COSMIC RAY AND NEUTRINO ASTROPHYSICS WITH THE ANITAIII TELESCOPE

A DISSERTATION SUBMITTED TO THE GRADUATE DIVISION OF THE
UNIVERSITY OF HAWAIʻI AT MĀNOA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN

PHYSICS

DECEMBER 2017

By

Benjamin J. Rotter

Dissertation Committee:

Peter Gorham, Chairperson
Gary Varner
Jelena Maricic
John Learned
Richard Mills

Thank you to everyone who helped me and loved me.

I can't express how important you all are.

I couldn't have done it without you.

ABSTRACT

Ultra-high energy cosmic rays (UHECRs), those with energies above 10^{18} eV, have been observed colliding with matter in the Earth's atmosphere by a variety of experiments since the 1960's. These particles are among the highest observed locally on earth, with energies many orders of magnitude higher than those produced in terrestrial accelerators. Due to the extreme energies of these rare particles, much study has been poured into determining their cosmic source accelerators, a search which continues today. This dissertation explores the third flight of the ANtarctic Impulsive Transient Antenna (ANITA) telescope payload in the 2014-2015 Antarctic season, which has a unique opportunity for novel observations of UHE particles. It covers the theory, payload, detector calibration, signal simulation, and presents results of a search for UHECRs and UHEν_τs with the third ANITA flight.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
INTRODUCTION AND MOTIVATION

The universe is vast, surprising, and has still just begun to be explored by humanity. As telescopes become larger and more sensitive, we as a species can peer further away from our tiny blue speck with increasing clarity. However, the absorption and scattering of high energy photons in the interstellar medium clouds much of the energetic universe from our gaze. Nevertheless, there is a constant flux of hadronic particles incident on earth from various cosmic accelerators which survive the attenuation over universe scale distances. These messenger particles provide a window into regions of the universe that are otherwise inaccessible, and their observation remains a frontier in astrophysics.

1.1 Theory of Ultra-High Energy Cosmic Particles

1.1.1 Cosmic ray detection history

Cosmic Rays (CRs) were first observed over a century ago, when an unexpected increase in ionizing radiation was observed as altitude increased (Figure 1.1). Previously, it was expected that terrestrially measured radiation was caused by radioactive element decays within the earth's crust, however this additional radiative component was indicative of a source far from the Earth [15]. This discovery spawned a search for the sources of these particles raining down from the heavens. In the subsequent years, hundreds of experiments have taken up the task of measuring and characterizing CRs incident on the earth, utilizing a variety of techniques [2]. However, the search for cosmic accelerators able to produce the highest energy CRs that have been experimentally measured remains an unsolved problem in physics that persists to the present day.

1.1.2 Cosmic ray physics and cosmology

While the sources of CRs remain unknown, their nature opens up the doors to many novel measurements of cosmological phenomenon that occur at extreme distances and energies. Much of the universe is opaque to the traditional astronomical messenger particle, the photon, at high energies. High energy gamma rays undergo electron-positron pair production in the presence of a magnetic field, effectively halting their travel before they can reach Earth.(Figure 1.2 [16]). Additionally, the exceptionally high energies of particles detected by observatories, on the order of 10^{20} eV, dwarf the energies produced at terrestrial particle accelerators. These particles can shed light on the world in the ultra high energy (UHE) regime.

The observed energy spectrum of these hadronic particles introduced additional puzzles. As energy increases, the number of candidate sources for CRs diminishes, leaving few to no candidates

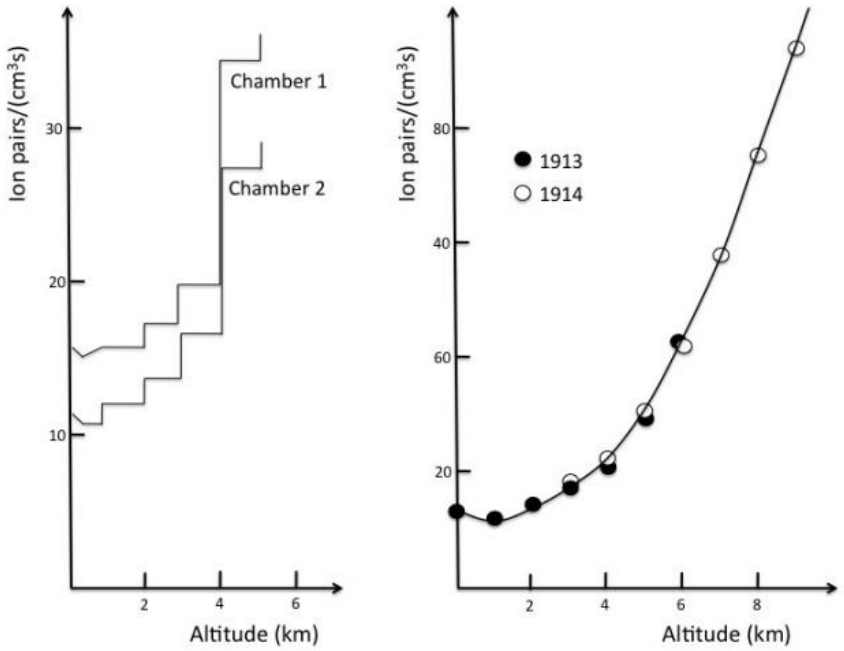

Figure 1.1: Measurements of increasing ionizing radiation as a function of energy, by Hess in 1912 (left) and Kolhrster in 1913 (right), which provided strong evidence for an extra-terrestrial particle flux [1].

within our galaxy [3]. Flux density as a function of energy, shown in Figure 1.4 is well measured up to the EeV energy scale with modern experiments. This spectrum tends to follow a power law, $N \propto E^{-\gamma}$, with decreasing flux as energy increases. Multiple smooth transitions between distinct regions of varying hardness give evidence for multiple source accelerators with varying distances from Earth [17]. In the UHE region, those particles with energies above 10^{18}eV, statistical limitations caused by the incredibly low flux, on the order of one per square kilometer per century, prevents the compilation of statistics required for an identification of sources. Additional measurements are required to determine the transition between CRs emanating from galactic and extra-galactic sources.

Using CRs as astronomical messenger particles presents new difficulties not present in classical astronomy. At lower energies, traveling charged particles will have their courses bent by the Lorentz force while transiting the magnetic fields of stellar objects. Subsequently, inverting their incoming detection vector will no longer yield accurate pointing back to the source accelerator. This is depicted in Figure 1.3. Higher energy CRs are proportionally less effected by magnetic fields due to their increased rigidity, and chargeless particles do not suffer this effect at all. This increases the desirability of detectors that can measure messenger particles with a low flux or a small cross section.

The composition of CRs at the highest energies remains under investigation as well. An ex-

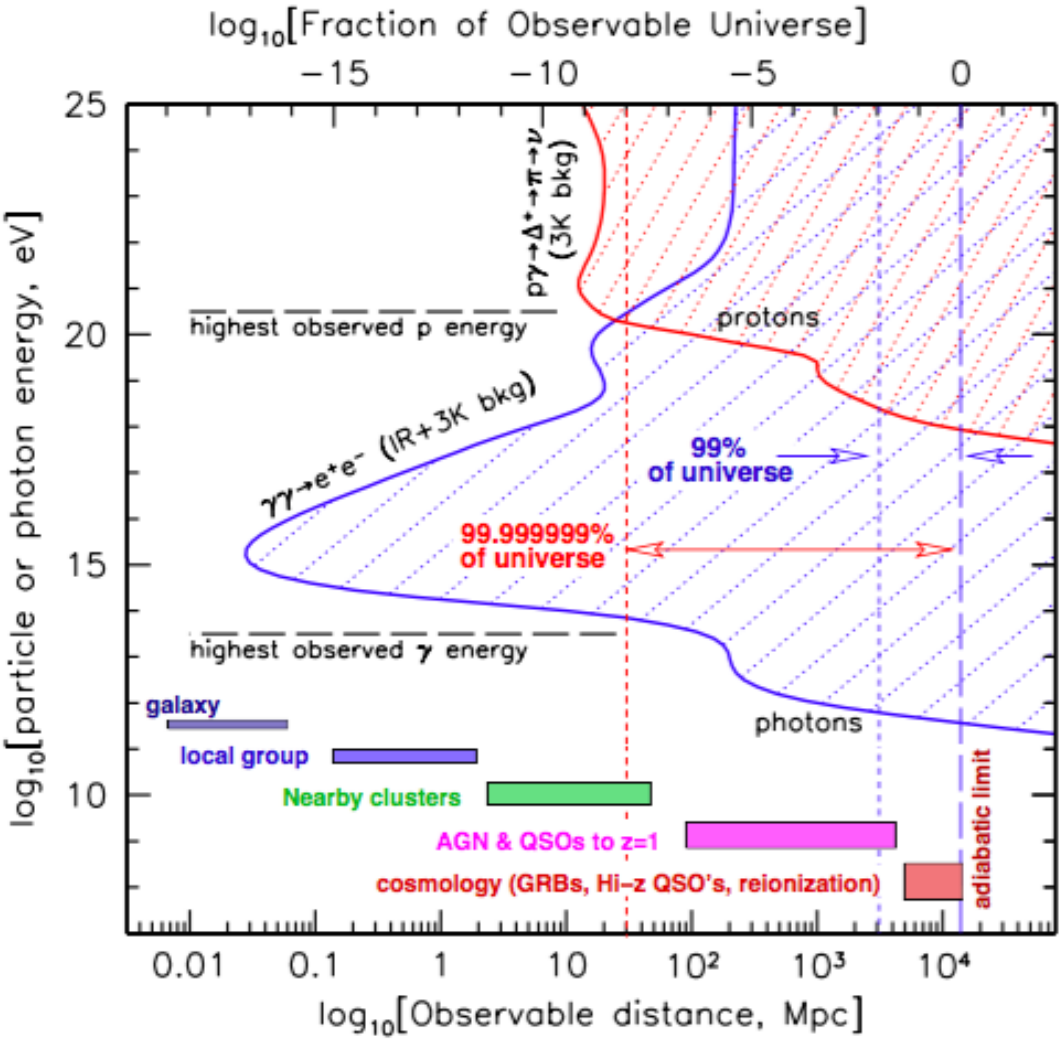

Figure 1.2: Interaction distances for photons and protons. Shaded regions represent regions of the universe which are opaque to an astronomical particle. Plot courtesy of Dr. Peter Gorham.

tensive air shower produced from a CR interacting with an atmospheric molecule has identical secondary detection characteristics whether the initial source particle was heavy ions or protons. However, acceleration mechanisms and evolution models for various compositions vary significantly, impacting cosmic evolution and the resulting cosmic ray flux models [18]. Recent measurements studying the maximum shower depth of air showers suggests a mixed composition, and possibly a hint at a transition from galactic to extragalactic source accelerator at EeV energies [19][20], however disagreement currently exists within the community [21]. Additional experimental observations of these cosmological particles will yield a better understanding of both the structure of matter and energy within the universe at large, as well as fundamental physics processes energetically unachievable from existing collider experiments. The creation and propagation of Ultra High Energy Cosmic Rays (UHECRs) through the universe opens a window to understanding of extreme high energy physics.

1.2 Neutrino astrophysics and the GZK interaction

The steep decline in observed flux measured by detectors alludes to an interaction mechanism that opens up a new detection prospect for UHE Cosmic Rays. A hardening spectrum at 10^{19} eV (see Figure 1.4) corresponds to the cutoff energy for an inelastic collision between a relativistic proton with a low energy gamma ray, via the delta resonance. UHECRs propagating from their source accelerators are subject to this interaction as they transit the cosmic microwave background (CMB) radiation that exists isotropically throughout the universe [22]. This interaction, predicted by Greisen, Zatsepin and Kuzmin [23] and colloqually known as the GZK limit, represents a theoretical high energy limit on particles coming from outside of the galaxy. Particles with energies above this predicted $5\text{x}10^{19}$ eV threshold will weakly scatter off the CMB and decline in energy through emission of secondary particles [23]. Observatories measuring UHE cosmic rays have detected a decrease in the quantity of observed particles consistant with a GZK induced suppression [24]. Since particles are detected up to the limit, it would also be expected that other regions of the universe also contain accelerators capable of creating cosmic rays in excess of the GZK limit, thus motivating a search for a corresponding UHE neutrino flux.

The GZK process has multiple channels, each of which produce an ultra high energy neutrino (UHEν) messenger particle that could subsequently be detected on earth (Figure 1.5 [23]). The cross section of neutrinos has been measured in accelerator facilities to be vanishingly small at modern accelerator scale energies [25]. However, at energies out of reach of modern accelerators, cosmic accelerators could illuminate our understanding of the neutrino cross section. The neutrino cross section at energies above those measured by accelerator facilities has been estimated by using standard model particle physics [26]. This small cross section allows the GZK interaction

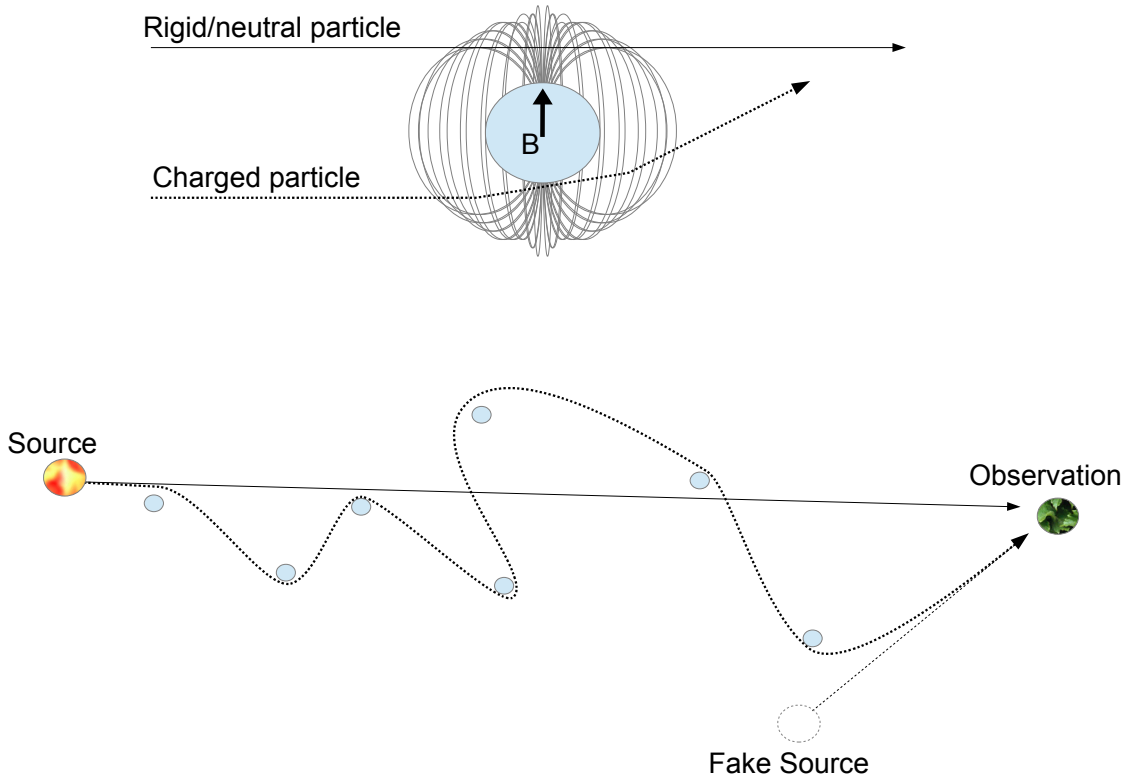

Figure 1.3: Top: A simplified diagram of lorentz-force induced curvature in low energy cosmic rays. Bottom: An example of why magnetic deflection distorts the observed CR source location

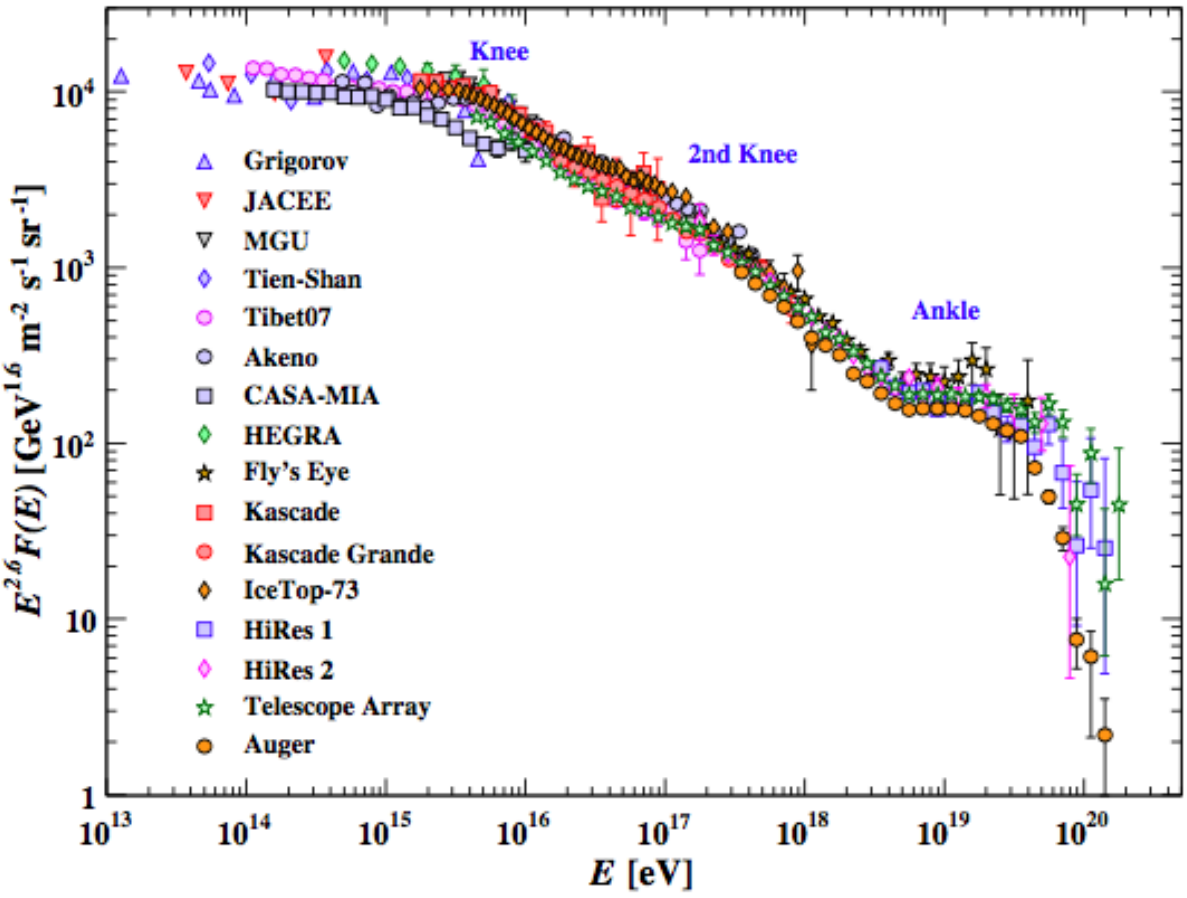

Figure 1.4: The all-particle spectrum as a function of energy-per-nucleus from measurements [2]

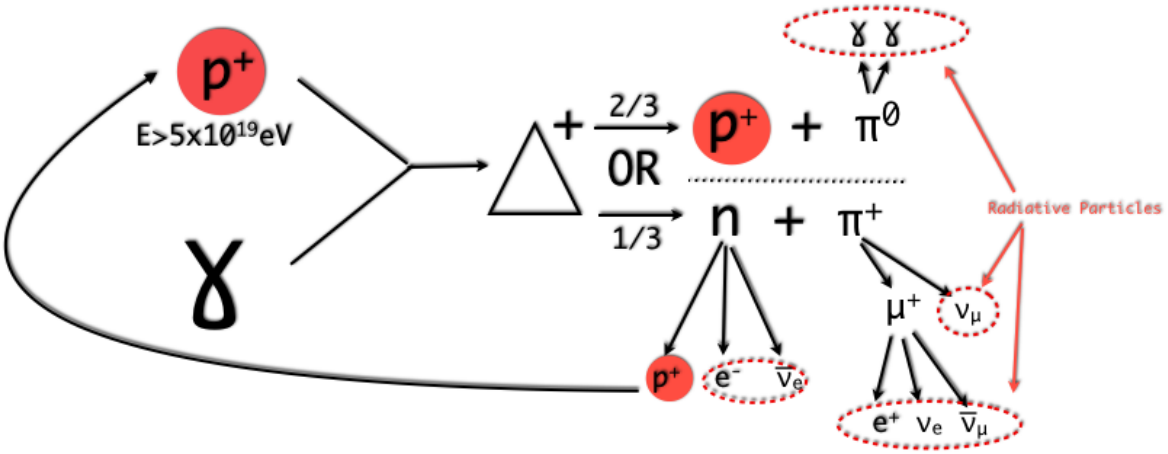

Figure 1.5: A diagram of the possible resultant particles of a GZK interaction between a photon and the CMB.

messenger neutrino to traverse unimpeded through the intergalactic medium before subsequently interacting with Earth. Additionally, the uncharged nature of neutrinos allow them to travel in a straight line from their source locations without Lorentz deflection from intervening magnetic fields. Neutrinos have an extremely small cross section, such that their mean free path length through the universe is essentially infinite. Experiments with the capability to detect these particles will allow an unparalleled view into the depths of the cosmos.

1.2.1 Unexplained UHE Source Mystery

Though there are numerous low to mid-scale energy cosmic ray sources that have strong theoretical motivation and experimental evidence, above $10^{17.5}$ eV the number of candidates becomes constrained and extra-galactic source hypotheses are required [3]. Due to the higher flux of nearby low energy cosmic ray sources, many experiments have gathered evidence to support source candidates within our galaxy as accelerators. However, above the GZK suppression limit there remains little collected experimental evidence that describes the transition to extra-galactic sources. These higher energy objects include Active Galactic Nuclei (AGN), supernova, quasars, gamma-ray bursts (Figure 1.6). However, the small statistics of incident particles at the UHE end of the spectrum makes identifying their acceleration mechanisms much more difficult, and requiring extremely voluminous detectors.

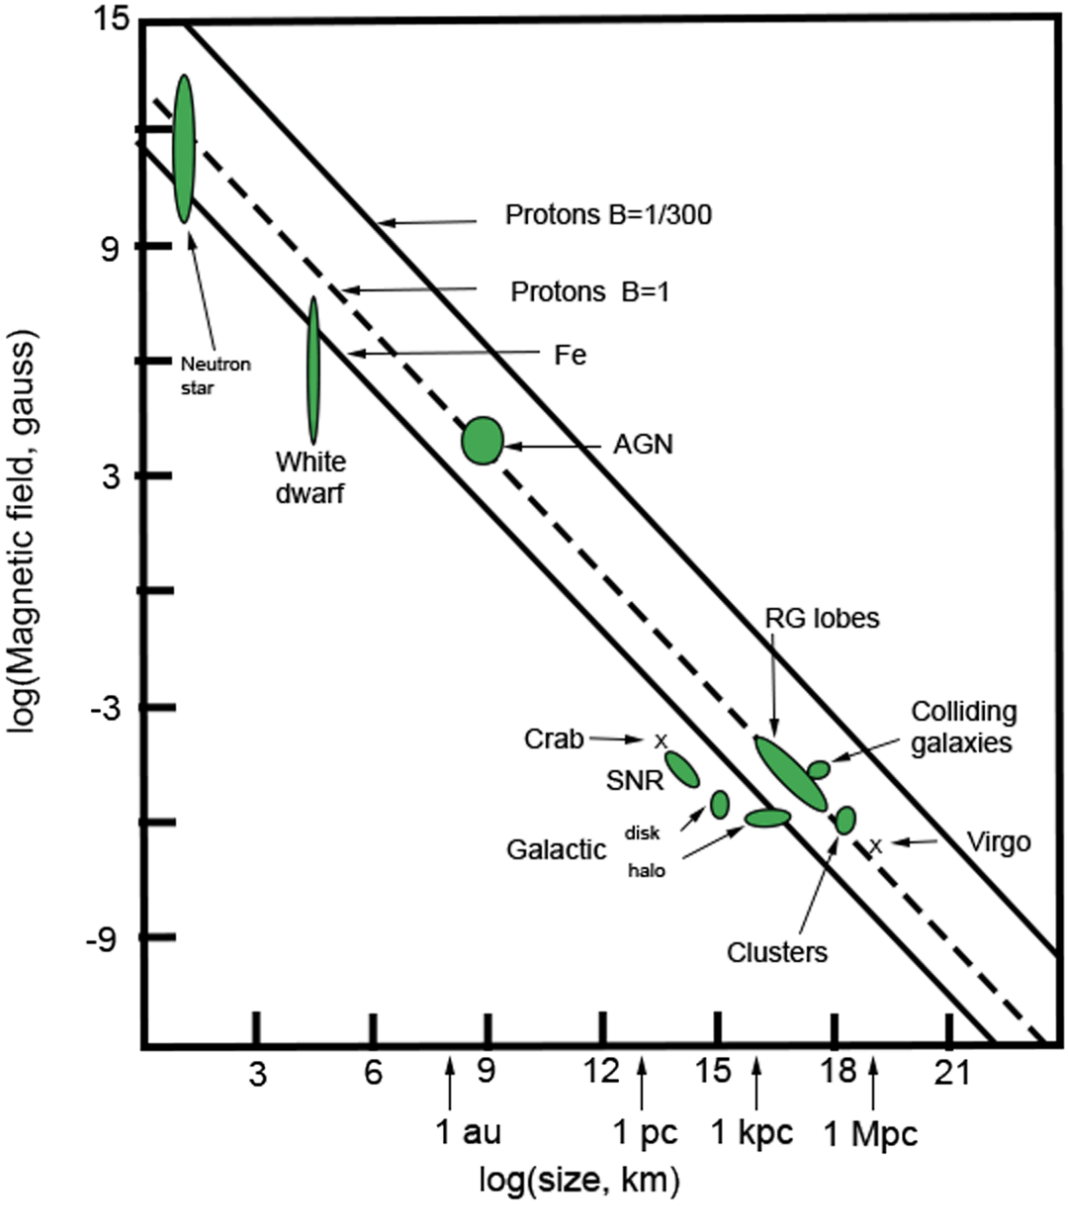

Figure 1.6: Theorized cosmic accelerators plotted by their size and peak magnetic field. The dashed line denotes a field strength capable of accelerating protons to 10^{18} eV. Plot vectorized from [3].

1.3 Neutrino Detection on Earth

The neutrinos produced in GZK interactions are messenger particles that carry information about UHECRs produced outside the galaxy to earth at energies above the GZK suppression effect. Astronomical neutrino telescope observations have been attempted out at lower energies since the discovery of the particle itself [27]. Recently however, there has been a notable increase in detector scale and energy sensitivity. The term "cosmic ray" describes any high energy particle incident on earth (including gamma rays) and thus, for the purposes of this dissertation, neutrinos fit the description of cosmic rays.

The largest notable difference between observations of UHE neutrinos and CRs is the medium in which they are most likely to interact. While CRs have large cross sections and are known to interact with even low density media, practical observations of neutrinos require a medium with far higher density. The primary observation location for CRs is within the atmosphere, which is sufficiently dense to cause an interaction, while neutrinos pass through the atmosphere unimpeded and must be observed in a denser media on the ground. Both particle interaction lengths can be described by the traversed density, or specifically for the case of CRs, the atmospheric depth, with units of $\mathrm{g\,cm^{-2}}$. For UHEνs however, this term becomes larger than the integrated total density of the atmosphere regardless of incoming slant angle. The neutrino will therefore either skim the atmosphere entirely if shallow enough, or interact within the earth if too steep. However, if a dense dielectric solid is introduced or utilized at the Earth's surface by an opportunistic experimenter, it is possible to capture neutrino interactions within the field of view of detection equipment. This equiptment can capture the induced electromagnetic showers from a high energy particle shower.

1.3.1 Charged current and Neutral Current interactions

Neutrinos have two distinct interaction classifications, Neutral Current (NC) and Charged Current (CC) [25]. Feynman diagrams depicting these interactions can be seen in Figure 1.7. Both interactions can occur with either baryons, such as protons and neutrons, or leptons, such as electrons. The two interactions have differing cross sections and result in different resulting particles with a large fraction of the incident energy. These different processes will both affect the initial UHE particle propagation, but only the CC interaction will produce a detectable signal.

A neutral current interaction occurs through the exchange of a Z^0 boson, and transfers a fraction of its energy and momentum to the particle with which it interacted. The resulting particles from the interaction are the same as the initial particles, with the incident neutrino retaining a majority of the momentum. Incident UHE neutrinos can "skip like a stone" off the dense matter of the earth, continuing to propagate, with much of their initial energy, through the earth. The resulting energy transfered to the opposing particle is insufficient to cause a detectable radiative shower.

A charged current interaction involves the exchange of a $W^{\pm}$ boson, and results in the neutrino

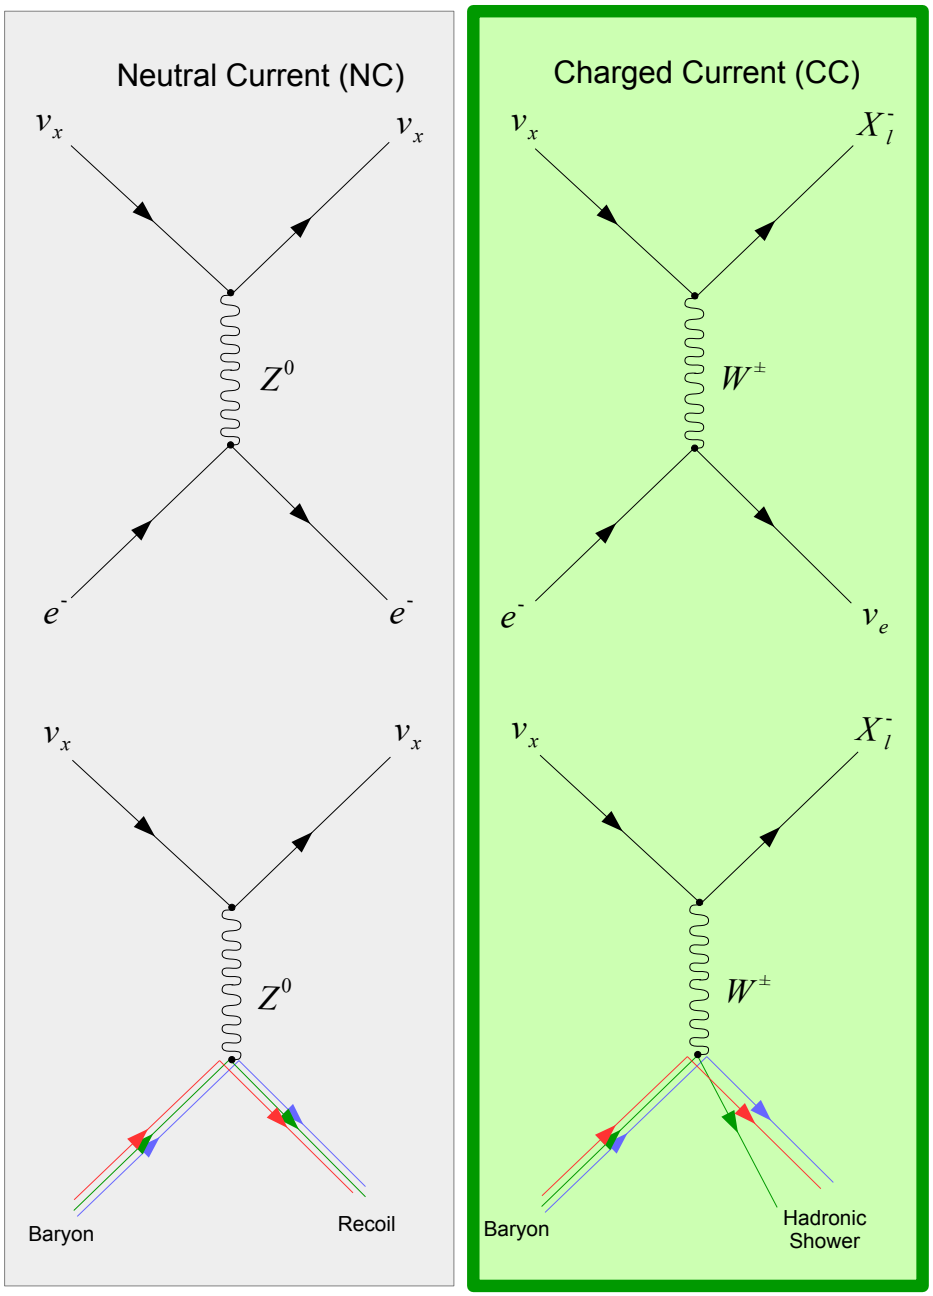

Figure 1.7: Feynman diagrams showing neutrino interactions. The weak force boson mediating the interaction defines whether an interaction is charged current ($W^{\pm}$) or neutral current (Z^0). Neutral current interactions do not result in secondary showers that produce detectable electromagnetic radiation. Charged current interactions, marked in green, can produce extensive showers with and electromagnetic component detectable by instruments. X_l^- denotes a charged lepton with the flavor of the incident neutrino ν_x.

being converted into a similarly flavored and charged lepton. The charged lepton, which retains a large fraction of the energy from the incident UHEν and has a much larger cross section, quickly proceeds to scatter again and spark a large shower of secondary particles. The CC interaction thus leads to a shower that produces a strong electromagnetic radiative signal that, if given a sufficiently transparent interaction medium, can be detected by physics instruments.

1.3.2 Extensive Air/Ice Showers

A highly energetic particle interacting with the hadronic matter present in Earth's atmosphere, or in any dense material, creates an extended series of particle production, scatterings, and decays. The extensive shower of particles scattered or created from a single high energy source particle has been appropriately named an Extensive Air Shower (EAS). These showers produce both traveling messenger particles that can be observed through weakly interacting secondary scatterings in instrumented transparent scintillating media, as well as through primary radiation, which is the topic of this thesis. This primary radiation is stimulated by the appearance of moving charged particle pairs that travel along the principle shower axis. There are two main electromagnetic radiative effects, detailed below and depicted in Figure 1.8, that can be observed in the VHF (3 MHz to 300 MHz) and UHF (300 MHz to 3 GHz) bands of the electromagnetic spectrum. This region of the spectrum has very good transmissive properties in the atmosphere, as well as within many dense solids [28][29][30].

1.3.3 Geomagnetic Radiation

The primary and strongest radiative effect in an EAS is from a longitudinal charge separation created from a Lorentz deflection as the newly created grouping of charged particles move through the geomagnetic field of the Earth.[31] The entirety of the shower is required to have zero net charge (A CR will have a total incident charge of +1 and a neutrino has no charge), and as such an equal number of positively and negatively charged particles are generated. The Lorentz force, $F = q(\vec{E} + \vec{v} \times \vec{B})$, acting with opposite sign on oppositely charged particles, will induce a spacial separation. This dynamically varying charge separation emits radiation with a linear polarization orthogonal to both the magnetic field and the shower axis. The intensity of this radiation is quadratically proportional to both the incident particle energy and local geo-magnetic field. This effect has recently been measured using an accelerator beam induced shower transiting an induced magnetic field. [32]

1.3.4 Askaryan Effect

A secondary radiation component, theorized by Gurgen Askaryan in 1962, is created by a longitudinal charge separation in the shower. The creation and forward scattering of electrons during

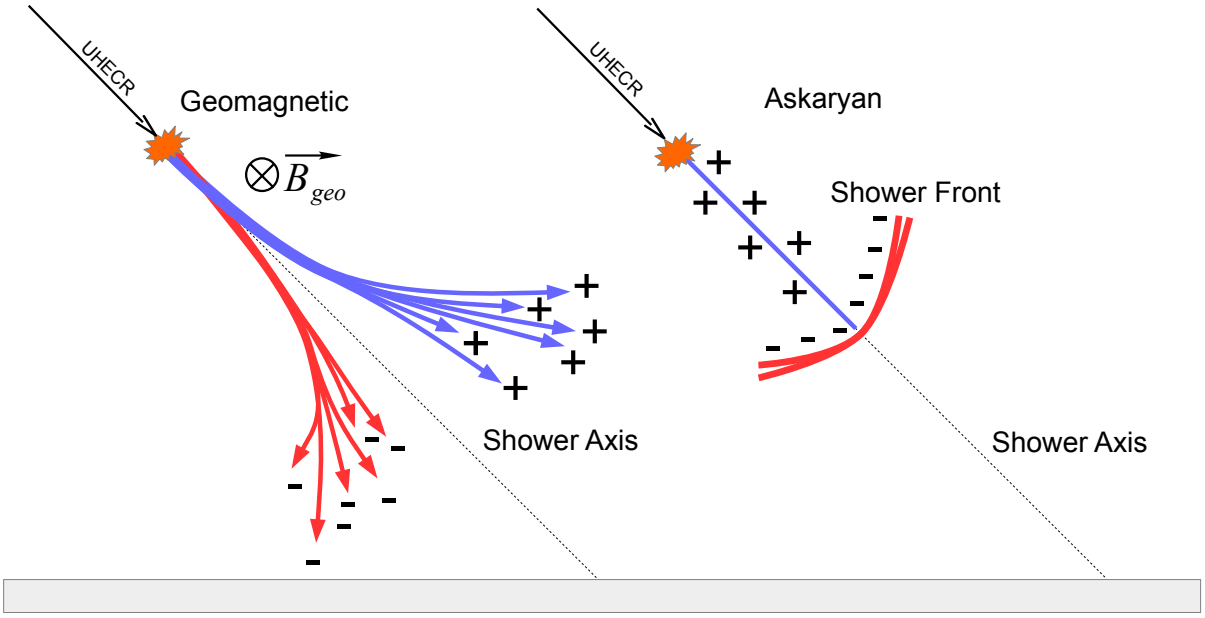

Figure 1.8: Simplified schematic of two EAS radiative trasmission mechanisms, geomagnetic (left) and charge-excess, or Askaryan, radiation (right)

shower development causes a build up of negative charge at the leading edge. Meanwhile, positrons annihilate with electrons within the media, leading to a net positive charge, occurs within the shower core [33]. These two effects create a charge separation with an associated electric field. This relativistically moving dynamic field will radiate coherently at critical off-angles where the wavefronts constructively interfere. This is visible in the time domain as a sharp, broad spectrum impulse as the stationary viewer observes an enormous instantaneous charge flux, as seen in Figure 1.11 [4]. This effect has much more recently been measured at particle accelerators in a variety of materials, including ice,[5] salt [34], and silica sand [35]. An example accelerator-initiated Askaryan pulse in ice measured by the ANITA electronics is visible in Figure 1.10.

1.3.5 Tau-νpecific Detection Prospects

Though the Earth becomes opaque to neutrinos at high energies, secondary particles produced in primary interactions may have a higher mean free path and additional detection prospects. A Tau lepton (τ), produced from a similarly flavored neutrino undergoing a charged-current interaction, have an extremely short lifetime, 2.9×10^{-13} s, and can leptonically decay back into a tau neutrino (ν_τ) which can then continue to propagate. The expected flavor production ratio from typical sources is expected to be $\Phi_{\nu_e} : \Phi_{\nu_\mu} : \Phi_{\nu_\tau} = 2{:}1{:}0$, with few to no produced ν_τ messenger particles. Due to flavor mixing however, the cosmologically long distances traversed by a UHEν is expected

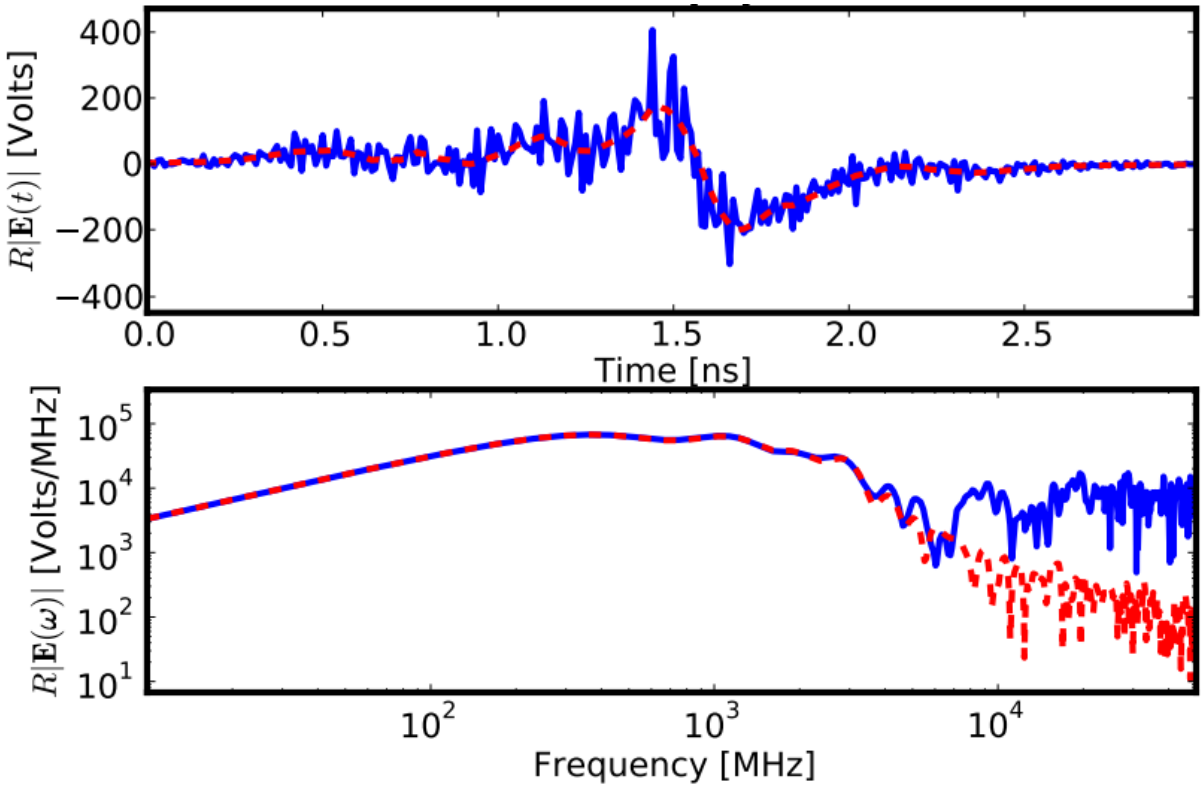

Figure 1.9: The electric field of an Askaryan pulse observed close to the critical angle in the far field generated by a mathematical model(red), and with the ZHS simulation package(blue) [4]

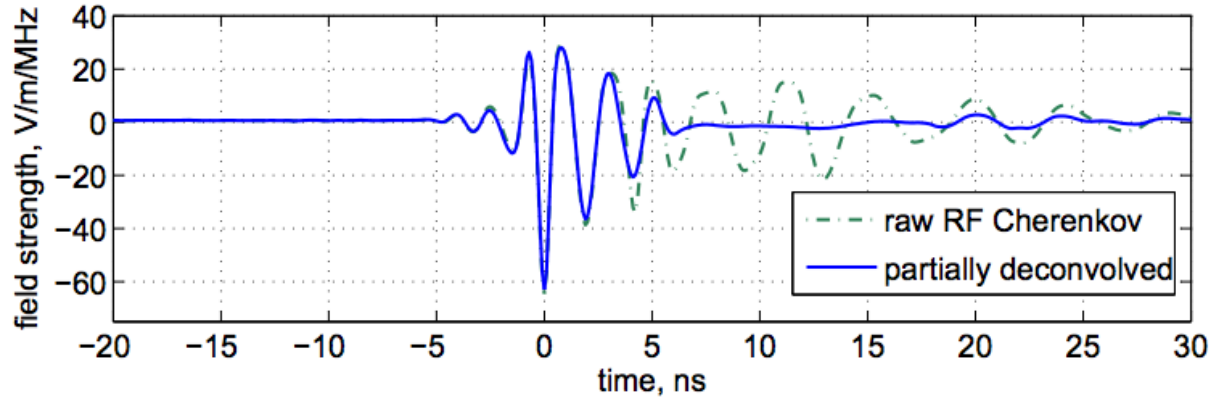

Figure 1.10: An example accelerator initiated Askaryan pulse in ice measured by the ANITA electronics The dispersion in the signal is introduced by the band-pass filtering in the RF signal chain [5]

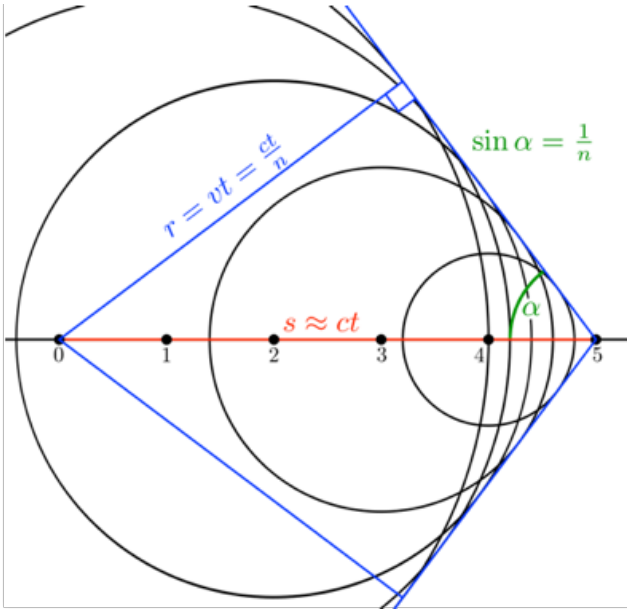

Figure 1.11: An example diagram of the Cherenkov critical angle where the radiative components of a CR shower experience a signal boost.

to provide an even distribution fraction of each flavor for detection at earth, $\Phi_{\nu_e} : \Phi_{\nu_\mu} : \Phi_{\nu_\tau} =$ 1:1:1.

This adds an additional propagation channel to be combined with the probability that the neutrino undergoes a NC interaction, decreasing a catastrophic decay (and resulting electromagnetic pulse) of the incident UHEν. Taus decay hadronically 67% of the time, where the majority of the tau energy is then carried away in the form of pions. The remaining third of the time however the tau decays leptonically, producing a tau neutrino with a large fraction of the initial incident UHE energy. Repeated regenerations of a UHEν_{tau} will increase the propagation length of the initial particle, dramatically extending its travel distance through the earth [36] The particle can then undergo a hadronic decay in the atmosphere, initiating an EAS that travels upwards from the surface of the Earth. A diagram of the proposed, increased lifetime, UHEν_τ interaction timeline resulting in a detectable radio frequency electromagnetic signature is shown in Figure 1.12.

Detection of UHEν_τs with this type of interaction history has recently been attempted by other experiments [37], though no events were identified. A search for UHEν_τ particles was recently done for the ANITA-I data set as well, with one candidate found [38]. ANITA has a unique opportunity to make observations of this type of regeneration

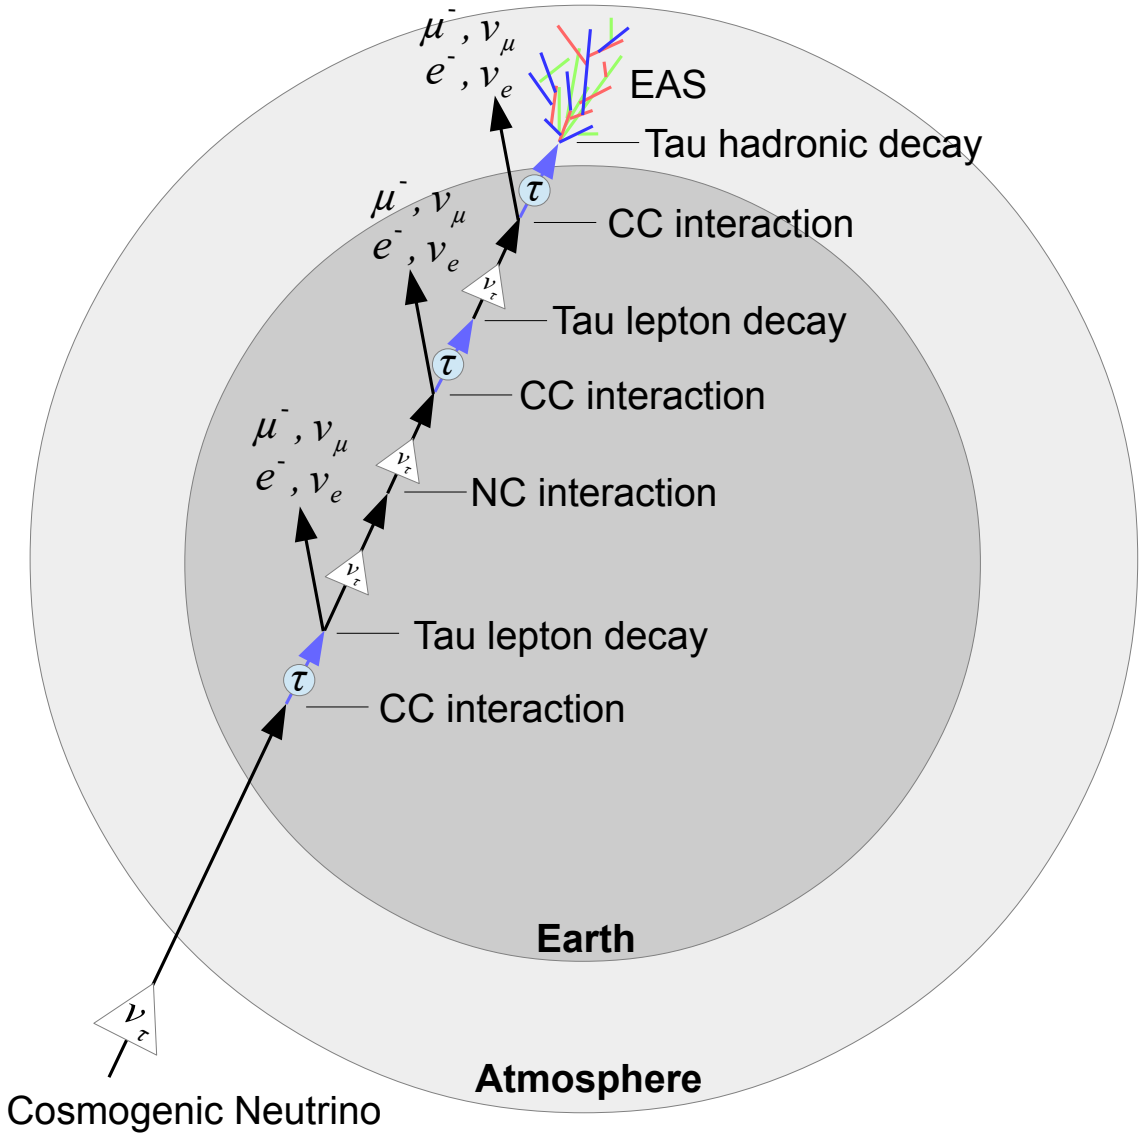

Figure 1.12: An example of an incident UHEν_τ undergoing multiple scatters within the earth before finally starting a detectable EAS in the atmosphere. This type of interaction timeline is proposed as a detectable steeply up-going electromagnetic signature. Figure not to scale.

1.4 ANITA

A detection platform for the physics phenomenon described so far in this chapter is the ANtarctic Impulsive Transient Antenna (ANITA). EASs present in the atmosphere or in the dense, radio transparent, dielectric solid ice sheet covering the souther continent of Antarctica will produce unique and characteristic signals through the above particle interactions and radiative effects. A high altitude telescope platform would have an extremely large field of view, thus increasing the possibility of observing one of these rate interactions.

The following chapters detail the instrument, calibration, and an analysis and results of a CR search with data from the third ANITA (ANITA-III) flight.

CHAPTER 2
ANITA INSTRUMENT AND PAYLOAD

2.1 Overview

The third iteration of the ANITA instrument, seen in Figure 2.1, is a broad-band radio frequency electromagnetic field transient detector and time-domain voltage digitizer. There have been four ANITA flights at the time of this thesis, numbered I through IV. ANITA-I flew in the 2006-2007 Antarctic balloon campaign [39], ANITA-II in 2007-2008 [40], ANITA-III in 2014-2015 (the topic of this thesis) and most recently ANITA-IV in 2016-2017.

2.1.1 Signals and backgrounds

The ANITA instrument's original scientific goal was to observe UHEν interactions within the ice sheet of Antarctica specifically. An analysis of the ANITA-I flight discovered events that were consistent with UHECR interactions in the atmosphere, expanding the scope of detectable UHE particle interactions. The hadronic and electromagnetic showers produced by a collision between a high energy particle and the atmosphere emits radiation which can propagate with low attenuation to ANITA's stratospheric altitude. However, a low flux density has thus far prevented their detection. Null result measurements for neutrino detections from the first two ANITA flights set experimental limits on some phenomenological models, which remain as the strongest limits as of this writing.

ANITA's additional detection capabilities for UHECR interactions within the atmosphere above the continent provide an expected, previously measured, population of astrophysical signal events. The radiative emission of the resulting showers initiated by UHECRs has two detection signatures, a direct signal and a signal first reflected off the ice-air boundary at the surface of the continent.

Lastly, UHEν_τ particles interacting within the earth which produce a τ that decays and initiates an upgoing shower in the atmosphere, as described in 1.3.5, can be detected by the ANITA instrument.

CR signatures, dominated by geo-magnetically induced radiation, are expected to be primarily horizontally polarized with respect to the payload, while in-ice neutrino signatures are expected to be primarily vertically polarized. However, since there have been no UHEν detections, the ultimate characteristics of the shower remain unknown. However, the ANITA instrument is optimized for detection and digitization of any of these signatures.

There are two major background noise sources for the ANITA instrument: thermal Johnson-Nyquist radio frequency background present throughout the field of view, and human activity generating non-thermal anthropogenic electromagnetic radiation. The expectation for Johnson-

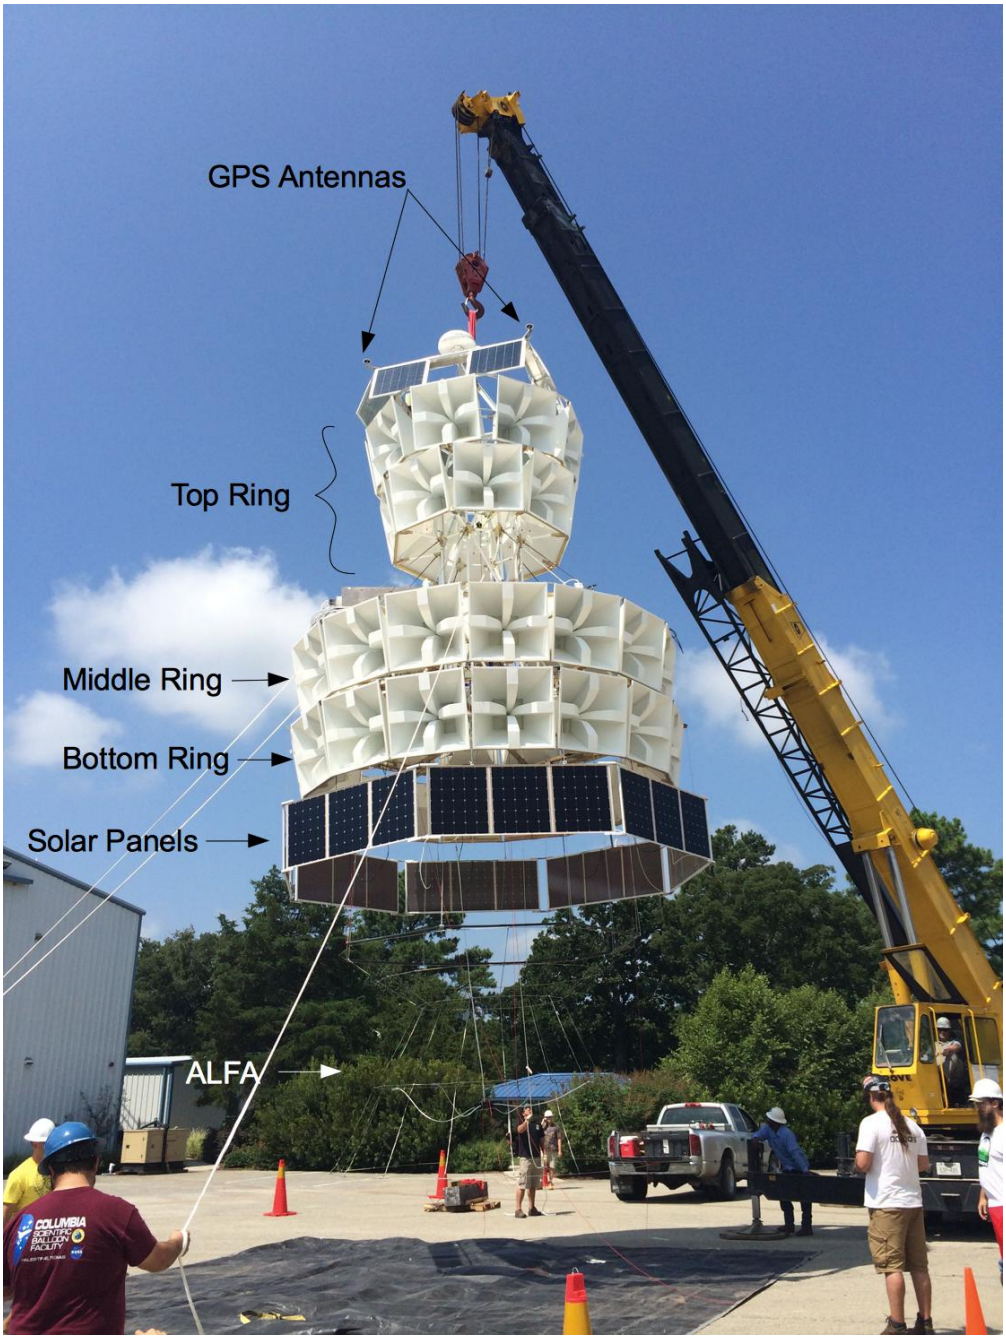

Figure 2.1: An image of the ANITA-III telescope during the test deployment in Palestine TX, Summer 2014. Marked from top to bottom are two of the nine GPS antennas, the three vertically separated rings, the deployed solar panel array, and a deployed ALFA antenna. The solar panels visible at the top of the payload are those used by the NASA Small Instrument Package (SIP). Not visible are the ANITA instrument crate and other objects on the deck of the experiment (between middle and top antenna rings).

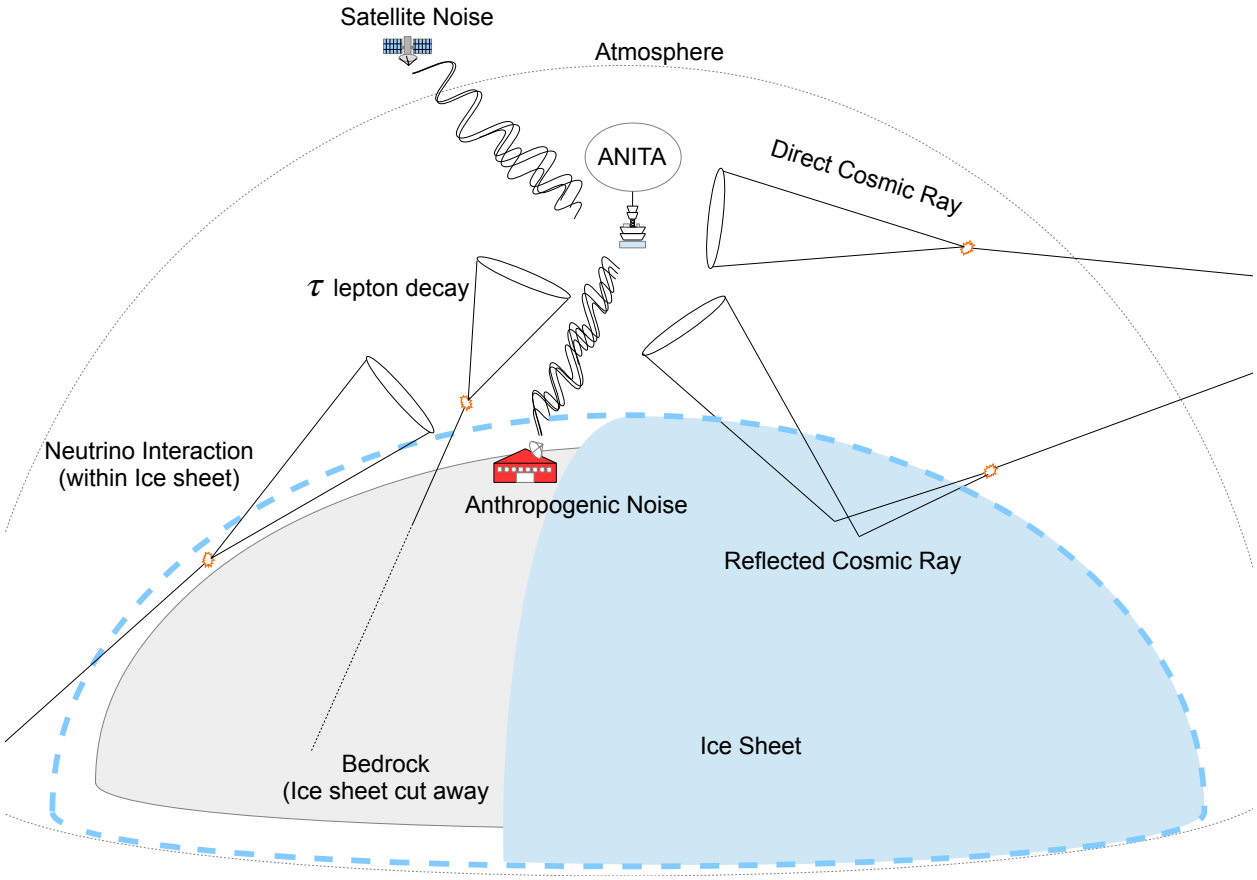

Figure 2.2: A cartoon of the three distinct impulsive observable signals present in the ANITA-III flight data addressed in this thesis. Not pictured is thermal background noise, which is emitted from all matter within the field of view of the instrument. Not to scale.

Nyquist noise is discussed in subsection 2.4.3. Anthropogenic background, which is effectively impossible to model completely, is discussed more fully in section 5.4.

A simplified diagram showing the signal and background incident on the payload is found in Figure 2.2.

2.1.2 Instrumentation

The main instrument structure consists of a collection of radially positioned, outward facing, broad-band, directional quad-ridge horn antennas. These antennas couple the incident electromagnetic radiation into a coaxial signal line, which is then amplified through a cascaded low noise-figure amplifier chain, band-pass filtered via analog compound filters into the relevant frequency band, before finally being digitized and stored to a non-volatile hard disk. (Figure 2.3). The antennas

are separated into three different vertically spaced rings, with each ring separated into 16 different azimuthally facing "phi-sectors." (Figure 2.5) These co-pointing antennas allow multiple observations of a signal event, each with unique physical distance baseline offset between them. The digitization is triggered by a parallel combinatoric square-law-integrating power detector circuit which arrives at a decision with minimal analog waveform buffering. After digitization, waveforms are read out via a CPCI interconnect backplane to a ruggedized, conductively-cooled CPU writing to on-board redundant storage devices. The payload is capable of reading out 50Hz of full payload waveforms with low (<5%) deadtime to disk. Geographic position and relative orientation data, as well as multiple diagnostics systems, are recorded and precisely time-stamped by a GPS-disciplined clock. Selections of the waveforms and housekeeping measurements are telemetered down to ground servers for in-flight diagnostics, however full recovery of the data vaults is desired for a complete analysis. The remainder of this chapter will expand upon these systems and the theoretical basis for their design and construction.

2.2 Antennas

The ANITA instrument utilizes a collection of dual-polarization quad-ridge horn antennas to couple the electromagnetic field incident on the payload into an electrical signal that can be digitized and stored for later analysis. The antenna design specifications prioritize a flat gain and phase response over a gigahertz of bandwidth with a 6:1 upper-to-lower frequency ratio, high directionality to reduce noise and boost signal fidelity, two orthogonal polarizations with roughly co-located phase centers, and light weight. An in-progress image of the top ring antenna assembly is shown in Figure 2.7 in which two sides of the antenna can be seen, and where both the RF ports and attachment mechanisms are visible.

2.2.1 Geometry and Orientation

ANITAI began with only two vertically separated rings and 32 antennas, however subsequent flights have packed more into the available space dictated by the flight vehicle launch envelope. Between ANITAI and ANITAII an additional eight deployable nadir ring antennas were added that unfolded into flight position immediately after launch. Between ANITAII and ANITA-III an additional 8 antennas were added to the lower ring, equalizing the number per phi sector and bringing the total count to 48. ANITAIV, being a fast-turn-around re-flight, flew the same number of antennas as ANITA-III.

Antennas are located to maximize geometric baseline distance between pairs covering the same field of view in order to maximize angular resolution recovered from interferometric pointing. The benefits of increased baselines are discussed further in section 4.5. The additional antennas added in each ANITA flight increase the number of baselines possible for this interferometric pointing.

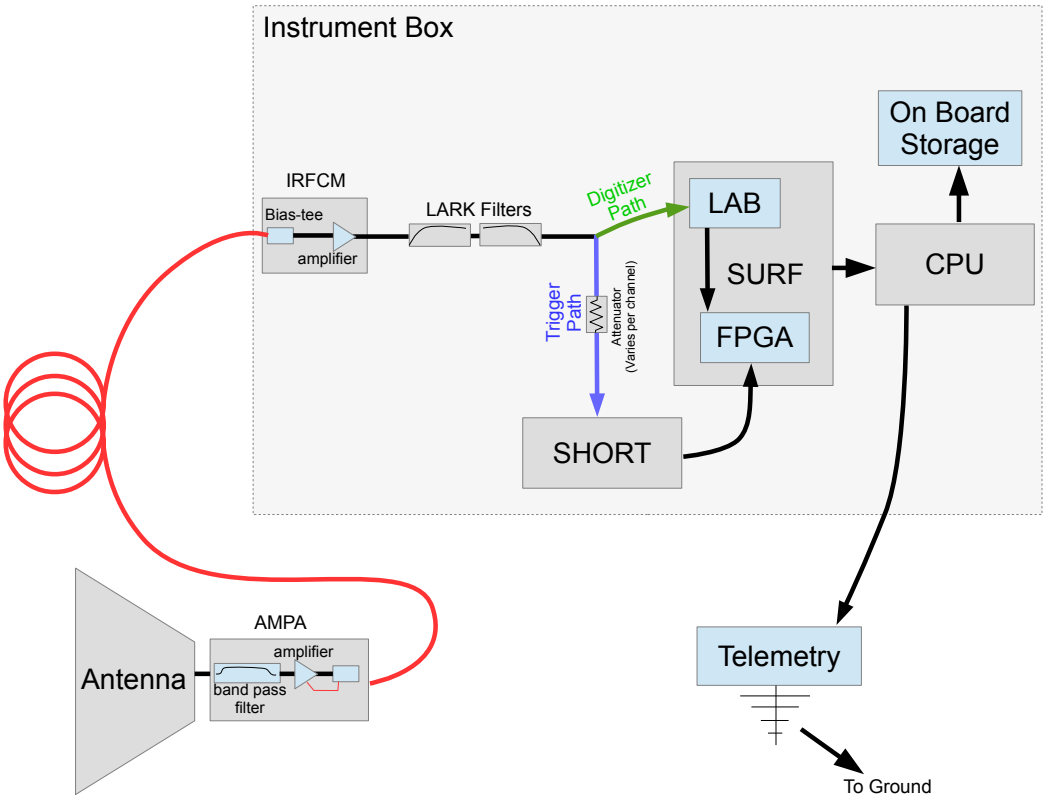

Figure 2.3: A simplified diagram of the ANITA-III signal chain from antenna to data storage or telemetry.

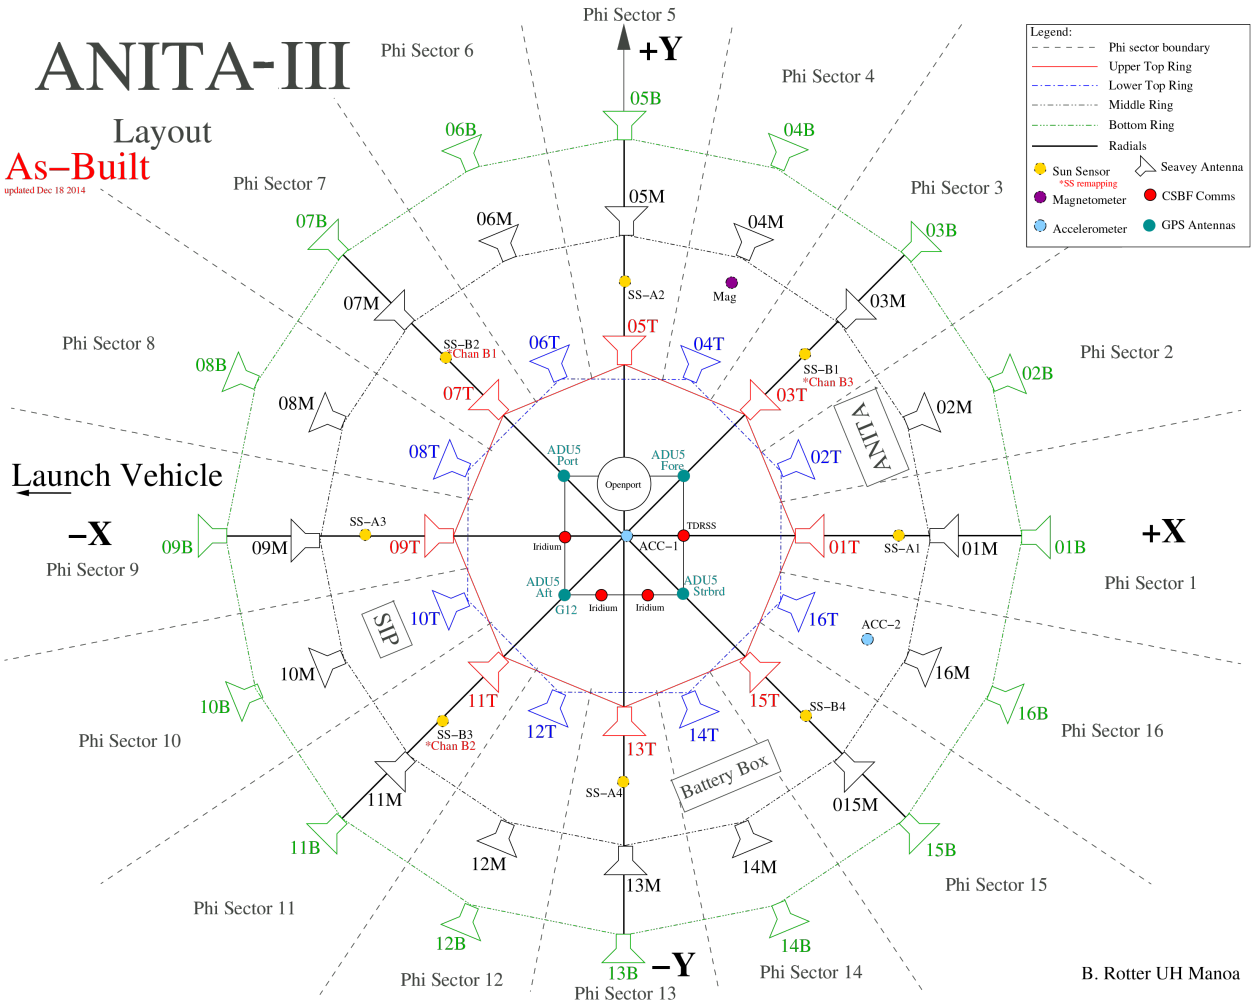

Figure 2.4: A top-down diagram detailing the locations of various components on the ANITA-III instrument as it was flown. Visible are the locations of the NASA Support Instrument Package (SIP) and ANITA instrument box as they relate to the GPS antennas and quad-ridge antenna phi sectors.

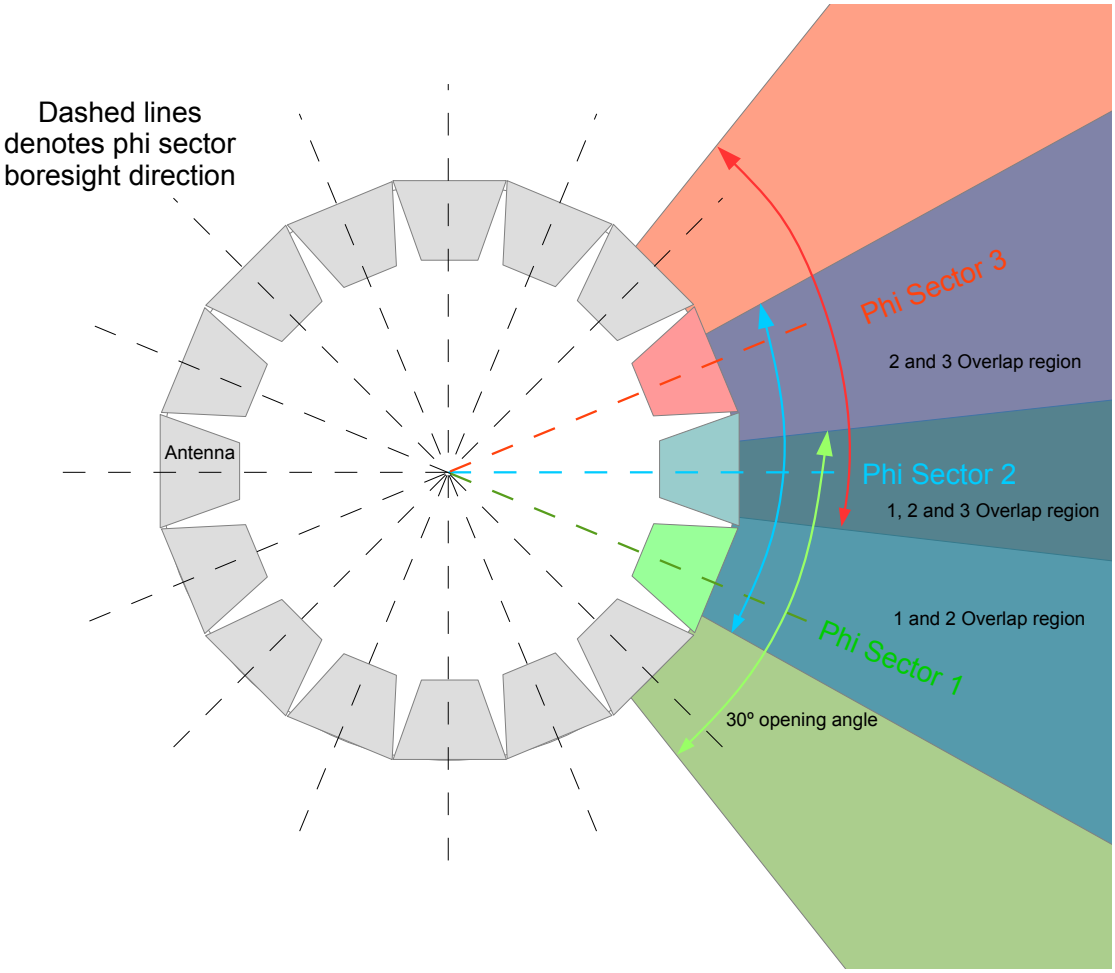

Figure 2.5: A top view of a single ring of horn antennas achieving complete azimuthal coverage. The antennas have a 30° opening angle. The antenna gain pattern is discussed further in the Calibration section.

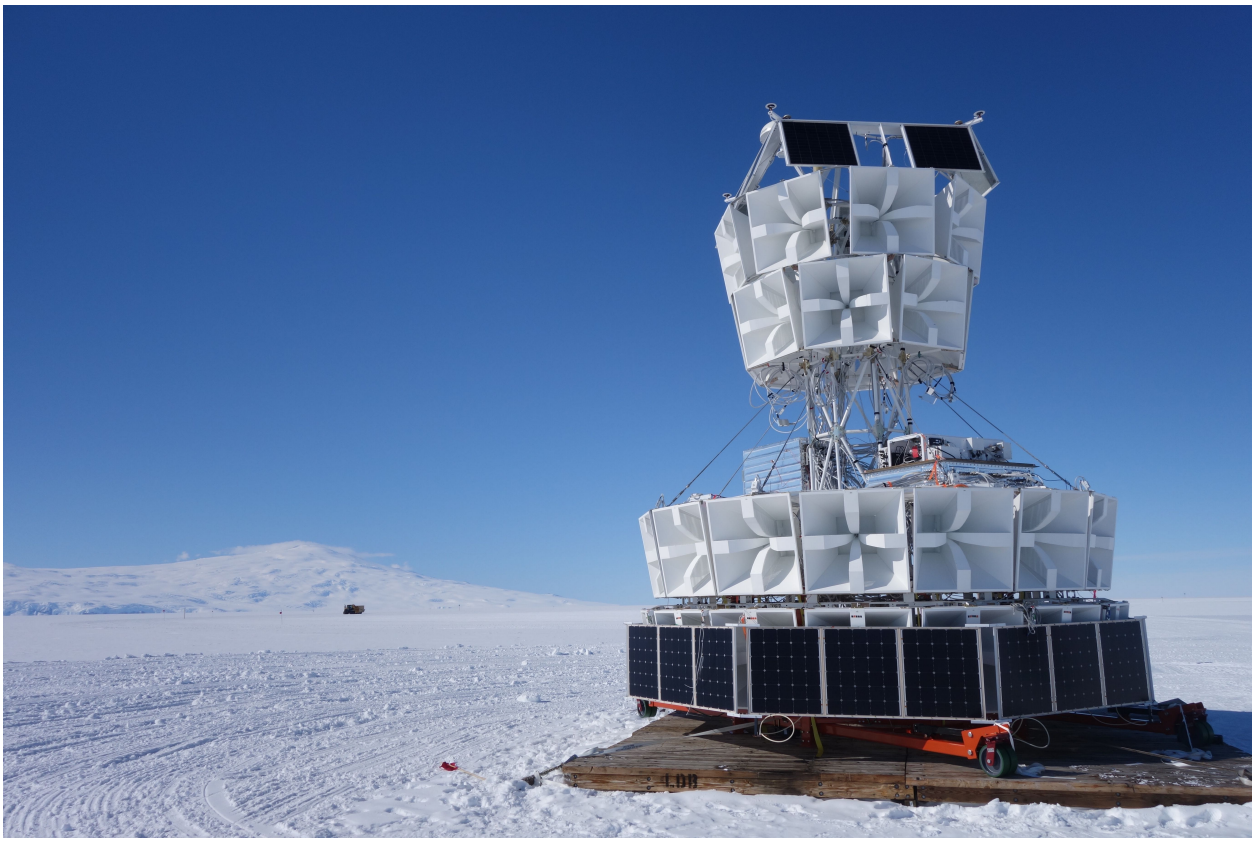

Figure 2.6: ANITA-III during GPS calibration on the "dance-floor" at the Long Duration Balloon (LDB) facility at McMurdo Station in Antarctica, December 2014. In the background is Mt. Erebus. In this image the solar panels are in their retracted state, the instrument crate is visible on the left hand side of the central column, and the NASA Support Instrument Package (SIP) is visible on the right hand side. The orange structure seen under the payload is a stand used during construction and testing of the instrument.

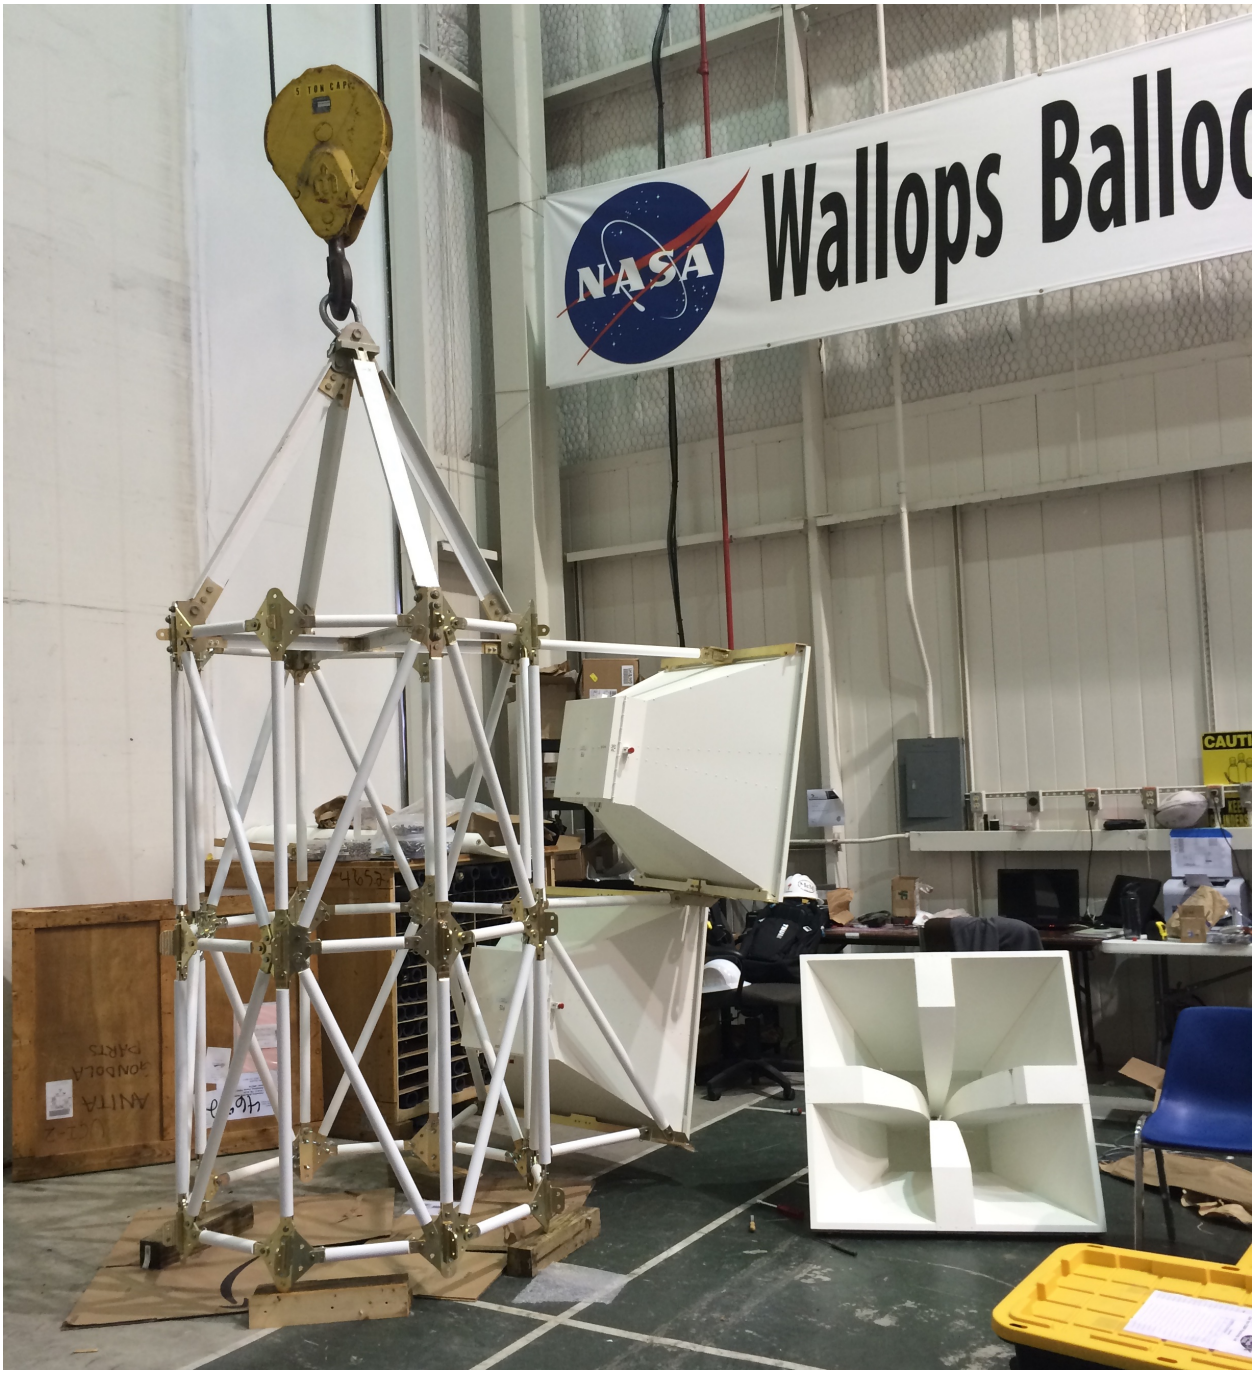

Figure 2.7: The top ring and apex structure of the ANITA-III payload during assembly at the Columbia Scientific Balloon Facility in 2014. Visible is the center column of the gondola structure, three of the dual-polarization antennas, two of which have been mounted. On the antennas the RF ports are visible, capped by red protective plastic covers. The top ring of antennas consists of two rows of eight antennas each, which are staggered to match the two lower ring boresight directions.

This provides a $\frac{1}{\sqrt{N}}$ incoherent noise reduction for coherently summed waveforms, increasing the signal to noise ratio of the final measured event.

The antennas are pointed at a 10° downward angle to put their maximum response slightly below the horizon of the Earth, where signals are most likely to appear. The 40km expected float altitude of the payload places the horizon at a 6 degree downward angle. Signals of any category have their most probable observation location near the horizon, where the integrated atmospheric volume is at a maximum.

The size of the antennas is dictated by the minimum desired frequency, a value dominated by both physical payload launch envelope constraints as well as anthropogenic CW noise, such as FM radio transmissions, prevalent in the VHF band. The use of a quad ridge horn type antenna, used to increase antenna directivity without adding significant dispersion, explains the additional width of the antenna geometry.

The dual, orthogonally facing polarization measurements from each antenna are desired to map out the complete Stokes parameters of any incident signal. These two polarizations allow each antenna to contribute two signal channels covering the same solid angle to the total payload waveform readout array. As all expected detection channels will be linearly polarized at characteristic angles, a thorough understanding of event polarization is required for efficient candidate identification. A diagram describing the primary field lines for each polarization on the antenna is shown in Figure 2.9.

2.2.2 Antenna Directivity Justification

The ANITA instrument can utilize its geometric symmetry and full azimuthal coverage to restrict the response angle of the antennas. Increased antenna directivity improves signal response while keeping noise power constant. ANITA is limited by thermal background noise created from the 220-250 K surface of the Antarctic continent, the 3K CMB, and cascaded electronic amplifier noise. Antenna thermal noise power is calculated by integrating the average temperature over the field of view of the antenna and amplifier noise is cascaded after the antenna, neither of which will vary with directivity. However, signal power at the antenna port scales linearly with directivity. Assuming a constant efficiency, an increase in the directivity of an antenna is linearly related to its gain by Equation 2.1, where G is the antenna gain, $E_{antenna}$ is the efficiency and D is the directivity.

$$G = E_{ant} * D \tag{2.1}$$

Additionally, scattered and possibly unknown anthropogenic sources permeate ANITA's field of view. Smaller antenna opening angles reduce the fractional effect that these small solid angle anthropogenic sources will have on the payload. Maximizing antenna directivity comes at a cost of requiring additional antennas and digitization channels, increasing power consumption and complexity.

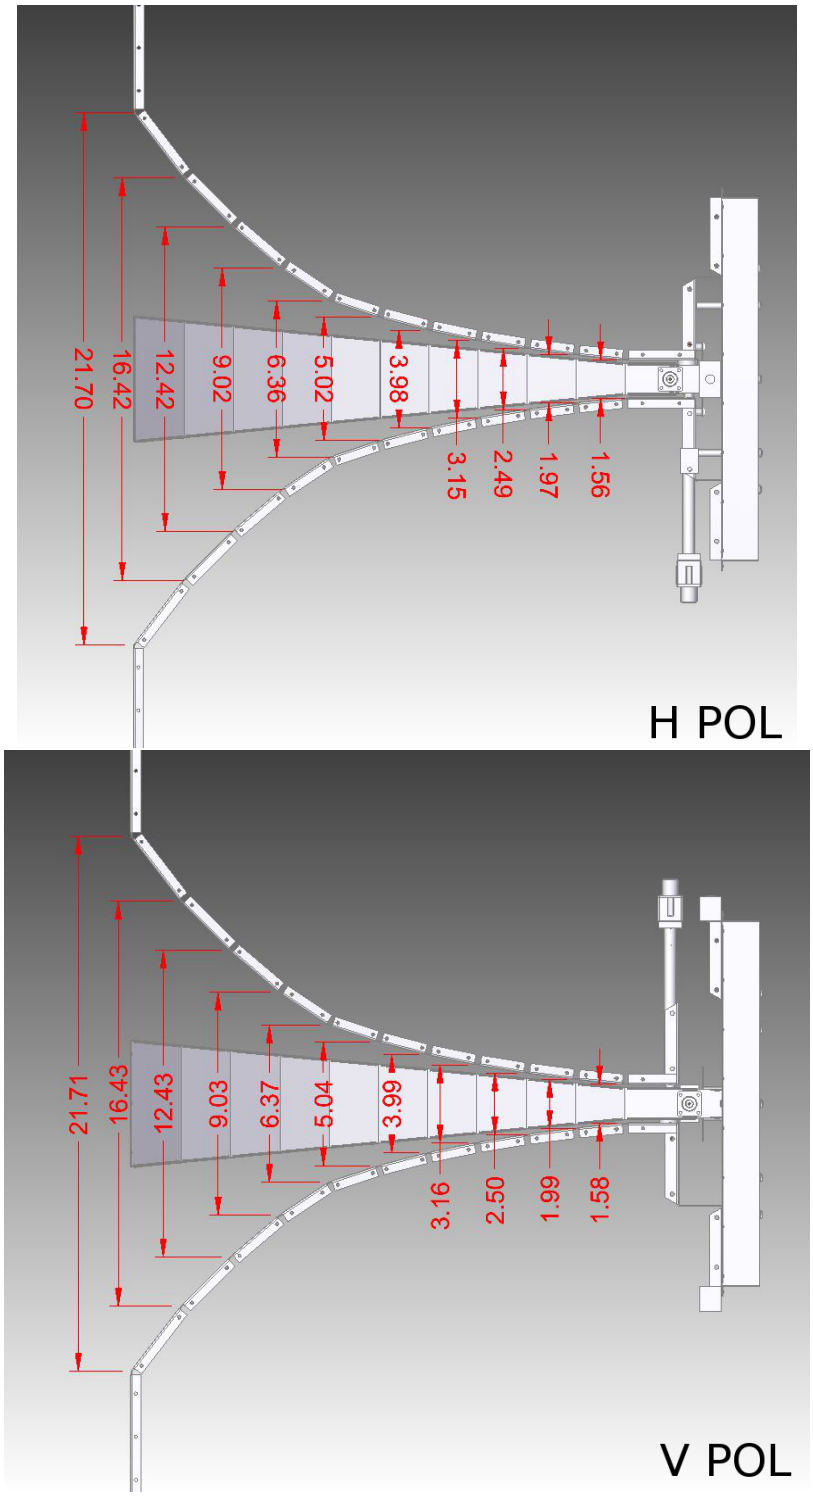

Figure 2.8: Measurements of the antenna ridge separation generated from the manufacturing CAD diagram by Christian Miki. Units are in inches.

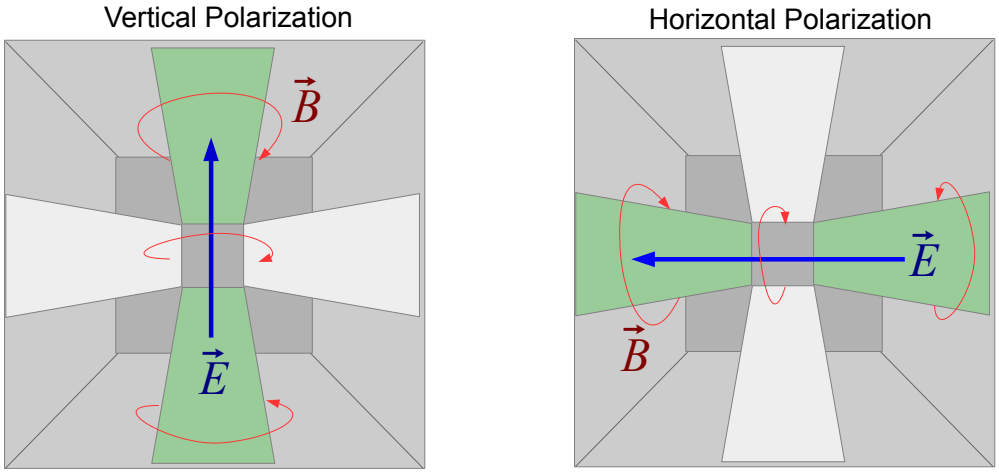

Figure 2.9: A diagram of the field line components for the Seavey dual polarization broad band horn antennas. The E field is the principle measurement for each polarization, while the B field is the incidental cross-polarized measurement. The antennas are specifically designed to minimize cross-polarized signal acceptance.

There are two maximum directivity constraints that each antenna must satisfy. First, each phi sector's response must overlap with neighboring phi sector pairs to establish azimuthal directionality from interferometric baselines. Second, their elevation must encompass both up-going reflected cosmic ray and neutrino signals, as well as have some sensitivity to earth-skimming and slight down-going air showers. As each phi sector is 22.5° separated from its nearest neighbor, a 60° opening angle allows each sector to cover slightly over half the field of view of its neighbor. Each polarization of the antenna has approximately symmetrical gain, however the E-plane and H-plane responses for each polarization are not equivalent, leading to a requirement for off-axis calibrations.

Measurements and calibrations of these antennas is explored further in subsection 3.5.2.

2.2.3 ANITA Low Frequency Antenna (ALFA)

In addition to the quad ridge horn antennas, ANITA-III flew with an additional VHF deployable "quad-slot" antenna with an omni-directional azimuthal response mounted underneath the main structure, pictured in Figure 2.10. The frequency response of the antenna was designed to sit below that of the main instrument, capturing waveforms from 40 to 80 MHz. As many radio frequency EAS experimental measurements are done at lower frequencies, this antenna was added to allow direct comparison between ANITA's observed particle track radiation events and those at ground based observatories. Due to a lack of additional digitization channels in the system, as well as on-board high-pass analog filtering on the SURF board, this antenna was heterodyned with a 900MHz Local Oscillator (LO) to up-convert the signal to a measurable frequency before being combined with channel 05TH. 05TH in turn was low-pass filtered to remove its signal and noise contribution in this signal region. The data from this antenna is not used at any point in this thesis, however its presence is notable and the asymmetry introduced in the system must be accounted for in all analysis steps.

2.3 Filtering

The radiative power from an Askaryan or EAS geomagnetic signal both have a wide bandwidth, however there are a few considerations that require the signals read by the ANITA telescope be band-limited.

For UHECRs, the individual radiative particles in the shower core begin to be resolved, resulting in a loss of coherence of radiated power above that band. This is visualized in Figure 2.11. This effect is stronger at angles further from the peak Cherenkov angle away from the shower axis.[4] This lack of coherent signal power in the UHF region provides a high frequency suppression to the total signal bandwidth.

On the low frequency end of the spectrum, the radiated power from the shower for UHEνs is expected to rise as a function of frequency up to 1GHz. Manufacturing and design limitations of a

Figure 2.10: Photo of the deployed ALFA antenna in the 2014 hang test of ANITA-III in Palestine Texas. The photo has been modified to emphasize the outline of the antenna.

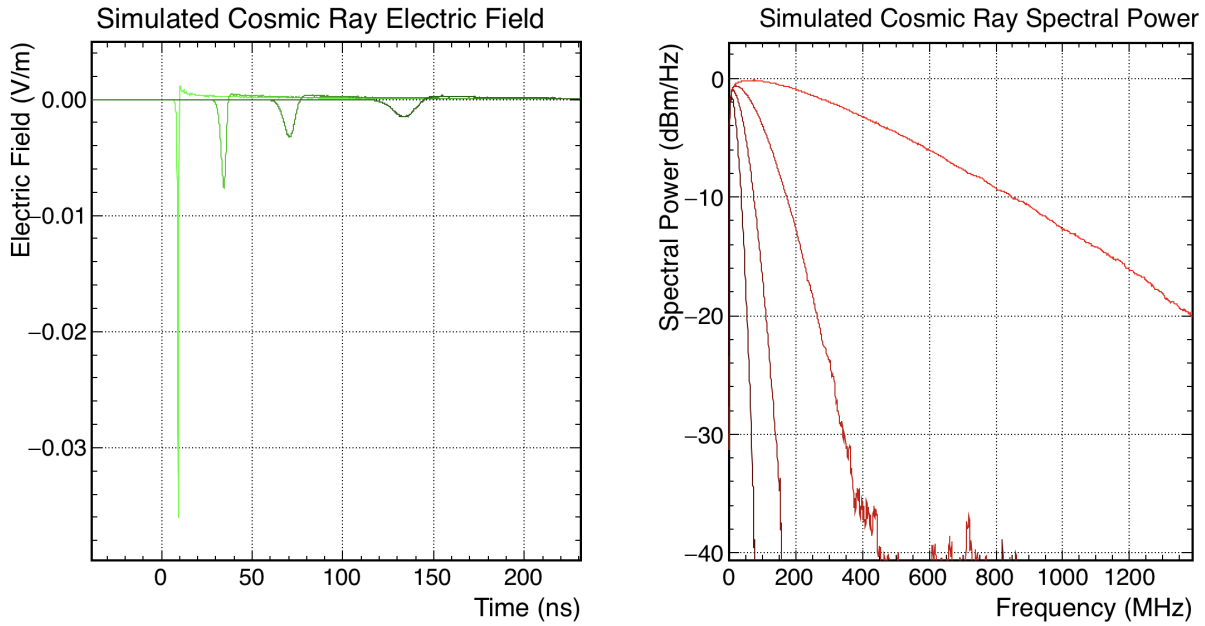

Figure 2.11: ZHAires [6] simulated UHECR electric field radiation transients at, and off, the critical coherence angle. Left: The peak coherence angle of the shower radiation, with increasing distances offset from that peak in darkening hues. The highest amplitude pulse bright green is the peak of the Cherenkov ring. Right: The frequency power spectrums in corresponding hues of red. As can be seen, high frequency power is significantly reduced while observed at locations away from the critical angle,

long wavelength antenna coupled with the high utilization of the VHF band by radio transmitters on satellites and ground stations lead to a requirement that lower frequencies, below 180MHz in the case of ANITA-III, be removed.

2.3.1 Multiple Stages

This band pass filtering is done in two stages, one immediately after the antenna, before the pre-amplifier in order to prevent saturation of the pre-amplifier from out of band signals, and one after the amplification chain in order to remove out of band amplifier noise. The primary filter must have an extremely low in-band loss, as any loss introduced by the filter is observed as backfilled broad spectrum noise temperature, a value which is cascaded through the entire amplification chain. This was accomplished with a custom band pass filter designed and manufactured by Lark Engineering. The secondary filtering was accomplished with two discrete low-pass/high-pass filters that do not require such a low in band loss characteristic due to their position in the signal chain.

2.3.2 Digitizer Bandwidth

The digitizer, which converts the electromagnetic field incident on the payload into a digital value that can be stored, also influences the selection of the band edges. The maximum sampling rate of the Analog to Digital Converter (ADC) chip used in ANITA, 2.6GS/s, yields a 1.3GHz Nyquist sampling frequency. However the 3dB analog bandwidth is limited to 800MHz[8]. Any out of band power will be aliased into the signal band and present itself as increased in-band noise[41]. Further discussion of the digitizer chip in occurs in section 2.5.

2.4 Amplification

The extremely low expected signal power incident on the payload must be significantly amplified in order to have a detectable amount of power. The amplification for the ANITA-III instrument was accomplished with two stages, one close to the antenna within the custom built module, and one upon entering the instrument crate itself within the iRFCM (Internal Radio Frequency Conditioning Module).

2.4.1 First Stage Amplifier: The AMPA and DDAMPA

The front end amplification was accomplished by two similar custom built modules named the Antenna Mounted Pre-Amplifier (AMPA) and the (historically named) Drop Down AMPA (DDAMPA). The internals for the AMPAs and DDAMPAs are diagrammed in Figure 2.12 and images of the modules are shown in 2.13. Each enclosure contains components for filtering, amplification, and a bias network power trasmission component. A summary of the gain and noise figure measurements for the ensemble of units is shown in Figure 2.14.

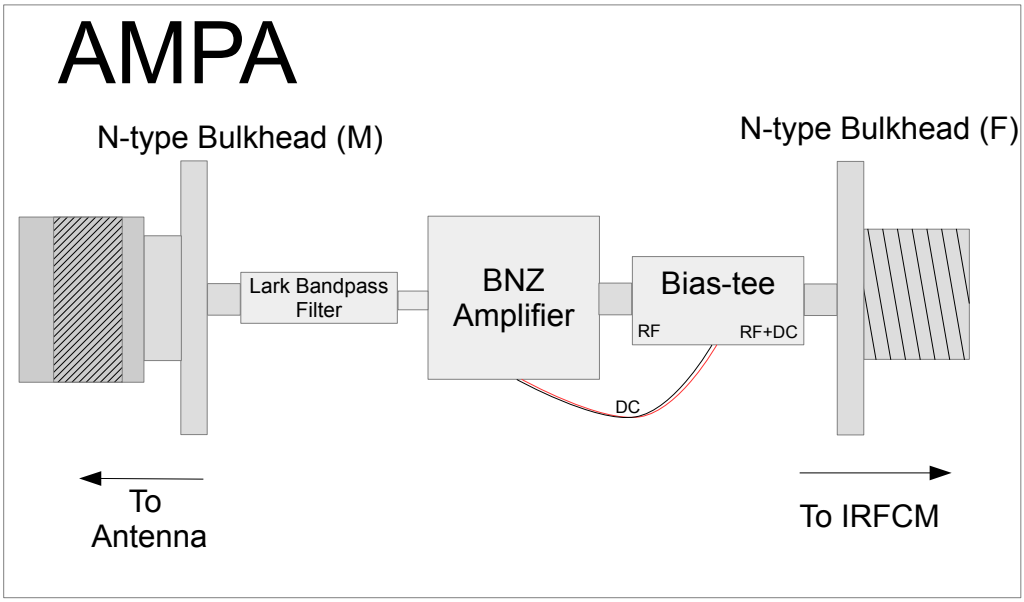

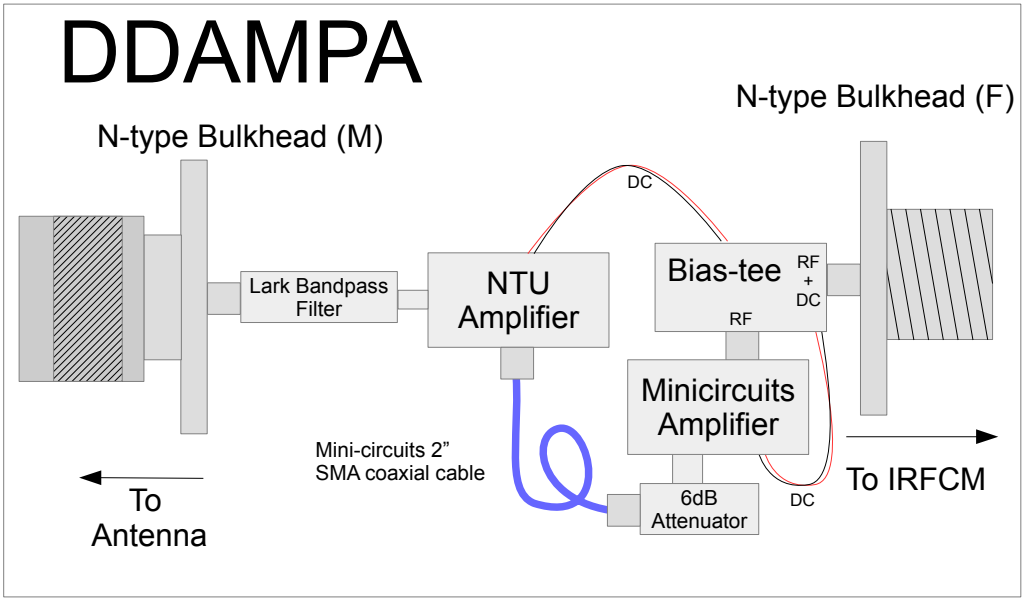

Figure 2.12: Block diagram for AMPA and DDAMPA front end amplification units

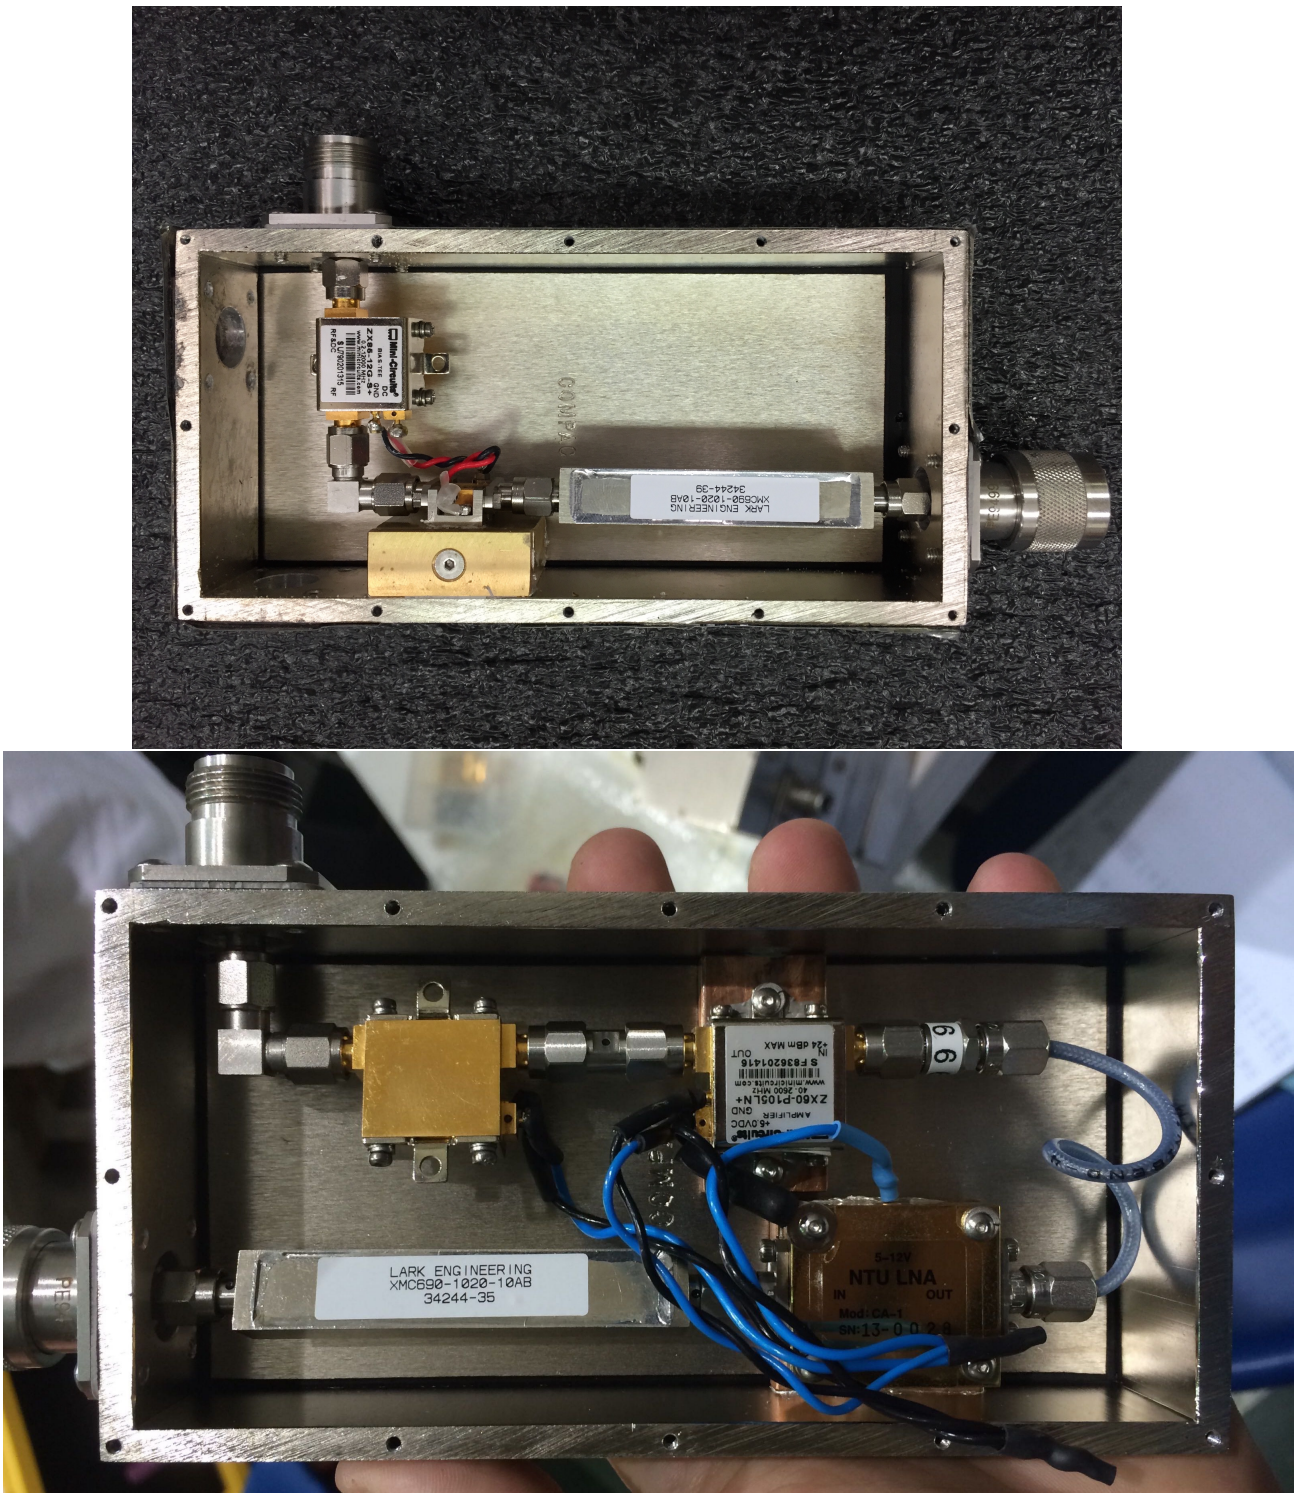

Figure 2.13: The AMPA (top) and DDAMPA (bottom) internals. AMPA photo courtesy of Jarred Roberts.

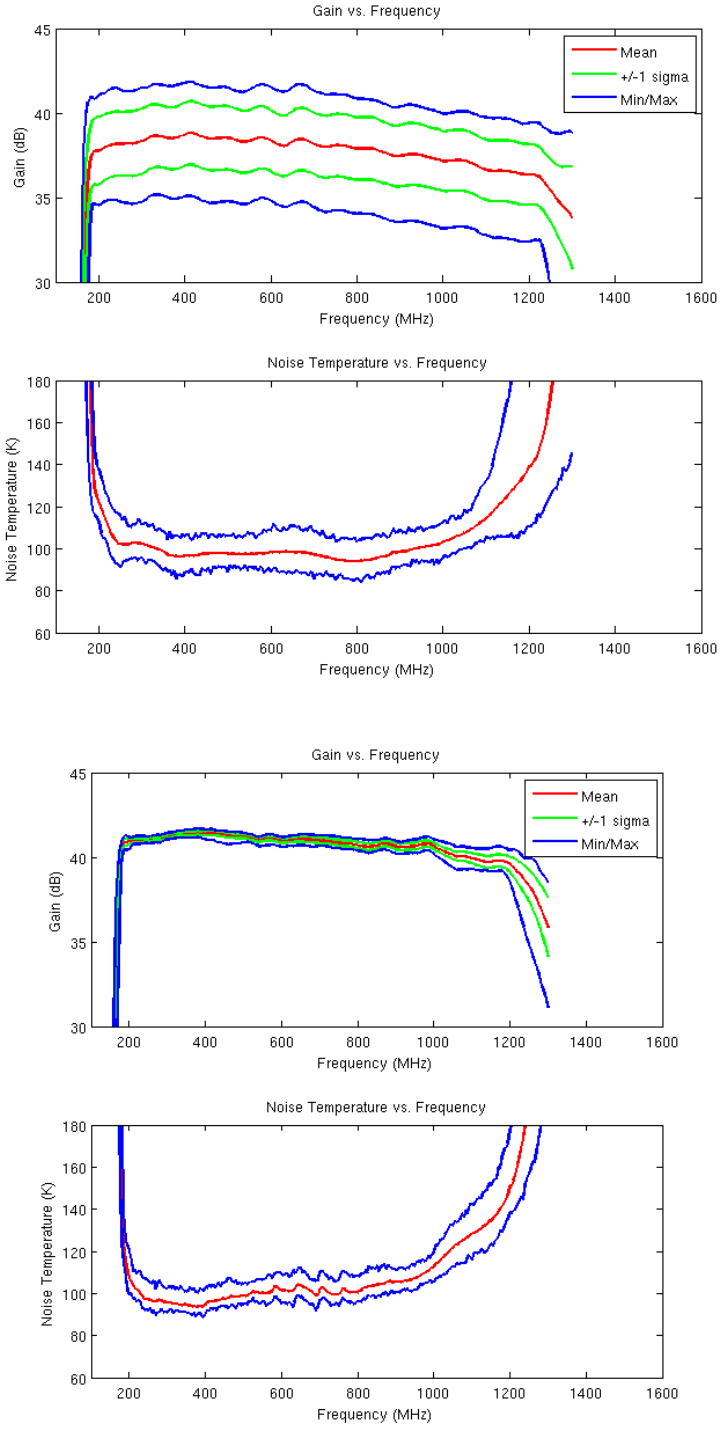

Figure 2.14: Mean and standard deviation of gain and noise figure measurements for the AMPA (top) and DDAMPA (bottom) amplifier modules, as reported within ANITA

2.4.2 iRFCM

The four custom built Internal Radio Frequency Condionting Modules (iRFCM), the internals of which can be seen in figure 2.15, handle second stage amplification and bias network supply for 24 of the RF signal channels. The block diagram for a single iRFCM channel is presented in Figure 2.16. There are three major active RF components for each signal chain present within the module. The AmpLite, a high gain, high dynamic range amplifier built by Patrick Allison of OSU, a tuning attenuator for gain balancing of the 96 different channels, and a bias-tee which, as its name suggests, adds a DC bias to the center conductor of the coaxial transmission cable between the AMPA and the iRFCM. This bias is used to power the AMPA and DDAMPA amplifier modules. The power gain for all 96 iRFCM channels is shown in Figure 2.17.

2.4.3 Expected electromagnetic field variations

Depending on the radiation mechanism being observed, the power of an EAS signal measurable at the detector is quadratically proportional to the energy of the incident cosmic ray particle[7]. This is shown in Figure 2.18 This signal power is convolved with the thermal and anthropogenic antenna noise present within the field of view of the detector. The constant thermal emission of matter in the field of view of an antenna introduces a Johnson-Nyquist thermal noise component, which is calculated in Equation 2.2,

$$P = kBT \tag{2.2}$$

where P is the thermal noise power observed across a termination resistor in Watts, k_b is Boltzmanns constant, T is the absolute temperature observed temperature in Kelvin, and B is the frequency bandwidth in Hz. The antenna, creating a smooth transition between the impedance of free space and that of the coaxial transmission network, has a noise temperature dictated by the integral of the temperature of objects in its field of view ($T(\theta, \phi)$), weighted by the gain pattern of the antenna ($G(\theta, \phi, \omega)$), as shown in equation 2.3.

$$\int_{-\pi}^{\pi} \int_{\theta}^{2\pi} \int_{0}^{\infty} G(\theta, \phi, \omega) T(\theta, \phi) sin^2\theta d\theta d\phi d\omega \tag{2.3}$$

2.4.4 Noise Figure

Since amplification is required, a minimum of additional electronics noise introduced by the detector. Besides a background of thermal noise, the dominant continuous noise source is electronics noise from the amplifiers. This noise, unlike the thermal noise incident on co-pointing antennas, is not coherent as each amplifier adds its own uncorrelated noise component. The noise figure is additionally increased by any loss of power in front of the amplifiers. Since the noise figure cascades

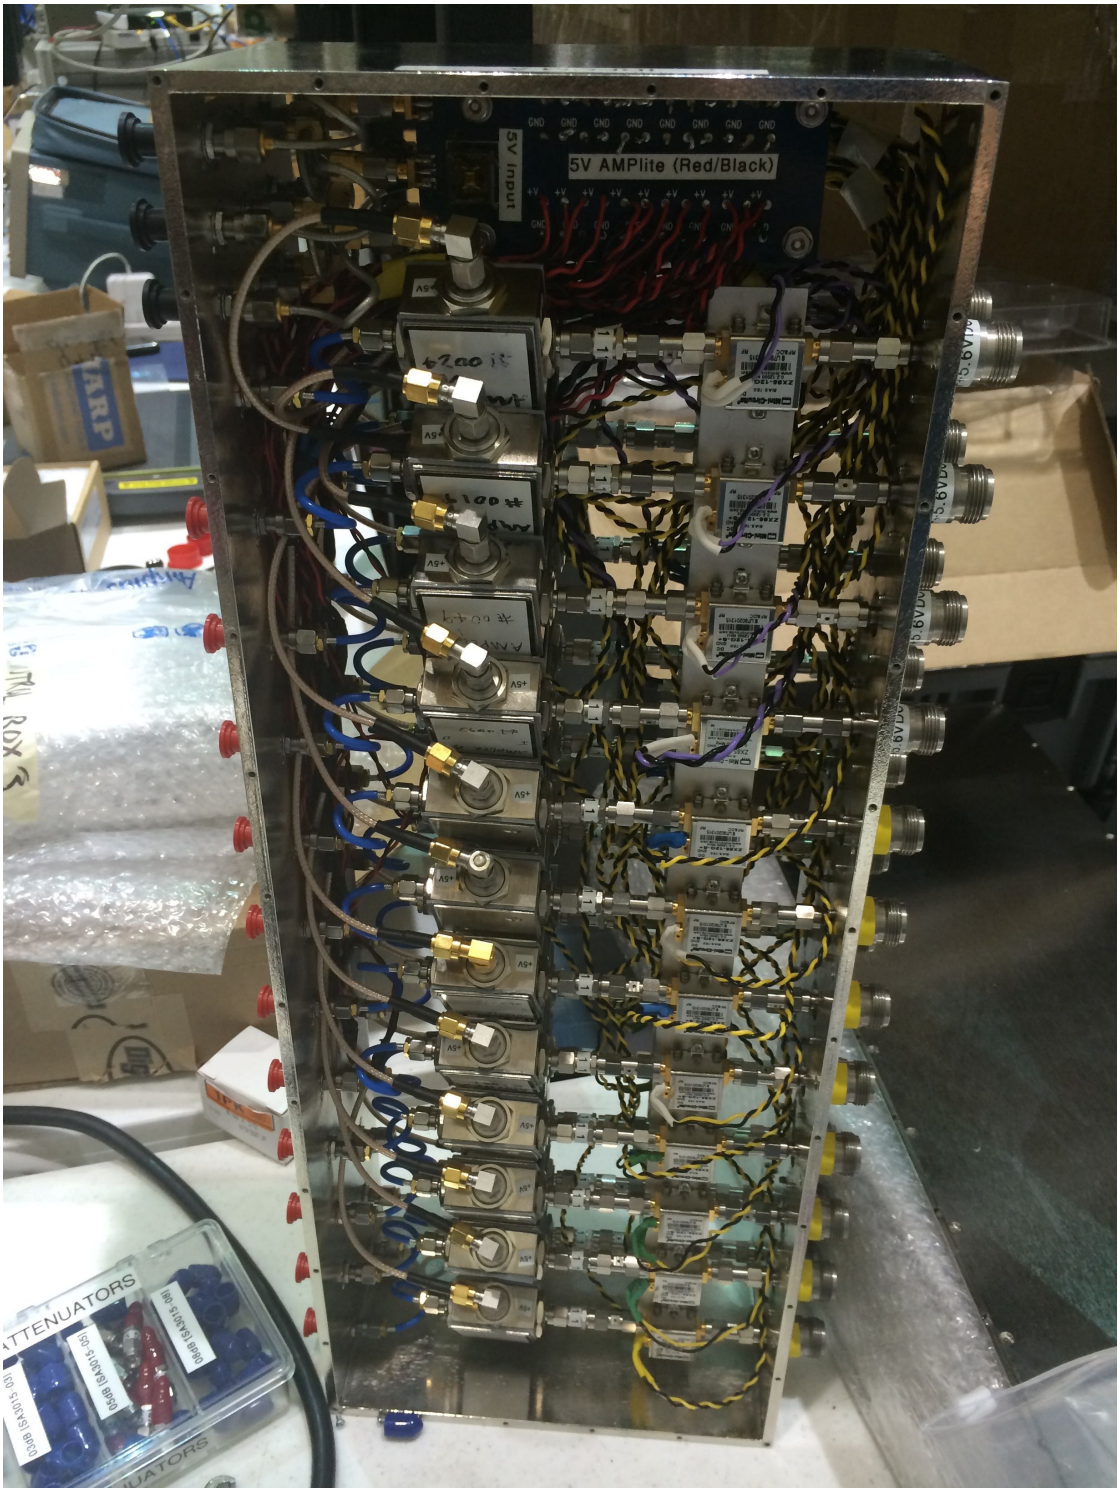

Figure 2.15: An image of the internals of an iRFCM module after assembly

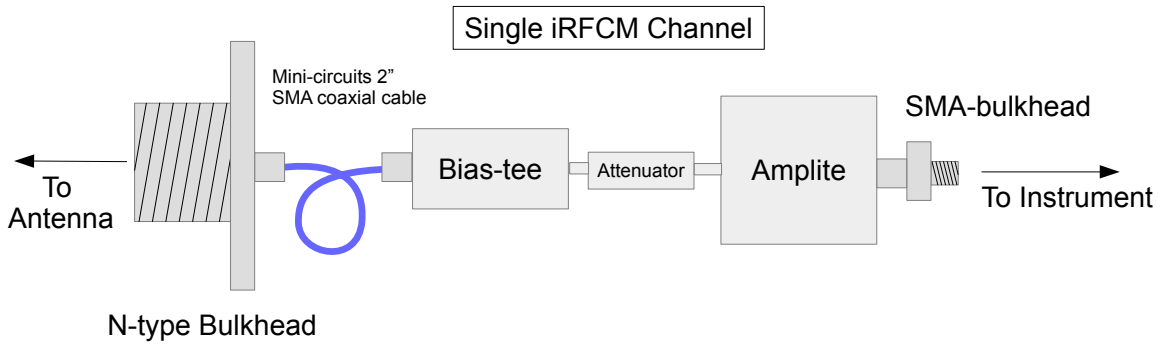

Figure 2.16: A block diagram of a single iRFCM RF channel. The bias-tee, as its name suggests, adds a DC bias to the signal in order to power the pre-amplifier modules.

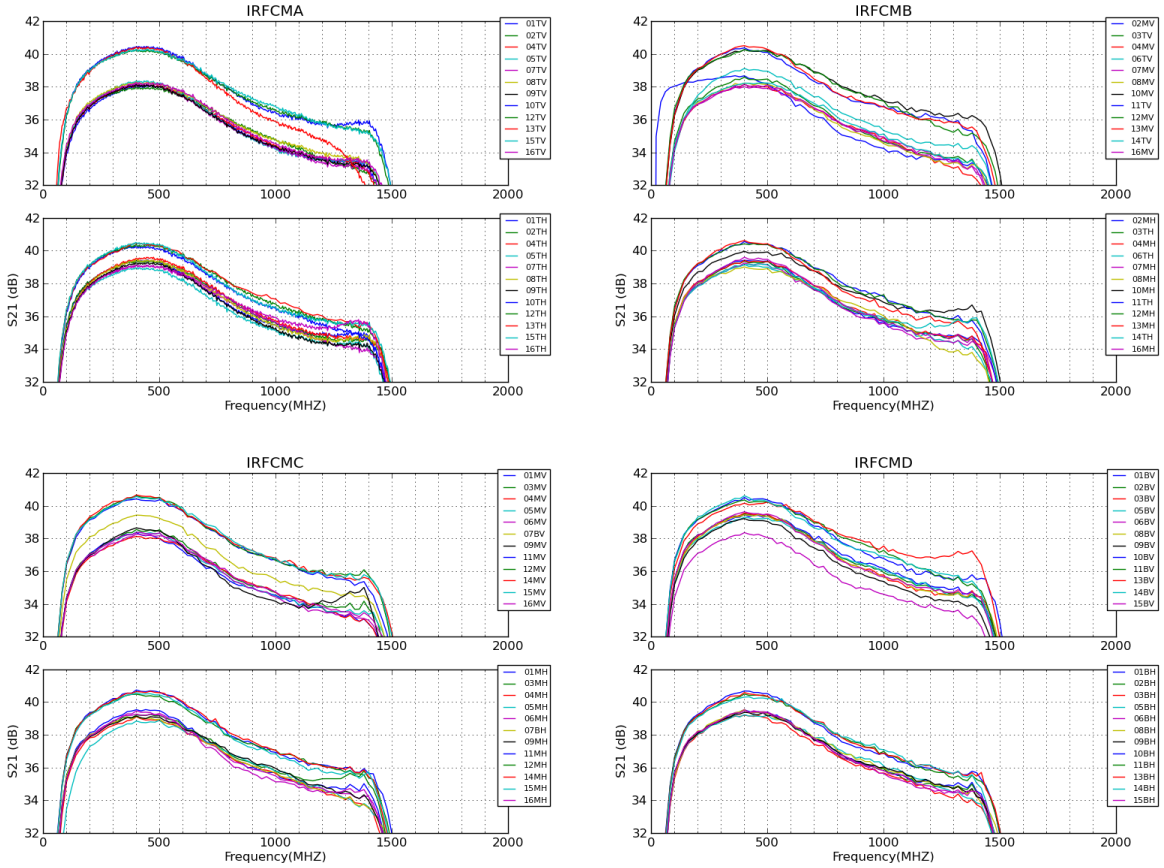

Figure 2.17: Gains for all 96 iRFCM channels. The two distributions are intentional, used to equalize for the differences in gain between the AMPA and DDAMPA electronics. iRFCM channels were paired with specific pre-amplifier modules to equalize total system gain across all channels.

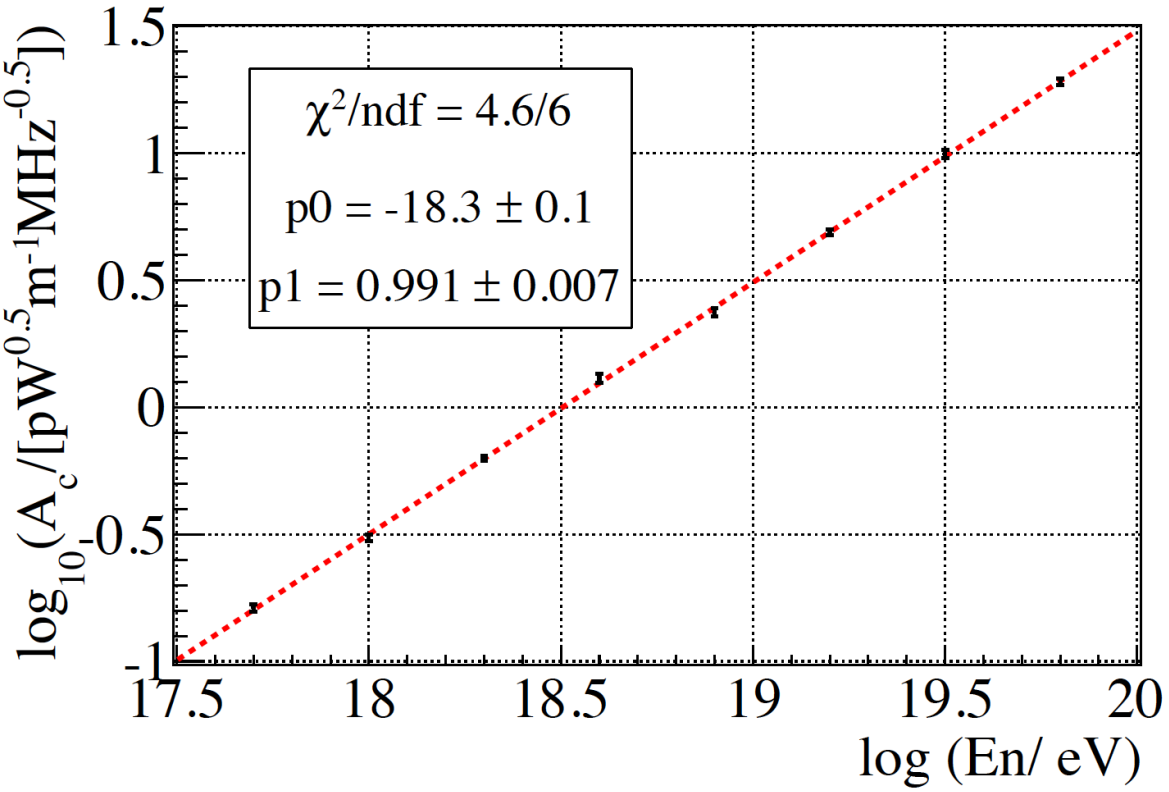

Figure 2.18: The relationship between EAS signal power at the peak coherence angle and primary particle energy from simulations. Note the logarithmic scale and square root which makes the relationship appear linear [7].

through an amplifier chain, it is imperative to reduce any additional noise figure at the beginning of the signal chain, and less important in the subsequent amplifiers. The cascading noise can be calculated using Frii's Formula (Equation 2.5), where T_{total} is the resulting noise temperature of the entire signal chain, T_n is the noise temperature of a specific element, and G_n is the gain magnitude of a specific element.

$$\bar{v}^2 = 4k_bBRT \tag{2.4}$$

$$T_{total} = T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1G_2} + ... + \frac{T_n}{G_1G_2...G_{n-1}} \tag{2.5}$$

From this, one can see that the +38dB of gain from the first stage amplifier reduces any subsequent noise figure from components by a factor of 6309, as shown in Equation 2.4.4.

$$38dB = 10 * log(G)$$
$$G = 10^{38/10}$$
$$G = 6309$$

This large reduction alleviates the requirement for low noise amplifiers in subsequent stages.

2.4.5 Dynamic Range and Gain

Each instrumental system measuring the electromagnetic field operates at a different optimal power level. For example, the digitizer can cover an approximately 2V full dynamic range. An observed voltage with peak to peak amplitude significantly below that level would never exercise the full range of the digitizer and allow bit level digitizer noise to become dominant over the antenna noise background. In contrast, too high an input power signal would become distorted at high peak amplitudes, as near full scale voltages the response of the digitizer becomes nonlinear, or "compressed".

There are three terminal systems connected to the coaxial feed line from the antenna, which will be discussed in detail later: the LABRADOR digitizer, the tunnel diode before the trigger circuit, and a broad spectrum RF power monitor. The required gain of the system can be arrived at by determining their input power range, then calculating what amplification of expected signal is required so that the power sits within that specification. The total system gain is driven by the instrument that requires the highest power, and any instruments that require lower input power are given in-line, flat spectrum, fixed attenuator components. Table 2.1 contains the three terminal components, their dynamic ranges, and the required gain necessary to put observed thermal noise from the antenna into a measurable power region. The justification for components not yet

Component	Maximum Power	Gain Requirement
LABRADOR digitizer	20dBm	80dB
RF Power Monitor	-10dBm	50dB
SHORT and Tunnel Diode	0dBm	60dB

Table 2.1: A description of the required input power and gain requirements for the three terminal instrumentation components of the ANITA-III system

mentioned in this thesis are covered in subsequent sections.

2.5 Digitization

The dynamic electromagnetic field variations incident on the payload are digitized using an array of custom designed fast analog to digital converter (ADC) application specific integrated circuits (ASICs). The ANITA-III instrument utilizes the third iteration of the LABRADOR (Large Analog Bandwidth Recorded And Digital Ordered Readout), the LAB3, designed by Gary Varner[8]. This chip was used on the first three ANITA flights because of its high precision, large dynamic range, Gigahertz of analogy bandwidth, and extremely low power requirements. It is a 12-bit, 2.6GS/s, 256 sample long ADC, yielding a window size of 100ns. Each chip has 9 channels, consisting of 8 RF analog inputs, as well as a 9th global clock channel propagated to all LAB chips for alignment. The SURF (Signal Unit for Radio Frequency) Board consists of 4 LABRADOR chips in order to allow multi-hit buffering, as well as paths for the trigger signal to a Field Programmable Gate Array (FPGA), discussed later.

2.5.1 LABRADOR3 ASIC

The LAB3 converts the analog signals to digital values through means of a sample and hold switched capacitor array read out with a wilkinson clock comparing the stored charge in a storage capacitor against a ramp signal driven by a constant current source. The sampling time base is driven by a square pulse that propagates through a starved transistor array internally to the chip, toggling each capacitor sequentially. This continuous sampling is halted when a trigger is formed and a HOLD is issued to the chip. The LAB3 is Wilkinson ADC, which measures the time to threshold of a constant current voltage ramp and converts the stored charge in each capacitor to a digital value, then read out in parallel to an FPGA for further storage (see Figure 2.19). The time-to-threshold circuit uses a 100MHz clock and a Gray code counter. Due to the dead time related to the digitization process, four LAB3 chips are placed in parallel in order to allow for multi-hit readout capabilities. When a HOLD is issued, the LAB3 chip returns both an estimate of which RCO phase it is in (which due to propagation latency is inverted near the wraparound region), as well as the locations of the "HitBus", or samples that were left in tracking mode (connected to

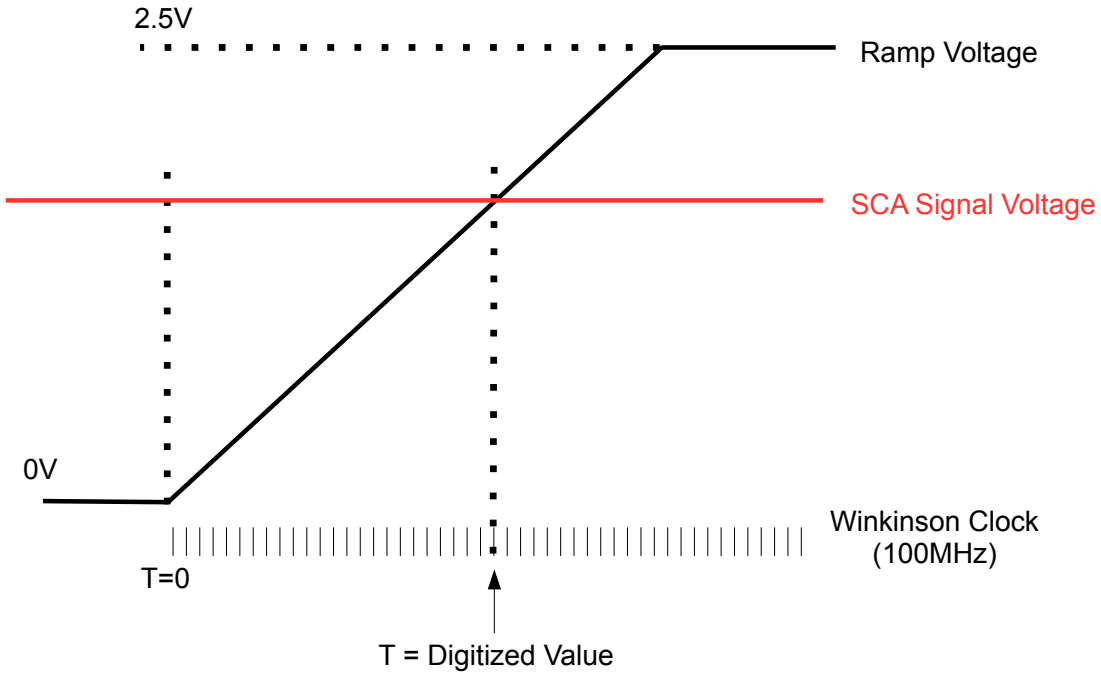

Figure 2.19: A simplified diagram of the LABRADOR digitization technique. A linearly increasing voltage ramp signal is generated by driving a constant current into a set capacitance, which is compared versus the voltage present on each SCA capacitor. At the beginning of the ramp signal, a 100MHz Wilkinson clock is started, which increments a 12 bit counter at every clock period. At the time of the threshold crossing, the value is latched and read out as the digitized value.

the RF input) when sampling was halted. The procedure and results of calibrating this readout is detailed in the Calibration section.

2.5.2 Limitations

After a hold is issued to a chip, the digitization freezes the ring buffer leaving the chip unable to sample until digitization is complete, a process that can take many microseconds. During a readout, no new data can be captured by the LAB3 chip. To minimize the frequency of this dead period, the RF input chain is split into four separate chips which are read out one at a time. This provides a buffer depth of four, which allows the instrument to remain live even while reading out several concurrent events.

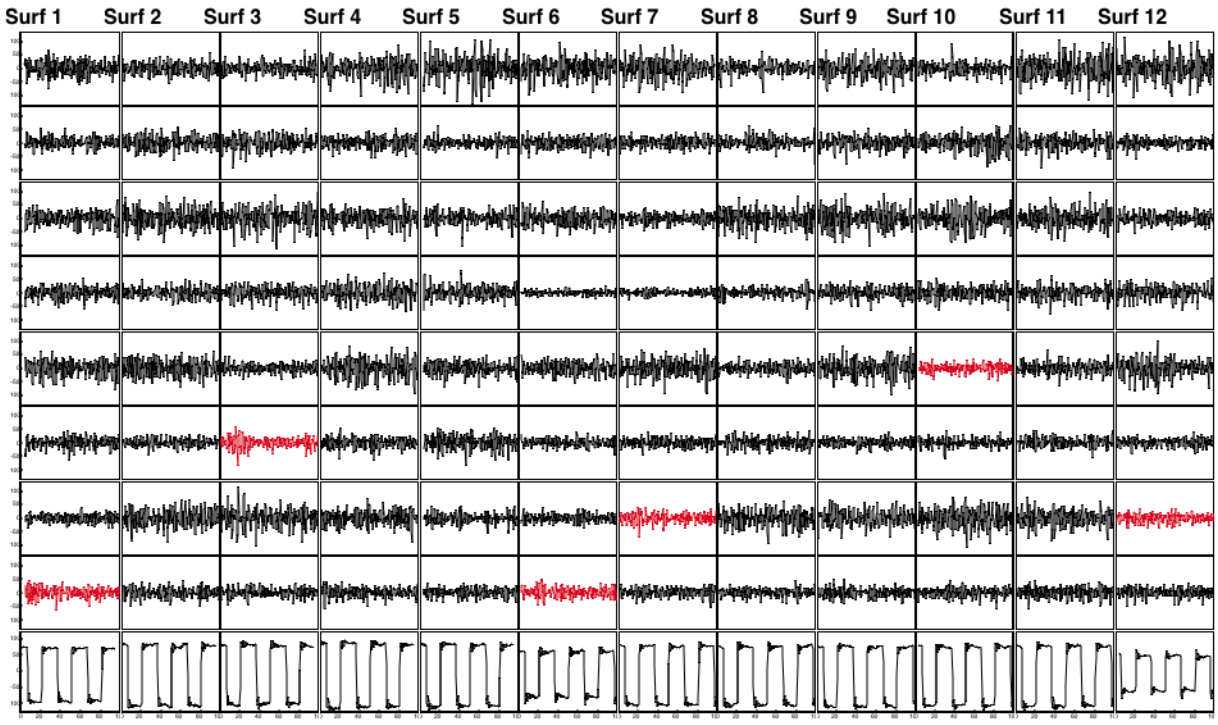

Figure 2.20: An example readout of all 96 channels of the full ANITA-III detector, organized by SURF (columns) and channels (rows). Bottom row is the synchronization clock inserted into all LAB chips. In red are phi sectors that triggered the event. This particular waveform is low SNR, self-triggered, West Antarctic Ice Sheet (WAIS) calibration pulser (event 61326092). Image was made with MagicDisplay data visualization code.

The analog bandwidth of the LAB3B (used in the ANITA-III and ANITAII experiments) did not fully cover the full bandwidth of the antennas and signal chain, which yields a drop-off in sensitivity for high frequency signals. The total system bandwidth is shown in the calibration Chapter, however a measurement of the analog bandwidth of the LAB3B specifically is shown in Figure 2.21

The time between samples is controlled by a charge-starved transistor chain that controls the connection between the sampling capacitor and the input RF signal. Due to process parameter spread in the manufacturing of the ASICs, the timing between subsequent samples is not well controlled. This yields an uneven time sampling that needs to be corrected in calibration of each chip individually. In addition, it leads to an unevenly sampled time domain waveform which introduces a difficult-to-correct frequency response, and requires interpolation between points to restore even sampling. Interpolation is CPU intensive, an issue that makes accomplishing any sort of in-flight analysis more difficult.

Since each LAB chip can only digitize eight channels concurrently, a total of twelve total SURF boards are required to measure all 96 channels of each event simultaneously. The propagation of the readout HOLD signal from the FPGA, and its subsequent latching within each LAB, adds jitter that needs to be corrected in analysis. The ninth channel on each lab observes a single 33.3MHz analog clock signal propagated through the CPCI backplane to each SURF. By measuring the relative phase of this clock between SURFs, this jitter can be corrected. This is discussed further in subsection 3.2.7.

Additionally, the 100ns SCA buffer length limits the minimum full period oscillation observation frequency to 10MHz, which is below the antenna frequency cutoff. Though not expected in any astrophysical channel, the effect of the limited trigger window would also be problematic for an elongated or repetitive signal that would be cut off at its trailing edge.

2.5.3 Impulse Response

The filtering and amplification chain has a dispersive effect on the signal, which requires a calibration of the system impulse response in order to fully understand the incident electromagnetic field. An impulse response is a measure of the phase and gain distortion created by the RF network to an input signal. In the ANITA, a band-limited system, the group delay is not constant as a function of frequency near the poles of the filters, which causes the signal power to be spread temporally. Any comparison between waveforms from multiple channels is dependent on the full complex parameters of the signal, significant variation between the impulse responses of channels will result in a reduced maximum correlation value, regardless of the electromagnetic impulse present at the antennas. The ANITA instrument is designed to be identical and symmetric across all 96 channels, making this effect small and constant. Measurements of this response and calibration are done in the Calibration section, and the effect and deconvolution is detailed in the Analysis section.

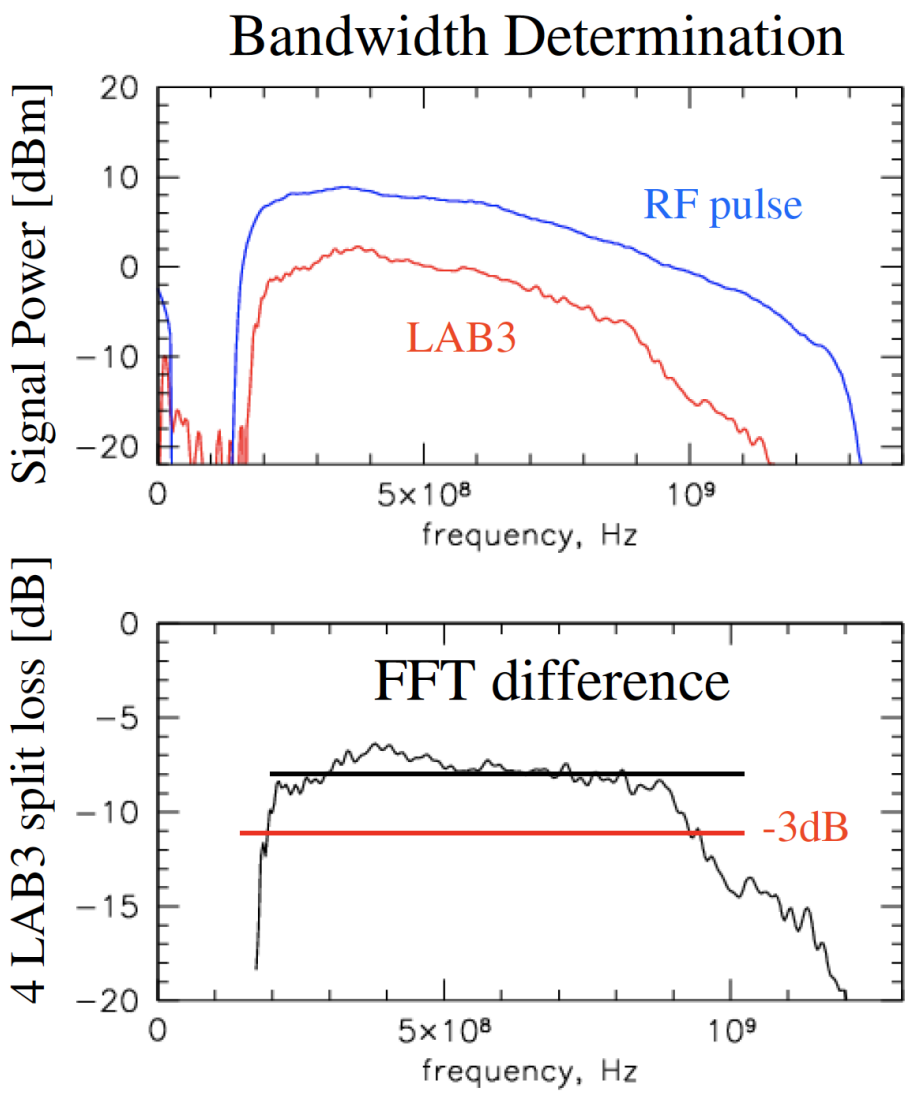

Figure 2.21: Bandwidth measurement of the LAB3 digitizer chip measured through injection of a impulsive signal to a test board. The signal undergoes a 4-way split on the test board to emulate the SURF board. In the top figure, the spectral power of a (blue) RF reference pulse and (red) LAB3 readout pulse are shown. At bottom is the difference, where a -6dB board loss would be expected with perfect coupling(black line). The -3dB point occurs several hundred MHz below the 1.2GHz desired instrument high frequency limit.[8]

2.6 Triggering

As the time domain digitization window is extremely small (100ns is one ten millionth (1e-7) of a second) and ANITA is limited to a 50Hz readout rate, it is necessary to selectively trigger on segments of time that have a high probability of containing a signal event. The physics signal created by a UHE particle interaction within the field of view of, and directed towards, the instrument, exhibits itself as a picosecond duration impulsive electric field propagating as a plane wave. The backgrounds are random incoherent thermal noise, and impulsive or single band carrier wave (CW) anthropogenic sources. As there is no digitizer buffering or analog delay lines, the triggering system must combine together information from several RF channels without full digitization and form a decision quickly, before the waveform is overwritten in the ring buffer of the LABRADOR digitizer. For the ANITA-III system, these triggering decisions were made in the Triggering Unit for Radio Frequency (TURF), which receives trigger information for each channel from FPGA comparator electronics on each SURF through the CPCI backplane passthrough connectors.

2.6.1 SHORT square-law power integrator

A solution to triggering on fast impulsive signals while minimizing incidental background events is to utilize a square-law integrating power detector. An Ezaki diode (or tunnel diode)[42], is a nonlinear semiconductor circuit that, in a specific input power region, rectifies squares and integrates an incoming RF signal through the quantum mechanical effect of electron tunneling. The square law response of the diode can bee seen in Figure 2.22. A capacitor in the diode circuit is capable of taking the extremely short duration impulsive signal, with a pulse width of under 1ns, and distributing its power over a longer time scale. In conjunction, these will increase the voltage signal to noise ratio to allow for detection of weaker signals. The signal from the diode must be amplified to bring the power within the dynamic range of the FPGA comparator circuit. The signal is then transformed from single-ended to differential pair and travels over a high occupancy cable to the SURF board and FPGA. A SURF High Occupancy RF Triggering (SHORT) unit consists the diode, amplifier, and transformer for 8 channels. 12 SHORTS are mounted to the top of the instrument box allowing all 96 channels to contribute to the trigger.

A clocked comparator circuit on the FPGA run at 4ns/cycle latches and propagates a logic signal for any electrical signal from the diode that crosses a certain voltage threshold. The first level L1 triggering rate can be tuned by altering the comparison voltage with an on board Digital to Analog Converter (DAC). With this capability, each channel's L1 rate can be stabilized, despite time and heading dependent power fluctuations, through a PID loop. Comparing the timing of these L1 triggers between antennas with similar pointing directions allows a decrease in overall trigger rate dictated by combinatorics.

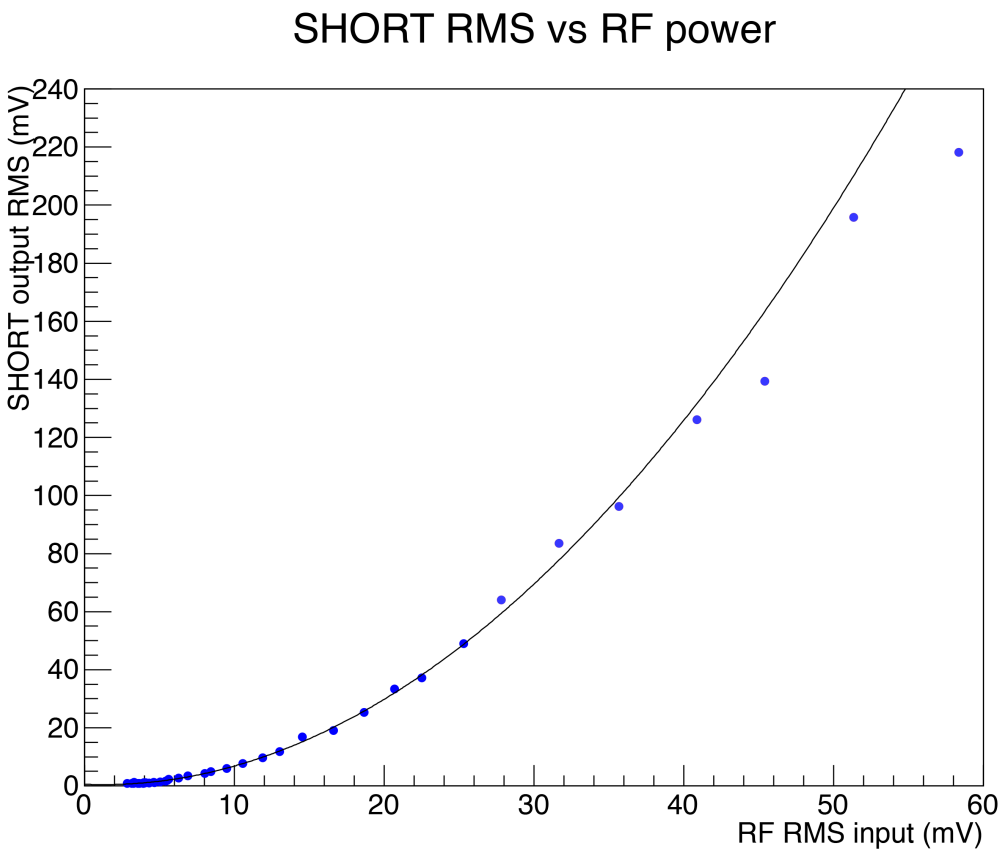

Figure 2.22: Measurements of the response of the SHORT tunnel diode "square law" RF power detector taken in Palestine TX in 2014. Plot courtesy of Katie Mulrey.

2.6.2 Trigger Hierarchy Overview

The triggering system is divided into several sub-triggers that layer hierarchically to arrive at a global trigger decision. These range from the individual power law thresholding to a time-windowed combinatoric trigger between multiple channels. In ANITA-III, the trigger is split into a vertical (Vpol) and Horizontal (Hpol) trigger system, as the two differing particle shower sources (in ice neutrino interactions vs geomagnetically dominated air showers) will have mostly orthogonal observed polarizations. These two trigger paths are identical and have equal weighting in the final global trigger decision. A diagram of the trigger hierarchy can be seen in Figure 2.23.

The first trigger level (Level 1 or L1), is the comparator circuit on the FPGA between the output of the SHORT and a DAC. It is possible to adjust each channel individually to give an even weighting in the ultimate trigger decision. This is accomplished by measuring the number of threshold crossing counts per time observed by each comparator circuit, called the pseudo-scalar rate, which is used as the input to a software Proportional Integral Derivative (PID) servo loop that controls the comparator voltage.

The next trigger level (L2) is a phi sector specific time dependent windowing function that combines the L1 triggers from all three channels in a single phi sector. A real plane wave signal incident on a phi sector will have a causal time separation between the physically separated antennas, while incidental noise threshold crossings will be uncorrelated in time. To fire an L1 trigger, two causal signals are required to fall within this window. The window is achieved by synchronously firing different length one shot pulses and logically combining them. The bottom ring opens a 16ns window, the middle a 12ns window, while the top only opens a 4ns window. These timing offsets are used a a quickly determined metric for further discriminating on noise, and is limited primarily by the timing precision provided by the FPGA trigger electronics.

The final trigger (L3) is a requirement that two neighboring phi sectors decide a L2 trigger within an 8ns window. This leads to the ultimate requirement that four out of six like-polarized antenna channels with an overlapping field of view observe a transient power fluctuation. There are two L3 triggers, one for each polarization, either of which firing issues a HOLD to a LAB3, beginning the digitization process.

Additional information about the trigger is detailed below.

2.6.3 L1 Tiggering efficiency and quality

The trigger must make a trade off between efficiency and quality. Efficiency is defined as the ratio of real signals selected for digitization over the total number of signals incident on the detector, and quality is the ratio of the number of true signal events over total selected events. A perfect detector would have an efficiency of 1 and a quality of 1. However in reality these values are both a function of the total maximum readout rate for the detector, which places a maximum limit on the number of events capable of being selected. As an example, if the trigger constantly selected all moments

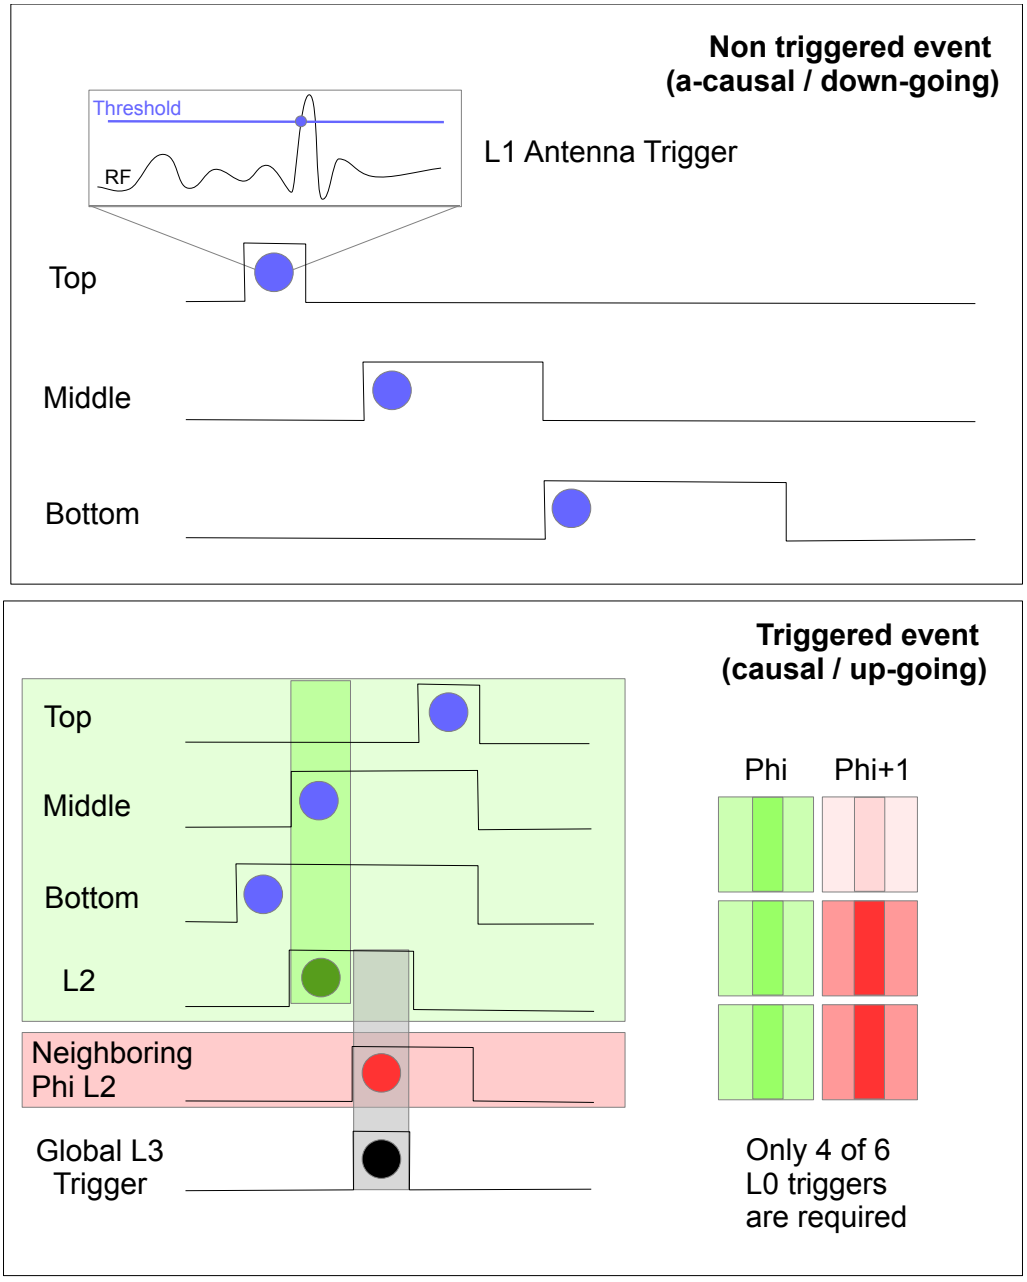

Figure 2.23: A depiction of the trigger hierarchy for a non-triggered (top) and triggered (bottom) series of L0 triggers. The blue dots denote an L0 trigger, and the green box represents the time of a L1 trigger. Note the 4ns top ring window, the 12ns middle ring window, and the 16ns bottom ring window. Two neighboring phi sectors with L2 triggers within a 8ns window cause a global L3 trigger (black). Note also that, though not the case of depicted event, only two of the three rings are required to be in coincidence leading to a 4 out of 6 channel L1 trigger in two neighboring phi sectors

in time it would be perfectly efficient, but the quality of the signal would be abysmal. Inversely, if the trigger was set to have perfect quality and only select extremely high power impulsive events, it would have poor efficiency for low Signal-to-Noise Ratio (SNR) signals. The efficiency of the detector can be statistically measured by injecting signals of varying SNR at a known rate and observing the ratio of triggered events to total events injected. The resulting characteristically shaped curve is shown in Figure 2.24. Using the peak-to-peak amplitude divided by two over the RMS of the noise to measure SNR, the 50% efficiency point (at SNR=3.5). The 50% efficiency point is used as the defining metric for the efficiency of the system. This point will be directly related to the allowable accidental noise trigger rate, which was decided to be set at a rate where the payload would be "dead" for approximately 5% of the time.

2.6.4 L1 Optimization

The only optimization parameter accessable for the tunnel diode L1 trigger rests with the input power to the tunnel diode circuit. The inverse-resistance region of the Esaki diode, where a decrease in voltage yields an increase in current flow, has a square law response. This occurs at an extremely low power level, far below the minimum quantization level of the LAB3, and thus the power must be attenuated in the signal chain after the split between the trigger and digitizer paths. If this attenuation is too great, and the signal power is reduced to a small fraction of the dynamic range of the DAC, changing the trigger threshold setting a single bit for a specific channel would cause a large jump in the overall trigger rate of the system. This could possibly allow a single channel to dominate (or be excluded from) the global triggers. However if the attenuation is too small, and the diode is pushed out of the inverse-resistance region, the response of the diode becomes linear, reducing its signal discrimination potential.

2.6.5 L2 Trigger Window Delay Limitations

As the physical offsets between the rings of antennas exceeds one 4ns FPGA clock period for most elevation angles of interest, it is possible to require a timing separation between the arrival of pulses between the separate antennas. The expected delays between rings as a function of elevation angle is shown in Figure 2.25. This modification, made between the ANITA-II and ANITA-III flights, decreases the incidental rate of the trigger, increasing the quality, while maintaining a low signal rate and high efficiency. The requirement for an L2 trigger is that two windowing pulses overlaps (see Figure 2.23). This has the side effect of allowing the bottom ring of antennas to have a higher weight in the trigger. A noise L1 trigger generates a pulse with four times the time weight as one from the top ring, and is thus more likely to trigger the system.

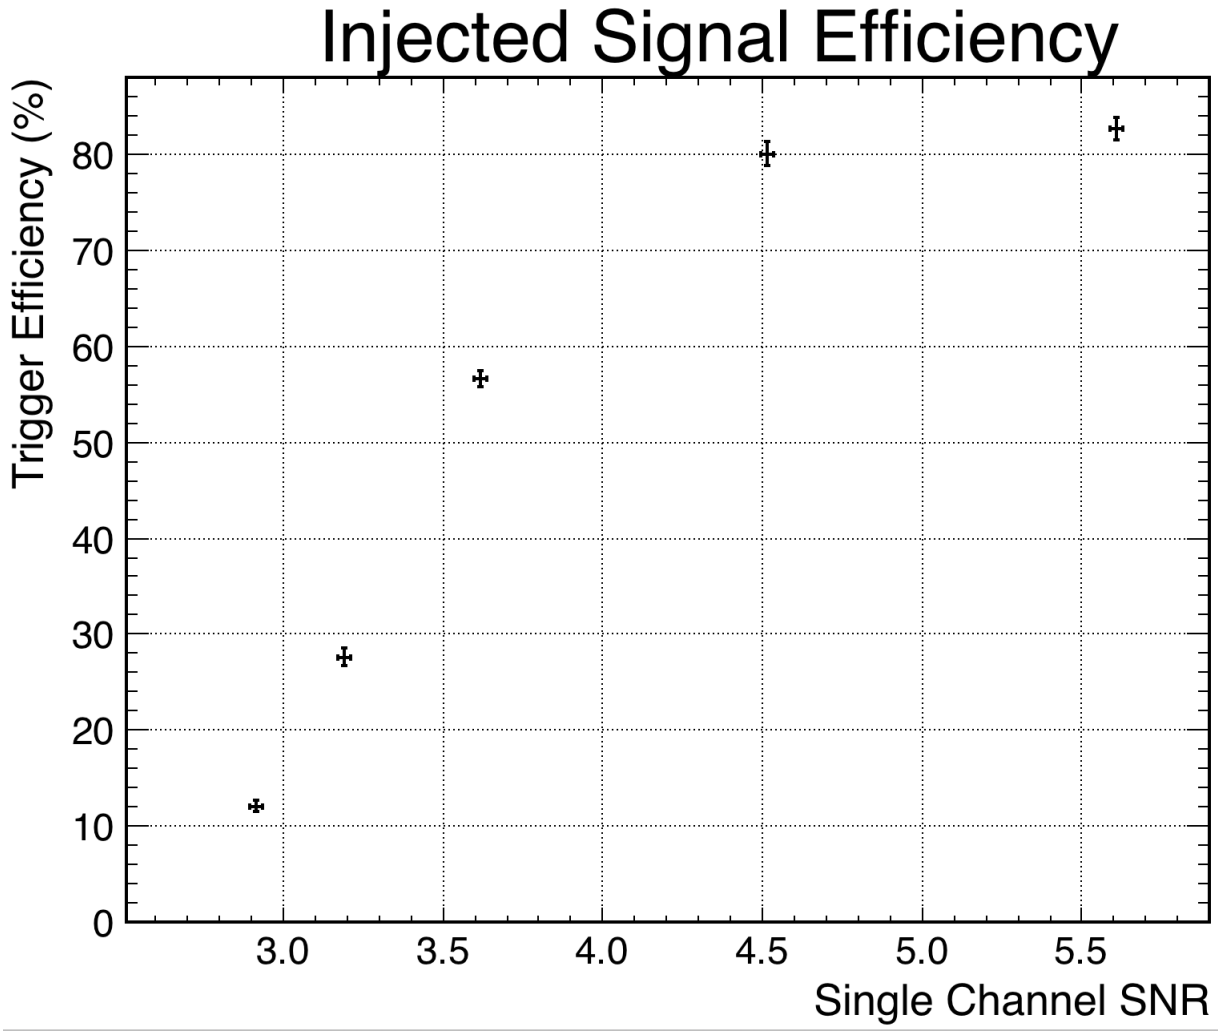

Figure 2.24: The payload global trigger efficiency for injected impulsive signals, as measured in Palestine in the summer of 2014. Identical signals, each with the same SNR, were injected at a GPS stabilized 1Hz into all 6 channels in two phi sectors that share a polarization. The signals pass through the entire signal chain and efficiency was calculated by averaging trigger rate. The L1 PID rate goals for this measurement were 100kHz. The 50% efficiency point occurs at an SNR of roughly 3.5. SNR is measured as averaged peak to peak signal voltage over noise RMS on a seventh signal propagated through an ANITA signal chain but measured on a laboratory oscilloscope.

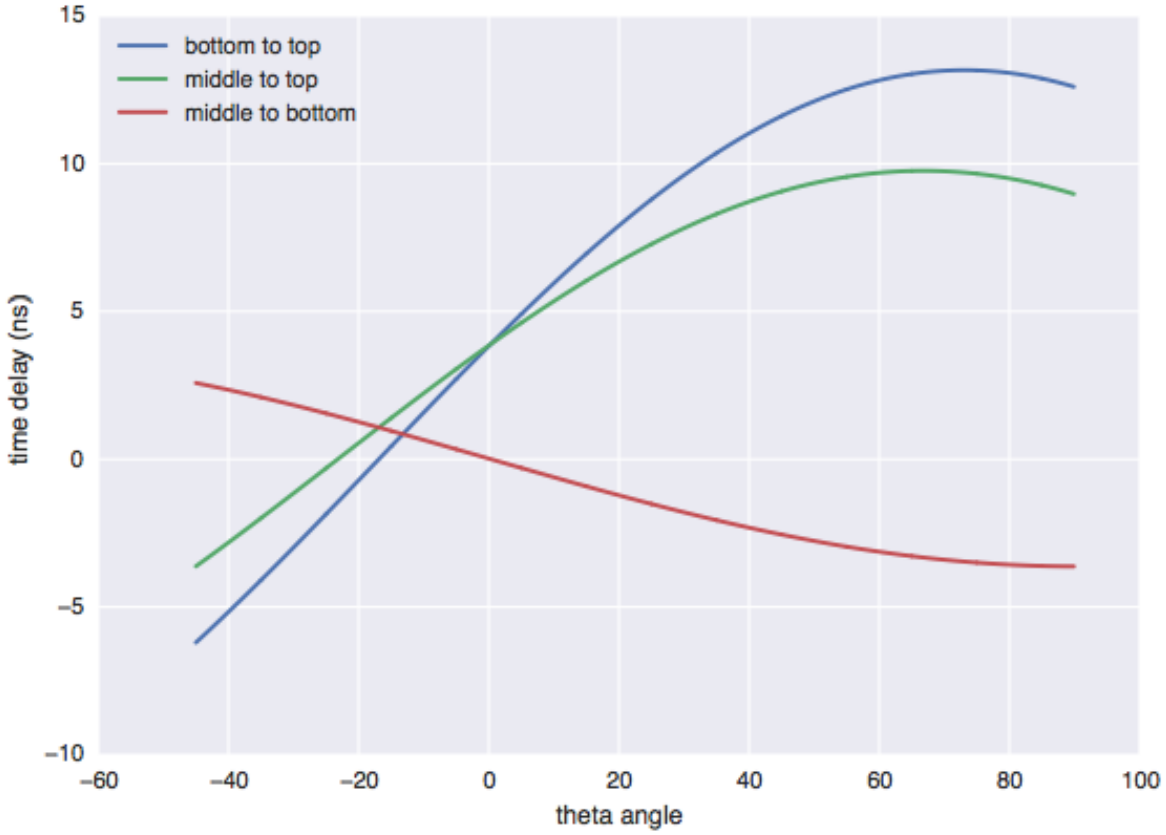

Figure 2.25: Relative propagation delays between antennas for an incident plane wave as a function of elevation angle. Delays were derived using the antenna positions measured using photogrammetry. These delays are used to determine the L2 trigger window locations.

2.6.6 Phi Sector Masking

In the ANITA-I flight, it was discovered that noise sources on the continent and in the sky had the capability to dominate the trigger rate and effectively blind the entire payload despite a large fraction of the payload observing a thermal environment. To alleviate this issue, a system where an antenna or collection (phi sector) of antennas experiences a relative increase in trigger rate in comparison to the rest of the payload, those affected phi sectors will be excluded from participating in forming the global system trigger for a user-adjustable time frame. This allows the instrument to continue observing quiet areas of the continent and increasing the total livetime of the system. There are additional nuances to this subsystem, including the possibility of very high power noise (or signal) events continuing to leak into the back and side lobes of the antennas from non-masked phi sectors, that need to be taken into account when making a measurement of the total instrumented area for a flight. During the ANITA-III flight, several unexpected in-band satellite CW signals caused the phi masking to be heavily utilized, which is discussed further in section 4.4.

2.6.7 Unbiased triggers

Any waveform readout caused by a global trigger has the systematic bias that it passed the trigger selection parameters. To measure the thermal noise environment during the flight, two forced triggers that were uncorrelated to the external signals were employed. These were a software trigger (also know as the soft-trig) that was issued by the CPU at a frequency of 1Hz, as well as a GPS synced 1Hz trigger controlled by the Pulse-Per-Second (PPS) output of the G12 GPS unit.

All events captured by the payload were marked with a trigType integer that described the source of their trigger decision. The GPS and soft-trig signals can be easily selected in this way by excluding RF triggers.

2.7 RF Power Monitor

An additional measurement of the RF signal is done by a broad band radio frequency power detector circuit. The fast waveform data only provides 200ns of non-triggered, unbiased noise data per second (a 1Hz software trigger and a 1Hz GPS trigger), so a monitor that digitized a larger fractional percentage of time was desired. This was accomplished on the SURF board using a commercial RF power monitor IC that converted the RMS voltage of the input signal into a log-magnitude proportional DC output that was digitized by an ADC and read out with the housekeeping diagnostic data.

2.8 GPS and orientation sensors

The payload is attached to a freely floating balloon with no control over rotational orientation or geographic location. This requires that the attitude be constantly monitored. Multiple GPS systems, including two co-located ADU5 GPS heading receivers in conjunction with a single G12 GPS system, recorded the location of the payload throughout the flight. The nine total GPS antennas, read out once per second, provide an absolute reference for the location and orientation of the payload with a resolution of a fraction of a degree.

2.8.1 Magnetometer

In addition to the heading information provided by the GPS systems, a magnetometer was attached to the payload and used to measure the vector components of the magnetic field throughout the flight. This system, coupled with a model for the geomagnetic field, could be used as a functional compass, assuming an accurate position, that would corroborate the pointing information returned by the GPS systems. As the GPS system remained functional throughout the flight, the major contribution from the magnetometer is as a validation of geomagnetic models in comparison to the GPS attitude measurement.

2.8.2 Sun Sensors

Many space-borne experiments that require very high pointing and location precision are above the altitude where GPS begins losing accuracy and precision. Their nominal orientation information is provided through the use of star sensors. ANITA's flight path over Antarctica during the austral summer, when there is continual daylight, is able to use just one star to obtain accurate heading information. With an accurate time and knowledge of the Earth's rotation around the Sun, one can use the four mounted angular sun sensors to determine the heading information for regions of time when the GPS is unavailable or unreliable. As the sun is a single source, the full attitude of the payload can not be determined with this measurement alone. However for brief GPS outages, where payload location does not change dramatically, sun sensors can be used to reconstruct payload heading, which has a much faster time scale of change compared to pitch and roll. As the GPS system remained functional throughout the flight, the sun sensor data was not used.

2.9 CPU, CPCI, and Data Readout

The custom LABRADOR digitizer ASICs and various control and command FPGAs are located on custom designed Printed Circuit Boards (PCBs) connected to a central processing unit (CPU) over a compact Peripheral Component Interconnect (cPCI) bus interface. An image of the conductively cooled cPCI crate, in which all digitizing and readout electronics were housed, is visible in Figure

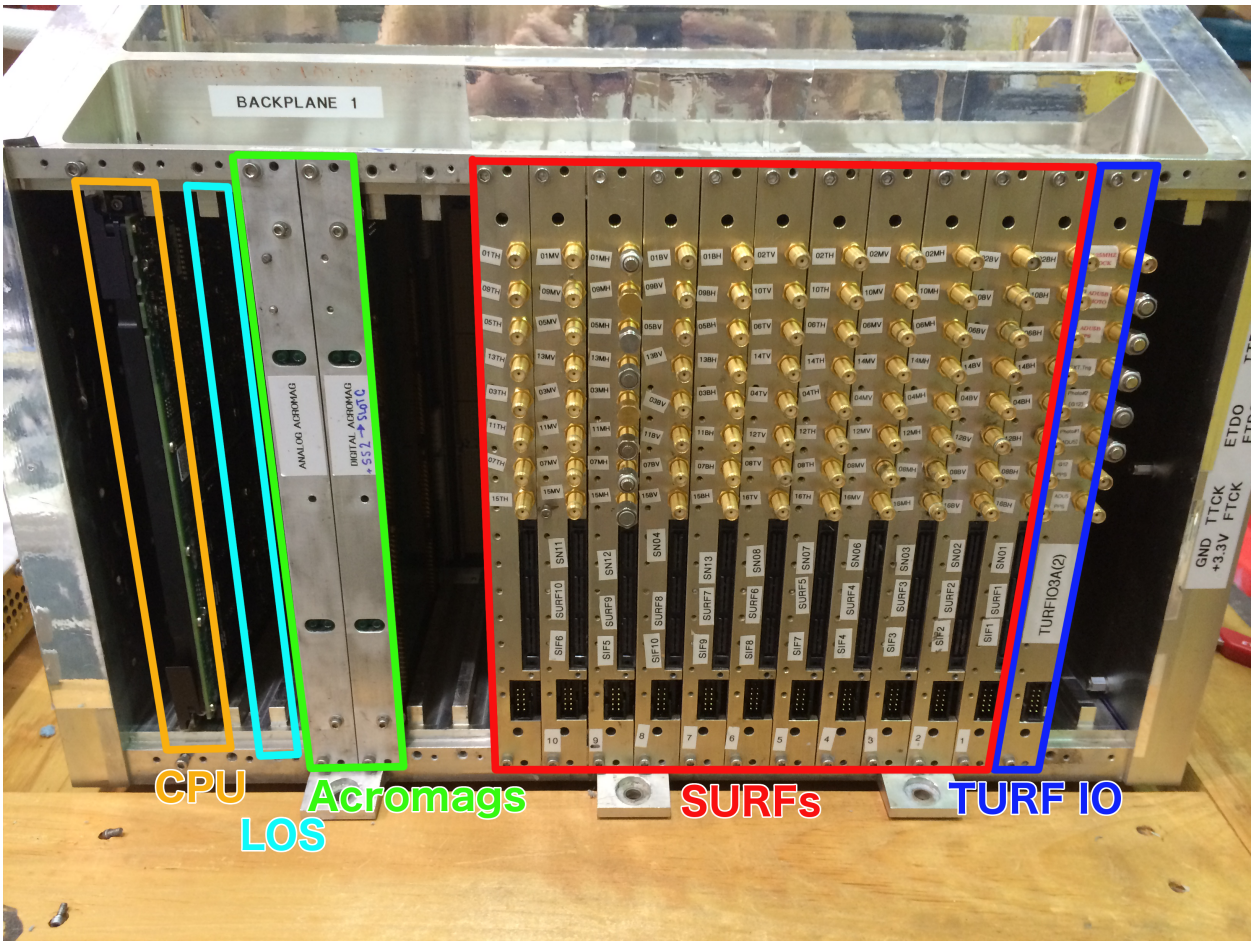

Figure 2.26: An annotated image of the cPCI crate prior to integration into the ANITA-III instrument box. Labeled positions from left to right: the CPU, LOS board, Acromag boards, SURF boards, and TURFIO board. Additional protective faceplates have yet to be installed in image.

2.26. The Triggering Unit for Radio Frequency (TURF) board, which communicates with the SURFs to form the global trigger and propagate hold signals, is mounted to the back of this crate, connecting to the front mounted boards through cPCI Rear Transition Modules (RTMs). Present within the cPCI crate is the CPU, a Line Of Sight (LOS) telemetry board, a digital and analog Acromag sensor readout board for reading various diagnostic subsystems, the 8 SURF boards, and a TURF Input/Output (TURFIO) board, which provides access to the TURF from the front panel.

2.10 Data Storage and Telemetry

An important piece of any experiment is the storage of the measurements for later analysis. The rate at which data is taken during flight is designed to maximize the throughput of the electronics used to read out and store it. It is not feasible to downlink the entirety of the stored information

during flight due to limited telemetry bandwidth. This necessitates the recovery of the payload, or at least the storage vaults, immediately after the flight is terminated.

2.10.1 Redundant data storage

The storage requirements for the flight were met by three separate digital storage subsystems, each employing a different method to ensure the maximum likelihood of full data recovery. This redundancy ensured that even in the event of a single failure the data from the flight would be preserved. The first storage method was two helium-filled conventional spinning hard disk drives, with data written simultaneously and identically by the flight CPU over SATA. The second method, named after National Taiwan University (NTU) who designed and assembled it, was a separately-housed array of six Solid State Drives (SSDs), written to by a dedicated single board computer which received the flight data over an Ethernet link. The third storage method, which ended up being flown in an inactive state, was the Redundant Independent Flash Repository for Antarctic Flight (RIFRAF), a microcontroller-steered multiplexed array of 48 128 GB commercial flash drives visible to the flight CPU over a USB interface. During flight, a failure of the Ethernet link lead to the loss of the SSD array subsystem, leaving only the He Drive storage format intact. Both drives survived the flight however, and two identical copies of data were recovered. In ANITA-III, the NTU device was a standalone CPU that connects with the flight CPU over ethernet and connected to six 1TB SSDs. When one filled up, it would flip to the next. It stopped communicating roughly a third of the way into the flight (probably because the ethernet power cable wiggled loose and all ethernet connectivity ceased), and was not needed due to the success of the He drive system.

2.10.2 Telemetry

Downlink of data taken during flight and uplink of commands to change flight configuration are transmitted to and from ground receivers over several systems. In order of decreasing data trasmission speed, they are the Line Of Sight (LOS) transmitter, an Iridium Pilot® Openport UDP link, a NASA Tracking and Data Relay Satellite System (TDRSS), and an Iridium® low rate downlink. These systems allow in-flight diagnostics of the ANITA-III instrument, as well as an ability to make pre-defined alterations to instrument control configurations, depending on need. Shortly after flight, the Iridium Pilot® Openport link failed for unknown reasons. However, the remaining systems were successful in allowing flight operators to debug and fix several issues that cropped up during the flight.

2.10.3 GPU prioritization

The data limitations of the various telemetry downlink systems allow only a fraction of the total data rate to be transmitted. This limitation motivates a prioritization system in order to only send down events that pass a limited set of post-digitization analysis cuts. This would allow a limited set of analyzable signals in the event of a catastrophic system failure or non-recovery of the payload. The two major figures of merit for discriminating between incidental thermal noise events and impulsive signal-like events is the normalized peak height of an interferometric pointing map, and the peak of the Hilbert envelope of the coherently-summed waveform for the respective map peak incidence angle. These time-consuming, computationally-intensive, physical-baseline-dependent, multiple channel correlation processes are nominally done offline in the analysis phase of the experiment, and is likewise further discussed in subsection 4.5.2. To accomplish the task in real time, as events are being logged by the instrument, a Graphical Processing Unit (GPU) linked with the flight CPU, was utilized to parallelize the computations and reduce any specific event readout's required processing time to that of the minimum instrument readout time. These values were then used to determine a priority for each event and assigned a priority for telemetering. The main work on this system was done by Ben Strutt of University College London, and details on its operation and results can be found in his doctoral thesis [9]. Since the instrument was recovered, these computations are re-done for this thesis, and the priority value was not used.

2.11 Balloon and Flight

The ANITA instrument is suspended from a NASA zero-pressure long-duration high altitude balloon over the continent of Antarctica. The partially inflated balloon and launch team immediately before launch of the ANITA-III payload can be seen in Figure 2.27. The 34 million cubic liter balloon, when filled with Helium gas, has a nominal float altitude of 130kft (40km) and has been rated to survive for up to 60 days. As the balloon ascends, it expands as the gas equilibrates to the surrounding atmosphere, which decreases in density with altitude. The balloon can be seen, partially expanded, in Figure 2.28. The balloon floats freely after launch, with no aeronautical controls with the exception of a ballast hopper which can be dumped to increase free lift as needed. The circumpolar winds present in Antarctica provide a nominally circular trajectory around the continent. The requirement of a balloon recovery for a complete physics dataset necessitates terminating the flight when the balloon heads northwards towards the ocean.

The ANITA-III flight launched on December 17th 2014 at 19:23 UTC and flew for 22 days before being terminated near Australia's Davis station as it moved close to the coast. The flight path of the payload is shown in Figure 2.29.

Figure 2.27: Image of the launch vehicles, the balloon, and the ANITA-III instrument immediately proceeding launch on December 17th, 2013. On the left of the image is the LDB launch vehicle, named The BOSS, suspending ANITA. Attached to ANITA and the launch vehicle laying across the ice is the termination package, including a parachute for decent, and balloon. The tanks in the background are transportation canisters for the Helium gas, which is filling the balloon through the transparent hoses attached on the sides.

Figure 2.28: Image of the ANITA-III payload approximately 5 hours after launch. Image was captured with aid of a Celestron telescope and captured from the LDB facility.

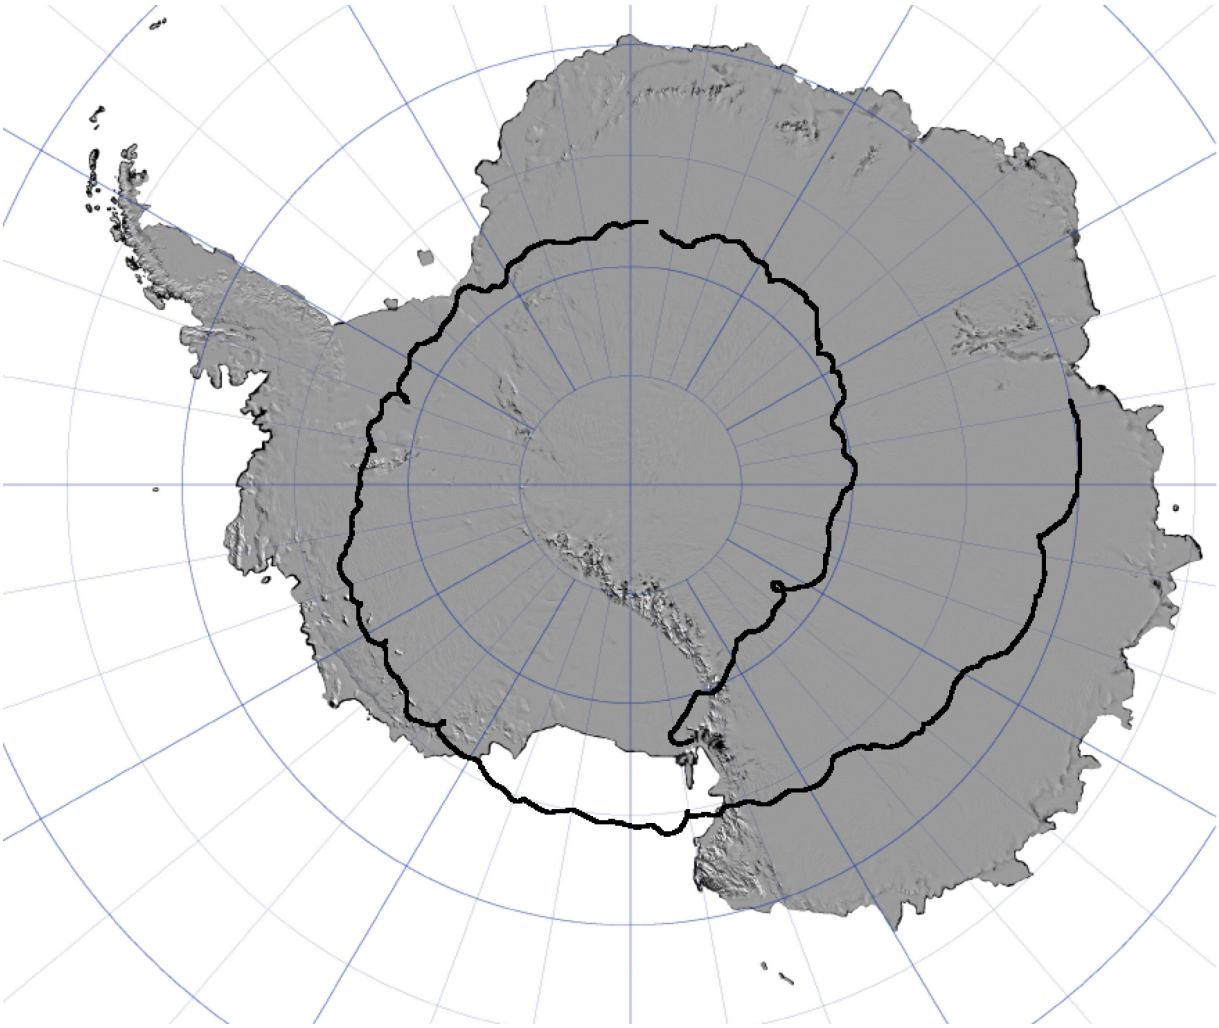

Figure 2.29: Flight path of the ANITA-III instrument as recorded by the onboard GPS. The gap at the "top" of the arc (grid north) corresponds to a 11 hour downtime experienced by the payload

CHAPTER 3
INSTRUMENT CALIBRATION

The ANITAIII digitizing electronics require calibration in order to relate recorded digital signals to an incident electromagnetic field, and possible physics event, being measured. There are several major considerations that must be taken into account in order to relate the recorded values to physical electromagnetic phenomena. The LAB3 reports values in ADC counts, which need to be scaled to represent voltage at the input to the chip. The sampling time base must be measured to determine its non-uniformity and correct for it after digitization. The gain and phase effect of the signal chain impulse response must be deconvolved to determine the absolute electromagnetic field present at the detector. Precision spacing of the antenna phase centers must be determined for accurate interferometric pointing and mapping, and for coherently summing waveforms. The in-flight trigger efficiency of the system must be determined. Data taken both on the ground and in-flight is used to determine the calibration of the ANITAIII instrument. The following chapter details the steps taken and results of calibrating these aspects of the ANITAIII instrument.

3.1 LABRADOR voltage calibration

The LAB3 ASIC, as described in section 2.5, uses an SCA to store many sequential snapshots of the charge on a capacitor from a time-varying electrical potential arriving at the chip. This captured analog charge is then digitized (an analog-to-digital conversion, or ADC) utilizing a Wilkinson ADC, which compares the voltage to a constant current source voltage ramp signal that stops a Gray code counter at the time of comparator activation (see Figure 3.1). The digital values read out by the acquisition software are thus a measurement of the time between when the ramp signal begins, and when its voltage passes that of the voltage of the stored charge in that particular SCA bin. This ADC count must be converted into a voltage.

3.1.1 Chip to Chip Variation

Each chip suffers from process parameter spread introduced in the manufacturing process. Due to this, each chip must be calibrated independently. The write pointer strobe propagation which generates the time base, the RCO phase dependence, as well as the ADC-to-voltage values must all be independently measured for each chip.

3.1.2 Pedestal Subtraction

The signal propagated through the ANITAIII RF system is AC-coupled by design, and as such, must have a DC offset reintroduced to it before being coupled to the SCA within the LABRADOR ASIC. The ramp voltage and comparator internal to the chip operate only at positive voltages

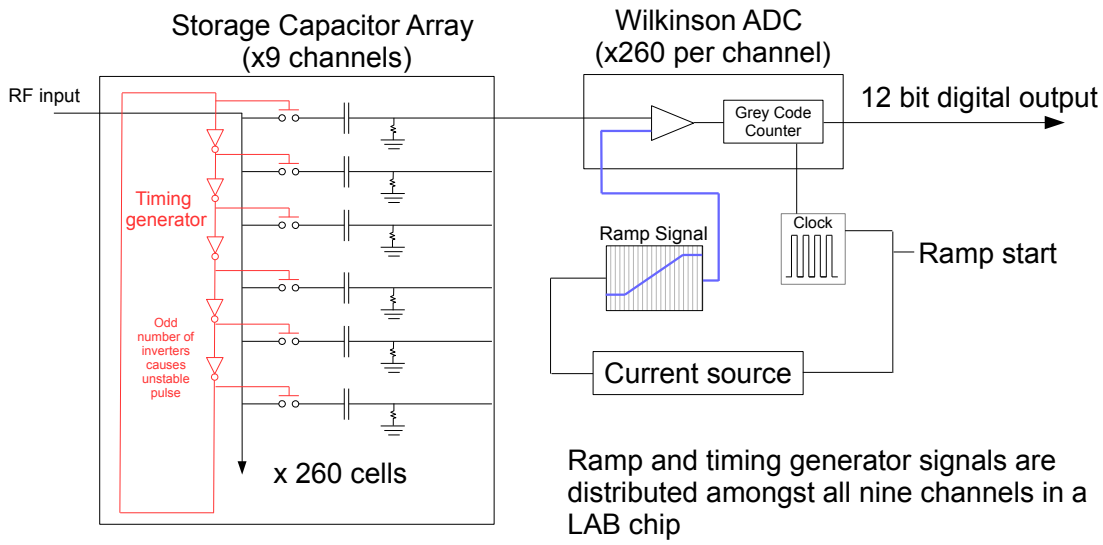

Figure 3.1: Simplified diagram of the digitization chain for the LABRADOR3 digitizer

(0 to 2.5V), and so, to get a measurement of both positive and negative voltage fluctuations, the signal must nominally be set at mid-range of the ramp signal, or about 1.25V. This is accomplished through a fixed voltage regulator external to the LABRADOR on the SURF board. This "pedestal" voltage, visible in the data as a constant ADC count offset, must be later subtracted to correctly zero mean the signal.

A "pedestal run," in which large numbers of samples were measured and averaged to determine a mean value per storage cell, was done at the beginning of the flight and stored to the redundant storage drives. Since generating the pedestals involves averaging together multiple values, their resolution exceeds that of the digitizer itself. The final pedestal values are stored as 11 bit integers. These pedestals are then subtracted immediately after being read by the acquisition software daemon(Acqd) before being stored, and are thus included in the raw data. The pedestals used in ANITAIII are plotted in Figure 3.2

3.1.3 Least significant bit masking

Though the LABRADOR chip reads out a 12 bit digitized ADC value, the value written to disk is only 11 bits. The least significant two bits of the LAB are known to be systematically biased and are thus masked. The entire value is then bit shifted downward once, leaving a single continuously masked least significant bit and 10 bits of effective dynamic range. The final ADC value stored by the acquisition software is an 11 bit signed integer, with the least significant bit determined by the

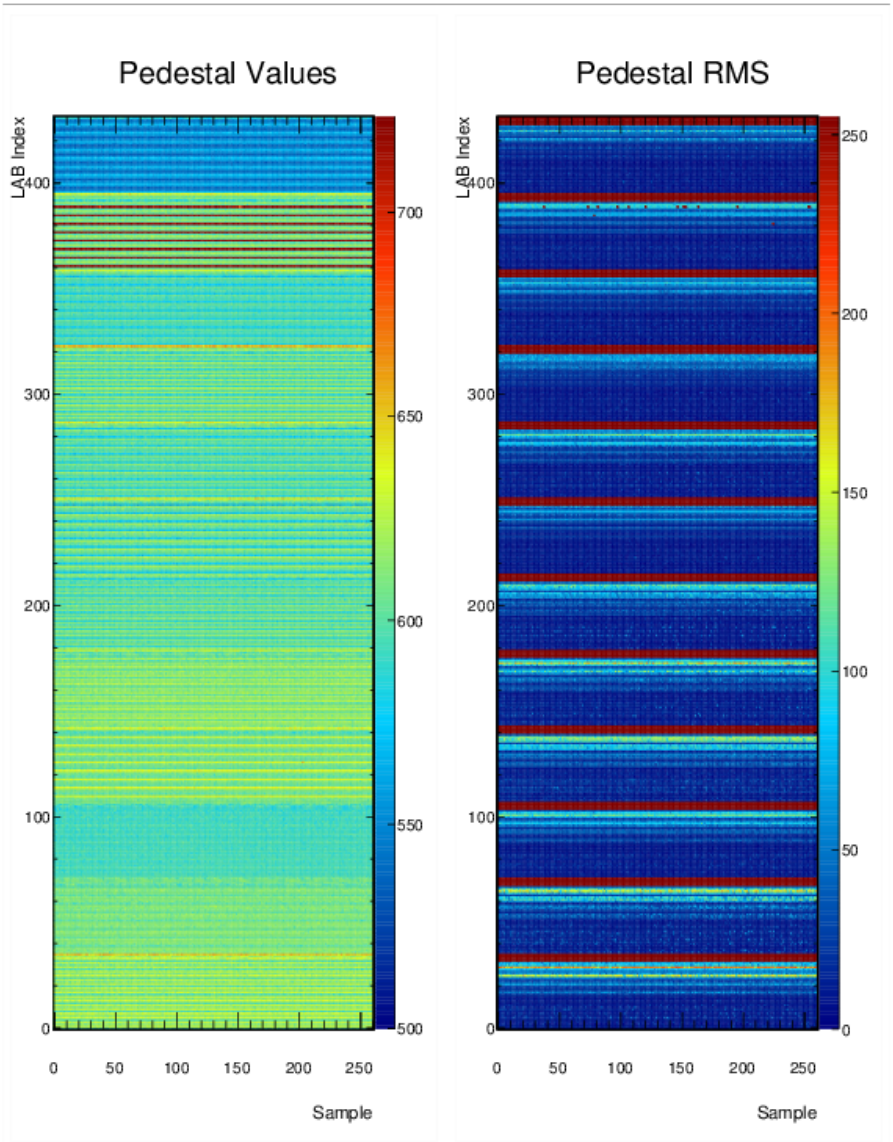

Figure 3.2: Two-dimensional histograms of the mean pedestal values used for subtraction(left) and RMS of bin in pedestal run (right) for all LABRADOR chips for the ANITAIII flight. The Y axis is an index describing the LABRADOR in question (SURF*36 + Chan*4 + LAB). The X axis is the capacitor bin number within that surf. The Z color axis is the value of the Pedestal or RMS in ADC counts. From this plot one can visibly see the small variance of pedestals within a single LAB, but moderate variation between SURFs. Also visible is the clock channel and its associated large pedestal RMS, as the clock remains on during pedestal runs.

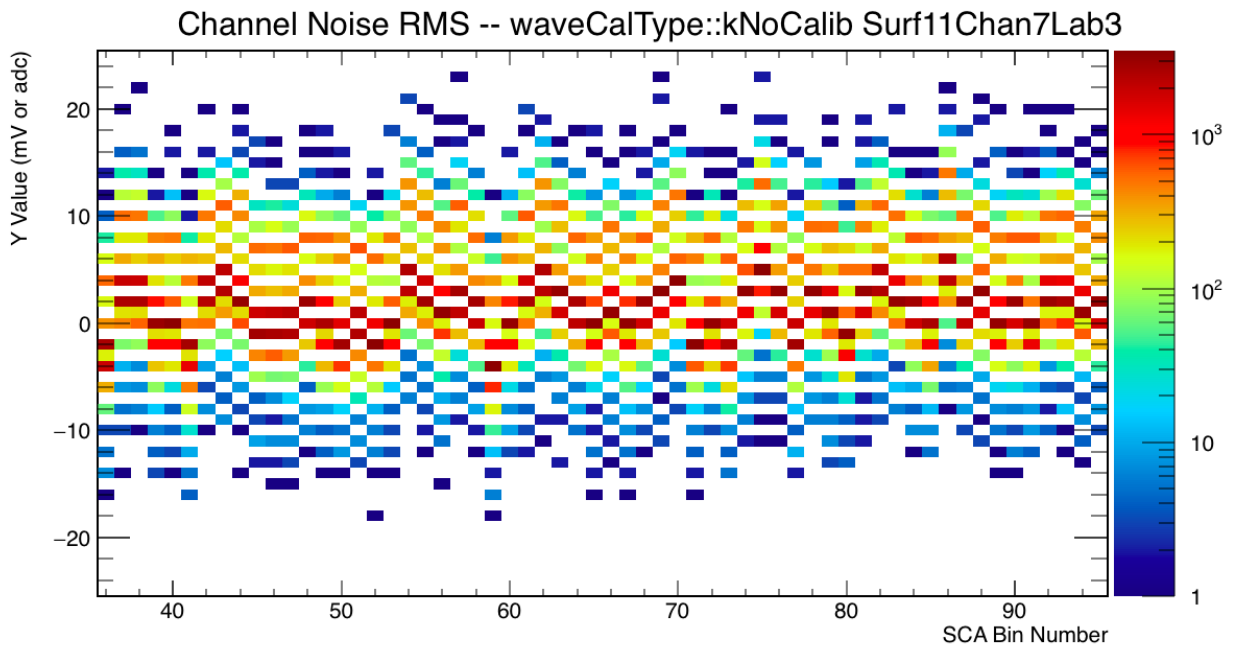

Figure 3.3: A two-dimensional histogram of ADC counts as a function of storage capacitor bin. The least significant bit from the LABRADOR is masked to zero, however the constant pedestal subtraction has a LSB addition that causes an offset (detailed in text).

pedestal correction (see subsection 3.1.2). This means that a single capacitor bin will have either even or odd values, but not both. This effect can be seen in Figure 3.3. Functionally it allows the stored pedestal values to have half bit resolution while not requiring float arithmetic for any real time CPU computations.

ADC count voltage normalization

To initially measure the scaling between observed voltage and recorded ADC counts of the LABRADOR chips, an identical RF pulse was inserted into two preamplifier channels, then measured both after the full signal chain immediately before entering the SURF RF input with a calibrated Tektronix oscilloscope, and again from the LABs using the ANITA readout software. The calibration set up is shown in Figure 3.4, note that this figure is used in multiple calibration analyses. The peak-to-peak height of the measured signal from the LAB chips was then compared against the pulse read out by the acquisition software and used to determine a scaling on a chip-to-chip basis. While keeping the reference channel constant, the test signal was moved throughout the remaining 95 channels.

One drawback of this method is that only the reference channel is coupled out from after the signal chain for comparison with the SURF output pulse. This pulse is dispersed by the amplification and filtering of the signal chain, and thus comparing it to the input pulse requires the

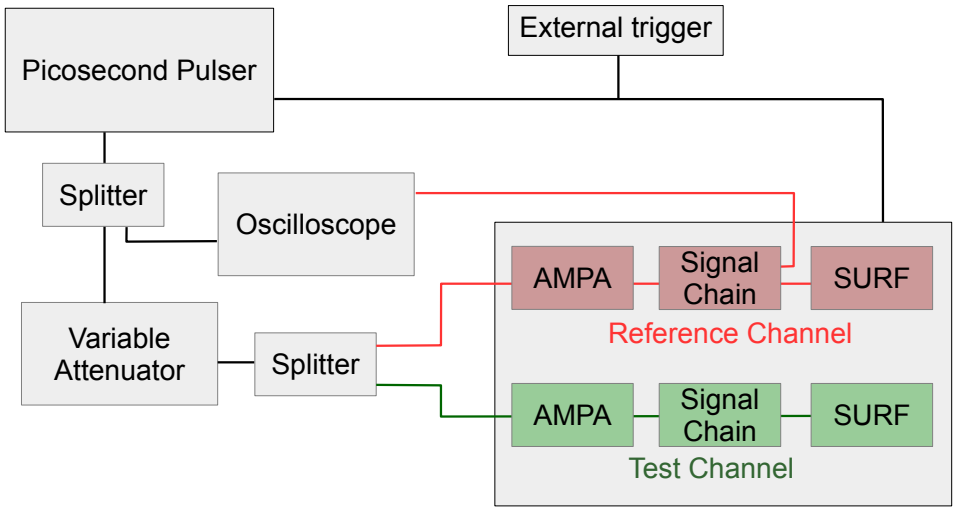

Figure 3.4: A schematic of the pulse insertion calibration setup used immediately before flight with the full signal chain in Antarctica in 2014.

full complex impulse response. The impulse response considers both gain and phase as a function of frequency for the system, an analysis that is described later in this chapter. For a simple first-order voltage calibration, one assuming voltage to be a linear function of ADC counts with no frequency dependence and an identical complex impulse response between signal chains, the peak-to-peak amplitude of each test channel compared to the reference channel is sufficient. However, since the magnitudes of the gains of the signal chains is not equal, this calibration serves best as a normalization between the full signal chains of all the channels. In other words, after applying the simple volts to ADC counts calibration, all channels are expected to have equal peak to peak voltages for an identical input signal regardless of signal chain gain or noise figure, though will not necessarily have the same noise power.

The ADC counts to mV conversion resulting from this analysis can be seen in Figure 3.5.

3.2 LABRADOR timing calibration

The LAB3 storage capacitors' time bases are controlled by a series of current starved transistors propagating a pulse down the sampling array. This pulse runs continuously and asynchronously to the global digital clock. The pulse is also output to the FPGA as a diagnostic logic signal called the Ripple Carry Out (RCO). This circuit allows each capacitor cell to be periodically and

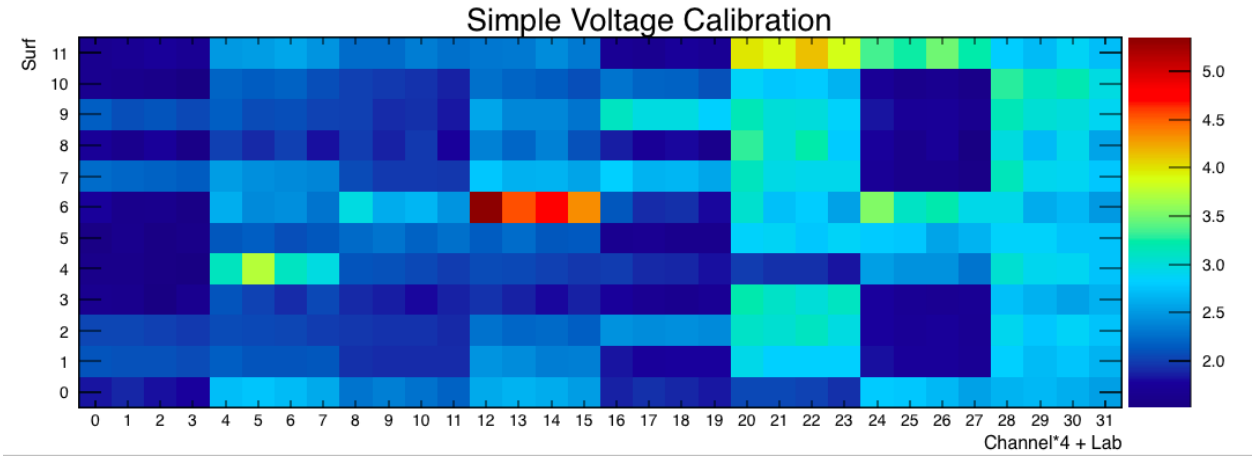

Figure 3.5: Two-dimensional histogram of simple voltage calibration values for all 96 channels of the ANITAIII instrument (taken from simpleVoltageCalibrationHarm.txt). The units of the Z-axis color bar are in mV/ADC counts. Note that each channel is digitized by four separate LAB chips. However, the largest effect on the voltage calibration is from variations in the signal chain.

sequentially connected to the input signal for a fixed time interval before being disconnected, leaving a stored charge on each capacitor related to the voltage at the sampling time which can then be later digitized. While the capacitor is connected to the input signal, it is referred to as being in "tracking mode". This timing generator circuit is effected by process parameter spread in the ASIC manufacturing process, and thus does not create equally spaced time separations between capacitors in the sampling array within a single chip. The timing separation, or delta-times (Δ Ts for short), of each sample is shared between all eight channels of each LAB chip, since it is generated by the same timing generator, however it can have significant variance between the different chips within a single SURF as well as between the SURFs. There are 96 total LAB chips actively recording data in the ANITAIII instrument, and each must be calibrated for this ΔT variance.

3.2.1 Time Domain Bin Width Matrix

The time bin width calibration constants, ΔTs, must be generated for each LABRADOR chip, of which there are 48 in ANITAIII, for each of the 260 samples per chip, as well as for each phase of the write strobe (up-going or down-going), leaving us with 520 total calibration constants per LAB and 24960 total values. There are several methods that can be utilized to determine the timing separation between any two capacitor bins, three of which were used to check the calibration of the ANITAIII instrument. The nominal LABRADOR sampling frequency of 2.6GHz, under perfect conditions, would yield a ΔT of 384ps for all bins. The distribution of values as determined through multiple calibration techniques is detailed below.

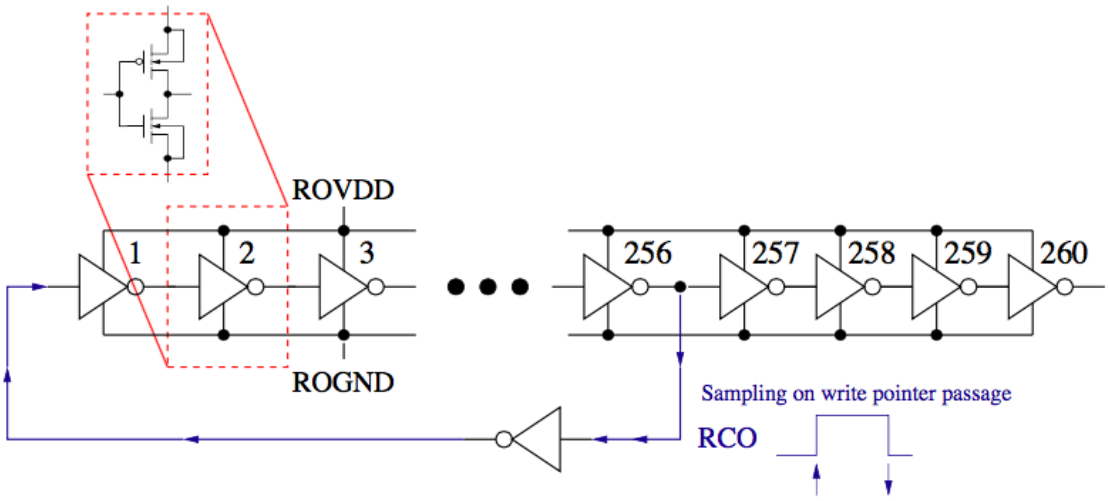

Figure 3.6: Diagram of LAB3 timing generator circuit [8].

Bin Occupancy Fraction Method

The first calibration technique will be referred to as the Bin Occupancy Fraction (BOF) method. This method involves injecting a constant periodic signal, either a digital clock or a constant frequency sine wave, into each of the LAB chips. By recording a large number of events, we can measure the resulting number of zero crossings that occur between each capacitor bin pair. By comparing this occupancy vs the expected occupancy for a uniformly sampled time window, in which each bin observes the same number of zero crossings, it is possible to determine the fractional width of each bin. A drawback of this method is that it requires very large statistics to reduce the uncertainty on the resulting value. A diagram summarizing this method is shown in Figure 3.7. The equation for calculating a nominal bin width is in Equation 3.1.

$$\tau_x = \frac{N_{meas}}{f N_{samples}} \tag{3.1}$$

Where τ_x is the timing width of the bin, N_{meas} is the number of zero crossings measured within that bin, f is the frequency of the calibration signal, and $N_{samples}$ is the number of waveforms used to generate the value.

This calibration was done for ANITAIII with two timing sources, an injected 432.1MHz sine wave and the 33.3MHz square wave sync clock injected into the system for the entire flight. A BOF calibration using an injected sine wave was preformed by Ben Strutt, and detailed in his thesis ([9]). This is the nominal calibration whose values are used for the remainder of this thesis. The array and distribution of bin widths can be seen in Figure 3.8. All other bin width matrix timing calibrations detailed below were done as a check on those values, and were found to be in

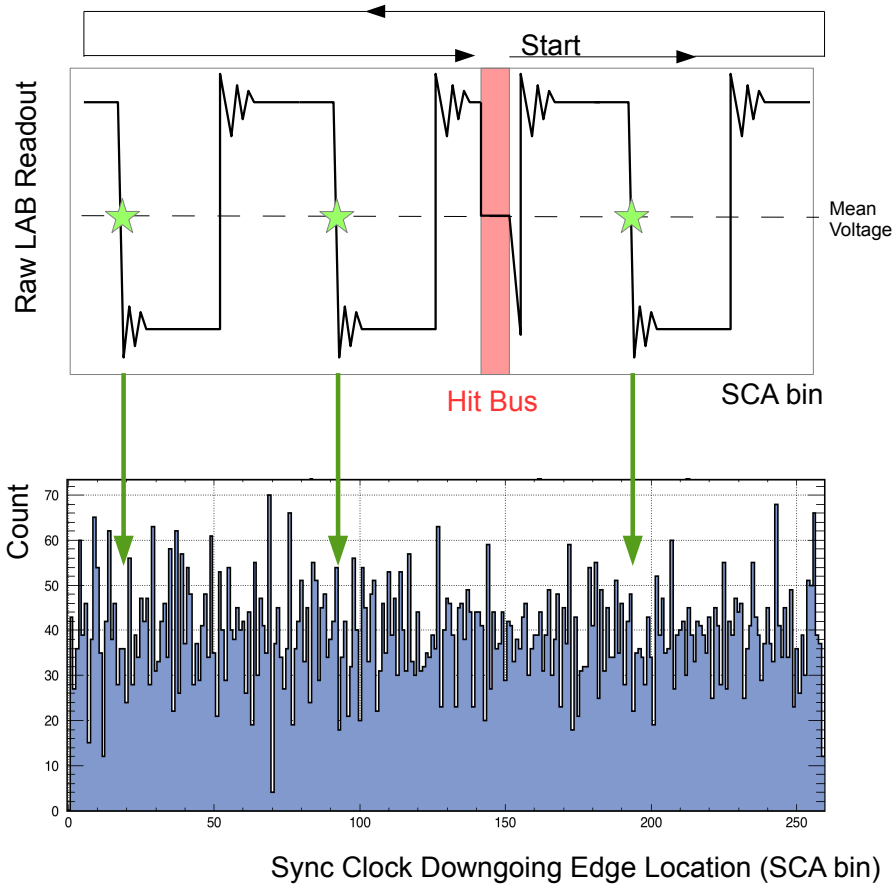

Figure 3.7: A simplified diagram showing the Bin Occupancy Fraction method. The raw, unwrapped, LABRADOR waveform readout is shown on top. The x-axis is capacitor number, and the y-axis is ADC counts recorded by the digitizer. The Hitbus region marked in red. The points immediately following the hitbus are the oldest samples, and would be unwrapped to the "front" of the remaining data (visualized by the arrows). The green stars mark zero crossings, which are used to fill the bottom histogram. The zero crossing and histogram filling process is repeated for as many waveforms as possible. The final bin-width calculation is then done by taking the fraction of zero crossings recorded in each capacitor bin against the expected number of crossings as dictated by the number of waveforms and clock frequency (see text).

agreement.

The second possible calibration source, a 33.3MHz square wave clock, is also useful for this calibration method. Though the duty cycle of the clock has undesired variability, full period lengths are stable. Using the falling edge of the clock provides a periodic signal with which to populate timing bins and derive their width. It has the additional benefit of being constantly present throughout the entire flight, as well as throughout the pre-flight period in which amplifiers were off yet data was continuously being taken for stability validation. Results of this analysis using only pre-flight runs where amplifiers were off are compared to the nominal values in Figure 3.9.

Ellipse calibration method

The second calibration technique is slightly more clever, and involves utilizing a characteristic of a collection of sine wave captures with random phase offsets, and will be called the Ellipse method [43]. Plotting the sum and difference of sequential samples of a pure sinusoidal function yields an ellipse whose semi-major and semi-minor axes determine the fractional time of each sample.

Using this method it is possible to determine the width of each bin by comparing the sum and difference of digitized values from neighboring bins that are offset by some time ΔT both observing the same sine wave. It is useful to consider the edge cases of such a system, for example two bins that have a negligible ΔT separation. In this case the measured voltage difference between them will always be small if not zero, while the addition axis can occupy a large range of values, the maximum being the crests of a sine wave. This will yield an ellipse with a very large eccentricity. Inversely will have the inverse eccentricity, bins that are separated by exactly one half wavelength of the injected frequency will always have wide range of possible differences, and sums always approaching zero. Nominal sampling between these extremes will provide an ellipse approaching a circle, as seen in Figure 3.10. This method is useful in that it requires far less data than the Bin Occupancy Fraction method.

Utilizing this method, all capacitor bin ΔTs were calibrated, and the results compared with the previously measured constants provided in this thesis and other published results [9]. The values are presented in Figure 3.11.

3.2.2 Sine Wave Fitting Calibration

There exists a third method for calibrating each digitizer bin's timing width, which is applying an iterative least-squares fit of the inserted sine wave and determining the relative phase offsets of the voltage of each recorded point versus the same value of the fit. This was done as an additional method for determining the validity and event to event stability of the ΔT constants. This method has the additional benefit of working well even with low statistics, with a measurement of the width possible on a single captured waveform. Of course, the four LAB buffer and dual RCO

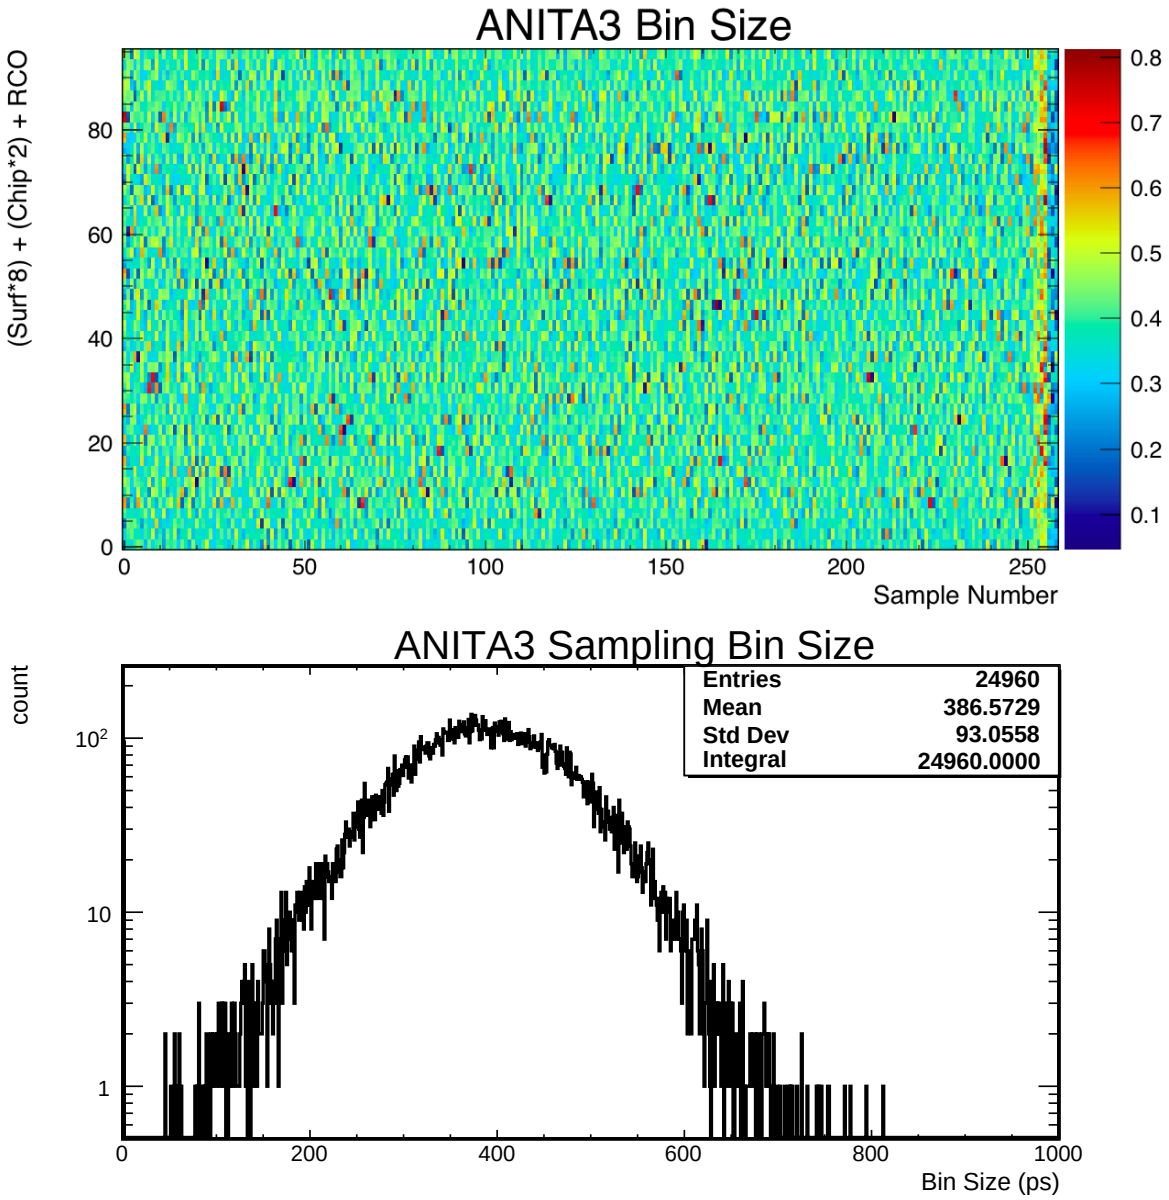

Figure 3.8: Nominal LAB3 Time Domain Bin Width "ΔT" values generated using a 432.1MHz injected sine wave calibration signal and the bin occupancy fraction method. Values were determined by Ben Strutt[9]. (Top) A bin-by-bin 2d histogram of the sampling widths for all capacitor bins, organized by SURF, LAB chip, and RCO phase. (Bottom) The distribution of all 24960 calibrated timing widths.

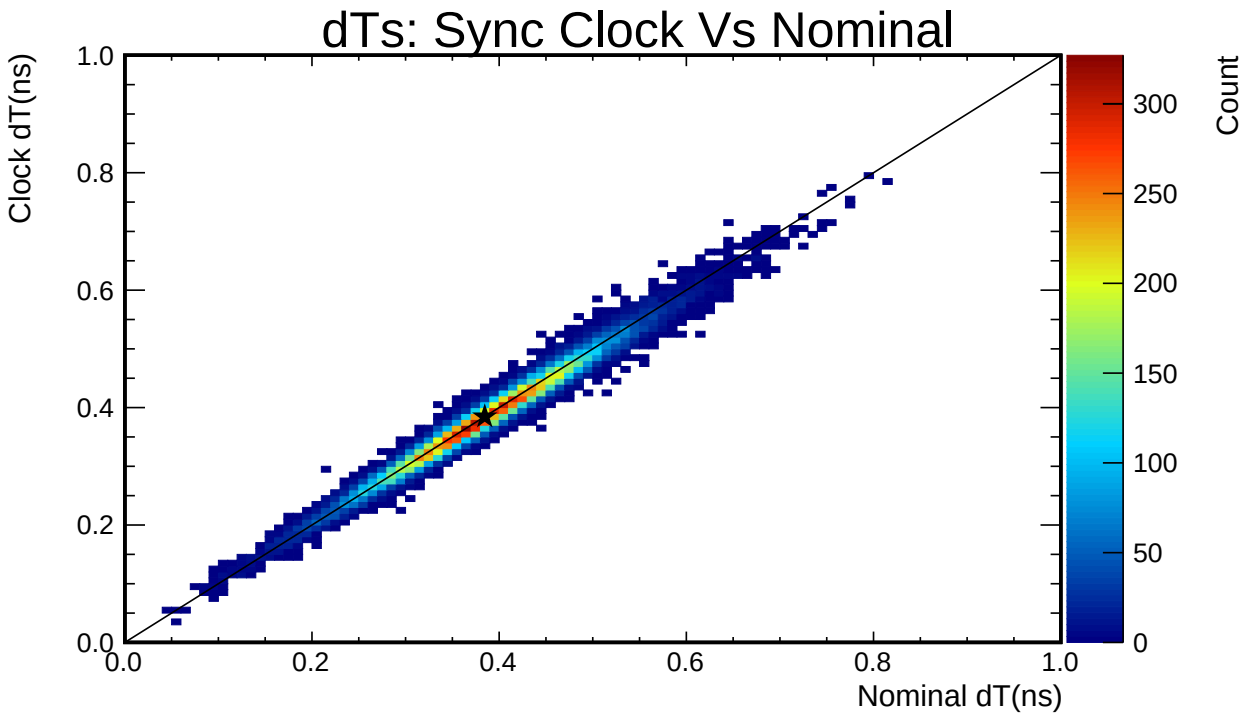

Figure 3.9: Comparisons of the derived bin widths for all channels using injected sine waves (nominal) and the 33MHz synchronization clock. The subtle black star marks the 2.6GS/s design sampling rate, and the black line marks where the two derived values would be identical.

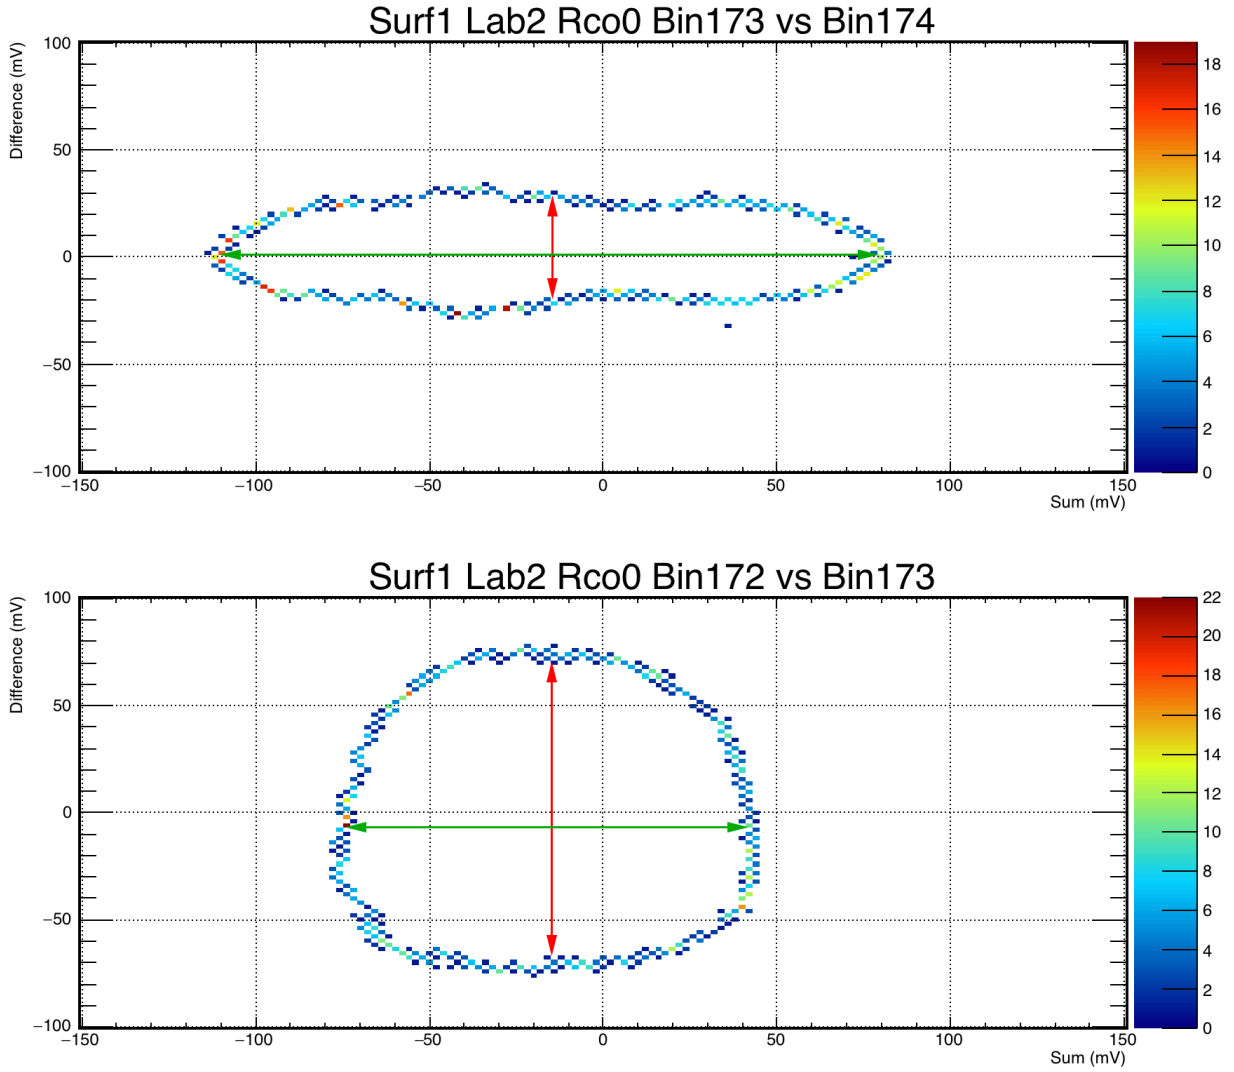

Figure 3.10: Two plots showing the variation of the ellipse created by two different capacitor bins with different timing separations. Using the ratio of the semi-major (green) to semi-minor (red) axes, and knowledge of the injected sine wave frequency, can be used to determine the bin width per capacitor. The bin pair shown in the top graph has a smaller time separation than the bin pair in the lower graph.

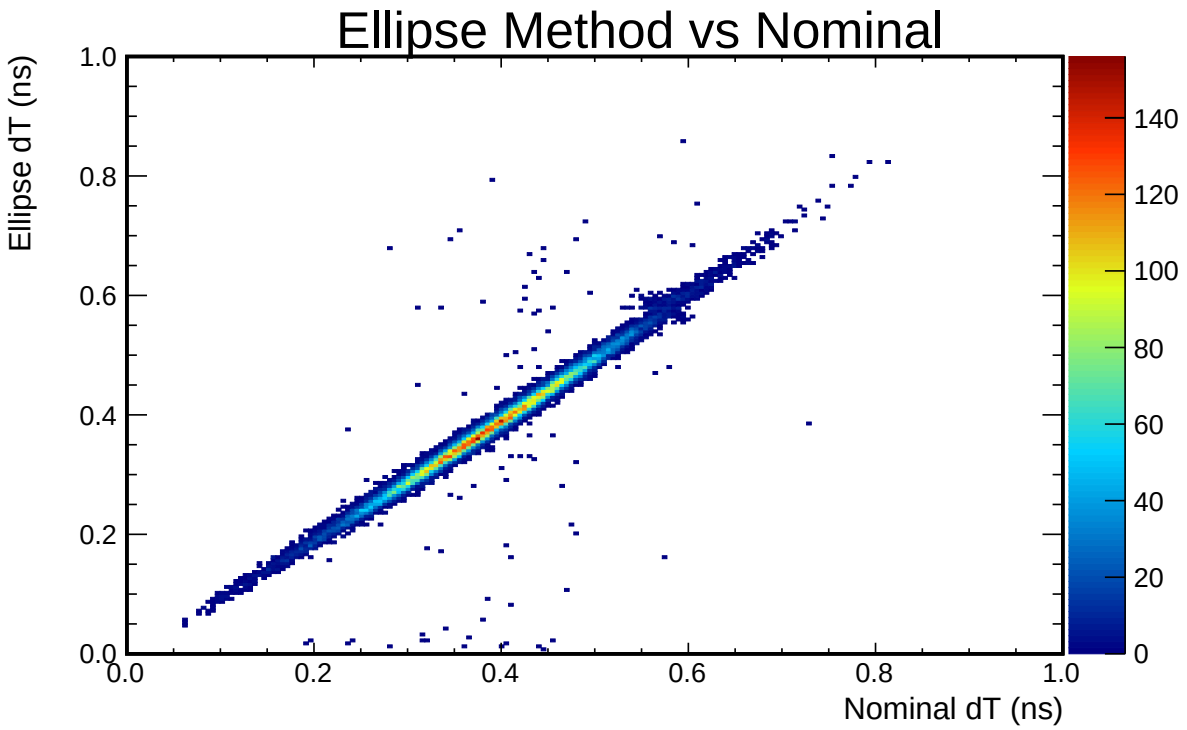

Figure 3.11: Comparison of calibration constants determined by the ellipse method and the nominal values. Single bins with large variations are the result of a poor ellipse fit. The high level of correlation between the two distributions lends support for the accuracy of the nominal calibration.

phases results in each bin only being measured one eighth of the total event read out. Additionally, the only periodic input function was a sine wave which has a large fractional region where small voltage offsets effect large timing offsets, or regions where the derivative is small. The bin occupancy fraction method is in a way a simplification of this method, using only times where the derivative is largest. Additionally, it is noteworthy that this method is already used to determine the Epsilon wraparound time (discussed later in this chapter).

3.2.3 dT Variance and Timing Results

The variance in the timing from bin to bin has a strong effect on the measured bandwidth of the system. The nominal sampling variance between the ANITAI and ANITAII flights increased by a factor of 3.5 to 25%, which reduces the signal power in the high end of the spectrum . The cause of this change between the ANITAII and ANITAIII flights is unknown, and was the reason for employing several different validation techniques of the measured ΔT width values.

3.2.4 Wraparound Time (ϵ

The propagation time of the write strobe between the end of the ring SCA and the beginning is larger than a single bin. As the SCA acts as a ring buffer, and the pulse can occur at any time within the ring, the write strobe must begin its transition back to the first bin before the end of the SCA at bin 256. The remaining four bins are then used to fill in the "wraparound" time, which will be described with ϵ. These epsilon values, like the ΔT bin widths, are dependent on the polarity of the RCO signal, and so each chip receives two ϵ correction constants. In addition, four ΔT periods is often longer than the wraparound time, and as such the 260th sample is often treated as a duplicate and discarded.

To measure these wraparound values, a sine wave is fit to the bin-width-corrected and unwrapped samples. This fit is then used to derive the offsets between the last and first bins in the array. The results of this analysis are shown in Figure 3.12

3.2.5 RCO Phase Determination

The RCO phase is sent as an output to the FPGA as a buffered copy of the wraparound voltage. Using this, it is possible to determine the current phase of the chip when a hold was placed. At the boundaries of the SCA, when the RCO has recently changed, the FPGA may not have the most current RCO value latched to a register when returning the digitized values. A correction for this can be made by measuring the period of the ninth channel sync clock, which will differ between each RCO phase. Using the firmware reported value in conjunction with the sync clock period allows for correctly determining the RCO phase with high accuracy.

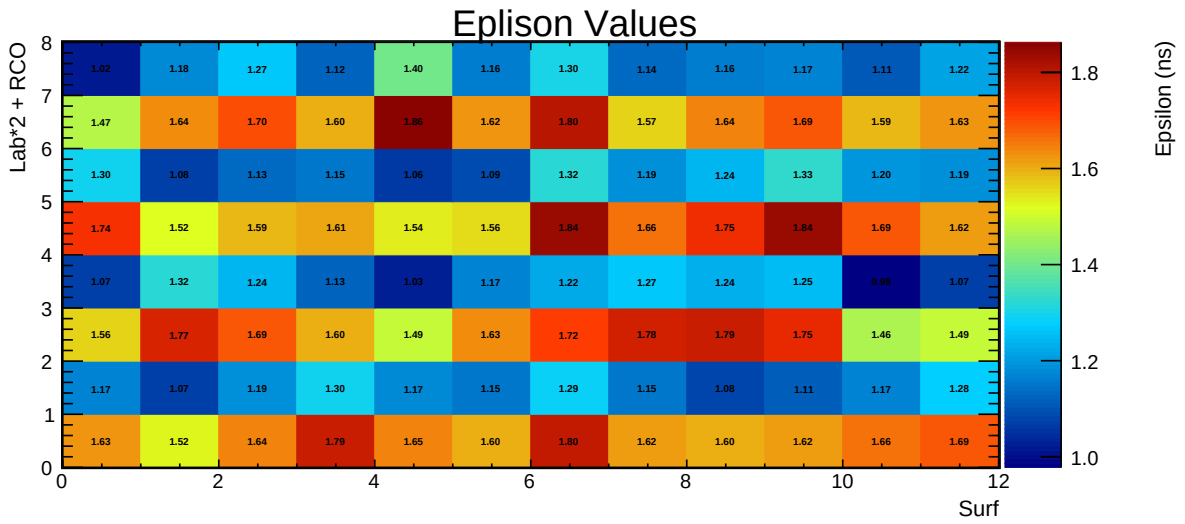

Figure 3.12: Measured values of the wraparound time (Epsilon) as a function of LAB and RCO phase

3.2.6 Temperature Dependence

All transistors react differently to temperature, and those present in the time base generator are no exception. As the payload cools and warms due to rotation, time of day, and latitude, the sampling frequency will increase and decrease. To alleviate this, a temperature correction must be applied to the final data. This is done via an average of the recorded number of cPCI clock transitions reported by the FPGA over multiple events.

3.2.7 Inter-SURF Timing for Waveform Alignment

As each LAB3 runs independently of each other and each channel has different signal chain group delay characteristics, each channel must be aligned in time with each other to determine the precise timing of any digitized waveforms. The delay between channels has both an invariant delay and a dynamic "jitter" time offset that changes on an event by event basis.

The invariant term, caused by the path lengths of the digitization control signals between the TURF and SURF boards in addition to any group delay offsets introduced by the signal chain, can be initially calibrated out using coincident pulses inserted into multiple channels. To measure and correct for this offset, an impulse was sequentially injected into each channel while a coupled impulse was simultaneously injected into an unchanging reference channel. By measuring the difference between the observed arrival times of these pulses at each SURF, it is possible to develop a matrix of constant delay offsets between channels. This calibration setup is described in Figure 3.4. The results of this calibration are shown in Figure 3.13.

The event-by-event time variance introduced by the nature of discrete-clocked digital logic and

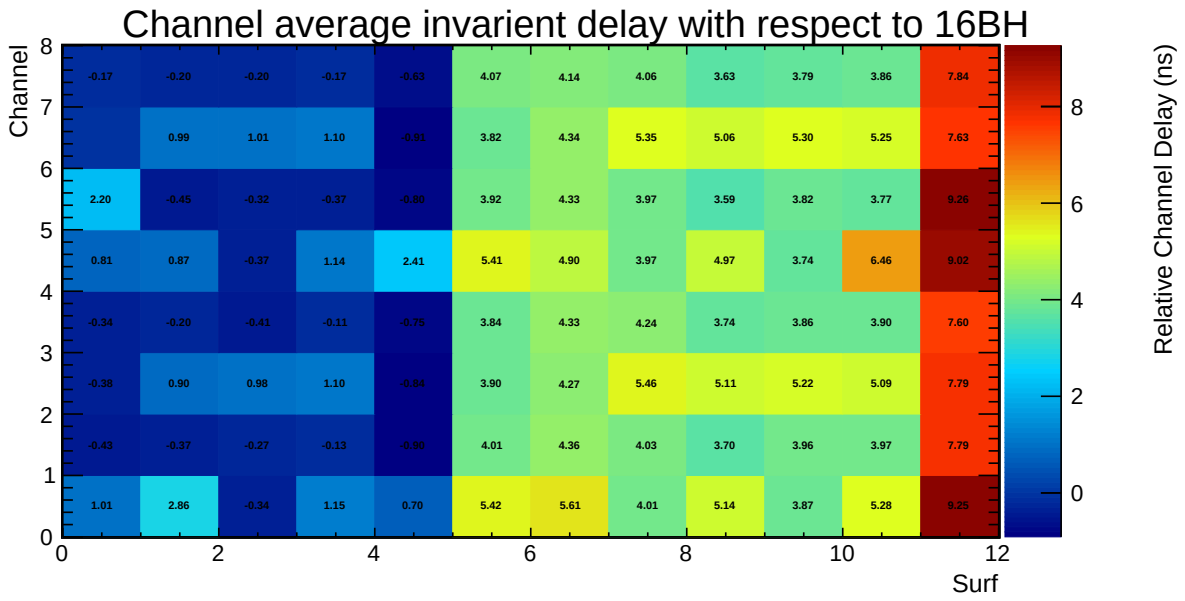

Figure 3.13: Measured static delays between each digitization channel. Caused by propagation delay of HOLD signal to digitizers, as well as cable and other physical delays in the RF signal chain.

the free-running sampling controls of the LABs, must be done at each event. This is done with a "sync" clock that is propagated amongst the boards and constantly digitized in the 9th LAB3 channel. This constant 33.3MHz digital clock is independent of any digital logic and thus should have a non-varying analog group delay offset between channels. Thus, any measured offsets of the digitized sync clock between LAB channels must then be an effect of the jitter in the propagation of HOLD command and subsequent response by the LAB. Aligning these clocks post-digitization can be used to correct for this invariant term.

3.3 Ground-based in-flight calibration pulses

Three ground pulser stations, and one trailing balloon source (HiCal), were set up on the Antarctic continent for the ANITAIII flight. These transmitted a high power, GPS synchronized, broad spectrum pulse that closely mimics an EAS impulsive radiation signature. These pulses can be used both for determining the overall sensitivity and efficiency of the trigger system, but can also be used as known-source measurements for doing precise in-flight antenna phase center location calibrations.

Two autonomous pulser stations were deployed at field camps in east Antarctica, and an additional manned pulser pulser was deployed at the launch site, also known as the Long Duration Balloon (LDB) facility. The remote stations were located at WAIS Divide and Siple Dome, however

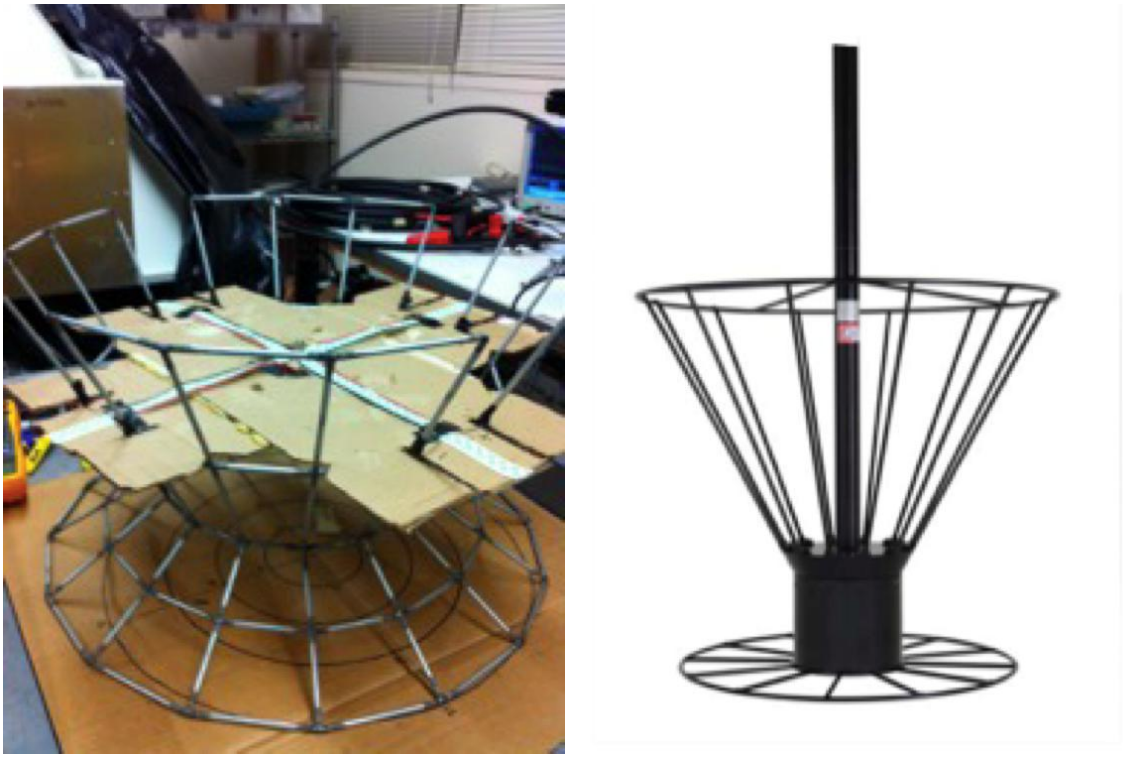

Figure 3.14: Images of the calibration pulser antennas deployed at WAIS divide and Siple Dome for calibration of the ANITAIII instrument in flight. Right: Custom designed, quad-slot "Mini-ALFA" antenna for horizontally polarized signals. Left: a Telewave ANT280S discone vertically polarized antenna. [10]

the payload only passed within line-of-sight of the WAIS divide pulser location. They employed an azimuthally symmetric, quad-slot "Mini-ALFA" antenna custom designed by A. Romero-Wolf for horizontally polarized signals, and a Telewave ANT280S discone antenna for vertically polarized signals. These are shown in figure 3.14. Their S_{11}, a measure of antenna efficiency, is seen in Figure 3.15.

All ground based calibration pulsers and the payload timing were synchronized to the GPS second. This was done to ensure that any calibration pulse signature could be discriminated from physics signal events easily. The HiCal pulser system was unable to employ a GPS synchronization due to cost and weight limitations, and their existence can only be derived from the reported location telemetry from its payload.

3.3.1 WAIS Divide

The autonomous station deployed at WAIS divide consisted of two antennas, one vertically and one horizontally polarized, pulsing at a repetition rate of 1Hz. For unconfirmed reasons, only horizontally polarized pulses measured at the ANITAIII instrument. The WAIS calibration pusles

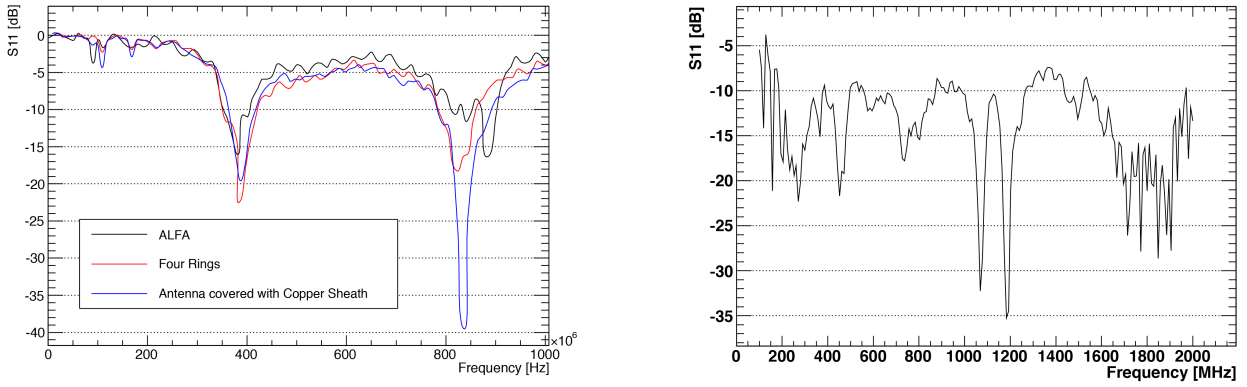

Figure 3.15: Measured S_{11} antenna efficiency for the horizontally (left) and vertically (right) polarized antennas deployed at remote locations for the ANITAIII in flight calibration pulsing. [10]

give us our best estimate for the locations of in flight antenna horizontally polarized phase centers.

Discriminating between WAIS pulser events and noise or physics signal events was done by making a cut on the TriggerTimeNs quantity stored in the header section of the event data. Since this quantity is reset on the GPS second, and the ground based calibration pulser was fired on the GPS second, the arrival time of the signal at the payload is predictable. Figure 3.16 shows a histogram of the occupancy of TriggerTimeNs values as a function of time. Using the GPS coordinates of both the payload and pulser to determine and correct for the time of flight and expected arrival time can be seen in Figure 3.17.

Reconstructed WAIS calibration event locations (a process described in subsection 4.5.5 and ANITA's flight path past the station can be seen in Figure 3.18.

3.3.2 LDB

The LDB pulser was observed twice during both the initial launch period, as well as during the second pass on the subsequent orbit around the continent. This pulser provided the only vertically polarized calibration signals. A drawback of this pulser location is the high level of non-thermal anthropogenic noise generated by McMurdo Station and the surrounding field camps. Subsequently, a high power level and high repetition rate of the pulser was required in order for the payload to efficiently self trigger on the signals. Additionally, as the signals were arriving from the same direction as the noise sources, trigger phi-masking had to be turned off for these measurements, a parameter not characteristic to the majority of the flight. Figure 3.19 shows the distribution of arrival times as a function of nanoseconds from the GPS PPS signal.

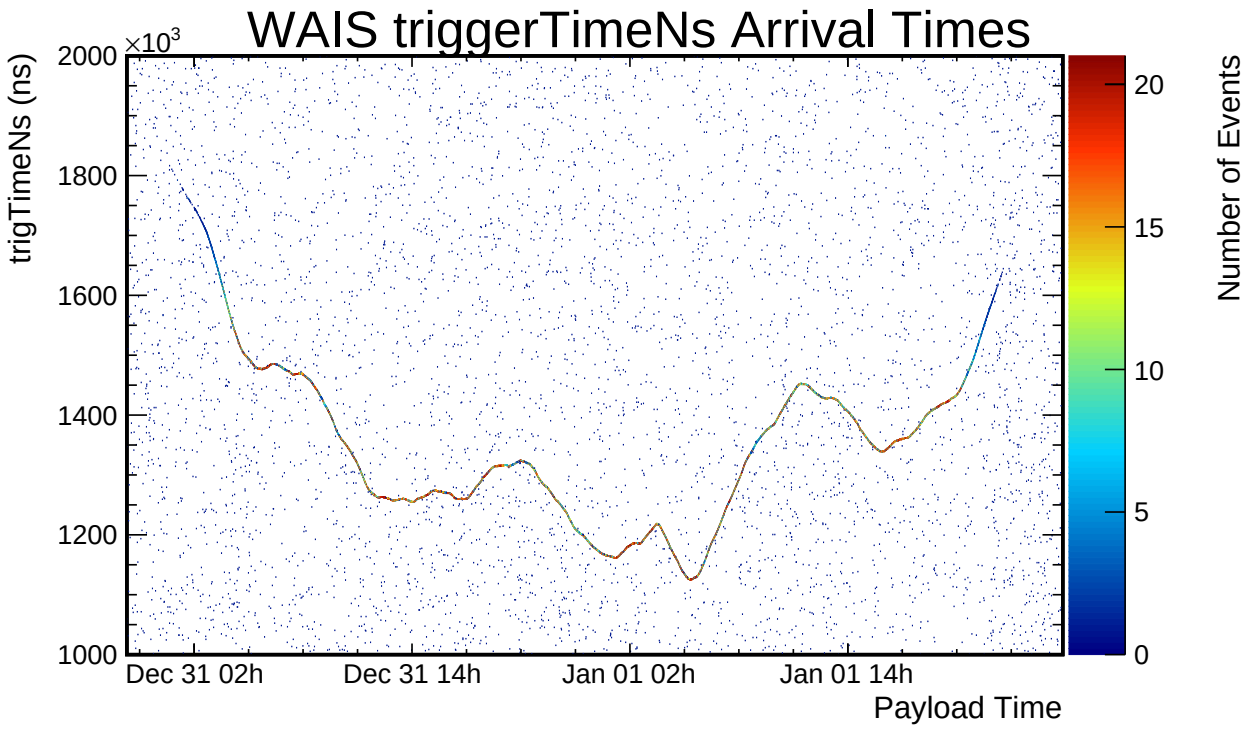

Figure 3.16: Stored triggerTimeNs trigger decision time distribution for RF triggers in the approximately two day period during which ANITAIII was within line of sight of the Wais Divide calibration pulser. The visible line is an increase in the number of events per bin well above expectations. The variation in the line is movement of the payload towards and away from the pulser location and the subsequent time of flight variation.

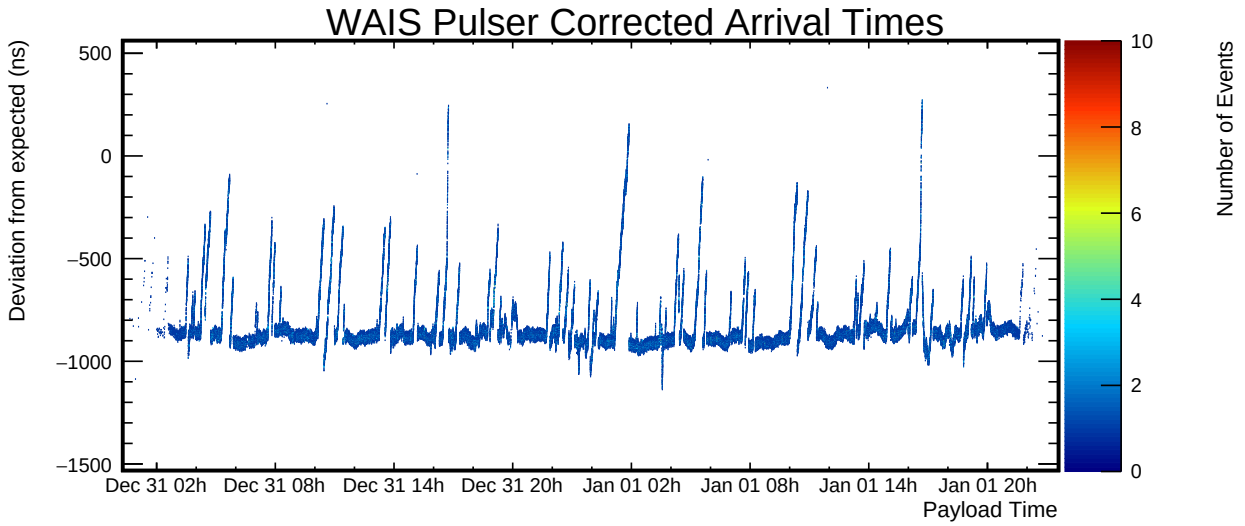

Figure 3.17: Stored triggerTimeNs trigger decision time distribution for RF triggers in the approximately two day period during which ANITAIII was within line of sight of the Wais Divide calibration pulser, corrected for time of flight to the balloon payload. Deviation spikes from the expected arrival time, which can be seen are as large as 1μs, are likely the result of the remote pulser station losing its GPS second synchronization lock.

3.3.3 HiCal

As real cosmic ray EAS radiation signals are emitted in the atmosphere, it is desired to have a calibration source that mimics this source location. HiCal utilizes a piezoelectric spark generator driven by a small DC motor attached to a horizontally polarized fat dipole antenna hanging from a balloon payload. After a failed launch attempt immediately after the ANITAIII launch, a second module was successfully launched as ANITA passed by LDB on its second orbit on January 6th. The flight path of HiCal can be seen in Figure 3.20. See [44] for more information about the HiCal flight.

3.4 Antenna Phase Center Location Measurements

To accurately reconstruct the arriving electromagnetic wavefront it is extremely important to know the positions of the antennas to very high precision. Using $\Delta x = \tau c$ for determining the distance traversed by a travelling wave, and unscientific imperial units, a wavefront traveling at the speed of light will traverse approximately one foot per nanosecond of time.

The timing limitation imposed by the inter-surf timing discussed above is on the order of 10ps, which would correspond to a 3.6mm maximum uncertainty on the phase centers of the antennas before the become the limiting measurement for coherent reconstruction of several digitized waveforms. Measuring the relative positions of these with such accuracy, or assuming their positions

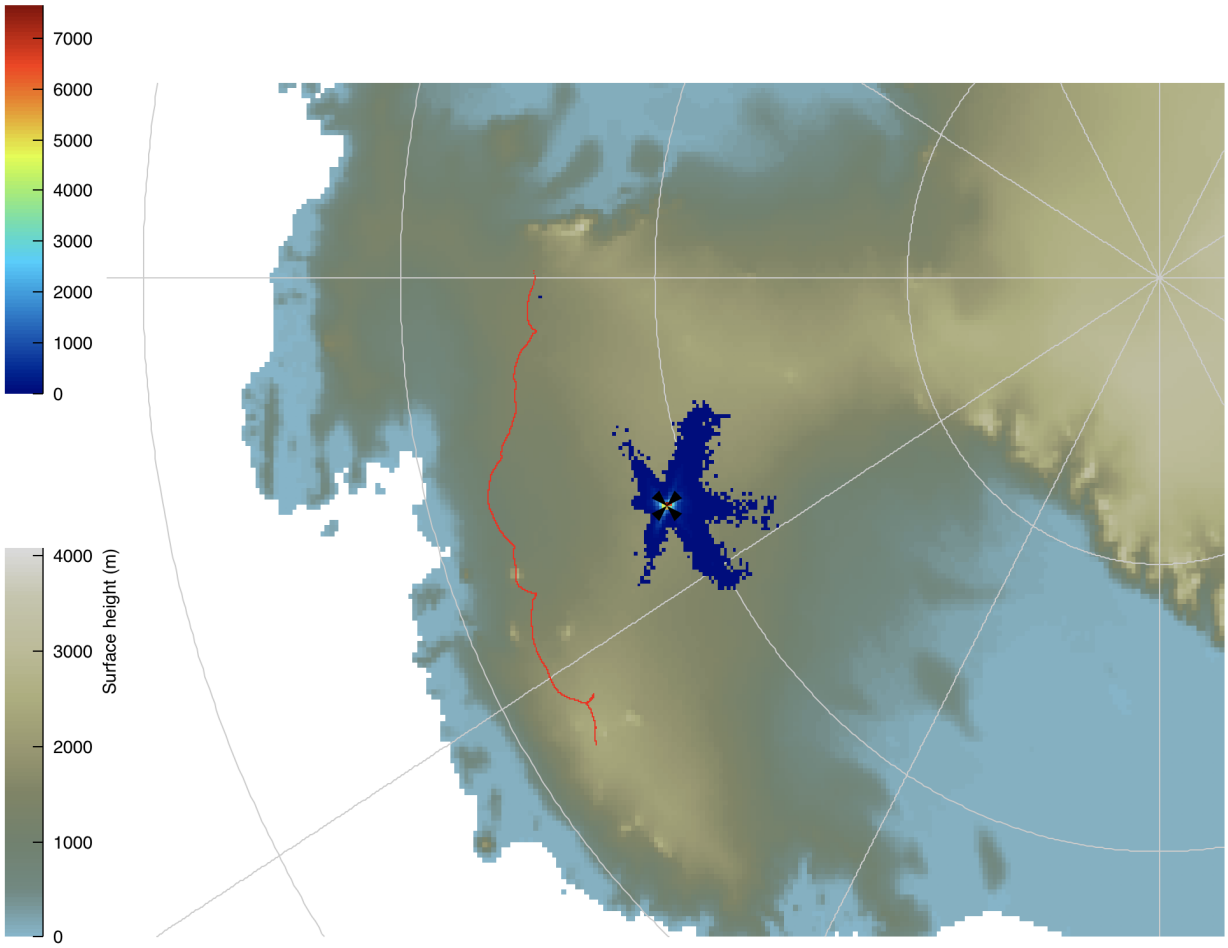

Figure 3.18: Map of Antartica with WAIS calibration event reconstructed direction overlaid as a two-dimensional histogram. Black "X" denotes the location of WAIS as measured with GPS by the calibration deployment team. Red line denotes the flight path of the ANITA instrument during the time in which WAIS calibration pulses were observed.

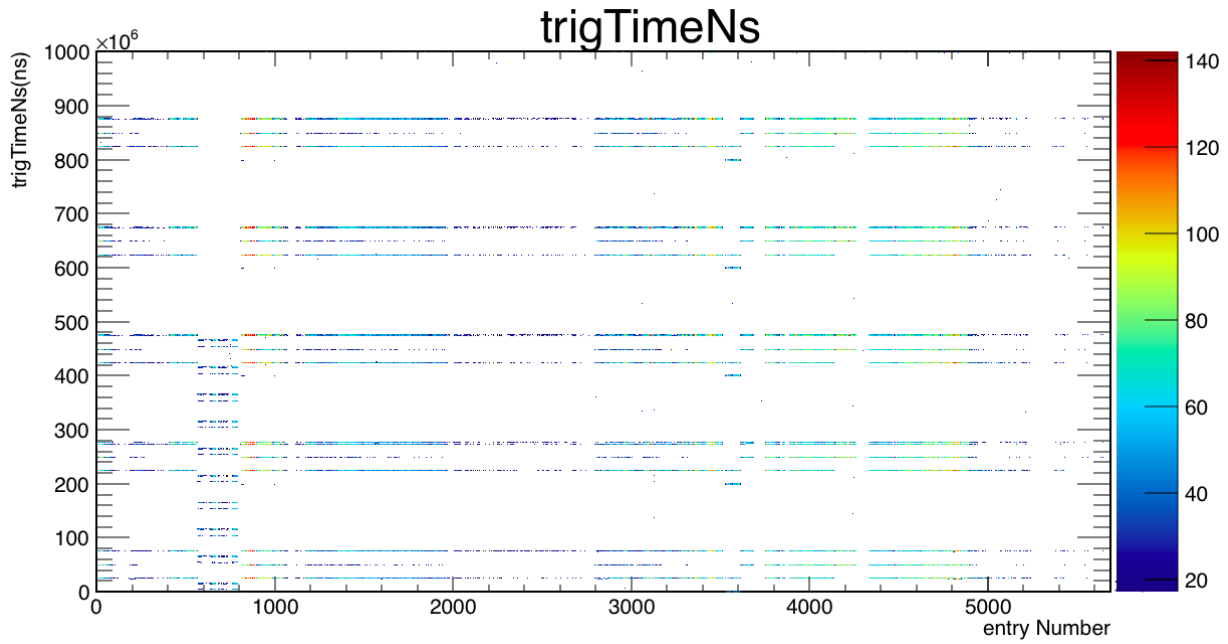

Figure 3.19: Stored triggerTimeNs trigger decision time distribution for RF triggers immediately after launch while ANITAIII was within line of sight of the LDB calibration pulser.

from a design model, is unrealistic, and so several other methods have been employed to determine their positions.

3.4.1 Photogrammetry

An initial measurement of the antenna locations with respect to each other, to the instrument box, and to the GPS antennas (and thus recorded attitude heading, and position), was done by analyzing a series of standardized photos of the payload taken from a variety of angles and rotations. Utilyzing this data coupled with precise measurements of single antenna dimensions, it is possible to arrive at a first order estimate of the antenna positions. An example image and subsequent wiremesh model of the payload can be seen in Figure 3.21.

3.4.2 Calibration Pulse Timing Optimization

Determining the apparent physical location at which the electromagnetic field couples to the antenna, called the phase center, requires an additional calibration beyond the simple geometric location of the physical antenna. The phase center is a function of frequency and incident wave direction, so the calibration must consider a large range of variables. Though directly related to the physical position and orientation of the antenna, precise interferometric reconstruction requires accuracy beyond photogrammetry. The calibration measurements are done in-flight by a ground

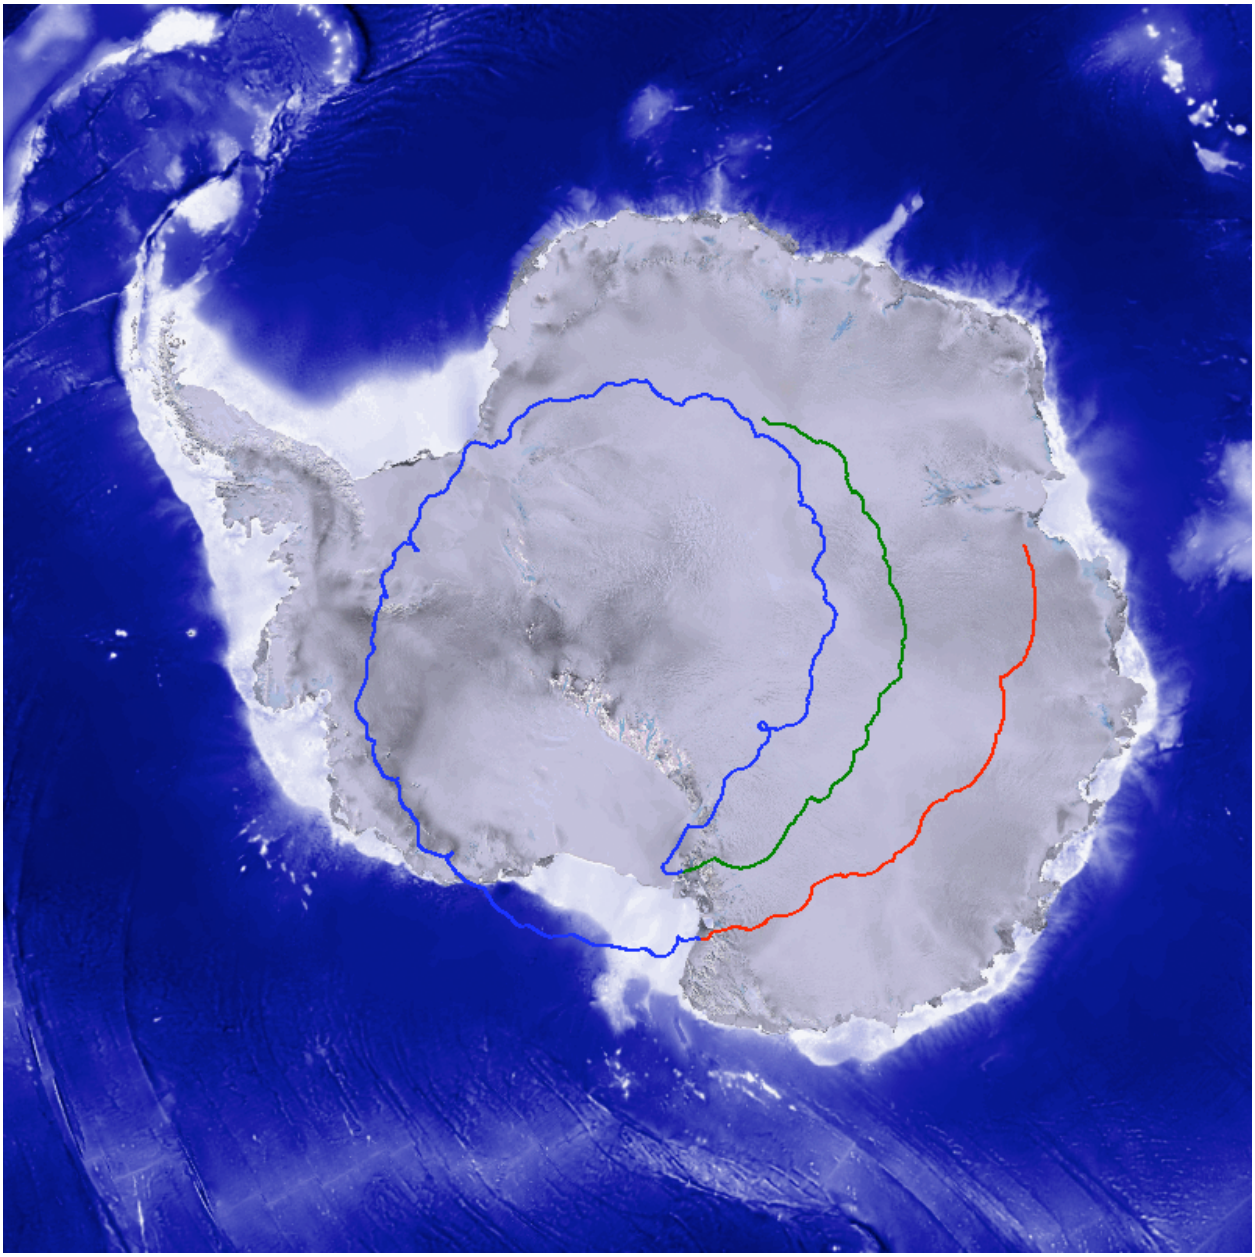

Figure 3.20: HiCal flight track over Antarctica (green). ANITAIII's first orbit (blue) and second orbit (red) are also visible.

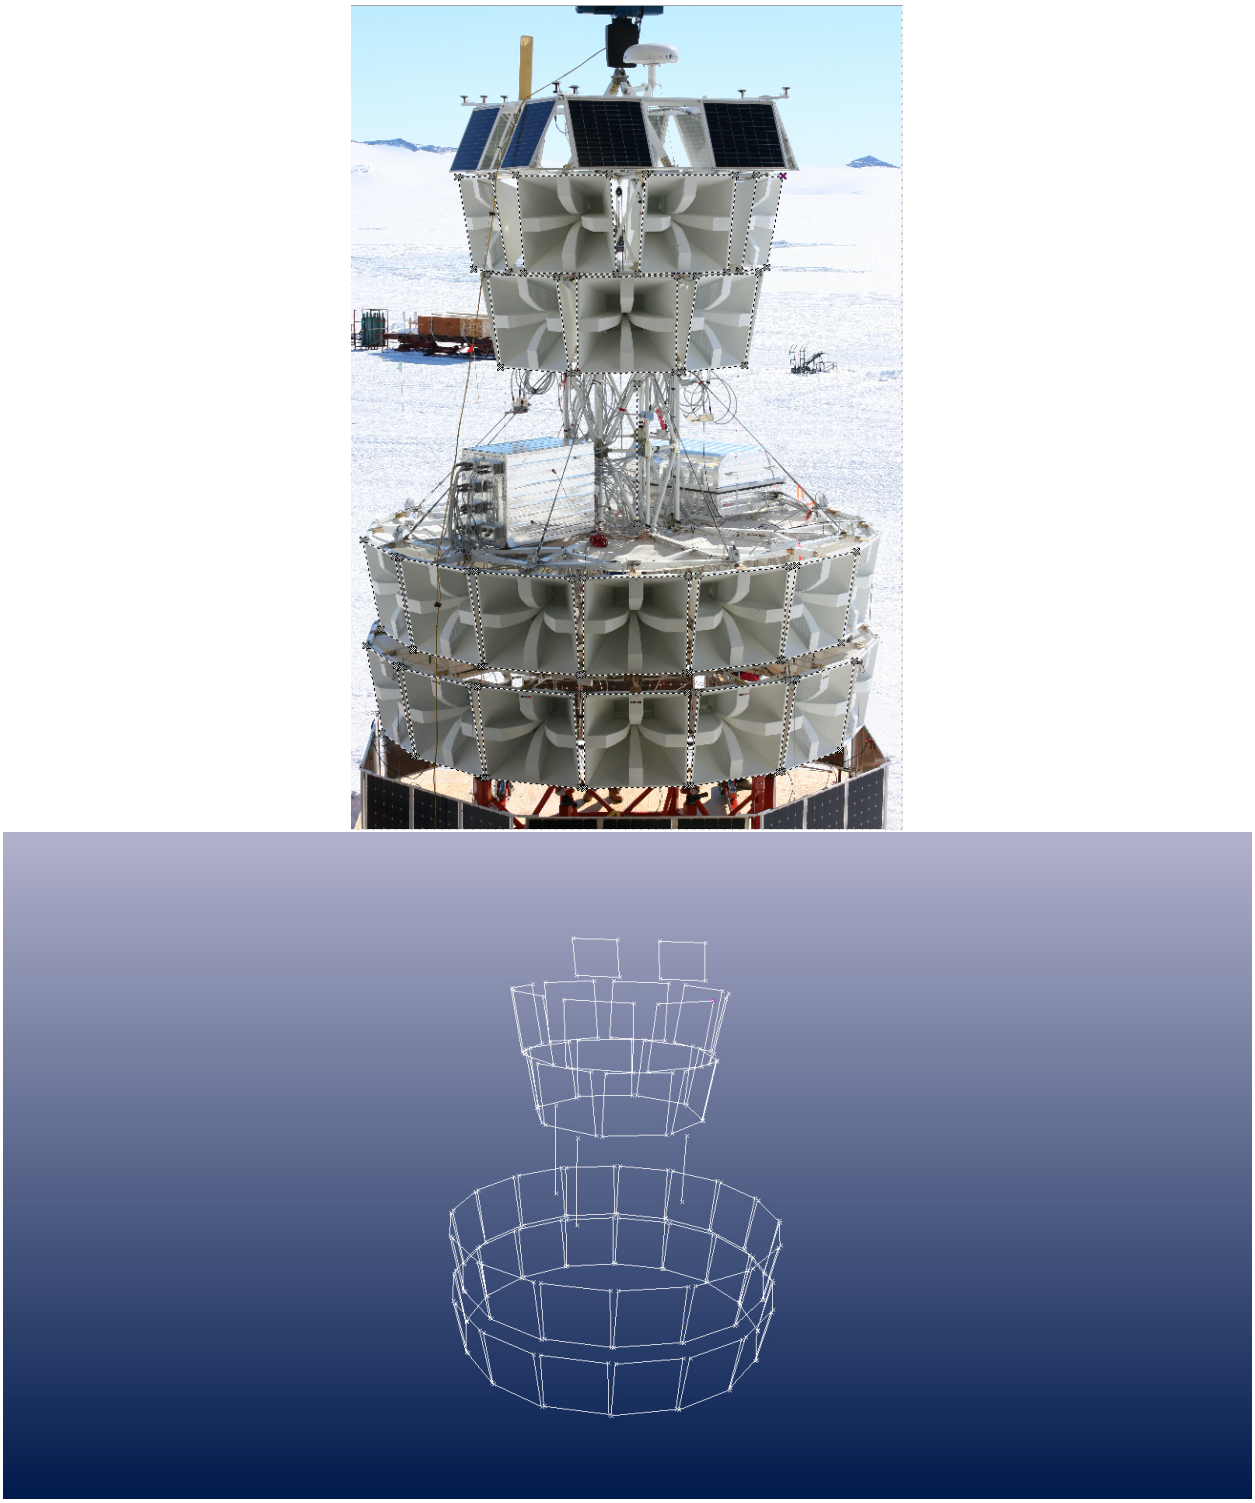

Figure 3.21: Top: Example photogrammetry image of the ANITAIII payload with markers drawn around the antenna face outlines. Bottom: Subsequent wiremesh model of antenna positions derived from numerous combined images. Images courtesy of Constantine Belov.

pulser transmitting an impulsive signal to the payload as it passes overhead. This pulse, able to be simultaneously measured by multiple antennas, can be used to determine the relative phase center offsets between nearby antennas. The positional error of the phase centers relates directly to the pointing error for events seen later in the flight. A diagram showing this method can be seen in Figure 3.22. Measurements of the vertical and horizontal polarization channels must be done separately, as there is a physical coupling offset required by antenna construction. Both the WAIS and LDB calibration pulser data is required, as the WAIS v-pol pulser was non-operational. These multiple measurements can also be used to validate the phase centers.

The peak correlation time offset of measured impulsive signals over all antennas can be used as an input to a minimization routine that determines the phase centers for the entire instrument. The variable parameters in the fit are the static cable delays, and the cylindrical coordinates of the phase center (height, radius and azimuth angle). Limits from photogrammetry and cable measurements can be used to constrain the phase space in the minimization.

3.5 Instrument Impulse Response

The signal observed by the digitizers is strongly effected by the various components in the signal chain. The largest effect on signal shape is the group delay introduced by the band pass filters. To properly reconstruct the electric field that was present at the payload at the time of digitization, a full complex phasor array representing the signal chain response and antenna height as a function of frequency must be developed. This is discussed at length in Appendix B.

3.5.1 Signal Chain

Impulse response measurements for each of the 96 signal chains flown in the ANITA3 instrument were taken in Antarctica immediately preceding the flight. A broad band impulse was injected into each channel sequentially and measured by the on board flight electronics. This signal was split at the source and simultaneously measured by an oscilloscope, providing an input and output reference signal which we can compare to determine the full complex transfer function.

3.5.2 Antenna Response

An important effect on the signal measured by the instrument is the angular response of the antennas. First, there are a few considerations for determining what the optimal antenna response, the ratio of induced voltage on the 50-ohm output terminal of the antenna as a function of electric field on the front face, as a function of angle. Firstly, it is important for the antennas to have some response outside of their specific 22.5 degree wide phi sector observations region in order to establish baselines between antennas with measurable signal power. The antennas also need to have a high gain, aka a high directivity, in order to increase signal power as a function of noise power.

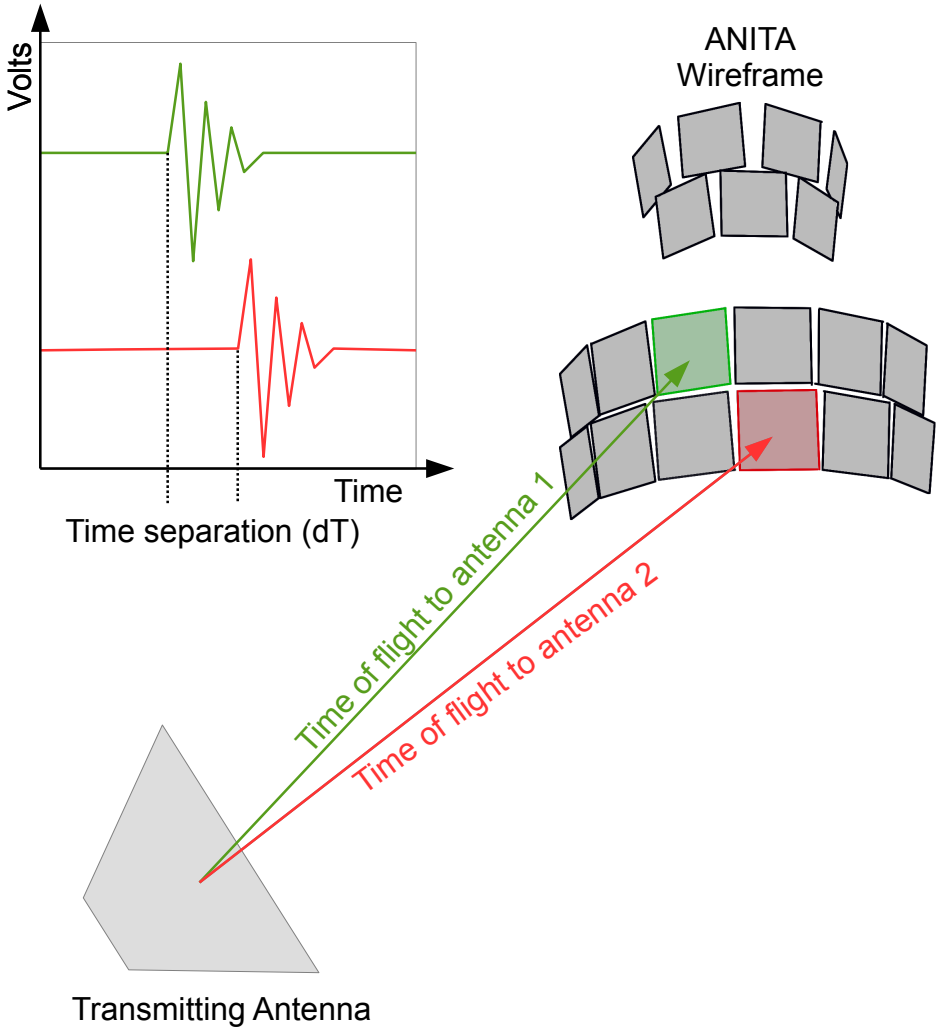

Figure 3.22: A diagram showing the process used to calibrate the electrical phase centers of each antenna by comparing waveform arrival times between antenna pairs. The pulser used was the WAIS divide pulser, which was an azimuthally symmetric autonomous station. The timing delay was determined by finding the peak correlation of the waveforms, and repeated iteratively for neighboring antenna pairs and all antennas.

If we only used omni-directional antennas, the interferometry would have additional baselines (as additional antennas would observe any specific signal source), however each channel would have a reduced signal power with an identical noise power. This also doesn't take into account the occlusion of signal by the gondola structure and discrete instrument packages residing ont the deck. In addition, there exists no perfect antenna with a step function response, and as such we need to characterize the non-ideal antennas that fly on ANITAIII.

Several different sets of measurements were made that describe the complex response (antenna gain and group delay as a function of frequency) of the flight antennas. One was done by the manufacturer (shown in Figure 3.23), and several were done on the ground in various configurations. An additional calibration was possible in flight using the calibration pulsers and the random rotations of the gondola with respect to it. The absolute measurements are most accurate in gain, as in the far field calibration measurements the absolute timing of the pulser was not synchronized. However, both gain and phase have relative measurements as a function of angle.

There are two measurements that go into the response of the antenna. The first is the S_{11} (or Voltage Standing Wave Ratio, VSWR) of the antenna. This describes the amount of input power is actually transmitted radiatively out of the antenna as a function of frequency. The second is the actual gain pattern, which describes *where* that power is transmitted in space. Both are invertible, and together can be used to determine the antenna effective height, h_{eff}, which relates the electric field incident on the antenna to the voltage transmitted to the transmission line.

Since we have interest in the polarization characteristics (described by Stokes parameters) of the signal, the antennas are designed to be measure both vertical and horizontal polarizations independently. Each polarization of the antenna has some leakage response in the opposite polarization however, especially in signals that have a combined vertical and horizontal component, which must be measured and factored out when for determining incident signal polarization. We have few to no measurements of this however. As a result, the uncertainties for polarimetry for off-axis signals are larger than are desired. In the future, it may be possible to further calibrate with in flight observations of known or constant polarization anthropogenic sources. For the goals of thesis, a search for UHECRs, these precision measurements are not required.

3.5.3 Effect of Uneven Time Sampling

The uneven sampling produced by the non-idealities of the LABRADOR sampling timebase requires the waveform to be evenly resampled before it can be properly compared to other waveforms in each event. In addition, the timing offsets induced by the physical baseline delays between antennas as well as the internal asynchronous hold times are not quantized to the sample bins, and thus must be up-sampled to be properly correlated and averaged. There are several methods that can be used to accomplish this, each with advantages and disadvantages.

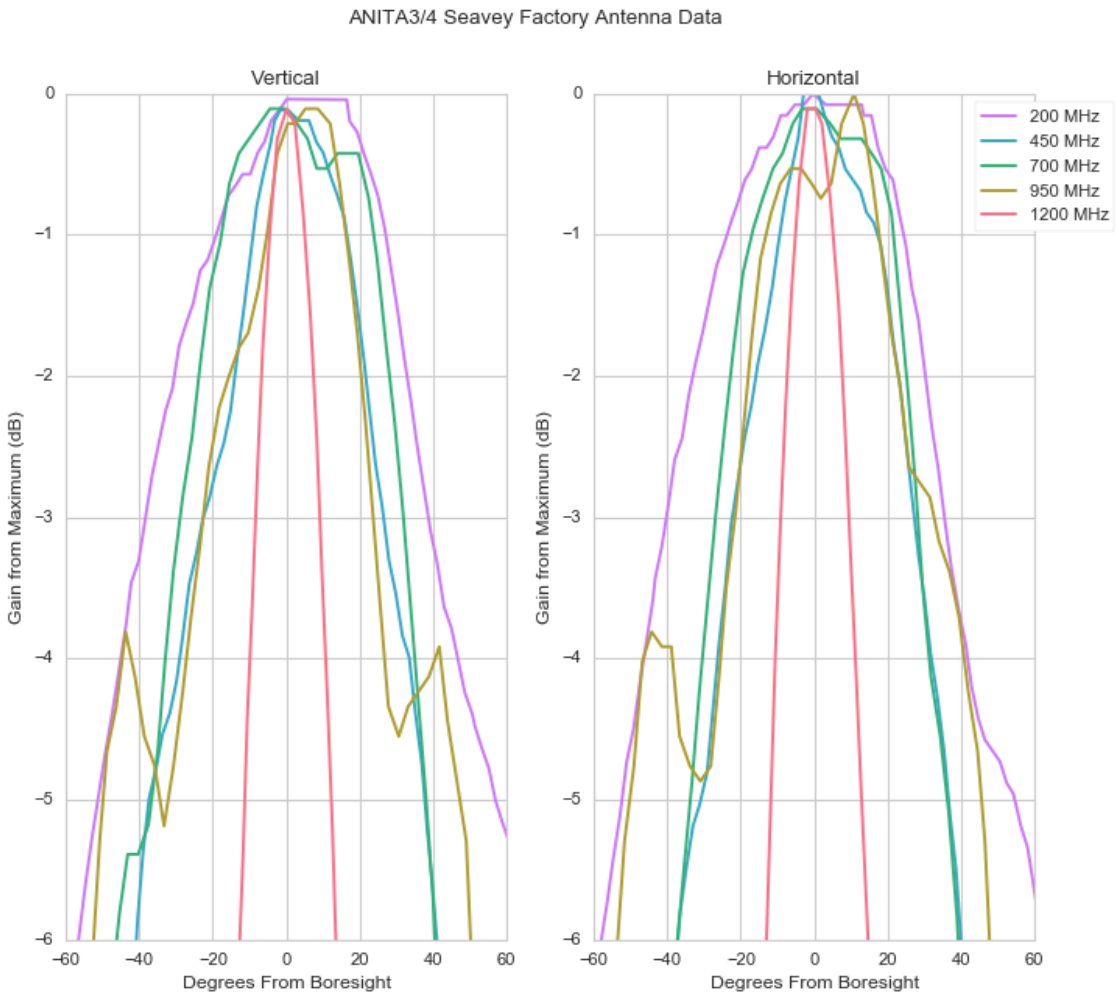

Figure 3.23: Replotted off angle antenna responses received from Seavey measured in post-manufacturing quality assurance. Gains are referenced to the maximum received power, which is not required to be the boresight angle. Responses are in reference to the electric field incident on the respective polarization ridge of the antenna.

Component	Peak Gain (dB)
AMPA	+38
iRFCM	+38
Dig/Trig split	-3
LAB split	-6
Misc. (cables, insertion loss, etc.)	-2
Total Gain	+65

Table 3.1: Link Budget. Mean values across all channels for the peak gain for each component and resulting average total gain for the signal chain.

3.6 Absolute Instrument Height

The impulse response described in section 3.5 gives the relative frequency dependent phase and gain of the instrument, however for determination of true observed signal power, the gain needs to be carefully normalized. Since the calculations to generate the impulse response require significant amounts of digital signal processing, double checking the measured signal power against sources of known power will allow us to reduce the uncertainty of eventual astrophysical particle energy estimations. This is done via two methods, first a link budget combining measurements for individual components present in the signal chain, and secondly a measurement utilyzing calibration noise sources conducted prior to flight.

3.6.1 Link Budget

A preliminary method for determining the gain of the signal chain can be done by summing the individually measured contributions of each component in the chain. Mean values across all channels for the peak gain for each of the components, as well as the resulting gain for the signal chain, can be found in Table 3.1. The expected peak gain in the signal chain from this budget is +65 dB.

3.6.2 Y-Factor analysis

Y-Factor analysis is a widely used technique for measuring the noise figure and gain of an RF network [11]. It compares the measured power of a resistor at two known temperatures with the expected Johnson-Nyquist noise power given by $P = kBT$. By calculating the slope and y-intercept of the line drawn between the two measured points, it is possible to derive both the gain and noise figure of the system. A graphical representation of this method, as well as the equations used to derive the gain and noise figure of the network, are shown in Figure 3.24. For this thesis only the gain will be discussed, as it is most relevant to the subject analysis.

For ANITA-III, two noise diodes were used as calibration sources previous to flight, a +15dB$_{ENR}$ source and a +5dB$_{ENR}$ source. These diodes have two states, "on" and "off", with a large calibrated difference in thermal Johnson-Nyquist output power between them. For reference, the units of

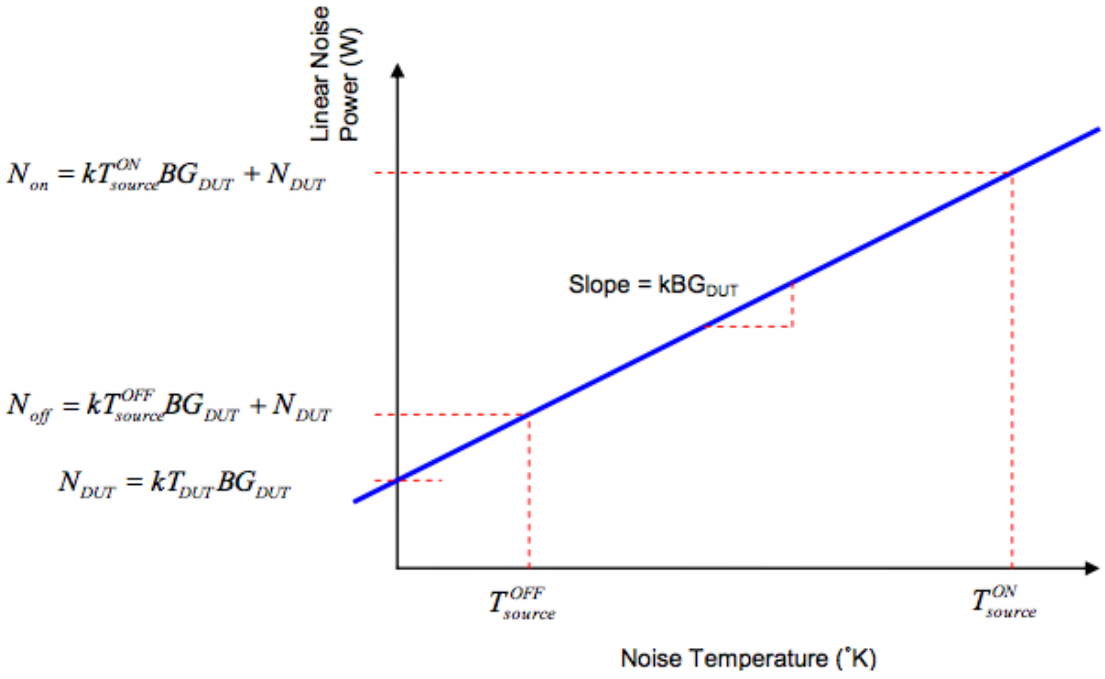

Figure 4: Diagram showing the Y factor variables.

Figure 3.24: Graphical representation of the Y-factor analysis method for determining the noise figure and gain of an RF network [11].

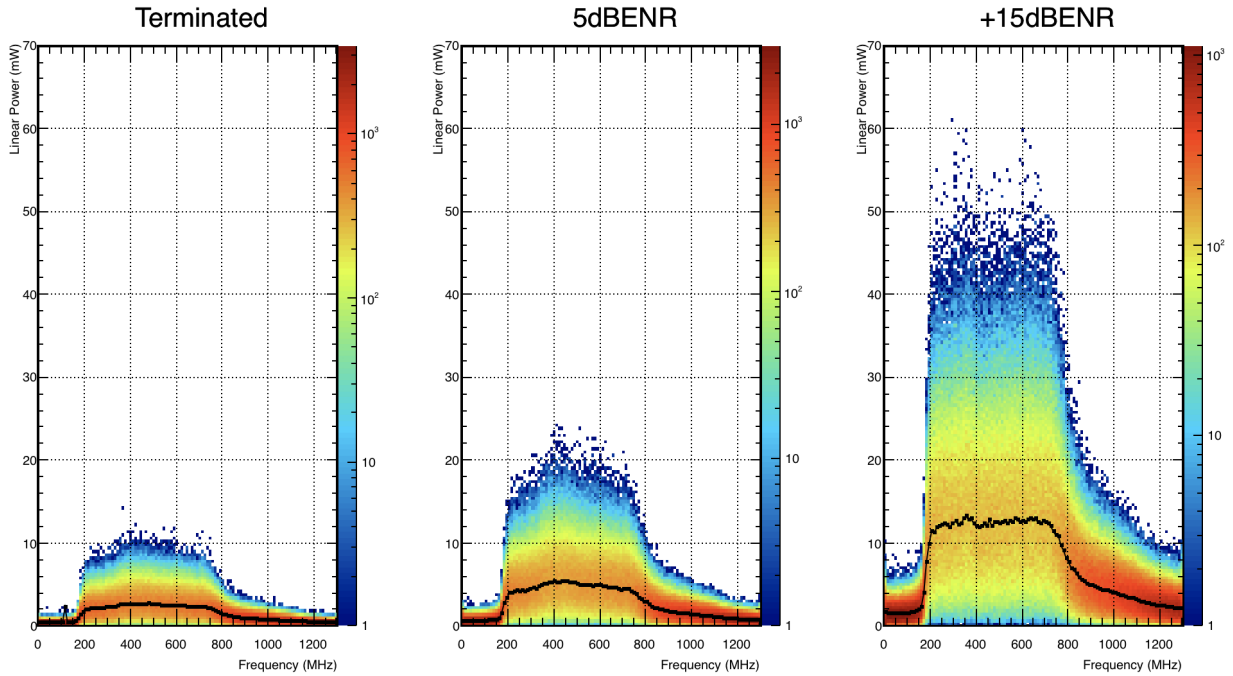

Figure 3.25: Two-dimensional histogram of values from pre-flight noise diode waveform measurements for a single example channel. Measurements are done with the full waveform calibration applied. Waveforms are interpolated via an Akima spline to an evenly sampled 2.6 GS/s. The black line denotes the peak of a Rician fit to the one-dimensional histogram of each frequency bin. This fit value is used to calculate the noise figure and gain via the y-factor method. Waveforms were captured at 10Hz, and this measurement was done for all 96 channels.

dB_{ENR} correspond to the logarithmic flat spectrum gain between the output noise power when the noise diode is in its activated state, and when it is powered off and is at room temperature. For example, a $+5\mathrm{dB}_{ENR}$ source would have an output power while off, over the 1 GHz of ANITA bandwidth at a room temperature of 20C, of 4.04 pW (-83.9 dBm). The source, when on, would then be 5dB higher than that value, or 12.8 pW (-78.9 dBm). Translating from a noise power to a noise temperature can be done using the Johnson-Nyquist relations, yielding $T = \frac{P}{kB}$. Using this, a $+5\mathrm{dB}_{ENR}$ noise source has an effective temperature of 927 K.

Since we have a time domain digitized waveform, it is possible to use Fourier analysis to split each digitized waveform into multiple frequency bins, each of which can be used to calculate the respective gain and noise figure. A plot showing the resulting measured noise power for each of the two diodes, as well as a 50Ω terminated measurement, can be seen in Figure 3.25.

Equation 3.2 gives the formula for calculating the gain G of the system using the y-factor method, where N is the measured linear noise power in mW, and T is the effective temperature of the noise diode.

$$G = \left(\frac{N_{on} - N_{off}}{T_{on} - T_{off}} \right) \frac{1}{kB} \tag{3.2}$$

Using this equation and the calibration measurements, the system gain was found. An example of the results is shown in Figure 3.26. The peak mean gain for all channels is +63.9dB. A two-dimensional histogram of the gains of all channels can be seen in Figure 3.27. This measurement is in agreement with the values determined via the Link Budget.

3.6.3 Instrument Height

The final instrument height can be found by taking the convolution between the absolute signal chain gain and the antenna height, discussed in subsection 3.5.2. This is a function of the incidence angle of the propagating electric field, and must be calculated on an event by event basis. If multiple antennas are used to create a summed waveform the calculation must be done for each one separately, then convolved to determine the resulting conversion between measured voltage and electric field.

3.7 Triggering Validation

Determining whether the trigger circuit is effectively discriminating between signal and noise-type events and reliably latching the digitizer circuit is important to determine the overall sensitivity of the instrument. The first level L1 triggers are dependent on signal strength, making any trigger efficiency likewise SNR dependent. The L2, and L3 triggers to an extent, are timing-dependent, creating an elevational dependance to the trigger efficiency. Calibration measurements were taken on the ground before flight, with additional measurements from ground based calibration pulsers taken during flight.

3.7.1 Windowed trigger testing

Before flight, a programmable delay line was used to inject impulses at various picosecond-precision-delayed impulses into all like polarized channels of two neighboring phi sectors to test for the true angular response of the payload. Due to the complexity of firmware, it is important to discern the exact L0 timing coincidence regions for which a global trigger will be generated. Results of these measurements can be seen in Figure 3.28.

3.7.2 Overall Trigger Efficiency

The trigger efficiency of the payload can be determined by measuring the number of recorded WAIS calibration and dividing it by the number of expected signals during that same time period. The results of this, as a function of SNR, can be seen in Figure 3.29. The reduction in efficiency between

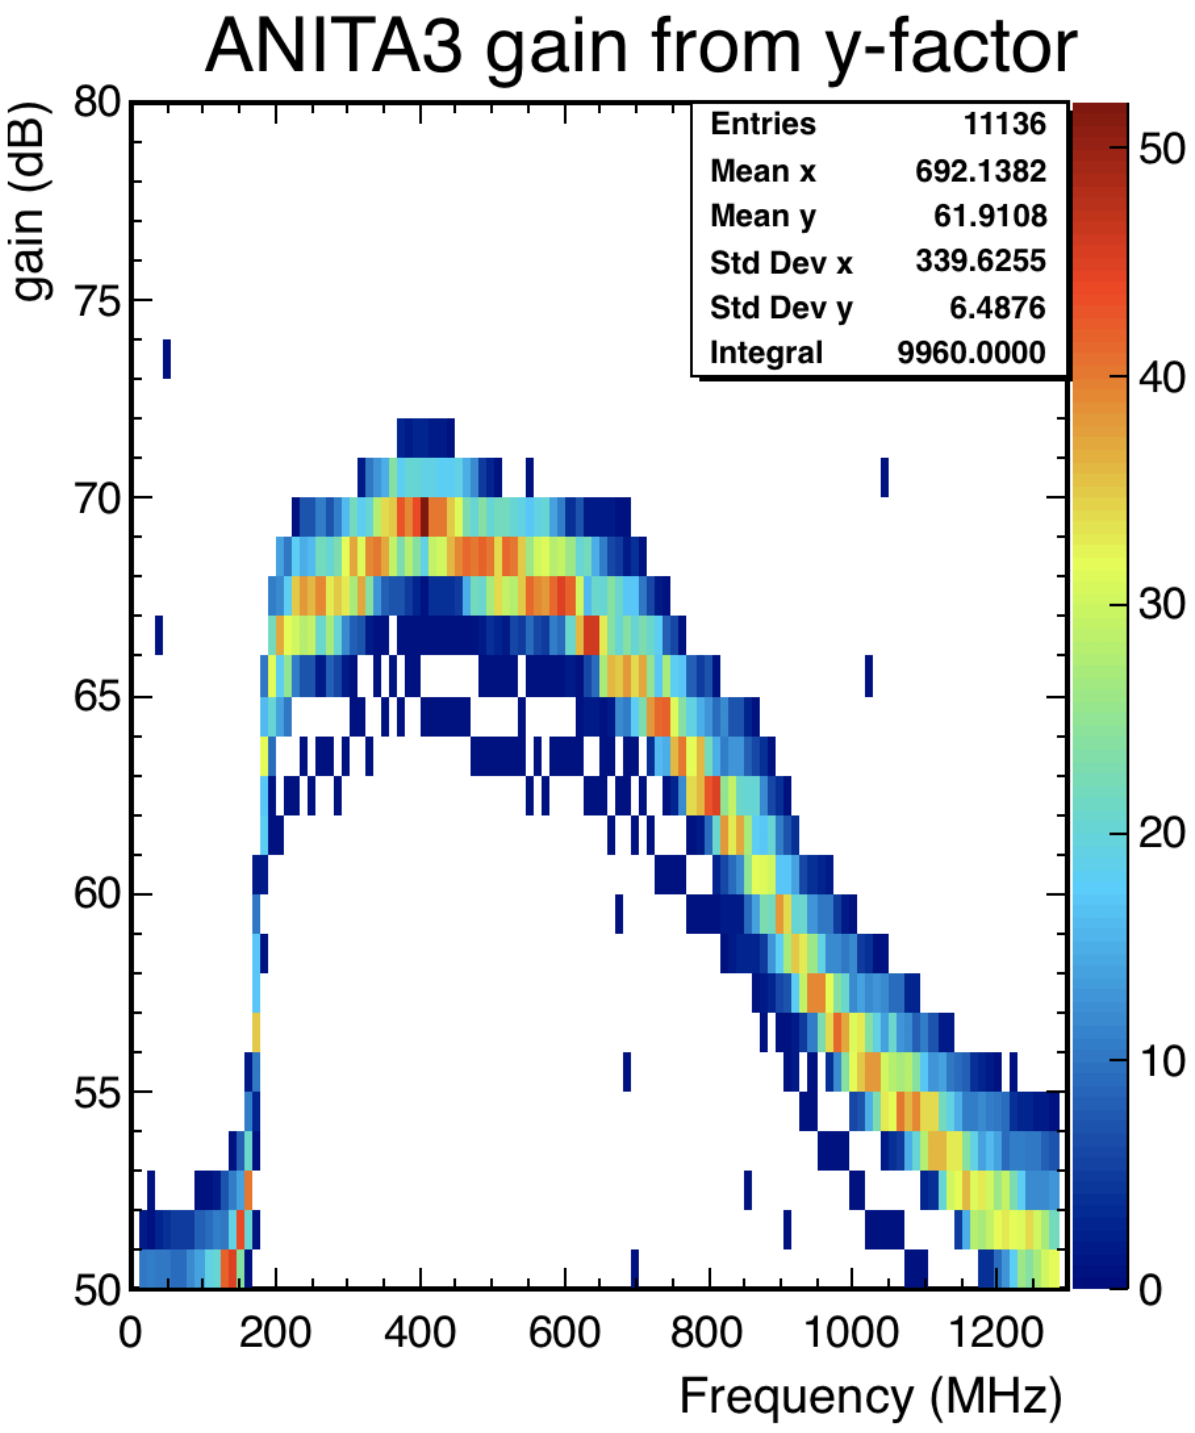

Figure 3.26: Two-dimensional histogram of values for the gain of a single channel found using y-factor analysis and a calibration data taken prior to the flight.

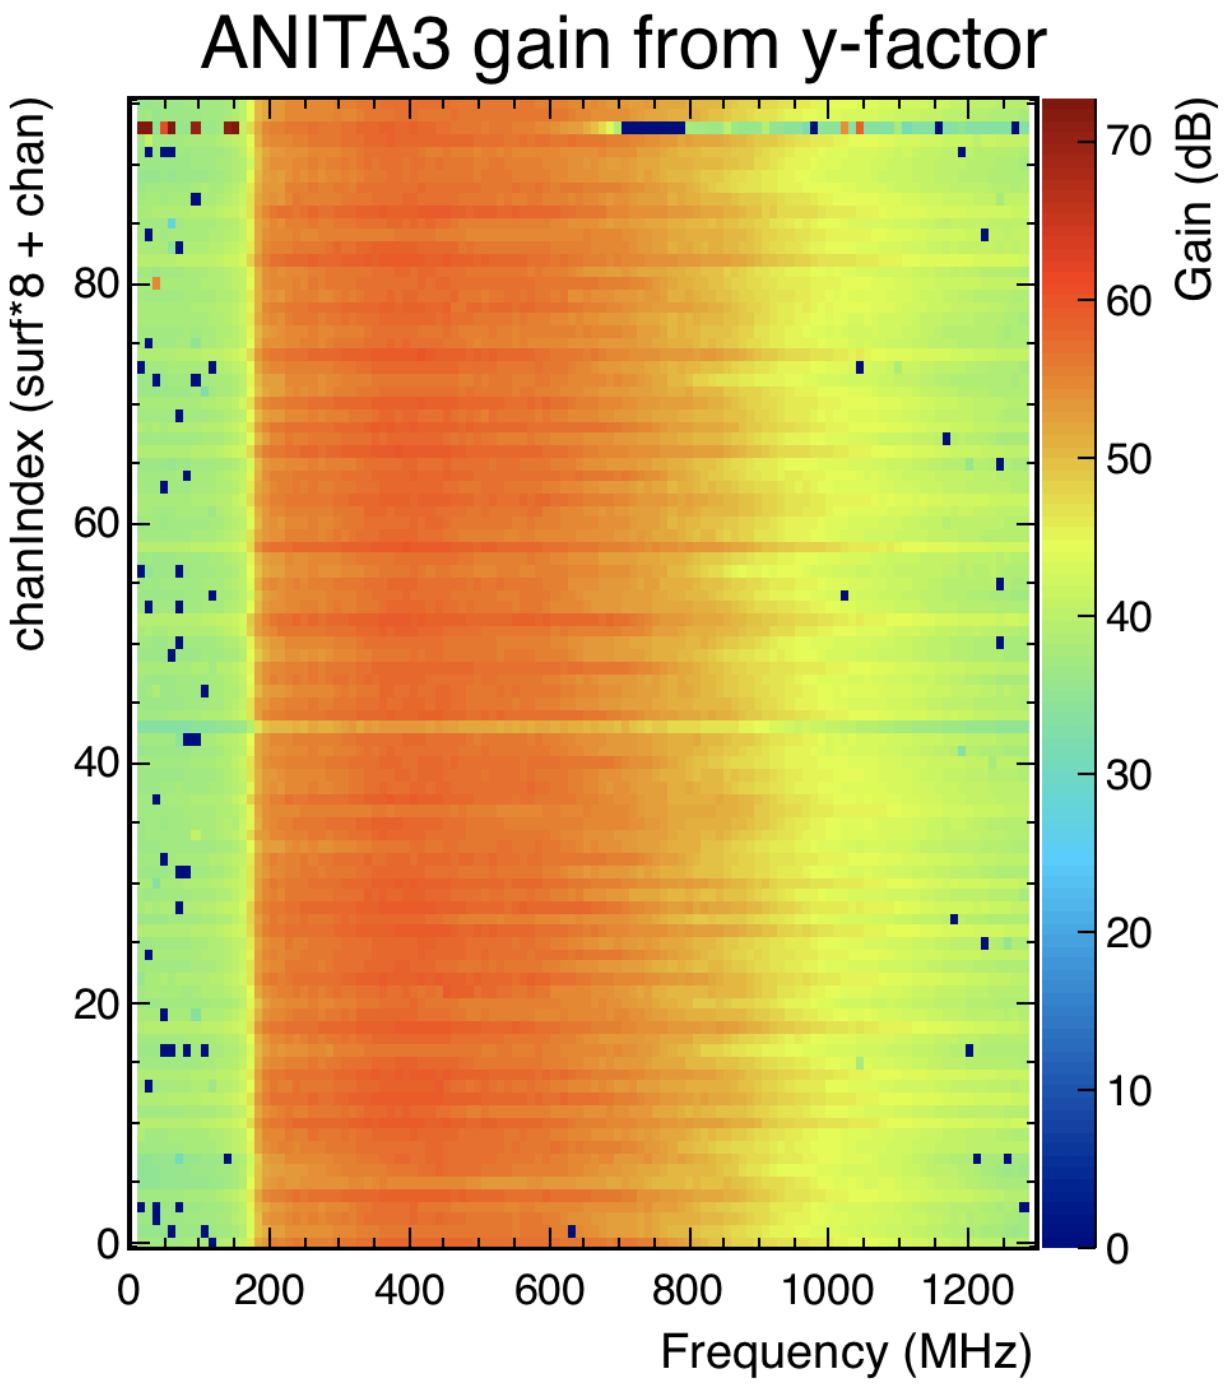

Figure 3.27: Two-dimensional histogram of values for the gain of all channels found using y-factor analysis.

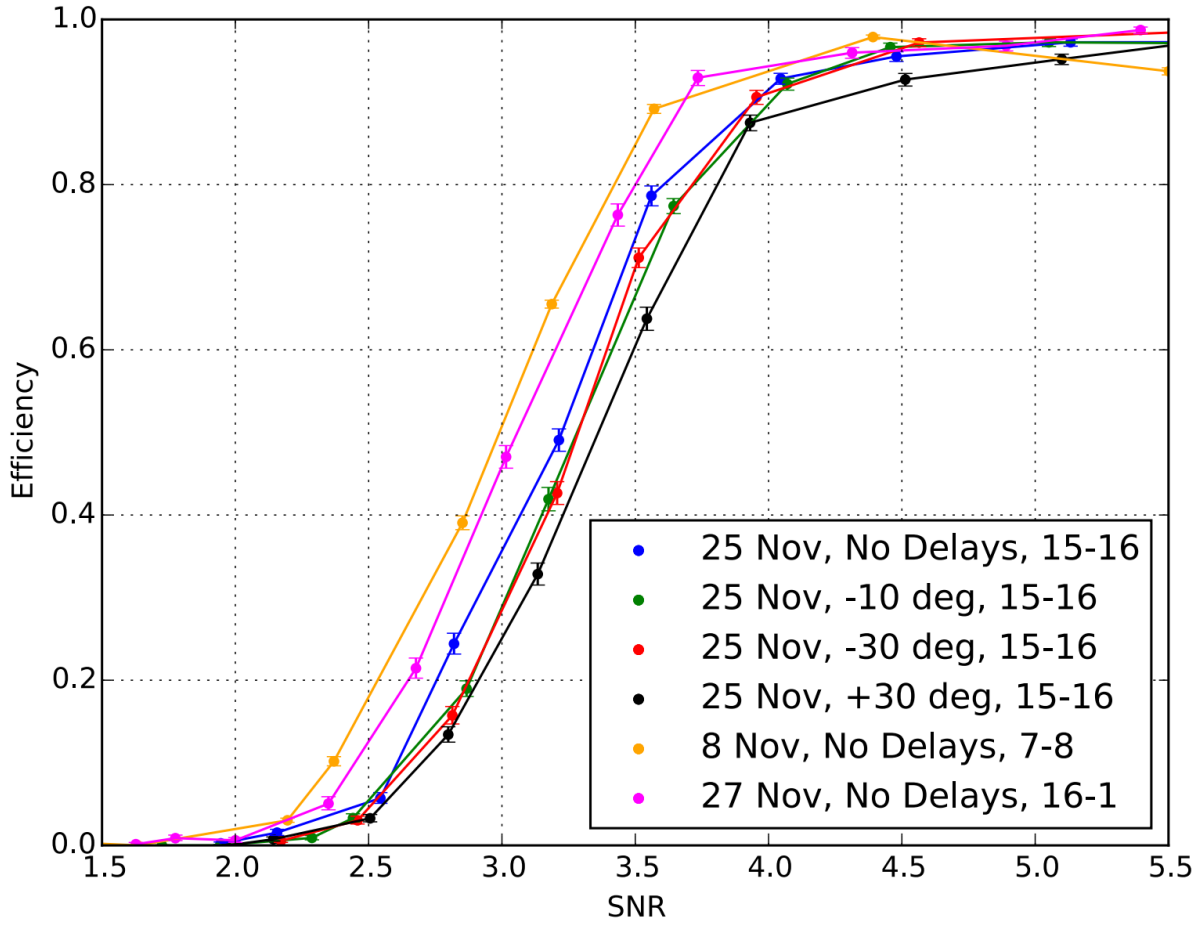

Figure 3.28: Global trigger efficiencies for injected impulsive signals with varying timing delays pre-flight. Identical signals were injected at 10Hz into six channels in two neighboring phi sectors with varying timing delays. The efficiency of any signal that falls within the triggering window is shown in this plot to be approximately equal across the entire window. Plot courtesy of Keith Bechtol

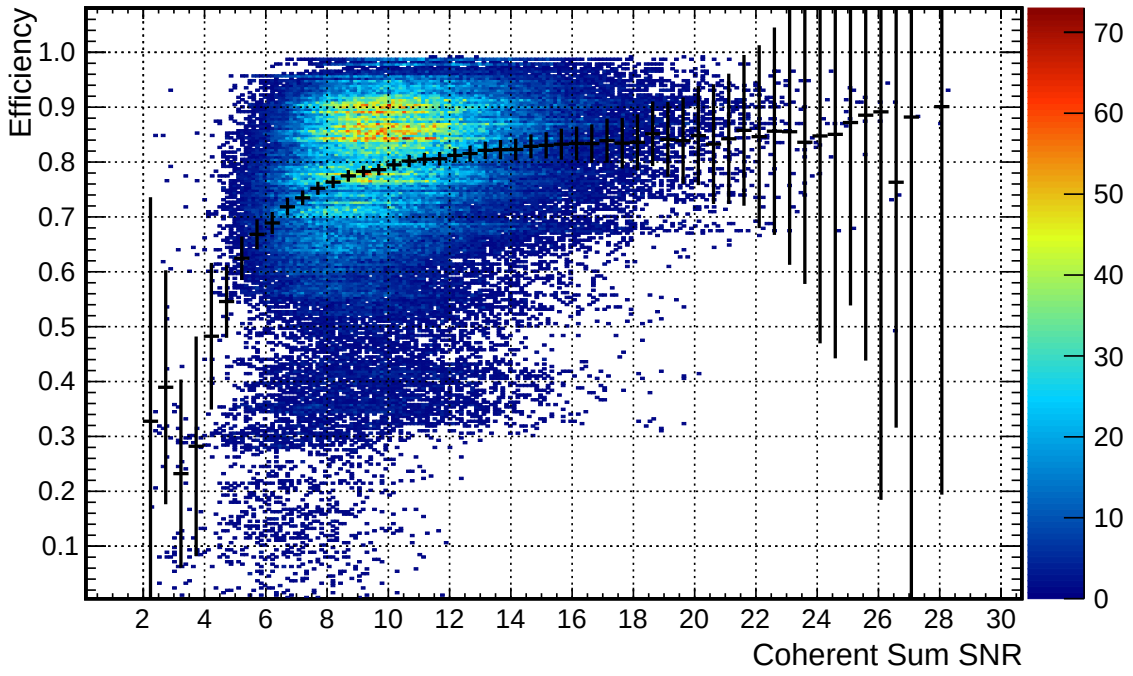

Figure 3.29: The global trigger efficiency to WAIS calibration pulser events as a function of SNR. Error bars represent the standard deviation of the underlying histogram. The color map is the number of events in each bin. The 50% efficiency point occurs at approximately 4.5

pre-flight injected signal calibration (3.5) and in-flight measurements (4.5) can be attributed to anthropogenic noise, both CW and impulsive, causing large variations in L1 thresholds.

3.8 Non-uniform Channels and Outliers

There are several channels in the ANITAIII instrument that are non-uniform. The first, and most obvious, is the channel in which the ALFA is digitized. The desire to have a 97th low frequency drop down omni-azimuthal responce digitizing antenna required an additional channel, however there were none available. To accommodate this, a modified filter was placed on channel 05TH that limited the nominal input signal to 800Mhz. The low frequency (40-80MHz) antenna was then heterodyned with a 900Mhz local oscillator (LO) which produces a upsampled beat pattern at $f_{LO} - f_{ALFA}$ and $f_{LO} + f_{ALFA}$. This channel must be treated slightly differently than the rest in a few ways. First, the signal must be digitally filtered to remove the ALFA component before it is correlated or averaged with similar antennas. Secondly, in simulation the reduced bandwidth (and subsequent lower noise and signal power) must be taken into account.

There are also 5 “bypass” channels that, while functionally equivalent to the others, should be

considered with a bit of caution. In these channels, some component of the signal chain within the iRFCM module failed, yielding a digitized value that had significantly lower, or essentially zero, power during calibration. As removing the entire iRFCM box so close to the launch date was untenable, additional second stage amplification units were built and installed along side the iRFCMs utilyzing spare bulkhead connectors on the outside of the instrument box. These channels were then re-calibrated, and do not differ noticeably from the original population.

From the results of the calibration, there exists one additional outlier channel. The time domain impulse of 13BH differs from the remainder of the channels for unknown reasons, though a filter is suspected.

3.9 Interferometric Pointing Resolution

An ultimate check on the calibration of the instrument can be seen in the precision and accuracy of interferometric pointing to known impulsive sources. Using the WAIS pulser events, distributions of event arrival direction and their separation from the true location of the calibration pulser. The pointing resolution is a function of the SNR of the signal. Distributions of the pointing accuracy and precision can be seen in Figure 3.30.

The full interferometric reconstruction method used to determine these values is described in more detail in the Analysis chapter.

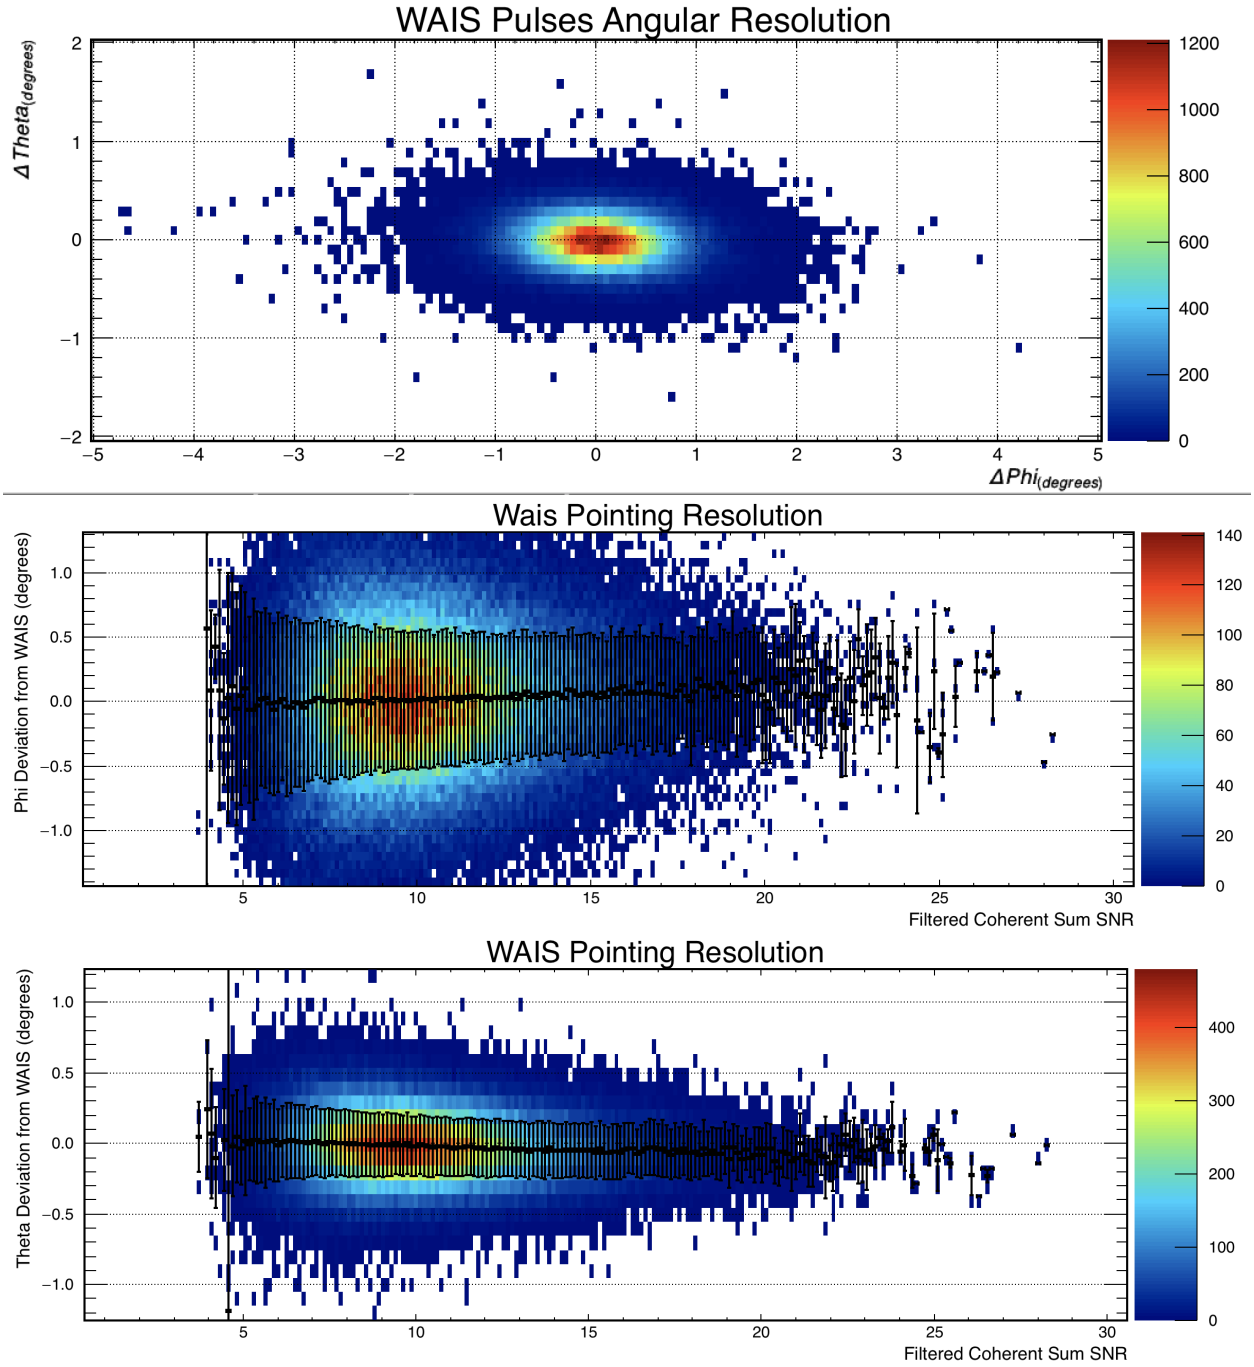

Figure 3.30: (Top) Interferometric pointing angular spread for horizontally polarized WAIS pulses. Additionally plotted as a function of SNR for azimuth (middle) and elevation (bottom). Error bars represent the standard deviation of the distribution. The color map shows the number of events captured with that SNR and angular deviation from expectation.

CHAPTER 4
ANALYSIS OF ANITAIII FLIGHT DATA

4.1 ANITAIII Analysis Overview

4.1.1 Analysis goals

The physics event search analysis detailed in this thesis was designed to find "cosmic ray like" events specifically. The flux and energy spectrum of CRs has been well measured, and there are already published observations of cosmic rays from previous ANITA flights. The blinding requirements and associated drawbacks of doing a full blind analysis for previously unobserved impulsive signals, such as in-ice ν_e and ν_μ interactions, can be greatly reduced by approaching the analysis in this manner. This analysis is designed to provide the highest possible exposure to a very specific signal, that of modeled and observed CR EAS impulses. This has the consequence of reducing sensitivity to signals that do not fit into the expectation provided by measurements and models of CR EAS radiation patterns.

This search is expected to be sensitive to two separate astrophysical particle interactions described in 2.1.1. These are UHECRs interacting in the atmosphere and viewed directly or as a reflection off the continental ice sheet, and upgoing UHEν_τ showers initiated by tau lepton decays arising from ν_τ interactions in the ice or rock. The shower profiles of the two events are expected to be very similar [45], with the primary difference being absolute polarity of the signal. An upward traveling UHEν_τ observed below the horizon would have radiation measured directly, whereas a UHECR viewed below the horizon would be observable in a reflection. This reflection would induce a polarity flip in the electric field oriented parallel to the surface of the earth. Since the observed signal is dominated by geomagnetically induced radiation, which at the poles are primarily vertically aligned, the strongest E-field component would be horizontally polarized.

The results of this analysis will be an accounting of the signals found using this search method and the estimated exposure the ANITA-III instrument had during its flight to this type of astrophysical event. The polarity of the signals will also be identified, offering a measurement of the UHEν_τ flux. Finally, a background estimate that is novel to the ANITA experiment will be introduced and applied to the data.

4.1.2 Analysis outline

Analysis of the data recorded by the ANITAIII flight instrument requires a significant computational effort to extract a meaningful physics result. The digital values recorded by the instruments on-board the payload, when properly calibrated and used in conjunction with one another, can establish whether any electromagnetic signals were traversing the payload at the time that the trigger

system latched the digitizers. The combination of geometric antenna offsets and their corresponding timing delays allows the discrimination of signal versus noise, as well as precision reconstruction of any incident signal. Using a variety of mathematical methods and an understanding of the expected signal, the nearly 85 million events, consisting of 8.16 billion waveforms, captured during the flight can be pared down to a handful of likely CR candidates, which can then be examined in greater detail. This chapter details the calculated values, their motivations, and the methods used to conduct the cosmic ray search.

A diagram that depicts the order in which analysis steps are taken, as well as their relavant software packages and intermediate files, is shown in Figure 4.1.

4.2 Software Package Overview

The analysis was done with a variety of software packages created by the ANITA collaboration. The relevant packages used for the analysis described in this thesis are described in the following section. They range from tools for interfacing with raw data written to disk, to ANITA specific calibration and interferometric mapping, to this-thesis specific packages for template correlation, clustering, and background estimation. The packages are described in order of dependency, with most independent packages listed first.

The cosmic ray search was done primarily in C++ using libraries from the ROOT Data Analysis Framework [46]. ROOT, developed at CERN, provides a variety of data storage objects and plotting functions optimized for physics experiments.

4.2.1 ROOT FFTW interface: libRootFftwWrapper

Fourier transforms were done using the the Fastest Fourier Transform in the West (FFTW) package. This package is not designed to be used with the ROOT objects that the remainder of the packages are based around, and so an interface library, named libRootFftwWrapper is used [47]. This package not only provides an interface however, it also includes a large number of ANITA-related numerical recipes.

4.2.2 Event Reader Software: eventReaderRoot

The data taken by the ANITAIII instrument is written in a compressed C++ binary class format to disk. This data format requires minimal computational load and very fast write speeds, with minimal storage space overhead. This data format requires a standardized class structure that is stored in the eventReaderRoot code library [48] in order to accurately read and extract any values. The raw data taken in flight is converted to ROOT objects, which increases data read speed and enforces data quality, with the tradeoff of slightly increased storage space and a one time computationally complex conversion. The eventReaderRoot library also stores information about

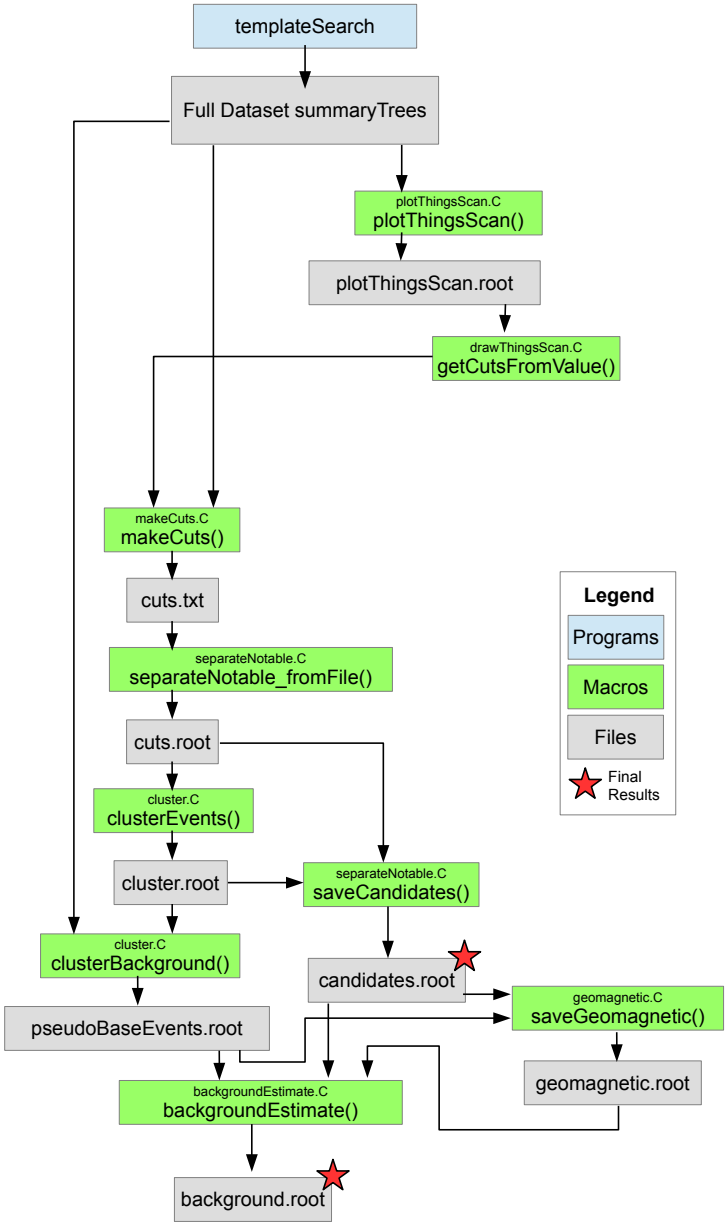

Figure 4.1: A diagram depicting the order in which analysis steps are taken, the intermediate storage files saved to disk, and the function names and ROOT macros where the bulk of the code is located. The calculations and algorithms contained within this code-base is described in detail throughout this chapter. The storage locations of final results are noted with red stars.

each flight, such as the number of antennas and their mapping to RF channel, as well as tools for applying the calibration constants to waveforms.

4.2.3 Coordinated Analysis Framework: AnitaAnalysisFramework

In order to facilitate active collaboration between multiple analyzers, much of the analysis software was consolidated into a library named the Anita Analysis Framework [49]. It provides a standardized set of C++ classes for storing and calibrating digitized LABRADOR waveforms in memory, implementing filtering algorithms, implementing agreed-upon common analysis tools, and a shared output format.

This framework, most importantly, provides a C++ summary class with agreed upon reduced quantity values of interest. This object, named the AnitaEventSummary, is created for each event taken during the flight, and is designed to be filled by any analysis routine with all values required to do a full physics analysis of the flight data. This prevents computationally burdensome calculations used in calibration and interferometry from being repeated multiple times.

4.2.4 University of Chicago Event Analyzer: UCorrelator

The data taken by the ANITA instrument requires processing to be useful in a physics analysis, and many of the computational techniques can be shared across multiple analyses. For this analysis, I have used, and collaborated, to the UCorrelator software package developed primarily by Cosmin Deaconu at the University of Chicago [50]. With the addition of my own bug-testing and feature expansion, this package has been used to do the base-level numerical computations for this analysis.

4.2.5 Event Display: anitaMagicDisplay

The anitaMagicDisplay package [51] allows a user to quickly and easily display many of the most relevant raw and reduced data values for any event in the flight. The package provides a Graphical User Interface (GUI) that can simultaneously display waveforms captured from the entire instrument sorted in various ways. It also includes dynamically selectable filtering strategies, maps and summaries of multiple interferometric reconstruction techniques, and payload orientation information.

4.2.6 Miscellaneous utility functions: anitaEventCorrelator

A separate library that includes some additional important functionality is anitaEventCorrelator [52]. The most important class in this code is UsefulAdu5Pat, which takes raw GPS sensor information, as well as Earth geometry and Antarctic topographical information, and provides a method for mapping spherical payload coordinates to a point on the continent, and vice-versa.

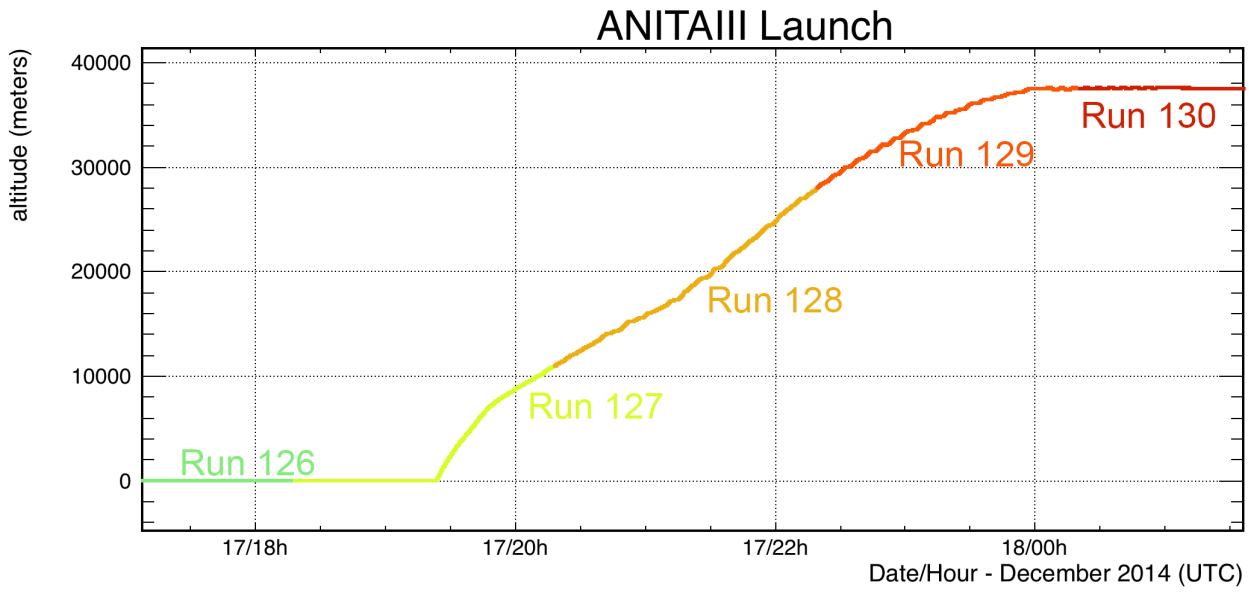

Figure 4.2: Altitude data captured by the on board GPS on launch day. Colors refer to the run number of that data. Run 130 is the first run where the instrument was in full data-taking configuration and at float altitude.

4.2.7 Thesis analysis specific: benCode

The eclectic accumulation of code used in this analysis is stored in a variety of libraries, macros, and compiled programs collectively named benCode [53]. Though there are many offshoot routines for debugging and the determination of specific values of interest, the main path for the analysis is summarized in Figure 4.1.

4.3 Dataset and Blinding Procedure

4.3.1 Dataset for analysis

The ANITAIII instrument was put into flight configuration on December 9th 2014, which marks the beginning of the run numbering. However, due to non-optimal launch conditions, the payload was not launched for another week. During this time, the system remained on, taking data at a reduced rate for last minute data validation and bug fixes. The successful launch on December 17th 2014 at 19:23 UTC is shown in Figure 4.2. The pre-amplifiers were turned on midway through the ascent in order to protect them from the high powered anthropogenic noise at the launch facility. Run 130 is the first run at altitude with all systems operating in final flight configuration. Run 439 was the final run recorded to disk before the instrument was powered down for descent. The final event was taken on January 9th 2015 at 05:52 UTC.

During the 22 days and 10 hours of active data collection at altitude, between run 130 and

run 439, ANITAIII captured 78,630,542 events. Runs 257 through 263 contain no data due to a bug in the Prioritizer daemon causing the Ramdisk to fill, and preventing the acquisition software from recording events. This outage occurred on Christmas day, and resulted in an 11 hour data blackout, from December 25th at 17:13 to December 26th at 04:28.

For the cosmic ray analysis, only events that fired the horizontally polarized L3 trigger are considered. The geomagnetic radiation that is predicted to dominate the electromagnetic emission of an EAS is expected to be mostly horizontally polzarised. 41,566,966 events fall into this dataset. 360,055 events fired both the horizontal and vertical L3 trigger, however they are not treated in any unique manner.

Additionally, the payload initialized a digitization twice per second regardless of the RF trigger, once initiated by the GPS second and once initiated by the CPU. There were 3,551,509 of these "minimum bias" events recorded during the flight. These events are also processed and used to determine the background noise environment during the flight and set cuts on thermal background.

4.3.2 Philosophical reasoning for blinding and past methods

Much discussion took place within the collaboration in order to prevent unwanted analyzer bias from being a significant effect of any analysis result. Unconscious bias is an inherently human trait, the cold unfeeling rules of nature do not bend to the whims of our emotions or desires. However, by allowing analyzers to pick and choose methods and procedures and cuts, it is possible that a result may unwittingly reflect those same desires and emotions. For example, an ambitious physicist, who desires to further his career, may see the opportunity to do so through an exciting and controversial result that makes other scientists take note of his work. On the other extreme, an exhausted physicist, who does not want to defend a controversial result, could "play it safe" and tune his analysis to find a result in agreement with other experiments. Both these biases are undesirable, but both could be subtly effecting results though no fault of either scientist.

The solution to this problem is to introduce a method in which the result is entirely outside of the experimenters control until all analysis procedures are complete. In previous ANITA analyses, this was accomplished by "hiding" 90% of the signal data, only allowing analyzers to use a 10% fraction of the full set to tune their cuts. Once the cuts were set to meet a desired background rejection fraction, the reminder of the events that did *not* meet the cut requirements would be revealed. A statistical method, known as the ABCD method, could be used to estimate the number of background events within the signal box from purely random statistical fluctuations. Any number of events in excess of this background estimate would then be considered signal. This method had several drawbacks. First, an inability to determine signal from background on an event by event basis leaves a possibility that the most interesting events are background. Additionally, it allows anthropogenic background to escape observation within the 10% data set, requiring it to be removed manually by hand later, re-introducing bias.

4.3.3 ANITAIII CR blinding technique

For the ANITAIII CR search, it was decided that a 10% - 90% blinding requirement was unnecessary. The energy spectrum and flux density of CRs are measured well up to the UHE range and the estimated number of observable events for the flight is much larger than one. Additionally, both ANITAI and radio extension projects at terrestrial cosmic ray observatories have measured the radio EAS signature of a CR interaction with the atmosphere. The ANITAIII results have also been subject to a fully blind analysis previously in a neutrino search that reported 4 transient events matching signal. For this reason, a novel and more lax blinding procedure was agreed to for this search.

The signal of greatest interest in this analysis is a regenerated tau neutrino transiting through the earth and hadronically decaying in the atmosphere. Though the primary particle is different from a CR EAS, the shower will evolve in a similar fashion and produce a similar radio pulse. The defining characteristic of this signal is a polarity with the same sign as a direct EAS shower observation, but with an incidence angle below the horizon (beyond misreconstruction uncertainty). To prevent any bias in this signal channel, the polarity of all events has been randomly inverted. This will create a random distribution of polarities, and prevent an analyzer from determining whether any particular event is of interest.

Additionally, only signals that triggered the horizontally polarized L2 trigger are analyzed to preserve the integrity of the neutrino search. Observations of CR EAS signals are expected to be primarily horizontally polarized, while in-ice neutrino interactions are expected to have primarily vertically polarized observed signals. This also helps to reduced the computational workload, as each polarization trigger accounted for approximately half of the overall events.

4.4 Digital filtering

The stored digital waveforms can be improved by the removal of well understood background signals. Anthropogenic digital communications often take the form of circularly polarized CW sources, and are presently superimposed in much of the data. Their sources include both manned and unmanned bases distributed throughout the continent, as well as orbiting satellites visible above the horizon. These CW signals have a significant effect on both the trigger circuit and the post-flight analysis of digitized data. A spectrogram of a period of the flight can be seen in Figure 4.3, with the CW peaks from satellite noise clearly visible at 220MHz and 380MHz.

The effect these CW signals have on the trigger is twofold. Firstly, they compress the tunnel diode by increasing the total power of the input signal, pushing the diode response out of its square-law range. Additionally, CW sources force an increase in the threshold voltage used on the comparator circuit used to trigger on the signals in order to maintain an acceptable global trigger rate. In many cases, the amount of power from CW sources forces mulitple phi sectors to self-mask,

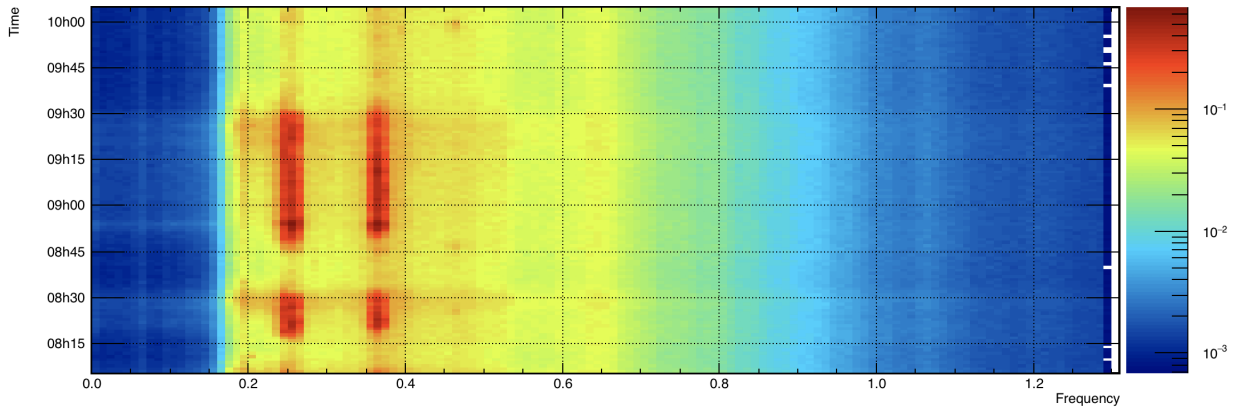

Figure 4.3: A spectrogram of the measured RF signal from a single antenna pointed north during the beginning of run 371. The Z-axis is linear power. The red regions at 220MHz and 380MHz are frequency bands with high power due to orbiting satellites.

as no DAC threshold setting can appropriately limit the L2 trigger rates from a specific phi sector. These negative effects cause a decrease in the sensitivity of the instrument to physics events with low signal power, and decrease the overall instrumented volume.

Despite these two effects, impulsive candidate signals can still be measured while superimposed on the CW source. In these cases, the CW causes the interferometric image to be "pulled" towards the source of the signal while also presenting false peaks in the image. To remediate the effects of these signals on the pointing and post-flight analysis, we can apply a digital filter that removes specific frequency bands.

4.4.1 Sine Subtraction Filtering

The filtering technique used in this analysis has been named sine-wave subtraction. Digital Finite Impulse Response (FIR) filters introduce dispersion, which will negatively effect the reconstruction accuracy and final SNR of any incident physics signal. To alleviate this, a method was developed which involves fitting the phase, amplitude, and frequency of a sine wave to each waveform, then subtracting that sine wave from the signal. This is repeated iteratively until the resulting subtracted wave has a negligible fractional power decrease. This acts to remove single frequency bands from each channel without adding a significant amount of dispersion, nor diminishing physics signal or thermal noise power. The drawback of this method is that it requires multiple fits per waveform, which is computationally expensive. The computations for this filtering strategy takes approximately half of the total computation time. The framework and coding of the SineSubtract algorithm was done by Cosmin Deaconu.

The sine subtraction works by iteratively fitting a sine wave to each captured waveform individually, minimizing on the amplitude, phase, and frequency. To aid in computational efficiency, the

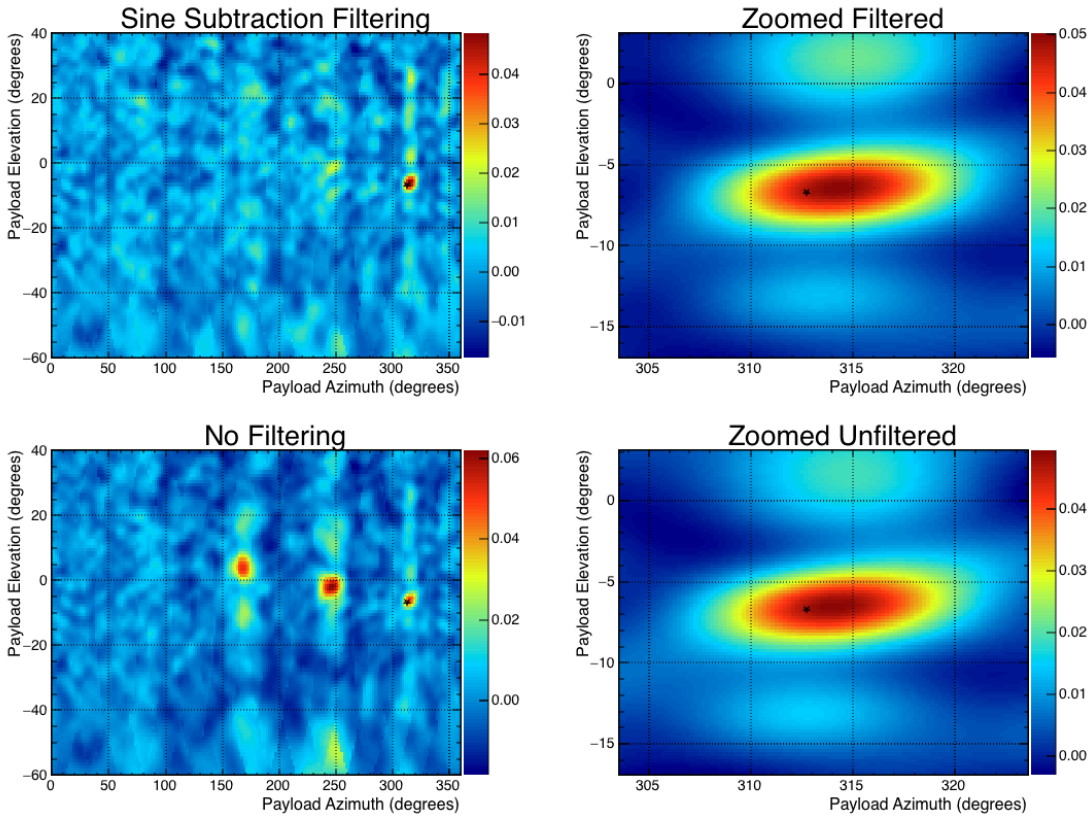

Figure 4.4: The effect that sine subtraction filtering has on interferometric mapping. This particular event, 56765803, is a WAIS pulser event. 35.9% of the initial power was removed in the filtering process. The filtering removes the two largest CW peaks from the map, leaving only the impulsive signal peak.

range of frequencies and amplitudes that the fit could minimize over were adaptively altered based on a previously generated full flight frequency spectrogram. This adaptive component has a configurable "exponent" variable that controls how strictly This best-fit function is then subtracted from the waveform, yielding data with a single frequency component removed. This process is repeated until the ratio of the power subtracted from the waveform over the total power in the original waveform is below some configurable threshold a specified sequential number of times. For this analysis, a ratio of 0.1 was used, and 3 failed iterations were allowed. The key used to generate this filter in the AnitaAnalysisFramework is "sinsub_10_3_ad_2".

An example of the effect that the filtering has on the interferometric map of a calibration pulser event can be seen in Figure 4.4. An example of the sine fit and removal over five iterations for a single waveform of the same event can be seen in 4.5

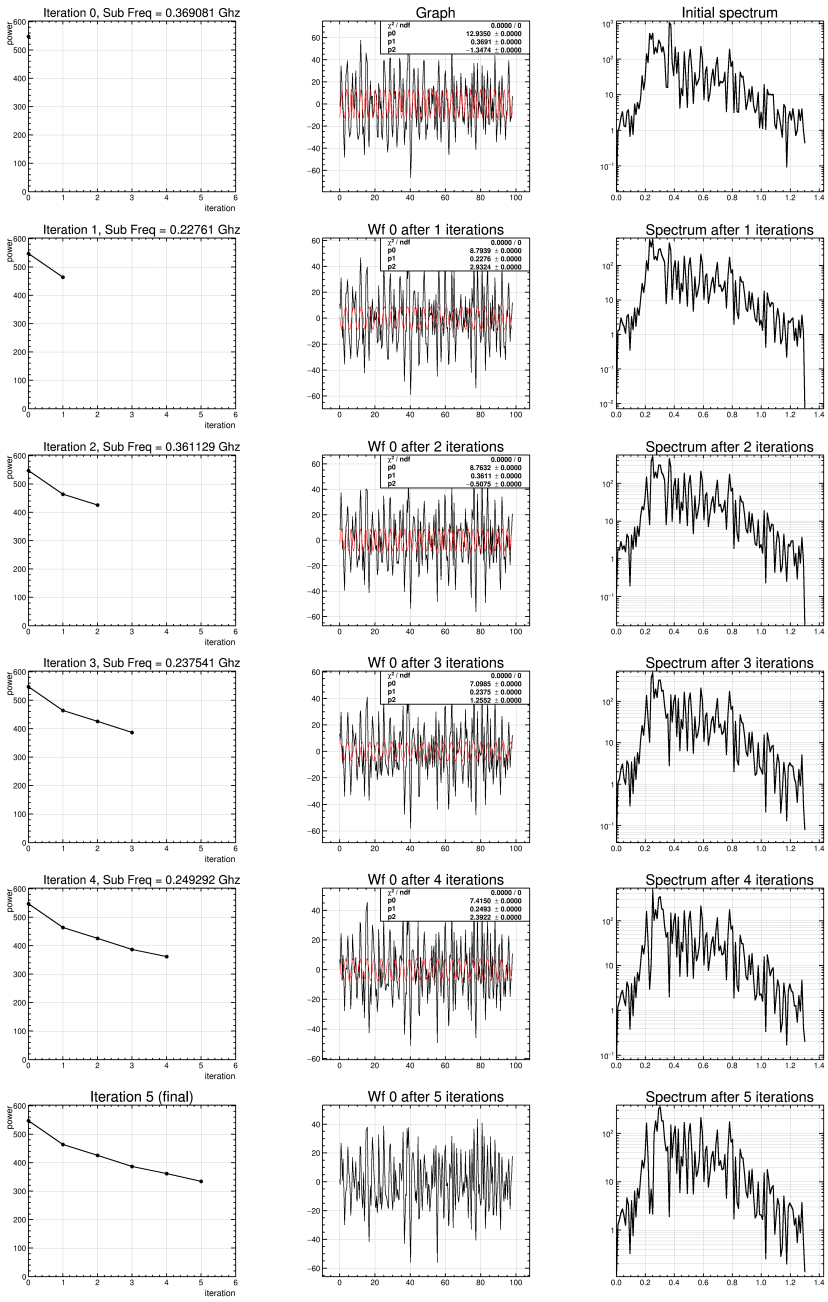

Figure 4.5: The effect that sine subtraction filtering has on a single waveform for event 56765803. The top-most set of plots is the initial waveform, and the bottom-most set is the final "filtered" waveform. (Left) The amount of power removed per iteration. (Middle) In black is the waveform used for the fit, and in red is the fit sine wave function. Also visible are the final fit parameters. From top to bottom they are the amplitude, frequency (in GHz), and phase. (Right) The spectral power of the corresponding waveform prior to the sine wave being subtracted.

4.5 Waveform Reconstruction

To extrapolate the measured signal on the ANITA instrument back to a physics signal or background noise event, we must combine the various waveforms captured around the instrument. By summing the waveforms with specifically determined time offsets, it is possible to determine the level of coherence of the electromagnetic environment. Once an arrival direction is determined it is possible to reconstruct the signal, averaging out the random thermal and electronic noise.

By correlating the waveforms captured between various antenna pairs it is possible to establish time delays in which multiple channels maximally constructively interfere. Known as interferometry, a radio astronomy technique, it makes use of the known physical separation vector between the antennas and the propagation delay for any incident plane wave.[54] By overlaying the correlation values for multiple baselines, it is possible to generate a map of likely pointing directions. ANITA has three vertical tiers of antennas, each separated from each other, and at least three co-pointing phi sectors that are expected to observe events from any angle. Using the combined baseline offsets from these antennas, it is possible to overlay $\frac{1}{2}N(N-1)$ single baseline interferometric maps to create a full pointing map, where N is the number of antennas used to form the map. By finding the peak of this map, a vector pointing away from the payload can be traced to the ice, where a map of the vertical height of the continent can be used to create a map of events on the ground.

4.5.1 Interpolation

As discussed earlier, the waveforms have a nominal sampling rate of 2.6GS/s, yielding a mean timing separation of 385ps. This presents an issue, namely that baseline timing delays between antenna pairs are not quantized in multiples of this value. In order to accurately overlay the recorded waveforms of different channels, upsampling the signal to a higher rate with a lower time separation between points is required. Additionally, the variation in time separation between samples caused by process parameter spread in the LABRADOR ASICs presents computational difficulties for Fourier analysis of signals.

There are several methods that can take unevenly sampled waveforms as an input, and output an evenly sampled approximation of the signal by extrapolating between points, a process called interpolation. An Akima spline iteratively fits a cubic polynomial to a range of sampled points, then uses that fit to determine the intermediate points [55]. Akima splines do a good job at determining an equivalent evenly sampled waveform, but introduces spectral distortion when used for increasing the sampling frequency. Another method of interpolation, Fourier interpolation, computes the complex Fourier transform of a waveform, appends zeros to the end until the desired frequency is reached, then calculates the inverse Fourier transform of that extended frequency domain signal. This method introduces no in-band spectral distortion, though for computational simplicity it must be performed on an evenly sampled waveform. A comparison of these two methods in interpolating

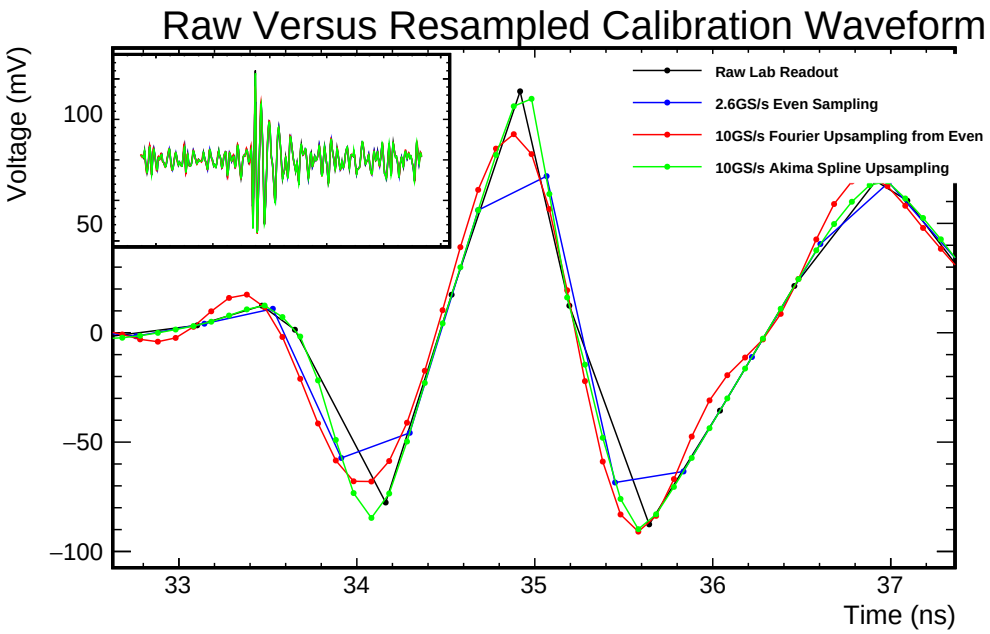

Figure 4.6: An example from channel 01BH of even resampling and up-sampling the original unevenly sampled LABRADOR digital readout. The peak of the impulse is shown zoomed, while a view of the entire waveform can be seen in the inset frame showing the macro-consistency of the methods. The method used in this analysis is FFT up-sampling. Despite undershooting some points, it minimally distorts the spectral power and phase response of the signal.

a raw, 2.6GS/s, unevenly sampled LAB output waveform to an evenly sampled 10GS/s waveform is shown in Figure 4.6.

The methodology used in this thesis for interpolation is to use an Akima spline to evenly sample the recorded waveforms, then using the Fourier domain upsampling technique to increase the sampling rate by a factor of three, to 7.8GS/s.

4.5.2 Radio Interferometric Mapping

Interferometric mapping is a method developed to compare measured time domain waveforms captured by antennas with a physical baseline offset with the goal of determining likely incoming plane wave direction.[54] This is accomplished by calculating the timing offsets expected from each direction of a pair of antennas in elevation and azimuth, computing the correlation function of the two waveforms, then building an interferogram by weighting each direction in relation to the correlation value at that specific timing offset. The number of unique permutations for N antennas choosing two at a time is given by the binomial coefficient, shown in Equation 4.1

$$\binom{N}{2} = \frac{N!}{2(N-2)!} = \frac{N(N-1)}{2} \tag{4.1}$$

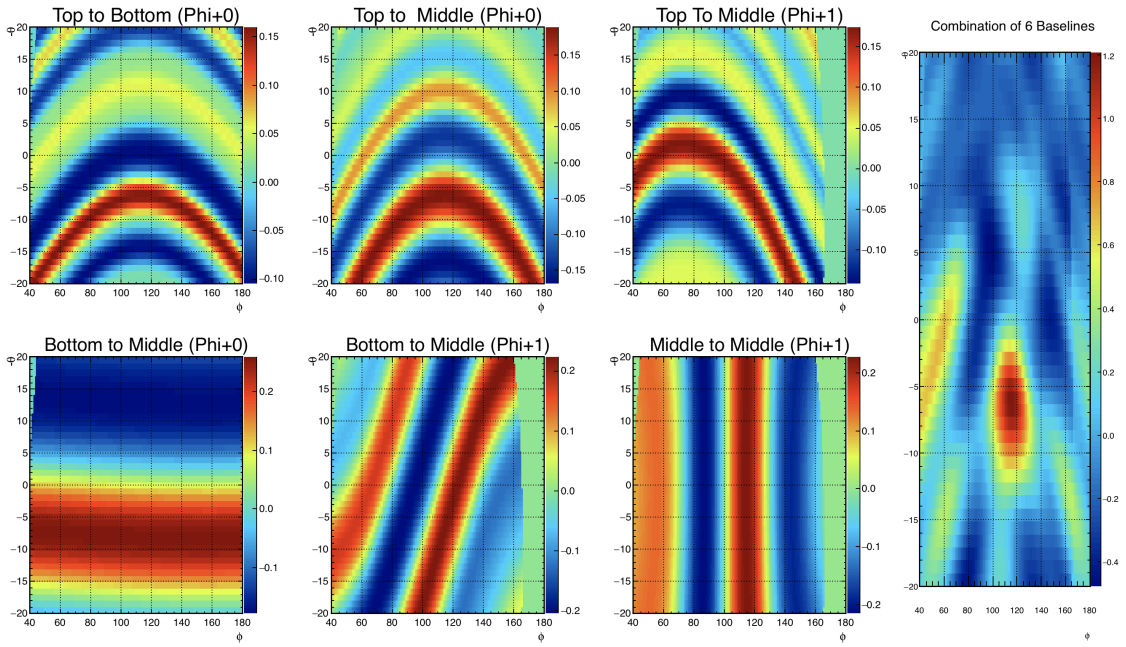

Figure 4.7: Six interferograms for single baseline correlation offsets between four antennas, and their combined full interferometric map. The signal being reconstructed for this image is a WAIS pulser event. The four antennas used are all three in the phi sector pointing nearest the incident pulse, as well as the middle ring antenna in the neighboring phi sector. Maps created in the full analysis combine interferograms from 15 antennas

For the analysis in this theis, $N = 15$ antennas were used, yielding 105 baseline pairs per map point. An example of four antennas yielding six single baseline pairs and the subsequent interferometric map can be seen in Figure 4.7.

4.5.3 Beamforming

By coherently summing multiple waveforms at precise offsets in a process beamforming, it is possible to improve the signal fidelity. The electronic and thermal noise present in each channel limits the sensitivity of the instrument to EAS signals with low power. Since the power within the radiation of an EAS shower varies linearly with the energy of the particle, low energy showers, which have a higher flux rate, will have signal to noise ratios approaching one. Additionally, polarization information and frequency content are effected by any noise present in the system. Averaging waveforms from multiple channels will reduce this incoherent background noise by a factor of $\sqrt{N}$, with N being the number of antennas used to form the new coherent waveform. After determining the peak interferometric pointing direction, it is possible to align the channels that received signal from an incident radiation field and average them together.

For this analysis 15 antennas, comprising 5 phi sectors, are used to construct the coherently

summed waveform. The 15 antennas are chosen based on the peak of the interferometric map, so that the closest 5 phi sectors are included.

The waveforms are then delayed based on the physical offsets between the phase centers of the antennas and a zero-point defined at the center of the payload, and added together with that timing offset preserved. This preserves the absolute timing information of the signal, since all coherently summed waveforms will have the same reference point.

4.5.4 Signal to Noise Ratio (SNR)

The Signal to Noise Ratio (SNR) is a measure of the ratio of powers between the coherent signal present in the waveform and the average noise power of that same waveform. This value is calculated by dividing the average of positive peak (V_+) and negative peak (V_-) voltages of the coherently summed waveform, and dividing it by an average of the root mean squared (RMS) of the past 600 waveforms for the channels that are used in the coherent sum. Equation 4.2 describes this calculation.

$$SNR = \frac{V_+ - V_-}{\sum_{i=0}^{60} V_{rms,i}} \tag{4.2}$$

In this thesis when I refer to SNR, I mean the value that is calculated from the voltage and timing calibration signal as stored to disk. This includes the dispersion of the signal chain and antennas, as well as the additional cascaded electronic noise introduced by the amplifiers. The SNR of the incident electric field will have a different value.

4.5.5 Ray Tracing to Continent

The location that an interferometric map peak points to on the ice can be determined by tracing a vector pointing in the peak direction back to the continent. To do this, it is important to have an accurate topographical map of both the Earth's surface, as well as the ice sheet that covers it. These maps are provided by the Radarsat Antarctic Mapping Project Digital Elevation Model Version 2(RAMPDEMv2) [56]. RAMPDEMv2 provides topography as referenced to the geoid model of the earth. Without a correct topological information, events would point to areas that are shadowed from view by mountainous terrain. An example of pointing to various topographical models, and description of the technique, are shown in Figure 4.8

4.5.6 Immediate Pointing Cuts

Real physics signals have a small range of elevation angles where their detection is most likely. The highest probability incident direction for an EAS is near the horizon, however noise events will point more evenly in all directions. Using this knowledge, we can cut out events that have very steep elevations. A distribution of event pointing direction can be seen in Figure 4.9.

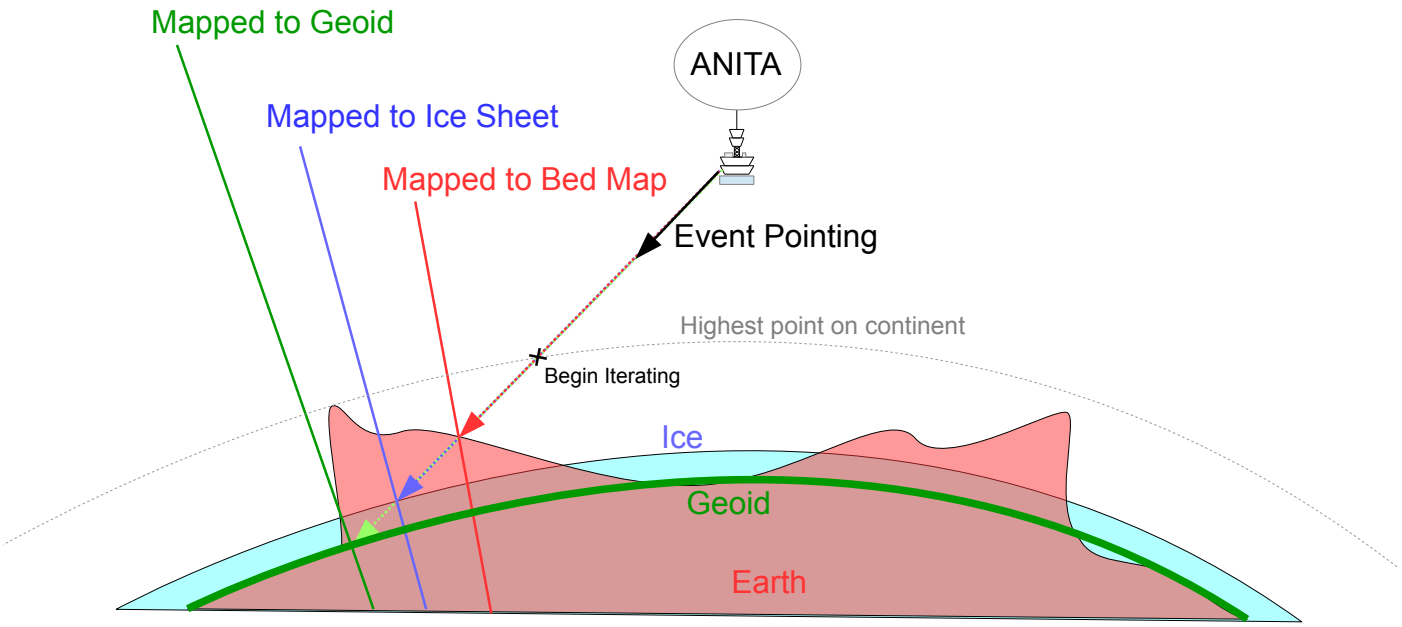

Figure 4.8: A diagram showing how events are traced back to the continent. First the latitude and longitude of an event at the highest altitude on the continent (Mt. Vincent at 5000m) is determined. Then the vector is iteratively lengthened until it intersects with a topographical dataset provided by RAMPDEMv2. This is done for both a map of ice field heigh, and a bedrock map. The latitude, longitude, and altitude trace back to the continent is then saved for the event for the top two peaks.

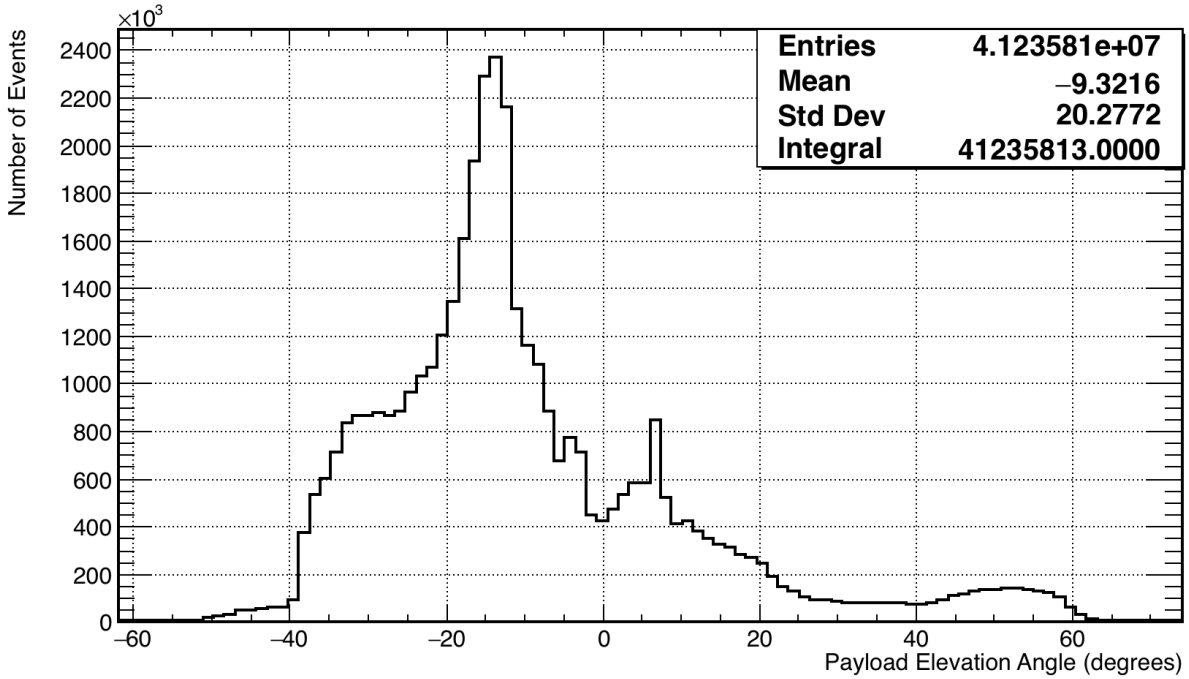

Figure 4.9: The distribution of elevation angles of the first interferometric peak for all waveforms. Zero points perpendicular to the vertical payload vector, and negative values denote events that point "downwards" from the payload. The horizon, which depends on surface height and payload altitude, has an average value of around -6 degrees. The most common source above the horizon is the sun.

4.6 De-dispersion

The antenna, filters, cables, and amplifiers are responsible for introducing a frequency dependent gain and group delay to the observed signals. These elements act to disperse the power in the signal across tens of nanoseconds from what was emitted from the relativistic shower at the critical angle as a highly impulsive, picosecond length essential delta-function like electromagnetic field transient. The impulse response, or complex phasor representation of this dispersive effect, was measured for the signal chain immediately preceding the flight for each of the 96 channels, and for all 48 antennas in a controlled manner in Palestine the summer beforehand. These two effects are then combined and used to determine a whole system impulse response that relates the measured ADC values at the LABRADOR digitizer to an electric field incident at the payload.

4.6.1 All-Pass Deconvolution

The numerical recipe used to deconvolve waveforms is named the all-pass deconvolution method. An ideal deconvolution, one that uses both phase and amplitude information to remove the instrument response, is poorly defined outside the instrument bandwidth. Since the instrument blocks any signal outside this frequency range, reversing the procedure introduces what is in essence a "divide by zero" error that grossly amplifies signal power outside the band. The simplest way to avoid this is to disgard the frequency dependent amplitude information, and only correct for the phase. While the ANITA amplifier chain does not have flat gain within the band, this is a good first approximation for signals with power occurring mainly below 600 MHz, such as UHECRs. The resulting waveform "stacks" signal power from all in-band frequencies on top of one another, increasing the observed peak to peak signal to noise ratio and offering a larger separation between signal and noise events during analysis cuts.

4.6.2 Generation of transfer function

The transfer function of the system was painstakingly developed and utilized a variety of both time domain and Fourier transformed frequency domain manipulations. These each introduce their own errors, as many of them require assumptions about the incoming signal. The full process for generating the transfer function is detailed in Appendix A.

4.6.3 Signal to noise ratio of impulse response

Any band limited deconvolution process requires a knowledge of the signal to noise ratio of the transfer function as a function of frequency ($SNR(f)$). Since the transfer function merely relays the amplitude and phase differences between an input and output signal (represented in either a complex phasor or a time domain waveform), one must also have an understanding of the total power contained within the input and output signals. If, for example, both the input and output

signal contained very little power out of a specified band pass, it could be wrongly assumed from a transfer function that the signal chain was able to pass frequencies that are out of the band pass.

4.7 Template Correlation

Though any physics experiment that aims to separate background events from signal events requires that a general understanding of characteristics inherent to either group, measurements and simulations of UHECR induced EAS signatures provide a waveform template that can be used to strongly cut on signal events. The simplest assumption, that an EAS will produce a broad spectrum, short duration, electromagnetic pulse, drives both the overall design of the payload, as well as several of the analysis cuts. Determining a normalized correlation between a template and the coherently summed waveform provides an additional constraint beyond simply requiring impulsivity, and also requires that the gain and phase characteristics of an incident wave are in agreement with an EAS signal.

4.7.1 ZHAires Shower Modeling

Simulating the radio emission of cosmic ray air showers is accomplished by using the ZHS extension to the AIRshower Extended Simulations (AIRES) package.[6] Concatenated, these two packages make up the ZHAires simulation framework. This framework, developed by Jaime Alvarez-Muniz, Washington Rodrigues de Carvalho Jr. and Enrique Zas, generates a time varient electric field for a configurable incident particle at some instrumentation point in the atmosphere. Recent work, with ANITA specifically in mind, has further extended the simulation to accurately depict reflections off the ice sheet.

Simulations provide an important method to determine the overall sensitivity of the ANITA instrument to air shower events as a function of incident angle, energy, and payload location. Due to the coherent beamed characteristic of air shower radio emmission, specifically a peak at the critical Cherenkov angle from the shower axis, ANITA is only sensitive to a small fraction of showers that occur within its field of view.

The spectral content of EAS radiation falls off sharply as a function of offset from the critical Cherenkov angle. Without being within a few degrees of the peak coherence angle, low frequency radiation, below the band of the ANITA instrument, dominates the spectrum. This was shown in Figure 2.18

4.7.2 Templates

It is possible to generate several different templates that accurately represent an observable EAS signal on the payload. Additionally, templates can be developed from averaging together calibration pulser or other anthropogenic signals of interest. For this analysis, the template used for a cut value

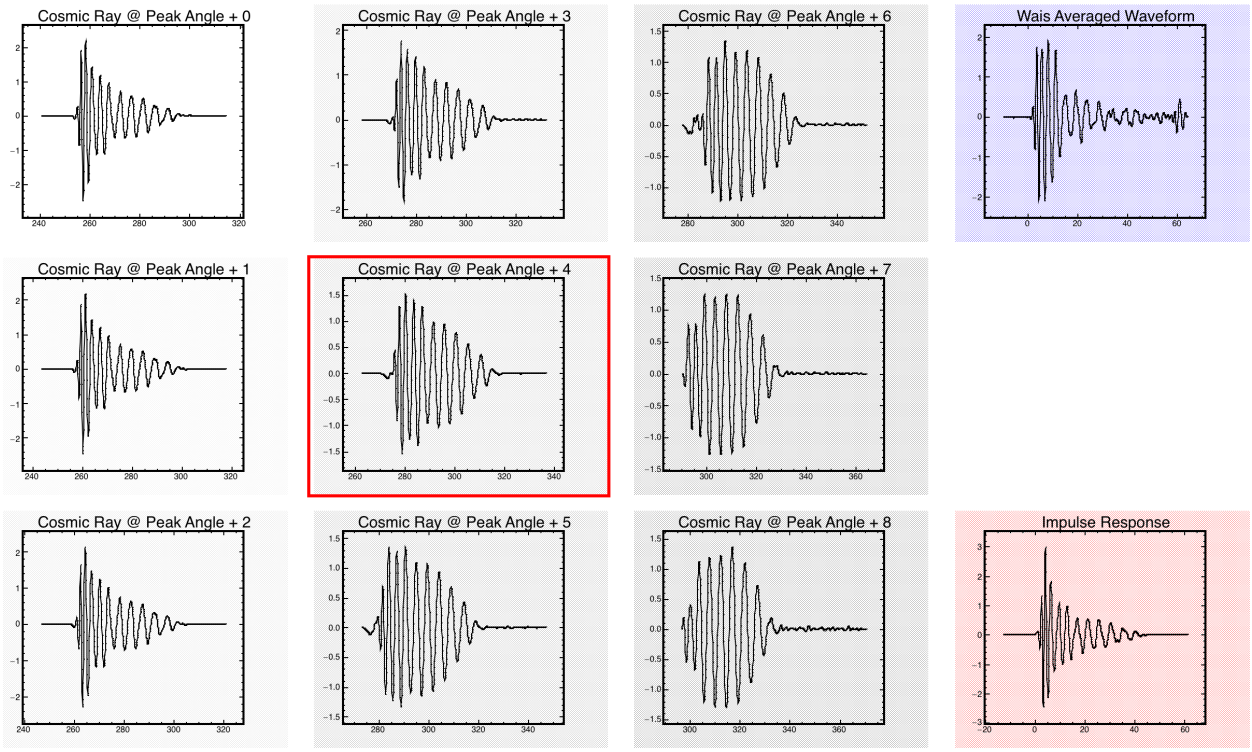

Figure 4.10: The 11 templates used in this analysis. The leftmost 9 templates were generated from ZHAires simulations convolved with the system impulse response. The center waveform, in the red box, was used for setting cuts for the analysis. The peak coherence angle for this simulated shower occurred at 0.675° from the shower axis, and each waveform was measured at steps of 0.04° from that peak. The top right is a template derived by correlating and averaging 10,000 WAIS pulser events, and is used for determining reasonable cut parameters based on calibration pulser efficiency. The bottom right is the impulse response of the system. All waveforms are normalized so that their autocorrelation is one.

was a simulated EAS electric field measured at 0.175 degrees outside the critical coherence angle, convolved with the instrument impulse response. Figure 4.10 shows a comparison between the 9 CR templates, as well as the impulse response template and averaged WAIS pulser template. The WAIS pulser template was used to establish cut efficiency on WAIS pulses, as they do not have the same phase structure as EAS signals.

4.7.3 Auto-correlation normalization

In order to make the template correlation value of different events directly relatable to one another, it is important to normalize both the template and the coherently summed waveform. As this measurement is most interesting in determining the amount that any given event "looks" like the template, the final value should be a fraction from 0 to 1, where 1 would be a waveform that is an exact copy of the template.

4.7.4 Issues with filtering

Ideally these templates would be compared against the coherently summed waveforms that have been filtered for CW interference on a channel by channel basis. However, the distortion of the phase structure induced by all of the filtering strategies, including sine subtraction, reduces the correlation to values below that of un-filtered waveforms. For this reason, only the template correlation value for unfiltered coherently summed waveforms is used.

4.8 Polarimetry

The expected signal from a cosmic ray air shower has a characteristic polarization that can be used as a further discriminator between random noise, which will have a random polarization. Specifically, the geomagnetically induced radiation component will be linearly polarized and orthogonal to both the shower axis and the magnetic field. This allows two cuts to be made, one of the linear polarization fraction and one of the linear polarization angle. Both these rely on the calculation of Stokes parameters made possible by the dual polarization antennas, which capture both these time domain electric fields simultaneously.

4.8.1 Stokes Parameters

The Stokes parameters are a mathematical method that describe the polarization of electromagnetic radiation with a four-vector of values. The four parameters in the Stokes vector are I, which describes the total intensity, V which describes the circular polarization, and U and Q, which describe the linear polarization in two rotated bases. Q corresponds to linear polarization in coordinate basis that is either parallel or perpendicular to the complex electric field measurements, and U is the basis rotated 45° from that. A visual representation of the Stokes parameters is shown in Figure 4.11. These values can be determined for each digitized event by comparing the temporally dynamic electromagnetic fields of the orthogonally positioned Hpol and Vpol channels from each antenna. The equations for calculating these values are shown in Equation 4.8.1.

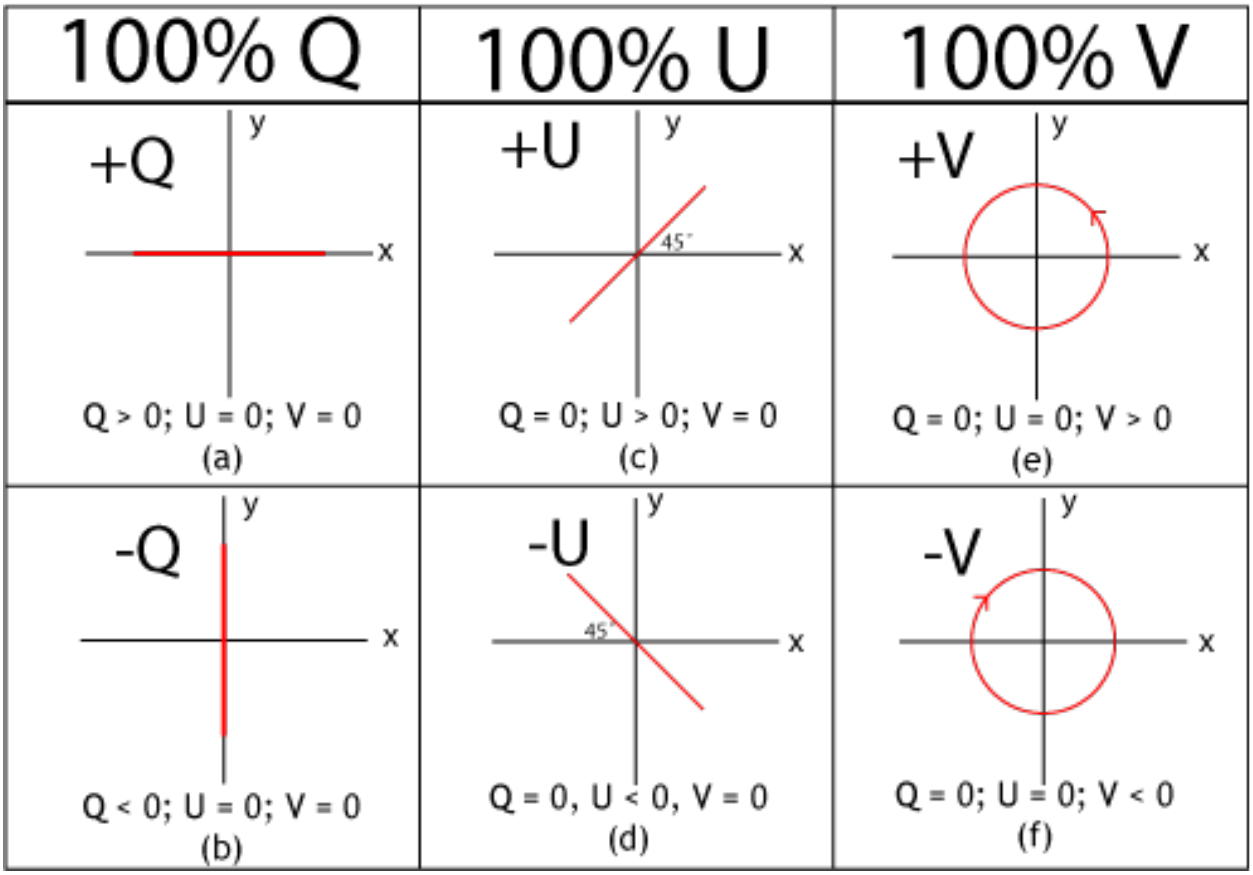

Figure 4.11: Visual representations of the polarizations as described by Stokes Parameters for waveforms with only a U component (a and b), a Q component, (c and d) or a V(component (e and f) [12]. Signals detected by the ANITA instrument have non-zero terms for all three.

$$\varepsilon_i = E_i + i\hat{E}_i$$

$$I = \frac{1}{n}\sum_{0}^{n-1}(|\varepsilon|^2_{i,H} + |\varepsilon|^2_{i,V})$$

$$Q = \frac{1}{n}\sum_{0}^{n-1}(|\varepsilon|^2_{i,H} - |\varepsilon|^2_{i,V})$$

$$U = Re(\frac{1}{n}\sum_{0}^{n-1}(\varepsilon_{i,H}\varepsilon_{i,V}*))$$

$$V = Im(\frac{1}{n}\sum_{0}^{n-1}(\varepsilon_{i,H}\varepsilon_{i,V}*))$$

$\hat{E}$ is the Hilbert transform of the electric field waveform, E_H is the signal from the horizontally oriented antenna channel, and E_V is the vertically polarized signal. The i is the sample number from the evenly spaced waveform. ε is then the complex voltage of the signal, where $i = \sqrt{-1}$. Re and Im denote taking the real or imaginary part of the value, and $*$ is the complex conjugate operator. This equation has axes referenced to the horizontal, perpendicular to the vector normal to the Earth. [57]

To reduce the effect that background noise has on the calculation of the Stokes vector, the waveform was truncated after 500 points ($n = 500$ in Equation 4.8.1), or 50ns after the beginning of the coherently summed waveform. This was chosen to capture the entirety of measured WAIS calibration pulses, but remove the trailing edge.

Determining the uncertainty on the calculation of these values was done by calculating the Stokes parameters for each waveform used in the coherent sum and determining the spread of the distribution of calculated values around the value calculated for the coherent sum.

4.8.2 Linear polarization fraction

The fraction of the signal that is linearly polarized can be calculated using Stokes Parameters. The Q and U components of the vector represent the linearly polarized portions of the signal, which can be normalized to the total coherent radiation I and added in quadrature to determine the fraction of power that is linearly polarized, p, as shown in Equation 4.3.

$$p = \sqrt{U^2 + Q^2} \tag{4.3}$$

It is expected that signal-like events will be mostly linearly polarized, while incoherent thermal sources will have low values of L. This can be used as a cut parameter for a cosmic ray search.

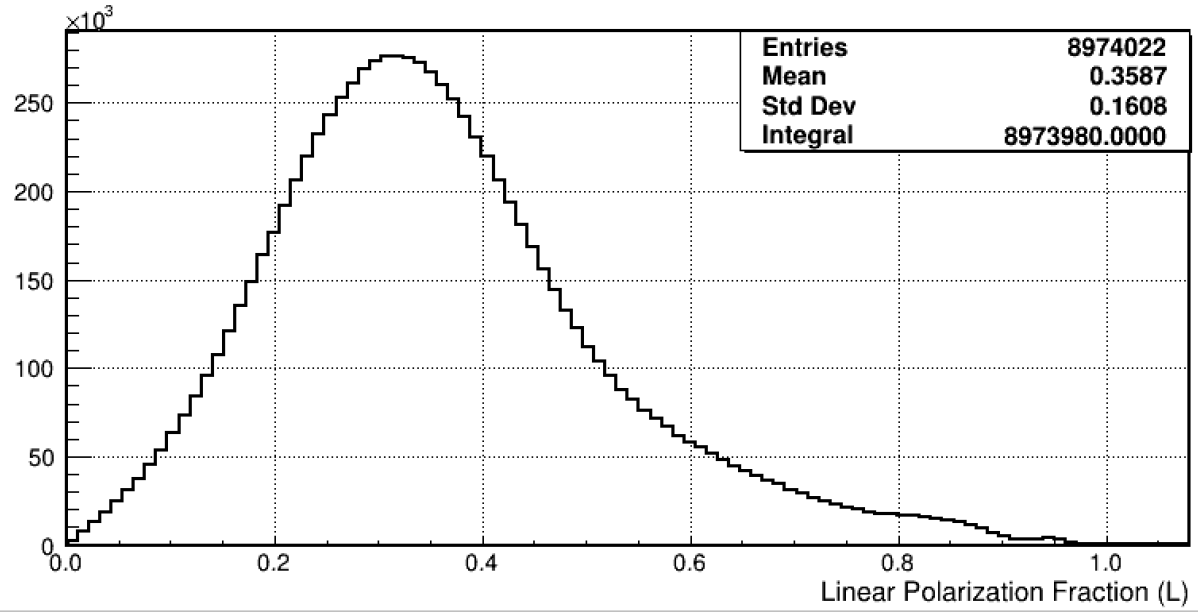

Figure 4.12: The distribution of linear polarization fraction for events measured during the ANITAIII flight that pass quality cuts. As expected, the majority of events are thermal in nature, and have low measured values of p.

The distribution of p for events that pass quality cuts can be seen in Figure 4.12.

4.8.3 Total polarization fraction

The two major contributors to CR induced EAS radiation, radially polarized Askaryan and linearly polarized Geomagnetic, can combine to create a signal with a circularly polarized element. It is therefore useful to calculate the total polarization fraction of the measured events. Since U, Q, and V all measure components of polarized coherent radiation, the quadrature sum of the total coherent radiation normalized parameters gives the total polarization fraction I_p, shown in Equation 4.4.

$$I_p = \frac{\sqrt{U^2 + Q^2 + V^2}}{I} \tag{4.4}$$

4.8.4 Plane of linear polarization

The predominant geometric plane of linearly polarized radiation is referred to the plane of polarization. The angle between the horizontally polarized antenna electric field axis and this plane is the linear polarization angle, and will be denoted by Ψ_L. This angle can be calculated by computing the angle between projections of the U and Q Stokes parameters onto their corresponding coordinate bases, shown in Equation 4.5.

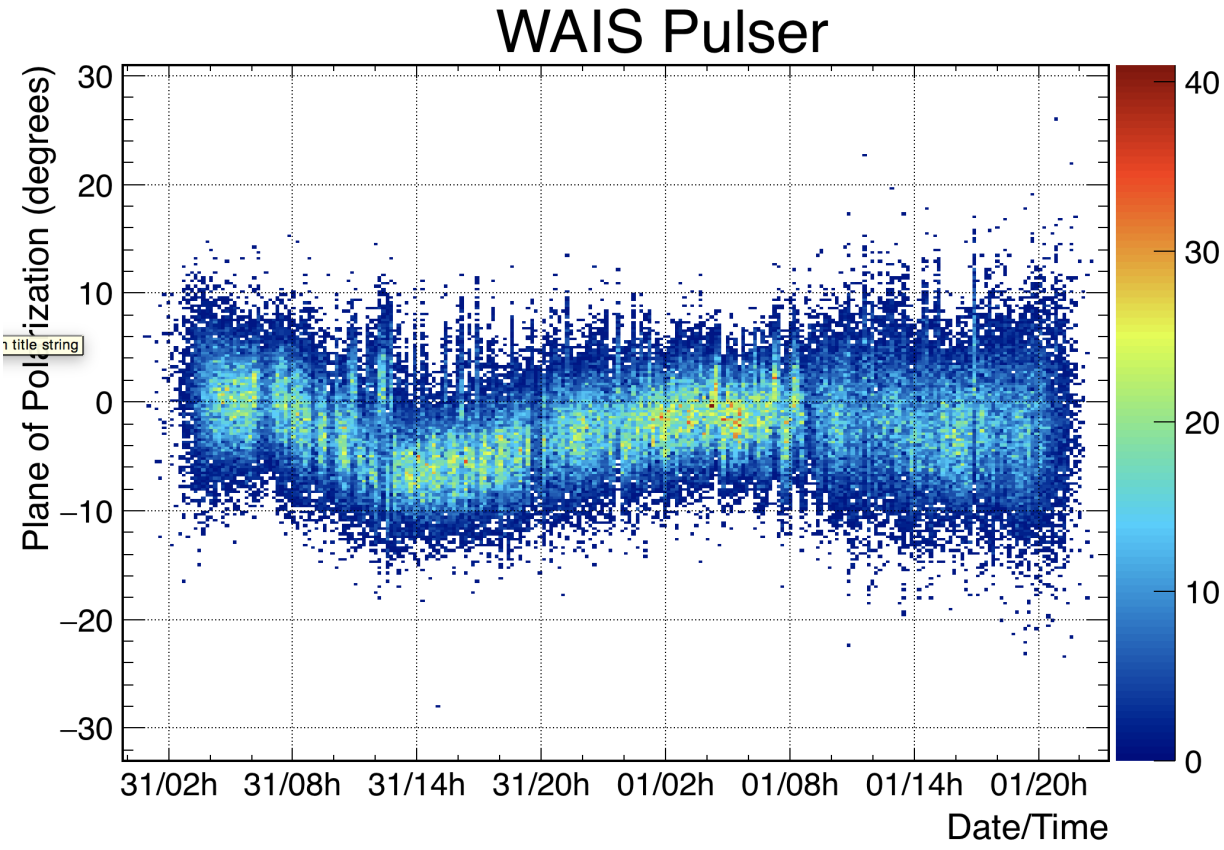

Figure 4.13: Two-dimensional histogram of calculated plane of polarization via Stokes parameters for all tagged WAIS pulser events. Long period fluctuations are likely caused by transmitting antenna and surrounding reflective and absorptive structures. As is seen in subsection 3.3.1, the payload passed by the WAIS encampment on a fairly straight heading, and time is directly correlated with transmit antenna azimuth direction.

$$\Psi_L = \frac{1}{2} tan^{-1}\left(\frac{U}{Q}\right) \tag{4.5}$$

Since the convention of the sign of the polarization of the signal is a matter of definition, for simplicity the values of Ψ_L are restricted to angles between 0° and 90°.

Calibration of this parameter was done through analysis of WAIS pulser events. The distribution of polarization angles can be seen in Figure 4.13.

4.9 Geomagnetically Induced Polarization Angle

The expected linear polarization angle of radiation emitted by an EAS will correspond primarily to the cross product between the shower axis and the geomagnetic field vector at the shower maximum. Comparing the expected and measured polarization angle for observed isolated impulsive events can

be used as a final discriminator on the likelihood that they are from an astrophysical interaction.

Calculating this expected angle requires a knowledge of the pointed event location on the continent of Antarctica, its reconstructed incidence elevation angle, and the geomagnetic field at the location in the atmosphere where shower max occurred. A map of the geomagnetic field over Antarctica Additionally, a majority of detectable EAS events will have reflected off the ice sheet of the continent, requiring the the calculation to include the Fresnel reflection coefficients for components of the incident radiation Poynting vector perpendicular and parallel to the ice-air boundary. This reflection also inverts the polarity of the parallel components of the electromagnetic field.

Up-going tau neutrino candidates are also expected to be dominated by linearly polarized geomagnetic radiation, however the location of the shower maximum will be above the ice sheet between the pointed source location and the payload. These events will then have different expected geomagnetic polarization angles than if the event was a reflected down-going cosmic ray induced extensive air shower. In order to preserve both astrophysical signals in the final candidate list, the expected geomagnetic angle is calculated for both cases, and the measurement falling close to the expectation of either case is treated as sufficient to pass cuts.

4.10 Polarity Estimator

In order to identify up-going tau neutrino candidates within the detected cosmic ray candidate set, it is important to have a method for determining the sign of the polarity for any particular event. Since a reflection off the ice will invert the polarity of the electromagnetic field impulse, it is clear that an event with a steep elevation angle but polarization that is inverted from the bulk of the CR candidates will be from an astrophysical particle exiting the continent and creating a shower within the atmosphere.

The polarity estimation for this thesis was accomplished by determining whether the absolute value of the peak of the correlation is the maximum or minimum, i.e. if the sign of the peak is negative or positive. Since detection of a tau neutrino would be both exciting and controversial, the mis-identification of this method is desired to be as small as possible. The sign of the peak of the correlation is recorded for both the coherently summed and the de-dispersed waveforms. When this method is applied to the 118,268 tagged WAIS pulser events, which all have the same polarity, only one was identified as the incorrect polarity by using the coherent sum, a misidentification rate of 8.46×10^{-6}. However, using the de-dispersed waveform, 200 events (including the single mis-identified event from the coherent sum) were found, a misidentification rate of 1.69×10^{-3}. This is shown graphically in Figure 4.14. To be considered properly identified, the estimates for the dispersed and coherent waveforms are required to be identical. This yields 1 (8.46×10^{-6}) WAIS event that is mis-identified, and 199 (1.68×10^{-3}) events of indeterminate polarity.

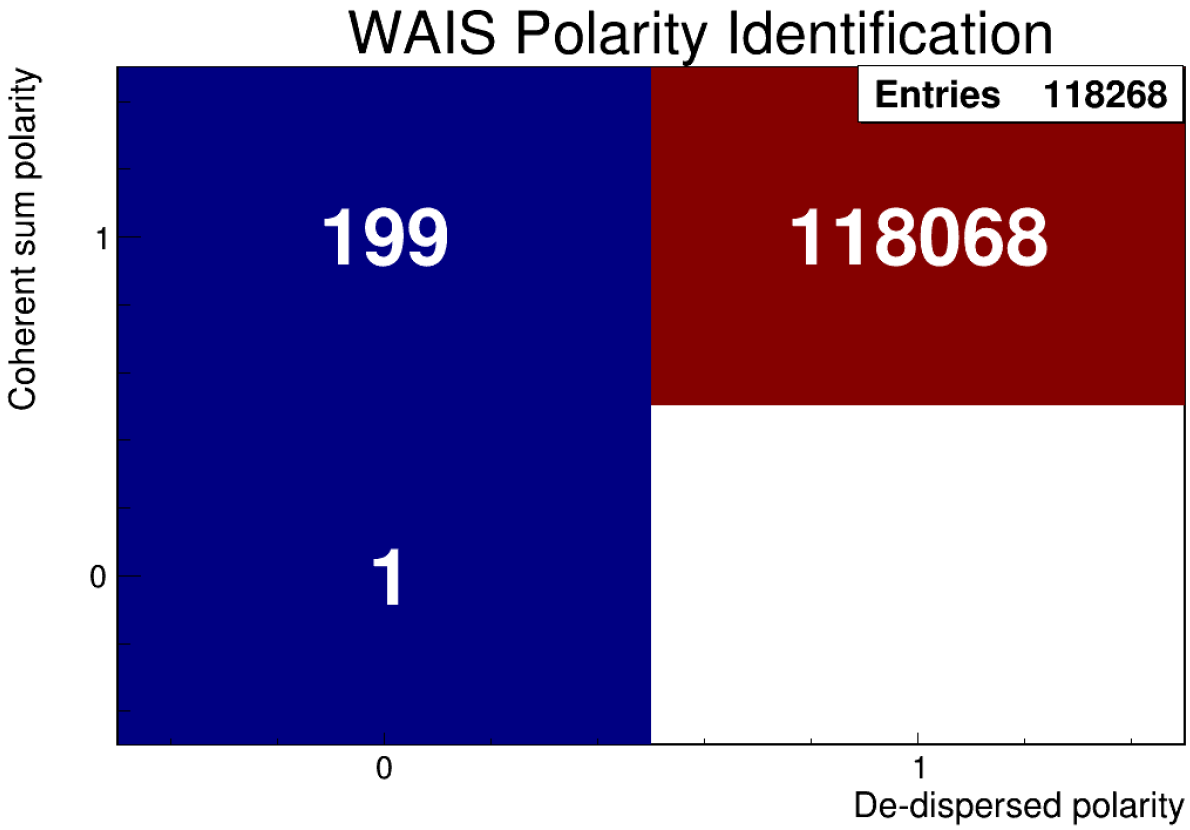

Figure 4.14: The results of the polarization identification on the 118,268 WAIS pulses. The algorithm has a 8.46×10^{-6} misidentification rate, and a 1.68×10^{-3} probability of returning an indeterminate polarity.

4.11 Known and Measured Source Identification

Many bases on the Antarctic continent are already cataloged by the various national programs that operate them. Using these known bases we can eliminate events that interferometrically point to objects of expected anthropogenic noise. There also likely exist bases that are, for whatever reason, not included in our catalog. We also use a list of "pseudo-bases" generated from clustered event lists from previous ANITA flights to eliminate possible anthropogenic interference. The resulting effect these excluded regions have on the flight can be determined by its flight path and a log-likelihood method for generating pointing error elipses around the bases.

4.11.1 Sun and Its Reflection, Thermal Noise Effect

The Sun does pass below the horizon during the austral summer, and it stays within the field of view of the ANITAIII instrument throughout the flight. While the thermal noise from the sun is random and unpolarized in nature, it presents as a correlated noise source to the antennas where it is within their field of view. For many events measured in the flight, the Sun presents itself as the highest interferometric map peak. However, since the sun is above an elevation angle of zero these events will not pass quality cuts. However, when the sun is at high elevation angles far from the antenna beam patter, the reflection of the sun off the ice, which also produces correlated thermal noise, can dominate the interferometric image.

4.11.2 Satellites, bases, and other "known" anthropogenic sources

The continent of Antarctica is under constant study by many scientific groups spread across the ice. Additionally, humanity has launched orbiting artificial satellites that are equipped with radio frequency communications. Both these sources are known to exist, and can be compiled into a list and used to mask portions of the field of view.

4.11.3 Pseudo-bases: unknown anthropogenic sources

Despite the best efforts of the collaboration, it is not possible to compile a complete list of Antarctic camps during the flight to a high degree of certainty. Because of this, it becomes necessary to use measurements from the payload itself to discern locations of human activity.

4.12 Event Quality Cuts

All events taken with the ANITA instrument must first pass basic checks on their ability to be considered in a sample of astrophysical impulses. ANITAIII does not seem to suffer from significant issues related to the digitization or storage of waveforms, however there are several observed event types that can be identified and removed prior to analysis.

4.12.1 RF Triggered

The first requirement for an impulse to be considered as a CR candidate is that the RF trigger must have initiated the readout. The 2Hz of minimum bias triggers are therefore excluded from the analysis.

4.12.2 Calibration Pulsers

Accurate tagging of calibration pulsers allows us to remove them from the sample considered in the physics analysis. This is done using the recorded nanosecond precision trigger time, a precise knowledge of the time in each second where each pulser fired, and the time of flight between the payload at the event record time and the calibration pusler. These events are known, do not contribute to either the background or can be considered as signal, and are thus not included in the final analysis dataset.

4.12.3 Payload Blasts

The ANITAIII flight was plagued by high power signals that seem to emanate from the payload. These events, called payload blasts, have a large amplitude signal in several neighboring phi sectors and tend to happen in rapid succession. A sample payload blast event can be seen in Figure 4.15. The blasts characteristically have high total RF power in the bottom two rings of antennas, but very little signal power in the top ring, suggesting that these events have a source on-board the payload. Blasts also do not have high interferometric peaks, adding further evidence that they are spherically propagating electromagnetic waves with a source on the payload, and suggesting that they will likely be removed through signal cuts later in the analysis process.

The primary metric used to select and remove these events is this ratio of total waveform power between top and bottom rings, R_{T2B} (Equation 4.6). This distribution can be seen in Figure 4.16. The cut for these events was placed at a value of 3, which eliminates approximately 731k events (1.77%) of events from the Hpol data set, while effecting a trivial number of non-RF triggered thermal events. The source of these blasts is unknown, though there is much speculation to their cause, and were present in both the ANITAI and ANITAII flights. Removing these events from the dataset is done prior to all other analysis steps.

$$R_{T2B} = \frac{\sum_{n=0}^{N} V_{bottom}^2}{\sum_{n=0}^{N} V_{top}^2} \tag{4.6}$$

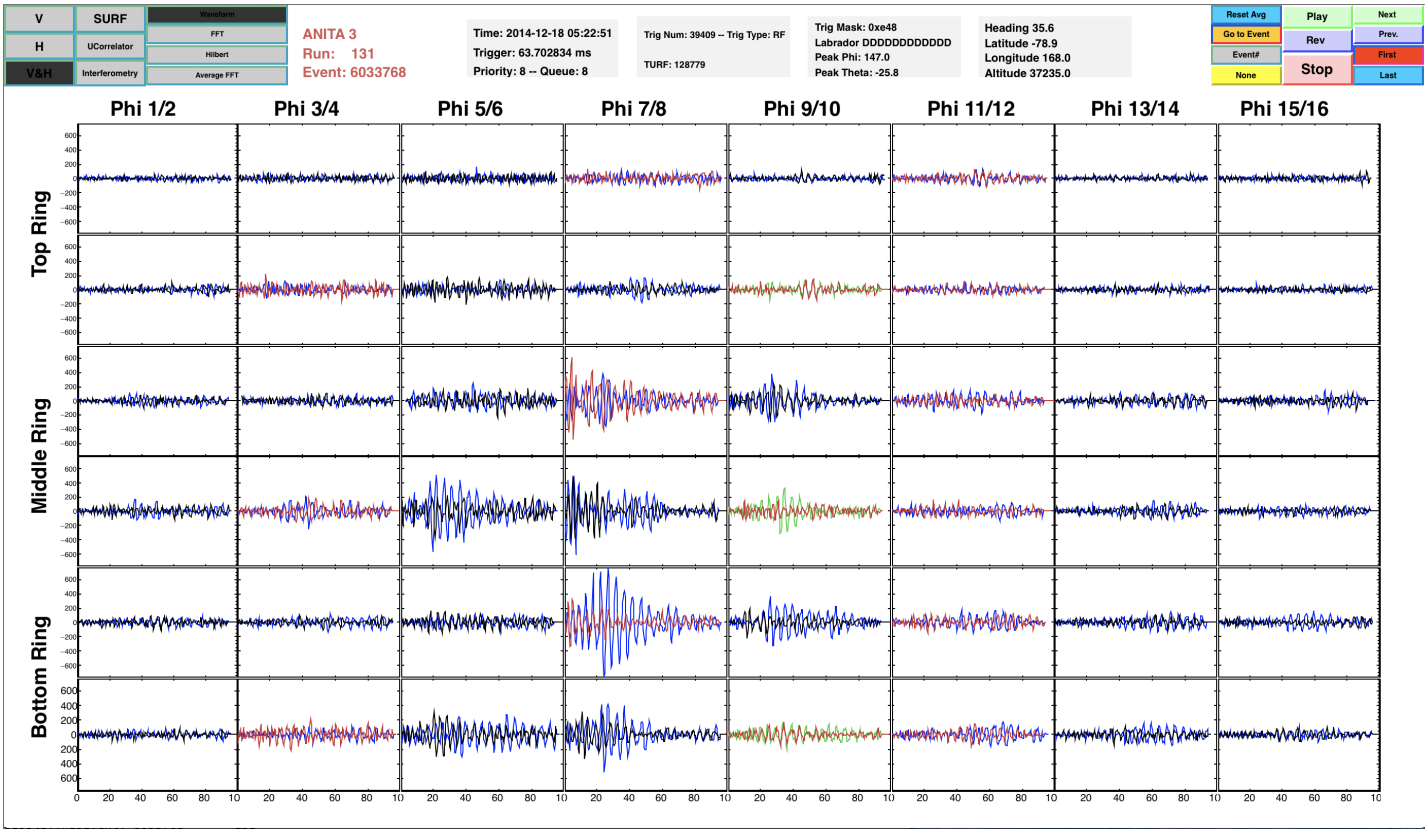

Figure 4.15: A magicDisplay screen capture of a suspected payload blast event. This event was selected due to the high ratio of power between the top and bottom antenna, 6.15 for this particular event. This pulse is additionally part of a 14 event train of pulses

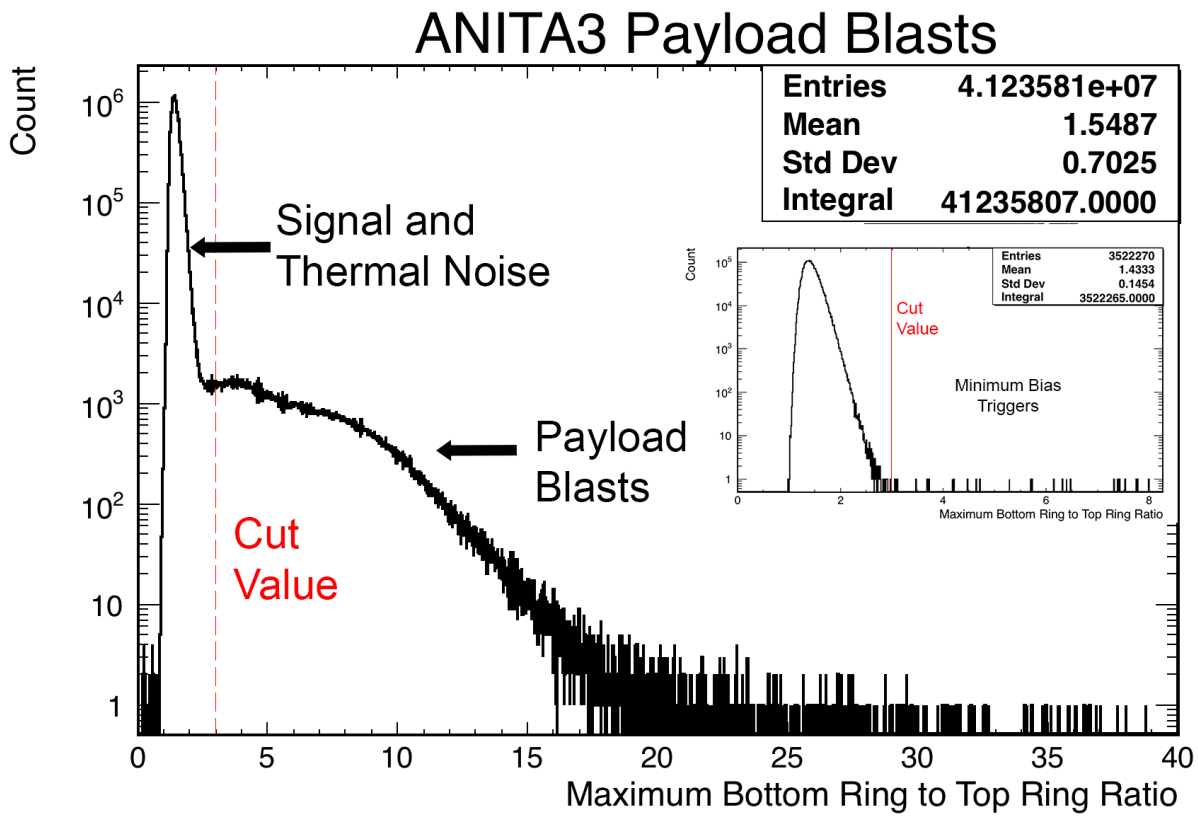

Figure 4.16: The event distribution of the ratio between the power measured in top and bottom rings of antenna for all horizontal polarization triggered phi sectors. The distribution for non-RF triggered events is displayed as well, within the full distribution plot. Red dashed lines denote the location of the quality cut.

4.12.4 Events above horizontal

Events that have an interferometric map peak that points upwards cannot be created by any of the astrophysical event types under consideration for this analysis. ANITAIII floats at the top of the stratosphere, above which there is sparingly small amounts of atmosphere with which a particle can interact. Even direct CRs will have impulsive signals that emanate from within the atmosphere. Because of this, all events that have peak elevation angles above zero can be safely discarded.

4.12.5 Delayed waveforms: reconstructed angle to hardware trigger angle cut

The dynamic phi sector masking introduces a class of events in which the waveform "falls off" the end of the digitization record, subsequently losing information about the earliest structure in the waveform. This is caused when impulsive events hit the payload in a masked phi sector, but have a sufficient amount of power to trigger channels with antennas in un-masked phi sectors, despite the large angles from their peak gain boresight. The majority of the signal power will still be observed in channels pointing towards the impulsive source, however the free space propagation time between the masked and unmasked phi sectors delay the trigger formation enough so that the waveform is truncated and information is lost. An example of this non-ideality can be seen in Figure 4.17.

Since these waveforms are missing important information, and the incident electric field cannot be reliably reconstructed, they can not be included in a final signal sample. To ensure that the at least one antenna involved with the coherent sum is included in the reported hardware trigger, this cut is set at 45°.

4.12.6 Major Bases

WAIS Divide and McMurdo Station are, for this analysis, defined as major bases which are removed prior to analysis. These bases have some of the highest activity on the continent, and the sheer number of events that come from these two bases has the effect of biasing the remainder of the analysis. Any event that reconstructs near either base is excluded. Nearness is conservatively defined as the map peak being within a geometric angular sum of 6° of the geographic coordinates of either base. This check is only done if the payload is within 700km of either base. Improperly time-tagged calibration pulser events would be excluded from the analysis data set via this cut.

4.12.7 Effect of quality cuts on number of events

The quality cuts remove a large fraction of the events recorded by the payload. The number removed by each cut can be found in Table 4.1. The resulting quality cut event list contains 2,670,191 events.

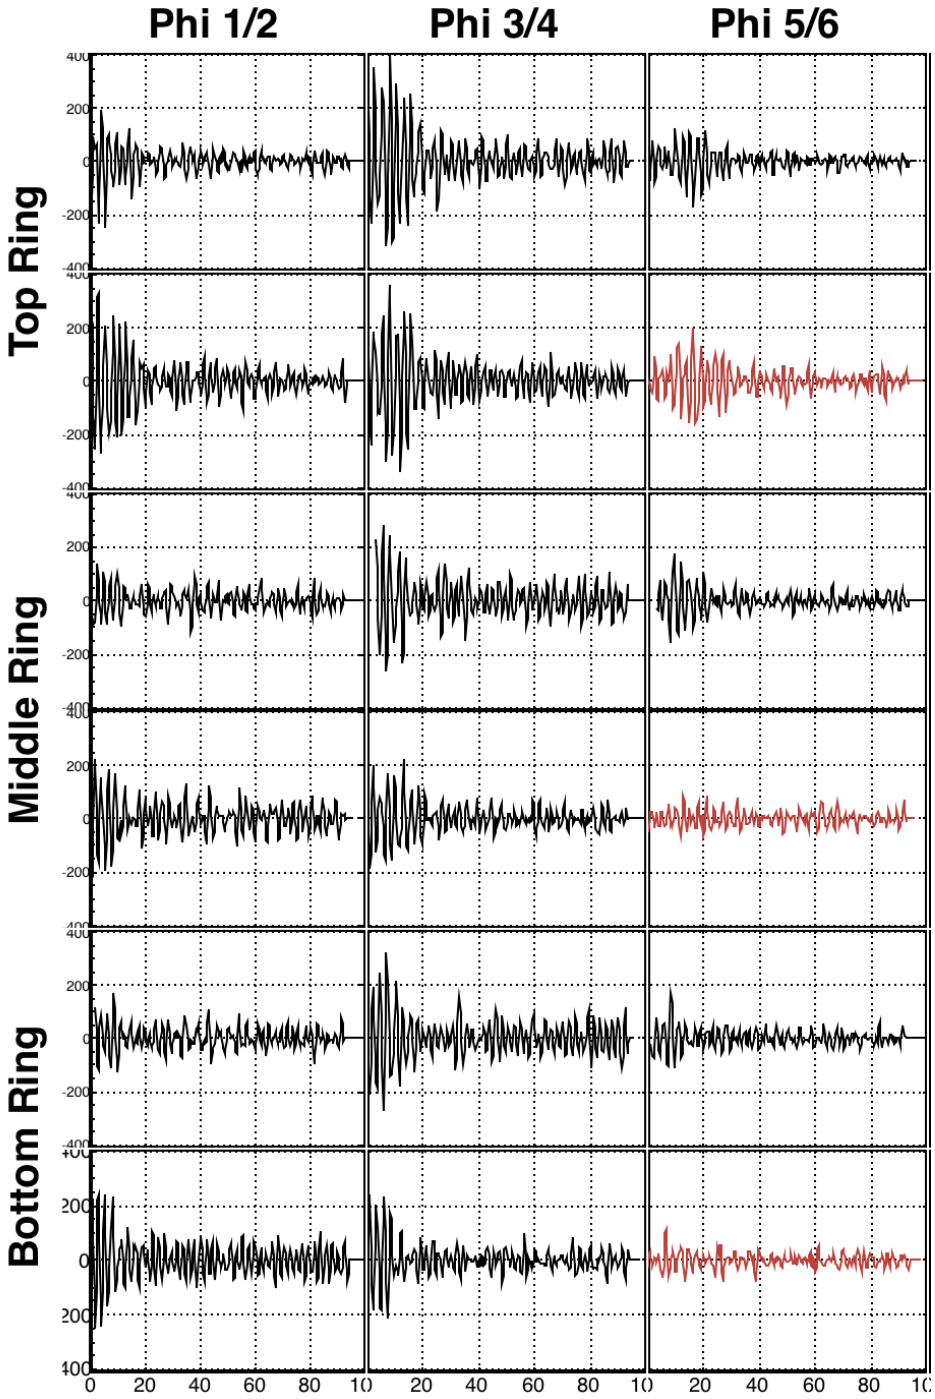

Figure 4.17: An example of a waveform where the promptest radiative component has fallen off the record and is not recorded. The red waveforms note channels which were reported to be involved in the trigger by the hardware. This event, 5877965, has a difference between the peak map angle and the hardware trigger angle of 77.5°.

Quality Cut Parameter	Events Failed (Single)	Events Failed (Cumulative)
LDB Pulser	342,108	342,108
WAIS Pulser	118,458	118,458
Points to LDB	740,683	314,438
Points to WAIS	271,846	17,188
Above Horizontal	10,106,687	
Hardware Trigger Angle	28,022,208	
Payload Blast	733,458	
Not RF Triggered	3,571,726	

Table 4.1: The number of events removed by each quality cut from the set of events with H-pol triggers. The cuts remove any event with a single value below those in Table 4.2. Events that are removed by several cuts are double counted, unless denoted by a "*", when tagged calibration pulser events are removed prior to determining removed event quantities.

4.13 Analysis Cut Methodology

Cut values that separate astrophysical sources from thermal and non-astrophysical anthropogenic background are determined by examining the distributions of various measured cut parameters for triggered events. The most desirable cut value will be one that maximizes the quantity CR-like while minimizing the number of thermal or anthropogenically triggered events. Two cuts are made to aid in identification, and removal through clustering, of strong anthropogenic sources: a signal cut that will define the final candidate events, and a lower impulsivity cut which will include many anthropogenic signals.

4.13.1 Impulsivity Cut

First, distributions of values for events that point at known major Antarctic bases are examined to determine a level which would capture weaker bases to high likelihood in order to identify and exclude them. The number of events that survive this cut will be inversely related to the final analysis exposure. This first more general "impulsivity" cut will capture information about locations on the continent that produce impulsive signal distributions that may leak into the signal. By searching for events that are geographically clustered, it is possible to remove generators of impulsive signals. This method is described in more detail in section 4.14.

4.13.2 Signal Cut

Secondly, a signal cut designed for maximum detection efficiency is used to classify candidate events. Events that meet this signal cut and are isolated on the continent will be labeled as CR candidates. WAIS calibration pulses were examined as a representative signal source, however with their specific expected waveform template in place of the CR template. In order to extract the highest number

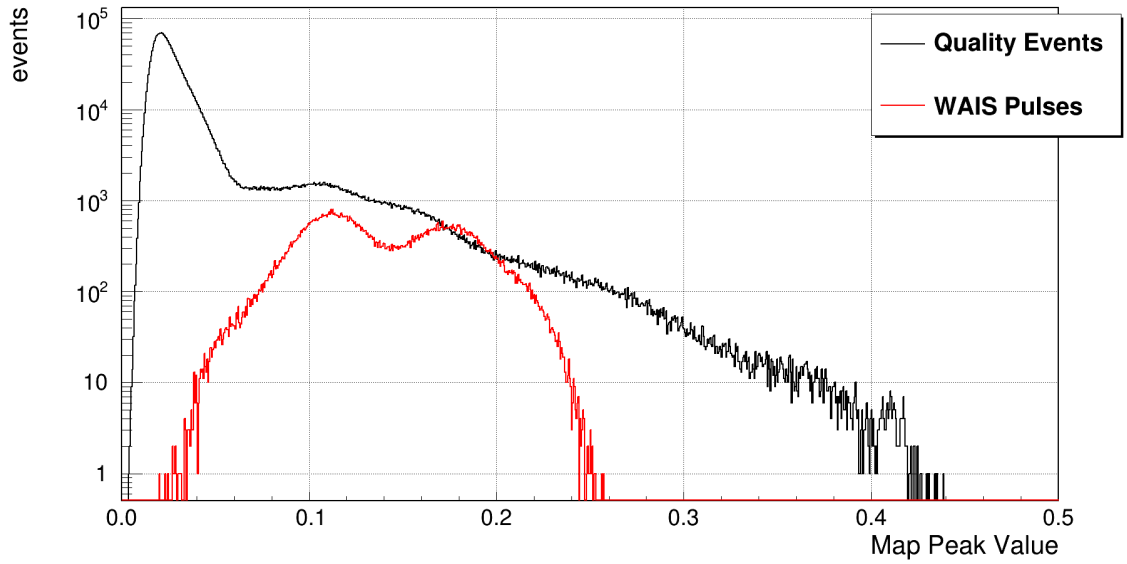

Figure 4.18: Distribution of interferometric map peak values for events that pass quality cuts and WAIS calibration pulser events.

of events with measured quantities most similar to an expected CR EAS radiation pattern, cuts were selected primarily to maximize efficiency on signal, and allow for a maximum exposure for the analysis. An a posteriori background estimate (discussed later in section 5.4) that examines the underlying localized distributions of events will be performed on all CR candidates in order to further constrain and quantify the background that would have leaked into the signal region.

Six calculated values were considered for the cuts. These, and the names they will be referred to, are: the peak value of the interferometric map (map peak), the absolute peak of the Hilbert envelope of the coherently summed waveform (hilbert peak), the linear polarization fraction averaged over 50ns of the coherently summed waveform, and the template correlation peak value for both the coherently summed (template) and for the de-dispersed waveforms (de-dispersed template). These values, whose calculations are described above, require that the measured signals both are impulsive and have a shape comparable to a CR, and singularly remove the same fraction of WAIS events. It is additionally desired that the final quantity of non-RF triggered thermal events passing these cuts is zero, with an expected rate of less than one tenth for the flight. Plots of these values for the non-RF triggered events, for events that point at major bases, and for the WAIS pulses are shown in Figures 4.18, 4.19, 4.20, 4.21, and 4.22. As a note, the template used for the analysis of WAIS pulses was the coherently summed and averaged WAIS pulse, derived from measurements.

The values for the cuts are shown in Table 4.2. The final signal cut values were set so that 99.9% of WAIS pulser events pass each cut individually.

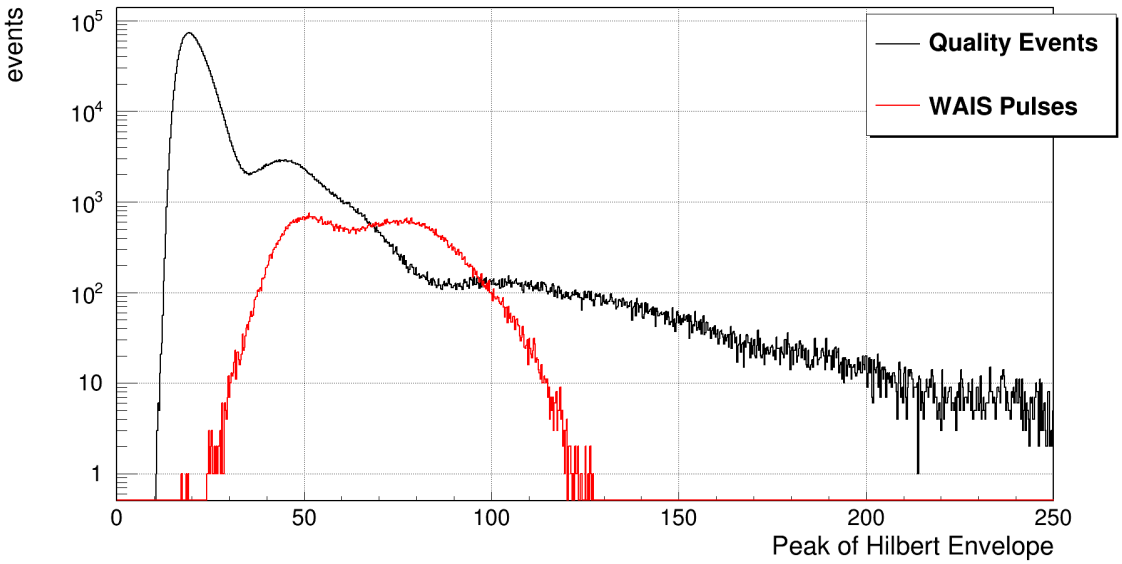

Figure 4.19: Distribution of Hilbert envelope peak values for events that pass quality cuts and WAIS calibration pulser events.

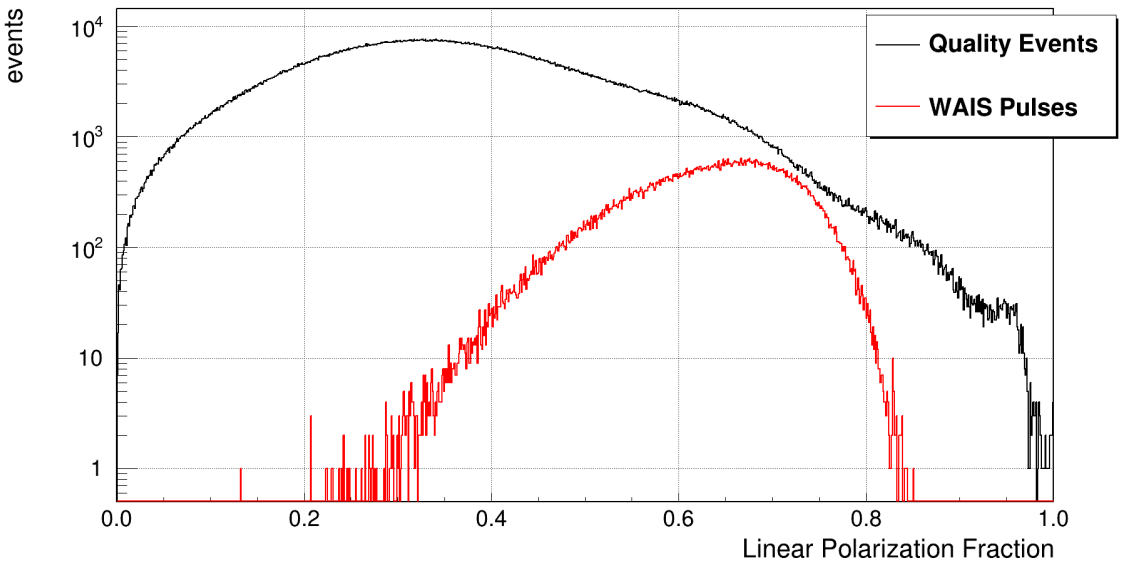

Figure 4.20: Distribution of linear polarization fractions for events that pass quality cuts and WAIS calibration pulser events.

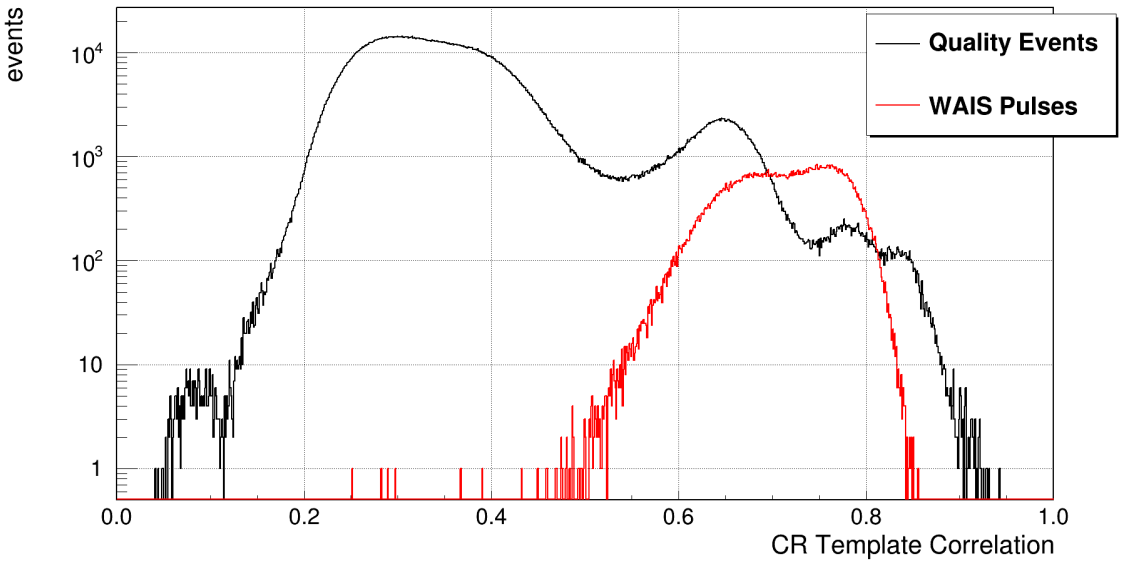

Figure 4.21: Distribution of peak cosmic ray template correlation values for events that pass quality cuts and WAIS calibration pulser events.

Figure 4.22: Distribution of peak cosmic ray template correlation values for events that pass quality cuts and WAIS calibration pulser events.

Quantity Name	Impulsivity Cut Value	Signal Cut Value
Map Peak	0.0435	0.0435
Map SNR	9.05	9.05
Hilbert Peak	25.0	31.1
Linear Polarization Fraction	0.50	0.60
Coherent Template Correlation	0.50	0.666
Deconvolved Template Correlation	0.50	0.666

Table 4.2: Values used for impulsivity and signal cuts. Impulsivity cuts are derived from measurements of known major bases, which produce a wide range of anthropogenic signals. Signal cuts were determined from WAIS pulser measurements. Determinations of the cuts is detailed in the text.

4.14 Geographic Event Clustering

After enriching the event sample by removing events with a high probability of being thermal noise or non-impulsive anthropogenic background, geographical clustering of the events can be done to further discriminate against backgrounds. Astrophysical signal events are expected to be isotropically distributed evenly across the continent, as there is no preferred cardinal direction for UHECR or UHEν_τ detections. Human activity, on the other hand, tends to cluster around regions of logistical or scientific interest. Anthropogenic background sources are thus expected to point at other captured events, which can be used to discriminate against them and further refine the data sample.

4.14.1 Base Clustering

Through communication with various international Antarctic institutions, a list of active camps and other human activity on the continent was compiled for the time frame of the ANITAIII flight. An image of these bases mapped to the continent can be seen in Figure 4.23. However these bases are not used to exclude any signal events. Determining the validity of this base map is plagued with issues, as they are run an managed by a large collection of different nations and groups, many of whom do not wish or care to make their activities publicly available. Additionally, camps that are reported by groups may be inactive, or have moved from when they were last reported. Due to this non-statistical and very difficult to model uncertainty, the relationship between any impulsive event and a camp present on the list does not significantly enhance our understanding of the source of such event. We must therefore employ alternate means to estimate the likelihood that a specific event is anthropogenic or astrophysical in nature.

4.14.2 Event Clustering

Impulsive events that point back to a geographically isolated source location have the highest likelihood to be astrophysical in nature. The metric that has been used in past ANITA experiments to determine the "closeness" of one event to another has been called the Log Likelihood, and is represented by the variable L. This measurement must take into account both the pointing uncertainty of the instrument, as well as the constantly changing payload location. Since each event can be seen from a different angle, the angular separation between the projected location of one event on the ice from the perspective of another event may not be the same as the inverse. This is shown graphically in Figure 4.24.

Lets assume that we have two events that are projected nearby onto the continent, A and B. From the payload location where event A is captured, you can measure the angular separation between the two projected positions in elevation(θ) and azimuth (ϕ). These angular separations will be denoted as $\Delta\theta_{AB}$ and $\Delta\phi_{AB}$ Equivalently, these separations can be measured from the payload

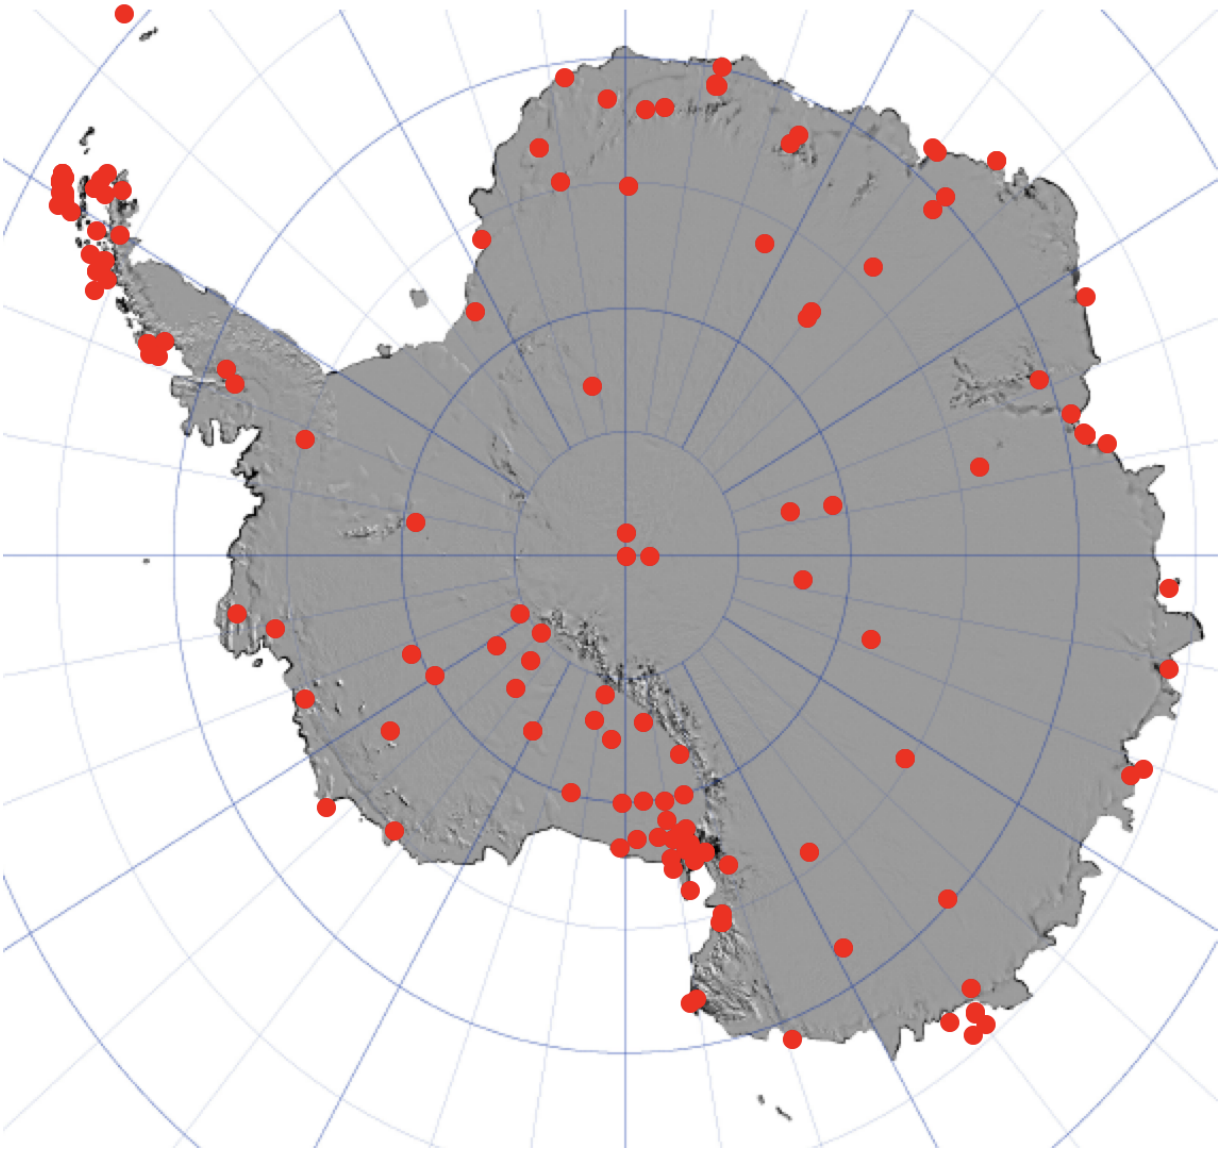

Figure 4.23: A map of the recorded human activity on the continent of Antarctica during the ANITAIII flight.

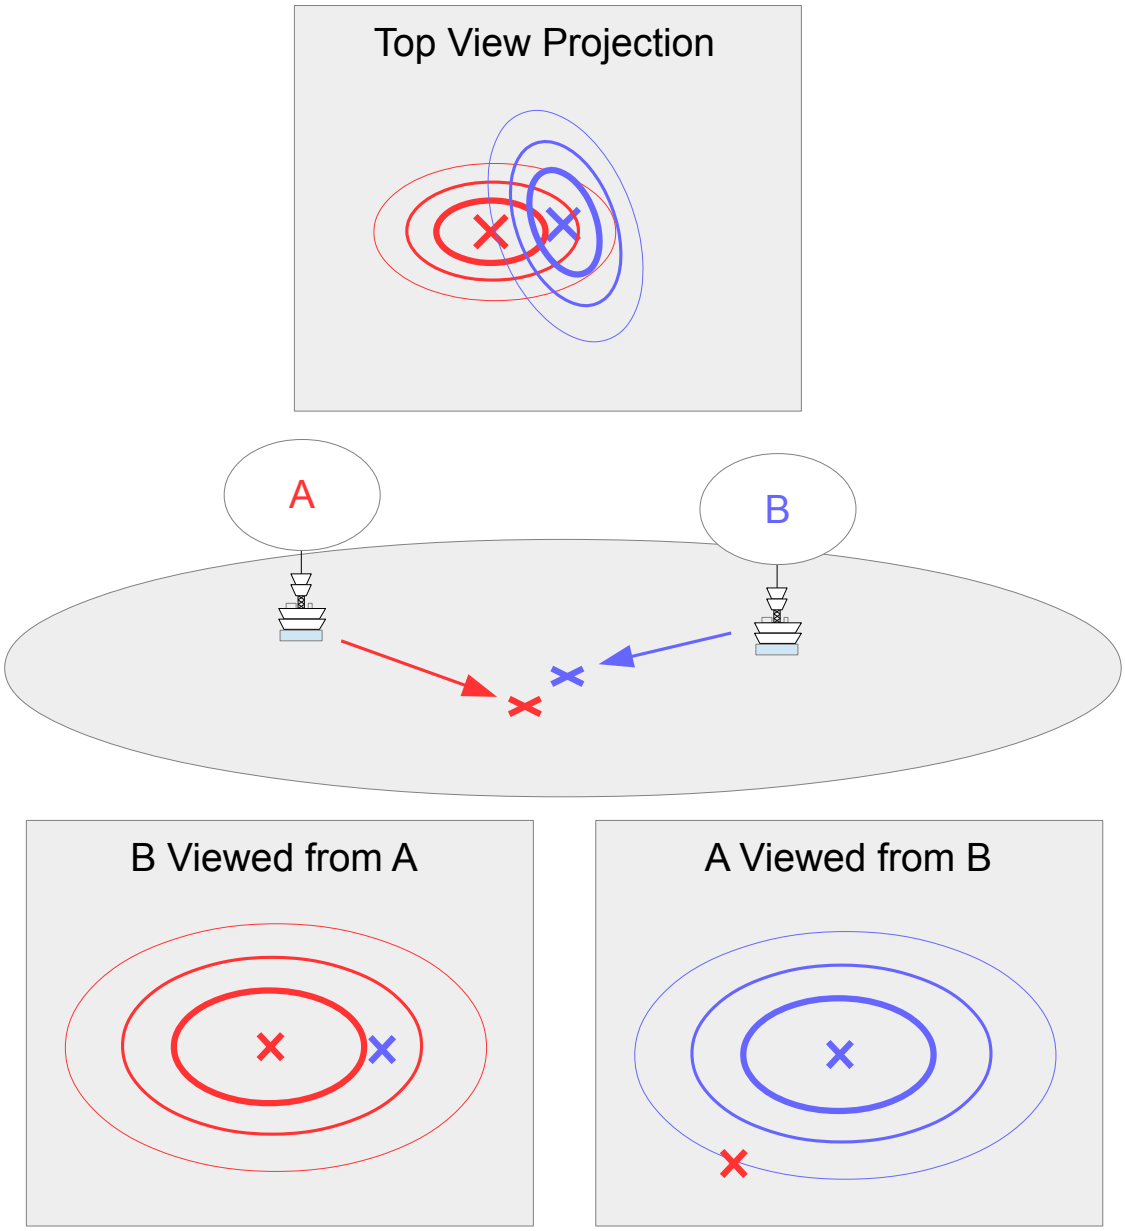

Figure 4.24: A simplified diagram showing the projection of two events recorded from two different locations into the field of view of each other. The Xs denote the location on the continent where the event pointed. The elliptical paths drawn around each location represent the 1σ, 2σ, and 3σ pointing confidence intervals. In this example, event B (blue) is less than2σ away event R (red) from the point of view of event R. However event R is more than 3σ away from event B from the perspective of event B.

A_ϕ	0.68726
B_ϕ	0.338725
A_θ	0.132516
B_θ	0.1557

Table 4.3: Fitted parameters for inverse relationship between SNR and pointing uncertainty. θ refers to elevation angle, and *phi* refers to azimuth angle.

position from which event B was captured, which will be denoted as $\Delta\theta_{BA}$ and $\Delta\phi_{BA}$. Since the elevation and azimuth angles are orthogonal, the total angular separation would be Equation 4.7

$$\sqrt{\Delta\phi^2 + \Delta\theta^2} \tag{4.7}$$

This equation however does not factor in the pointing uncertainty difference between azimuth (σ_phi) and elevation (σ_theta). Due to the longer baseline differences, elevation angles can be pointed with higher certainty than azimuthal angles. The normalized angular separation, which has units of angular uncertainty, is shown in Equation 4.8.

$$\sqrt{\frac{\Delta\phi}{\sigma_{phi}(SNR)}^2 + \frac{\Delta\theta}{\sigma_{theta}(SNR)}^2} \tag{4.8}$$

Note that the pointing resolution of of any single event is a function of the SNR, and thus any clustering analysis must factor this into its calculation. The pointing resolution as a function of SNR was found from analyzing self triggered WAIS pulses and their subsequent pointing to the known calibration pulser location on the continent. This was discussed earlier in the Calibration chapter. Events with SNR values lower than the lowest WAIS pulse, a value of approximately 4.0, are treated as if they had the pointing uncertainty of an impulse with an SNR of 4.0. An inverse relationship between SNR and pointing resolution is assumed based on a fit to the data. This is shown in Equation 4.14.2.

$$\Delta\theta = \left(\frac{A_\theta}{SNR}\right) + B_\theta$$
$$\Delta\phi = \left(\frac{A_\phi}{SNR}\right) + B_\phi$$

The values determined from the fit are shown in table 4.3.

The figure of merit is then the geometric sum of the pointing uncertainty normalized angular separation from each event observation location, shown in Equation 4.9.

$$L = \sqrt{\frac{\Delta\phi_{AB}^2 + \Delta\phi_{BA}^2}{\sigma_\phi(SNR)} + \frac{\Delta\theta_{AB}^2 + \Delta\theta_{BA}^2}{\sigma_\theta(SNR)}} \tag{4.9}$$

Using this equation we can determine the relative "closeness" of any two events, as derived from the measurements and including the systematic uncertainties.

4.14.3 Distance clustering

Events that reconstruct with steep downward elevation angles will not have high log likelihood values even to events that occur close in physical distance. To prevent events that may be related physically from avoiding clustering cuts, an additional condition for considering two events to be clustered is that they fall within 50km of each other on the ice. This is calculated using the great circle distance between the mapped locations of the events onto the ice.

4.14.4 Impulsive source identification

In order to exclude events that are likely to be anthropogenic background, it is necessary to determine locations on the continent from which impulsive RF transients are produced. Since events created by Humans are likely to emanate from discrete sources, this can be done by determining locations in which many impulsive events occur close to one another. This is done by marking events that fall within a log likelihood of $L < 40$ from any other event. Events with a single other event falling within this threshold are marked as pseudo-bases. By keeping the information for all events that fall within a cluster, instead of combining them to a single mean point, clustered impulsive source locations become elongated and exclude a larger region of the instrumental exposure. This is done conservatively to keep anthropogenic background from leaking into the signal region.

4.14.5 Event Clustering

The final set of signal cuts will provide a set of events that are most likely to be signal. These events are then clustered with all events tagged as pseudo-bases, and removed from the signal sample if they do. The remaining events will then be isolated impulsive transients that closely match the simulated radiation from an EAS, and will be the final candidates of the analysis.

CHAPTER 5
RESULTS OF ANALYSIS CUTS, CANDIDATES, BACKGROUNDS, AND LIMITS

Using the reduced quantities, clustering algorithm, and a priori determined cut values described in the previous chapter, a set of final CR candidates was determined. A quantitative accounting of these results is detailed in the following chapter.

5.1 Results of Impulsivity Analysis Cuts

After making the impulsivity cuts on the quality event dataset, only 5997 events remain. This represents 1.44×10^{-4} of the initial Hpol triggered dataset. Table 5.1 describes the number of events that failed each reduced quantity.

5.1.1 Above horizon events

Of the events that pass the impulsivity cuts, 641 occur above the horizon and do not trace back to the continent, but are below zero elevation. Included in this set would be CRs interacting in the atmosphere being observed directly. These events cannot be clustered in the same way as events that land on the continent. Since they are individual observations from a single location, there is no ability to use parallax to determine their source location depth. Because of this, a method of temporal and pointing clustering is used to determine if they are part of the same source distribution. This is not done until the final signal cuts are made.

Cut Parameter	Events Failed (Single)	Events Failed (Cumulative)
Map Peak	2238151	2238151
Map SNR	2247560	25364
Hilbert Peak	2056059	39710
Linear Polarization Fraction	2356046	90502
Coherent Template Correlation	2380525	12160
Deconvolved Template Correlation	2318737	738

Table 5.1: The number of events removed by each impulsivity cut from the set of events that has already been reduced with quality cuts, mentioned in Table 4.1. The cuts remove any event with a single value below those in Table 4.2. The column marked Single is the number of events that fail that specific cut parameter if applied to uncut quality data. The column marked Cumulative is if each cut was taken sequentially, and the event removed if failing, before moving to the next.

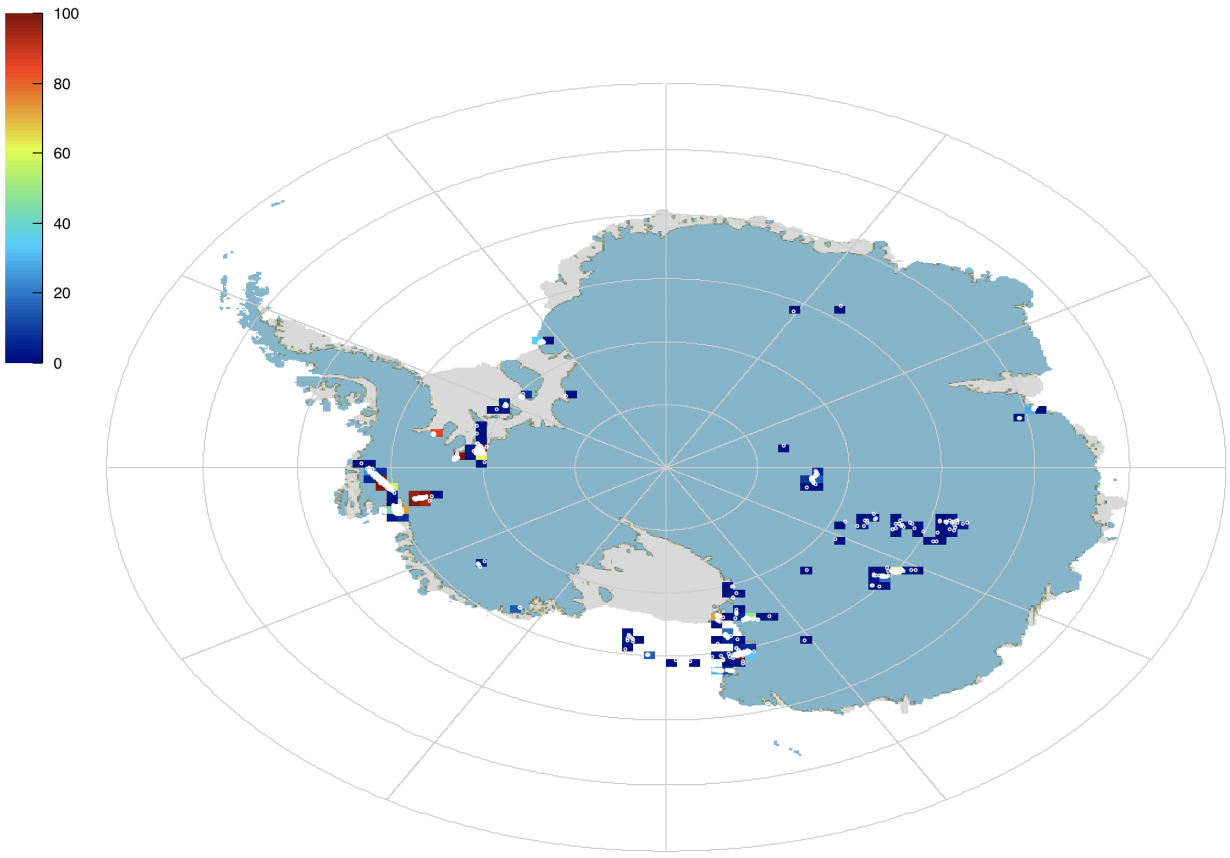

Figure 5.1: A 2D histogram overlaid on a map of Antarctica, showing the locations of the 5324 clustered events. White circles denote the location of an event that was determined to have clustered. Events are only required to have one other event with $L < 40$ to count as a cluster. The Z axis shows the number of events that fall into each bin.

5.1.2 Event pointing results

Of the 5356 events that point below the horizon, when compared against themselves, only 32 fall at least $L > 40$ away from any other event in the set. This clustering result is shown in Figure 5.1. These events are impulsive in nature, and contain the final candidate set.

5.2 Results of Signal Analysis Cuts

Making the final CR signal cuts on the weak event list, regardless of clustering status, yields 882 events. 873 of these project onto the continent, and 9 point above the horizon. Of the nine that point above the horizon, 7 can be removed due to either clustering with events that fall at the horizon, temporal and true azimuth clustering. These removed events, and the clusters they associate to, can be found in Table 5.2. The below horizon CR candidates can be determined by

Event Number Removed	Clustering description
48160214	Clusters with other events on the continent with 0.12 downward θ adjustment
58112462	Clusters with other events on the continent with 0.02 downward θ adjustment
69050305 69050312 69050331	Cluster with other points on the continent with $< 0.1^\circ$ downward θ adjustment
80134362 80299371	Part of a 401 above horizon event burst lasting 624 seconds. Event numbers 79995831 - 80745194.

Table 5.2: Above horizon clusters that contain events that pass final signal cuts. 7 events are excluded from the final candidate list due to events in close time proximity, pointing in the same direction.

Cut Parameter	Events Failed (Single)	Events Failed (Cumulative)
Hilbert Peak	147	147
Linear Polarization Fraction	3060	2948
Coherent Template Correlation	4179	1920
Deconvolved Tempalte Correlation	4265	97

Table 5.3: The number of events removed by each signal cut from the set of events that has already been reduced with impulsivity cuts, mentioned in Table 5.1. The cuts remove any event with a single value below those in Table 4.2. The column marked Single is the number of events that fail that specific cut parameter if applied to the impulsivity data. The column marked Cumulative is if each cut was taken sequentially, and the event removed if failing, before moving to the next.

choosing events within this, and results in 18 downward pointing CR candidate events, and 2 events within the atmosphere, for 20 events total.

Table 5.3 describes the number of events from the impulsive set that failed each reduced quantity. Note that interferometric map peak and map SNR do not change between the cuts, so they are not included in this table.

5.2.1 Candidate locations on continent

The locations of the final 18 below horizon and 2 above horizon candidates can be seen in Figure 5.2.

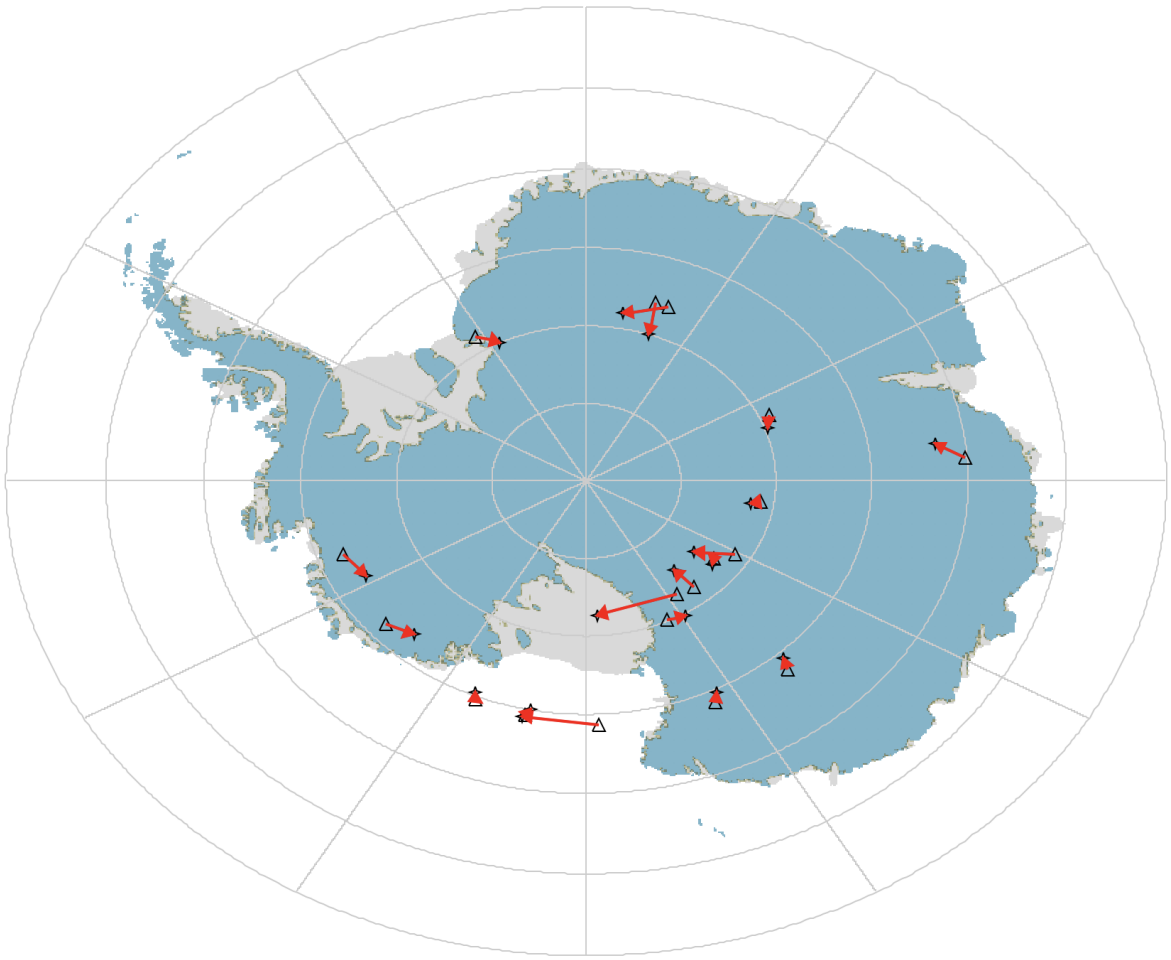

Figure 5.2: Locations of the CR candidate events discovered in the analysis. The triangle is the location of ANITA, the X denotes the location on the continent where the event traces, and arrow is the vector between the two. Above horizon events, which do not project to the continent, are shown with blue arrows indicating their observation direction.

5.3 Geomagnetic Analysis

The final characteristic of CRs that can be used to check the events categorized as candidates by cuts and clustering is a comparison between expected and measured geo-magnetically induced plane of polarization. Using the process detailed in section 4.9, the expected geo-magnetic polarization expectation was generated using MATLAB and code generated for the ANITA-I CR analyses. Note that events that are observed above the horizon are not reflected off the ice-sheet, and have a slightly different calculation. The resulting plot of measured vs. expected can be seen in Figure 5.3. The angular difference between measured and expected, along with a gaussian fit to the distribution, can be seen in Figure 5.4.

5.4 Background Estimates

Estimating the number of background events that passed all signal cuts and are present in the final candidate list is import for establishing the confidence in the final result. The strong cut on template correlation is expected to eliminate the vast majority of non-astrophysical thermal background signals, however anthropogenic sources may still produce impulsive events that leak into the signal box. This is addressed in several estimates that determine the likelihood of any individual event being anthropogenic.

5.5 Diffuse Source Background Estimate

It is possible to determine whether signals look like other observed signals coming from regions of the continent determined to be base-like from the clustering. This is accomplished using statistical inference about the probability that non-physics events would satisfy all cut requirements [58]. We attempt to divide the events into two sets, one background and one signal, then use the arithmetic of probability to determine the likelihood that an event from the background set could be found within the signal set.

In order to relate the measured data to uniform

Bayes' theorem is stated in Equation 5.1, and can be used to determine the estimate of the background.

$$P(B|S) = \frac{P(S|B)P(B)}{P(S)} \tag{5.1}$$

Where P(B) is the number of events in the background set

Two different philosophies can be referenced for determining what goes into the background set. The signal set can be defined as either the expected observation rate of UHECR particles during the flight, or as a single count for each measured candidate event individually which is then

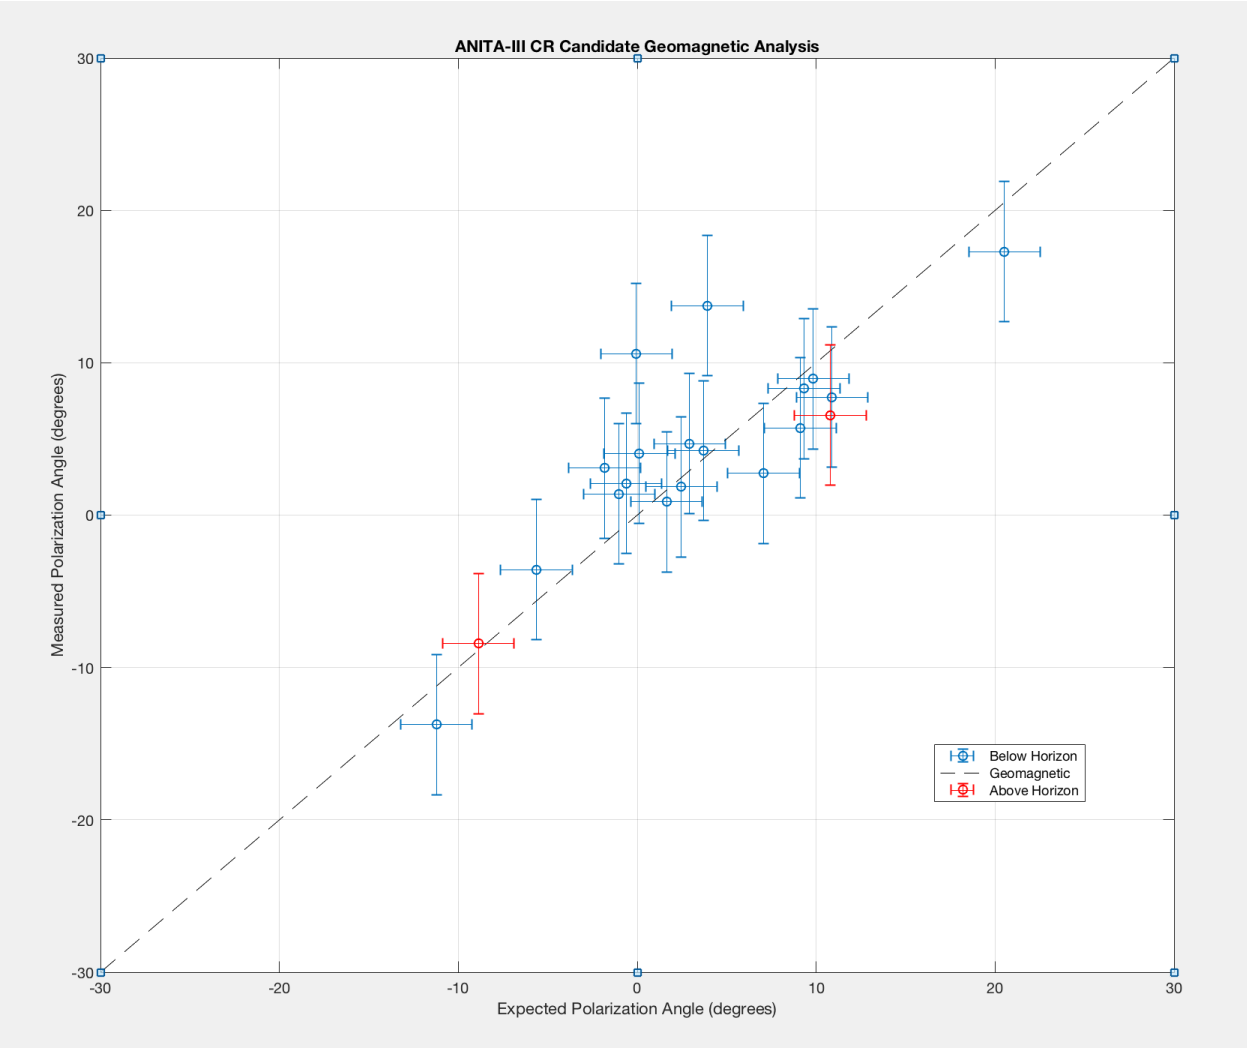

Figure 5.3: Plot of the expected geomagnetically induced polarization angle compared to the measured polarization angle for all 20 candidate events. Red markers denote events which occurred above the horizon, which needed to be treated differently as they do not undergo a reflection off the ice. Blue markers are events that are observed below the horizon. Uncertainty on the expected angle of 2° is dominated by the lack of information as to where on the cone the event was observed. Uncertainty on the measured angle of 4.6° is determined through analysis of WAIS pulses and is detailed in subsection 4.8.4.

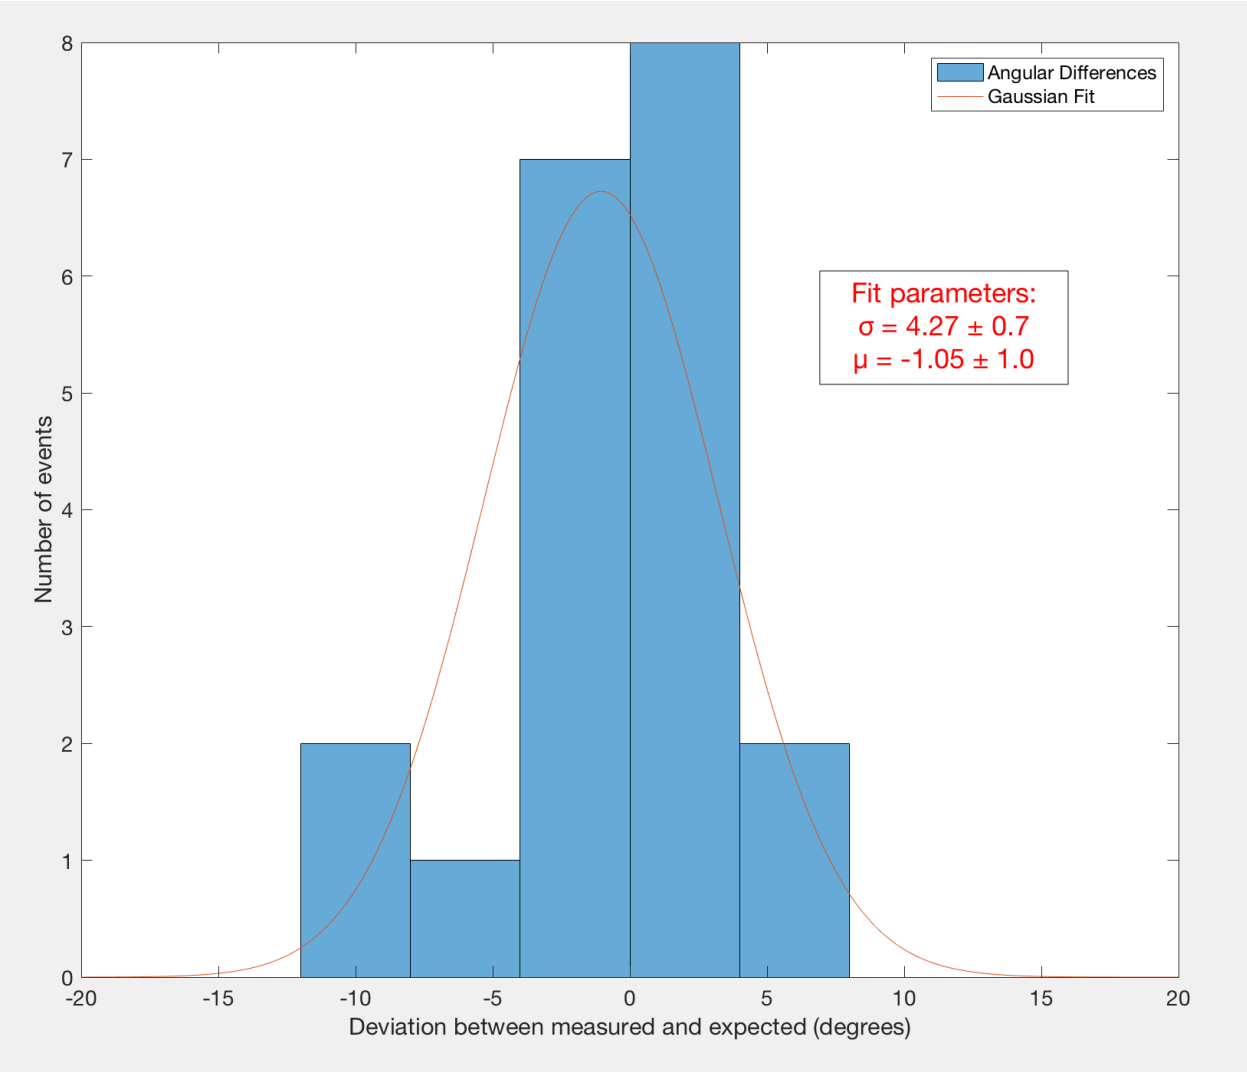

Figure 5.4: Histogram of the difference between expected geomagnetically induced polarization angle and measured polarization angle for all 20 candidate events. A gaussian fit is overlaid on the data, and has a mean of 1° and a width of 4.27°.

summed to find the total background probability.

5.6 Local Source Background Estimate

Another method to determine the likelihood that any particular candidate is generated by a terrestrial source is to compare it to the other events that are geo-associated with the candidate events. If a background source generator exists, one that produces events that are non-thermal but do not pass the impulsivity or signal cuts, it should elevate the distributions of the calculated values used to create cuts. This does not account for background sources that produce single CR-like signals but behave thermal otherwise.

This estimate requires an assumption that such sources are anthropogenic in nature, and that anthropogenic sources do not produce single high powered impulses without an associated amount of detectible background signal. Validation of this assumption can be accomplished by repeating this estimate with isolated events that pass impulsivity cuts, but not signal cuts. The number of observed events of this type should be replicated by this quantitative estimate.

5.6.1 Geo-associated events

To determine events that are geo-associated with the candidate events, the clustering algorithm described in section 4.14 is utilized. All events that pass quality cuts, regardless of signal parameters, are compared versus each candidate event and flagged if they cluster with the event. For events that are steeply inclined, the 50km geographic radius dominates the clustered events. A resulting map of events geo-associated with a single candidate (62273732) can be seen in Figure 5.5.

5.6.2 Joint product (JP) statistic

To directly compare events, each with multiple variable quantities that describe their similarities to expected signals, a single metric is needed. The simplest way to do this is by simply taking the product of the quantities that are most likely to be largest for signal. This new value, which will be called the joint product (JP) statistic, can then be used to quantitively compare different event populations.

To determine which values will be maximal for signal events, but suppressed for both thermal and anthropogenic background noise events, a qualitative look at candidate distributions and full quality dataset distributions was done. From this, it was determined that many parameters used in cuts were actually larger for anthropogenic noise events than for candidate events. These included interferometric map derived values, such as mapPeak and mapSNR, as well as waveform information, such as the peak of the Hilbert envelope. Since these values are proportional to total coherent waveform power, it is not unexpected that large signals from man-made sources have high values. However, the template correlation value for both the de-dispersed and coherently summed

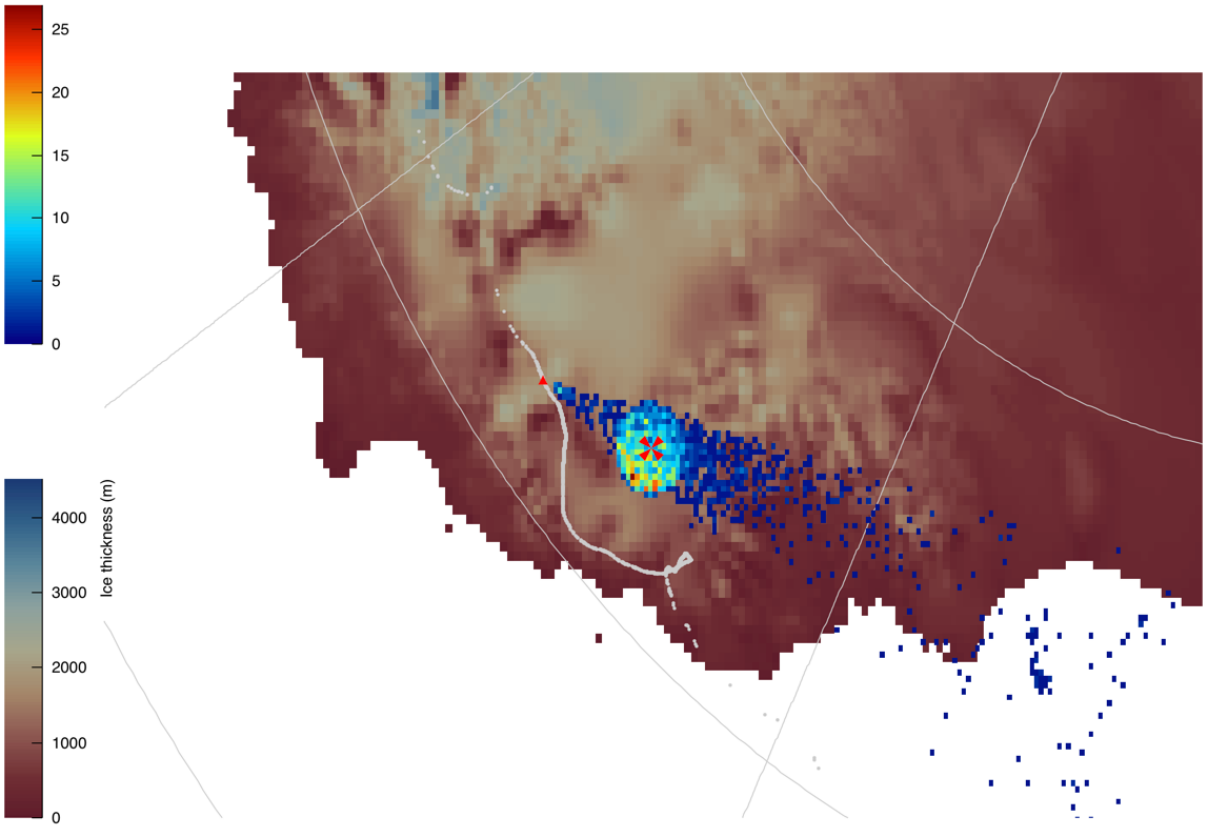

Figure 5.5: Locations projected onto the Antarctica continent of quality events that geo-associate with event number 62273732. Visible is the cone of events that geo-associate via the log-likelihood clustering method and the circle of events that fall within 50km of the candidate. The 2d histogram color maps to the number of events in each grid rectangular bin, which do not have equal sizes when projected onto the Earth geoid.

waveforms, required to be large for candidate identification, remained small for anthropogenic populations.

An additional parameter not used to make cuts, the difference between the expected geomagnetically induced plane of polarization and that measured in any particular event, is used as a quantity in determining the JP statistic. This calculated value, as noted in section 4.9, was excluded from cuts in order to preserve a quantity that is both highly correlated to CR signals and uncorrelated to the final purified candidate list. Due to the uncertainty in both the expected calculation and event plane of polarization measurement, for the purposes of the JP statistic, it is desirable to avoid penalizing events that fall within a few standard deviations of an angular difference of zero. To accomplish this, a value of 1 was assigned to events with a difference between measured and expected plane of polarization less than 2.5σ. Outside this range, the value falls off as $\frac{2.5}{\sigma_{diff}}$, where σ_{diff} is the number of sigma for each event. 2.5 is used to prevent a discontinuity at the boundary of the function. The 2.5σ value is treated as a final cut parameter on the candidates.

The final JP statistic is then shown in Equation Equation 5.2.

$$JP = T_c \times T_d \times G \tag{5.2}$$

T_c and T_d represent the template correlation values for coherently summed and de-dispersed signals respectively. G represents the value determined from the expected and measured plane of polarization.

5.6.3 Background from candidate regions

By combining all events that geo-associate with the candidate events, we can determine an expectation for the likelihood that an event would fluctuate higher than those local distributions. For the 20 candidate events, there were 59,498 quality events that geo-associated with them. The distribution of their JP values can be seen in Figure 5.6

To determine an expected count of events from the background distribution that would surpass the signal cut value, we must integrate from the cut value out to infinity. Since, by construction, there can be no background events that surpass the cut value, a fit and extrapolation must be made for the background distribution, for which an exponential was fitted for the tail of the distribution only. The value of the integral will be properly normalized to then establish the number of background events expected within the geo-associated region. The background estimate for the full flight can be extrapolated from that by multiplying by the ratio of events within the geo-associated region to events that are not excluded by clustering in the remainder of the flight.

A plot highlighting the behavior of the distribution at the tail as well as the exponential fit is shown in Figure 5.7.

For an exponential described by $y = exp(A * x + B)$, the resulting values of the fit are A=10.02 and B=-36.36. Integrating this from the cut value to one results in an estimated background count

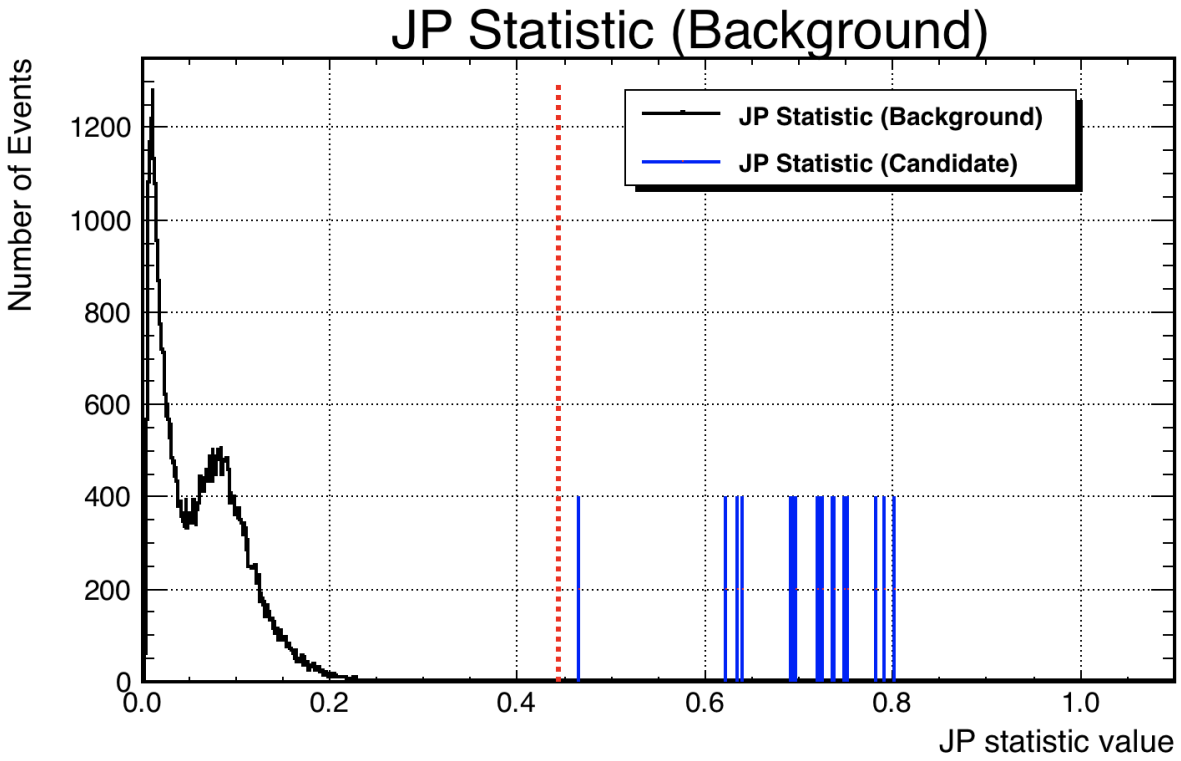

Figure 5.6: Histogram of the distribution of calculated JP values for events geo-associated with the candidate events (black). The red dotted line denotes the location of the signal cut parameters. Events with values higher than the red line would, assuming they survive the clustering phase of the analysis, be classified as candidate events. Blue points denote the JP statistic values of the candidate events. Their y-axis heights are scaled for visibility, and are not meaningful.

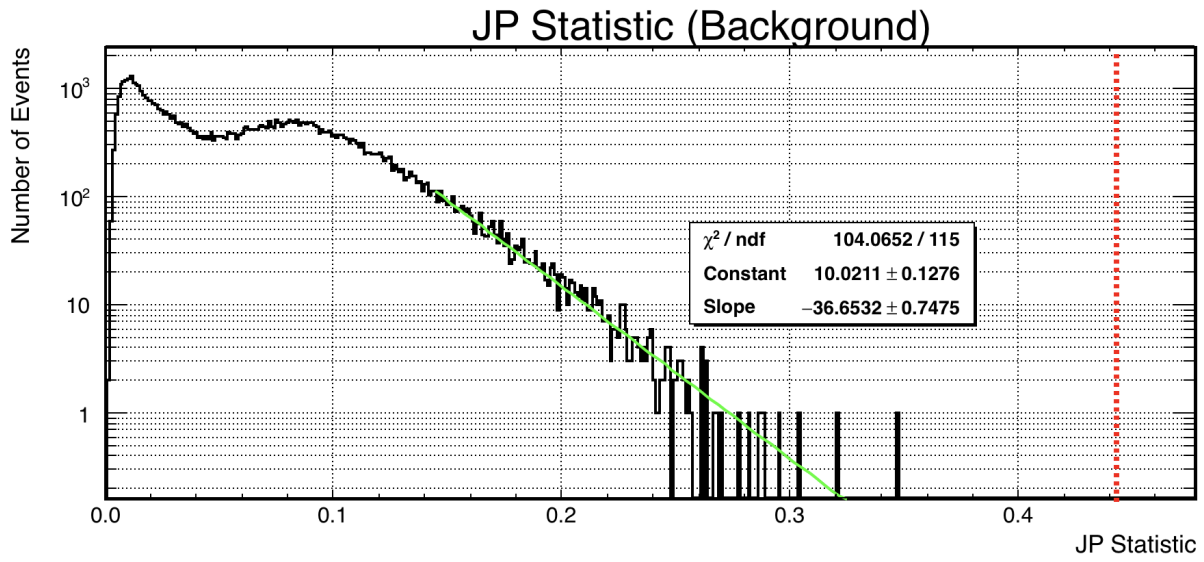

Figure 5.7: An exponential fit to the tail of the distribution of candidate geo-associated JP values, shown with the y-axis in log scale to highlight the behavior at low count rates. This fit can be extrapolated to determine the likelihood that an event will randomly exceed the cut threshold, denoted by the red dotted line.

of 6.133×10^{-5} for the geo-associated region.

There are 1,174,842 events that are neither impulsive themselves, nor cluster with any impulsive events (See "The Wasteland" in Figure 5.55). These events represent the portion of the flight in which CR signals could be detected, but where there was no signal present. This region contains 19.6 times more events than the region geo-associated with candidates. Scaling the estimated background found from the candidate region by the fraction of the flight it represents yields a background event count of 1.20×10^{-3}.

5.6.4 Check with "Dirty Dozen" impulsive isolated events

We can check the validity of the previously described estimated rate by determining the likelihood for an impulsive, but not signal, event to be found. Integrating the exponential fit of the JP statistic for the geo-associated candidate region from the impulsivity cut (0.5^2) to the signal cut (0.666^2) yields a value of 6.964×10^{-2}. Once again multiplying this value by the fraction of the flight it represents yields an expected impulsive background rate of 1.36 events.

There were 10 observed impulsive isolated events in the flight. Of these, 5 have a plane of polarization within 2.5σ of the geomagnetically induced expectation. This suggests that the wasteland region of the flight has a higher average background than the regions geo-associated with candidate events by a factor of 3.67. Scaling the background estimate established from the geo-associated region by this factor yields a final signal background event rate expectation of 4.4×10^{-3}.

5.7 Cosmic Ray Candidate Results

There were 18 events that passed all signal cuts and did not cluster to any other impulsive event on the continent, and are hereby known as the cosmic ray candidates. Of these 18, 2 point above the horizon but still within the atmosphere, indicating that they are directly viewed EAS signals that do not reflect off the continent. The remaining 16 point downwards and project onto the continent. These events, their waveforms, their locations on the continent, and the distributions of nearby events are presented in the following section. A map summarizing the results of the analysis, including the locations of the events on the continent, along with events that cluster and events that pass impulsivity cuts but fail signal cuts, is shown in Figure 5.8. An overlay of measured horizontally polarized waveforms for all candidates, aligned to their peak correlation and polarity flipped relative to their estimated polarity, is shown in Figure 5.9. The vertically measured waveforms have significantly less signal but are presented with alignment determined through H-Pol correlation in Figure 5.10. Additionally, the coherently summed waveforms are presented side by side in figure 5.11.

Locations of the events traced to the continent, along with the location of ANITA at the time of detection and the elevation angles and true azimuth with respect to north, can be found in Table 5.4. The calculated cut values of each event is shown in Table 5.5. Plots of the waveforms and interferometric maps for all the candidates can be seen in Figures 5.12 through 5.51.

Event 9097075

5.8 $\nu_\tau$andidate Polarity Unblinding

After setting the cuts, determining the background estimate, and compiling a list of candidate events, permission to look at the results of the polarity estimator was asked of the collaboration. The results of the unblinding are shown in Table 5.52.

After unblinding the data, it was determined that 15 of the below horizon events shared the same polarity, opposite of the two above horizon events. 1 event that points below the horizon, event number 15717147, has an inverted polarity. A plot overlaying H-pol waveforms for below horizon events, with their measured polarities, is shown in Figure 5.53. A plot overlaying the flipped polarity below horizon event with the observed above horizon events is shown in Figure 5.54.

This flipped polarity event is consistent with a directly observed EAS shower. The measured elevation angle of the event was 35° below the horizontal, well below the horizon. A directly observed EAS shower with incident particle emerging from the ice is the expected signal from a UHEν_τ. However, additional study is required to confirm this result.

Event Number	Event Latitude	Event Longitude	Event Alt (m)	ANITA Latitude	ANITA Longitude	ANITA Alt (m)	Elevation (°)	True Azimuth (°)
9097075	-79.8303	148.814	2131	-80.0817	154.602	36323.4	16.8661	-78.8563
11116669	-81.3294	176.136	6	-81.2609	146.761	36355.2	6.43217	105.382
11989349	-82.6084	141.396	2405	-81.1009	140.177	35438.7	11.8167	174.066
15717147	-81.3986	129.016	2621	-81.6109	126.507	35861	34.9699	61.41
16952229	-82.676	128.674	2742	-80.8189	121.333	34681.6	8.73916	154.019
19459851	-81.189	99.517	3544	-80.6972	98.2029	35995.6	28.8256	157.821
23695286	-79.8056	70.1918	3851	-79.4478	65.9819	35390.9	19.0219	117.305
27142546	N/A	N/A	N/A	-79.1585	50.943	35760.8	5.5537	150.221
32907848	-79.0226	10.1545	2891	-77.9829	21.1706	36263.5	8.25657	-120.653
33484995	-79.9982	19.3425	3089	-77.9758	17.7245	36873.6	9.48746	172.1
39599205	N/A	N/A	N/A	-78.2385	-17.9381	36120.6	3.61121	151.327
41529195	-79.9325	-26.977	362	-78.9955	-32.116	36947.3	14.553	137.408
58592863	-76.9118	-117.691	1803	-76.395	-110.151	35015.4	10.2213	-110.165
62273732	-76.5879	-137.478	1913	-76.0299	-131.171	35943.1	11.632	-113.518
66313844	-75.2196	-156.846	-59	-74.815	-157.437	36484.6	37.3596	159.579
68298837	-75.0579	-168.67	-62	-74.6522	-168.028	36419.3	36.8965	-157.819
70013898	-74.5006	-167.459	-63	-74.3386	177.572	36245.1	6.64026	99.4938
73726742	-74.8095	153.33	2149	-74.2488	154.376	36291.1	26.3685	-154.026
75277769	-74.5808	137.642	2708	-73.9239	138.908	34941.4	21.6866	-152.964
83877990	-71.5575	82.5361	2801	-70.1421	85.785	35007.7	10.2116	-144.515

Table 5.4: Pointed observation locations for each CR candidate, along with ANITA observation location, elevation, and true azimuth. Directly observed events do not point to the continent, and so their source hypothesis is undefined without reconstructing the shower vertex.

eventNumber	TemplateCorr Coherent	Template Corr De-dispersed	mapPeak	mapSNR	LinPolFrac Coherent	HilbPeak Coherent
9097075	0.848129	0.867255	0.275075	20.6464	0.892743	119.846
11116669	0.879512	0.888349	0.204412	21.3007	0.842908	69.5003
11989349	0.81368	0.851099	0.0863109	15.0991	0.621034	50.2824
15717147	0.852206	0.83971	0.142453	13.605	0.782663	51.9856
16952229	0.834765	0.83505	0.102279	15.6861	0.651961	48.931
19459851	0.822059	0.771824	0.165427	15.8911	0.737391	67.2559
23695286	0.85406	0.843391	0.0781213	14.53	0.634915	39.3999
27142546	0.870272	0.858413	0.176355	20.7221	0.834333	66.9543
32907848	0.899383	0.891038	0.171207	19.5985	0.644866	65.6926
33484995	0.793178	0.805227	0.173698	20.3076	0.802018	63.692
39599205	0.751907	0.749293	0.16824	22.0372	0.820529	96.5364
41529195	0.84424	0.855839	0.120433	17.3735	0.697294	49.7193
58592863	0.832529	0.829167	0.148024	17.7801	0.772304	49.8597
62273732	0.682009	0.682559	0.198354	19.5202	0.846162	58.9087
66313844	0.850128	0.867374	0.150474	16.0031	0.686217	47.5222
68298837	0.868809	0.865949	0.092407	11.1013	0.646708	51.914
70013898	0.800858	0.775827	0.165466	23.1731	0.776458	76.0259
73726742	0.85167	0.852433	0.141211	14.8105	0.795029	44.4288
75277769	0.887369	0.891011	0.159072	15.9893	0.846339	67.6613
83877990	0.827815	0.839115	0.180813	20.301	0.783432	66.958

Table 5.5: Calculated cut quantities for all events. They are abbreviated for space. They are, in order, the template correlation fraction between the coherently summed waveform and signal chain convolved simulated EAS impulse, the template correlation fraction between the de-dispersed measured waveform and the simulated E-filed of an EAS pulse, the peak of the interferometric map, the SNR of the interferometric map, the linear polarization fraction of the coherently summed waveform, and the hilbert peak of the coherently summed waveform.

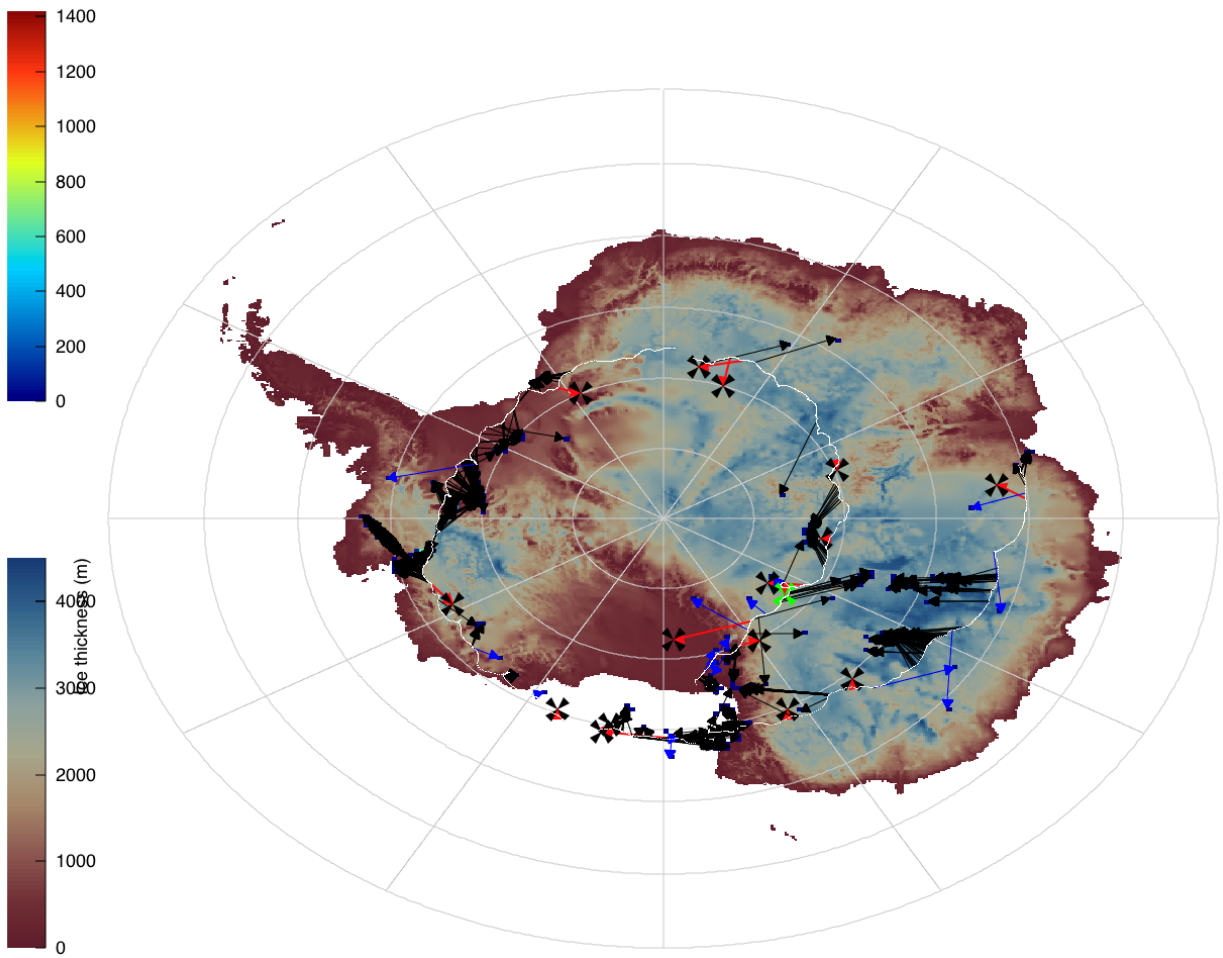

Figure 5.8: A map of Antarctica summarizing the results of the analysis. "X"s with red arrows denote the mapped location on the continent for the 18 CR candidates. The blue arrows show locations of isolated events that did not meet signal cuts. The black arrows and underlaying histogram shows events that were discovered with the impulsivity cuts, but failed either signal cuts or clustering cuts.

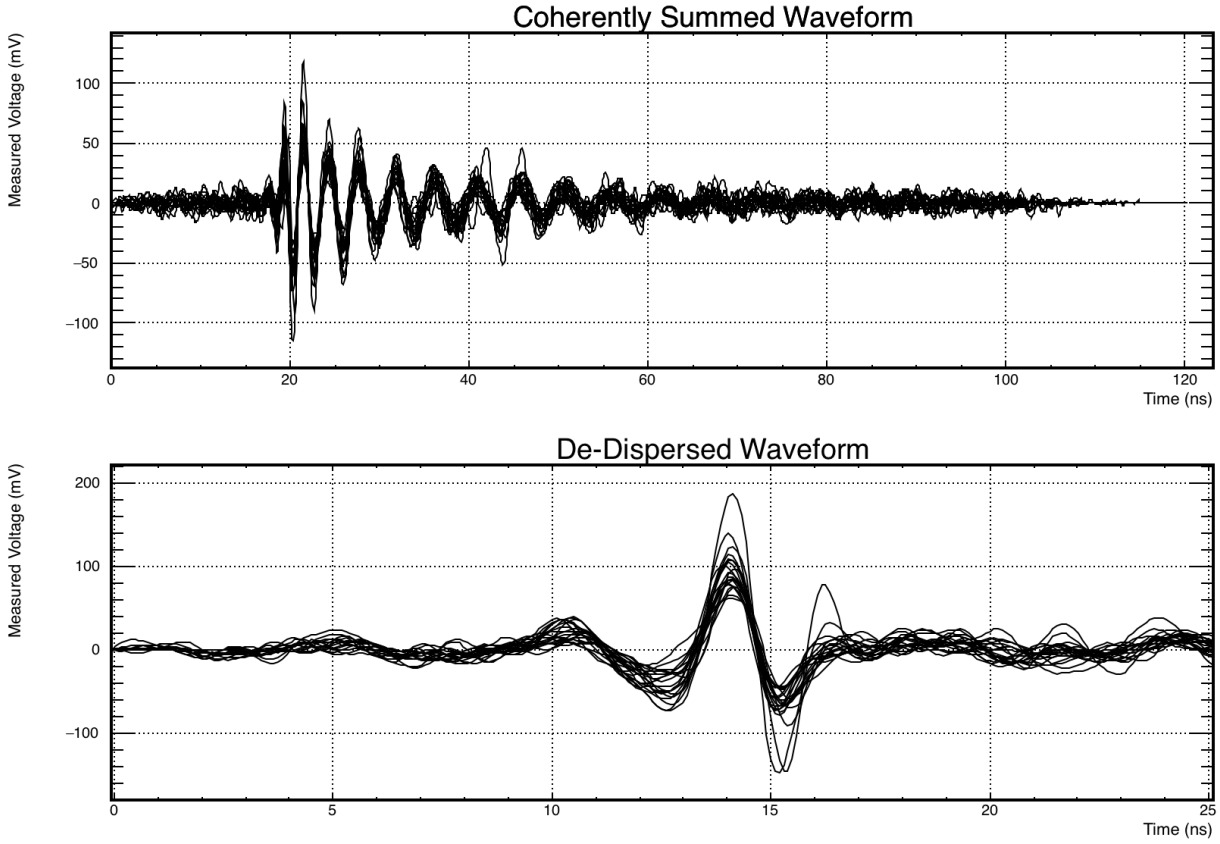

Figure 5.9: All 20 CR candidate H-pol measured signals events overlaid. Their blinded polarity has been aligned with respect to their estimated polarity. They have aligned to their peak correlation. Top: Coherently summed waveform. Bottom: All Pass Filter De-dispersed waveform.

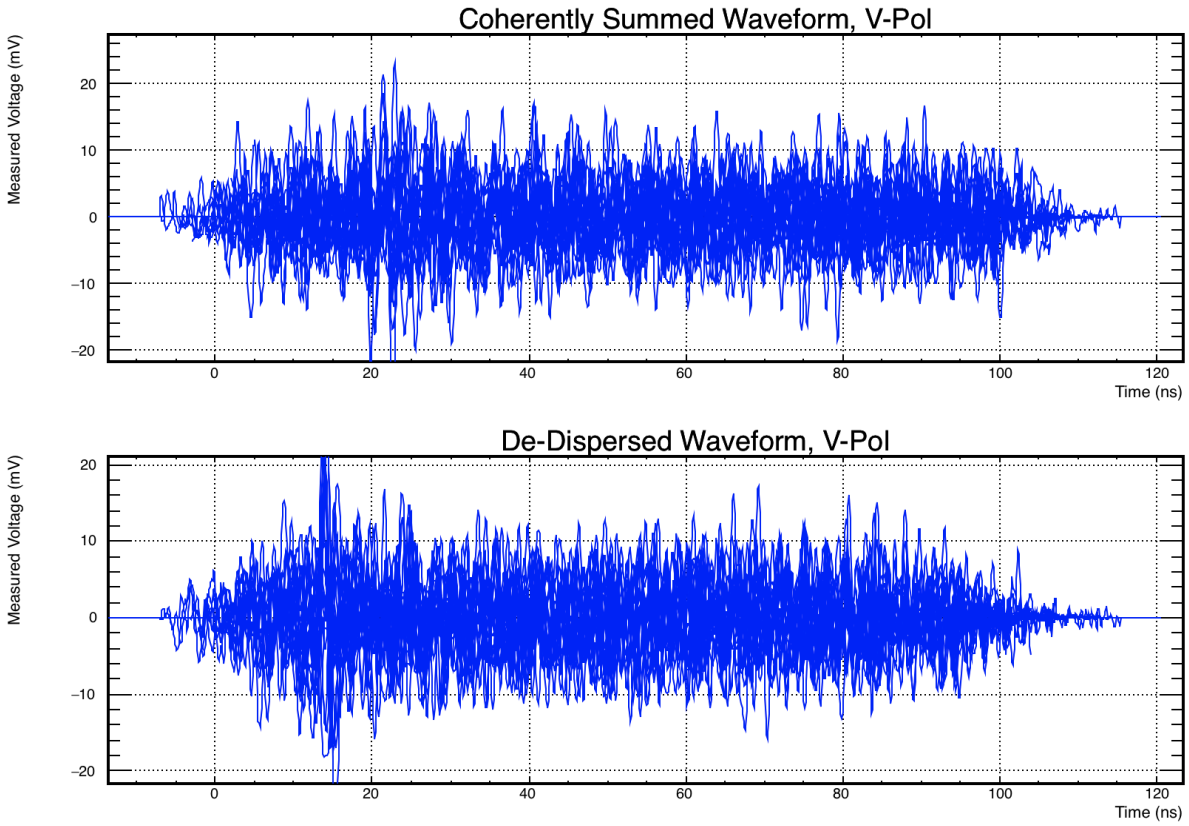

Figure 5.10: All 20 CR candidate V-pol measured events overlaid. They are aligned to the peak of the correlation of the H-pol signal, which has significantly more power. The amount of signal expected in this polarization is dependent on location in the geomagnetic field, and the incidence angle of the UHE particle. Top: Coherently summed waveform. Bottom: All Pass Filter De-dispersed waveform.

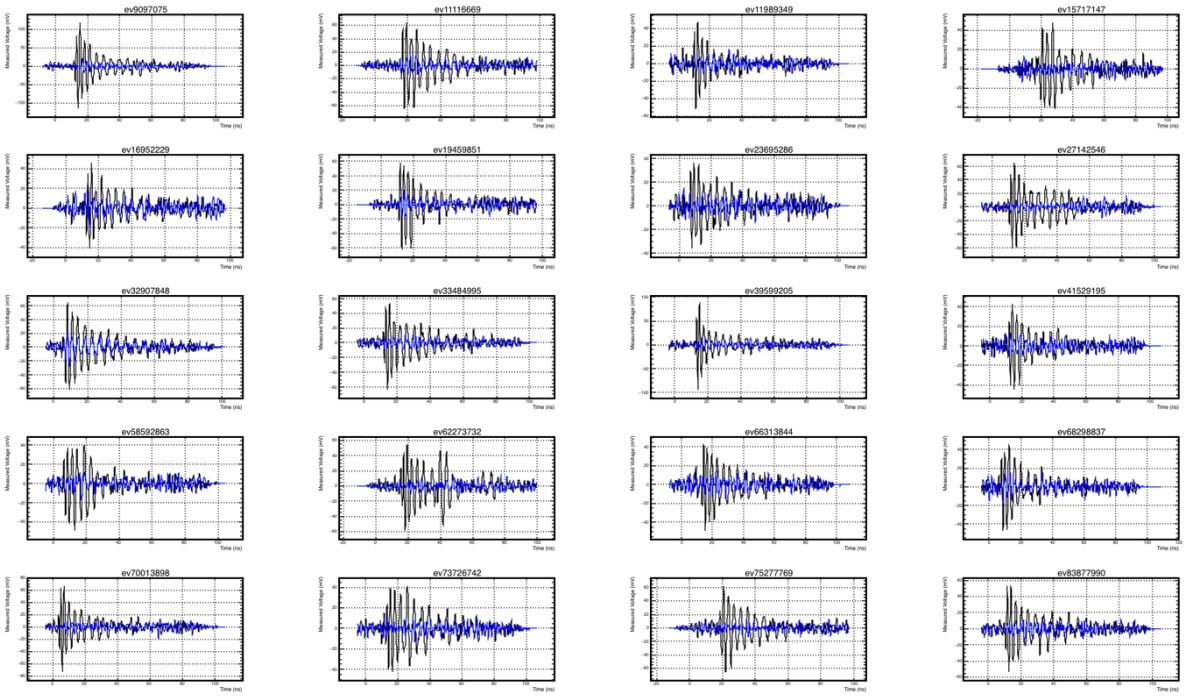

Figure 5.11: All 20 CR candidate events. Their blinded polarity has been aligned with respect to their estimated polarity. Black lines denote the H-Pol signal, while blue lines denote the V-Pol signal.

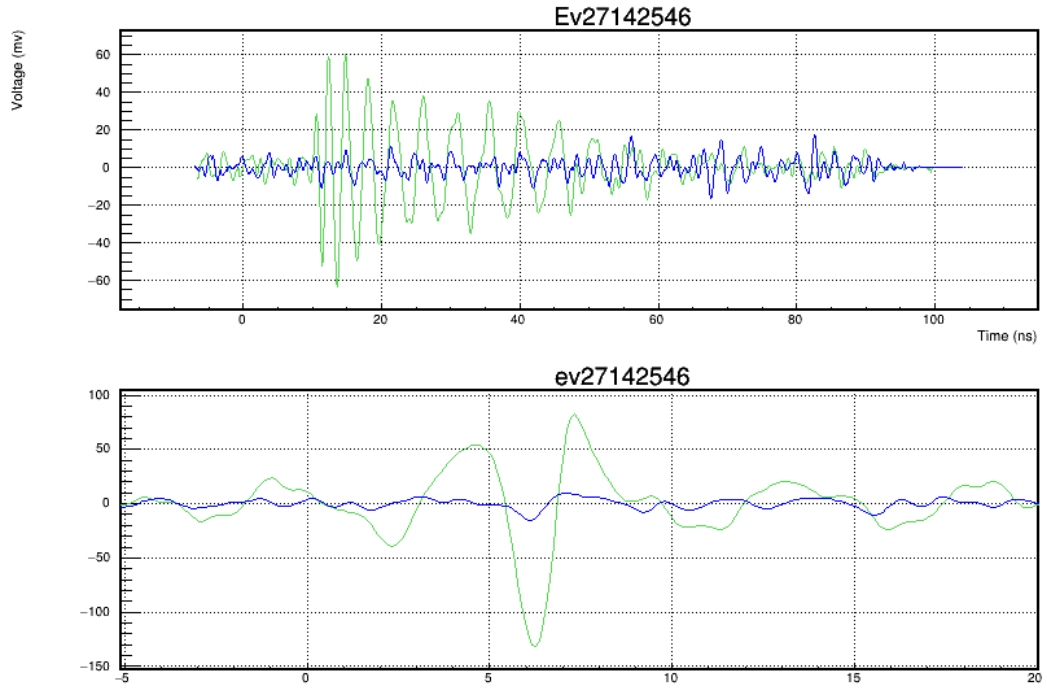

Figure 5.12: Top: Coherently summed waveform. Bottom: De-dispersed waveform

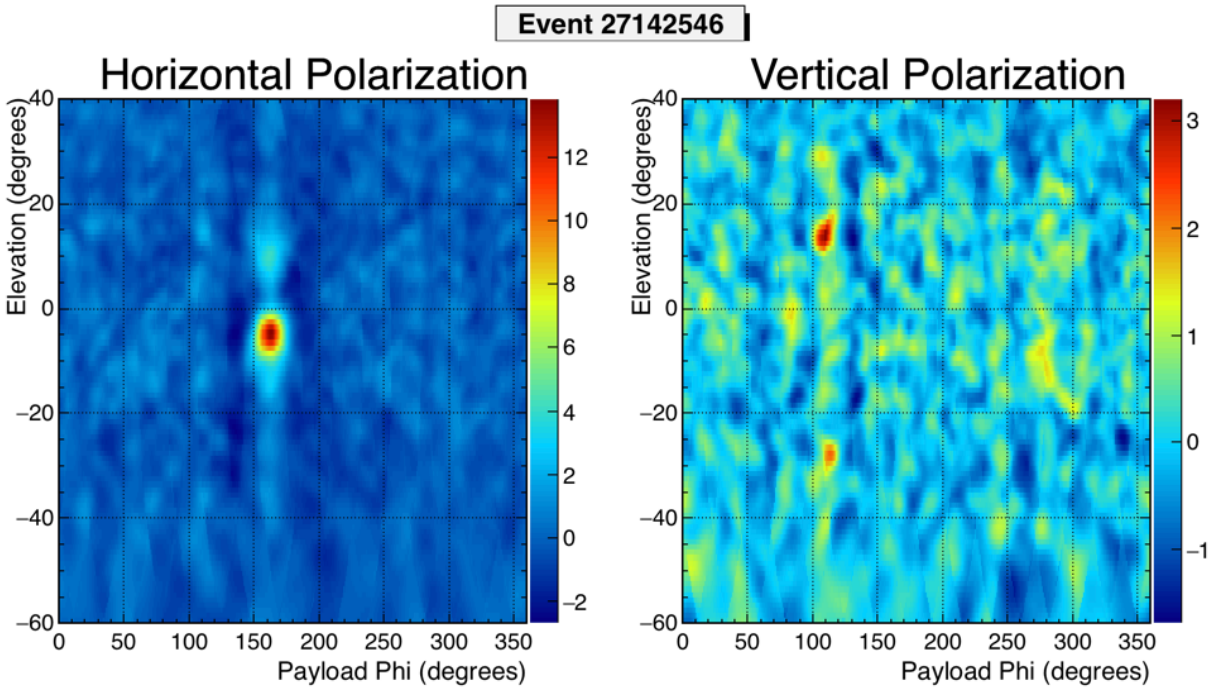

Figure 5.13: Interferometric maps for event 27142546

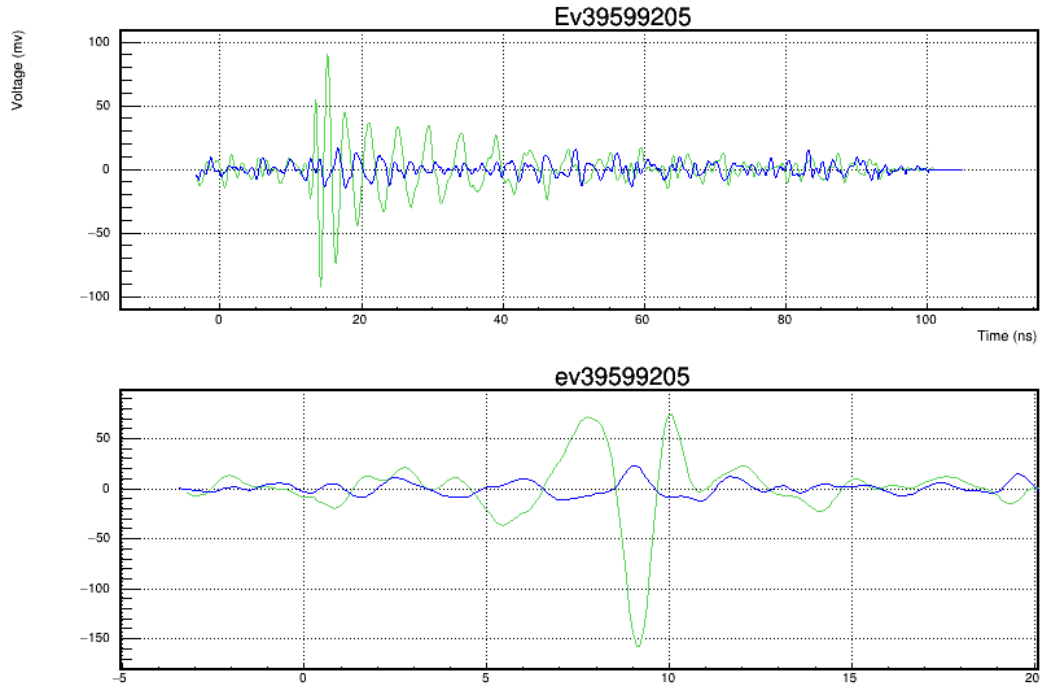

Figure 5.14: Top: Coherently summed waveform. Bottom: De-dispersed waveform

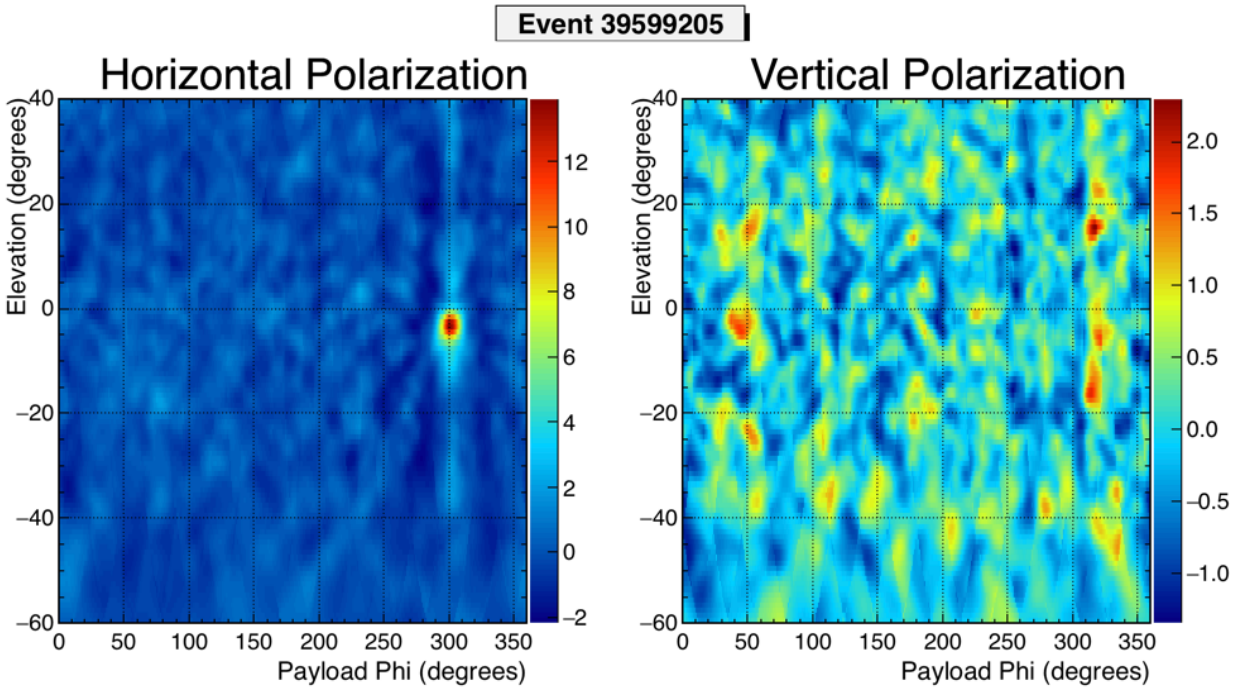

Figure 5.15: Interferometric maps for event 39599205

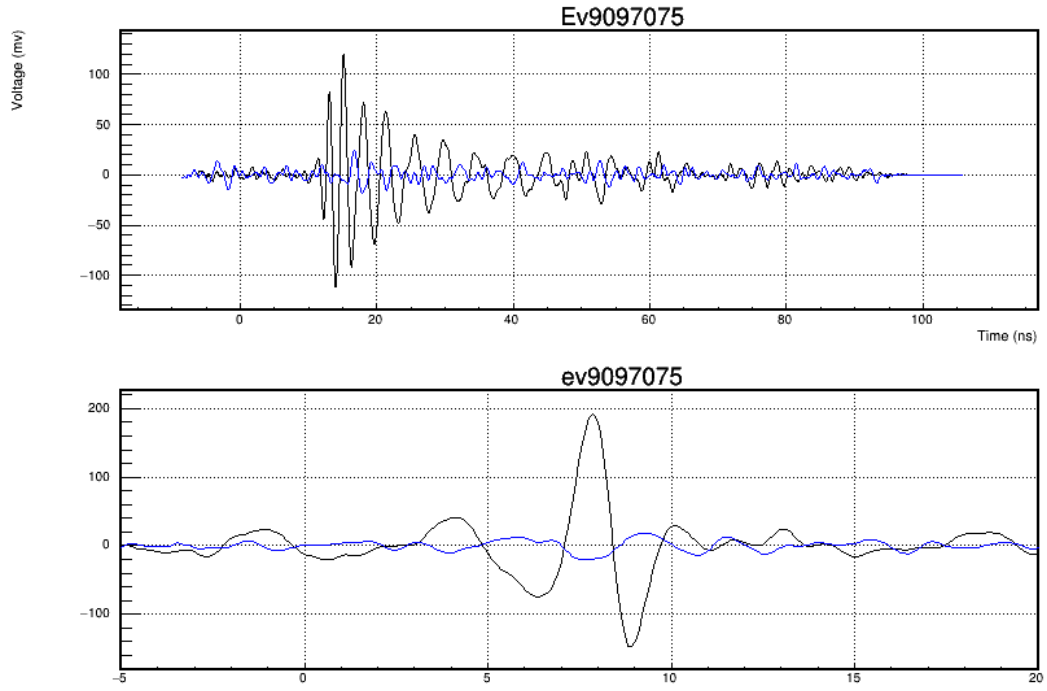

Figure 5.16: Top: Coherently summed waveform. Bottom: De-dispersed waveform

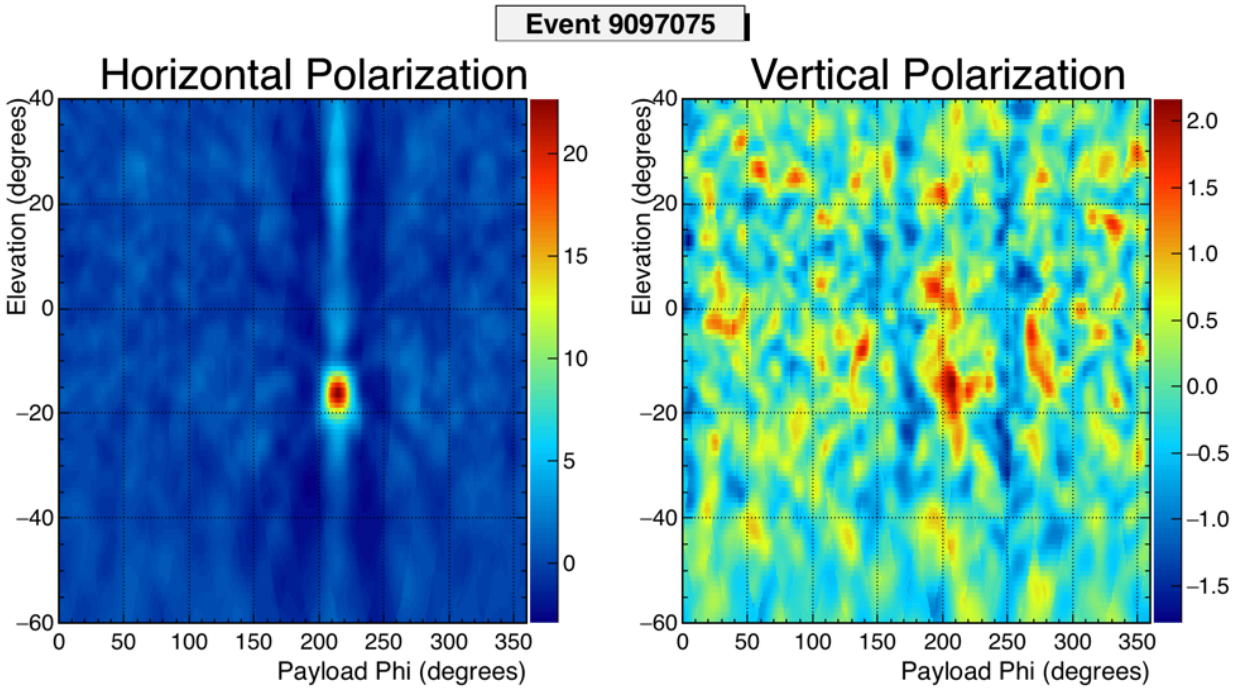

Figure 5.17: Interferometric maps for event 9097075

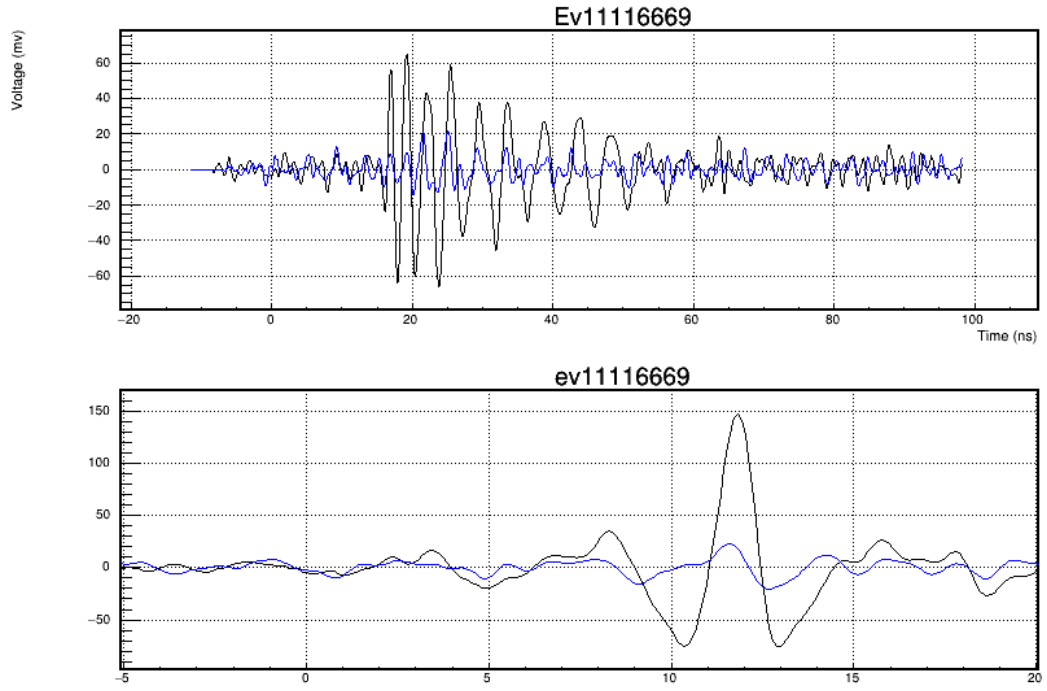

Figure 5.18: Top: Coherently summed waveform. Bottom: De-dispersed waveform

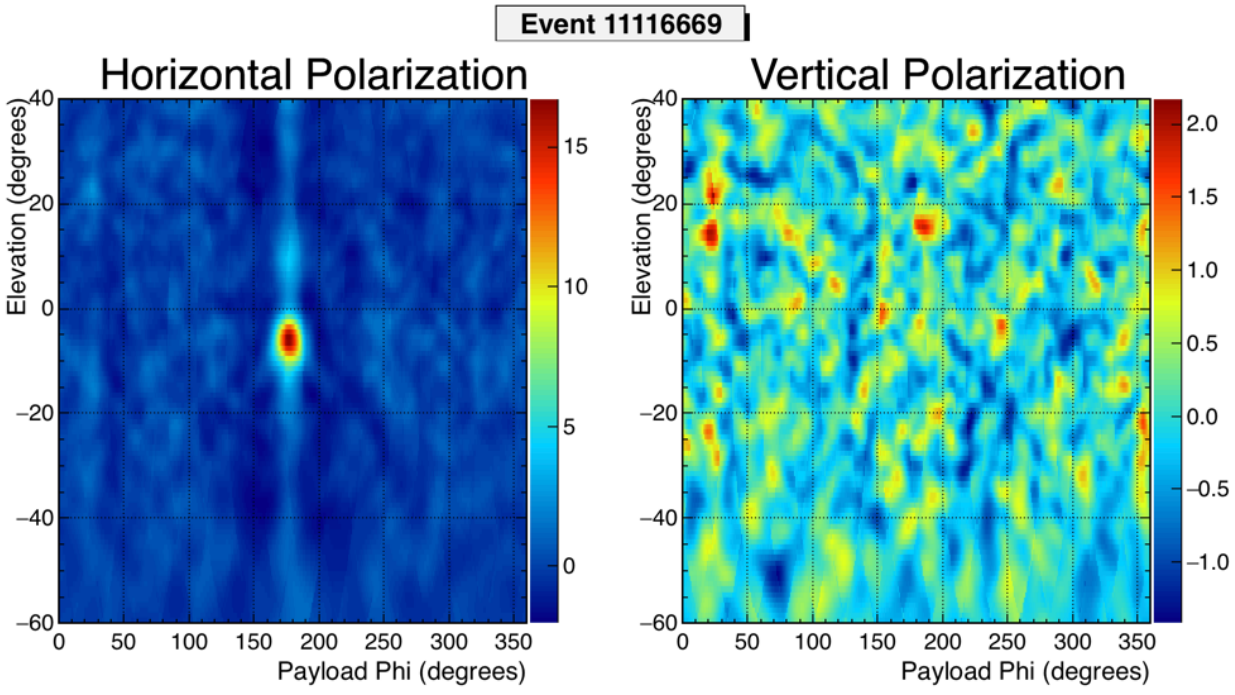

Figure 5.19: Interferometric maps for event 11116669

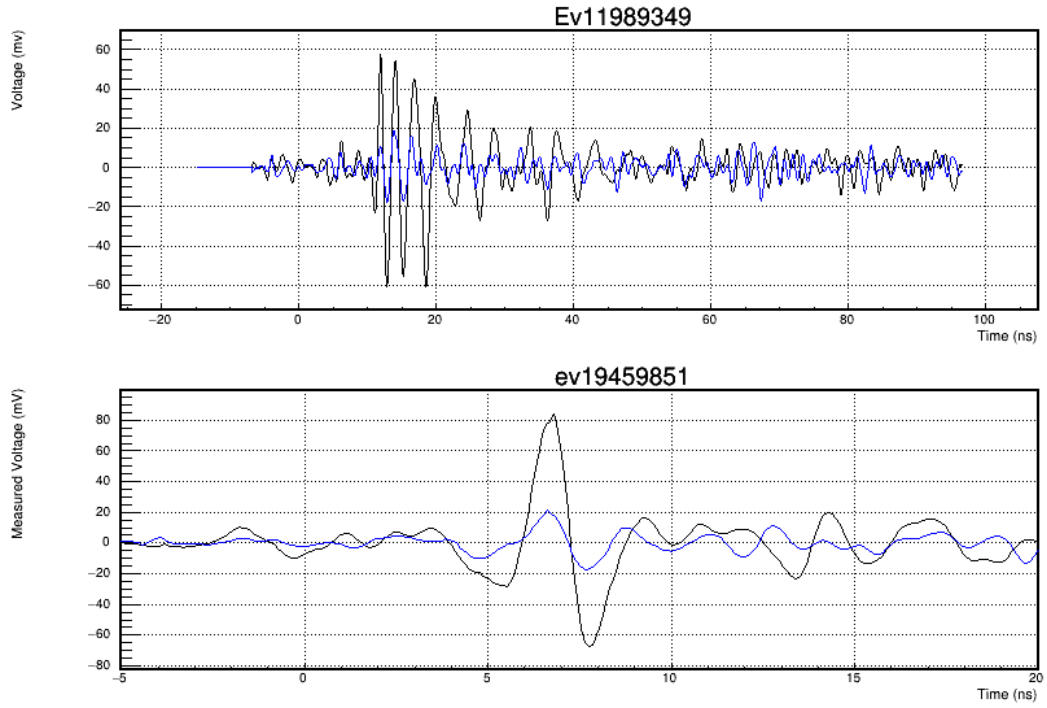

Figure 5.20: Top: Coherently summed waveform. Bottom: De-dispersed waveform

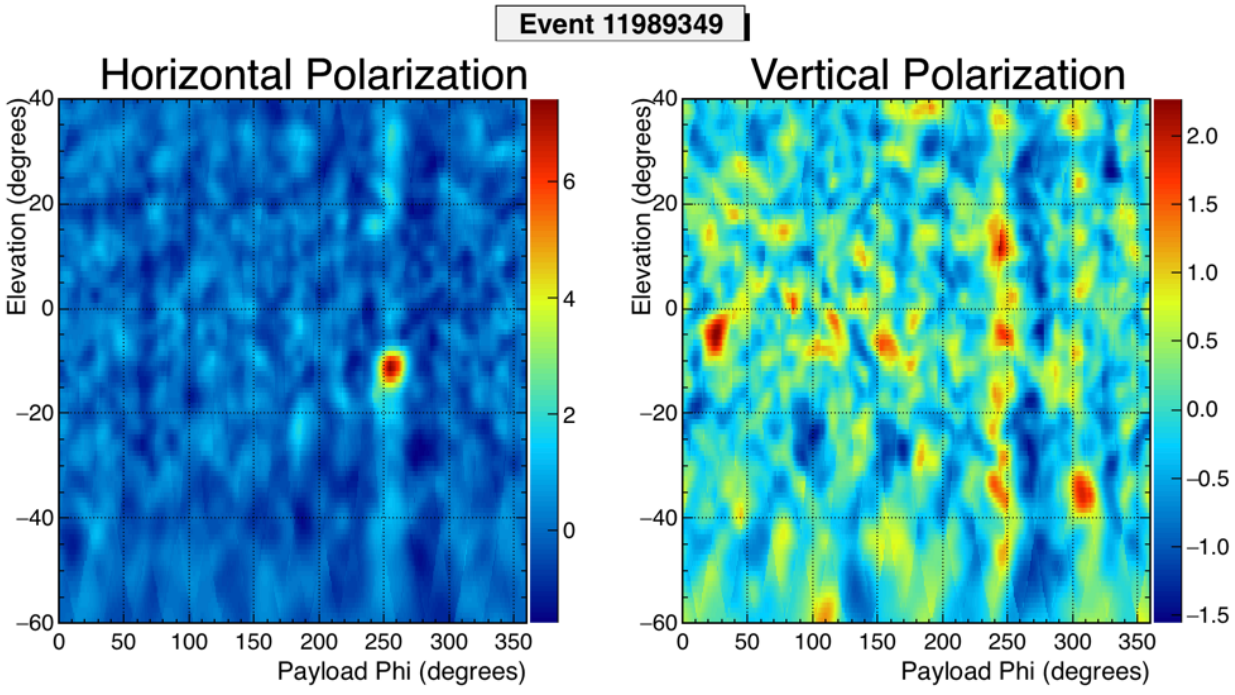

Figure 5.21: Interferometric maps for event 11989349

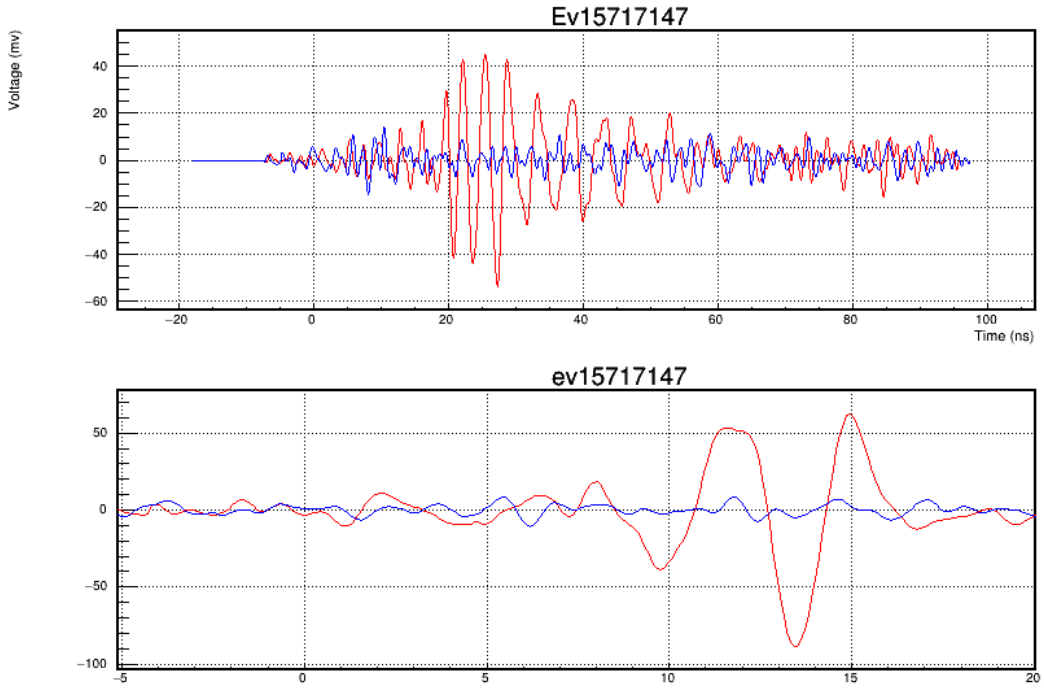

Figure 5.22: Top: Coherently summed waveform. Bottom: De-dispersed waveform

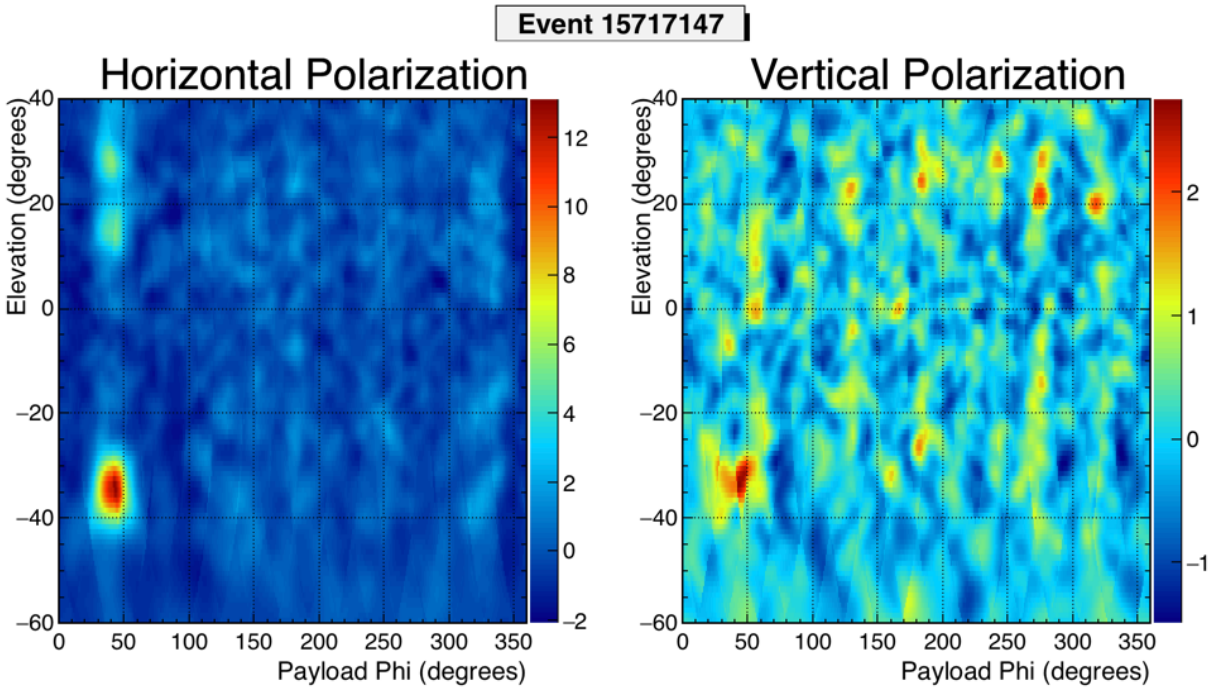

Figure 5.23: Interferometric maps for event 15717147

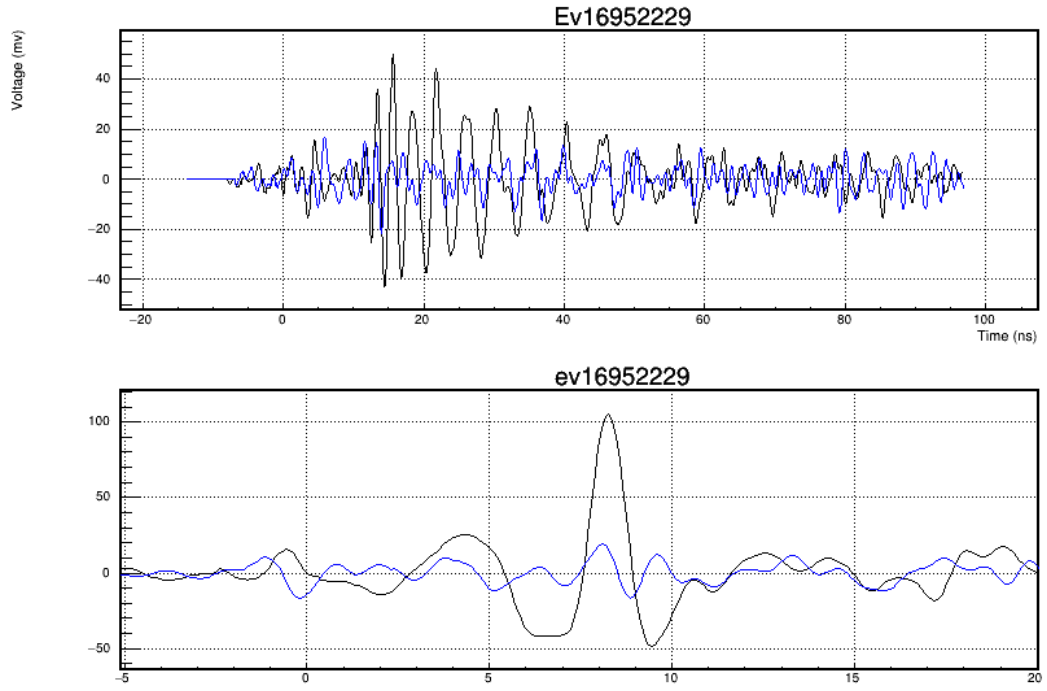

Figure 5.24: Top: Coherently summed waveform. Bottom: De-dispersed waveform

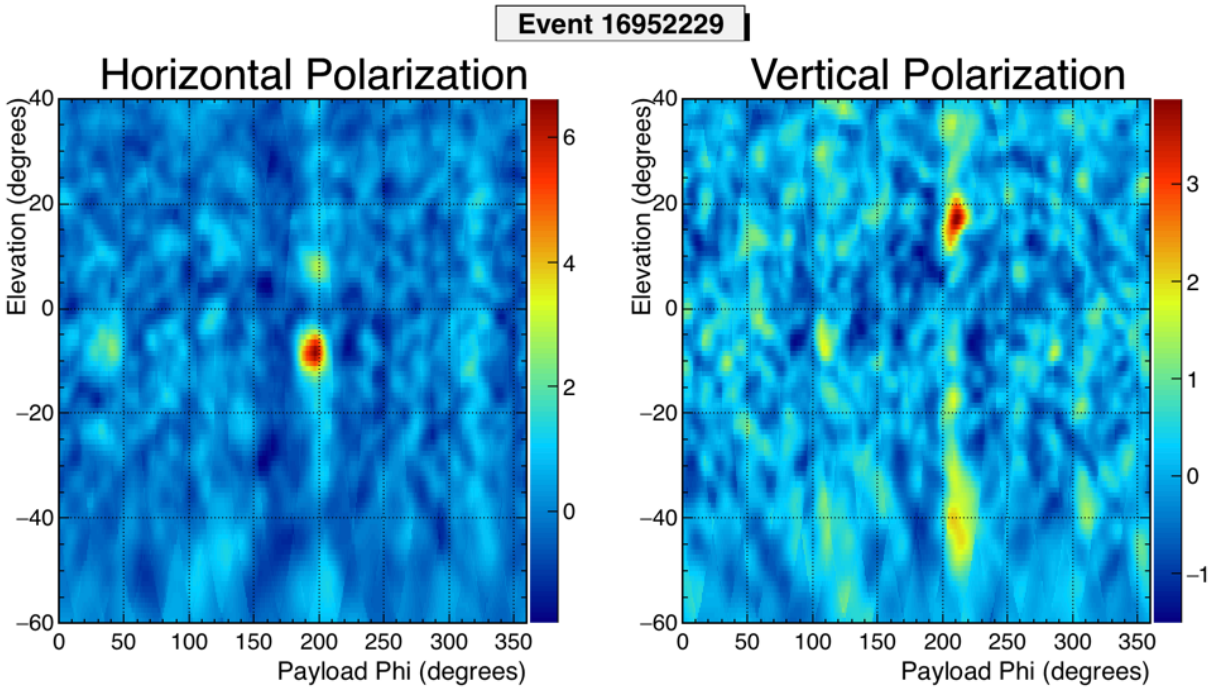

Figure 5.25: Interferometric maps for event 16952229

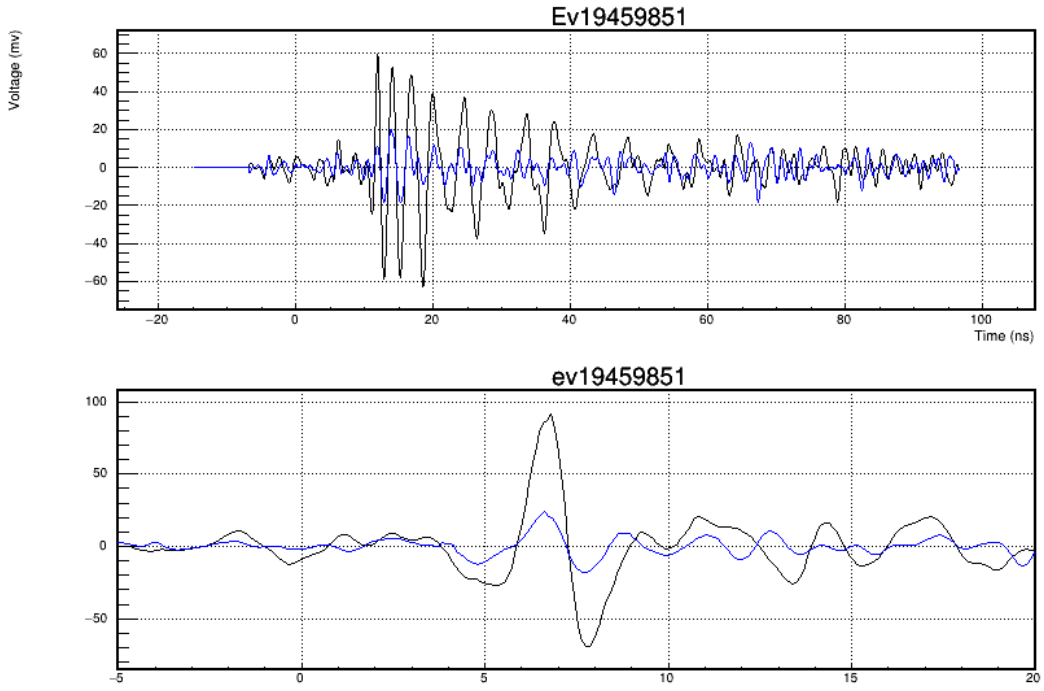

Figure 5.26: Top: Coherently summed waveform. Bottom: De-dispersed waveform

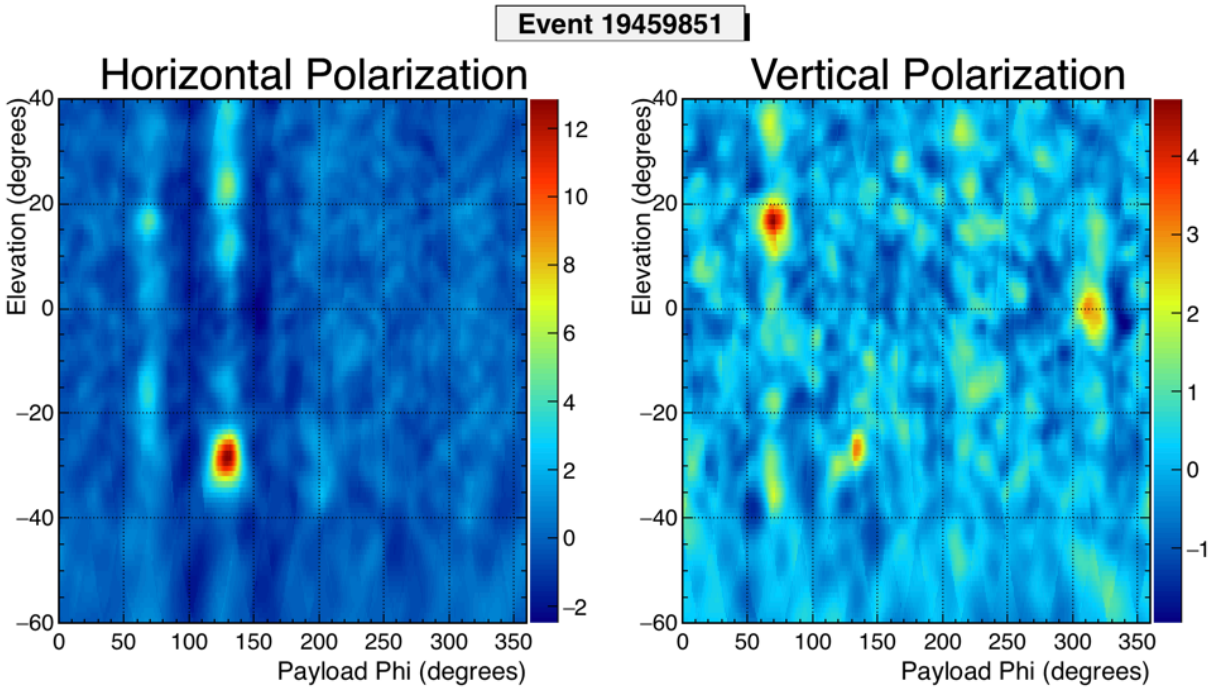

Figure 5.27: Interferometric maps for event 19459851

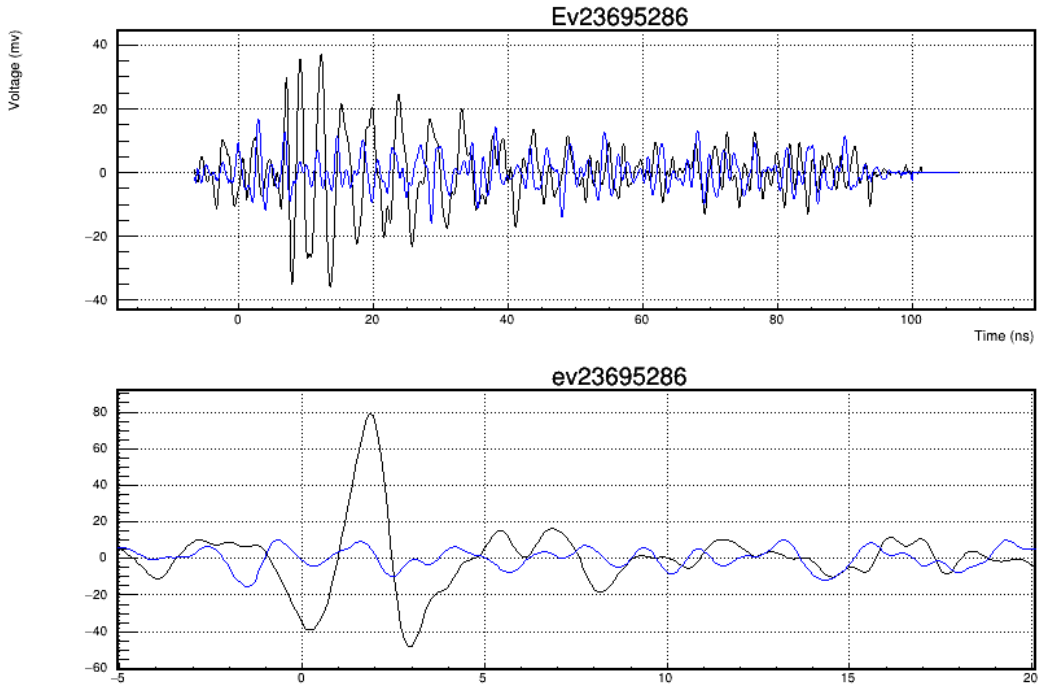

Figure 5.28: Top: Coherently summed waveform. Bottom: De-dispersed waveform

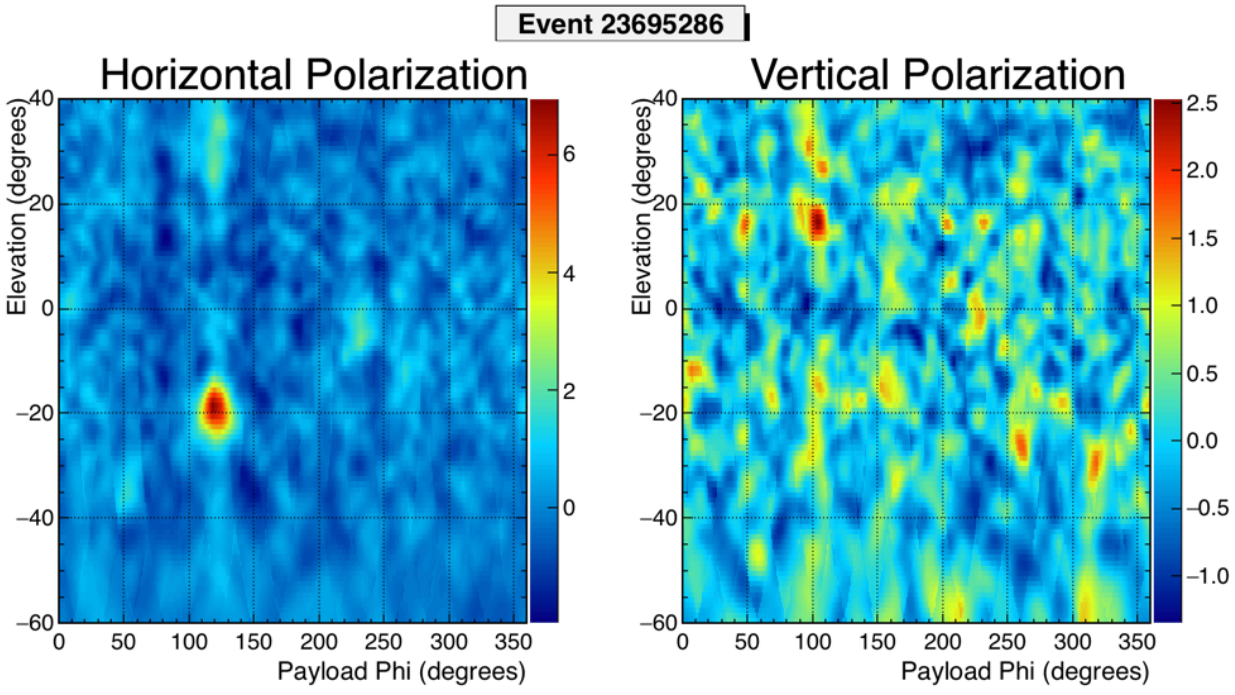

Figure 5.29: Interferometric maps for event 23695286

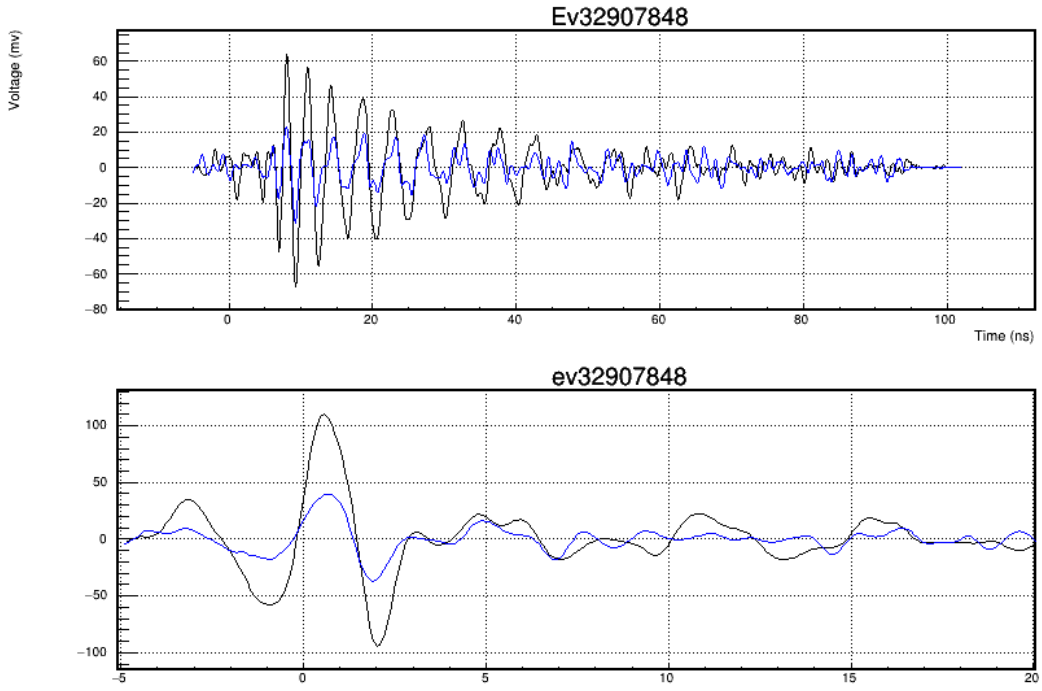

Figure 5.30: Top: Coherently summed waveform. Bottom: De-dispersed waveform

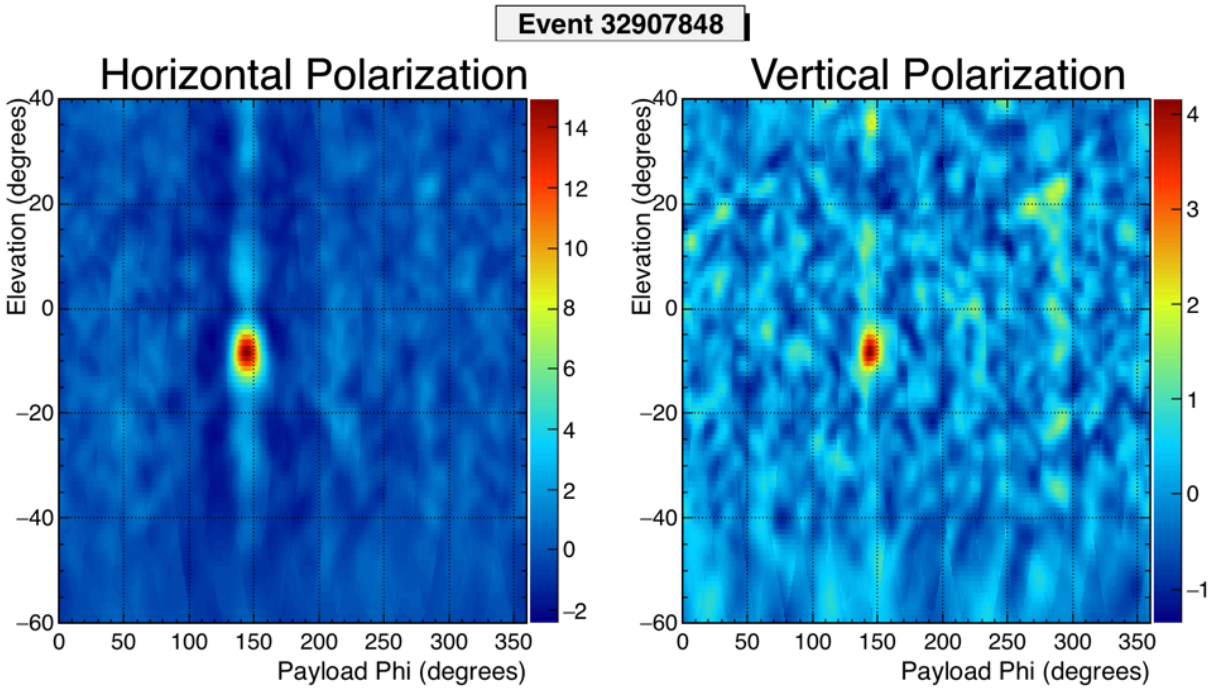

Figure 5.31: Interferometric maps for event 32907848

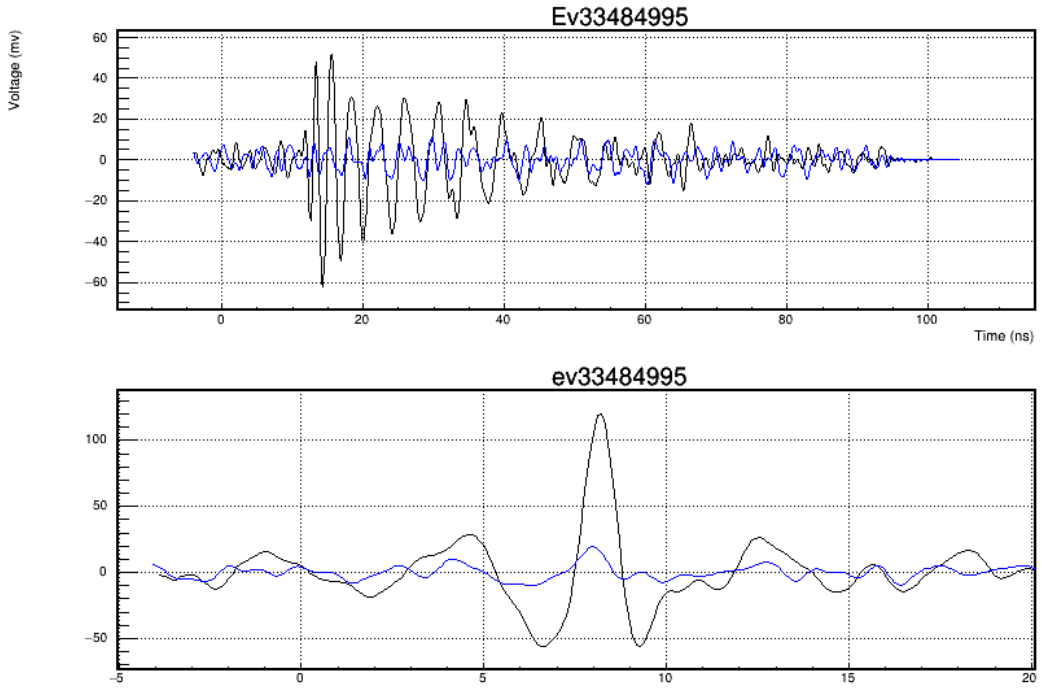

Figure 5.32: Top: Coherently summed waveform. Bottom: De-dispersed waveform

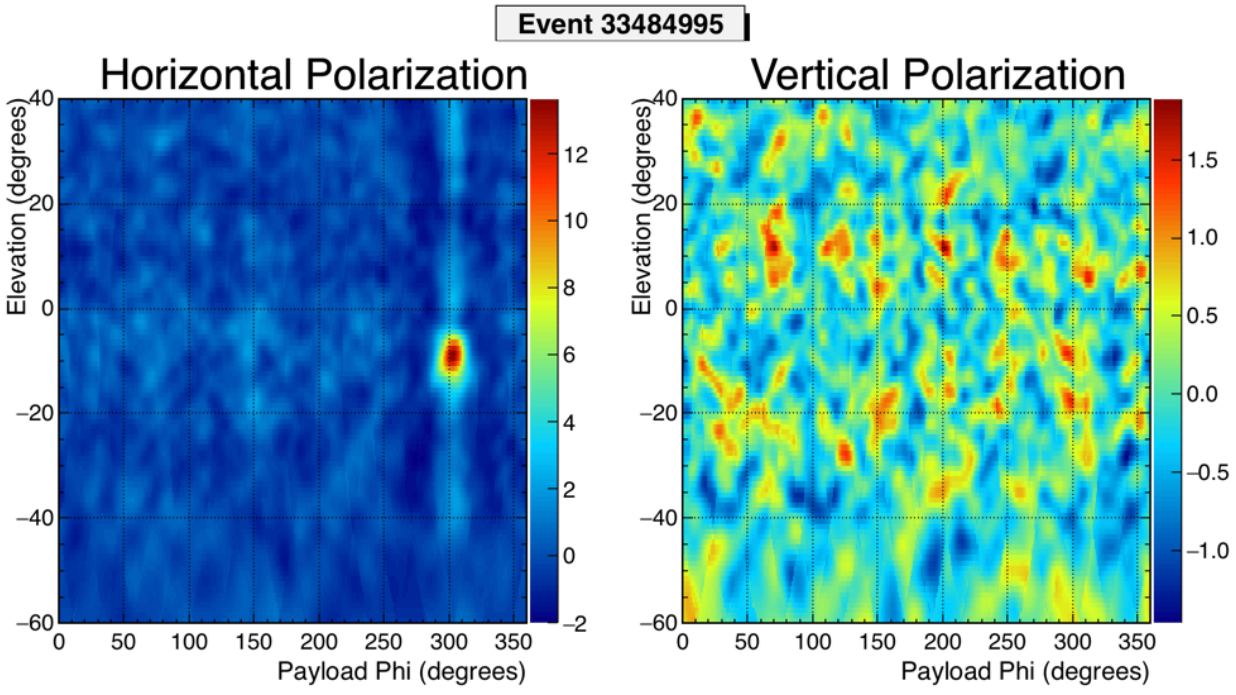

Figure 5.33: Interferometric maps for event 33484995

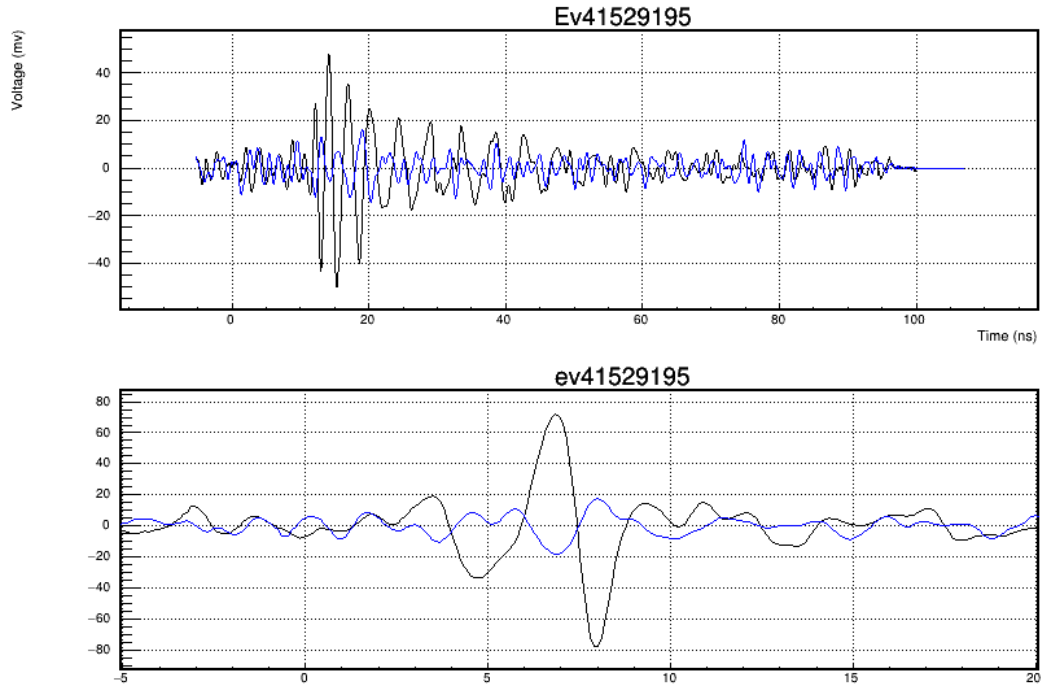

Figure 5.34: Top: Coherently summed waveform. Bottom: De-dispersed waveform

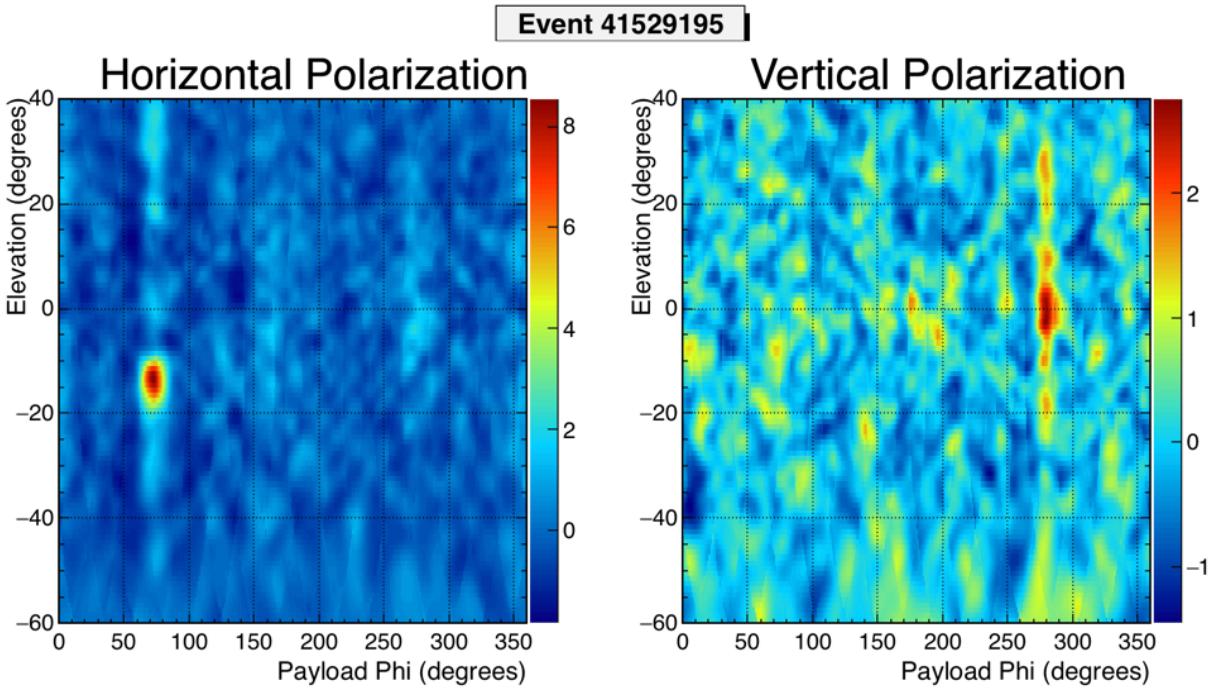

Figure 5.35: Interferometric maps for event 41529195

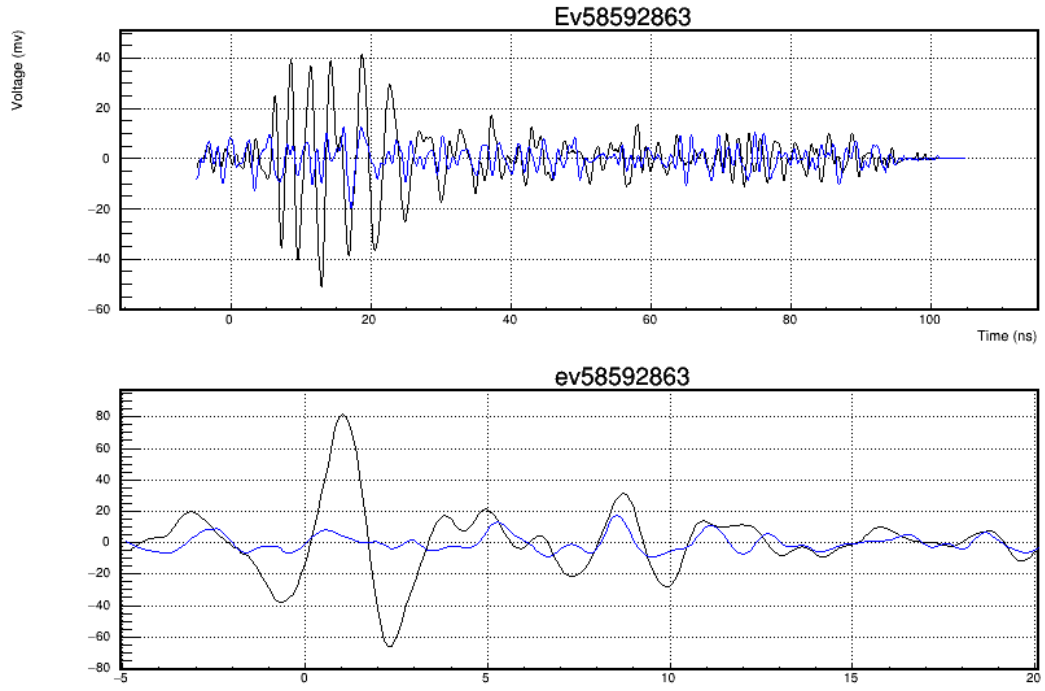

Figure 5.36: Top: Coherently summed waveform. Bottom: De-dispersed waveform

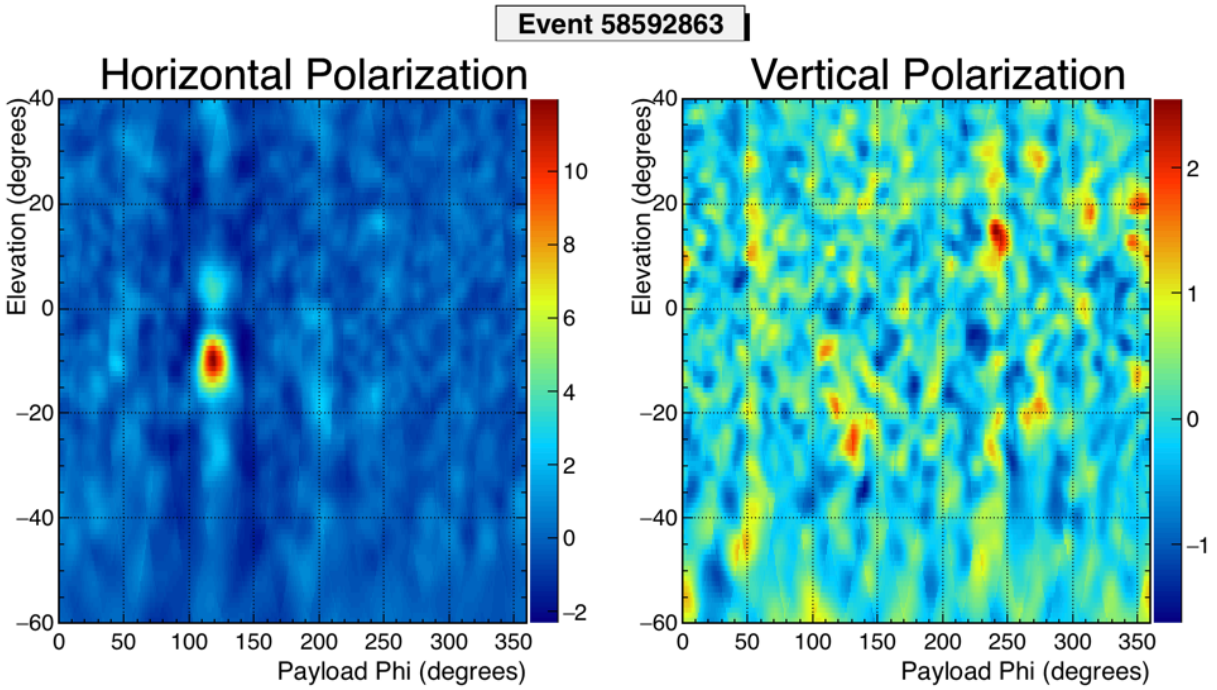

Figure 5.37: Interferometric maps for event 58592863

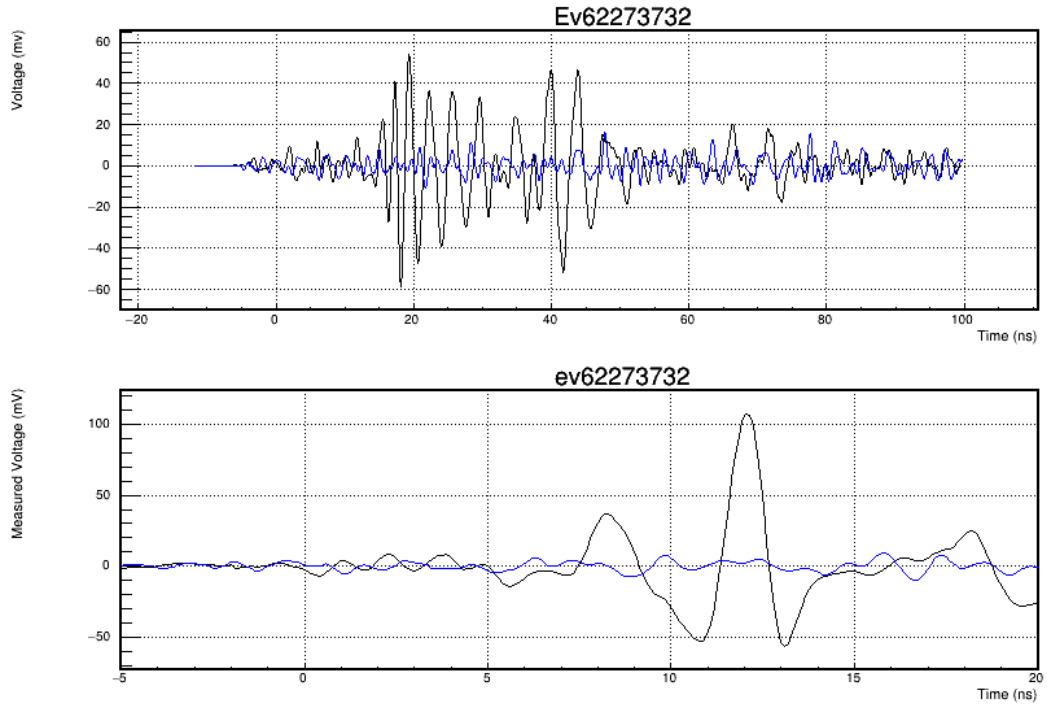

Figure 5.38: Top: Coherently summed waveform. Bottom: De-dispersed waveform

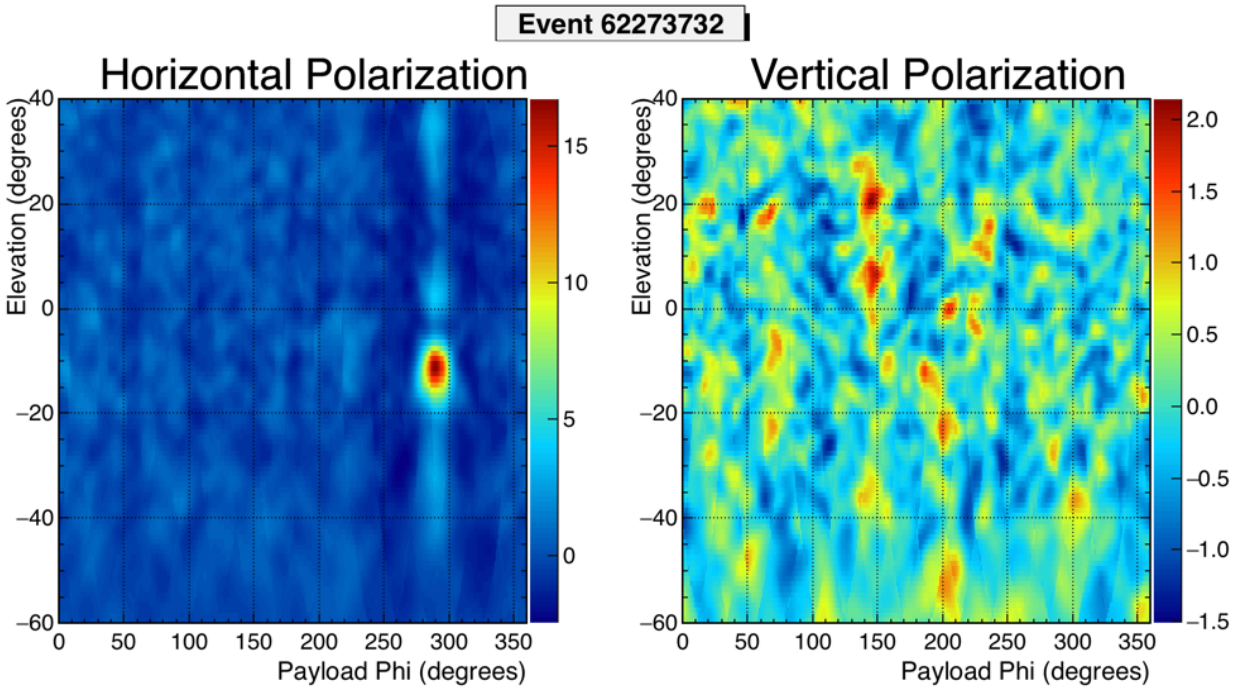

Figure 5.39: Interferometric maps for event 62273732

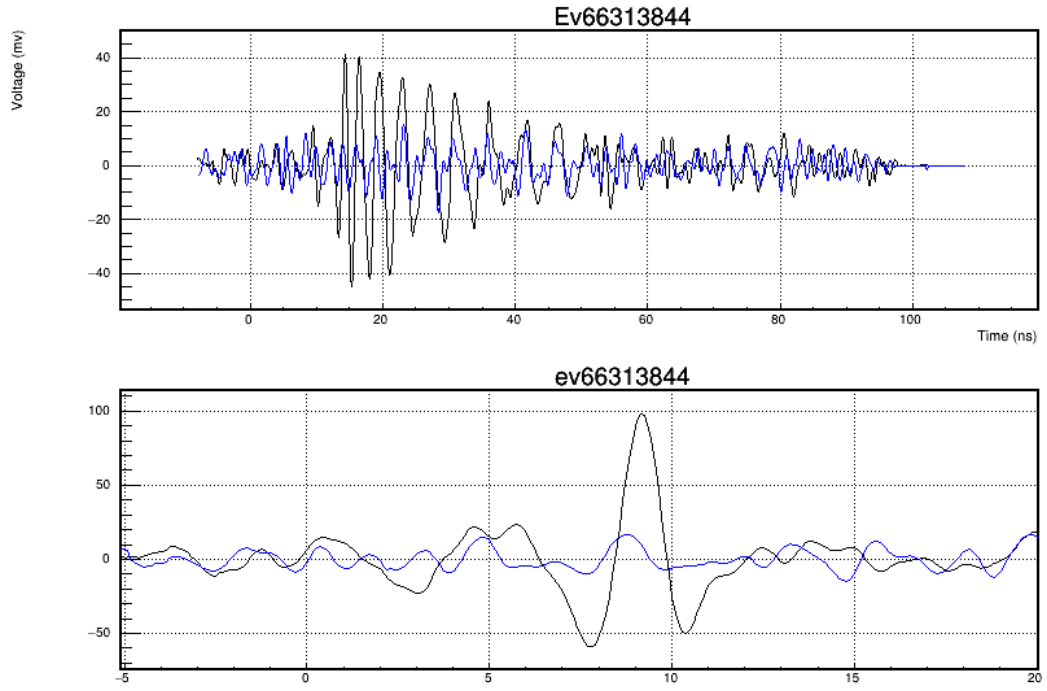

Figure 5.40: Top: Coherently summed waveform. Bottom: De-dispersed waveform

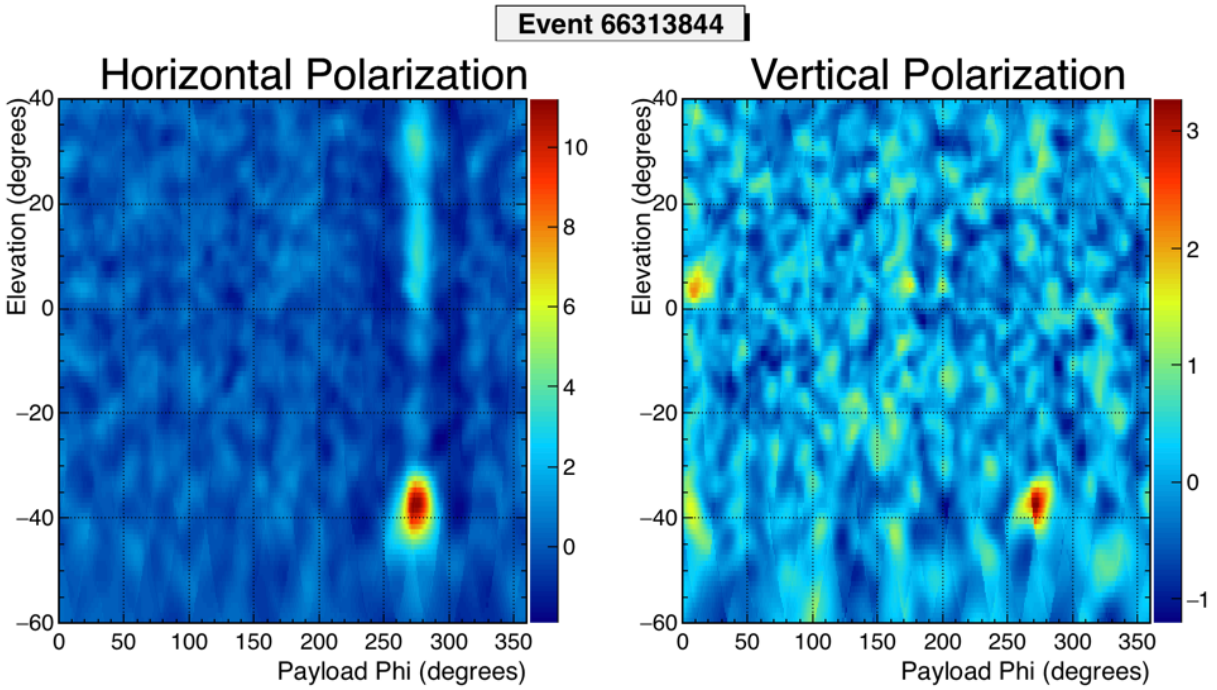

Figure 5.41: Interferometric maps for event 66313844

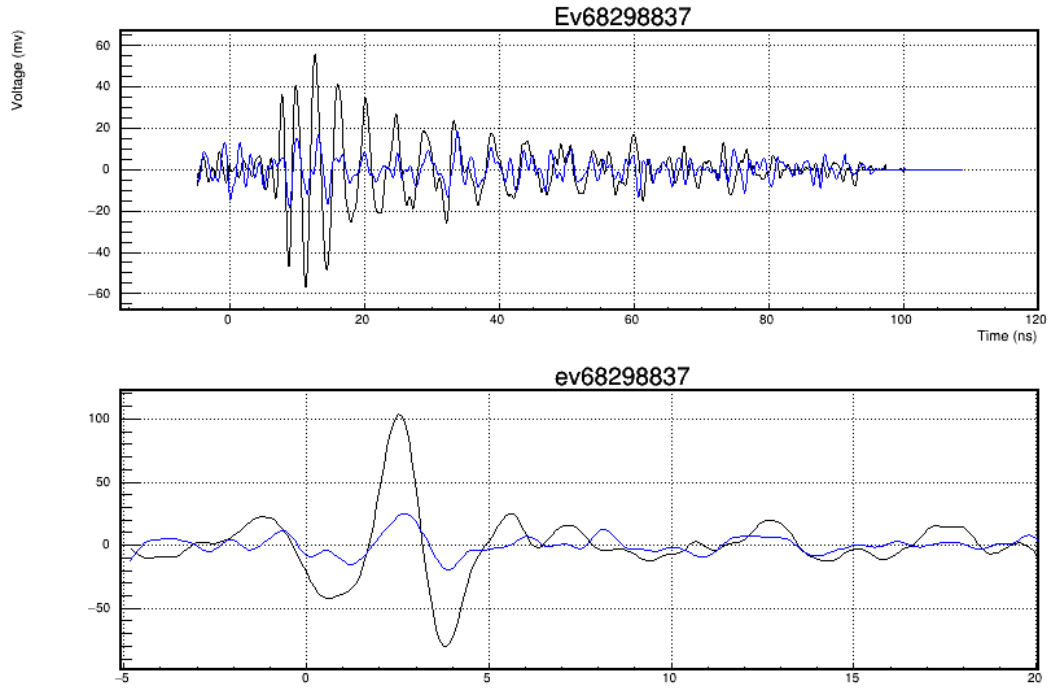

Figure 5.42: Top: Coherently summed waveform. Bottom: De-dispersed waveform

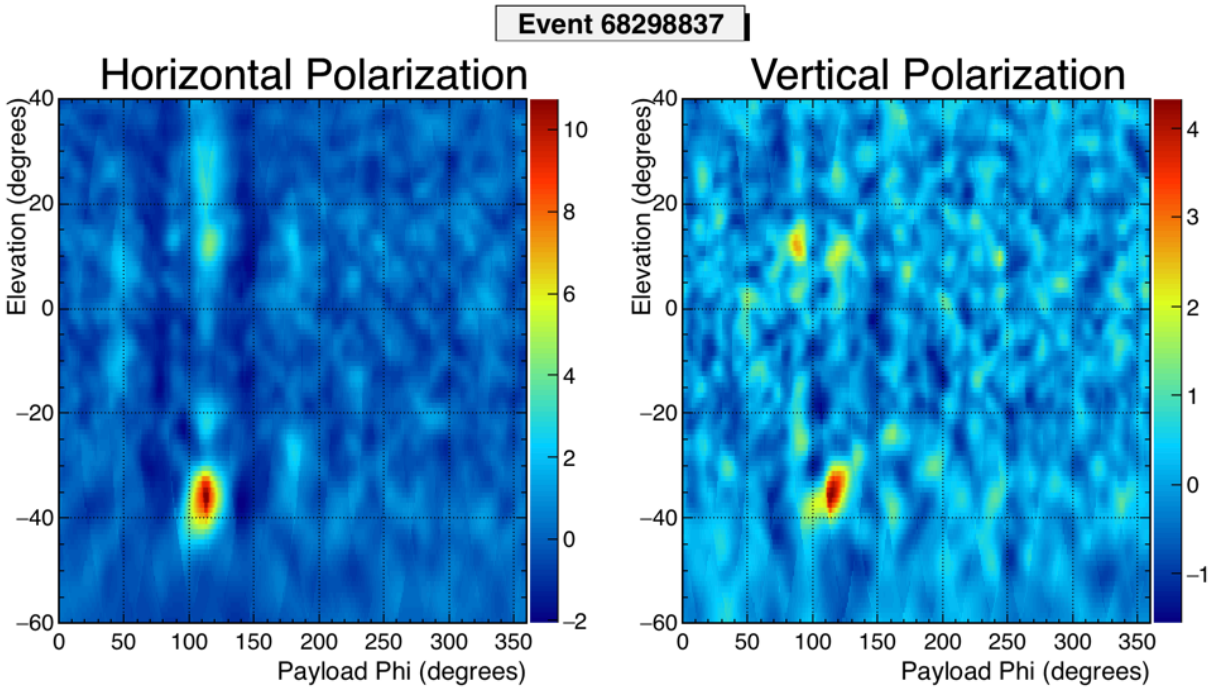

Figure 5.43: Interferometric maps for event 68298837

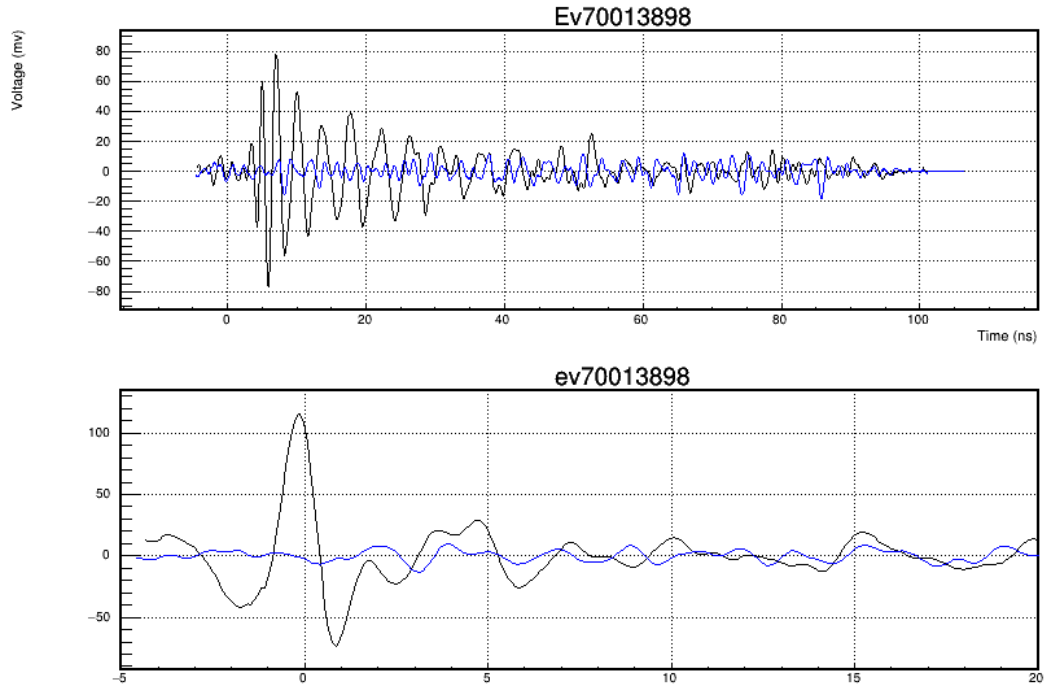

Figure 5.44: Top: Coherently summed waveform. Bottom: De-dispersed waveform

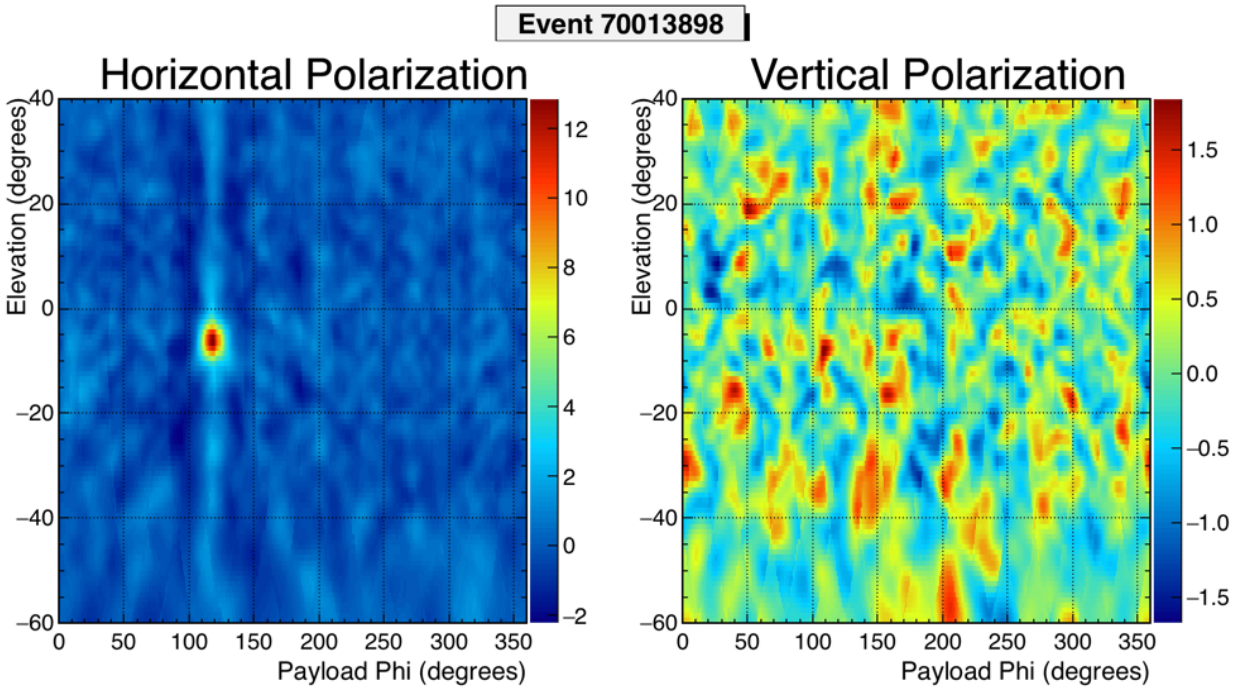

Figure 5.45: Interferometric maps for event 70013898

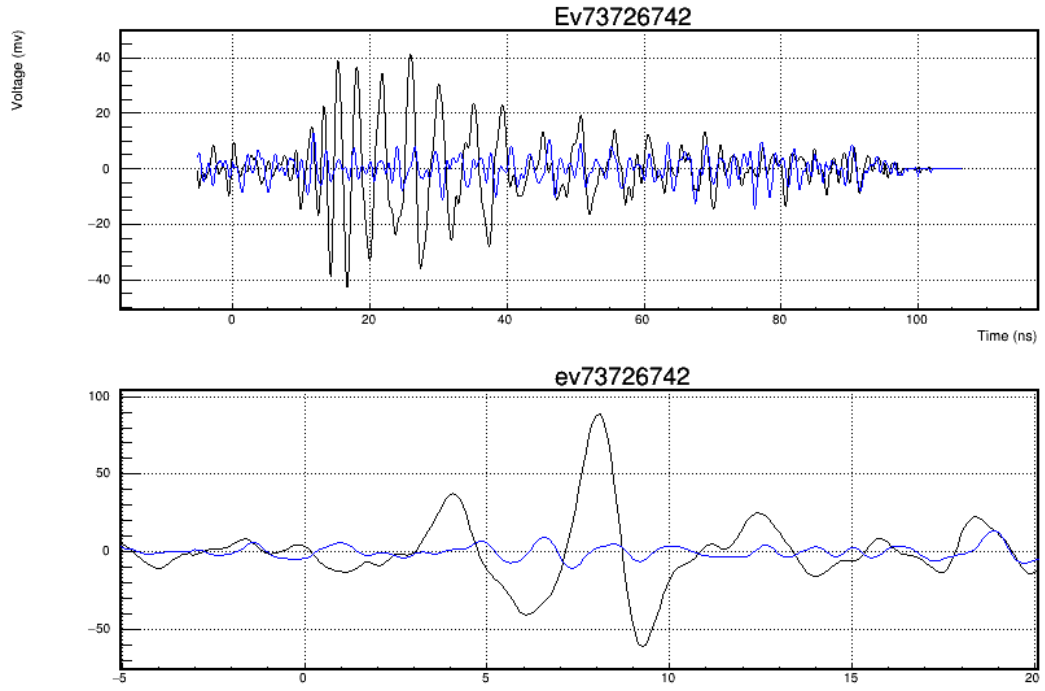

Figure 5.46: Top: Coherently summed waveform. Bottom: De-dispersed waveform

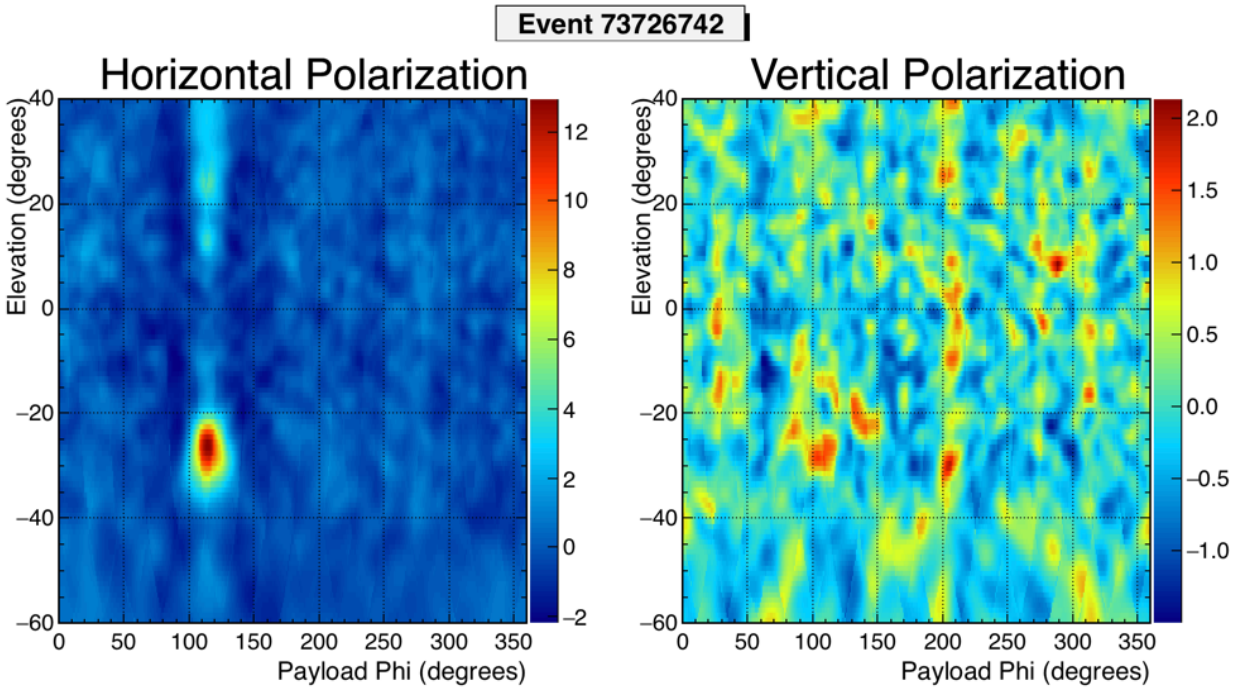

Figure 5.47: Interferometric maps for event 73726742

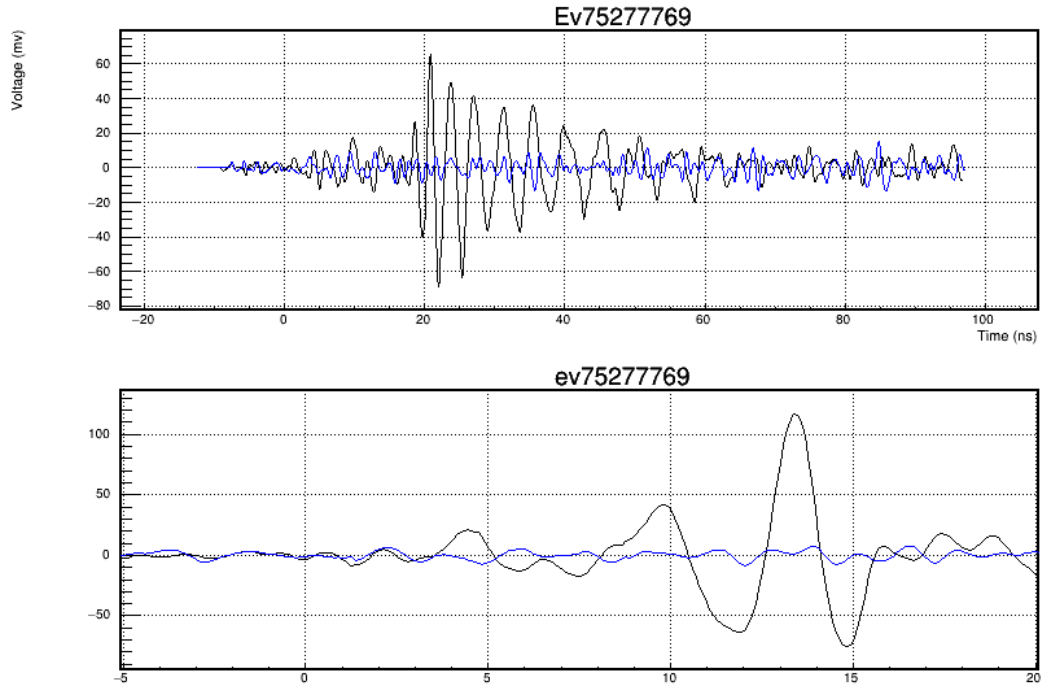

Figure 5.48: Top:: Coherently summed waveform. Bottom: De-dispersed waveform

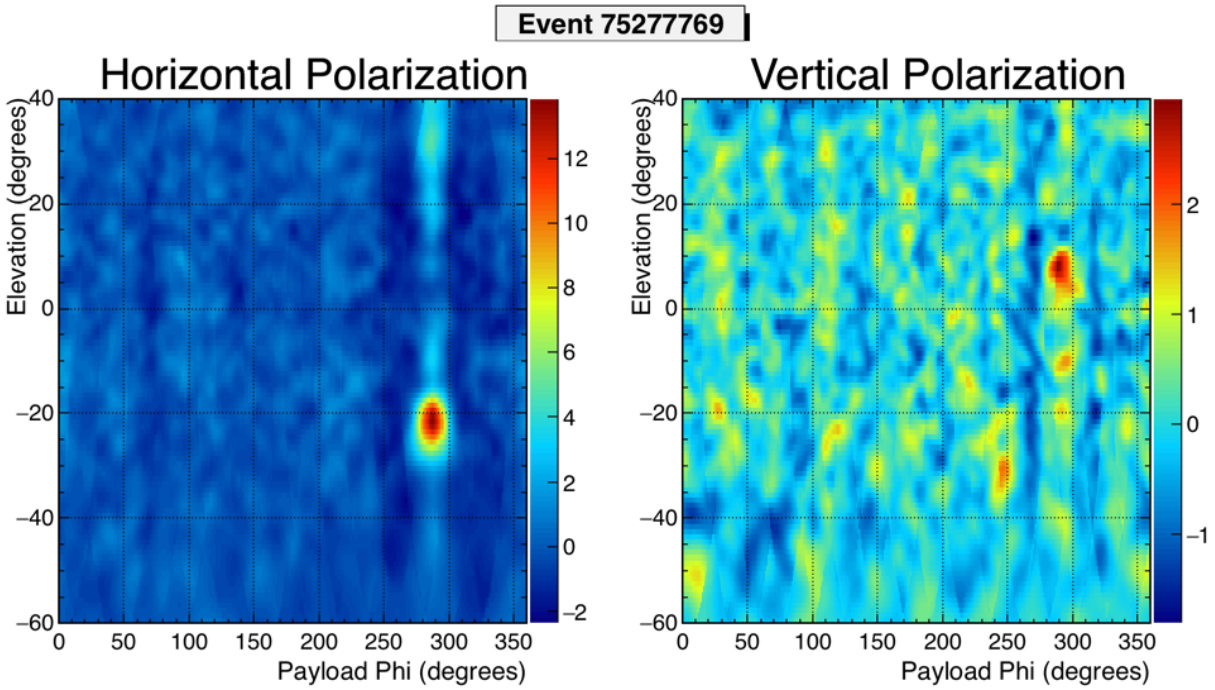

Figure 5.49: Interferometric maps for event 75277769

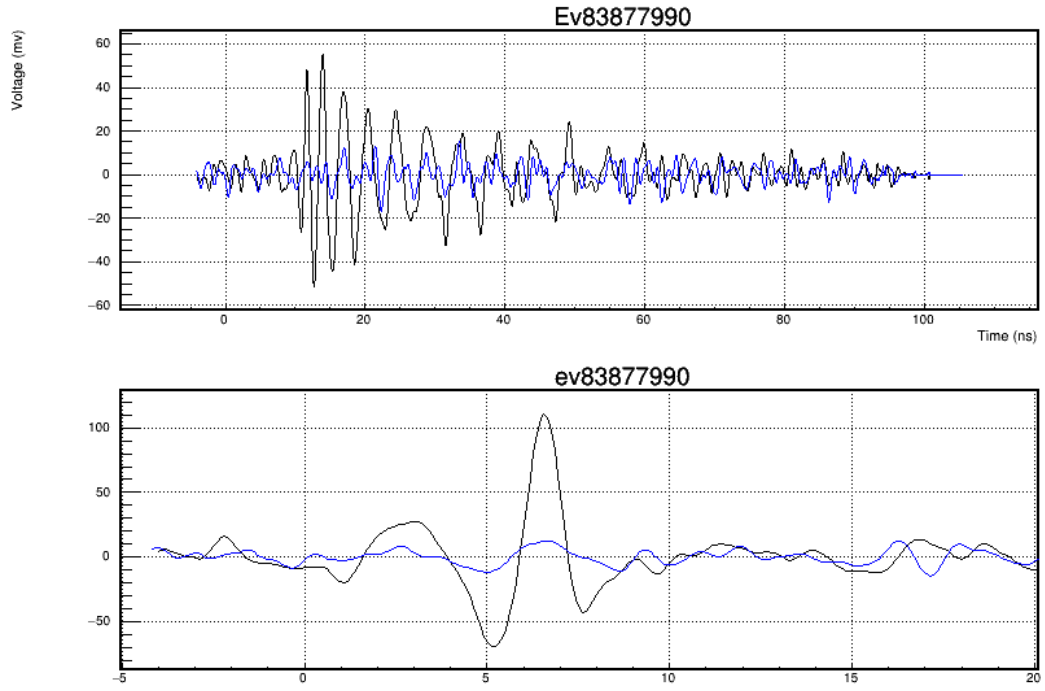

Figure 5.50: Top: Coherently summed waveform. Bottom: De-dispersed waveform

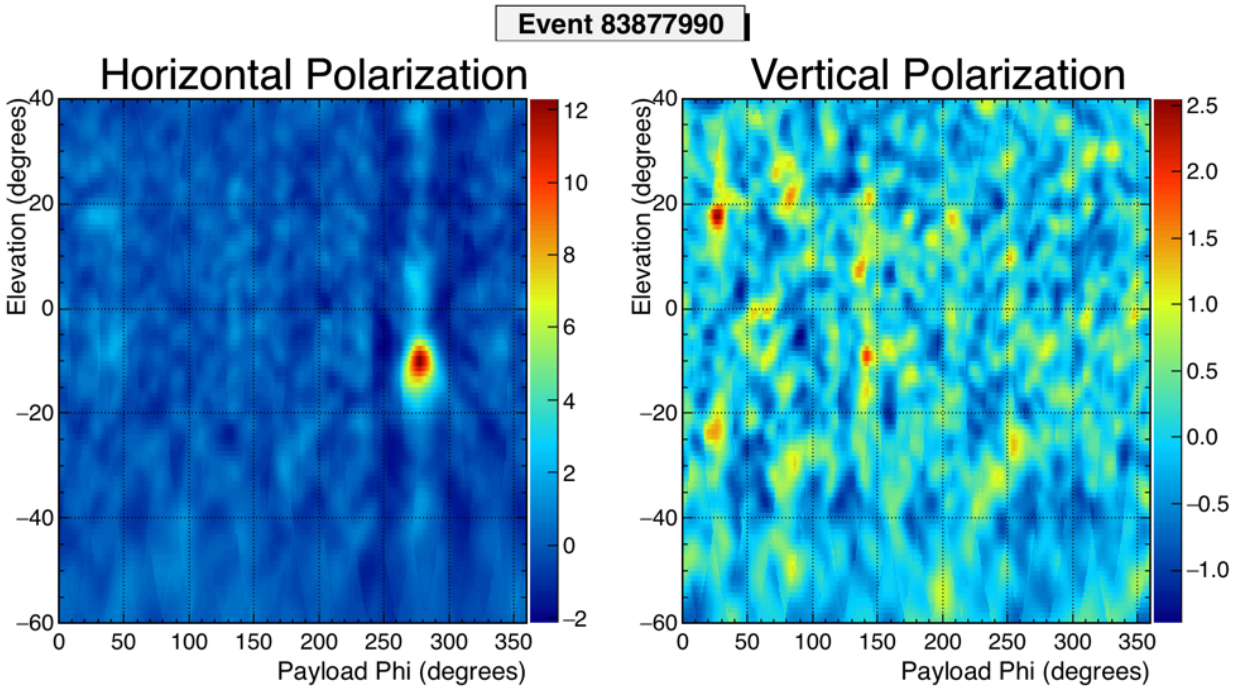

Figure 5.51: Interferometric maps for event 83877990

Event Number	Above/Below Horizon	Coherent Polarity	Deconvolved Polarity
9097075	Below	0	0
11116669	Below	0	0
11989349	Below	0	0
15717147	Below	1	1
16952229	Below	0	0
19459851	Below	0	0
23695286	Below	0	0
27142546	Above	1	1
32907848	Below	0	0
33484995	Below	0	0
39599205	Above	1	1
41529195	Below	0	0
58592863	Below	0	0
62273732	Below	0	0
66313844	Below	0	0
68298837	Below	0	0
70013898	Below	0	0
73726742	Below	0	0
75277769	Below	0	0
83877990	Below	0	0

Figure 5.52: Results of the polarity unblinding. Events are classified into above horizon directly observed events, or below horizon reflected events. The value for the polarity estimate is a binary value representing the sign of the absolute peak of the highest valued simulated CR template correlation. 15 of the reflected events exhibit the same polarity, which is opposite that observed from the 2 directly observed events. 1 below horizon event is observed to have an inverted polarity.

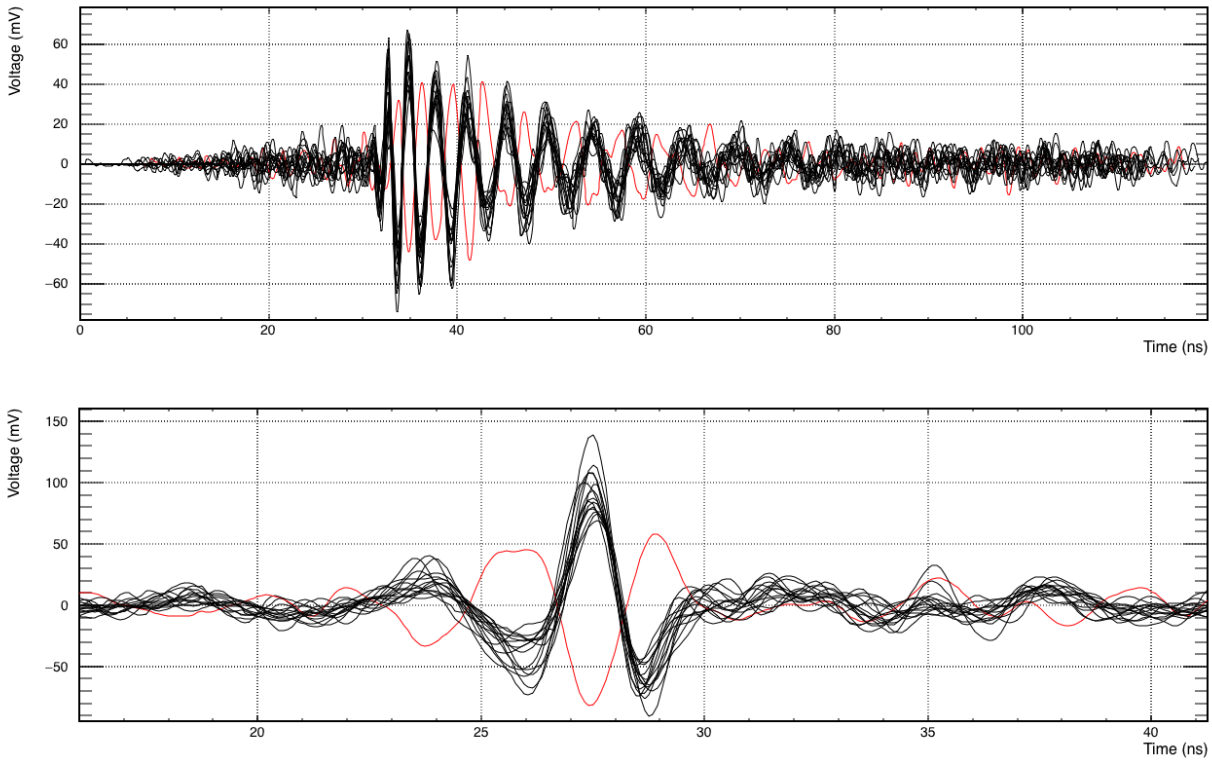

Figure 5.53: 18 below horizon CR candidate events overlayed. These waveforms show the true measured polarity without blinding. They have been aligned to their peak correlation, except the flipped polarity event, which is aligned to the peak anti-correlation. The flipped waveform is highlighted in red. Top: Coherently summed waveform. Bottom: All Pass Filter De-dispersed waveform.

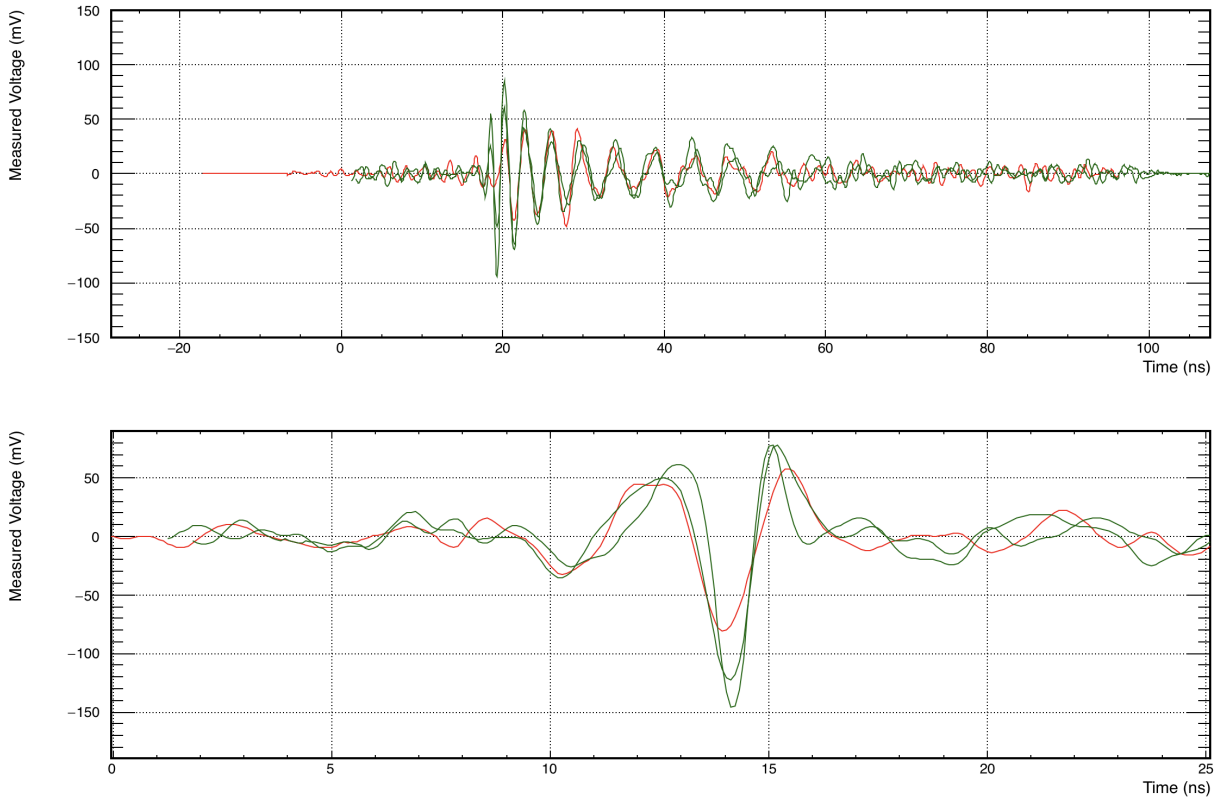

Figure 5.54: Both above horizon CR candidate events overlaid with the one flipped polarity below horizon event. These waveforms show the true measured polarity without blinding. They have been aligned to their peak correlation. The flipped waveform is highlighted in red. Top: Coherently summed waveform. Bottom: All Pass Filter De-dispersed waveform.

5.9 Summary of Results

In summary, there were 20 CR candidate events detected during the ANITA-III flight with an estimate background count of 4.4×10^{-3}. The distribution of events into categories can be seen in Figure 5.55.

5.9.1 Cosmic Ray Energy Estimate

Determining the energy of the candidate CRs is done via a method developed for ANITA-I [7]. This method utilizes measurements of the measured electric field peak amplitude and spectral slope to establish the energy of the incident particle by comparing them to the same parameters generated from simulated air showers. These parameters are also dependent on the incident angle of the particle, which will effect the Fresnel reflection coefficients and radiation propagation losses of the impulse. Amplitude will be represented by the variable A and has units of the square root of spectral intensity ($\text{pW}^{0.5}$ m^{-1} $\text{MHz}^{-0.5}$). Spectral slope will be represented by a γ and has units of Hz $^{-1}$.

Determining the values of the amplitude (A) and spectral slope (γ) is done by fitting an exponential to the frequency spectrum of the coherently summed candidate waveforms between 300 MHz and 1000 MHz to match the work done in [7]. The fit function is shown in Equation 5.3.

$$A_f(f) = Ae^{\gamma(f-300MHz)} \tag{5.3}$$

Note that this describes the linear amplitude spectrum of the signal instead of the logarithmic power spectrum as is used in other locations of this thesis. To compare this to simulation, the measured and recorded voltage at the digitizer must have the impulse response of the signal chain and the antenna height deconvolved to establish the corresponding electric field. An example waveform, the logarithmic power spectrum, and the linear amplitude spectrum and subsequent fit for an ANITA-III candidate CR is shown in Figure 5.56.

Results of this fit for all 20 candidate events can be seen in Table 5.6.

To first order, higher energy CRs will have a linearly proportional higher amplitude electromagnetic impulse (Figure 5.57). The radiation, beamed due to coherence within the shower core, will have a maximum amplitude at a critical angle away from the principle shower axis, and will fall off away from that critical angle (Figure 5.58). Additionally, the conditions for coherence are stricter as wavelength decreases, and as a result the spectral slope of the electric field will steepen as the observation point moves further from the critical angle (Figure 5.59). All of these plots show expectations of the electromagnetic field at the ANITA payload assuming a reflection off the ice and a 70° incidence angle. It takes into account estimates of Fresnel reflection coefficients, Earth curvature, surface roughness, and propagation losses. Combining these relationships, a two-dimensional parameter space that uniquely relates spectral slope and pulse amplitude to particle energy can be

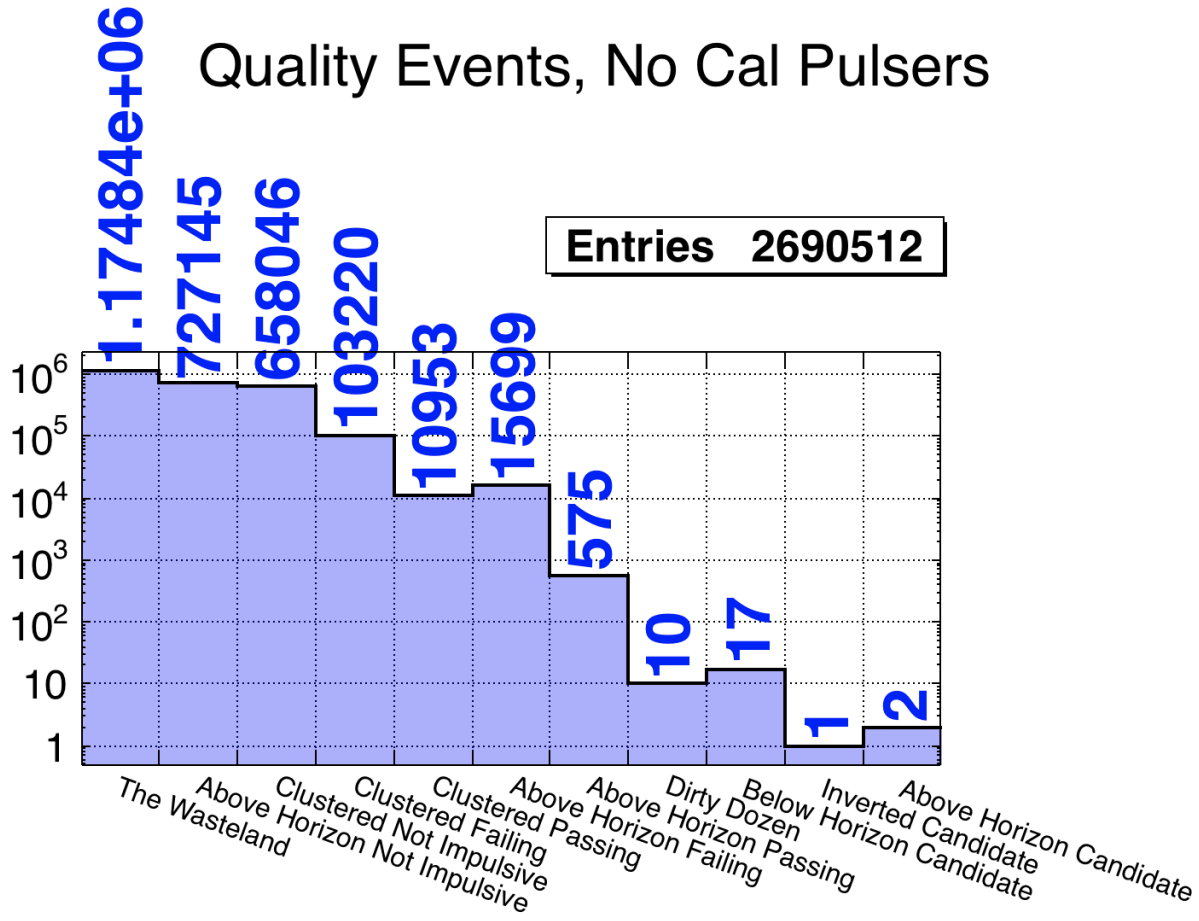

Figure 5.55: Final results and quantities of quality events that fall within each analysis category. "Dirty Dozen" events refer to impulsive isolated signals. "The Wasteland" refers to events that do not pass impulsivity cuts and do not cluster with any other impulsive events. Wasteland events quantitatively represent the amount of flight in which signal detection was not excluded by anthropogenic sources, but no signal was seen. Events are categorized into above/below horizon, then by whether they are geo-associated with any pseudo-base impulsive signal sources, then by their strength. Candidates are divided into above/below horizon events, and the inverted event earns its own category. Calibration pulsers are excluded prior to categorization to prevent artificial inflation of clustered event regions, however events that associate with major bases are included.

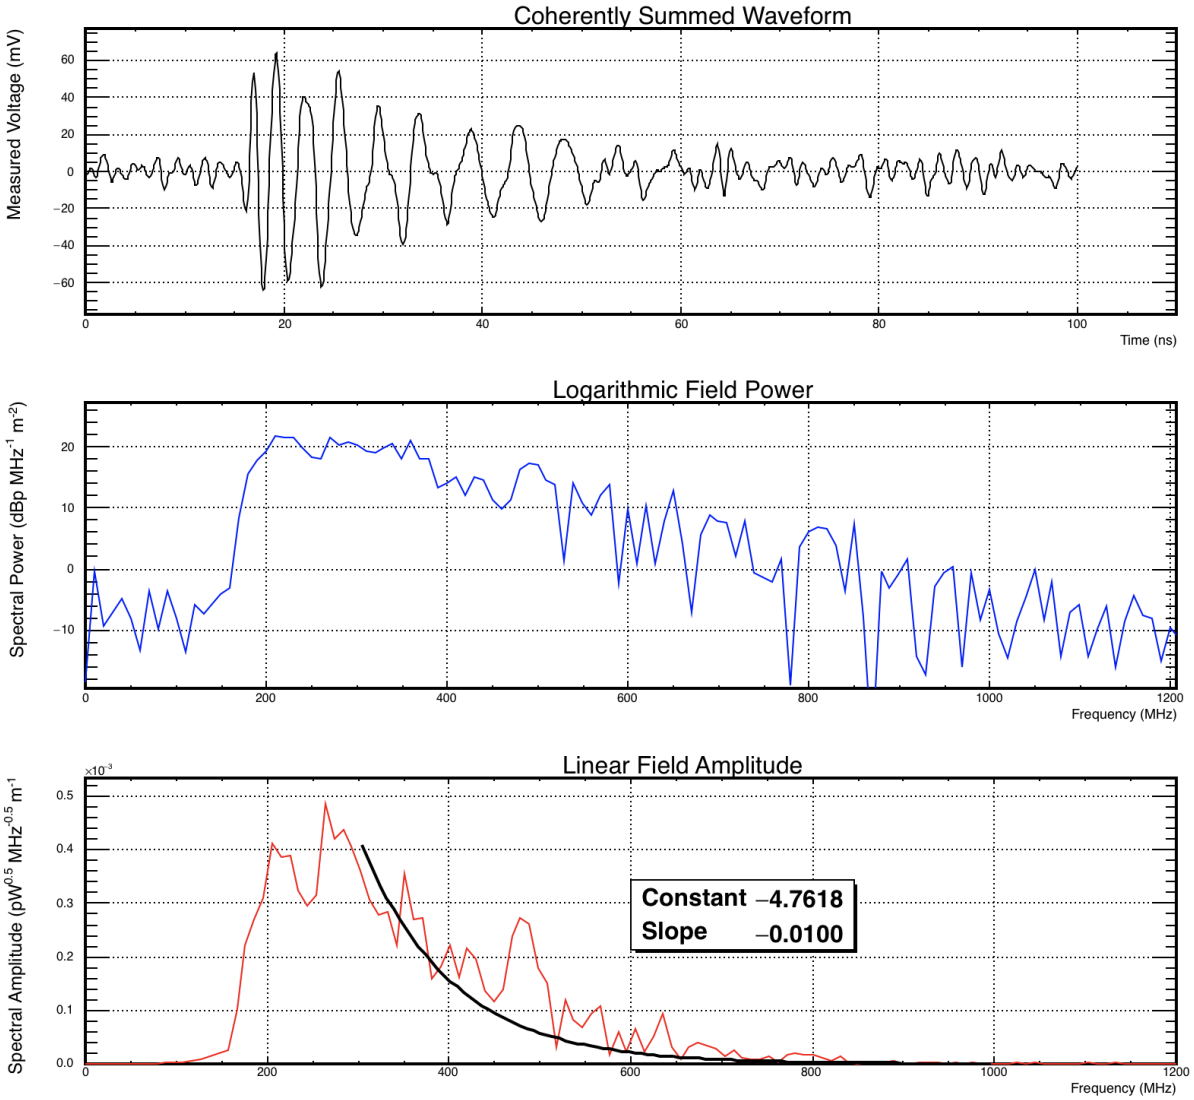

Figure 5.56: An example of the time domain waveform, electric field power spectrum, and amplitude spectrum with fit, for event 11116669. Fitting the amplitude spectrum with an exponential is done for all events to determine values to compare with simulation for the purpose of estimating the incident particle energy.

Event Number	Incident Angle degrees	Amplitude (A) $pW^{0.5}$ m^{-1} $MHz^{-0.5}$	Slope (γ) Hz^{-1}
9097075	73.13	0.68	-5.12×10^{-9}
11116669	83.57	0.46	-5.73×10^{-9}
11989349	78.18	0.27	-4.50×10^{-9}
15717147*	55.03	0.39	-8.68×10^{-9}
16952229	81.26	0.24	-4.14×10^{-9}
19459851	61.17	0.37	-4.68×10^{-9}
23695286	70.98	0.23	-4.84×10^{-9}
27142546	84.45	0.39	-5.25×10^{-9}
32907848	81.74	0.40	-6.04×10^{-9}
33484995	80.51	0.33	-4.58×10^{-9}
39599205	86.39	0.37	-3.70×10^{-9}
41529195	75.45	0.26	-4.89×10^{-9}
58592863	79.78	0.26	-5.06×10^{-9}
62273732	78.37	0.35	-4.97×10^{-9}
66313844	52.64	0.29	-5.03×10^{-9}
68298837	53.10	0.33	-6.00×10^{-9}
70013898	83.36	0.33	-3.91×10^{-9}
73726742	63.63	0.24	-4.65×10^{-9}
75277769	68.31	0.40	-5.99×10^{-9}
83877990	79.79	0.33	-4.82×10^{-9}

Table 5.6: Fits for the amplitude and spectral slope for the 20 CR candidates uncovered by this analysis. These values can be used, in conjunction with simulations of the induced electric field, to estimate their energies. Note that event 15717157 points below the horizon, but is viewed directly.

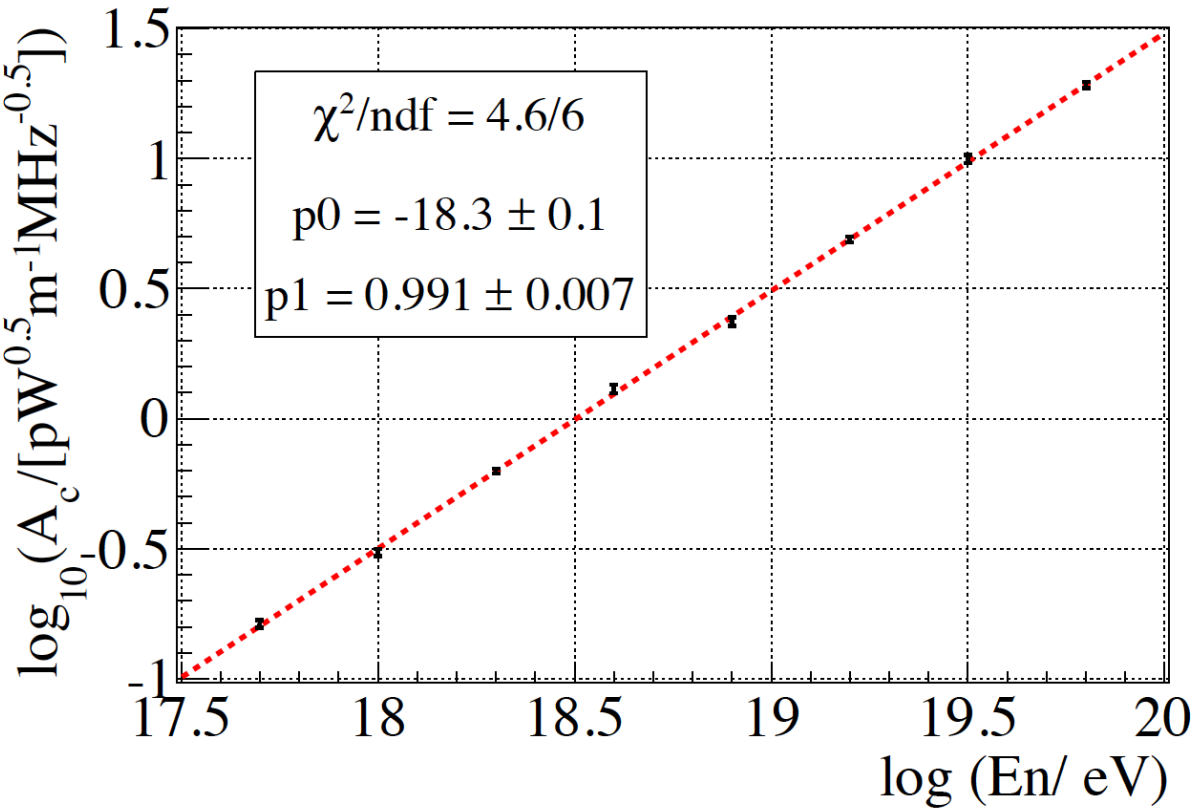

Figure 5.57: Relationship between the peak amplitude (A) of a simulated CR induced electromagnetic field at the ANITA payload and the energy of the incident CR [7].

seen in Figure 5.60 [7].

5.9.2 Comparison between ANITA-I and ANITA-III

Using the amplitude and spectral slopes for the events to estimate an energy requires simulations for each event, work which is beyond the scope of this thesis. However some qualitative conclusions may be drawn about the flight by comparing the fit values with those from ANITA-I. The values for ANITA-I can be found in [7]. In ANITA-I, there were 14 detected CR candidates, while in ANITA-III there are 19 events that can be classified as CRs, not including the one inverted event. The mean of the amplitudes for those events, $\bar{A_I}$, was 0.75. For the events covered in this thesis, the mean is a factor of two lower, $\bar{A_{III}} = 0.345$, which would correspond to a factor of $\log_{10}(2)$ reduction in particle energy. Additionally, the spectral slopes have different means when comparing the two flights, $\bar{\gamma_I} = -2.48$ and $\bar{\gamma_{III}} = -5.12$, which would correspond to an additional factor of two reduction in energy. Meanwhile, the average measured elevation angle between the flights is comparable, $\bar{\theta_I} = 73.1°$ and $\bar{\theta_{III}} = 73.5°$. This suggests that the observed energies of the cosmic

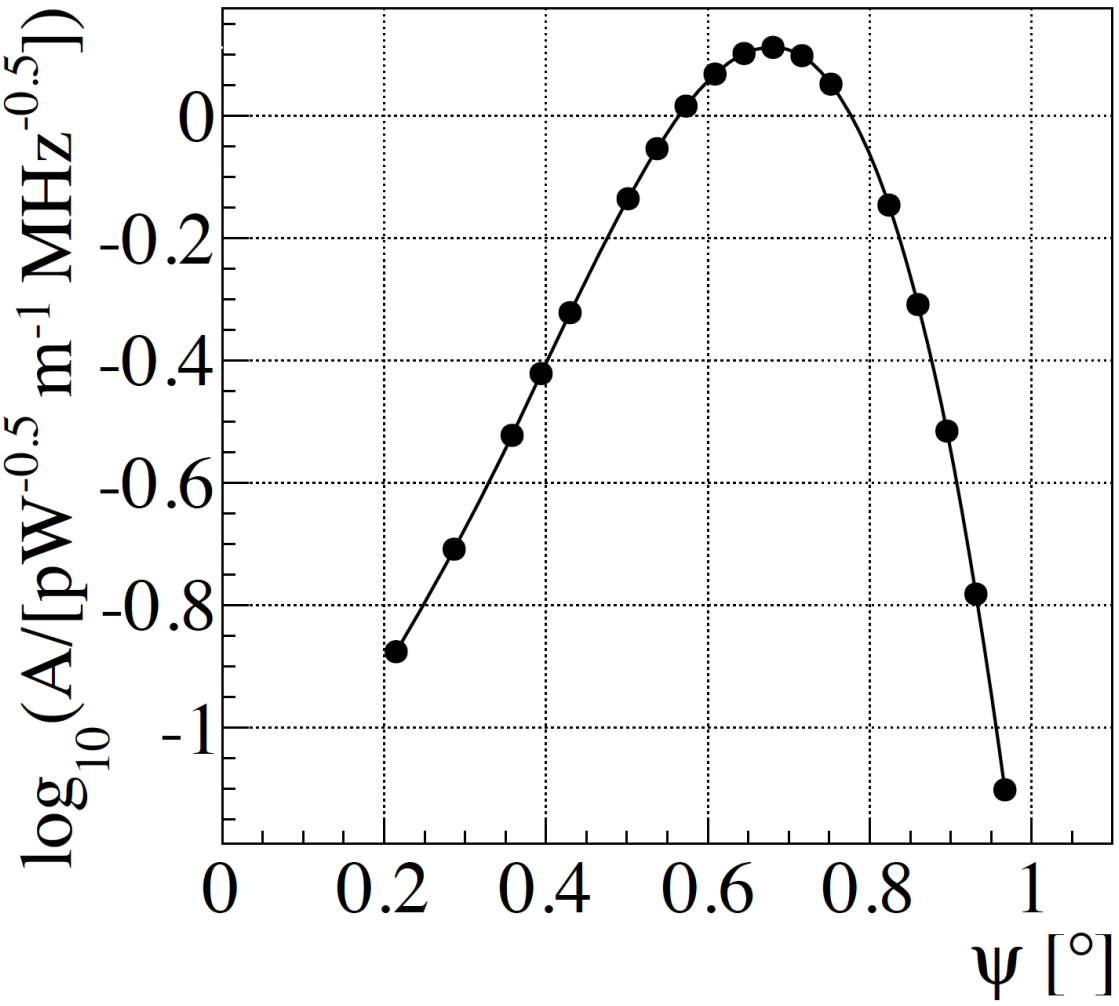

Figure 5.58: Relationship between the peak amplitude (A) of a simulated CR induced electromagnetic field at the ANITA payload and the angle from the shower axis where it is observed [7].

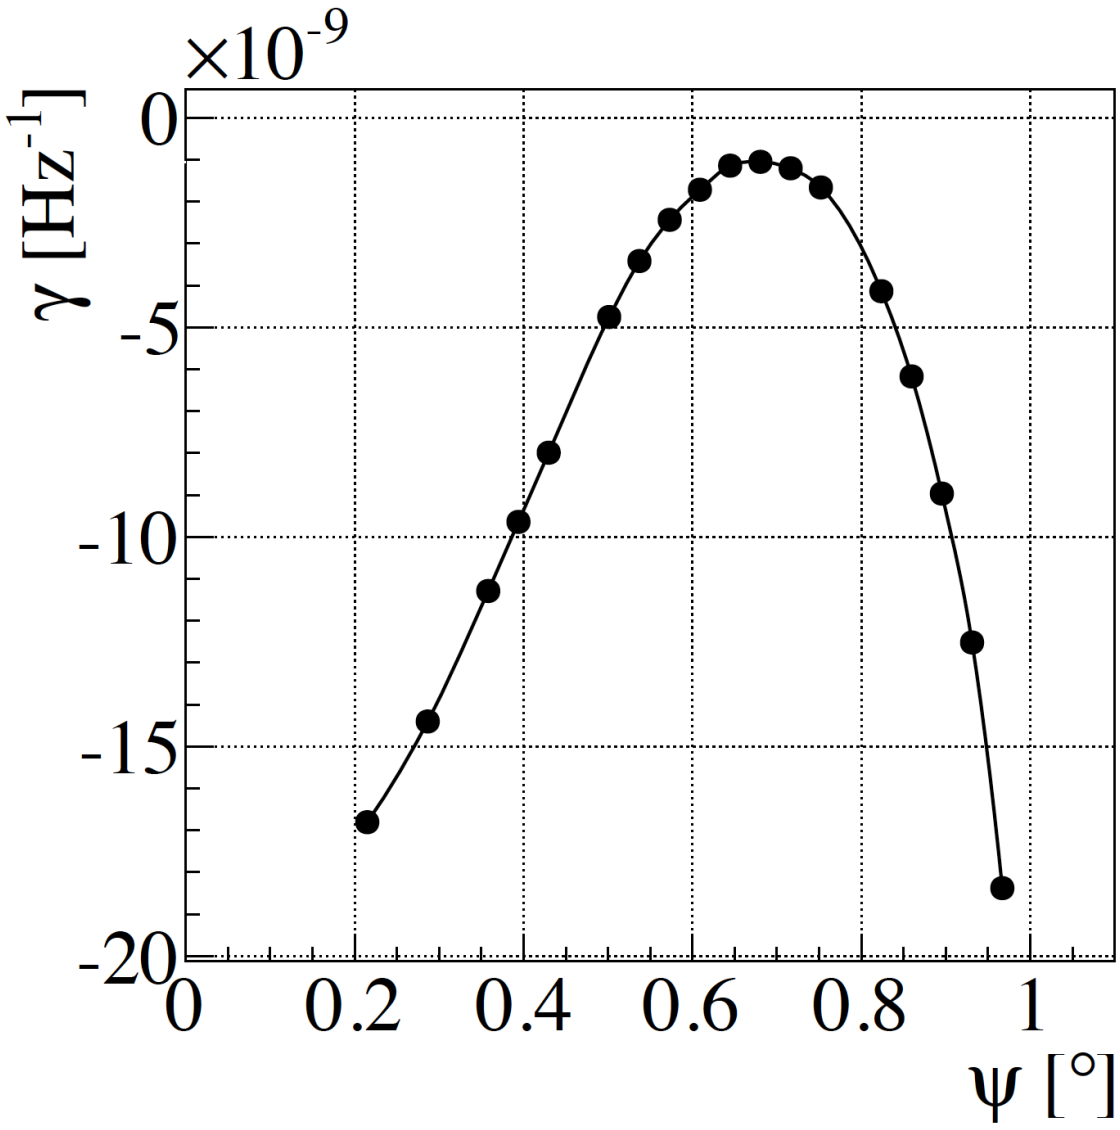

Figure 5.59: Relationship between the spectral slope (γ) of a simulated CR induced electromagnetic field at the ANITA payload and the angle from the shower axis where it is observed [7].

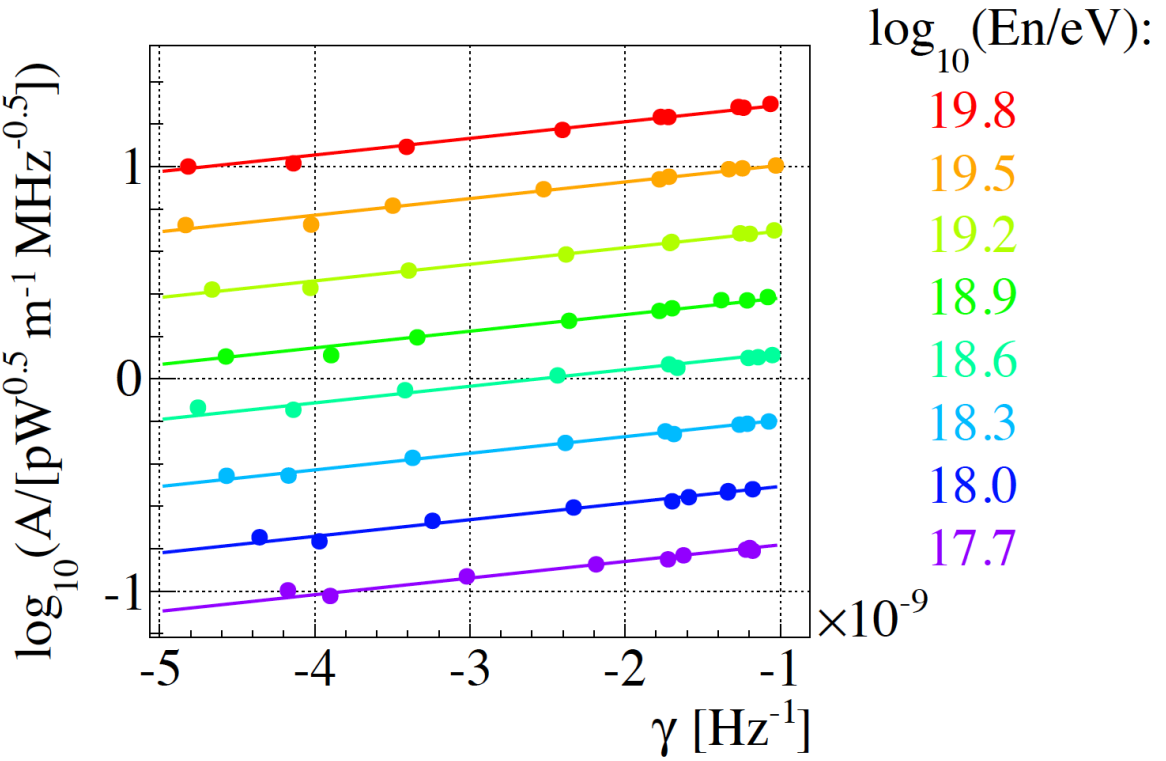

Figure 5.60: Expected CR air shower electric field impulsve spectral slope and amplitude as determined by simulation of a CR with a 70° incidence angle [7]. Each point in the two-dimensional space, which would correspond to measurements taken by the ANITA payload, relate to a unique associated energy.

rays detected in ANITA-III and detailed in this thesis are of lower energy than those detected in ANITA-I by on average a factor of $2log_{10}(2)$.

The exposure of the ANITA-III flight was lower than that of ANITA-I, however the sensitivity to low energy signals was higher. Exposure was negatively impacted by the trigger masking caused by CW interference. This masking reduced the aperture by approximately 50%, reducing the total exposure time from 22 days to an effective 11 days. This is in comparison to ANITA-I, which had 18 days of effective exposure, a 39% decrease. Analyzing the number of events in each category of 5.55 suggests that the clustering cut in the analysis cut out an additional 30% of the flight. The flux spectrum in the EeV regime has a power law of approximately 2.9 [59]. A decrease in energy sensitivity by a factor of $2log_{10}(2)$ along this power law would yield an increase in the number of observable by 80%. Combined, this suggests that this ANITA-III analysis should have observed approximately 40% more CRs than ANITA-I. The increase in the number of detected events from 14 to 19 between the two flights represents a 35% increase, which is in agreement.

5.9.3 Inverted CR candidate energy

Determining the energy of the single inverted event, 15717147, is important for understanding the likelihood that it is caused by a τ_ν incident on earth. Very roughly, we can back out the various losses that are included in ANITA-I energy estimates for the most steeply observed event in that flight.

The most steeply inclined ANITA-I event occurred at a 57° incidence angle, where the inverted event occured at a 55° incidence angle. It has a measured amplitude A_I of 0.66 $pW^{0.5}$ m^{-1} $MHz^{-0.5}$ and a spectral slope γ_I of -3.2×10^{-9} Hz^{-1}. The estimated energy for this event was determined to be 9.9 EeV. The ANITA-III inverted event has an amplitude A_I of 0.39 $pW^{0.5}$ m^{-1} $MHz^{-0.5}$ and a spectral slope γ_I of -8.6×10^{-9} Hz^{-1}. The inverted event has a weaker detected signal, but also further from the coherence angle.

The estimated reflection co-efficient for a incidence angle of 55° is only on the order of 25%, which can be used to scale the expected signal. Additionally, the up-going event will have been viewed directly, and will have a much lower expected propagation loss. After propagating through these corrections, the inverted candidate is found to have an estimated energy of 0.9×10^{18}eV.

CHAPTER 6
CONCLUSION

6.1 Summary of Results

My analysis of the ANITA-III detected 20 UHE CR EAS events over the 22 day flight over Antarctica. It also observed an inverted polarity CR signal that is consistent of an upward going τ decay from a UHE ν_τ transiting through the earth. Current models for the cross sections of these astrophysical particles at UHE are not in agreement with this detection. Further study of the event and the systematic uncertainties presented in this analysis must be undertaken before any conclusions can be drawn.

6.1.1 Impact of inverted event

The inverted event, if assumed to be a UHEτ_ν that interacted within the earth and created a τ which decayed in the atmosphere before being observed by the payload, presents issues for the current understanding of the standard model. Results have been published that utilize simulations to estimate the likelihood that such an event will occur [13]. The resulting likelihoods for a variety of particle incidence angles and energies can be seen in Figure 6.1. The detection of an event with an emergence angle of 35° is in strong disagreement with this estimate, especially at an estimated energy of 0.9×10^{18} eV.

6.2 Future Analysis Work on the ANITA-III payload

During my analysis I discovered several calibration issues which needed to be addressed. This includes the off axis response and cross-pol leakage of the antennas, which is not well constrained. Additionally, an absence of 45° polarized calibration signals presents challenges for accurage alignment of the phase centers for the two different polarization ridges. These uncertainties have a strong effect on energy reconstruction and plane of polarization angle measurement for the detected signals, and will need to be corrected for an accurate result. The geomagnetic angle is of particular interest, as it provides an additional parameter with which to discriminate between background and signal events.

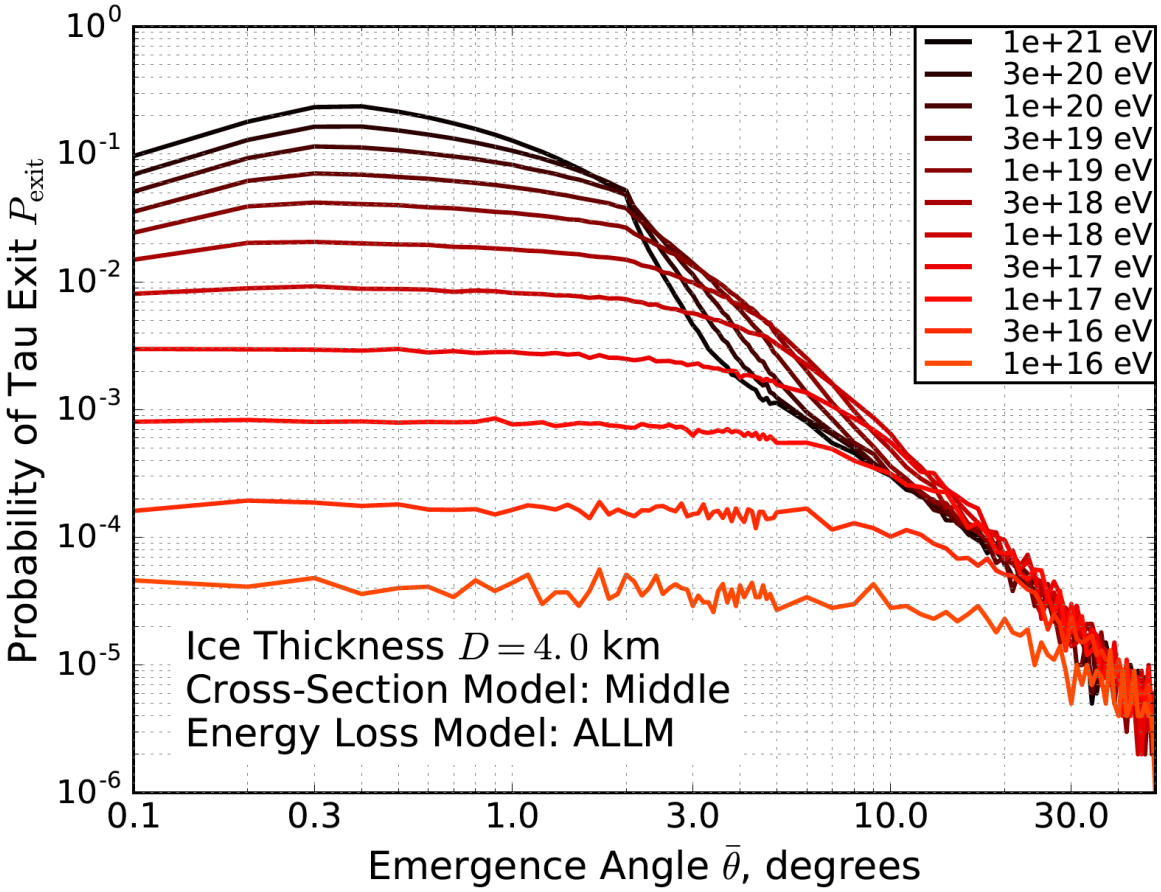

Figure 6.1: Expected likelihoods for tau lepton emergence after UHEν_{tau} interactions within the earth for a variety of angles and energies from simulations [13].

6.3 ANITA-IV Flight and Improvements

6.3.1 Instrument development

The ANITA-IV payload, an upgraded version of the ANITA-III instrument, flew in the 2016-2017 Antarctic balloon campaign. It corrected several issues that improve exposure, sensitivity, and provide additional constraints on measurement uncertainties that plagued ANITA-III. These include a tunable notch filter on each channel to eliminate satellite background, and additional calibration pulser measurements.

Since the time LAB3 was developed in 2005, there has been significant developments to the LABRADOR architecture. Currently, the most advanced generation of the ASCI, the LAB4D, is a single channel readout that improves upon the previous generations by both vastly increasing the storage buffer which increases the record length, or buffer depth of the chip. It also allows for the correction of dT offsets onboard the chip, minimizing the requirement for post-digitization correction of the waveform and the subsequent high-frequency signal loss. These improvements will allow for even greater sensitivity in any future ANITA flights.

There has additionally been significant development work on an interferometric trigger utilizing the Realtime Independent Three-bit Converter (RITC) ASIC, which would increase the sensitivity of a future ANITA flight, and allow for the observation of signals with lower incident particle energies [60].

6.3.2 Physics Outlook

The results of this analysis significantly enhance the value of the recovered ANITA-IV flight data. With increased exposure and sensitivity, additional inverted polarity detections would confirm the event found in this analysis.

The future looks bright for ANITA.

APPENDIX A
SYSTEM IMPULSE RESPONSE

The generation of a system impulse response for the ANITA3 instrument from pre-flight calibration measurements is required to accurately relate the digitized time series of analog voltages propagated to the detector electronics to the correlating electric field physically transiting the instrument. As the system is sensitive to a wide bandwidth, a full frequency dependent array of complex phasors is required to completely characterize the system. The ANITA3 (and more recent ANITA4) flights employ slightly different antennas than the previous ANITA1 and ANITA2 flights. The purpose of the change was to capture additional sub 200MHz frequency electromagnetic signals. Additionally, the amplification and filtering chains of the newer flights have undergone subtle modifications that preclude direct relation to those of earlier flights. This change necessitates a thorough analysis and characterization of the full ANITA3 signal chain, from antenna to digitizer.

A.1 Introduction

There are two distinct systems on the ANITA3 instrument that alter the received time domain electromagnetic signal prior to digitization: the Seavey horn antennas, and the filtering and amplification network. The antennas used on the ANITA flight are broad band, with a design bandwidth of 180MHz to 1200MHz, and whose dual orthogonal polarizations have co-located phase centers. These antennas couple the propagating electromagnetic (EM) radiation emitted by an extensive comic ray or neutrino air shower (EAS) into a 50Ω transmission line by uniformly transforming the impedance from that of free space to that of the RF system. The collection of amplifiers, filters, and other 50-ohm RF components that populate the path between the antenna and the LABRADOR chip is colloquially known as the signal chain, and adds additional frequency dependent gain and phase modulation before the signal is finally read out by a fast sampling broad band ADC. To calibrate these two systems, input signals were injected into the systems and measurements were taken of both the input and output signal on a calibrated oscilloscope. The methods and techniques used to generate an impulse response using these input and output signals will be discussed, as well as the results for the ANITA3 flight.

A.1.1 Discrete Fourier Transform

To determine the spectral content of a time domain signal, I have employed the use of a Discrete Fourier Transform (DFT). Equation A.1 [61] exhibits a transformation between a discretely sampled time domain signal $x_n(t)$ and the equivalent Fourier series $\mathcal{F}_k(f)$. Additionally, Equation A.2 [61] shows the inverse of the DFT, with a $\frac{1}{N}$ normalization to preserve reciprocity of the transformations.

$$\mathcal{F}_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i k n/N} \tag{A.1}$$

$$x_n = \frac{1}{N} \sum_{k=0}^{N-1} \mathcal{F}_k e^{-2\pi i k n/N} \tag{A.2}$$

As our measured signal is required to be purely real, the Fourier equivalent series will be hermitian, setting the limits for the sum to purely positive integers. Additionally, though a true Fourier transform takes the time or frequency integral over its respective range and thus introduces a time multiplication to the resulting unit, the DFT does not have a time normalization, leaving the units of both arrays as Volts.

The effect of transforming from a time series to a Fourier series and back should be zero, however any signal manipulations such as zero padding, up-sampling, or truncating, will have an effect on the opposing Fourier equivalent pair as well. This can be used to more easily manipulate signals in a manner that matches the physical boundaries of the system.

A.1.2 Units and Spectrum Discussion

Accurately representing the amount of power in a time domain waveform or network as a function of frequency requires both a normalized Fourier transform as well as the appropriate units. There are two main types of signals which we are considering in this analysis, absolute measurements, referenced to a measured voltage on constant impedance network, and relative measurements, relating two different measured signals.

An example of an absolute measurement is the time series waveforms captured on oscilloscopes. These can be described by decibles referenced to 1mW (dBm), seen in Equation A.3, or Power Spectral Density (PSD), represented by dBm/Hz and seen in Equation A.4. For the purposes of this analysis, the Power Spectral Density will be used.

$$Power(f) = 10log_{10}(\frac{|\mathcal{F}(f)|^2 * 1000[\frac{mW}{W}]}{Z}) \qquad [dBm] \tag{A.3}$$

$$PSD(f) = 10log_{10}(\frac{|\mathcal{F}(f)|^2 * 1000[\frac{mW}{W}]}{Z * df}) \qquad [dBm/Hz] \tag{A.4}$$

The ratio of spectral power between two different signals $x_A(t)$ and $x_B(t)$, can be done through division of their equivalent Fourier representations $\mathcal{F}_A(f)$ and $\mathcal{F}_B(f)$. The logarithmic magnitude of this ratio is what is most often recognized as the gain of the network. The calculation of the gain is seen in Equation A.5.

$$Gain(f) = 10log_{10}(|\frac{\mathcal{F}_A(f)}{\mathcal{F}_B(f)}|^2) \qquad [dB] \tag{A.5}$$

An example of relative measurements are network analyzer S21 measurements and transfer functions, which do not quantify an absolute power, only a ratio of powers. Note that the subtraction of two absolute measurements, regardless of normalization, yields a gain. Also note the distinction between Equation A.5, which measures the logarithmic increase in power, to other quantities commonly referred to as gain which measure the ratio between amplitudes. In this case, there is a factor of two difference between them.

A.2 Seavey Antenna Impulse Response

A.2.1 Measurement Goals

Many calibration measurements were taken for the antennas used in the ANITA3 flight. Each measurement probed a different characteristic of the antennas, and had various drawbacks and benefits. There are several major parameters of interest for the antenna measurements. First, the absolute phase and gain response for any incident EM field at the peak response angle. This angle of peak response, nominally pointed orthogonally to the antenna face, is known as the "boresight" angle, and the resulting quantity is known as the antenna gain, represented in units of dBi and discussed later. The directionality and gain pattern as a function of angle off the maximal transmission and receive direction is directly related to the antenna gain, and is needed to determine amplitudes for signals incident at angles off of the boresight. Lastly, a measurement of the cross polarization fraction is required to determine the full Stokes parameters of any incident field. This last measurement has currently not been studied for the ANITA3/4 antennas.

A.2.2 Measurements Summary

Three different measurements of the antennas were taken that are discussed in this paper, each of which probe a unique region of the full calibration measurement.

UH Anechoic Chamber

First, a measurement was taken at the University of Hawaii at Manoa in a copper lined, RF sealed, anechoic chamber using two "identical" (same model and manufacturing batch) ANITA3 flight antennas. This measurement allowed a maximally noise free environment in which to measure any small signal effects that may be present in the antennas. It also removes the requirement for correlation and averaging of multiple waveforms, decreasing uncertainty introduced by interpolation techniques. Unfortunately, the size of the anechoic chamber forced the distance between the antennas into the near field region for the lower side of the frequency bandwidth. This causes results below 200MHz, with a wavelength of 1.6m and minimum far field distance requirement of 3.2m, to be slightly distorted and require near-field Fresnel corrections. As this is the region in which the

antennas were supposed to perform differently than previous ANITA antenna designs, this experimental setup and required correction is highly problematic. Additionally, measurements of the gain pattern far off the maximal bore-sight angle have Fresnel zones occluded by the absorbing material that lined the walls of the chamber, once again introducing Fresnel interference and increasing the uncertainty in the measurements.

UH Rooftop Measurements

Additional tests were undertaken at the University of Hawaii with much larger separation distances. To accomplish this, the two antennas under test were placed on building rooftops separated by a moderately sized courtyard. The building separation was large enough to escape the near-field of the lower frequency antenna response, yet close enough to preclude any ground bounce interference or multi-path issues. These measurements were plagued by anthropogenic noise, however a large sample of waveforms uncorrelated to this noise was taken reducing the resulting background significantly. These measurements provide the best results for the absolute system gain and antenna directivity.

Palestine TX 2014 Antenna Consistency Checks

Finally, for the full ANITA flight the consistency and similarity of all 48 antennas needed to be measured. Bore-sight gain measurements were taken for all antennas at the Columbia Scientific Balloon Facility (CSBF) in Palestine TX during the 2014 integration campaign of the ANITA3 payload. This was done in a CSBF balloon hangar, is used in this analysis, and is detailed in the following section.

Due to the physical and engineering limitations of the experimental apparatus, one of the antennas had to be flipped in relation to the other, causing the approximation of identical antennas to lose some justification. This introduces several issues. First, there will be a polarity flip of the measured signal. This can be easily handled by multiplying the measured time series by -1, re-inverting the signal. A worse problem occurs when the peak power angle of the antenna beam "wanders" away from the boresight. In the case where the antennas were oriented symmetrically to each other, an antenna height that peaked at a location other than the boresight angle would not distort the measurement. However, when the antennas are oriented in a flipped configuration, any angular deviation of the peak antenna pattern will be squared as each antenna contributes its own angular deviation.

A.2.3 Absolute Boresight Complex Antenna Height

The relationship between the magnitude of an electric field vector and the induced potential on the transmission line output of the antenna is a combination of the radiation efficiency and beam

pattern, and is colloquially known as the complex antenna height, measured in units of volts per meter, and represented in this paper by $\mathcal{H}_{ant}$.

The equation for determining the complex antenna height of an antenna receiving an impulse in a transceiver pair, where the transmission antenna's height is known, is given by Equation A.6. [62][63]

$$\mathcal{H}_{Rx}(f) = \frac{cr(f)\mathcal{F}_{rec}(f)}{if\mathcal{F}_{src}(f)\mathcal{H}_{Tx}(f)} \tag{A.6}$$

In this equation, f is the radiation frequency, $\mathcal{H}(f)$ represents the complex antenna height of the receive (Rx) and transmit (Tx) antennas, c is the speed of light, $\mathcal{F}(f)$ is the Fourier equivalent of the time sampled waveform for the source (src) and received (rec) impulse, and $r(f)$ is the free space path distance between the phase centers of the antennas.

The location of the phase centers as a function of antenna geometry were not known prior to this analysis, and would require additional absolute timing measurements which were not available. For the purposes of this analysis, a linear relationship between frequency and phase center was derived assuming two points: the 180MHz phase center is at the antenna face, and the 1.2GHz phase center was at the feed point. The distance between the antenna face and the feed point is 21.86" (55.52cm) for the vertically polarized channel, and 22.65" (57.53cm) for the horizontally polarized channel, as seen in Figure A.1.

If the antennas are identical, i.e $\mathcal{H}_{Rx}(f) = \mathcal{H}_{Tx}(f)$, Equation A.6 becomes:

$$\mathcal{H}_{ant}(f) = \sqrt{\frac{cr(f)}{if}\frac{\mathcal{F}_{rec}(f)}{\mathcal{F}_{src}(f)}}$$

Note that using Equation A.2.3 results in a sign degeneracy, as there are two possible solutions that satisfy the square root term. This manifests itself as a lack of constraint on the absolute polarity of the emitted signal. Determining the absolute polarity of the antennas requires using a calibrated reference antenna of which the polarity is previously known.

Complex antenna height relates electric field to measured voltage according to the following equation:

$$\frac{V_{rec}(t)}{\sqrt{Z_{sys}}} = h(t) \circ \frac{E_{rad}(t)}{\sqrt{Z_o}} \tag{A.7}$$

Where $\circ$ is the convolution operator, Z_{sys} is the matched impedance of the transmission line (50Ω in our system) and Z_o is the impedance of free space (377Ω). Using these relationships it is possible to solve for the conversion between electric field and induced voltage.

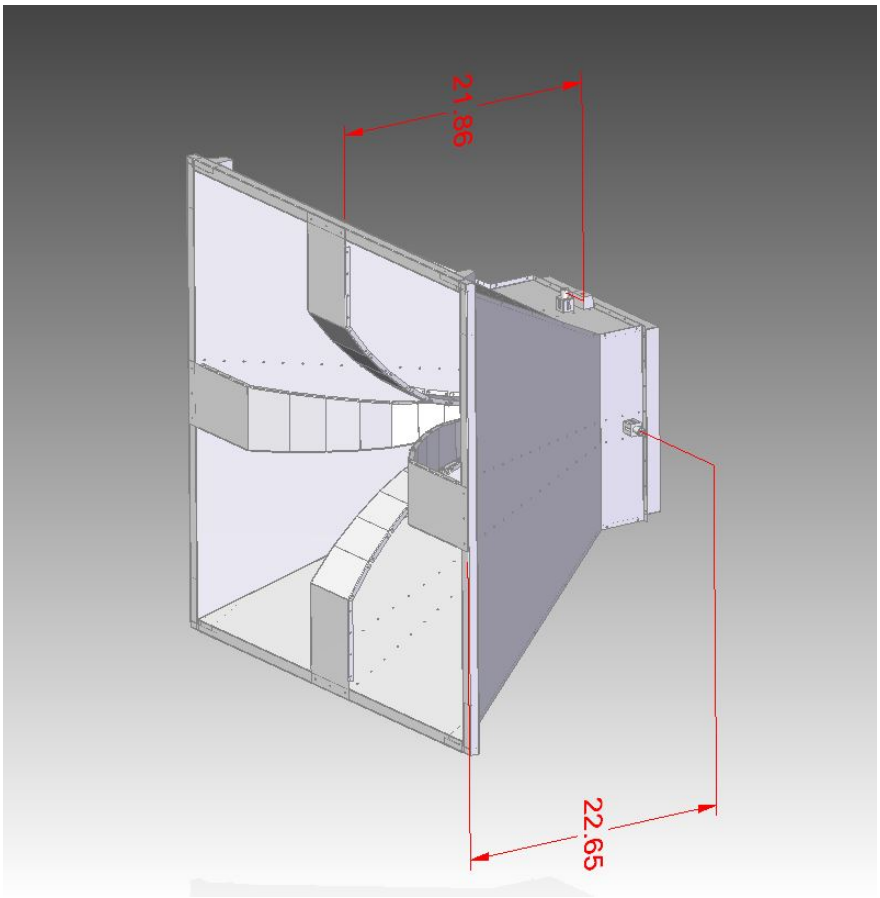

Figure A.1: CAD model of the Seavey quad ridge horn antenna. The distance between the face of the antenna and the feed point is annotated for use in estimating the phase center location as a function of frequency.[14]

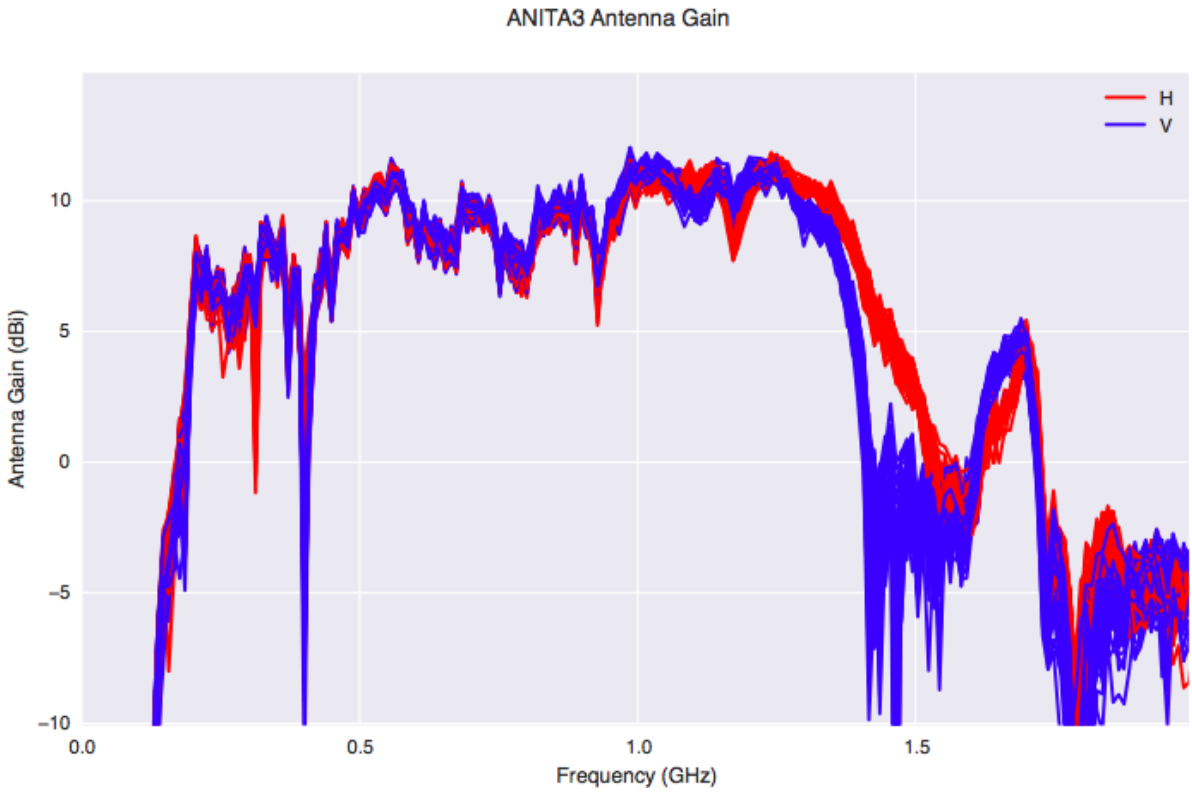

Figure A.2: Results for the antenna gain as a function of frequency for all 48 ANITA3 horn antennas.

A.2.4 Antenna Gain

The figure of merit for the absolute radiative power of an an antenna is called the Antenna Gain, and is calculated by taking the logaritm of the ratio between the measured peak radiated power of the antenna under test and that of an isotropically emitting antenna. This value takes on the units dBi. The equation to determine the antenna gain from the normalized complex antenna height is shown in Equation A.8.

$$G_{eff}(f) = \frac{4\pi f^2}{c^2}|\mathcal{H}_{ant}(f)|^2 \tag{A.8}$$

A.2.5 Results of Palestine Antenna Measurements

All 48 antennas for the ANITA3 flight were measured at the Columbia Scientific Balloon Facility (CSBF) in Palestine, TX during the 2014 integration campaign. The complex antenna height did not noticably vary between antennas, and resulted in an average antenna gain of approximately 10dBi. Figures A.2 through A.6 show the results for all antennas.

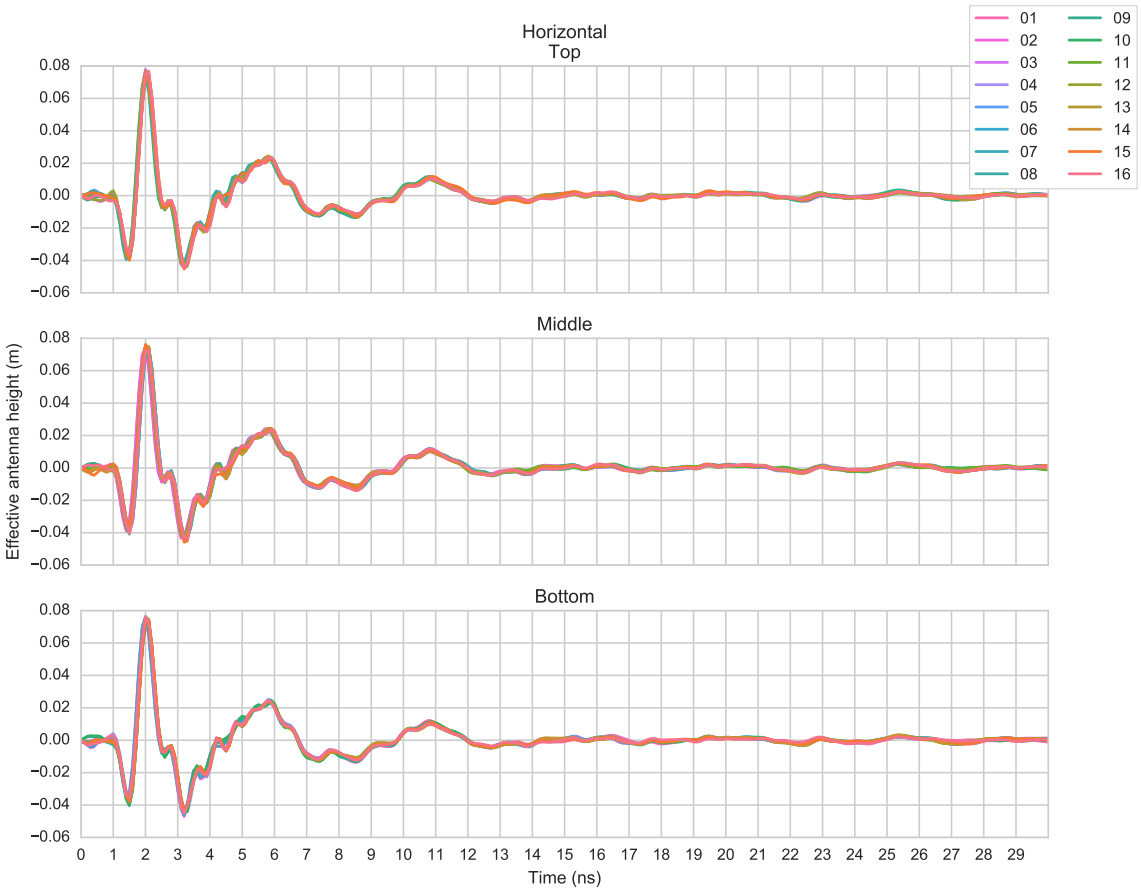

Figure A.3: Time domain impulse response of the horizontally oriented polarizations of the ANITA3 horn antennas.

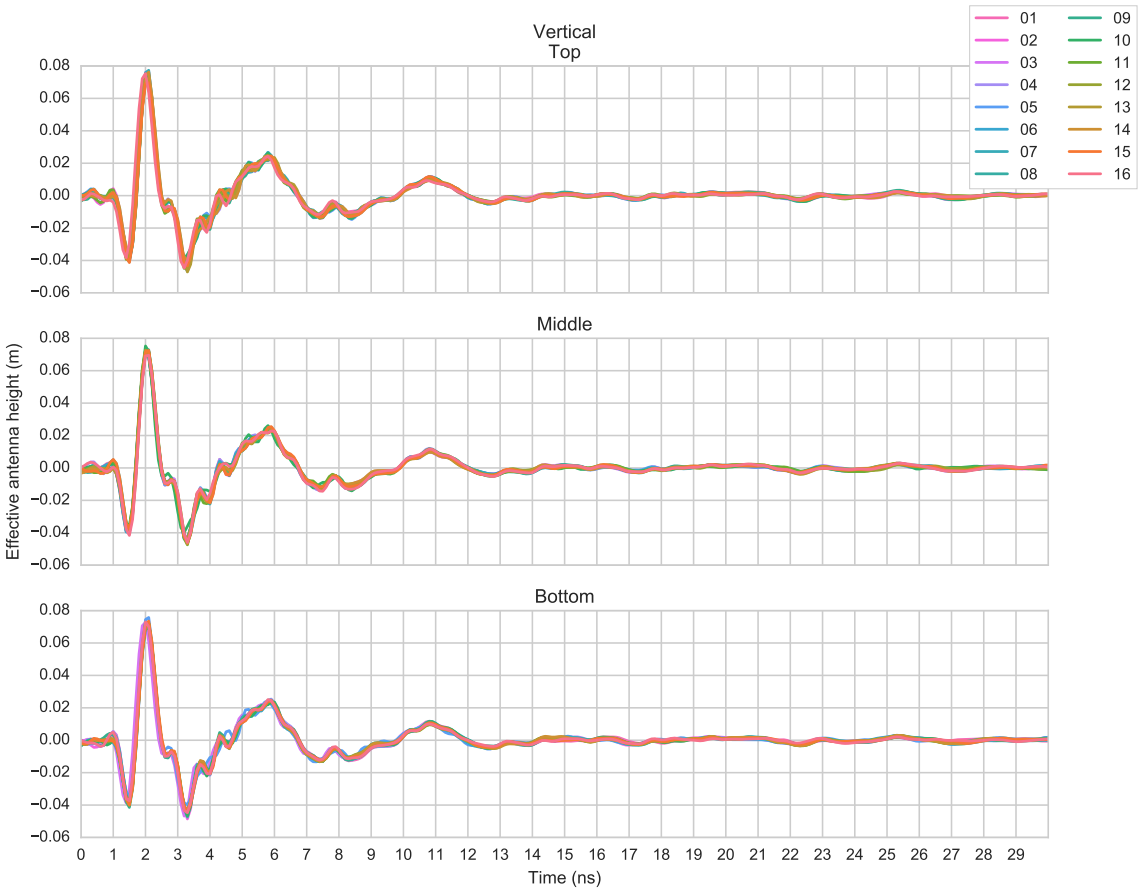

Figure A.4: Time domain impulse response of the vertically oriented polarizations of the ANITA3 horn antennas.

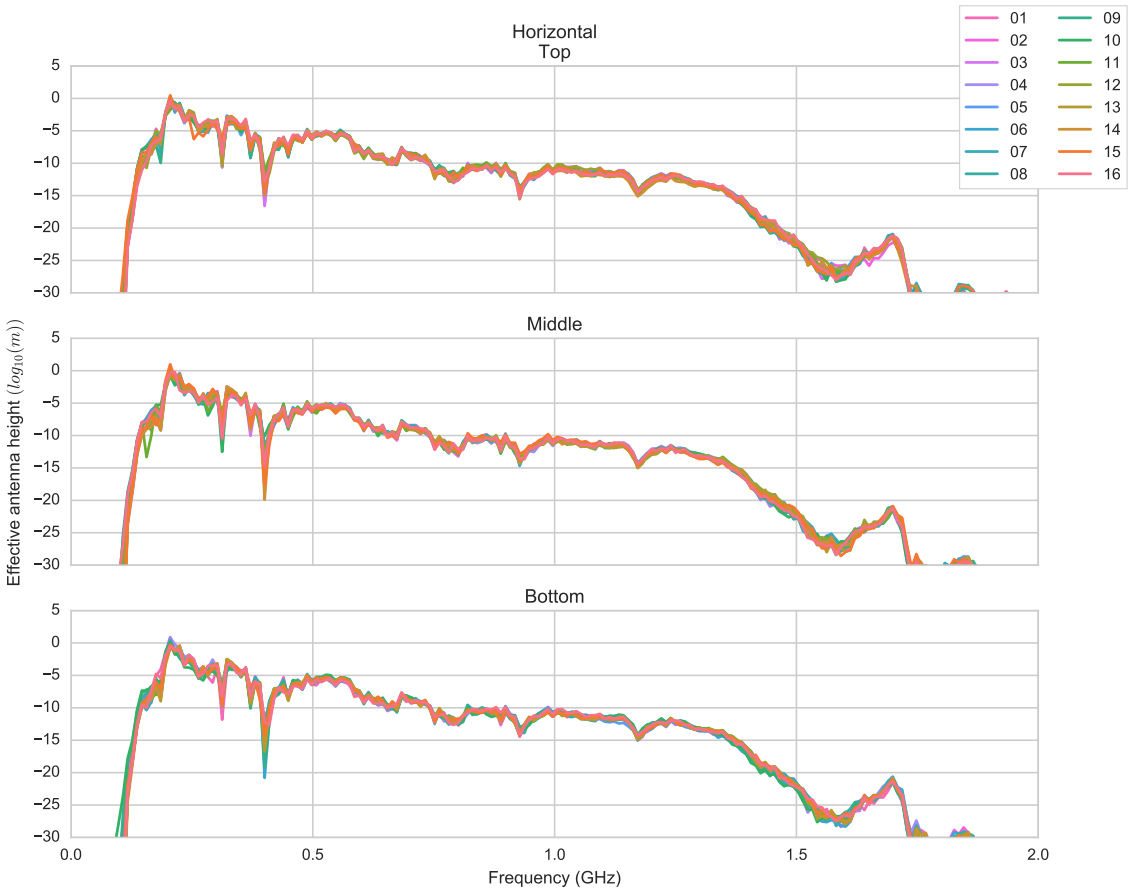

Figure A.5: Complex antenna height for the horizontally oriented polarizations of the ANITA3 horn antennas.

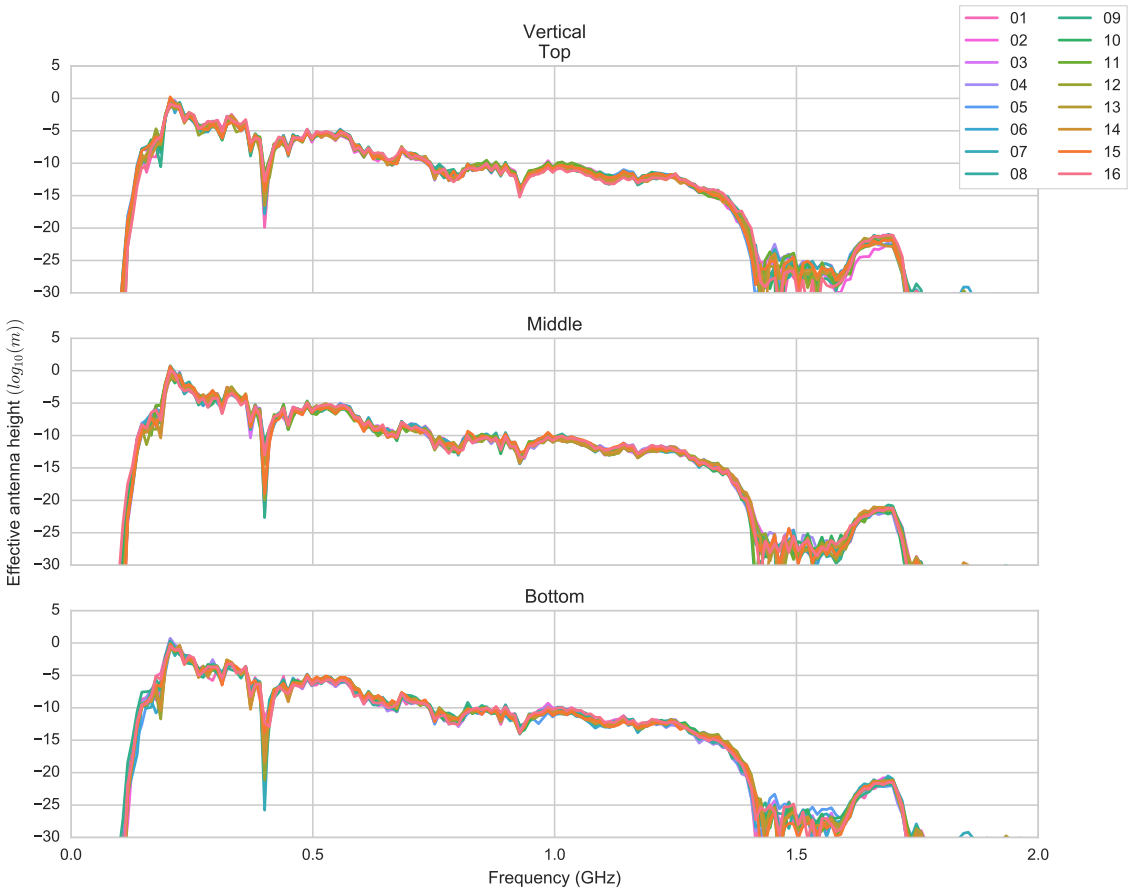

Figure A.6: Complex antenna height for the vertically oriented polarizations of the ANITA3 horn antennas.

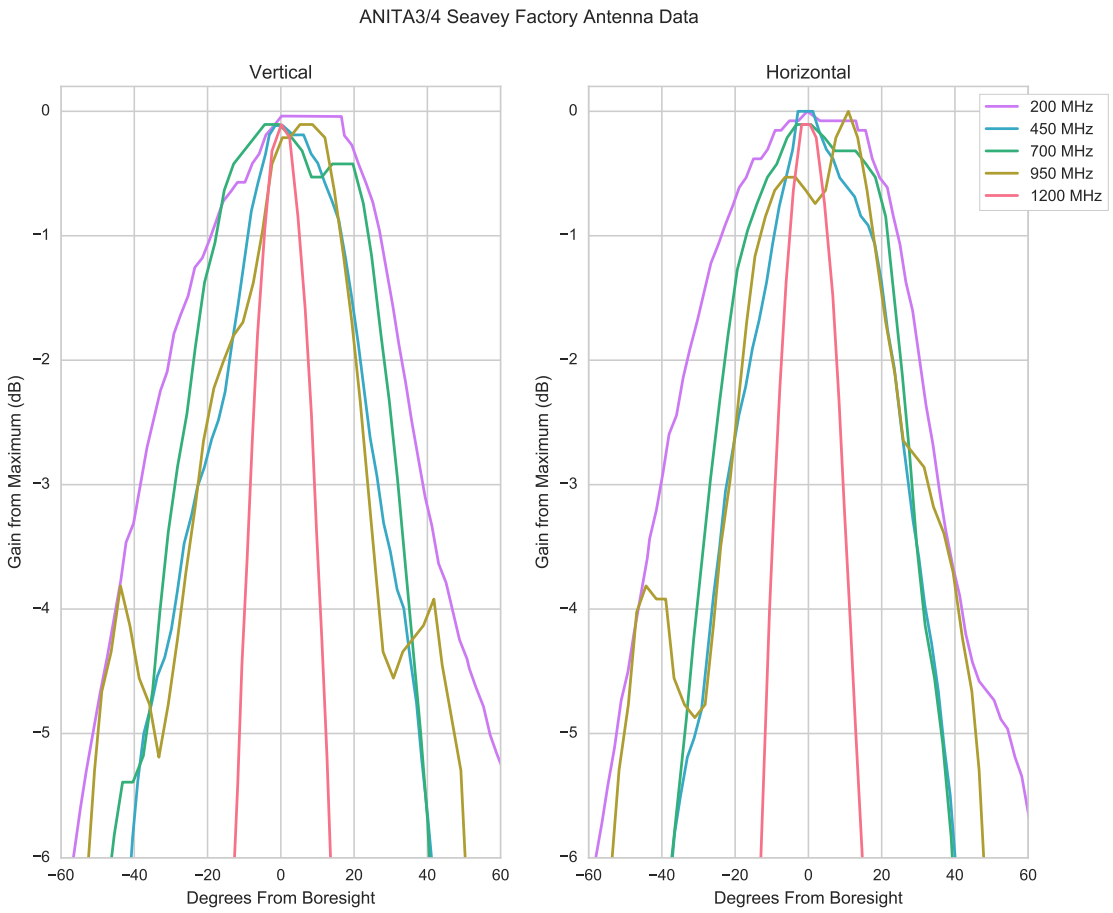

Figure A.7: Replotted off angle antenna responses received from Seavey after antenna manufacturing. Gains are referenced to the maximum received power, which is not required to be the boresight angle. Responses are in reference to the electric field incident on the respective polarization ridge of the antenna.

A.2.6 Off Angle Antenna Response

An additional characteristic of the antennas is their response to signals that are incident on the antenna at angles off the boresight. Seavey provided manufacturing test result plots for the off angle response at five discrete frequencies. These were manually digitized and replotted, and are presented in Figure A.7.

Off Angle Antenna Measurement

To independently measure the off angle response of the antennas, the UH anechoic chamber data was used, despite its low frequency near field limitations. The above procedure for developing complex antenna height was done for multiple rotation angles of the receive antenna while the transmission

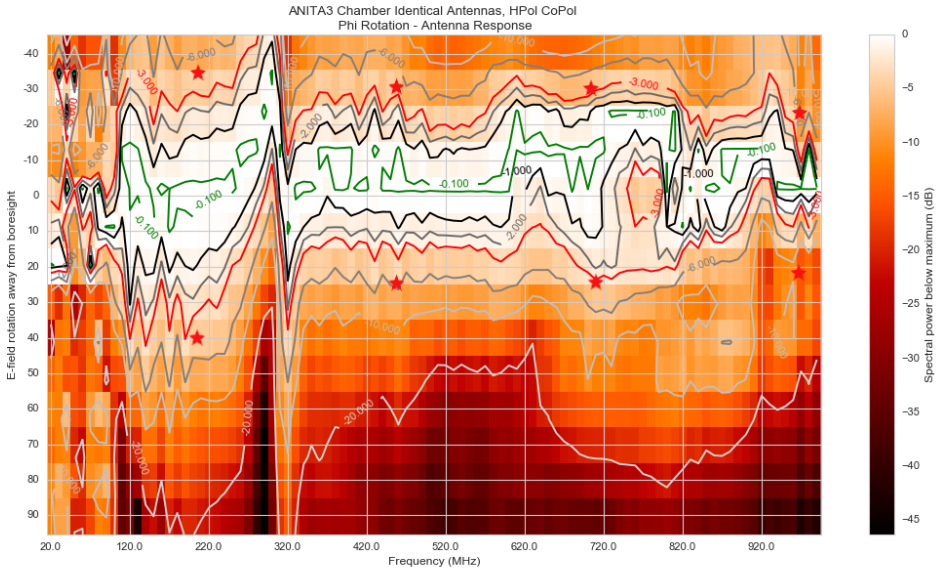

Figure A.8: Off angle antenna gains from the analysis of the Seavey horn antennas for the E-field of the Horizontally polarized ridge, rotated horizontally. This corresponds to the azimuthal response for the Hpol channels of the ANITAIII instrument. Measurements were taken in the UH anechoic chamber with identical receive and transmit antennas, with only the receive antenna being rotated. Counter-clockwise rotation beyond 45° is obstructed by the chamber wall. Red stars denote the -3dB points from the Seavey manufacturing data at the frequencies measured by them. A "wandering" of the beam pattern caused by the Vpol feed structure can be seen at 320MHz.

antenna was kept stationary. Because of this, Equation A.6 must be used with $\mathcal{H}_{TX}(f)$ set as the average complex antenna height measured wile the antennas were in the zero rotation angle configuration. Cables and propagation losses are removed from the measurement. The results of this analysis are shown in Figures A.8 and A.9.

Th

A.3 RF Signal Chain Impulse Response

A.3.1 Measurements Summary

Complex phasor response measurements for each of the 96 signal chains flown in the ANITA3 instrument were taken in Antarctica immediately preceding the flight. A broad band impulse was injected into each channel sequentially and measured by the on board flight electronics. This signal was split at the source and simultaneously measured by an oscilloscope, providing an input and output reference signal which we can compare to determine the full complex transfer function. A diagram of the setup is seen in Figure A.10.

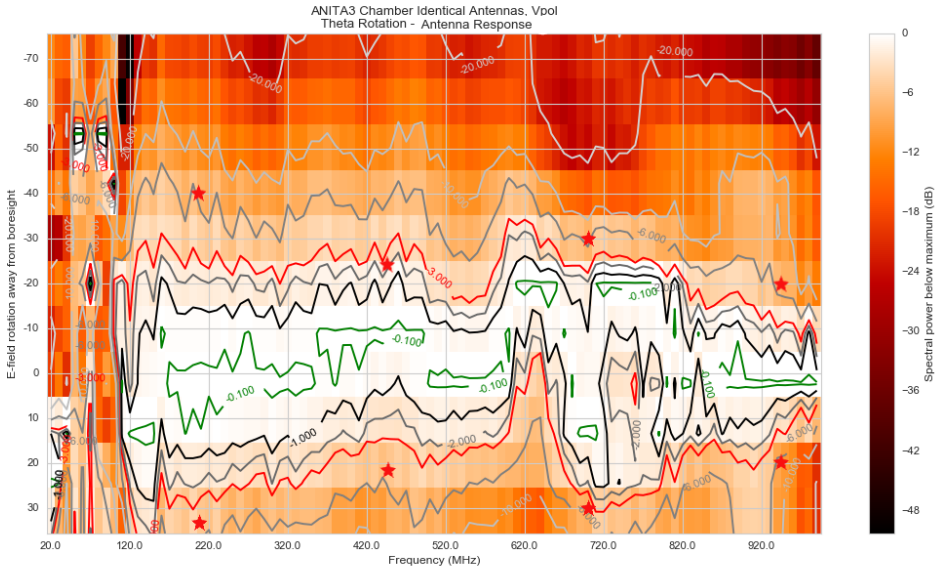

Figure A.9: Off angle antenna gains from the analysis of the Seavey horn antennas for the E-field of the Vertically polarized ridge, rotated vertically. This corresponds to the elevation response for the Vpol channels of the ANITAIII instrument. Measurements were taken in the UH anechoic chamber with identical receive and transmit antennas, with only the receive antenna being rotated. Downward rotation beyond 35° is obstructed by the chamber floor. Red stars denote the -3dB points from the Seavey manufacturing data at the frequencies measured by them.

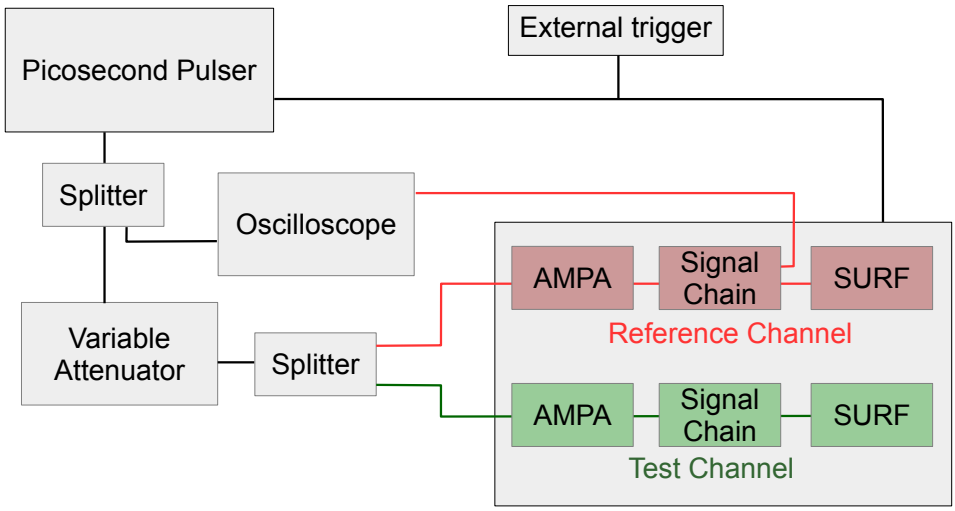

Figure A.10: A schematic of the pulse insertion calibration setup used immediately before flight with the full signal chain in Antarctica in 2014

Note that the cables connecting the various elements of the calibration test setup are not present in the final flight configuration, and needed to be separately measured and removed.

A.3.2 Complex Transfer Function

The full mathematical relationship between two waveforms that have traversed a non-linear signal network is called the complex transfer function, which I will denote with $h_{sig}(t)$ and its Fourier transform equivalent $\mathcal{H}_{sig}(f)$. The relationships between a measured input signal $V_{in}(t)$ and output signal $V_{out}(t)$, and their respective Fourier transform equivalents $\mathcal{F}_{in}(f)$ and $\mathcal{F}_{out}(f)$, are shown in Equation A.9.

$$\begin{aligned} V_{out}(t) &= h_{sig}(t) \circ V_{in}(t) \\ \mathcal{F}_{out}(f) &= \mathcal{H}_{sig}(f)\mathcal{F}_{in}(f) \end{aligned} \tag{A.9}$$

The $\circ$ denotes a convolution operator. The Fourier domain equivalent to convolution is complex multiplication, which can be seen in this equation pair. Using these equations and the measurements of input and output digitized time domain waveforms we can determine the complex transfer

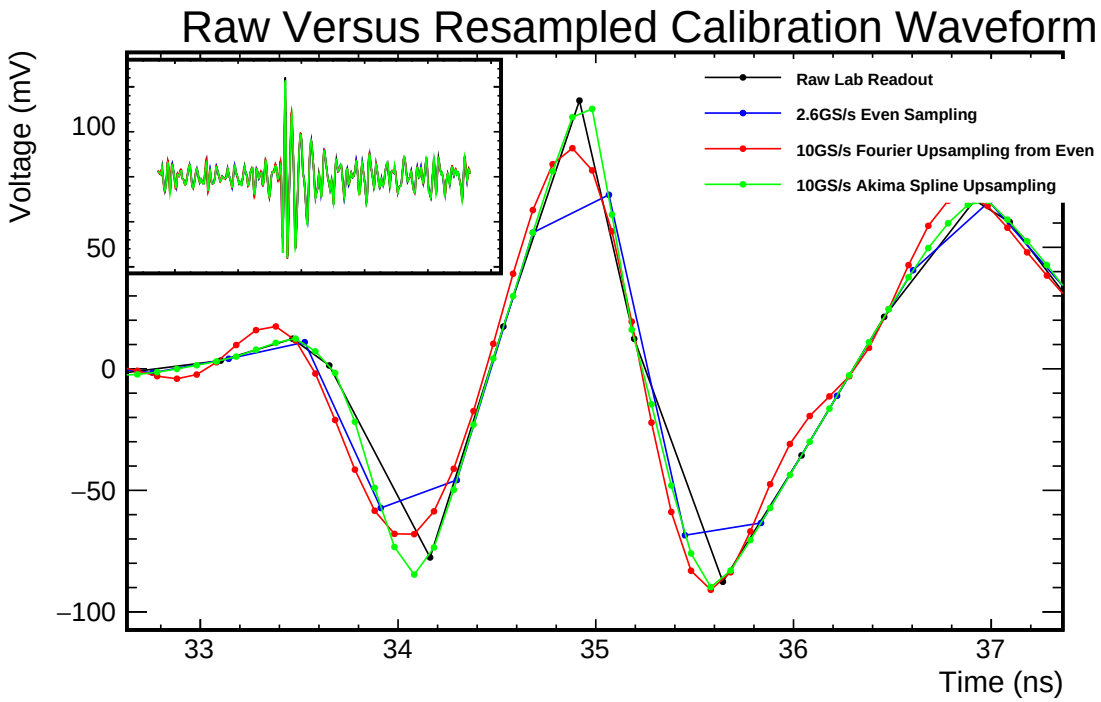

Figure A.11: An example from channel 01BH of even resampling and up-sampling the original unevenly sampled LABRADOR digital readout. The peak of the impulse is shown zoomed, while a view of the entire waveform can be seen in the inset frame showing the macro-consistency of the methods. The method used in this analysis is FFT up-sampling. Despite undershooting some points, it minimally distorts the spectral power and phase response of the signal.

function for each channel of the ANITA3 signal chain.

A.3.3 Methodology

The input and output waveforms were recorded by different devices with different sampling rates and record lengths, and so some digital processing must be done to make them directly relatable. The oscilloscope's 10GS/s (100ps/sample) digitized readout of the input pulse waveform was treated as the sampling rate reference for this analysis. The record length was chosen to be 1024 samples which allows for a 102.4ns window, comparable to the nominal window size of the LABRADOR.

The LABRADOR returns unevenly sampled waveforms at a nominal 2.6GS/s (385ps/sample) rate. To compare the two signals, first an Akima piecewise cubic spline was applied to determine an equivalent evenly sampled waveform. This waveform was either zero padded or cut in order to make the number of points equal to 256. Next, the Fourier equivalent series, which has $\frac{N}{2}+1 = 129$ samples was found. This was zero padded by a factor of $\frac{dT_{old}}{dT_{new}} * N_{\mathcal{H}} = \frac{384ps}{100ps} * 129$, up to 496 samples. Taking the inverse Fourier transform of this zero padded series results in an up-sampled time domain waveform at 10GS/s. An example of this up-sampling is shown in Figure A.11.

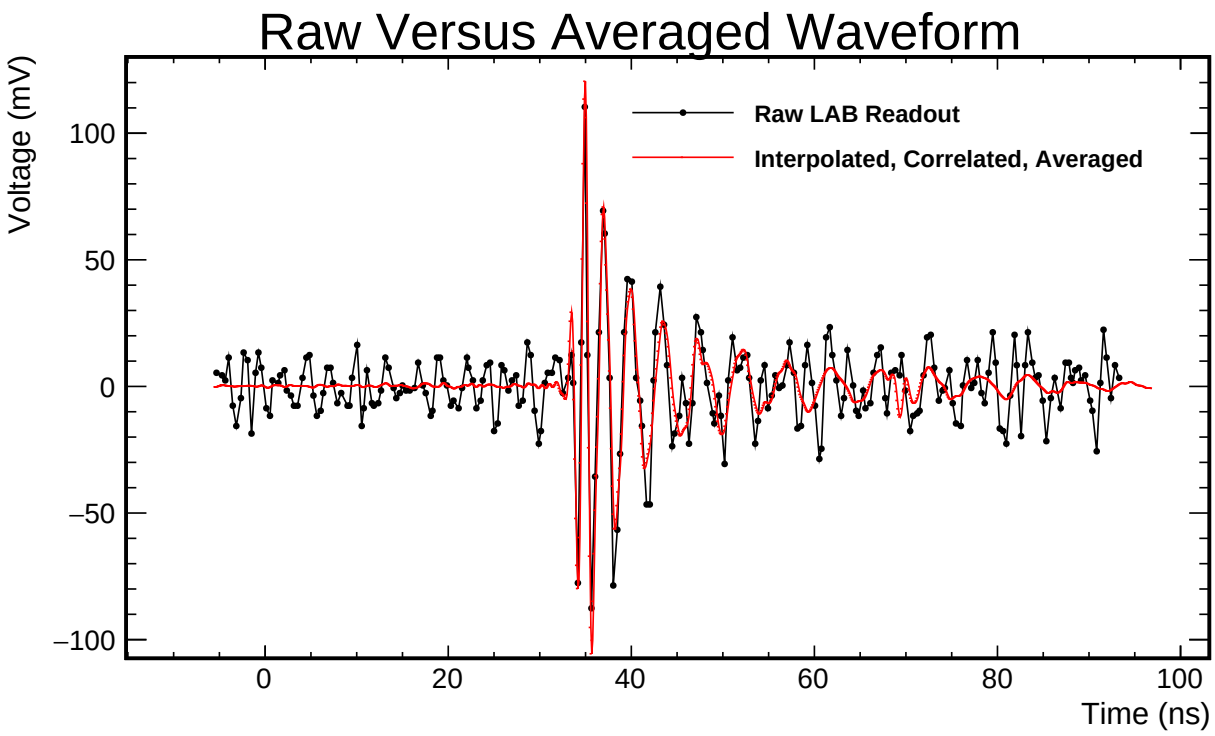

Figure A.12: An example from channel 01BH of the reduction in individually sampled waveform thermal and electronic noise versus the correlated and averaged signal result.

After the input and output time domain waveforms had the same waveform characteristics, they were separately correlated and averaged to reduce thermal and electronic noise introduced by the system. The pulse captured by the oscilloscope underwent no amplification and had an extremely high signal to noise ratio, and this averaging was done primarily to equivalently sample the interval over which output pulse data would be averaged. However, the signal chain's high gain specification, required for detection of the physics phenomenon under test, necessitated significant attenuation of the signal prior to it being fed into the front of the chain. This lowered the signal to a level only moderately above the electronic noise introduced by the first stage amplifiers and room temperature thermal noise present in the attenuator and cables. This allowed real physics signals to fall within the dynamic range of the digitizer. Aligning and averaging reduces the noise contribution by $\sqrt{N}$, leaving primarily the signal. The number of averages in this analysis for all waveforms was $N = 1024$. An example of the reduction of noise is seen in Figure A.12.

The units of the signal digitized by the oscilloscope have been nominally calibrated by comparing the peak to peak measurement of the LAB waveform and an oscilloscope measurement taken immediately before input the the SURF , however this analysis does not use this calibration for the LABRADOR ADC conversion. The normalized transfer function for the signal chain will include an implicit calibration of the ADC to voltage conversion. To match units, a scaling of 1 ADC count

per 1mV was used.

Next, the Fourier equivalent was computed for each waveform, which was then divided to determine the transfer function for the system. A band pass filter was placed on the resulting waveform, from 180MHz to 1.2GHz, in order to eliminate any power from regions of the spectrum not of interest.

A.3.4 Results for Signal Chain Transfer Function

The results of the analysis, separated into Vertical and Horizontally polarized channels, are presented in Figures A.13 through A.16. Of note are channels 05TH and 13TH. 05TH is the ALFA, a the low frequency drop down antenna that was heterodyned with a 900MHz LO and multiplexed into a digitizer. 13TH has a response slightly different than the others, mainly a higher gain at high frequencies. This issue persisted from the ANITA3 flight to the ANITA4 flight, and the cause is unknown.

A.4 Full Instrument Response

The final goal of determining the complex transfer functions for the antenna and signal chains is to relate measured ADC counts stored by the instrument into an incident electric field present on the payload. The electric field is the true physics quantity being measured by the instrument, and a precise absolute value is desired. Simulations of extensive air showers produced by terrestrial high energy particle interactions can generate electric field vector fields, which can be compared against our measurements only if we fully understand the full instrument response.

A.4.1 Convolution of antenna and signal chain

The convolution of the antenna normalized complex height and the signal chain complex transfer function yield the following equation:

$$V_{meas}(t) = \sqrt{\frac{50\Omega}{377\Omega}}[h_{sig}(t) \circ h_{ant}(t) \circ E_{inc}(t)] \tag{A.10}$$

This can be simplified by defining a global complex instrument response, $h_{inst}(t)$:

$$\begin{aligned} h_{inst}(t) &= \sqrt{\frac{Z_c}{Z_o}} h_{sig}(t) \circ h_{ant}(t) \\ V_{meas}(t) &= h_{inst}(t) \circ E_{inc}(t) \end{aligned} \tag{A.11}$$

Solving for the instrument response yields the results shown in Figures A.17 through A.20. As the signal chain response is a unitless ratio of powers per frequency bin, the full instrument response

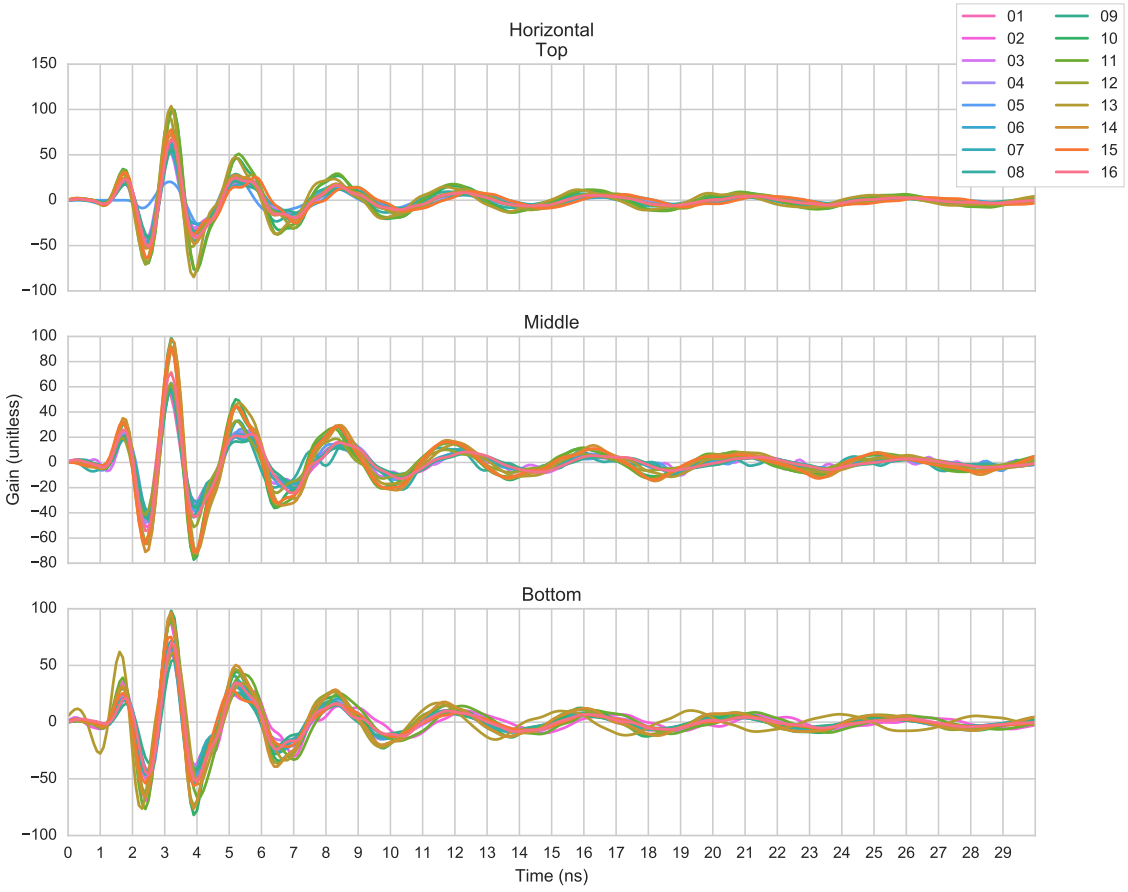

Figure A.13: Time domain impulse response of the horizontally oriented polarizations of the ANITA3 signal chain. Note 05TH, which is the ALFA antenna, and 13BH, which has been noted has an asimilar impulse response in ANITA3 and ANITA4

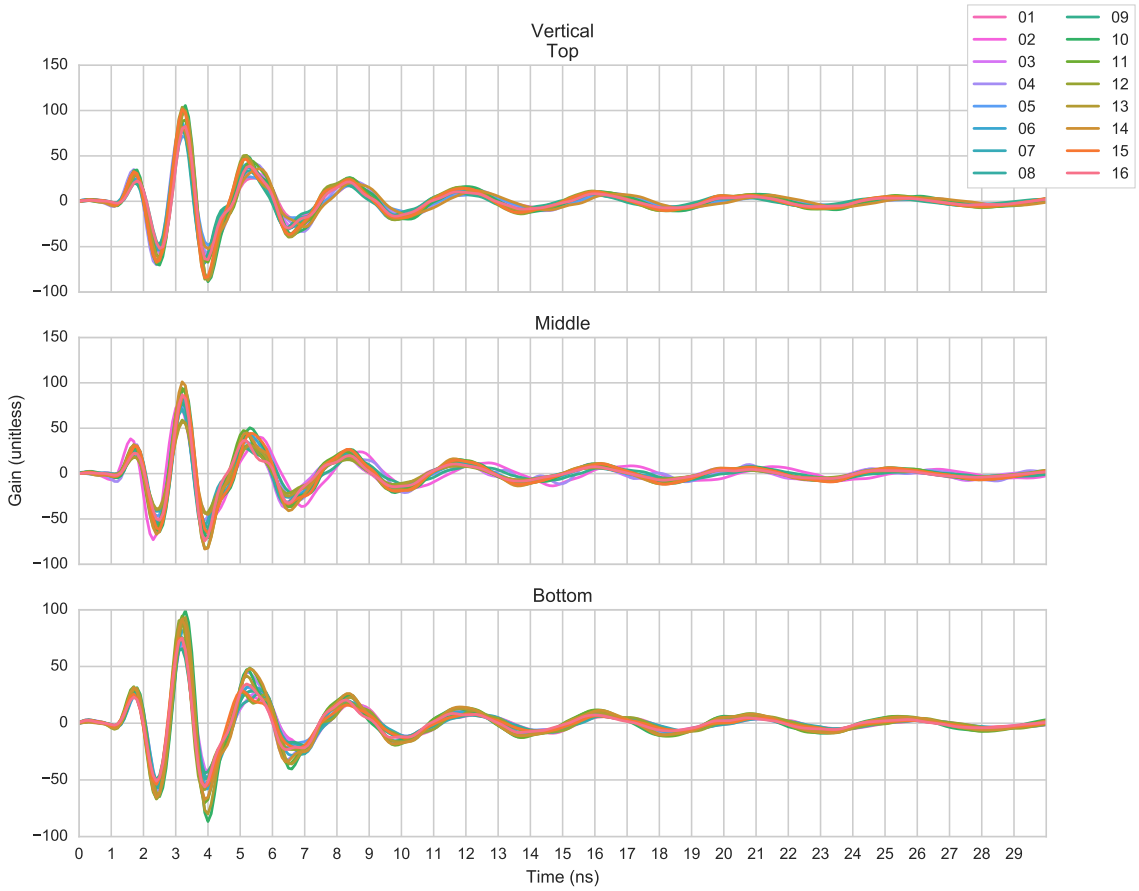

Figure A.14: Time domain impulse response of the vertically oriented polarizations of the ANITA3 signal chain.

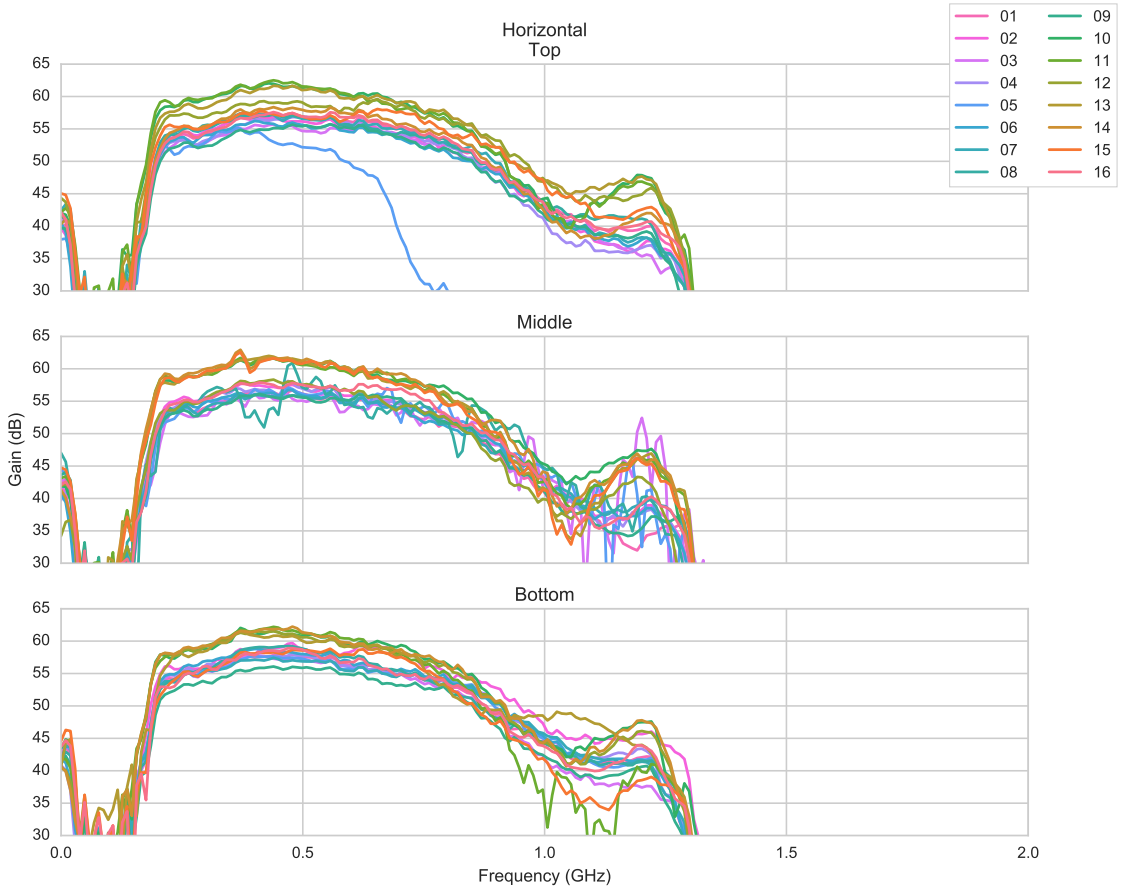

Figure A.15: Gain spectrum for the horizontally oriented polarizations of the ANITA3 signal chain. Note 05TH, which is the ALFA antenna, and 13BH, which has been noted has an asimilar impulse response in ANITA3 and ANITA4

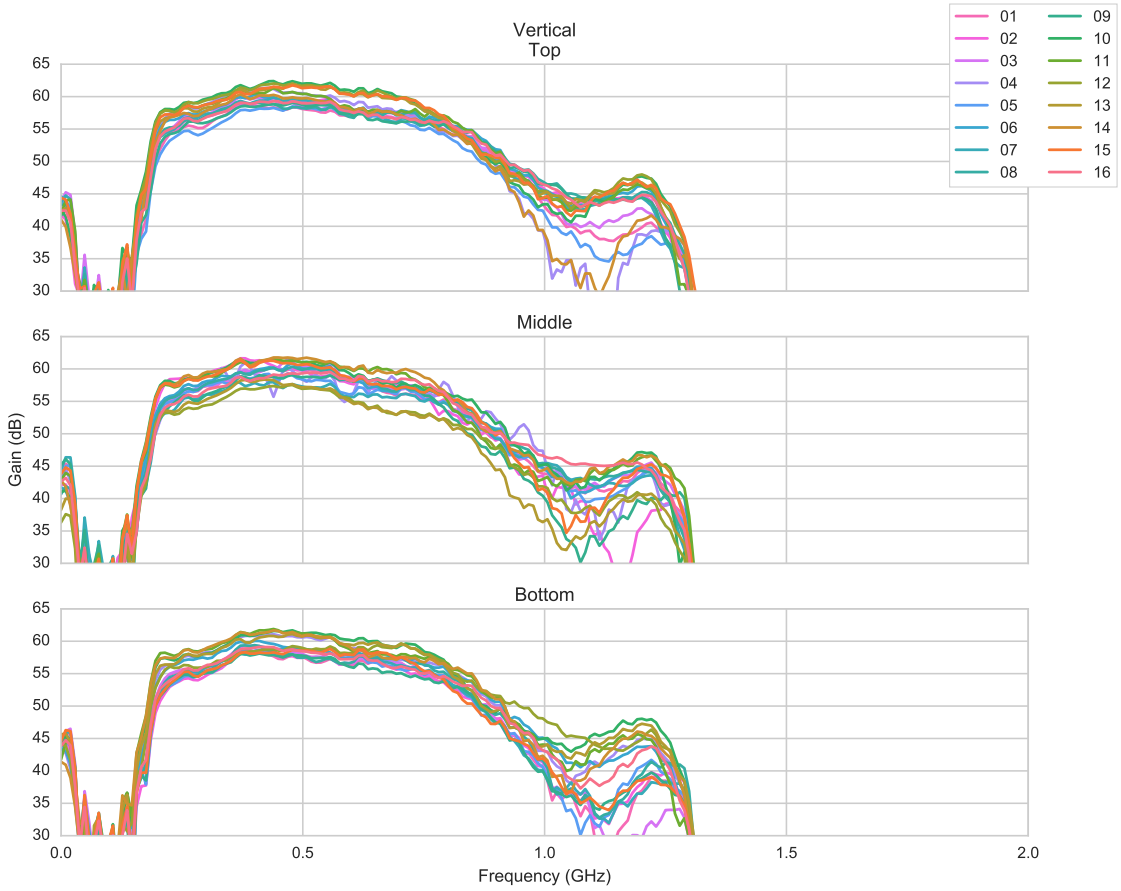

Figure A.16: Gain spectrum for the vertically oriented polarizations of the ANITA3 signal chain.

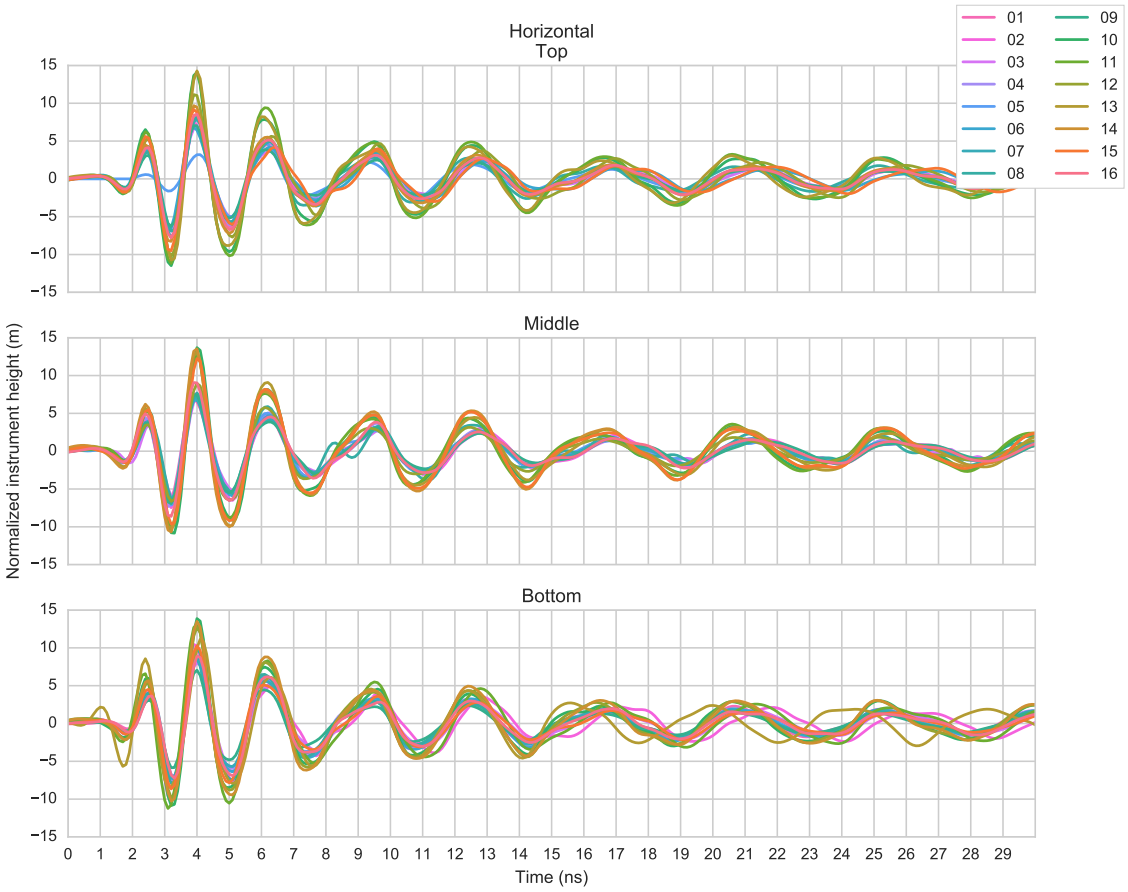

Figure A.17: Time domain impulse response of the horizontally oriented channels of the ANITA3 signal chain. Note 05TH, which is the ALFA antenna, and 13BH, which has been noted has an asimilar impulse response in ANITA3 and ANITA4

takes on the units of the antenna height. As such, I've used the naming convention normalized instrument height to describe it.

A.5 Conclusion

The ANITA antenna and signal chain complex phasor responses have an overwhelming effect on the shape and magnitude of the measured air shower signal incident on the payload. Using the processes and calibration data detailed above, that response is constrained, and signals measured by the instrument can be related back to the fundamental physics processes by which they were created.

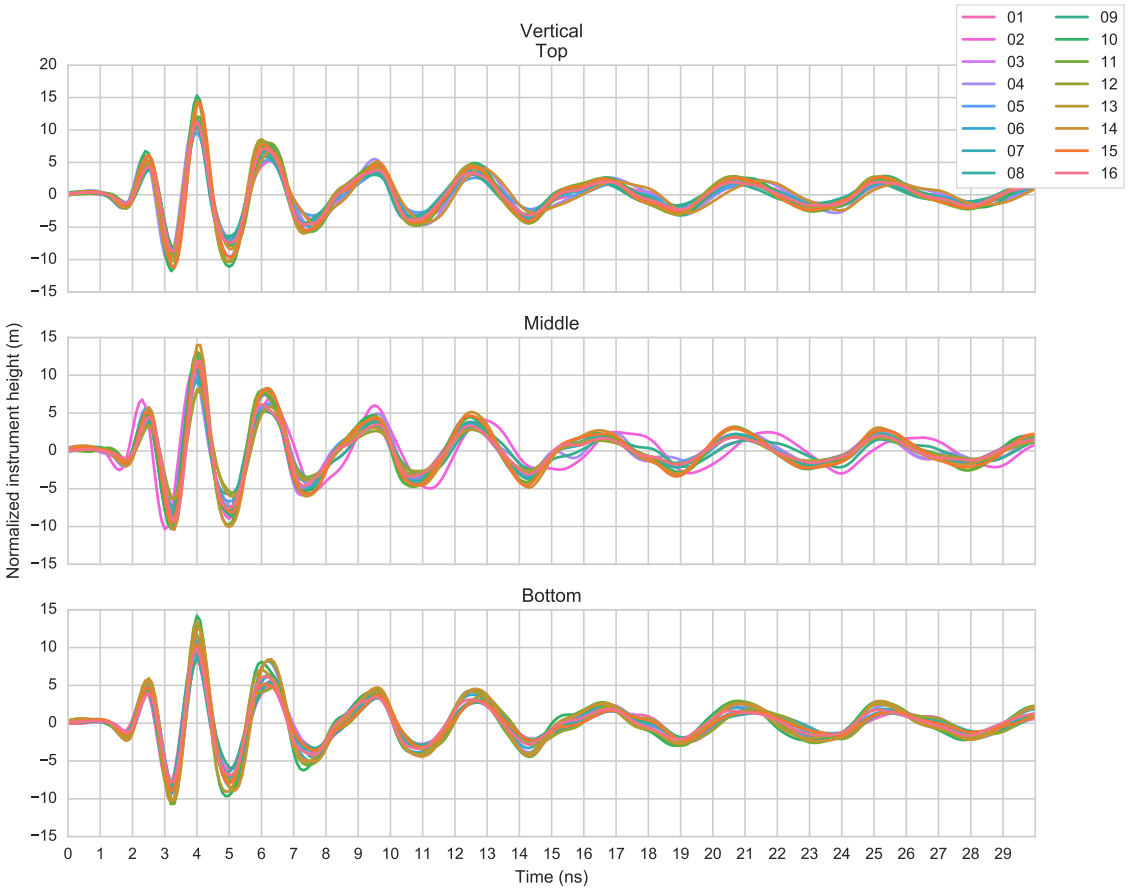

Figure A.18: Time domain impulse response of the vertically oriented channels of the ANITA3 signal chain.

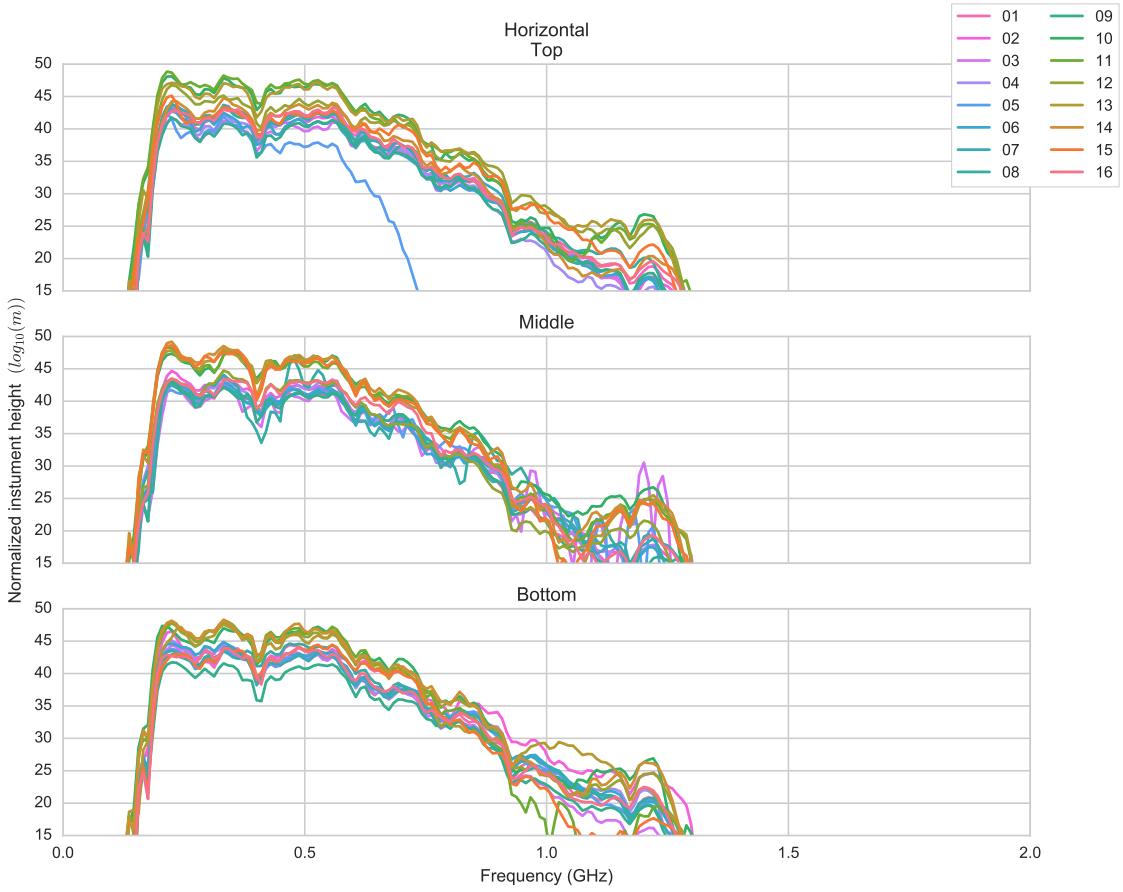

Figure A.19: Gain spectrum for the horizontally oriented channels of the ANITA3 signal chain. Note 05TH, which is the ALFA antenna, and 13BH, which has been noted has an asimilar impulse response in ANITA3 and ANITA4

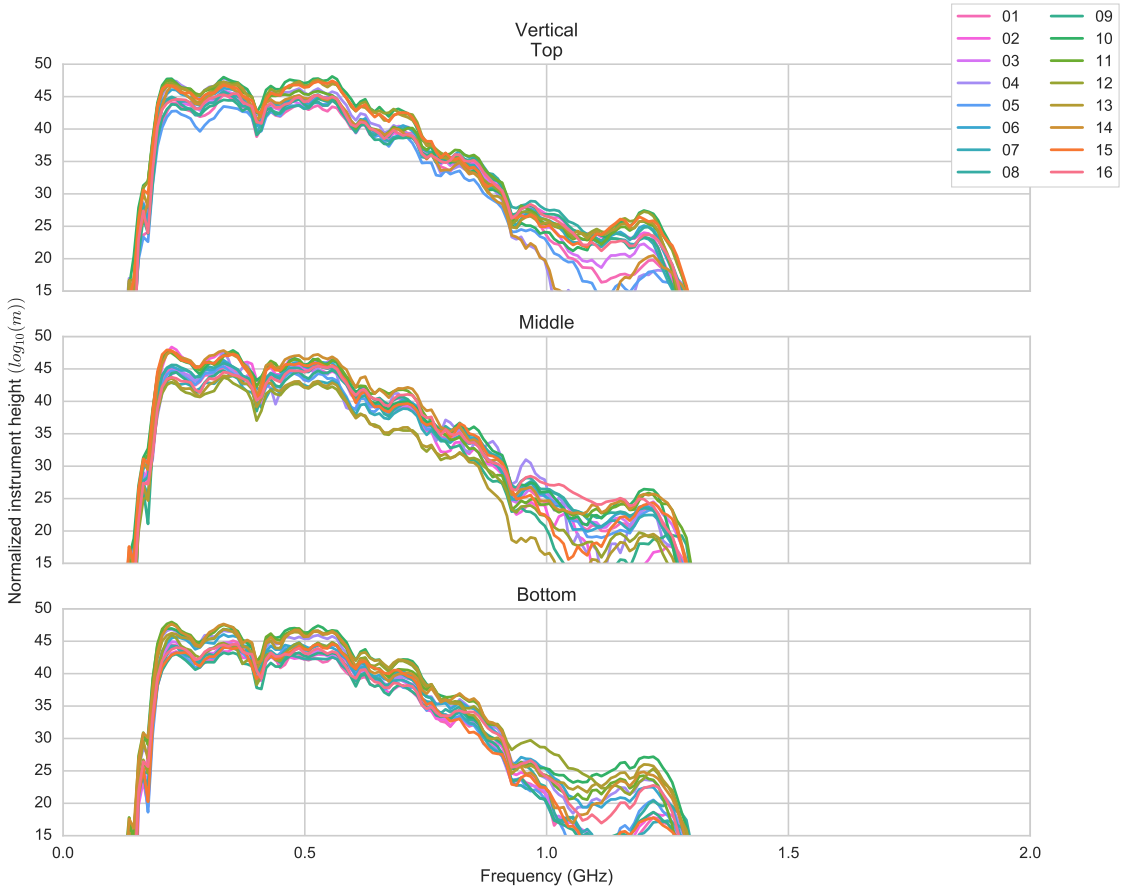

Figure A.20: Gain spectrum for the vertically oriented channels of the ANITA3 signal chain.

BIBLIOGRAPHY

[1] Alessandro De Angelis Wikimedia Commons. Cosmic rays: Increase of ionization with height as measured by hess in 1912 and by kolhrster in 1913 and 1914. `https://commons.wikimedia.org/wiki/File%3AHessKol.jpg` File: HessKol.jpg, 2011. Last visited on April 1st 2017.

[2] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.

[3] George Field. Astrophysics. *Rev. Mod. Phys.*, 71:S33–S40, Mar 1999.

[4] Jaime Alvarez-Muñiz, Andrés Romero-Wolf, and Enrique Zas. Practical and accurate calculations of askaryan radiation. *Phys. Rev. D*, 84:103003, Nov 2011.

[5] P. W. Gorham, S. W. Barwick, J. J. Beatty, D. Z. Besson, W. R. Binns, C. Chen, P. Chen, J. M. Clem, A. Connolly, P. F. Dowkontt, M. A. DuVernois, R. C. Field, D. Goldstein, A. Goodhue, C. Hast, C. L. Hebert, S. Hoover, M. H. Israel, J. Kowalski, J. G. Learned, K. M. Liewer, J. T. Link, E. Lusczek, S. Matsuno, B. Mercurio, C. Miki, P. Miočinović, J. Nam, C. J. Naudet, J. Ng, R. Nichol, K. Palladino, K. Reil, A. Romero-Wolf, M. Rosen, L. Ruckman, D. Saltzberg, D. Seckel, G. S. Varner, D. Walz, and F. Wu. Observations of the askaryan effect in ice. *Phys. Rev. Lett.*, 99:171101, Oct 2007.

[6] Jaime Alvarez-Muiz, Washington R. Carvalho Jr., and Enrique Zas. Monte carlo simulations of radio pulses in atmospheric showers using {ZHAireS}. *Astroparticle Physics*, 35(6):325 – 341, 2012.

[7] H. Schoorlemmer, K. Belov, A. Romero-Wolf, D. Garca-Fernndez, V. Bugaev, S.A. Wissel, P. Allison, J. Alvarez-Muiz, S.W. Barwick, J.J. Beatty, D.Z. Besson, W.R. Binns, W.R. Carvalho Jr., C. Chen, P. Chen, J.M. Clem, A. Connolly, P.F. Dowkontt, M.A. DuVernois, R.C. Field, D. Goldstein, P.W. Gorham, C. Hast, T. Huege, C.L. Heber, S. Hoover, M.H. Israel, A. Javaid, J. Kowalski, J. Lam, J.G. Learned, J.T. Link, E. Lusczek, S. Matsuno, B.C. Mercurio, C. Miki, P. Mioinovi, K. Mulrey, J. Nam, C.J. Naudet, J. Ng, R.J. Nichol, K. Palladino, B.F. Rauch, J. Roberts, K. Reil, B. Rotter, M. Rosen, L. Ruckman, D. Saltzberg, D. Seckel, D. Urdaneta, G.S. Varner, A.G. Vieregg, D. Walz, F. Wu, and E. Zas. Energy and flux measurements of ultra-high energy cosmic rays observed during the first {ANITA} flight. *Astroparticle Physics*, 77:32 – 43, 2016.

[8] G.S. Varner, L.L. Ruckman, J.W. Nam, R.J. Nichol, J. Cao, P.W. Gorham, and M. Wilcox. The large analog bandwidth recorder and digitizer with ordered readout (labrador) {ASIC}. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 583(23):447 – 460, 2007.

[9] Ben Strutt. *A Search for Ultra-High Energy Neutrinos and Cosmic Rays with ANITA-3*. PhD thesis, University College London, 2017.

[10] Wissel S, Saltzberg D, Belov K, Peck A, Borch K, and Fan A. Ground calibration for anita-iii. Internal Review for the FY14 Antarctic LDB Campaign Project Initiation Conference, 2013.

[11] Agilent Technologies. Noise figure measurement accuracy the y-factor method, 2010.

[12] Dan Moulton Wikimedia Commons. Stokes parameters in some degenerate cases. `https://commons.wikimedia.org/w/index.php?curid=3319458`. Last visited on December 1st 2017.

[13] Jaime Alvarez-Muiz, Washington R. Carvalho Jr, Kvin Payet, Andrs Romero-Wolf, Harm Schoorlemmer, and Enrique Zas. A comprehensive approach to tau-lepton production by high-energy tau neutrinos propagating through earth, 2017.

[14] C. Miki. private communication.

[15] Alessandro De Angelis. Spontaneous ionization to subatomic physics: Victor hess to peter higgs. *Nuclear Physics B - Proceedings Supplements*, 243:3 – 11, 2013.

[16] J. W. MOTZ, HAAKON A. OLSEN, and H. W. KOCH. Pair by photons. *Rev. Mod. Phys.*, 41:581–639, Oct 1969.

[17] IceCube collaboration. Measurement of the cosmic ray energy spectrum with icetop-73. *Phys. Rev. D*, 2013.

[18] Roberto Aloisio. Acceleration and propagation of ultra high energy cosmic rays. *PTEP*, 2017, 2017.

[19] J. Abraham and The Pierre Auger Collaboration. Measurement of the depth of maximum of extensive air showers above 10^{18}ev. *Phys. Rev. Lett.*, 104:091101, Mar 2010.

[20] P. Blasi R. Aloisio, V. Berezinsky. Ultra high energy cosmic rays: implications of auger data for source spectra and chemical composition. *Journal of Cosmology and Astroparticle Physics*, 1410, 2014.

[21] Valerio Verzi, Dmitri Ivanov, and Yoshiki Tsunesada. Measurement of energy spectrum of ultra-high energy cosmic rays, 2017.

[22] et al. C. L. Bennett. Nine-year wilkinson microwave anisotropy probe (wmap) observations: Final maps and results. *The Astrophysical Journal Supplement Series*, 208(2), Sept 2013.

[23] Kenneth Greisen. End to the cosmic-ray spectrum? *Phys. Rev. Lett.*, 16:748–750, Apr 1966.

[24] High Resolution Flys Eye Collaboration. First observation of the greisen-zatsepin-kuzmin suppression. *Phys. Rev. Lett.*, 100:101101, Mar 2008.

[25] J. A. Formaggio and G. P. Zeller. From ev to eev: Neutrino cross sections across energy scales. *Rev. Mod. Phys.*, 84:1307–1341, Sep 2012.

[26] Raj Gandhi, Chris Quigg, Mary Hall Reno, and Ina Sarcevic. Ultrahigh-energy neutrino interactions. *Astroparticle Physics*, 5(2):81 – 110, 1996.

[27] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968.

[28] D. Besson, R. Keast, and R. Velasco. In situ and laboratory studies of radiofrequency propagation through ice and implications for siting a large-scale antarctic neutrino detector. *Astroparticle Physics*, 31(5):348 – 358, 2009.

[29] Edward P. Tilton. *The radio amateur's V. H. F. manual.* American Radio Relay League, Newington, Co., 3rd ed. edition, 1972.

[30] Taylor Barrella, Steven Barwick, and David Saltzberg. Ross Ice Shelf in situ radio-frequency ice attenuation. *J. Glaciol.*, 57:61–66, 2011.

[31] J H Hough. Calculations on the radio emission resulting from geomagnetic charge separation in an extensive air shower. *Journal of Physics A: Mathematical, Nuclear and General*, 6(6):892, 1973.

[32] K. Belov, K. Mulrey, A. Romero-Wolf, S. A. Wissel, A. Zilles, K. Bechtol, K. Borch, P. Chen, J. Clem, P. W. Gorham, C. Hast, T. Huege, R. Hyneman, K. Jobe, K. Kuwatani, J. Lam, T. C. Liu, J. Nam, C. Naudet, R. J. Nichol, B. F. Rauch, B. Rotter, D. Saltzberg, H. Schoorlemmer, D. Seckel, B. Strutt, A. G. Vieregg, and C. Williams. Accelerator measurements of magnetically induced radio emission from particle cascades with applications to cosmic-ray air showers. *Phys. Rev. Lett.*, 116:141103, Apr 2016.

[33] G. A. Askar'yan. Excess negative charge of an electron-photon shower and its coherent radio emission. *Sov. Phys. JETP*, 14(2):441–443, 1962. [Zh. Eksp. Teor. Fiz.41,616(1961)].

[34] P. W. Gorham, D. Saltzberg, R. C. Field, E. Guillian, R. Milinčić, P. Miočinović, D. Walz, and D. Williams. Accelerator measurements of the askaryan effect in rock salt: A roadmap toward teraton underground neutrino detectors. *Phys. Rev. D*, 72:023002, Jul 2005.

[35] David Saltzberg, Peter Gorham, Dieter Walz, Clive Field, Richard Iverson, Allen Odian, George Resch, Paul Schoessow, and Dawn Williams. Observation of the askaryan effect: Coherent microwave cherenkov emission from charge asymmetry in high-energy particle cascades. *Phys. Rev. Lett.*, 86:2802–2805, Mar 2001.

[36] Edgar Bugaev, Teresa Montaruli, Yuri Shlepin, and Igor Sokalski. Propagation of -neutrinos and -leptons through the earth and their detection in underwater/ice neutrino telescopes. *Astroparticle Physics*, 21(5):491 – 509, 2004.

[37] X Bertou, P Billoir, O Deligny, C Lachaud, and A Letessier-Selvon. Tau neutrinos in the auger observatory: a new window to uhecr sources. *Astroparticle Physics*, 17(2):183 – 193, 2002.

[38] The IceCube Collaboration. Search for ultrahigh-energy tau neutrinos with icecube. *Phys. Rev. D*, 86:022005, Jul 2012.

[39] S. Hoover et al. Observation of Ultra-high-energy Cosmic Rays with the ANITA Balloon-borne Radio Interferometer. *Phys. Rev. Lett.*, 105:151101, 2010.

[40] P. W. Gorham et al. Observational Constraints on the Ultra-high Energy Cosmic Neutrino Flux from the Second Flight of the ANITA Experiment. *Phys. Rev.*, D82:022004, 2010. [Erratum: Phys. Rev.D85,049901(2012)].

[41] H. Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, April 1928.

[42] L. Esaki. New Phenomenon in Narrow Germanium p-n Junctions. *Physical Review*, 109:603–604, January 1958.

[43] Kurtis Nishimura and Andres Romero-Wolf. A correlation-based timing calibration and diagnostic technique for fast digitizing asics. *Technology and Instrumentation in Particle Physics*, 2011.

[44] P. W. Gorham et al. Antarctic Surface Reflectivity Measurements from the ANITA-3 and HiCal-1 Experiments. *J. Astron. Inst.*, 06(02):1740002, 2017.

[45] Shahid Hussain and Douglas W. McKay. Comparative study of radio pulses from simulated hadron-, electron-, and neutrino-initiated showers in ice in the gev-pev range. *Phys. Rev. D*, 70:103003, Nov 2004.

[46] Rene Brun and Fons Rademakers. ROOT - An Object Oriented Data Analysis Framework. `http://root.cern.ch/`, 2017.

[47] ANITA Collaboration. libRootFftwWrapper - A ROOT wrapper for FFTW3. `https://github.com/nichol77/libRootFftwWrapper`, 2017.

[48] ANITA Collaboration. eventReaderRoot - An ANITA Event Reader in ROOT format. `https://github.com/anitaNeutrino/AnitaAnalysisTools`, 2017.

[49] ANITA Collaboration. Anita Analysis Framework. https://github.com/anitaNeutrino/AnitaAnalysisFramework, 2017.

[50] Cosmin Deaconu. UCorrelator. https://github.com/anitaNeutrino/UCorrelator, 2017.

[51] ANITA Collaboration. anitaMagicDisplay - Marvellous Anita Graphical Interface and Classy Display. https://github.com/anitaNeutrino/anitaMagicDisplay, 2017.

[52] ANITA Collaboration. anitaEventCorrelator. https://github.com/anitaNeutrino/anitaEventCorrelator, 2017.

[53] Ben Rotter. benCode - ANITA Cosmic Ray Analysis Software. https://github.com/brotter/benCode.

[54] A. Romero-Wolf, S. Hoover, A. Vieregg, P. Gorham, and the ANITA Collaboration. An interferometric analysis method for radio impulses from ultra-high energy particle showers, 2013.

[55] Hiroshi Akima. A new method of interpolation and smooth curve fitting based on local procedures. *J. ACM*, 17(4):589–602, October 1970.

[56] H. Liu, K. C. Jezek, B. Li, and Z. Zhao. RAMPDEM: Radarsat Antarctic Mapping Project Digital Elevation Model, Version 2., 2015.

[57] O. Scholten, T. N. G. Trinh, A. Bonardi, S. Buitink, P. Correa, A. Corstanje, Q. Dorosti Hasankiadeh, H. Falcke, J. R. Hörandel, P. Mitra, K. Mulrey, A. Nelles, J. P. Rachen, L. Rossetto, P. Schellart, S. Thoudam, S. ter Veen, K. D. de Vries, and T. Winchen. Measurement of the circular polarization in radio emission from extensive air showers confirms emission mechanisms. *Phys. Rev. D*, 94:103010, Nov 2016.

[58] J. Neyman. X—outline of a theory of statistical estimation based on the classical theory of probability. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 236(767):333–380, 1937.

[59] J. D. Hague, B. R. Becker, M. S. Gold, and J. A. J. Matthews. Power laws and the cosmic ray energy spectrum. *Astropart.Phys.27:455-464*, 2006.

[60] Kurtis Nishimura, Matthew Andrew, Zhe Cao, Michael Cooney, Peter Gorham, Luca Macchiarulo, Lisa Ritter, Andres Romero-Wolf, and Gary Varner. A low-resolution, gsa/s streaming digitizer for a correlation-based trigger system, 2012.

[61] Eric W. Weisstein. Discrete fourier transform From MathWorld—A Wolfram Web Resource. http://mathworld.wolfram.com/Tree.html. Last visited on April 1st 2017.

[62] P. Miočinović, R. C. Field, P. W. Gorham, E. Guillian, R. Milinčić, D. Saltzberg, D. Walz, and D. Williams. Time-domain measurement of broadband coherent cherenkov radiation. *Phys. Rev. D*, 74:043002, Aug 2006.

[63] Avionics Department. *Electronic Warfare and Radar Systems.* Naval Air Warfare Center Weapons Division, 4 edition, 2013.