jayaratne, E. R., and C. P. R. Saunders (1984), The “Rain Gush”, Lightning, and the Lower Positive Charge Center in Thunderstorms, *Journal of Geophysical Research*, *89*(4), 816–818.
- Jayaratne, E. R., C. P. R. Saunders, and J. Hallett (1983), Laboratory studies of the charging of soft-hail during ice crystal interactions, *Quarterly Journal of the Royal Meteorological Society*, *109*, 609–630, doi:10.1002/qj.49710946111.
- Kabirzadeh, R., N. Lehtinen, and U. S. Inan (2015), Latitudinal dependence of static mesospheric E-fields above thunderstorms, *Geophysical Research Letters*.
- Krehbiel, P. R. (1986), The electrical structure of thunderstorms, *Tech. rep.*, New Mexico Institute of Mining and Technology.
- Krider, E. P., and C. Guo (1983), The peak electromagnetic power radiated by lightning return strokes, *Journal of Geophysical Research*, *88*(3), 8471–8474, doi: 10.1029/JC088iC13p08621.
- Kuo, C. L., A. B. Chen, Y. J. Lee, and L. Y. Tsai (2007), Modeling elves observed by FORMOSAT-2 satellite - Kuo - 2007 - Journal of Geophysical Research: Space Physics (19782012) - Wiley Online Library, ... *Space Physics (1978 . . .*
- Lang, D., D. Hogg, K. Mierle, M. Blanton, and S. Roweis (2010), Astrometry.net: Blind astrometric calibration of arbitrary astronomical images, *The Astronomical Journal*, *139*, 178–1800.
- Laux, C. O. (1993), Optical diagnostics and radiative emissions of air plasmas, Ph.D. thesis, Stanford University.

- Laux, C. O. (2003), Optical diagnostics of atmospheric pressure air plasmas, *Plasma Sources Sci. Technol.*, *12*, 125–138.
- Lay, E., S. Xuan-Min, and A. Jacobson (2014), D region electron profiles observed with substantial spatial and temporal change near thunderstorms, *Journal of Geophysical Research: Space Physics*, pp. 4916–4928, doi:10.1002/2013JA019430. Received.
- Liang, C. (2014), Electrodynamic and Thermodynamic Modeling of the Lightning Channel, Ph.D. thesis, Stanford University.
- Liang, C., B. Carlson, N. Lehtinen, M. Cohen, R. A. Marshall, and U. Inan (2014), Differing current and optical return stroke speeds in lightning, *Geophys Res Lett*, pp. 1–7, doi:10.1002/2014GL059703. During.
- Mach, D. M., and W. D. Rust (1989), Photoelectric return-stroke velocity and peak current estimates in natural and triggered lightning, *Journal of Geophysical Research*, *94*(D11), doi:10.1029/JD094iD11p13237.
- Marshall, R. (2009), Very Low Frequency Radio Signatures Of Transient Luminous Events Above Thunderstorms, *Thesis*, pp. 1–184.
- Marshall, R., R. Newsome, and U. S. Inan (2008), Fast Photometric Imaging Using Orthogonal Linear Arrays, *IEEE Transactions on Geoscience and Remote Sensing*.
- Marshall, R. A. (2012), An improved model of the lightning electromagnetic field interaction with the D-region ionosphere, *Journal of Geophysical Research*, doi: 10.1029/2011JA017408.
- Moss, G. D., V. P. Pasko, N. Liu, and G. Veronis (2006), Monte Carlo model for analysis of thermal runaway electrons in streamer tips in transient luminous events and streamer zones of lightning leaders, *Journal of Geophysical Research: Space Physics*, *111*, 1–37, doi:10.1029/2005JA011350.
- Newsome, R. (2010), Ground-Based Photometric Imaging Of Lightning EMP-Induced Transient Luminous Events, *Thesis*, pp. 1–167.

- Newsome, R. T., and U. S. Inan (2010), Free-running ground-based photometric array imaging of transient luminous events, *Journal of Geophysical Research*, 115.
- Nucci, C. A., C. Mazzetti, R. F., and M. Ianoz (1988), On lightning return stroke models for LEMP calculations, in *19th Int. Conf. on Lightning Protection, Graz, Austria*, pp. 463–469.
- Parzen, E. (1962), On Estimation of a Probability Density Function and Mode, *The Annals of Mathematical Statistics*, 33(3), 1065–1076.
- Pasko, V. P., U. S. Inan, T. F. Bell, and Y. N. Taranenko (1997), Sprites produced by quasi-electrostatic heating and ionization in the lower ionosphere, *Journal of Geophysical Research*, 102, 4529, doi:10.1029/96JA03528.
- Pasko, V. P., M. a. Stanley, J. D. Mathews, U. S. Inan, and T. G. Wood (2002), Electrical discharge from a thundercloud top to the lower ionosphere., *Nature*, 416(March), 152–154, doi:10.1038/416152a.
- Quick, M. G., and E. P. Krider (2013), Optical power and energy radiated by natural lightning, *Journal of Geophysical Research: Atmospheres*, 118(4), 1868–1879, doi:10.1002/jgrd.50182.
- Rakov, V. A. (2007), Lightning Return Stroke Speed, *Journal of Lightning Research*, 1(1935), 80–89.
- Rakov, V. A., and A. A. Dulzon (1987), Calculated electromagnetic fields of lightning return stroke, *Tekh. Elektrodinam*, 1, 87–89.
- Rakov, V. A., and W. G. Tuni (2003), Lightning electric field intensity at high altitudes : Inferences for production of elves, *Journal of Geophysical Research*, 108, 1–6, doi:10.1029/2003JD003618.
- Rakov, V. A., and M. A. Uman (1998), Lightning Return Stroke Models Including Some Aspects of Their Application, *IEEE Transactions on Electromagnetic Compatibility*, 40(4), 403–426.

- Rakov, V. A., and M. A. Uman (2003), *Lightning: physics and effects*, Cambridge University Press.
- Ratcliffe, J. A. (1959), *The magneto-ionic theory and its applications to the ionosphere*, University Press.
- Rosenblatt, M. (1956), Remarks on Some Nonparametric Estimates of a Density Function, *The Annals of Mathematical Statistics*, 27(3), 832.
- São Sabbas, F. T., D. D. Sentman, E. M. Wescott, O. Pinto, O. Mendes, and M. J. Taylor (2003), Statistical analysis of space-time relationships between sprites and lightning, *Journal of Atmospheric and Solar-Terrestrial Physics*, 65, 525–535, doi:10.1016/S1364-6826(02)00326-7.
- Said, R. K., U. S. Inan, and K. L. Cummins (2010), Long-range lightning geolocation using a VLF radio atmospheric waveform bank, *Journal of Geophysical Research*, 115(D23), D23,108, doi:10.1029/2010JD013863.
- Sentman, D. D., E. M. Wescott, D. L. Osborne, D. L. Hampton, and M. J. Heavner (1995), Preliminary results from the Sprites94 aircraft campaign: 1. Red sprites, *Geophysical Research Letters*, 22(10), 1205–1208, doi:10.1029/95GL00583.
- Shao, X.-M., E. H. Lay, and A. R. Jacobson (2012), Reduction of electron density in the night-time lower ionosphere in response to a thunderstorm, *Nature Geoscience*, 6(1), 29–33, doi:10.1038/ngeo1668.
- Sipler, D. P., and M. A. Biondi (1972), Measurements of O (1D) Quenching Rates in the F Region, *Journal of Geophysical Research*, 77(31).
- Taranenko, Y. N., and U. S. Inan (1993), The interaction with the lower ionosphere of electromagnetic pulses from lightning: excitation of optical emissions, *Geophysical Research Letters*.
- Taranenko, Y. N., U. S. Inan, and T. F. Bell (1992), Optical signatures of lightning-induced heating of the D region, *Geophysical Research Letters*, 19(18), 1815–1818.

- Tascione, T. F. (1994), *Introduction to the space environment*, 2nd ed., Krieger, Malabar, FL.
- Uman, M. A., and D. K. Mclain (1969), Magnetic Field of Lightning Return Stroke, *Journal of Geophysical Research*, *74*(28).
- Vallance-Jones, A. (1974), Aurora, in *Aurora*, chap. no. 9 in G, Springer-Verlag, New York.
- Wait, J. R., and K. P. Spies (1964), *Characteristics of the Earth-ionosphere waveguide for VLF radio waves*, US Dept. of Commerce, National Bureau of Standards.
- Wang, D., N. Takagi, T. Watanabe, V. a. Rakov, and M. a. Uman (1999), Observed leader and return-stroke propagation characteristics in the bottom 400 m of a rocket-triggered lightning channel, *Journal of Geophysical Research*, *104*, 14,369, doi:10.1029/1999JD900201.
- Wescott, E. M., D. Sentman, D. Osborne, D. Hampton, and M. Heavner (1995), Preliminary results from the Sprites94 aircraft campaign: 2. Blue jets, *Geophysical Research Letters*, *22*(10), 1209–1212.
- Wescott, E. M., H. C. Stenbaek-Nielsen, D. D. Sentman, M. J. Heavner, D. R. Moudry, and F. T. São Sabbas (2001), Triangulation of sprites, associated halos and their possible relation to causative lightning and micrometeors, *Journal of Geophysical Research*, *106*, 10,467, doi:10.1029/2000JA000182.
- Williams, E. R. (1998), The positive charge reservoir for sprite-producing lightning, *Journal of Atmospheric and Solar-Terrestrial Physics*, *60*, 689–692, doi:10.1016/S1364-6826(98)00030-3.
- Yair, Y., C. Price, P. Israelevich, A. Devir, and B. Ziv (2003), New Space Shuttle Observations of Transient Luminous Events During the MEIDEX, in *AGU Fall Meeting 2003*.