

THE UNIVERSITY OF CHICAGO

A LOW THRESHOLD NEUTRINO SEARCH WITH THE ASKARYAN RADIO ARRAY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY

KAELI AUTUMN HUGHES

CHICAGO, ILLINOIS

JUNE 2022

Copyright © 2022 by Kaeli Autumn Hughes
All Rights Reserved

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	xii
ACKNOWLEDGMENTS	xiii
ABSTRACT	xv
1 INTRODUCTION	1
1.1 Motivation	1
1.1.1 Neutrino Astronomy	2
1.1.2 Multi-Messenger Astronomy	3
1.1.3 Fundamental Physics	4
1.2 Production Mechanisms	5
1.2.1 Astrophysical Neutrinos	5
1.2.2 Cosmogenic Neutrinos	6
1.2.3 Science from a Non-Detection	8
1.3 Neutrino Detection	9
1.4 Neutrino Detector Summary	11
1.4.1 The IceCube Neutrino Observatory	12
1.4.2 ANTARES and KM3NET	13
1.4.3 EUSO-SPB2 and POEMMA	13
1.4.4 ANITA and PUEO	14
1.4.5 ARIANNA	15
1.4.6 RNO-G	16
1.4.7 BEACON	18
1.4.8 ARA	19
2 THE ASKARYAN RADIO ARRAY (ARA)	21
2.1 Station Design	22
2.2 Signal Chain	24
2.3 Current Status and Future Upgrades	25
3 THE PHASED ARRAY INSTRUMENT	27
3.1 The Phased Array Trigger Concept	27
3.2 The NuPhase Instrument at the South Pole	30
3.2.1 The RF Chain	31
3.2.2 Trigger Implementation	33
3.2.3 Matching Events Between the Phased Array and ARA5	34
3.2.4 Limitations	34

4	TIMING AND VOLTAGE CALIBRATION OF ARA STATION 5	36
4.1	The IRS2 Chip	36
4.2	Channel Mapping and Available Calibration Data	38
4.3	Pedestal Subtraction	40
4.4	Timing Calibration	42
4.4.1	Correcting the Overall Frequency	43
4.4.2	Correcting the Residual Timing for each Sample	43
4.5	Voltage Calibration	49
4.6	Validating the Timing and Voltage Calibration	53
4.7	Future Improvements	54
5	GEOMETRY CALIBRATION OF ARA STATION 5	55
5.1	Introduction	55
5.2	Determining the ARA Antenna Locations	55
5.3	Step 0: Defining the Calibration Dataset	58
5.4	Step 1: Determining the Phased Array Cable Delays and the Ice Model	59
5.5	Step 2: Determining the ARA Antenna Locations	64
5.6	Step 3: Use the common channel to find the location of the Phased Array relative to the ARA antennas	69
5.7	Event Reconstruction	71
5.8	Lessons Learned	77
6	THE PHASED ARRAY ANALYSIS	80
6.1	Simulation	80
6.2	RZ Correlation Maps and Analysis Variables	82
6.3	Continuous Wave (CW) Background Events	89
6.4	Livetime Cuts	93
6.4.1	Anthropogenic Events	93
6.4.2	Calibration Pulser Switches	95
6.5	Defining a Deep Region	96
6.6	Deep Region Analysis	98
6.7	Deep Region Results	103
6.8	Surface Region Results	106
6.9	Future Work	108
7	THE BEAMFORMING ELEVATED ARRAY FOR COSMIC NEUTRINOS (BEA- CON)	109
7.1	Motivation	109
7.2	White Mountain Site Study	110
7.3	Interferometry Studies	112
7.3.1	A Two-Antenna Short Term Interferometer	112
7.3.2	The BEACON Prototype: A Four-Antenna, Year-Long Interferometer	113
7.4	Understanding Local RFI	115
7.4.1	High-Voltage Power Lines	115

7.4.2	Galactic Noise	115
7.4.3	Airplane Tracks	116
7.5	Work Since 2018	117
	REFERENCES	119

LIST OF FIGURES

1.1	A plot showing the distances and energies for which the Universe becomes opaque to photons. Neutrinos and gravitational waves are the only way to explore the distant, highest energy Universe. From [1].	2
1.2	A plot comparing the flux of various messenger particles: γ -rays (left), neutrinos (middle) and cosmic rays (right). Also shown are various models for astrophysical neutrino production. From [2].	3
1.3	A calculation of the neutrino-nucleon charged-current cross section, calculated from the IceCube 6-year dataset. The models match the data within uncertainties up to 10^6 GeV, after which no neutrinos have been detected. From [3].	5
1.4	A measurement of the cosmic ray flux as a function of energy for two experiments, Auger and the Telescope Array. The discrepancies in the spectra at highest energies are unexplained and appear to be more than statistical. The cutoff visible near $10^{19.5}$ eV is due to the GZK effect. From the PDG.	7
1.5	A plot showing the electric field emitted at various angles away from the Cherenkov angle (left), and the associated differences in frequency content (right). From [2].	10
1.6	A plot of the current best limits on the neutrino flux. Included are limits from IceCube [4] [5], Auger [6], ARA [7], ARIANNA [8], and ANITA [9]. Also shown are theoretical neutrino production models. From [7].	11
1.7	A diagram of the IceCube detector, with the Eiffel Tower for scale. From icecube.wisc.edu.	12
1.8	A diagram showing the Askaryan emission channel as seen by PUEO.	14
1.9	A diagram showing the types of interactions an ARIANNA station is capable of detecting. From [8].	16
1.10	A diagram of an RNO-G station, including both surface and deep stations. From [2].	17
1.11	A diagram showing the concept of a BEACON array, including the geometry of detecting a decaying tau lepton.	18
1.12	A diagram showing the geometry of a neutrino interaction seen by the ARA experiment. From [7].	19
2.1	A diagram of a typical ARA Station, including close-up views of the VPol and HPol antennas.	21
2.2	A picture of the HPol antenna, along with its simulated gain in dBi. From [10].	23
2.3	A diagram of the signal chain of the ARA instrument. From B. Clark [11] . . .	24
3.1	An example of a one-dimensional vertically-spaced phased array trigger array (center), surrounded by additional antennas used for reconstruction. From [12].	28
3.2	A diagram of ARA Station 5, with the traditional ARA antennas in grey and the Phased Array antennas in red.	30
3.3	An illustration showing the RF chain for the nuphase instrument at the South Pole. From [12].	31

3.4	The in-situ efficiency of the Phased Array instrument, compared to the efficiency of the classic ARA trigger [12].	33
3.5	The expected beam widths for each of the beams in the South Pole Phased Array instrument. From [12].	33
4.1	An example of a Switch Capacitor Array. From [13].	36
4.2	Left: the average ADC offset for a specific sample over time. Right: the corrected mean value after subtracting the automatically-generated pedestal values. The red lines correspond to times when the pedestal value was recalculated.	40
4.3	An example of two blocks of calibration data, with the targeted region identified.	41
4.4	The distribution of frequency fits for all calibration events. The clock speed is scaled so that the mean value matches the true frequency of 218 MHz.	42
4.5	Uncalibrated and calibrated residual timing for electric channel 0. Before calibration, the samples are offset from the sine fit, causing a systematic offset every time that sample is measured. After calibration, the offset is approximately 0 ns.	44
4.6	Uncalibrated and calibrated residual timing for electric channel 1.	44
4.7	Uncalibrated and calibrated residual timing for electric channel 2. The even samples are discarded due to poor behavior.	45
4.8	Uncalibrated and calibrated residual timing for electric channel 3. The even samples are discarded due to poor behavior.	45
4.9	An example of extreme bimodality for electric channel one. Left: the calibrated residual timing for channel 1 as a function of sample number. Right: the calibrated residual timing as a function of the measured slope. The split on the left hand side is caused by the slope shown on the right.	46
4.10	The time between odd and even samples, for channel 0 (top left), channel 1 (top right), channel 2 (bottom left), and channel 3 (bottom right). The nominal spacing is 0.625 ns.	47
4.11	Uncalibrated data vs. calibrated timing data, for channel 0 and channel 1. The blue points fit especially well near 0 ADC, where the uncalibrated voltage has the smallest effect.	48
4.12	Example odd and even samples, comparing the ADC counts measured and the predicted voltage from the sine fit. In each example, the positive and negative points are fit with a cubic polynomial.	49
4.13	(a) An example sample with looping behavior. (b) The same sample, only plotting the positively-sloped data. (c) The same sample, only plotting the negatively sloped data.	50
4.14	The distribution of χ^2 for all 32,768 samples. Any fits with $\chi^2 < 1$ are kept; fits worse than this are substituted with the fit from the neighboring sample.	51
4.15	A comparison of timing-calibrated data and timing and voltage-calibrated data. Especially notable are the data points at high ADC values, which have been adjusted to proper voltage values.	52
4.16	The timing delays between calibration pulser events comparing channels 0 and 16.	53

5.1	Left: the relationship between the fractional wavelength offset and the sum between two antennas. The horizontal line indicates the mean incoherent addition. Right: The two antenna sum as a function of frequency and displacement. Areas with a sum greater than the mean incoherent addition value will add coherently. From C. Deaconu.	55
5.2	A map of places of interest near ARA Station 5.	57
5.3	A comparison of the raw calculated time delays from the Phased Array instrument, compared to the predicted time delays from the AraSim Raytracer. The mismatch is due to cable delays not yet being applied to the raw data.	60
5.4	Left: comparison of the various ice model fits, after correcting for the cable delay offset. Most of the models do not accurately predict the shape of the curve as a function of depth. Right: same as the left, after fitting the third parameter of the UNL 2016 model to better match the data.	61
5.5	The measured time delays against the predicted time delays using the modified UNL 2016 model.	62
5.6	A comparison of the predicted time delays after moving the SPIceCore hole by 40 m. The horizontal offset shifts by less than 0.2 ns for the longest baseline, and the shape of the curve stays the same.	63
5.7	Averaged calibration waveforms for each of the eight Vpol antennas.	65
5.8	The distribution of time delays across VPol channels for a given set of calibration pulser events. Channel 24 was not calibrated due to missing calibration data, resulting in its distributions being wider than the others.	66
5.9	An example of the changing time delays between Electric Channels 0 and 1 for the ARA instrument. The red points represent the interpolated values used in the Minit fit. The outliers are caused by the cross correlation being maximized between the wrong peaks.	67
5.10	The calibrated locations of the A5 antenna boreholes, compared to the planned locations.	69
5.11	The distribution of time delays across VPol channels for a given set of calibration pulser events. Channel 24 was not calibrated due to missing calibration data, resulting in its distributions being wider than the others.	70
5.12	An example low SNR event from the local calibration pulser, as seen by the ARA instrument.	72
5.13	An example low SNR event from the local calibration pulser, as seen by the Phased Array instrument.	72
5.14	Left: a correlation map for a high SNR local calibration pulse, using only the classic ARA antennas and the Hilbert envelope of the correlation waveforms. Right: the same, except for a low SNR local calibration pulse.	74
5.15	Left: a correlation map for a high SNR local calibration pulse, using the Phased Array and ARA antennas, generated with the Hilbert envelope of the correlation waveforms. Right: the same, except for a low SNR local calibration pulse.	74

5.16	Left: a correlation map for a high SNR local calibration pulse, using only the classic ARA antennas and the full correlation waveforms. Right: the same, except for a low SNR local calibration pulse.	75
5.17	Left: a correlation map for a high SNR local calibration pulse, using the Phased Array and ARA antennas, generated with the full correlation waveforms. Right: the same, except for a low SNR local calibration pulse.	75
5.18	Left: the reconstructed azimuth and zenith using only classic ARA baselines, and Hilbert enveloped-correlation waveforms. Right: the same, without applying a Hilbert envelope.	76
5.19	Left: the reconstructed azimuth and zenith using all available ARA baselines, and Hilbert enveloped-correlation waveforms. Right: the same, without applying a Hilbert envelope.	76
5.20	A plot comparing the resolution in both azimuth and zenith for various combinations of baselines.	77
6.1	The in-situ efficiency of the Phased Array instrument, compared to the efficiency of the classic ARA trigger [12].	81
6.2	The effective volume of the old Phased Array simulation compared to the new effective volume simulation.	81
6.3	Example correlation maps for four types of events. Upper left: the local calibration pulser, with a blue dot for the actual location. Upper right: the rooftop calibration pulser, located approximately 4 km away at the surface. Lower left: a thermal noise trigger. Lower right: a simulated double pulse neutrino event, with the blue dot indicating the true location.	84
6.4	Left: An example low SNR calibration pulser. Top right: The correlation map for the low SNR calibration pulser. Bottom right: The coherently averaged waveform and the de-dispersed coherently averaged waveform (CSW).	85
6.5	Left: The CSW of a local calibration pulser in blue, with the Hilbert envelope in green. Right: The CDF of the Hilbert envelope starting from the peak. The average value is marked with a horizontal line.	87
6.6	Left: The CSW of a thermal noise event in blue, with the Hilbert envelope in green. Right: The CDF of the Hilbert envelope starting from the peak. The average value is marked with a horizontal line.	87
6.7	Before and after applying the Sine Subtraction method to a channel with 405 MHz CW.	89
6.8	All of the frequencies removed during the 24-hour unblinded period, centered around 405 MHz.	89
6.9	The Maximum Correlation as a function of time before applying the Sine Subtract method (left) and after applying it (right).	90
6.10	An example of an event with CW Contamination from satellite communications. Typically thermal noise events have a maximum correlation value below 0.10; here the maximum is 0.12.	91
6.11	A comparison of the events with CW contamination at 137 MHz (Top) and the beam 4 threshold pattern (Bottom).	92

6.12	A plot of outlier events in the sample during a time with a lot of known station activity (October 2019). Four runs account for nearly all of the outlier events. . .	93
6.13	A plot of the surface correlation values for thermal noise events in the 10% sample. The red line indicates an exponential fit.	94
6.14	A plot of the calibration pulser rate for each run. Most of the drops in event rate correspond to know mode switches or alternative calibration pulser configurations.	95
6.15	The population of simulated neutrinos (left) and the 10% sample (right, middle) plotted with the two variables defining the deep region: Zenith Angle and Surface Correlation Ratio. The color axis on the left and middle plots are event counts; the color axis on the right is the maximum correlation in each bin.	97
6.16	The remaining 10% sample (left) and the simulated neutrino set (middle, right) plotted with two of the analysis variables: SNR and maximum correlation. . . .	99
6.17	An example of a dataset being split by a Fisher Discriminant using the Scikit-learn package. From [14].	100
6.18	Left: A comparison between the 10% sample distribution and the simulated neutrino flux, with the optimized threshold marked with a vertical line. Right: the tail of the 10% sample, with the exponential fit overlaid. The value of X_{min} is defined as the minimum value of the Fisher Discriminant in the tail of the 10% sample.	101
6.19	A plot of the estimated background distributions as a function of cut value. As the cut value decreases, the background distribution gets wider.	102
6.20	Top: the correlation map of the passing event. Bottom left: the Coherently Summed Waveform of the passing event, compared to a simulated event at the same incoming angle. The noise in the simulated waveform is normalized to have the same root-mean-square as the average from data. Bottom right: a table showing some of the values of analysis variables for the passing event.	104
6.21	The limit from this Phased Array analysis, using one station and six months of livetime, along with the projected sensitivity from the entire available livetime of 2.6 years. The expected sensitivity is calculated by directly scaling the livetime and expected backgrounds, selecting the median expected background, and calculating the 90% upper limit. Plotted for comparison are the previous published results from ARA (two stations, each with 4.0 station-years), along with results from ANITA [9], ARIANNA (seven stations, each with 4.5 station-years) [8], Auger [6] and IceCube[4] [5]. Additionally, theoretical production models of cosmogenic neutrinos are plotted [15] [16] [17].	105
6.22	Analysis efficiency of this analysis compared to the most recent ARA Station 2 analysis efficiency [7]. The blue curve is the efficiency on neutrinos in the deep region; the purple curve includes the efficiency loss from removing the surface region which results in a 21% loss in efficiency averaged over all energies. The analysis efficiency shown in the plot on the left is generated for each energy bin separately. The analysis efficiency shown in the plot on the right is generated based on a cosmogenic flux [18].The trigger efficiency vs. SNR is from [12], while the analysis efficiency from ARA Station 2 is adapted from [7].	106

7.1	Spectra and Rates from White Mountain Site Study.	110
7.2	A map of the source directions as found by the single baseline interferometer. Lines to the left of center represent data taken while the interferometer was placed overlooking the West Valley, and lines to the right are from the East. . .	111
7.3	View of BEACON prototype from above the mountain ridge. Antennas were camouflaged to minimize environmental impact.	113
7.4	Approximate beam directions of the BEACON prototype. Here due East corresponds to 0 azimuth.	113
7.5	A spectrogram showing the changing frequency content over the course of 100 minutes.	114
7.9	Time delays for each baseline as a function of readout time, using vertically-polarized channels. The sloped nature of the lines is evidence that the source is moving over the course of 25 seconds.	116
7.6	A comparison of the thresholds of each beam formed with the BEACON prototype system. One σ on this plot corresponds to the average power over 16 samples of a typical noise trigger.	118
7.7	Example of the 60 Hz structure that appears during data taking.	118
7.8	Evidence that the BEACON prototype is sensitive to radio from the galactic center. Here the altitude of Sgr A* is compared to the median RMS of noise trigger events after undergoing a low pass filter.	118

LIST OF TABLES

4.1	The channel mapping of all channels in the A5 baseline instrument. The Electric Channel uniquely describes a specific antenna and channel.	38
4.2	A list of the available calibration data taken for ARA Station 5 by Thomas Meures.	39
5.1	The list of pre-deployment locations of the A5 boreholes.	59
5.2	A list of previously-derived ice models for the South Pole.	62
5.3	The locations of each of the channels in the ARA Station 5 experiment. Note that channels PA 5 and ARA 24 refer to the same physical antenna. The Phased Array cable delays are reported relative to the Phased Array channel 0.	79
6.1	Table of analysis variables used in the Fisher Discriminant, along with their definitions and expected ranges where applicable.	88
6.2	The list of runs that were removed.	96
6.3	Table of cuts, background estimates, and analysis efficiencies from the 10% sample in the deep region. For the Fisher Discriminant, the median background is reported, as are the 16th and 84th percentiles.	103

ACKNOWLEDGMENTS

I would first and foremost like to thank my advisor, Prof. Abby Vieregk, for the invaluable support, advice, and guidance over the last five years. Thank you to Cosmin Deaconu for all of the time spent answering my questions, the impromptu office chats, and the ridiculously helpful debugging when I had problems with the Midway computer. Thank you to Eric Oberla for showing me the ropes during field work both at the South Pole and at White Mountains, and for answering all of my hardware questions. Thank you to all of my fellow graduate students in the Vieregk Lab, past and present, but especially to Andrew Ludwig, Dan Smith, and Dan Southall, for all of the coffee breaks at the Div School, for answering my most basic questions, and for just generally being fun to work with.

Thank you to the entire ARA Collaboration for their advice and input on my calibration and analysis. Thank you to Abby, Cosmin, Amy Connolly, Dave Besson, and Brian Clark for agreeing to another biweekly call to give input on my analysis; your input was extremely valuable to me and helped make the analysis stronger. I'd like to specifically thank Amy Connolly for taking me on as a 19-year-old freshman at Ohio State and investing so much time and energy into helping me succeed. Thank you to Ilya Kravchenko and Dave Seckel for all of your input as part of my internal paper committee.

Thank you to the BEACON Collaboration for so many opportunities to conduct field work, and thank you for putting up with my old iTunes songs for a five hour car ride through the desert after we got a flat tire.

Thank you to the UChicago Physics community for their support throughout my entire Ph.D. I'd like to specifically thank Young-Kee for asking me for honest feedback and listening to my answers. Thank you to the NSF GRFP and the Grainger Fellowship for funding.

Thank you to all of my friends in the 2017 UChicago cohort, without whom I would not have completed the first year and who made Chicago feel like home. Thank you for the pub trivia nights, PSD Happy Hour, Sunday brunches, Dungeons and Dragons, scone parties,

and karaoke celebrations. Special thanks to Tracy, Emily, and Claire for always offering your apartment for group gatherings those first few years.

Thank you to my cats, Leia and Westley, for their input on this thesis by helpfully sitting on my keyboard. Thank you to my parents for only freaking out a little bit when I said I was going to the South Pole for two months. Thank you to my mom for all the trips to the library to check out books about the planets when I was in elementary school, and for letting me watch NOVA past my bedtime. Thank you to my dad for encouraging me to apply for physics scholarships in high school and, when I was considering a physics major, argued, “Well, why don’t you try it for a year? How bad could one year of physics classes really be?” Thanks to my wonderful siblings for their continual support and for always making me laugh.

Finally, thank you to my husband, Lucas, whose support and encouragement has helped more than I can put in words. I’ll save the Dr. Beaufore title for you!

ABSTRACT

Detecting astrophysical neutrinos at the highest known energies will help answer critically important questions in both astronomy and particle physics. While the existence of ultra-high-energy particle accelerators in our universe is well motivated by the decades of cosmic ray detections up to 10^{21} eV, not much is known about the sources themselves. Neutrinos are the only known messenger particle capable of both traveling directly from the source and traveling cosmic distances to reach detectors on Earth. However, because neutrinos rarely interact with matter, detectors must be built to survey many square kilometers to have a chance at regularly seeing these rare events.

The Askaryan Radio Array (ARA) is one such detector, located at the South Pole and designed to be sensitive to the radio emission created when a neutrino interacts in the Antarctic ice. Over the last decade, five independent ARA Stations have been built, amassing decades of station years worth of data and setting some of the most competitive flux limits between 10^{16} and $10^{19.5}$ eV compared to other radio experiments.

This thesis is focused on the newest result from ARA: an analysis of one year of data from the newest station, ARA Station 5. This station is equipped with a low threshold trigger capable of triggering on weaker events than previous ARA stations, with a particular improvement in sensitivity at lower energies. This work discusses the hardware design, deployment, calibration, and first analysis of this new station, proving the feasibility of this type of station design, and motivating its use in future full-scale experiments such as the Radio Neutrino Observatory in Greenland (RNO-G), the Payload for Ultrahigh Energy Observations (PUEO) and potentially IceCube-Gen2.

This thesis will also discuss other projects; in particular, the Beamforming Elevated Array for Cosmic Neutrinos (BEACON), a mountaintop detector looking specifically for cosmic rays. With so many new experiments coming online in the next few years, this analysis will impact big decisions regarding future detector design and analysis techniques.

Contributing to these efforts has been exciting work, as we look to officially enter the era of ultra-high energy neutrino astronomy.

CHAPTER 1

INTRODUCTION

While astronomy has been a field of interest for millennia, it is only within the last century that the field of particle astrophysics began. Since the first discovery of the cosmic ray in the early 1910s, cosmic ray experiments have been carefully, constructed, analyzed, and improved over the course of decades. These experiments have been able to reliably detect cosmic rays across a wide range of energies, allowing the focus to shift to deeper questions about production mechanisms, composition, and astrophysical sources.

Neutrino astronomy is more recent, first undertaken by the Homestake experiment to study solar neutrinos in the 1960s [19]. However, in the nearly 60 years since the Homestake experiment, only three sources have ever been identified as potential neutrino sources: the sun, Supernova 1987(a) [20][21], and blazar TXS 0506+056 (with a 3σ significance) [22] [23]. Additionally, while the astrophysical neutrino flux has been successfully measured by IceCube up to a few PeV, the higher energy neutrino flux has not yet been directly detected. This leaves a wide range for future improvement in neutrino detector design.

In this chapter, I will discuss the motivation for ultra-high energy (UHE) neutrinos, defined for the purposes of this work as neutrinos above 10 PeV. I will then discuss the production mechanisms of neutrinos, neutrino detection methods, and current and future experimental designs.

1.1 Motivation

While neutrinos above 10 PeV have not yet been detected, their detection would contribute significantly to the fields of astronomy and particle physics. In this section, I will consider the questions neutrinos can answer about the Universe, divided into three categories: neutrino astronomy, multi-messenger astronomy, and fundamental physics.

1.1.1 Neutrino Astronomy

Neutrinos offer a unique view into the Universe, particularly at the highest energies. While photons at low energies can travel from any distance, around 10^{15} eV the universe becomes generally opaque to high-energy photons at most distances, as shown in Figure 1.1, due to scattering off of radiation fields like the Cosmic Microwave Background. Additionally, cosmic rays at high energies are limited to the nearby Universe (on the scale of 50 Mpc) and bend in the presence of magnetic fields, making them poor candidates for learning about distant sources. Neutrinos, with low cross sections even at high energies, are excellent candidates for learning about the high-energy, distant Universe.

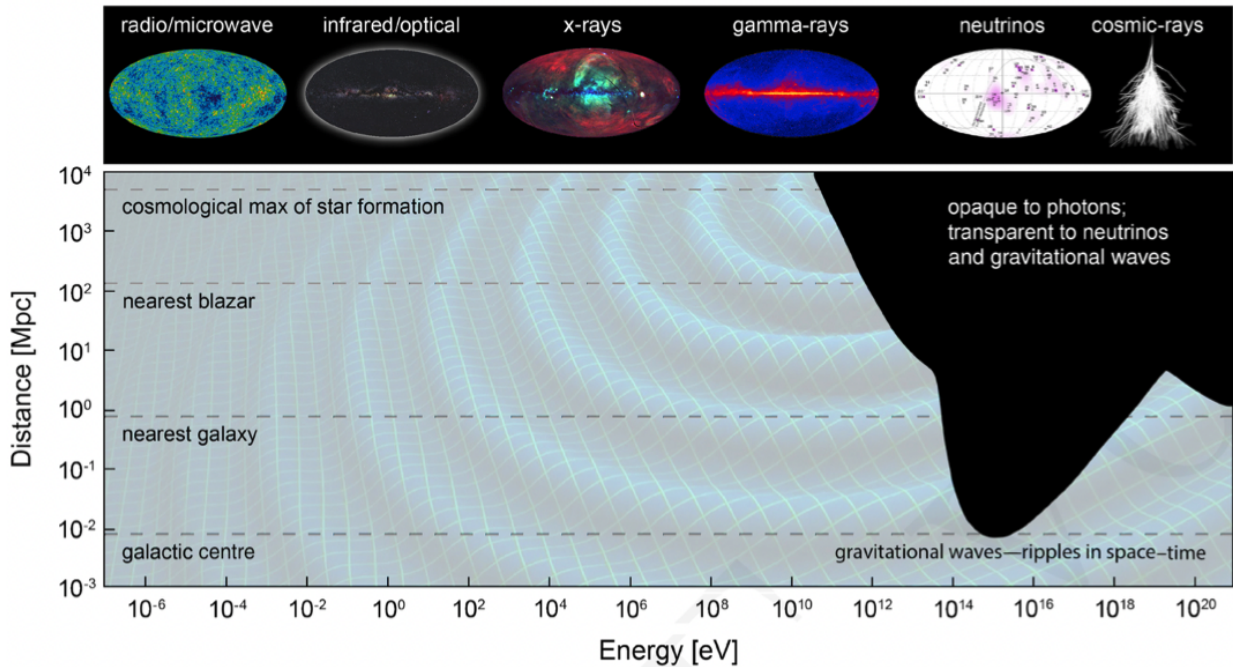


Figure 1.1: A plot showing the distances and energies for which the Universe becomes opaque to photons. Neutrinos and gravitational waves are the only way to explore the distant, highest energy Universe. From [1].

1.1.2 Multi-Messenger Astronomy

While neutrinos carry unique information about very distant and dense astrophysical sources, detecting them in concert with other messenger particles offers the most comprehensive study of the high-energy Universe. In Figure 1.2, I compare the fluxes of three high energy messenger particles: gamma rays, neutrinos, and cosmic rays [2]. The measurements from Fermi-LAT, IceCube, and Auger span many decades of energy, and yet the energy density of each particle is surprisingly similar. This could suggest a common origin for all three messenger particles.




Figure 1.2: A plot comparing the flux of various messenger particles: γ -rays (left), neutrinos (middle) and cosmic rays (right). Also shown are various models for astrophysical neutrino production. From [2].

Additionally, measuring one messenger particle can be used to infer information about other particles at sources as well. As an example, observing UHE neutrinos from a source would provide evidence of hadronic acceleration of cosmic rays at the same source, event without directly reconstructing a cosmic ray event. Similarly, the existence of cosmic ray

particles up to 10^{20} eV motivates the existence of neutrinos at slightly lower energies (discussed more in a later section).

The concept of multi-messenger astrophysics with neutrinos is further bolstered by IceCube's recent detection of a neutrino event coming from blazar TXS 0506+056 in temporal and spatial coincidence with a flare in the gamma ray spectrum measured by Fermi-LAT and MAGIC [22] [23], with a significance of approximately 3σ . This is both the first potential identified high-energy neutrino source and the first hint of a multi-messenger detection using high-energy neutrinos. By extending the measurement of the neutrino spectrum to even higher energies, we can further probe the relationship between neutrinos and other messenger particles.

1.1.3 *Fundamental Physics*

UHE Cosmic Rays (UHECRs) have been measured up to 10^{20} eV, confirming that cosmic accelerators exist that are capable of accelerating particles to energies approximately 10^7 times greater than the highest energy particle accelerators on Earth. This provides a unique opportunity to learn about how the Standard Model behaves at the highest energies. One example of this is the neutrino cross section as a function of energy. While IceCube has measured the neutrino-nucleon cross section for center-of-mass energies up to about a TeV [24] and found the results to be within reasonable error of the Standard Model cross sections, no measurements have been taken at higher energies. Discrepancies between cross section measurements and expectations for UHE neutrino interactions could point to Beyond the Standard Model (BSM) physics.

Another example of probing the Standard Model is flavor composition, which can be studied extensively with UHE neutrinos considering their long baselines. At the sources, neutrinos are expected to be produced in a ratio of $\nu_e : \nu_\mu : \nu_\tau$ of 1 : 2 : 0 due to charged pions decaying; however, with predicted neutrino oscillations, the expected ratio at Earth

is 1 : 1 : 1. This ratio would be impacted by BSM theories like neutrino decay [25] [26] or Lorentz-invariance violation [27] [28].

1.2 Production Mechanisms

In this section, I will discuss the production mechanisms of UHE neutrinos, which are generally broken down into two categories: astrophysical neutrinos and cosmogenic neutrinos.

1.2.1 Astrophysical Neutrinos

Astrophysical neutrinos are created via interactions with UHECRs at the sources. The UHECRs are thought to be accelerated in the sources via Fermi shock acceleration [29][30].




Figure 1.3: A calculation of the neutrino-nucleon charged-current cross section, calculated from the IceCube 6-year dataset. The models match the data within uncertainties up to 10^6 GeV, after which no neutrinos have been detected. From [3].

These cosmic rays, considered here as protons (p), interact with other particles in the source (p and γ). pp interactions result in approximately an equal number of π^+ , π^0 , and π^- , while the $p\gamma$ interactions are dominated by the $\Delta(1232)$ resonance and create pions:

$$p + \gamma \rightarrow \Delta^+ \rightarrow \begin{cases} n + \pi^+ \\ p + \pi^0 \end{cases} \quad (1.1)$$

with the first interaction occurring in approximately $\frac{1}{3}$ of cases, and the second interaction occurring in approximately $\frac{2}{3}$ of cases. Neutral pions decay via $\pi^0 \rightarrow \gamma + \gamma$, creating high energy γ -rays, while charged pions will decay via interactions like

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu \\ \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu \end{aligned}$$

In this way, sources are capable of producing a UHE neutrino flux. The exact energy density of the flux depends on the model of the source. Example models are shown in Figure 1.2, and include Active Galactic Nuclei [31] [32], blazars [33], gamma-ray bursts [34] [35], and pulsars [36].

1.2.2 *Cosmogenic Neutrinos*

Cosmogenic neutrinos are not produced directly at a source, but instead by interactions with UHECRs as they propagate through the universe. Specifically, UHECRs above approximately $10^{19.5}$ eV are expected to interact with the cosmic microwave background photons at a center-of-mass energy of a Δ^+ resonance, producing either a p and a π^0 , or a n and a π^+ :

$$p + \gamma \rightarrow \Delta^+ \rightarrow p + \pi^0 \quad (1.2)$$

$$\pi^0 \rightarrow \gamma + \gamma \quad (1.3)$$

$$p + \gamma \rightarrow \Delta^+ \rightarrow n + \pi^+ \quad (1.4)$$

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (1.5)$$

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (1.6)$$




Figure 1.4: A measurement of the cosmic ray flux as a function of energy for two experiments, Auger and the Telescope Array. The discrepancies in the spectra at highest energies are unexplained and appear to be more than statistical. The cutoff visible near $10^{19.5}$ eV is due to the GZK effect. From the PDG.

In the case of a neutral π^0 , no neutrinos are created, but a charged pion will decay into three neutrinos, each with an energy about 1/20 of the parent cosmic ray, meaning the resulting neutrino spectrum would peak around 10^{18} eV. The exact values depend on the incoming energy of the cosmic ray and the composition of the cosmic ray. Heavier cosmic rays are expected to lose most of their energy through photodisintegration, in which a cosmic ray interacts with a cosmic microwave background photon and reduces both the cosmic ray's nucleon number and energy, making it less likely that the resulting cosmic ray will be energetic enough to interact in the process above.

This effect is often referred to as the GZK effect [37] [38] and is the expected reason for the steep falloff in flux of cosmic rays above $10^{19.5}$ eV seen in Figure 1.4, as cosmic rays above this threshold are only expected to travel tens of Mpc before interacting in this way.

1.2.3 *Science from a Non-Detection*

Because the neutrino flux above 10 PeV has never been measured, it is possible that the various astrophysical and cosmogenic models are overly optimistic, potentially significantly so. In this scenario, analyses from neutrino experiments would set increasingly tighter limits on the neutrino flux, and would not report a detection. While at first glance this seems like a disappointing result, a non-detection still provides a significant window into the possible interactions at the highest energies.

Because the cosmic ray flux has been definitively measured up to about 10^{21} eV, a lack of corresponding neutrinos from the GZK effect could indicate a much heavier mass composition of cosmic rays at high energies. If heavier mass composition models were also ruled out, this would imply that our knowledge of particle physics at the highest energies is missing something fundamental. Similarly, if astrophysical models are also ruled out, the potential sources of the IceCube neutrino flux would become better constrained, leading to better

models of both the sources that are found to produce neutrinos as well as sources that do not.

1.3 Neutrino Detection

Because neutrinos interact weakly, it is very challenging to detect them directly. Instead, we design experiments to look for particle showers that occur after neutrinos have interacted in the Earth. The expected rate of UHE neutrinos is on the order of one event per cubic kilometer per century per steradian, meaning that experiments have to cover a huge amount of volume before being able to detect even a handful of UHE neutrino events.

Neutrinos will interact with matter through a deep-inelastic scattering interaction via either a charged-current (CC) interaction or a neutral-current (NC) interaction. For a CC interaction, the form of the interaction is $\nu_l + N \rightarrow l + N'$, where N is a nucleon and ν_l and l are a neutrino and its associated lepton. The interaction is mediated by a W^\pm boson, and the resulting shower is mostly electromagnetic. For a NC interaction, the form is $\nu + N \rightarrow \nu + N^*$, the interaction is mediated by a Z boson, and the resulting shower is caused by nucleon scattering, starting off as hadronic before becoming more electromagnetic.

As the electromagnetic shower develops, a combination of Compton scattering, Moller scattering, Bhabha scattering, and positron annihilation occurs. The net effect of positron annihilation and electrons scattering in to the shower is a negative charge excess of about 20%. This charge excess is moving faster than the speed of light in the media, creating Cherenkov radiation, radiating outward from the shower axis at an angle defined by

$$\cos(\theta) = \frac{1}{n\beta} \tag{1.7}$$

where n is the index of refraction and $\beta = \frac{v_p}{c}$. The shower has a finite width, called the Moliere radius and defined by the density of the material and the viewing angle. In ice,

for example, the Moliere radius is about 10 cm. The Cherenkov radiation occurs at all frequencies, but for wavelengths greater than the Moliere radius, the radiation is coherent; no matter where these wavelengths originate within the shower, they add coherently. This is called Askaryan radiation [39] and results in coherent radio emission in a dense dielectric material like ice [40]. The radiated Askaryan power scales quadratically with shower energy, unlike typical Cherenkov radiation which scales linearly. This coherence makes Askaryan emission the dominant emission type compared to optical wavelengths for shower energies above approximately 10 PeV. In addition to ice, the Askaryan effect has been observed in sand, rock salt, polyethylene, and the atmosphere [41] [42] [43] [44]. An example of how the Askaryan emission changes as a function of viewing angle is shown in Figure 1.5.

An additional benefit of Askaryan Radiation over optical Cherenkov is the improved attenuation length in ice of radio compared to optical. The attenuation length of optical signals in ice is on the order of 100 m, while the attenuation length of radio is closer to 1 km,




Figure 1.5: A plot showing the electric field emitted at various angles away from the Cherenkov angle (left), and the associated differences in frequency content (right). From [2].

allowing the radio detectors to be more sparsely instrumented than the optical experiments targeting lower energies.

1.4 Neutrino Detector Summary

In this section, I will briefly discuss the current astrophysical neutrino detectors, summarizing the current status of various neutrino experiments at high energies, including those discussed




Figure 1.6: A plot of the current best limits on the neutrino flux. Included are limits from IceCube [4] [5], Auger [6], ARA [7], ARIANNA [8], and ANITA [9]. Also shown are theoretical neutrino production models. From [7].

in 1.6.




Figure 1.7: A diagram of the IceCube detector, with the Eiffel Tower for scale. From icecube.wisc.edu.

1.4.1 The IceCube Neutrino Observatory

IceCube is mainly sensitive to neutrinos in the TeV-PeV range, and consists of thousands of optical Cherenkov detectors covering an instrumented volume of 1km^3 , with its main detectors buried in the ice at the South Pole at depths between 1450 and 2450 m. Figure 1.7 illustrates the detector design. IceCube is densely instrumented so that it captures the shape of a particle shower as it develops. By recording quantities like the trajectory, the energy deposited, and the length of the shower, IceCube is able to reconstruct the flavor, energy, and direction of the original neutrino particle. Even though IceCube is located at the South Pole,

they are most sensitive to events that arrive from the Northern sky, where the Earth blocks atmospheric muons and atmospheric neutrinos, some of the most prominent backgrounds. IceCube currently has measured the neutrino flux up to a few PeV, and has the best limit up to an energy of $10^{19.5}$ [4] [5]. Other experiments that have utilized similar detection methods to IceCube are AMANDA (the pathfinder for IceCube)[45] and ANTARES [46].

1.4.2 ANTARES and KM3NET

The Astronomy with a Neutrino Telescope and Abyss Environmental RESearch (ANTARES) experiment is a water Cherenkov detector deployed in the Mediterranean Sea and designed to be sensitive to Cherenkov light emitted by showers caused by neutrino interactions. It consists of 12 lines with 75 photo-multiplier tubes (PMTs) each, reaching a total instrumented volume of 0.1 cubic kilometers. ANTARES was fully constructed by 2008 and recently completed its data-taking in early 2022, publishing results that are compatible with the IceCube cosmic neutrino flux [46]. KM3NET is another underwater experiment currently under construction which builds off of the success of ANTARES, with a planned cubic kilometer of instrumented volume. KM3NET consists of two separate experiments: ORCA, which focuses on atmospheric neutrinos and lower energy neutrino studies, and ARCA, which is designed to detect cosmic neutrinos between tens of GeV and a few PeV.

1.4.3 EUSO-SPB2 and POEMMA

The Extreme Universe Space Observatory on a Super Pressure Balloon II (EUSO-SPB2) is a pathfinder instrument for the future Probe of Extreme Multi-Messenger Astrophysics (POEMMA) [47]. Currently planned to fly in 2023, it will be sensitive to both UHECRs (above 1 EeV) using a fluorescence detector technique and UHE neutrinos (above 20 PeV) by detecting Cherenkov emission from their interactions. EUSO-SPB2 and POEMMA look for optical Cherenkov, not coherent radio emission, and will be the first experiment to

measure air showers from extremely high altitudes. The instrument development and flight of EUSO-SPB2 will directly influence the design of the future POEMMA instrument, which is undergoing NASA-funded concept studies.

1.4.4 ANITA and PUEO




Figure 1.8: A diagram showing the Askaryan emission channel as seen by PUEO.

The Antarctic Impulsive Transient Antenna (ANITA) was a balloon experiment that flew four separate times around the Antarctic continent looking for the Askaryan radio signal. Because ANITA is flying 30 km above the surface of the ice, it is most sensitive to extremely bright, high energy neutrino events. ANITA's antennas targeted a frequency range of 200-1200 MHz, and although the exact design of ANITA changed with each consecutive flight, each flight was successful, setting increasingly competitive limits at the highest end of the

energy spectrum [48][49][50][9].

PUEO is the next generation balloon experiment targeting UHE neutrinos, and is currently under development for an expected flight date of December 2024. A schematic of the PUEO concept is shown in Figure 1.8. There have been quite a few updates from the old ANITA design, most notably a phased array trigger modeled after the successful prototype deployed by the ARA experiment. Additionally, the antennas are smaller with sensitivity beginning at 300 MHz, and a new low-frequency drop down instrument will be deployed after launch. The low frequency instrument will be particularly helpful for investigating air shower events, including the flipped polarity events seen by earlier ANITA flights with unknown origins [51]. These updates are expected to give PUEO the leading neutrino sensitivity above 10^{18} eV [52].

1.4.5 *ARIANNA*

ARIANNA is an in-ice experiment sensitive to Askaryan emission from neutrinos that reflects off of the ice-water boundary at the bottom of the Ross Ice Shelf, as shown in Figure 1.9. ARIANNA deploys its antennas at the surface of the ice near Moore’s Bay, Antarctica without drilling holes. This is an inexpensive and quick way to install a neutrino detector, especially considering each station only consists of four LPDA antennas each pointed down and oriented perpendicularly to each other. The LPDA antennas have a high gain and are not limited in their design by needing to be deployed down a hole, or attached to a balloon. They operate remotely in an area without cabled power, resulting in the ARIANNA stations operating for approximately five months of the year. These design choices generally make them less sensitive to neutrino events, as their trigger is run at a higher threshold and their effective volume is smaller. ARIANNA’s most recent analysis result is included in Figure 1.6 [8].




Figure 1.9: A diagram showing the types of interactions an ARIANNA station is capable of detecting. From [8].

1.4.6 RNO-G

The Radio Neutrino Observatory in Greenland (RNO-G) is a neutrino observatory currently under construction in Summit Station, Greenland [2]. Its location in the northern hemisphere sets it apart from the other radio experiments discussed so far, giving RNO-G the unique opportunity to investigate interesting Northern hemisphere sources like the Telescope Array hotspot [53] and the blazar detected by IceCube, Fermi-LAT, and MAGIC [22] [23].

RNO-G is planned to include 35 total stations, which would make it the largest in-ice radio detector. Each RNO-G station is designed as a hybrid between the deep stations of ARA and the surface stations of ARIANNA. Each station is also equipped with a phased array trigger, similar to the prototype trigger installed as part of the ARA experiment.

One of the challenges of RNO-G is the lack of cabled power or communications, which has necessitated each RNO-G to be autonomously powered from solar panels and battery banks. RNO-G is designed to operate in multiple data-taking modes at different power consumption levels, ranging from 6W-24W for data-taking modes and down to 70mW for winter operations.

As the number of stations continues to scale up, RNO-G will quickly become a detector-scale instrument, and could realistically be the first experiment to detect UHE neutrinos with the radio technique.




Figure 1.10: A diagram of an RNO-G station, including both surface and deep stations. From [2].

1.4.7 BEACON

The Beamforming Elevated Array for Cosmic Neutrinos (BEACON) is a new prototype experiment in the White Mountains of California. Although BEACON is fairly different from the previous experiments discussed in this section, it is mentioned here to introduce the concept so that the later section on BEACON is well motivated.

BEACON is sensitive to geomagnetic radio emission from cosmic ray air showers and UHE ν_τ . ν_τ will likely create τ leptons when skimming the Earth, which will decay and shower, generating synchrotron radiation in the radio frequency regime. The full BEACON




Figure 1.11: A diagram showing the concept of a BEACON array, including the geometry of detecting a decaying tau lepton.

instrument would include 100 independent stations with 10 antennas each, installed on moun-

tainsides across the globe [54]. These stations would use the high elevation to increase their viewing angle, and each station would use a beamforming trigger to further improve the station performance. Currently, a prototype BEACON experiment has been installed at White Mountain Research Station since summer 2018. Early development on this station and its current status will be discussed in more detail in Chapter 7.

1.4.8 ARA




Figure 1.12: A diagram showing the geometry of a neutrino interaction seen by the ARA experiment. From [7].

The Askaryan Radio Array (ARA) is an in-ice detector deployed at the South Pole, just a few kilometers from IceCube. The first ARA station was first deployed in 2011, and since then ARA has amassed many station-years of data across its five independent stations. A typical ARA Station consists of 16 antennas of mixed polarization deployed to a depth of 200 m at the South Pole, just a few kilometers away from IceCube. ARA has the most competitive limit of a radio experiment between 10^{17} eV and $10^{19.5}$ eV [7].

ARA has recently deployed a new station with a prototype phased array trigger [55] [12], expected to improve performance significantly. This station is the topic of this work. Chapter 2 will discuss the ARA station design and electronics. Chapter 3 will discuss the methodology behind the updated beamforming trigger, its design, and its implementation as part of ARA Station 5. Chapter 4 will discuss the calibration of the IRS2 digitizer chips installed as part of the classic ARA station system. Chapter 5 will discuss the method used to determine the antenna locations within the ice. And finally, Chapter 6 will discuss the analysis of one year of data from ARA station 5.

CHAPTER 2

THE ASKARYAN RADIO ARRAY (ARA)

The Askaryan Radio Array (ARA) is an in-ice neutrino detector at the South Pole built to detect Askaryan emission from neutrino interactions within the Antarctic ice. This chapter details the instrument design of the various ARA stations, up to and including the most recently deployed stations, Station 4 and Station 5. The details of the Phased Array instrument, deployed as part of ARA Station 5, are included in the next chapter.




Figure 2.1: A diagram of a typical ARA Station, including close-up views of the VPol and HPol antennas.

2.1 Station Design

ARA is designed to be sensitive to Askaryan radio emission induced by neutrino interactions with Antarctic ice, discussed in Chapter 1. This signal is expected to be broadband below 1 GHz, with the exact frequency content dependent on the viewing angle relative to the Cherenkov cone. With this in mind, the ARA instrument is designed to be sensitive to frequencies between 200-850 MHz, installing both vertically-polarized (VPol) and horizontally-polarized (HPol) antennas to record polarization information about incoming signals. Recording this polarization information allows the neutrino direction to be reconstructed.

The design of a typical ARA Station is shown in Figure 2.1. The main instrument consists of four holes drilled to a maximum depth of 200 m, with a set of VPol and Hpol antennas deployed at the maximum hole depth, and another set of antennnas deployed approximately 30 m above. This design is motivated by the characteristics of the ice at the South Pole. Snow accumulates at the South Pole and slowly compacts into ice over the course of many years, causing the index of refraction of the ice to vary in the few hundred meters closest to the surface, an area called the firn. The change in index of refraction is approximately exponential, with a surface index of refraction similar to that of snow, and the index of refraction below a few hundred meters stable at approximately 1.78. Thus, the radio signals propagating through the ice bend as they move through the ice, making it advantageous to deploy antennas as deep as possible to increase the effective volume of the experiment. For the ARA experiment, the deployment depth was chosen to be 200 m. The ice model will be discussed in more detail in a later chapter.

One challenge of deploying antennas into drilled holes with a 15 cm diameter is designing antennas with the appropriate size requirements while maintaining isotropic sensitivity in all directions. This is significantly easier to do for VPol antennas, which are sensitive along the same axis as the hole and thus can be designed to have the necessary length required

for the desired bandwidth. The VPol antennas are often referred to as “birdcage” antennas, due to their wire frame and enclosed dipole design; a picture of the VPol antenna used is shown in Figure 2.1. These antennas have broadband sensitivity in the 150-800 MHz range.

The HPol antennas are more challenging to design, as their sensitivity is meant to be in the direction perpendicular to the hole, yet their width is limited by the width of the hole. The HPol antenna is a quad-slot antenna, with a special magnetic material called a ferrite included in the feed to improve the performance at low frequencies. A picture of the HPol antenna and its simulated gain is shown in Figure 2.2.

The ARA Experiment currently has five separate stations deployed, and while there have been small changes in the station design as more stations have been built, the basic layout has remained the same. Each station is separated from its nearest neighboring station by about 2 km, each surveying its own independent area of ice. ARA Stations 4 and 5 are the most recently deployed stations, and each has increased separation between the strings (40 m instead of 20 m) to improve the effective volume further. Additionally, ARA Station 5 includes an additional string and an additional triggering system called the Phased Array. The Phased Array will be the topic of the next chapter.

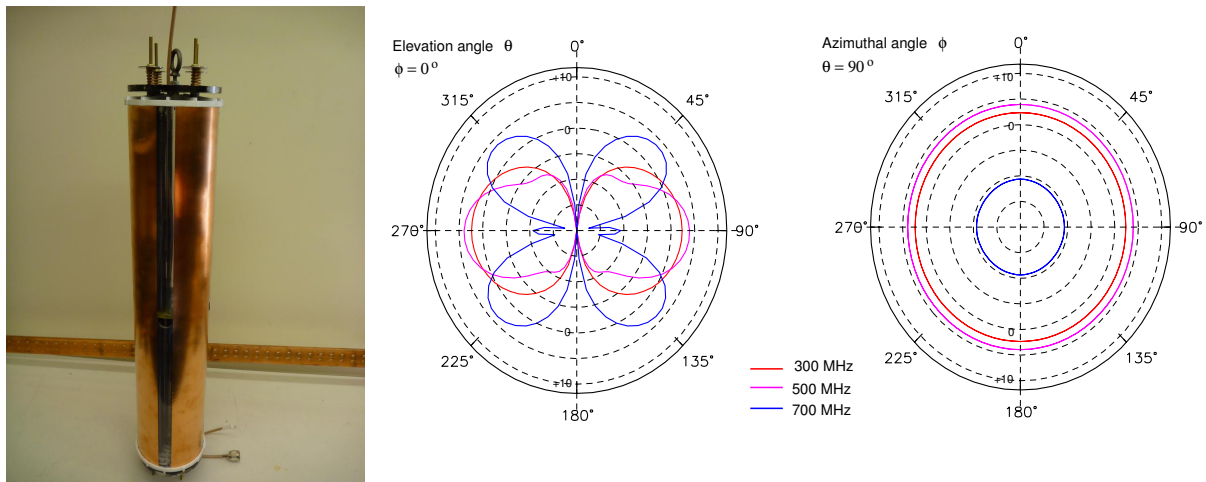


Figure 2.2: A picture of the HPol antenna, along with its simulated gain in dBi. From [10].

2.2 Signal Chain

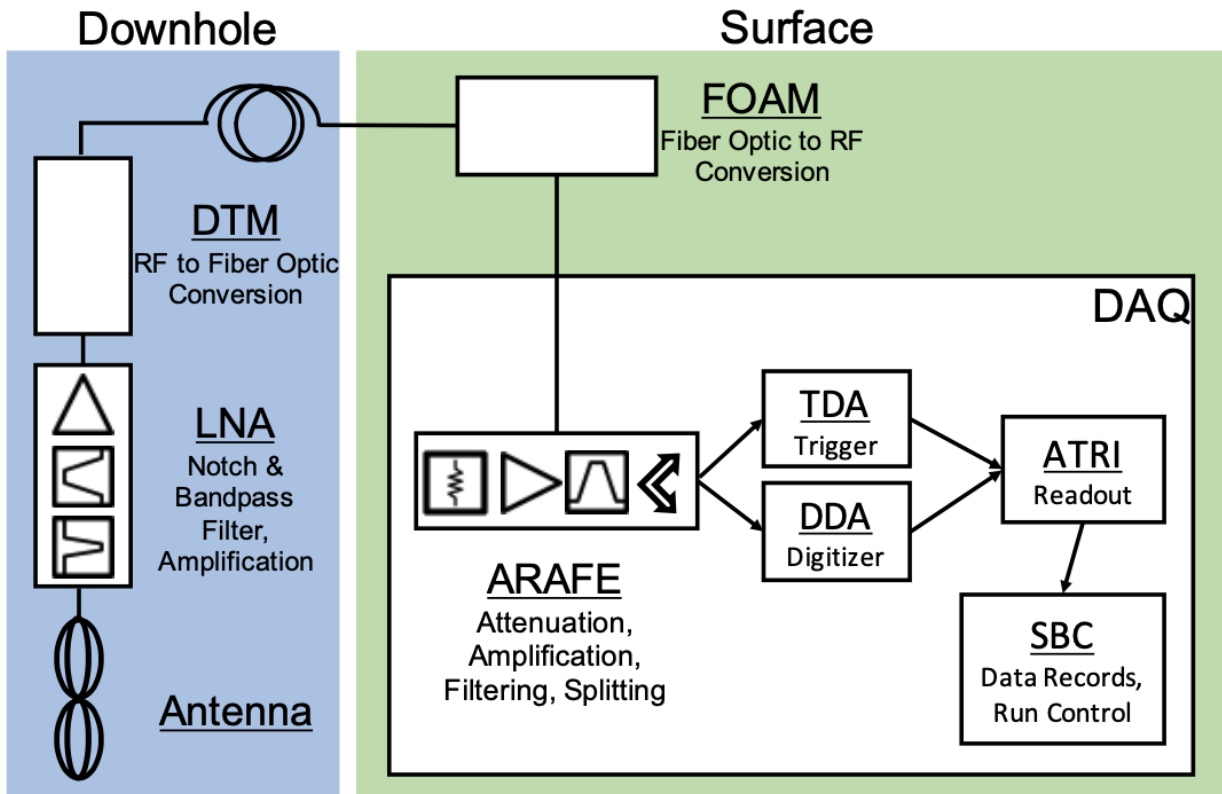


Figure 2.3: A diagram of the signal chain of the ARA instrument. From B. Clark [11]

The signal chain of the ARA instrument is shown in Figure 2.3. Before the radio frequency (RF) signal is transmitted to the surface, a notch filter is applied at 450 MHz to remove radio communications from the nearby South Pole Station. Additionally, a commercially-available bandpass filter from 150-800 MHz is applied to limit out-of-band noise, before boosting the signal by 40 dB in two stages of low noise amplifiers that are mounted on each antenna.

After the RF signal has been amplified, the electronic signal is converted from RF to RF-over-fiber (RFoF) cable by the Downhole Transmission Module (DTM) which limits the attenuation that occurs over the 200 m distance to the data acquisition (DAQ) box at the surface. Once the signal has reached the surface, the Fiber Optic Amplification Module (FOAM) converts the signal back to an electronic signal and amplifies the signal again by

approximately 16 dB.

For ARA Stations 1-3, the FOAM had additional amplification capabilities, and the DAQ was equipped with an additional set of commercial bandpass filters to reduce noise introduced by the amplifiers. This solution was costly, as commercial filters provided more gain reduction than was necessary, and added additional mechanical failure points on each of the channels. For ARA Stations 4 and 5, the ARA Front End (ARAFE) was introduced, combining the additional amplification and filtering into one printed-circuit board that was housed in the DAQ. This board also had an attenuator for each channel, allowing the gain on each channel to be matched.

At this point, the signal was split and sent through the Triggering Daughter Boards (TDAs), responsible for applying the trigger, and Digitizing Daughter Boards (DDAs), responsible for recording the waveforms of the triggered events. These communicate through the ARA Triggering and Readout Interface (ATRI) board. The digitizers on the DDA are Ice Ray Sampler 2 (IRS2) chips, discussed in detail in Chapter 4.

2.3 Current Status and Future Upgrades

As of this writing, ARA Stations 1, 3, and 5 are currently operational. Over the years, a failure mode was discovered in which the USB port on the FX2 chip on the ATRI failed, making it impossible to program the board and thus rendering the station unusable. This has happened to ARA Stations 2, 4, and 5, and while ARA Station 2 was successfully fixed with a replacement ATRI board, the spare has now failed as well. The traditional ARA trigger at ARA Station 5 is no longer operational, but the Phased Array instrument is, and after deployment work during the 2019-2020 season, the VPol channels from the traditional ARA instrument were added to the Phased Array instrument. This combination station has been operational since.

There has been some discussion about future upgrades to the existing ARA stations, both

to permanently fix the known USB failure mode and to upgrade other pieces of the DAQ box, such as replacing the IRS2-based DDAs with LAB4D-based boards. This discussion is still ongoing.

CHAPTER 3

THE PHASED ARRAY INSTRUMENT

Over the years, radio experiments like ARA and ARIANNA have been accumulating data and setting increasingly competitive limits. However, the effective volume of these experiments at energies around 10 PeV is low, making it difficult for these experiments to search for extensions of the IceCube flux without a significant increase in livetime, analysis efficiency, or number of stations. With this in mind, the phased array concept was developed, simulated, and eventually deployed as part of ARA Station 5 [55] [12]. This trigger is now the baseline design for current and future radio experiments, like RNO-G [2], PUEO[52], and IceCube-Gen2[56].

3.1 The Phased Array Trigger Concept

In many in-ice experiments, the gain of the antenna is a limiting factor due to physical constraints like the width of the hole that can be drilled. This can be improved by phasing multiple antennas together, a process done often in radio astronomy experiments [57]. For the purposes of this chapter, I will be describing phased array systems made up of antennas arranged vertically in one dimension with limited spacing between each individual antenna, similar to the design proposed in [55]. The benefits of this technique can be illustrated by considering the gain G_{array} of an array of N antennas, described by

$$G_{array} = 10 * \log_{10}(N \times 10^{\frac{G}{10}}) \tag{3.1}$$

where G is the gain of the individual antennas in the array. In this case, increasing N increases G_{array} , allowing the gain of the full system of antennas to be higher than the gain of a single antenna. The gain of the array and of the individual antennas is frequency and direction dependent. This improvement in gain translates directly to a lower energy

threshold and a higher effective volume for neutrino experiments.

Adding together signals across multiple individual antennas results in the noise adding incoherently compared to the signal, which adds coherently after applying the appropriate time delays for the direction of the signal. At the trigger level, multiple potential directions can be implemented at the same time, allowing the phased antenna array to cover the full directional range of a single antenna but with higher gain.




Figure 3.1: An example of a one-dimensional vertically-spaced phased array trigger array (center), surrounded by additional antennas used for reconstruction. From [12].

In practice, a phased array design is implemented in the time domain using a “delay and sum” method. In this method, a coherent sum $S(t)$ is formed according to

$$S(t) = \sum_{n=0}^{N-1} w_n y_n(t - \delta_n) \quad (3.2)$$

in which w_n is the weight of the n th individual antenna amplitude, y_n is an array of voltage measurements recorded over time by the antenna, and δ_n is the delay applied to the antenna signal. For our purposes, all antennas are given equal weights. In this way, many coherent sums can be formed for many sets of time delays. These sets of time delays are chosen to correspond to incoming directions, given by

$$\sin \theta_m = \frac{c m \Delta t}{n d} \quad (3.3)$$

where c is the speed of light, m is the beam number, Δt is the time spacing of the digitizer, n is the index of refraction, and d is the spacing between the antennas. By calculating the appropriate set of delays for all pairs of antennas that correspond to a given incoming direction, a beam is formed. The trigger can include multiple beams, each corresponding to its own incoming direction, allowing more complete coverage of the surrounding area.

The trigger assumes the incoming signal is plane-wave-like, meaning that the signal originates from a location far enough away that the wave front can be approximated as a plane instead of a sphere. It also assumes that the signal looks the same on all antennas, an assumption that is true if the signal hits identical antennas at the same angle and if their impulse responses match. Because of these assumptions, it is easiest to design a phased array system to be as compact as possible, with minimal spacing between the individual antennas. This is especially true for environments with a changing index of refraction, which magnifies the differences seen across antennas at different depths. However, the compact array has limited ability to reconstruct events, a task best done by antennas spaced by tens of meters. Thus, for many experimental designs, it is advantageous to design a compact array of antennas to be used in the phased array trigger, with additional strings of antennas separated by

tens of meters to aid in the reconstruction of the interaction location. An example diagram of this setup is shown in Figure 3.1.

The rest of the chapter discusses the specific implementation of the phased array trigger design as part of the ARA experiment.

3.2 The NuPhase Instrument at the South Pole




Figure 3.2: A diagram of ARA Station 5, with the traditional ARA antennas in grey and the Phased Array antennas in red.

ARA Station 5 was deployed during the 2018-2019 season at the South Pole, and includes both a traditional ARA trigger, discussed in Chapter 2, as well as a Phased Array trigger. An example station diagram is shown in Figure 3.2. In total, ten VPol antennas and two HPol antennas were deployed in the Phased Array borehole, although three VPol antennas

were inoperable after deployment, likely due to issues with the mechanical connection of the antenna feed exacerbated when lowered into the hole. This left seven VPol antennas deployed at depths between -172 m and -180 m, with 1 m spacing between the five most shallow antennas and 2 m spacing between the two deepest VPol antennas and the HPol antennas beneath them. The HPol antennas are not used for triggering but can be used for event reconstruction. The Phased Array string is located in the center of the ARA antennas to take advantage of the long baselines and superior reconstruction ability of the ARA antennas. With an external trigger from the Phased Array routed into the ARA DAQ system, the ARA system is able to record events that trigger its own antennas, or events that are detected only by the Phased Array.

3.2.1 The RF Chain

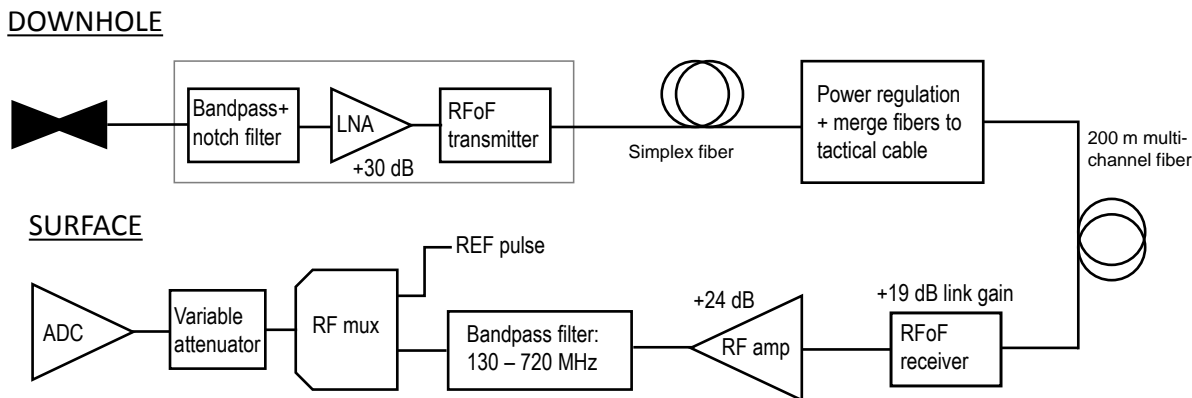


Figure 3.3: An illustration showing the RF chain for the nuphase instrument at the South Pole. From [12].

In this section, the components of the RF chain are described. The full chain is illustrated

in Figure 3.3.

The VPol antennas used by the Phased Array instrument are nearly identical to those used by the rest of the ARA instrument; they have the same basic geometry but a different feed point. The HPol antennas are identical to those used by ARA. Each antenna is equipped with a low-noise amplifier (LNA), bandpass filter, and RF-over-Fiber (RFoF) transmitter. As with the ARA system, the bandpass filter includes a 50 dB notch filter at 450 MHz to suppress anthropogenic noise from the nearby South Pole Station. The LNA, while adding 32 dB of gain to the system, also provides pass-through power to the antenna 1 m below, a crucial difference from the ARA system. This ensures that the complex impedances of each of the antennas are matched, resulting in nearly identical signals recorded by each of the VPol antennas. The RFoF transmitter is necessary to limit the loss that occurs when transmitting the signals to the surface. The downhole system in total requires approximately 25 W of power.

The downhole components were housed in a custom-built frame made from fiberglass rods and plastic faceplates, necessary to ensure the antennas did not rotate while deployed and to absorb any stress that would otherwise fall on the connections between the antennas, the LNAs, and the cables. The antennas were connected to each other via rope to ensure the coaxial cables were not weight bearing.

The RFoF fiber is transmitted directly into the NuPhase instrument box, which houses the RFoF receiver to convert the RFoF signal back to electronic RF on a coaxial cable. The signal then undergoes another round of amplification and bandpass filtering, necessary to suppress out-of-band noise above 750 MHz. At this point, the signal on each channel goes through an RF mux to switch the input from the downhole signal to an on-board calibration pulse using for timing calibration, and then through a variable attenuator, used to normalize the gain of each channel. The signals are then sent to custom analog-to-digital conversion (ADC) boards, after which the VPol signals are sent directly to the FPGA for beamforming.

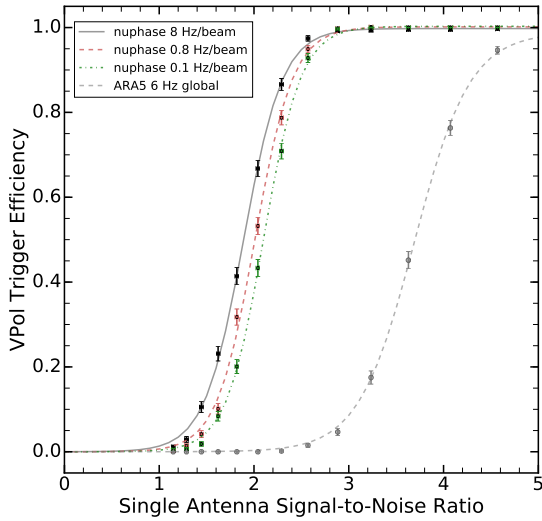


Figure 3.4: The in-situ efficiency of the Phased Array instrument, compared to the efficiency of the classic ARA trigger [12].




Figure 3.5: The expected beam widths for each of the beams in the South Pole Phased Array instrument. From [12].

3.2.2 Trigger Implementation

The Phased Array trigger includes 15 total beams, each made up of a primary and secondary sub-beam. These primary and secondary sub-beams are coherent sums that include antennas spaced 1 m and 2 m apart, respectively. This allows for each beam to have a full width half maximum beamwidth of approximately 7° , with full coverage between the elevation angles of -53° and 47° , roughly consistent with the beam pattern of the individual VPol antenna. The beam coverage is shown in Figure 3.5.

Each beam is given a trigger goal of 0.75 Hz to meet a global 15 Hz trigger rate, with the threshold in each beam adjusted live to meet this requirement. While the beams are directional, a signal does not have to be centered on a beam in order for a trigger to be issued. For signals with high enough signal strength, it is possible for a beam not associated with the direction to trigger, due to enough power coherently summing in the wrong beam. While this could be adjusted at the firmware level by recording the beam with the most power as the triggered beam, this feature is not necessary in an environment with minimal

anthropogenic contamination like the South Pole.

3.2.3 Matching Events Between the Phased Array and ARA5

ARA Station 5 contains two separate DAQs, which makes combining information between the two more difficult than would be expected. The Phased Array is capable of externally triggering the classic ARA 5 instrument so that full reconstruction of an event is possible. This is successful for up to 98% of events, although this varies depending on the operational mode of the classic ARA 5 instrument. However, the events are saved separately by each experiment and not easily matched between the systems. Additionally, the clocks of each experiment can float relative to each other, requiring a correction factor to be applied before the events can be aligned. After correcting the clocks, the spacing between the events on each system is the easiest way to match the events together..

3.2.4 Limitations

The following are limitations of the current Phased Array instrument which could be addressed in future iterations of such an instrument:

Temperature Regulation :One unanticipated difficulty of the Phased Array instrument is its inability to dissipate heat. The Phased Array instrument as a whole generates approximately 80 W of power, the majority of which occurs in the DAQ box itself. The DAQ box is surrounded by snow, effectively insulating the ARA 5 crate and minimizing the possibility of dissipating heat into the surrounding environment. While this will not be a concern for future Phased Array systems, which are designed to operate at lower power, this limitation requires the Phased Array system to be turned off between December and March, the warmest months of the year.

Misaligned Channels : The time-alignment process between different ADC chips is not

always successful and needs to be checked for each run. As discussed in [12], a random 1.5-GHz clock-cycle offset occurs at start-up for each of the ADC chips, which each are responsible for digitizing two channels. This has been solved at start-up by sending a pulse to each of the channels and aligning the clocks based on the location of the peak in each datastream. However, a bug in this code meant that occasionally, the peaks would be misaligned by one sample, causing two of the channels to be offset relative to the others. This problem affects approximately 5% of runs in 2019 and is easily solved in processing by re-running the alignment script without the bug.

Matching with ARA5 Events : Recording half of an event on one system and the other half on another is a system that is prone to errors or challenges. Both systems need to be separately monitored and any issues present in either system makes recovering the full reconstruction information for an event difficult if not impossible. Future experiment will be designed to have a single central computer for each station, which will solve problems related to this matching process.

CHAPTER 4

TIMING AND VOLTAGE CALIBRATION OF ARA STATION 5

4.1 The IRS2 Chip

In this chapter, I will discuss the calibration of the IRS2 Digitizer, necessary for complete waveform recovery for the classic ARA5 channels. The Phased Array digitizers are self-calibrating and are thus not discussed in this chapter.

A digitizer for a radio experiment like ARA needs to meet two main design criteria. One, the digitizer must be capable of recording multiple channels at rates of a few GHz. Two, the digitizer must have low power consumption, in order to operate successfully in environments where power may be limited, such as the South Pole or other remote areas.

With this in mind, the IRS2 chip was developed [58]. Fundamentally, this is a Switch

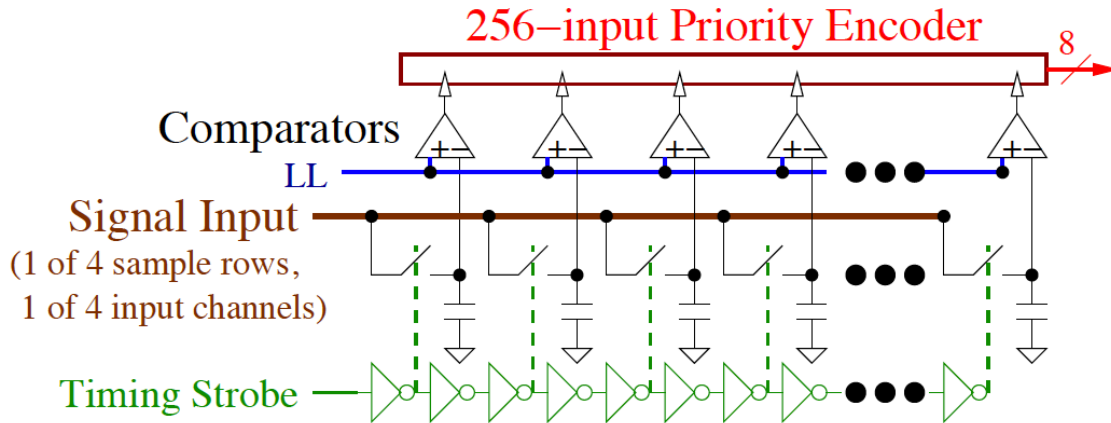


Figure 4.1: An example of a Switch Capacitor Array. From [13].

Capacitor Array [59], consisting of a sampling array of 128 capacitors, as shown in Figure 4.1. The switches open and close based on two timing strobes, timed so that the voltage coming in on the signal input line can be stored by one capacitor at a time. The timing strobes operate at 25 MHz, and the delay between them is 0.31 ns, meaning that each capacitor records the voltage within a 0.31 ns time window. The clock speed and delay offset can be tuned on the IRS2 chip to meet specifications up to a sampling rate of 4 Gs/s at maximum.

The capacitors in the sampling array are separated into two chunks of 64 capacitors each, called blocks. These blocks alternate between sampling the signal and storing the signal in the storage array. The storage array contains 512 blocks which can all act as storage simultaneously, creating an analog buffer 32,768 samples deep, translating to approximately 10 μ s of temporary buffer storage. If the IRS2 chip receives a trigger from the attached FPGA, the signal is digitized using the Wilkinson method [60] [13]: for each storage capacitor, the value stored is compared to the voltage on a comparator as it is ramped up, and the passing count is recorded as ADC counts.

There are two main calibrations that are required in order to use the data taken with the IRS2 chip. First is the timing calibration, which refers to the small differences in the timing constant for each of the 128 sampling capacitors in the sampling array. Second is the voltage calibration, which must be done individually for each of the 32,768 storage capacitors.

One of the challenges of the IRS2 chip is that the performance degrades as the channel number increases. While each IRS2 chip has eight total channels that can simultaneously record signals, only the first two are usable for digitization at a rate of 3.2 GHz; the next two channels are usable at half of that rate, and the final four channels are not usable. This means that for a station with sixteen total antennas that need digitizing, four IRS2 chips are needed. Each IRS2 chip is mounted on a Digitizing Daughterboard for ARA (DDA), which applies some additional bandpass filtering from 150-850 MHz to limit aliasing. The DDAs are each labeled with a number from 0 to 3 so that each channel can be uniquely identified.

Electric Channel	DDA	IRS2 Channel	Antenna Type	Antenna Location
0	0	0	VPol	Bottom
1	0	1	VPol	Top
2	0	2	HPol	Bottom
3	0	3	HPol	Top
8	1	0	VPol	Bottom
9	1	1	VPol	Top
10	1	2	HPol	Bottom
11	1	3	HPol	Top
16	2	0	VPol	Bottom
17	2	1	VPol	Top
18	2	2	HPol	Bottom
19	2	3	HPol	Top
24	3	0	VPol	Bottom
25	3	1	VPol	Top
26	3	2	HPol	Bottom
27	3	3	HPol	Top

Table 4.1: The channel mapping of all channels in the A5 baseline instrument. The Electric Channel uniquely describes a specific antenna and channel.

4.2 Channel Mapping and Available Calibration Data

Before deploying the ARA systems in the ice, calibration data is taken to record the behavior of the channels in a room temperature environment, which can be used to correct for features in the digitizer chips that affect how they record timing and voltage information for each waveform. It is also important to tag each channel with a unique identifier, so that the antennas in the ice can be identified in the data. The list of channel identifiers, DDA numbers, and eventual antenna type and location are recorded in Table 4.1. The locations of each antenna will be covered in the next chapter.

The calibration data, taken by collaborators at UW Madison, consists of sine waves taken with frequencies ranging from 218-800 MHz and input power ranging from -10 to -19 dBm. This power was split across 8 channels and amplified using the ARA amplifiers, adding 22 dB per channel. This results in amplitudes between 158-445 mV. The calibration

Run Number	Frequency (MHz)	Input Power (dBm)	V_p (mV)	Electric Channels
1402	218	-10	445	0,3,8,16,19,27
1403	353	-10	445	0,3,8,16,19,27
1404	521	-10	445	0,3,8,16,19,27
1405	702	-10	445	0,3,8,16,19,27
1406	800	-10	445	0,3,8,16,19,27
1407	218	-9	501	0,3,8,16,19,27
1408	218	-11	398	0,3,8,16,19,27
1409	218	-13	316	0,3,8,16,19,27
1410	218	-15	251	0,3,8,16,19,27
1411	218	-10	445	1,2,9,10,17,18,25,26
1412	343	-10	445	1,2,9,10,17,18,25,26
1413	521	-10	445	1,2,9,10,17,18,25,26
1414	702	-10	445	1,2,9,10,17,18,25,26
1415	800	-10	445	1,2,9,10,17,18,25,26
1416	218	-9	501	1,2,9,10,17,18,25,26
1417	218	-9	501	1,2,9,10,17,18,25,26
1418	218	-11	398	1,2,9,10,17,18,25,26
1419	218	-13	316	1,2,9,10,17,18,25,26
1420	218	-15	251	1,2,9,10,17,18,25,26
1421	218	-17	200	1,2,9,10,17,18,25,26
1422	218	-19	158	1,2,9,10,17,18,25,26

Table 4.2: A list of the available calibration data taken for ARA Station 5 by Thomas Meures.

data was taken on half of the channels at a time, to minimize cross talk across neighboring channels. Table 4.1 lists all available calibration data, including the frequency, input power, and channels. As of this writing, the data is available on the Cobalt server at Wisconsin under `/data/wipac/ARA/pre-deployment/ARA05-UW`.

When taking this data, two channels (11 and 24) were accidentally disconnected and thus the calibration data was not recorded. This data was retaken while at the South Pole in December 2017, and subsequently deleted due to a mix-up in the computer system. These digitizer channels remain uncalibrated, and the effects of this lack of calibration data is discussed in a later section.

The steps for calibrating the digitizer chips are as follows:

1. Perform pedestal subtraction.
2. Calibrate the timing of the digitizer sampling array (128 samples per channel).
3. Calibrate the voltage of the storage array (32,768 samples per channel).

These steps are described throughout the rest of the chapter.

4.3 Pedestal Subtraction

A bias voltage is applied to the RF inputs of the DDAs so that each sample always records a positive value in ADC. This can be corrected on a sample-by-sample basis by subtracting the average measured value for each sample, resulting in data that, in theory, should be

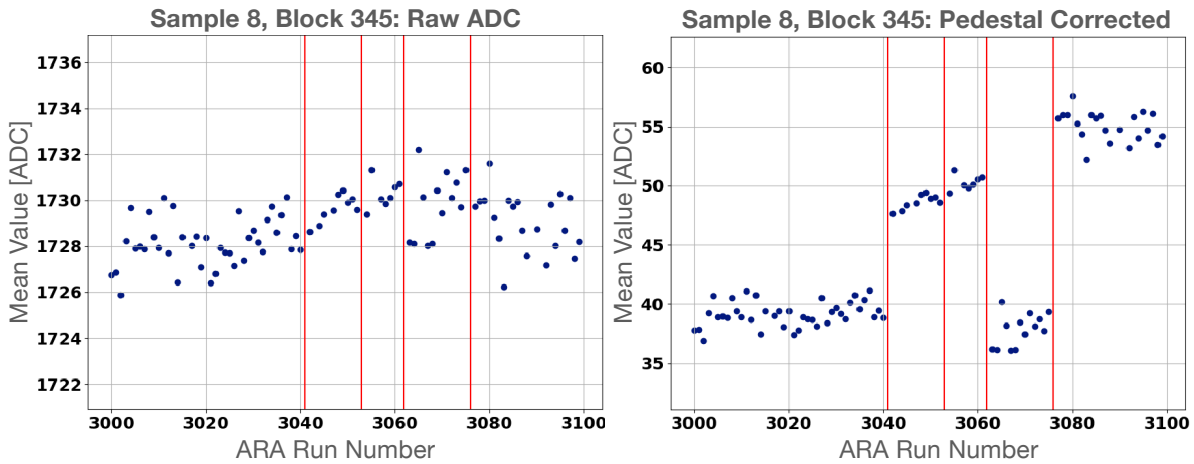


Figure 4.2: Left: the average ADC offset for a specific sample over time. Right: the corrected mean value after subtracting the automatically-generated pedestal values. The red lines correspond to times when the pedestal value was recalculated.

centered around 0 ADC. Traditionally, this subtraction has been done using the generated pedestal files taken once per day, which only records the average value for each sample. However, this average value can be skewed by the data recorded within the first block of data, which is known to be corrupted by artificially high values, skewing the average and leading to a pedestal file that does not produce the desired result when applied to the data. This is demonstrated in Figure 4.2; if the pedestal file worked, the values on the right hand side would all be close to 0 ADC. Instead, the pedestal-corrected mean value is offset from 0 ADC by an amount dependent on the pedestal file itself, not correlated with the actual raw ADC value recorded across the sample.

The alternative to this is to create new pedestal files for each run individually, by recording the average value of each sample based on the data within each run, rather than using the pedestal files generated by the DAQ software. This method has the advantage of being able to exclude the corrupted data from the first block from the average calculation.




Figure 4.3: An example of two blocks of calibration data, with the targeted region identified.

4.4 Timing Calibration

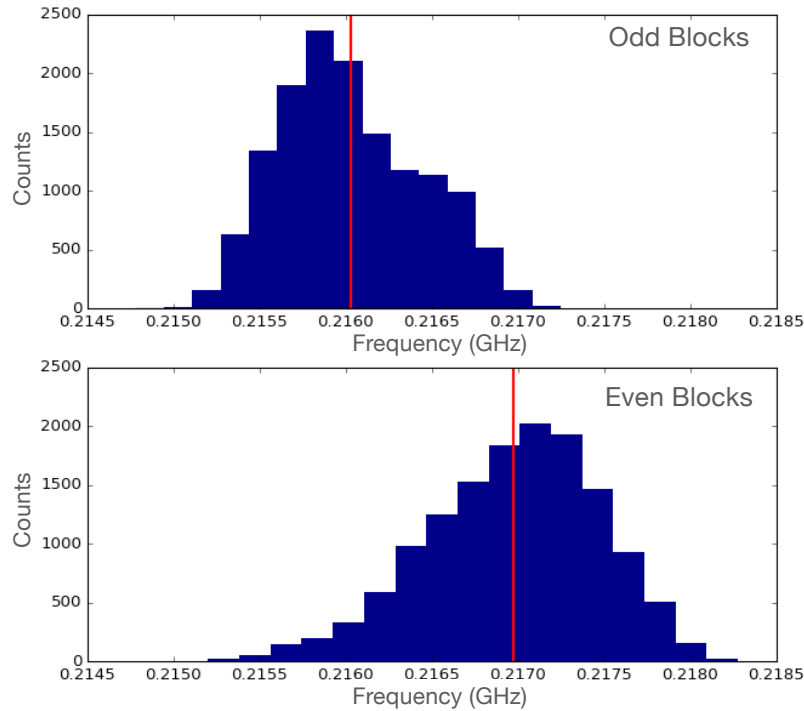


Figure 4.4: The distribution of frequency fits for all calibration events. The clock speed is scaled so that the mean value matches the true frequency of 218 MHz.

As mentioned in the previous section, sine wave calibration data was taken for 14 out of the 16 RF channels. The result is approximately 6000 events per channel per frequency, with each event 896 samples long. The timing calibration requires the 128 elements in the sampling array to be individually calibrated. In the following sections, the steps of timing calibration are described, which are completed up to five times recursively to ensure that the best result is found. The goal is to find the true timing values, t_i , for each of the 128 samples. The nominal starting values for t_i are $t_i = 0.3125 * i$ where i ranges from 0 to 127.

4.4.1 *Correcting the Overall Frequency*

Each event was individually fit to a sine wave using the function

$$f(t) = A * \sin(2\pi kt - \phi) \quad (4.1)$$

where A is the amplitude, k is the frequency, t is the time, and ϕ is the phase offset. Because the voltage is corrected after the timing is corrected, the amplitude of the waveform in ADC is chosen to be equivalent to the expected amplitude in mV. This works because the relationship between ADC and mV is approximately 1:1 in the range of 400-500 mV, which is what is used in the timing calibration fit.

Both the frequency and the phase of the sine wave are fit using SciPy's `optimize.curve_fit`, a non-linear least squares regression. While the phase is expected to vary across events, the frequency is expected to be a distribution centered around a single frequency, as all events should be measuring a sine wave with a frequency matching the input frequency of 218 MHz. The spread is caused by the timing jitters that have yet to be corrected. This is seen clearly in Figure 4.4.

For all even and odd blocks, the frequency of the sine fit is recorded, and this average frequency is used to scale the timing array:

$$t_i = t_i * \frac{218MHz}{f_{avg}} \quad (4.2)$$

This corrected timing array is then carried to the next step.

4.4.2 *Correcting the Residual Timing for each Sample*

After the overall frequency has been corrected, each event is fit again to a sine function. This time, the even and odd blocks are fit simultaneously, using only the odd sample on each type of block, as these are known to be the most stable of the samples and introduce

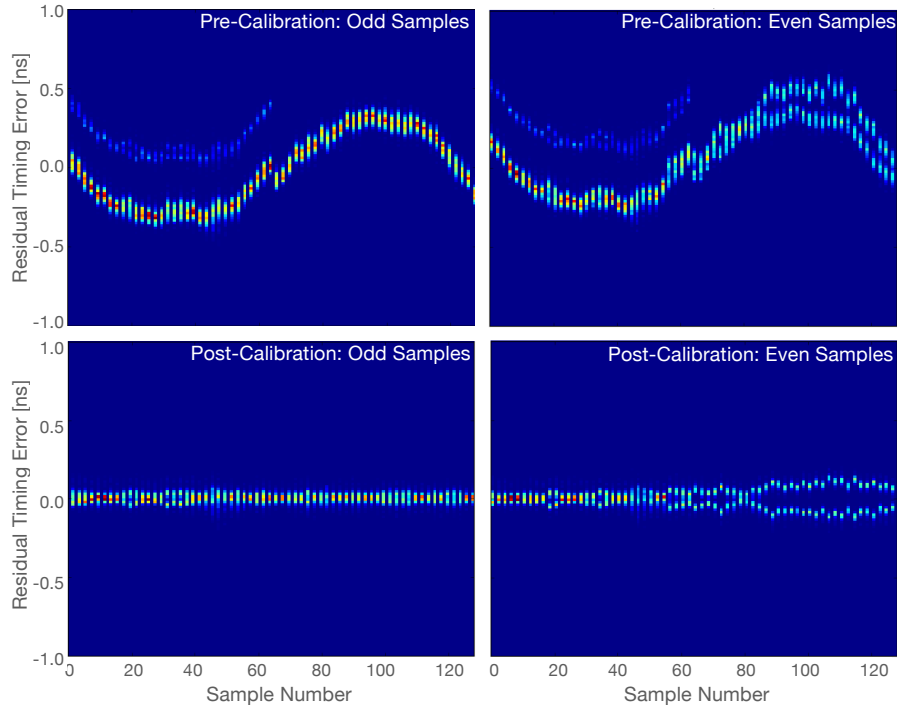


Figure 4.5: Uncalibrated and calibrated residual timing for electric channel 0. Before calibration, the samples are offset from the sine fit, causing a systematic offset every time that sample is measured. After calibration, the offset is approximately 0 ns.

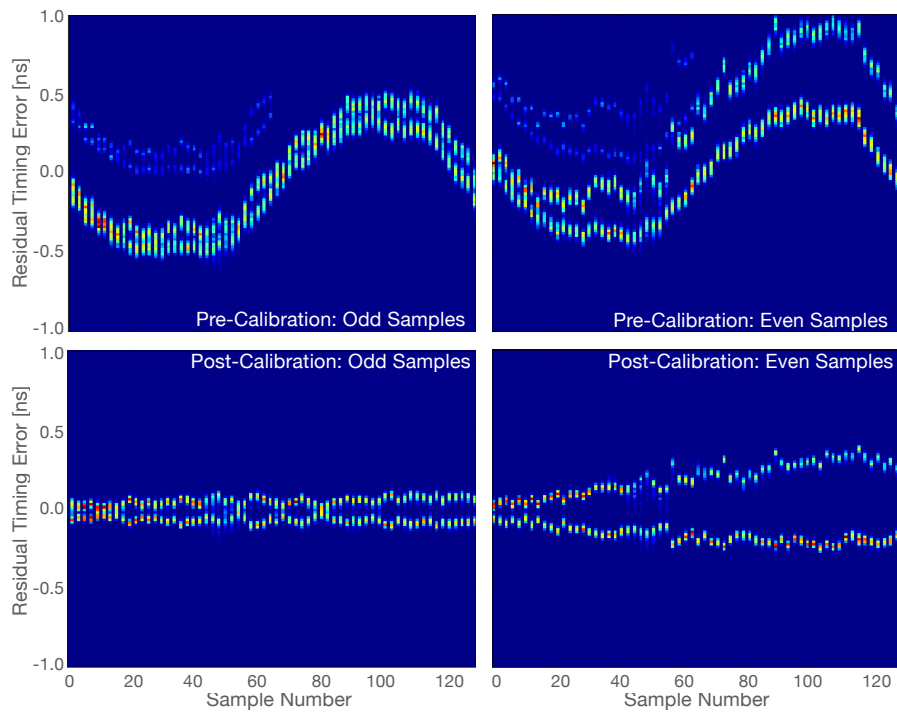


Figure 4.6: Uncalibrated and calibrated residual timing for electric channel 1.

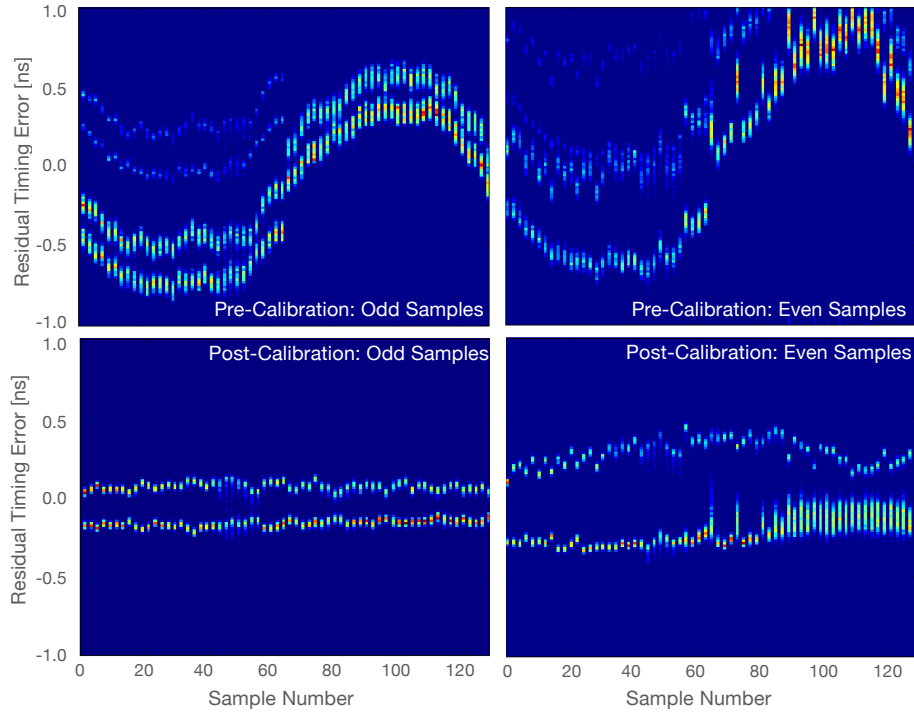


Figure 4.7: Uncalibrated and calibrated residual timing for electric channel 2. The even samples are discarded due to poor behavior.

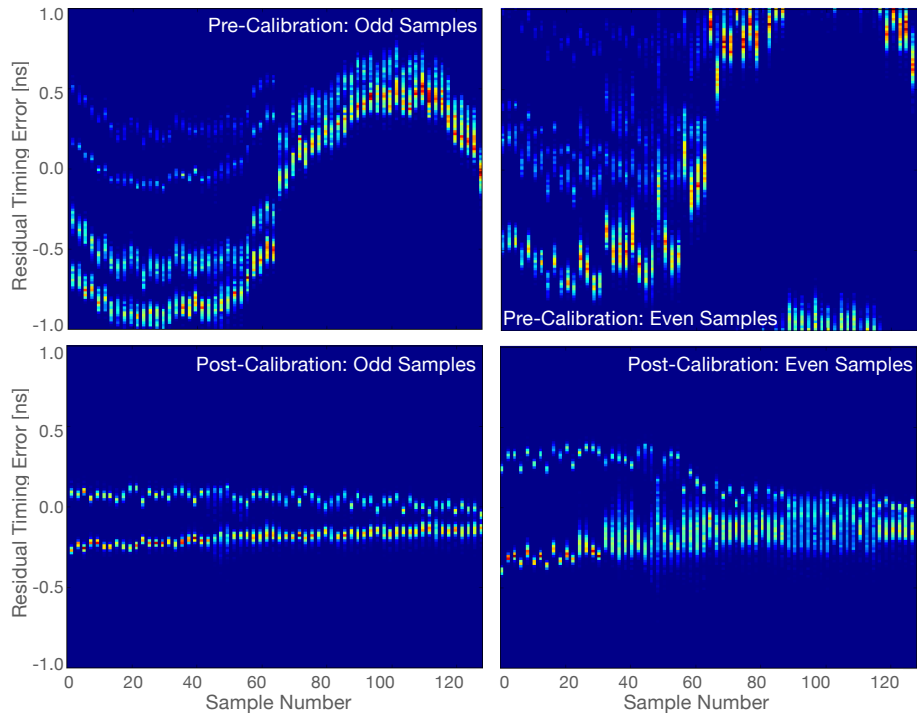


Figure 4.8: Uncalibrated and calibrated residual timing for electric channel 3. The even samples are discarded due to poor behavior.

less overall error. Recall that at this point, both the timing and voltage has to be calibrated. The shape of the sine wave means that voltage errors will be most prominent at the peaks and troughs, and least prominent near the center where the slope is greatest. Because of this, any sample that is recording an event within 30 ns of 0 ADC is selected to be used in this step, as shown in Figure 4.3

The recorded time of these samples is compared against the time expected from the sine wave fit. The difference between observed and expected is recorded and organized by sample number, and repeated for all events until every sample has accumulated a distribution of offsets between observed and expected.




Figure 4.9: An example of extreme bimodality for electric channel one. Left: the calibrated residual timing for channel 1 as a function of sample number. Right: the calibrated residual timing as a function of the measured slope. The split on the left hand side is caused by the slope shown on the right.

For a well behaving channel, the distributions of offsets for each sample will be generally Gaussian, which can be seen in Figure 4.5. For even samples, the distribution is not Gaussian

and instead shows some bimodality in the later samples. As in previously calibrated stations, the bimodality is caused by a slope dependence, in which the timing depends on the slope of the measured signal. This can be seen clearly in Figure 4.9. The underlying cause of this bimodality is unknown, but seems consistent across multiple IRS2 chips across multiple generations of ARA stations. This introduces a maximum error of 0.4 ns for some of the even samples, and generally is closer to 0.1 ns or less.

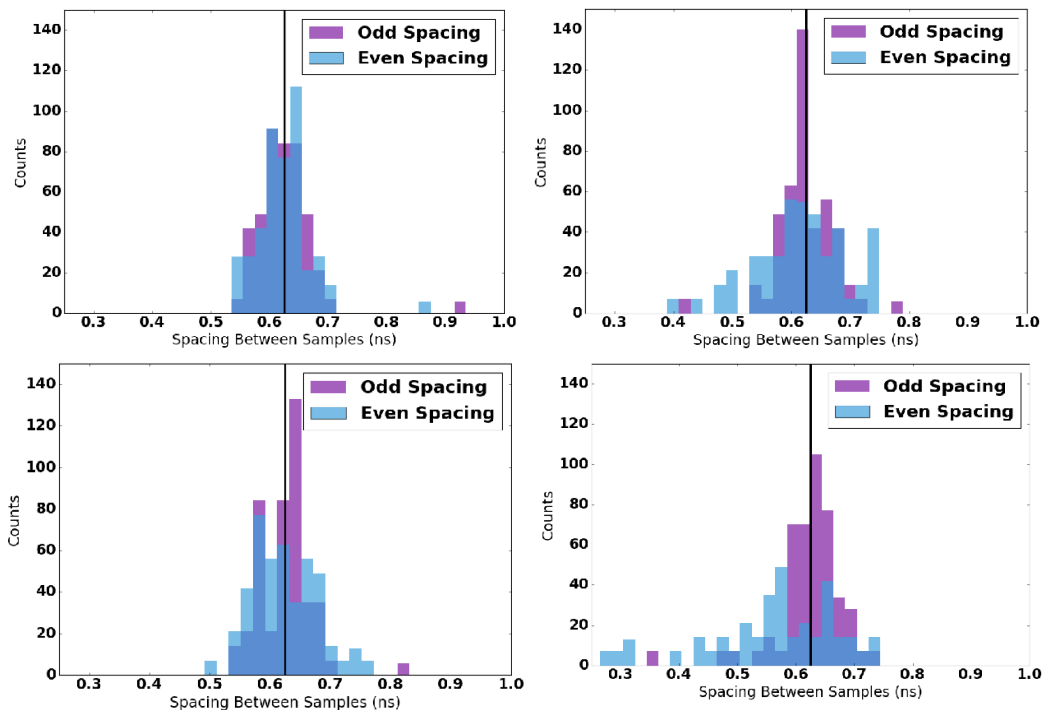


Figure 4.10: The time between odd and even samples, for channel 0 (top left), channel 1 (top right), channel 2 (bottom left), and channel 3 (bottom right). The nominal spacing is 0.625 ns.

The mean value of these distributions is calculated and added to the time array. After all samples have been adjusted, the process starts over, re-fitting the overall frequency, then adjusting individual time delays, in a recursive loop until the values have settled. In general, this recursive process takes anywhere from 1-5 times to settle.

The final distributions of offsets after the recursive process has finished is shown in Figures 4.5 through 4.8 in the bottom two panels. Additionally, the delays between odd

samples and even samples is shown in Figure 4.10; here it is clear that the delays between samples are generally close to their expected values of 0.625 ns. The degradation of the channel performance over time, discussed in the earlier part of this chapter, is seen very clearly, especially for even samples on channels 2 and 3. The even sample performance completely fails for channels 2 and 3, and the odd samples begin to show increasing levels of bimodality. To account for this, the even samples on channels 2 and 3 on all IRS2 chips are discarded, meaning that the HPol channels all use a sampling speed that is half that of the VPol channels.

For the code used to calibrate the timing of ARA Station 5, please see [61].

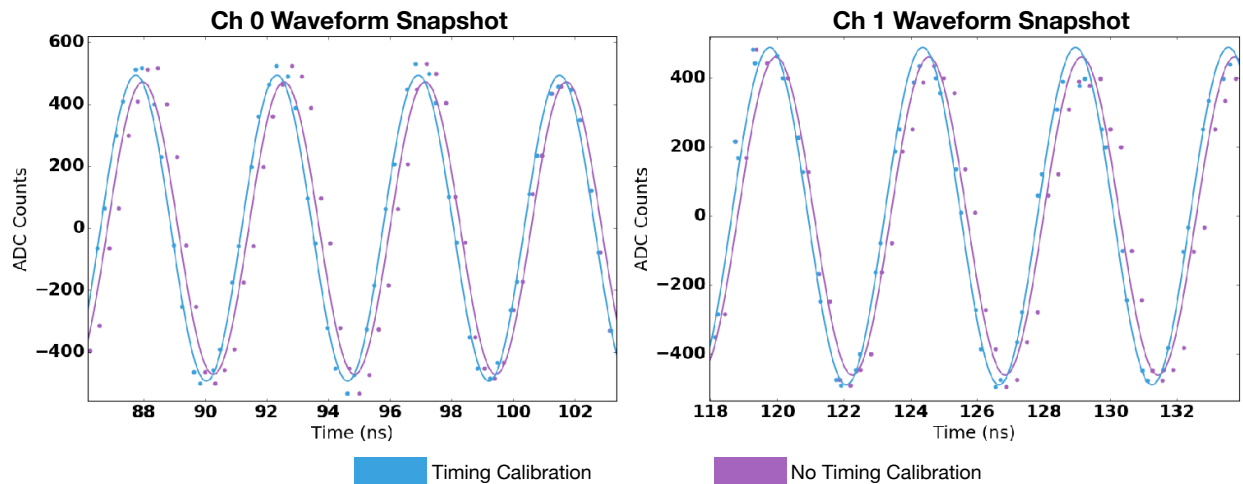


Figure 4.11: Uncalibrated data vs. calibrated timing data, for channel 0 and channel 1. The blue points fit especially well near 0 ADC, where the uncalibrated voltage has the smallest effect.

4.5 Voltage Calibration

The resulting timing calibration fit is shown in Figure 4.11, using Electric Channels 0 and 1 as examples. The absolute value of the points has shifted due to the frequency correction and the individual sample corrections, even though the frequency of the calibrated and uncalibrated sine waves only differ by less than a MHz. The timing calibration looks most successful near the center of the sine wave, as expected. The next step is to calculate the voltage.




Figure 4.12: Example odd and even samples, comparing the ADC counts measured and the predicted voltage from the sine fit. In each example, the positive and negative points are fit with a cubic polynomial.

Unlike the timing calibration, which needed to be done for all 128 sampling array capacitors, the voltage calibration must be done for each of the storage samples, of which there are 32,768. To characterize the behavior of each sample individually, the calibration data must first be converted from event-by-event sorting to sample-by-sample sorting, creating




Figure 4.13: (a) An example sample with looping behavior. (b) The same sample, only plotting the positively-sloped data. (c) The same sample, only plotting the negatively sloped data.

a database of measured ADC counts and predicted voltage values for each of the 32,768 samples. The predicted voltage value is calculated by fitting the events to a sine wave, using the amplitude from Table 4.2.

The result of sorting the calibration data based on sample number is shown in Figure 4.12. As with the timing calibration, there are inherent differences between the performance of odd samples and even samples. Figure 4.12 shows an example odd sample and an example even sample for reference. In this plot, the relationship between ADC counts and Voltage looks to first order linear, especially for the odd sample. However, the non-linearity near 0 ADC counts, previously seen during the calibration of ARA Stations 2 and 3, motivates using a cubic fit to describe the relationship between ADC and Voltage. This cubic fit is expected to be valid for the range used in the fit: -445 mV to 445 mV. Any ADC counts recorded outside of this range are instead converted to voltage assuming the relationship is 1 ADC count per 1 mV.

For each sample, the positive and negative values of ADC are fit with their own cubic function. The reduced χ^2 for each fit is calculated, with the resulting distribution shown

in Figure 4.14. For some samples, the χ^2 is quite large, sometimes reaching values of 10 or greater. This could be caused by low statistics or unpredictable behavior, but most often




Figure 4.14: The distribution of χ^2 for all 32,768 samples. Any fits with $\chi^2 < 1$ are kept; fits worse than this are substituted with the fit from the neighboring sample.

this is caused by slope dependence impacting the relationship between ADC and voltage to shift. An example of the slope dependence is shown in Figure 4.13. If samples record a χ^2 greater than 1, the fit from the neighboring odd sample is used. This is a reasonable substitution; due to the way the IRS2 chip was constructed, neighboring samples are likely to behave in similar ways. Approximately 10 percent of the samples use the fit from the neighboring sample.

After each of the 32,768 samples have been calibrated, the results are formatted into a file that can be read by AraRoot, the processing software that the ARA Collaboration uses to read in the data. The improvement after completing the voltage calibration is shown in Figure 4.15. The voltage calibration, perhaps unsurprisingly, does not change the parameters in the sine function fit. Instead, only the values of the voltage have changed, a feature most notable at the peaks and troughs of the waveform.




Figure 4.15: A comparison of timing-calibrated data and timing and voltage-calibrated data. Especially notable are the data points at high ADC values, which have been adjusted to proper voltage values.

4.6 Validating the Timing and Voltage Calibration




Figure 4.16: The timing delays between calibration pulser events comparing channels 0 and 16.

The final step in calibrating the IRS2 digitizers is validating the calibration against calibration pulser events. The easiest way to do this is to cross correlate the signals on two separate channels to calculate the best delay. When comparing the results before calibration to the results after timing and voltage calibration, an improvement in the distribution is expected. This improvement can be seen in Figure 4.16. More details on how the cross correlation is calculated will be included in the next chapter.

The time delays across all channel pairs can be seen in Figure 5.11. The channel pairs that perform the worst and the pairs that include Channel 24, which was the one VPol channel not calibrated. These pairs have a standard deviation of 0.1 ns or larger. The other channel pairs consistently have a standard deviation below 0.09 ns. All of these uncertainties are carried through to the uncertainties in the geometry calibration, discussed in the next

chapter.

4.7 Future Improvements

The IRS2 digitizer chip can be difficult to calibrate due to its degradation of channel performance, slope dependence, and frequency dependence. At the time of deployment, this digitizer chip was one of the only options that met the design requirements of a high sample speed, long buffer, and low power consumption. Since deployment, other radio projects have opted to use other digitizing options, from commercially available digitizers [12] or other custom built digitizers like the LAB4D [62]. I would recommend continuing to pursue these other digitizing options for future experiments.

CHAPTER 5

GEOMETRY CALIBRATION OF ARA STATION 5

5.1 Introduction

5.2 Determining the ARA Antenna Locations

The successful calibration of the antenna locations is critical to the overall success of an ARA Station. In the frequency range of 200-800 MHz, the individual antennas must be known to within 10 cm in all directions in order to phase together antennas in different locations, as shown in Figure 5.1. While the borehole locations are surveyed and the distances between antennas are measured during deployment, these measurements are often unreliable or incomplete. This story is further complicated at ARA Station 5, in which two separate instruments, the Classic ARA Station and the Phased Array, must be calibrated relative to each other.




Figure 5.1: Left: the relationship between the fractional wavelength offset and the sum between two antennas. The horizontal line indicates the mean incoherent addition. Right: The two antenna sum as a function of frequency and displacement. Areas with a sum greater than the mean incoherent addition value will add coherently. From C. Deaconu.

To fully appreciate the constraints of the problem, it is helpful to list the quantities that

must be solved for, and the quantities that are known. The quantities that must be solved for include: the (x, y, z) of each of the 16 ARA antennas and each of the 9 Phased Array antennas; the local ice model of the ice in the South Pole; the (x, y) of each of the calibration pulse locations; and the cable delays for each of the antennas on each of the instruments. To simplify the problem, I consider only the VPol antennas for each system, and assume that because the HPol antennas from the Classic ARA Station are all deployed 1 m above their respective VPol antennas at a distance that is measured prior to deployment, solving for the VPol antennas will effectively solve the HPol locations as well.

The known quantities in the system are limited to the sets of time delays that can be measured between channels. For a given pulse originating somewhere in the local environment, the path along which the signal travels will determine the time delay seen between any two channels. It's important to emphasize here that looking at a single channel is not enough to determine the time it took to travel from the source to the antenna; the only information is the difference in travel times between two channels.

Ideally, an ARA Station would have three local calibration pulsers installed near the experiment, each deployed in their own holes surrounding the array, with the calibration boreholes forming an equilateral triangle around the ARA boreholes. In this way, the location of each antenna could be easily determined by what is essentially triangulation. This method is especially advantageous because it depends very little on the ice model used to describe the local ice. However, three independent holes used only for calibration is costly, and instead previous ARA Stations have conducted geometry calibrations by using only two local calibration pulsers, set perpendicular to each other.

The ARA Station 5 system was deployed with only one calibration hole, due to delays and difficulties during drilling. This single calibration source is not sufficient for constraining the locations of the antennas. There are other calibration sources at the South Pole, namely, the deep pulser deployed on an IceCube string and the SPIceCore hole, both visible in the map

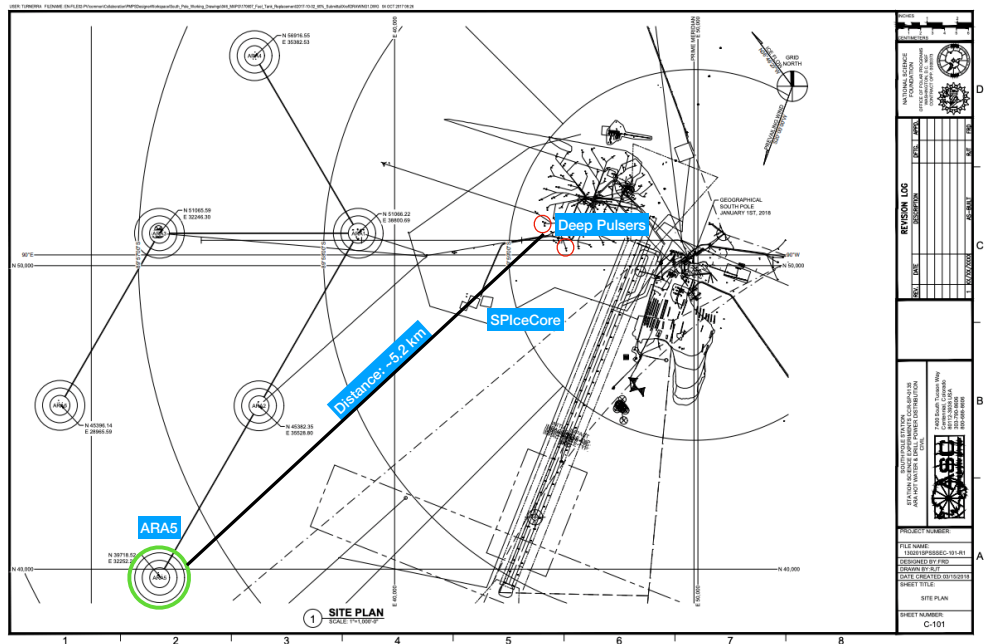


Figure 5.2: A map of places of interest near ARA Station 5.

in Figure 5.2. Both of these sources are 4-5 km away, necessitating the precise determination of the ice model.

Considering these limitations, some assumptions needed to be made to constrain the system fully. The first assumption was that the depths of the VPol antennas are perfectly known. This is a reasonable assumption, due to the fact that the Phased Array VPol antennas are deployed within 1 m of each other, with the exact distance between them measured to centimeter precision. Additionally, any error here is expected to be absorbed into the determination of the cable delays for the Phased Array. The next assumption is that the depths reported at the SPIceCore hole are accurate as measured by the SPIceCore team in December 2018 [63]. The final assumption is the SPIceCore borehole is relatively perpendicular to the ARA5 boreholes.

Using these assumptions, the following plan was developed for calibrating the ARA 5 Station:

1. Use the SPIceCore data recorded by the Phased Array to simultaneously determine the ice model and the cable delays.
2. Use the SPIceCore data recorded by the classic ARA antennas, and the calibration pulsers recorded by the classic ARA antennas, to simultaneously determine the (x, y, z) and cable delays for each of the ARA VPol antennas.
3. Use these channel common to the Phased Array and the classic ARA antennas to determine the location of the Phased Array borehole relative to the classic ARA boreholes.

In the following sections, this plan will be described in more detail.

5.3 Step 0: Defining the Calibration Dataset

Before beginning the calibration, a dataset was selected from the available calibration runs.

The available dataset includes:

The Local Calibration Pulser: this pulser is locally deployed near the ARA5 station and is set to trigger every second. Any run with calibration pulsers can be used for calibration, although there is evidence that the time delays between channels drifts by up to 50 ps over the course of the first two years of deployment. Because of this, a calibration run close in time to the SPIceCore runs was used for geometry calibration. The SPIceCore Pulser Drop: During the 2018-2019 Antarctic season, the SPIceCore borehole was used to take calibration data at a multitude of depths. The main difficulty during data taking was the Phased Array kept overheating during the calibration runs, meaning that it needed to be shut off intermittently to cool down. The highest quality dataset for both the Phased Array and the classic ARA Station was taken on December 24, 2018, and is associated with run 1784. The Deep Pulser: There are multiple deep pulsers installed onto two IceCube strings which are capable of pulsing in radio, although only the one on string 22 has operated for the last few years. A dataset from May 2018 containing deep pulser runs was used as validation to the

ice model fit; however, because this dataset was taken before there was a common channel, and because it is redundant to the SPiceCore (due to the two pulsers being in almost the identically same azimuthal direction) was it was not used for the later fits. The Rooftop Pulser: Rooftop pulsing data was taken in December 2018 as well, also pointing in the same direction as the Ice Cube Lab and the SPiceCore. The pulses from this dataset are used later on in the analysis, but are not used for calibration, partly because they have low SNR and thus are not ideal candidates for reconstruction.

5.4 Step 1: Determining the Phased Array Cable Delays and the Ice Model

Description	Easting [ft]	Northing [ft]
Phased Array	32355.0	39734.5
D1	32337.6	39795.1
D2	32420.0	39745.5
D3	32361.5	39654.5
D4	32290.0	39710.6
Cal Pulser	32272.2	39582.8
SPiceCore	42358.9	48974.2
Deep Pulser	44884.4	5144.8

Table 5.1: The list of pre-deployment locations of the A5 boreholes.

Focusing first on solving for the local ice model and the VPol cable delays of the Phased Array instrument cuts down the number of unknowns from 103 to 11. The remaining unknowns include 7 cable delays, 3 ice model parameters, and the lateral distance between the Phased Array and the SPiceCore hole. In this case, only the SPiceCore data is used because the signal propagates from multiple depths, meaning that the changes observed throughout a SPiceCore data drop are directly attributable to the ice model and not, for example, the cable delay, which is a constant offset. The first step is to calculate the raw time delay be-

tween the direct signal for each channel pair for all depths of the Phased Array, and compare those time delays to those predicted by the Raytracer using the position of the SPIceCore and Phased Array boreholes available in Table 5.1. This can be done by considering the correlation waveform defined as

$$\text{Corr}_{i,j}(\tau) = \sum_{t=0}^{N-1} \frac{V_i(t)V_j(t+\tau)}{\sigma_i\sigma_j}, \quad (5.1)$$

The time delay is the location of the maximum of this waveform. The set of time delays from the SPIceCore can be seen in Figure 5.3. As depth increases, the time delay measured across each antenna pairs will approach an asymptotic value equal to the time it takes for a signal to travel between the antennas.




Figure 5.3: A comparison of the raw calculated time delays from the Phased Array instrument, compared to the predicted time delays from the AraSim Raytracer. The mismatch is due to cable delays not yet being applied to the raw data.

The difference between the raw time delay and the predicted time delay in Figure 5.3 is caused by the cable delay not being accounted for in the plot. A cable delay can be added

to this plot by shifting the raw time delays horizontally by a fixed value. The difference between the observed and expected time delays increases as the distance between the pair of antennas increases, due to the longer cable delays present for the deeper channels.

It is also important to note that the absolute cable delays cannot be determined independently. Instead, the relative time delays can be determined, with some unknown bulk cable length not known, a quantity which is not necessary for pointing or reconstructing events. In this work, I report cable delays relative to the Phased Array’s channel 0.




Figure 5.4: Left: comparison of the various ice model fits, after correcting for the cable delay offset. Most of the models do not accurately predict the shape of the curve as a function of depth. Right: same as the left, after fitting the third parameter of the UNL 2016 model to better match the data.

After compensating for the missing cable delay by sliding the predicted time delays horizontally by a constant value, it is easy to compare the performance of various existing ice models, which can be seen in Figure 5.4. The ice models are listed explicitly in Table 5.2. Because some of these models are not linear, I require that the data and model overlap at a depth of -1400 m. The model that matches the data most closely is UNL 2016, although

even this model does not describe the data at shallower depths well. Because UNL 2016 is the best match, I use two of the three parameters (parameters A and B in Table 5.2) in this model and float the third parameter, optimizing it for the best χ^2 using the Channel 0/Channel 7 pair as comparison. I find the best value of C to be 0.0202.

Model Name	A	B	C
AraSim	1.78	1.35	0.0132
Gorham	1.788	1.325	0.0140
AraRoot	1.78	1.353	0.0160
UNL (2016)	1.78	1.326	0.0182
RICE (2004)	1.78	1.36	0.0132
SPIceCore 1	1.774	1.293	0.0154
SPIceCore 2	1.774	1.249	0.0163

Table 5.2: A list of previously-derived ice models for the South Pole.




Figure 5.5: The measured time delays against the predicted time delays using the modified UNL 2016 model.

This model works well on all antenna pairs, which can be seen in Figure 5.5. Because the

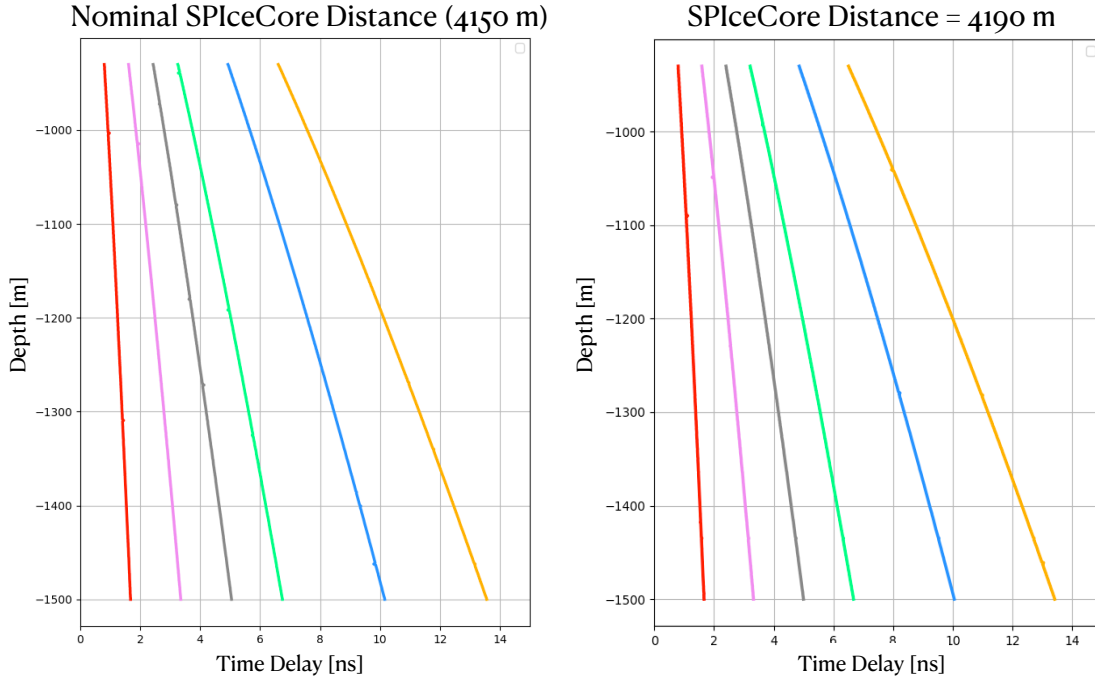


Figure 5.6: A comparison of the predicted time delays after moving the SPICECore hole by 40 m. The horizontal offset shifts by less than 0.2 ns for the longest baseline, and the shape of the curve stays the same.

ice model mainly impacts the shape of the function, I assume that the cable delay accounts for the entirety of the offset present in the dataset. This means that the ice model and the cable delays have been found simultaneously.

The uncertainties on the parameter C in the new ice model mainly come from the uncertainty in the lateral distance between the SPICECore borehole and the Phased Array borehole. There are conflicting reports from the South Pole surveyors and the SPICECore researchers about exactly where the SPICECore hole is, causing the lateral distance to differ by around 40 m. This has a minor effect on the ice model fit; the impact of moving the SPICECore hole by 40 m is shown in Figure 5.6. To estimate the uncertainties, the fit was repeated for lateral distances ± 25 m and resulting value of C with uncertainties was found to be $0.0202 \pm 0.0004 m^{-1}$.

This result was validated using data from the IceCube Deep Pulsar. While the phased

array cannot resolve distance on its own, a χ^2 calculation as a function of distance and depth has a minimum at a distance and depth of 5244 m and -1379 m respectively, only a few meters off from the true distance and depth of 5240 m and -1386 m. This is the most successful reconstruction of the deep pulser data with the ARA instrument so far.

5.5 Step 2: Determining the ARA Antenna Locations

After successfully determining the ice model from the Phased Array instrument, I have enough information to locate the ARA antennas within the ice. The first step is to methodically build the set of time delays that will be used for the fit, described below:

Local Calibration Pulser: An averaged local calibration pulser is shown in Figure 5.7. The different pulse shapes are caused by the different viewing angles of both the transmitting and receiving antennas, and by differences in impedance matching between the different antennas. This can make cross correlating the different antenna signals difficult, especially for the lower SNR channels. In particular, cross correlating individual waveforms will often select the wrong cycle to align, causing a difference in the time delays of a few nanoseconds.

For geometry calibration, a different method of cross calibration was developed, taking advantage of the higher SNR averaged pulses to better align the signals on different antennas. To illustrate this, I will consider two channels, Channel 0 and Channel 1. The first step is to create averaged pulses for each Channel using multiple events, boosting the SNR for each individual channel. From here, the cross correlation between the averaged Channel 0 and averaged Channel 1 signal is calculated and the time delay δ_{01} between the channels is calculated. Next, the time delays between each individual Channel 0 event and the averaged Channel 0 event are calculated, as are the time delays between each individual Channel 1 event and its average, creating two lists of time delays that each describe the jitter in trigger timing between the average signal and the individual events. Finally, the difference between these two lists of time delays is calculated, creating a single histogram that describes the

spread of time delays for the Channel 0 / Channel 1 antenna pair. The mean of the histogram will be 0 ns if δ_{01} is exactly right; otherwise, it will be slightly smaller or greater. This mean is subtracted from δ_{01} to get the true mean value of the delay, and the standard deviation

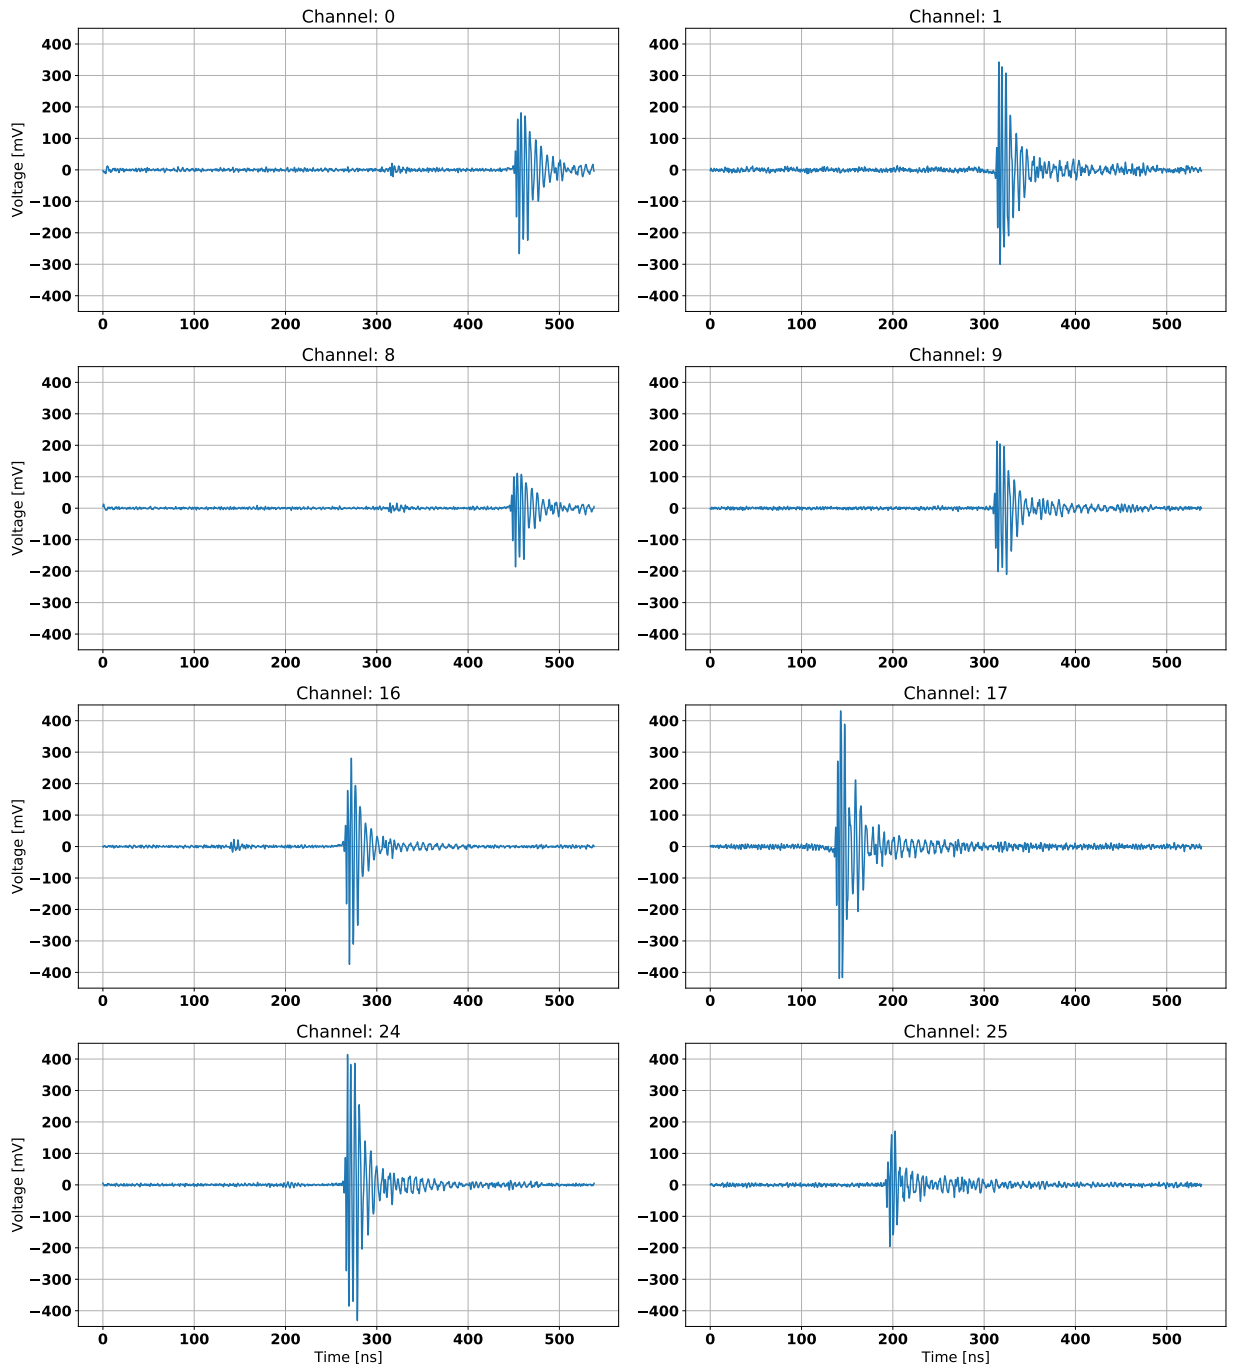


Figure 5.7: Averaged calibration waveforms for each of the eight Vpol antennas.

of the histogram is used as the error.

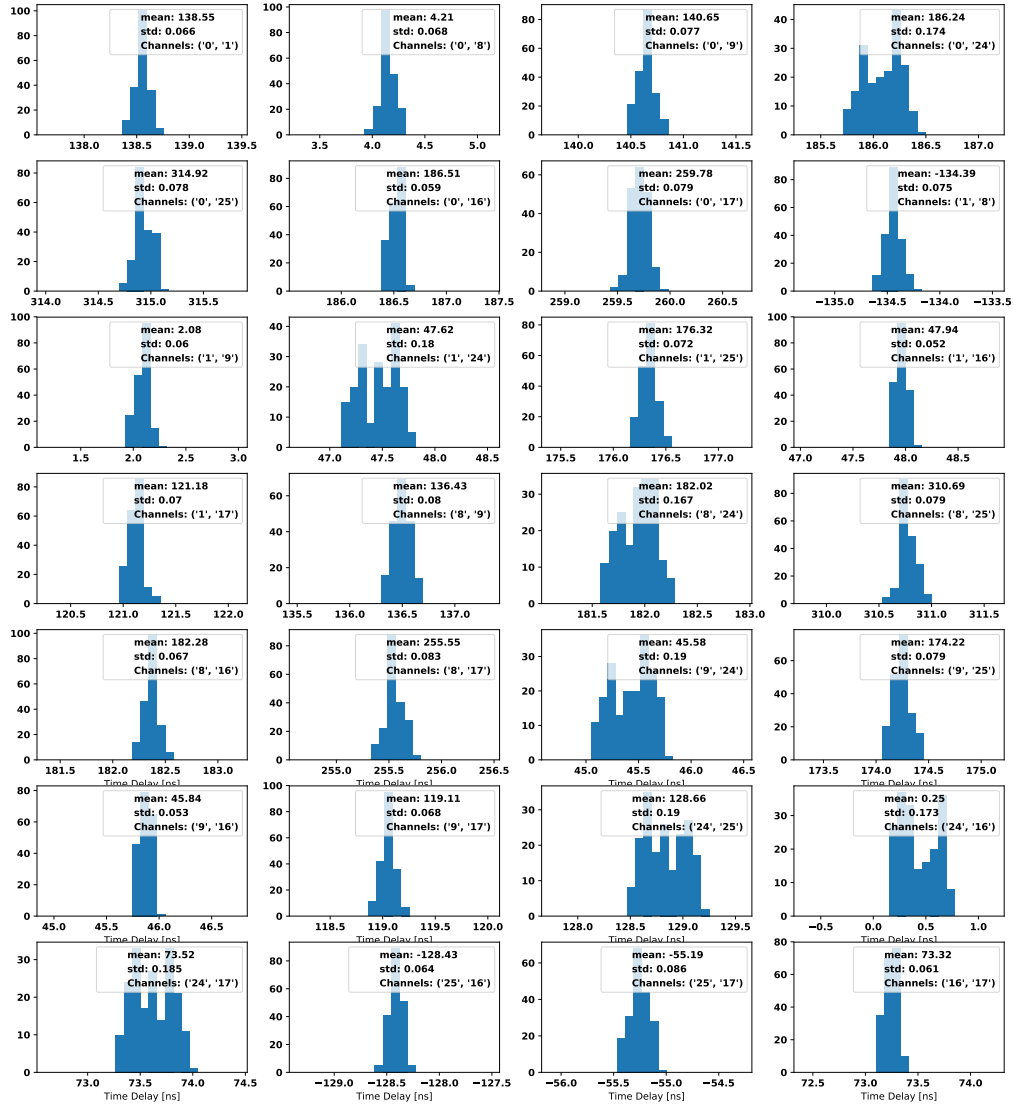


Figure 5.8: The distribution of time delays across VPol channels for a given set of calibration pulsar events. Channel 24 was not calibrated due to missing calibration data, resulting in its distributions being wider than the others.

Using this method of time delay calculation results in the histograms shown in Figure 5.11. This method of cross correlating only works for a stationary pulse that occurs multiple times, allowing it to be easily averaged together. The resolution achieved here is not directly transferrable to resolution in azimuth or elevation. However, for geometry calibration, where the most important feature is to have the correct time delay and its uncertainty, this method gives more confidence that the correct peak is selected.

SPIceCore Calibration Pulses: While the local calibration pulser can be averaged together to get a clearer signal, the SPIceCore calibration pulses happen as the pulser is dropping, meaning that pulses cannot be averaged together in the same way. Instead, the time delays for all SPIceCore events were calculated individually for each channel pair and plotted as a function of depth, seen in Figure 5.9 as an example for Channel 0 and Channel 1. The outliers here are caused by the algorithm choosing to align the wrong peaks, as discussed previously.




Figure 5.9: An example of the changing time delays between Electric Channels 0 and 1 for the ARA instrument. The red points represent the interpolated values used in the Minuit fit. The outliers are caused by the cross correlation being maximized between the wrong peaks.

Depending on the channels, there can be anywhere from a handful of outliers to two lines running parallel to the center line at all depths. To select data for the optimization, I select data points between -1050 m and -1450 m that are representative of the values along the true peak. These values are stored in a lookup table for every antenna pair.

After the time delay lookup tables for both the local calibration pulser and the SPIceCore pulser have been created, a Minuit optimizer is used to compare the time delays in the lookup table to what is predicted by the raytracer using the ice model found in the earlier section. The function that is optimized is the following:

$$\chi^2 = \sum_{pairs} \frac{(\Delta t_{1,2} - RayTracer(r_1, r_2, z_1, z_2))^2}{\sigma^2} \quad (5.2)$$

where r_1 is the lateral distance between the source and the first antenna, z_1 is the depth of the first antenna, r_2 is the lateral distance between the source and the second antenna, z_2 is the depth of the second antenna, σ is the error calculated from the standard deviation of the histogram, and $\Delta t_{1,2}$ is the measured time delay between the two antennas, corrected for the difference in cable delays between the two. r_1 and r_2 need to be calculated from the x and y values, which are the quantities used as inputs to the optimizer. I also assume that the depth of the calibration pulser is known based on deployment logs.

Minuit requires initial conditions to best find the fit for a given system. Because of this, two rounds of Minuit fitting are performed: one forcing the boreholes to be exactly vertical, and another using the output of the first round and then removing the requirement that the boreholes be vertical. The result is shown in Figure 5.10.

5.6 Step 3: Use the common channel to find the location of the Phased Array relative to the ARA antennas

At this stage, the ice model, the x , y , z and cable delays for each of the ARA VPols, and the cable delays of the Phased Array have been solved for. The final step is finding the x and y position of the Phased Array borehole relative to the rest of the ARA5 instrument. This is done by generating the same lookup tables that were created for the ARA instrument but using the Phased Array. An example of an averaged Phased Array calibration pulser and its timing distributions for each channel pair is shown in Figure .

Because the Phased Array and the ARA instrument are separate triggers, there is not typically a way to combine information across systems. However, in December 2018, electric channel 24 in the ARA instrument was split and sent into both instruments. This means that electric channel 24 on the ARA instrument and channel 5 on the Phased Array instrument




Figure 5.10: The calibrated locations of the A5 antenna boreholes, compared to the planned locations.

are the same, with some potential offset that depends on exactly when the trigger happened within the window. This allows time delays to be computed across the systems.

As an example, consider Channel 0 on the Phased Array system and Electric Channel 17 on the ARA system. Cross correlating these channels without any additional steps would not result in a physically meaningful value, because the trigger is not necessarily in comparable places for either event. Instead, first Channel 0 would be cross correlated with Channel 5 on the Phased Array to get a time delay. Then, Channel 24 on the ARA instrument would be cross correlated with Channel 17. As long as the cable delays of Channel 0 and Channel 17 are accounted for, the sum of these time delays will correspond to the time delay of Channel 0 with Channel 17. In this way, the lookup tables from both the ARA instrument and the Phased Array can be combined to have time delays for every pairing of antenna. This dataset is fed through a Minuit optimizer in much the same way as Step 2, with the x y of

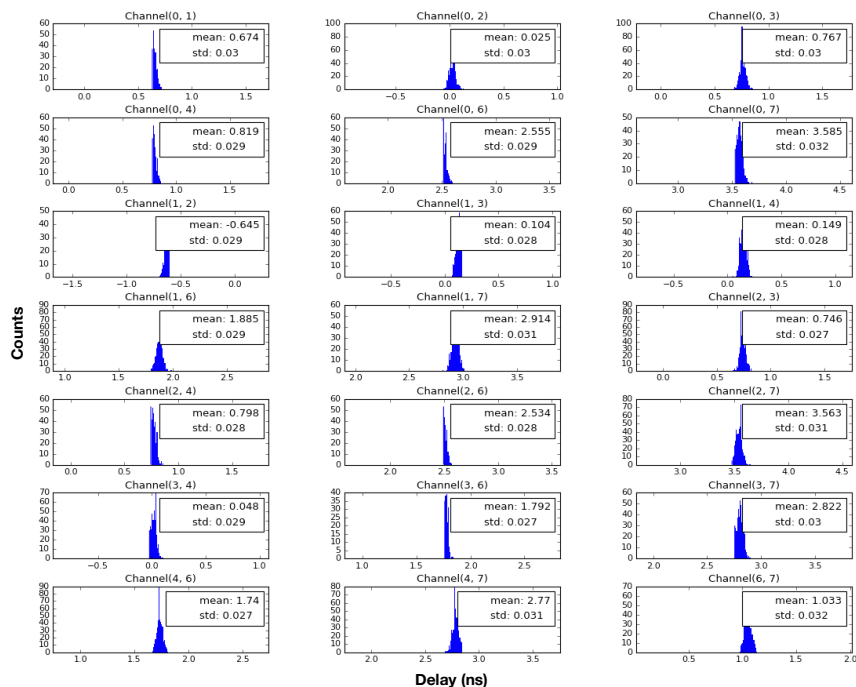


Figure 5.11: The distribution of time delays across VPol channels for a given set of calibration pulser events. Channel 24 was not calibrated due to missing calibration data, resulting in its distributions being wider than the others.

the Phased Array and the z of the local calibration pulser allowed to float.

The final locations of each of the antennas, and the cable delays, are shown in Table 5.3. Minit is able to report the error on the various values. Each value was reported to have an uncertainty of less than 10 cm, with the majority of parameters being less than 5 cm. This meets the requirements to use full phasing for correlation maps and analysis variables in the future.

5.7 Event Reconstruction

After successfully calibrating the full ARA5 instrument, the next step is to measure the pointing resolution of the instrument. Specifically, this process is done by measuring the azimuth and elevation precision. Unfortunately, there are no sources that have not already been used for calibration in some way, so it is difficult to perform an external validation of the calibration. Instead, the local calibration pulser is used, including data taken in a lower SNR mode. Examples of the low SNR pulses are shown in Figures 5.12 and 5.13.

When considering the pointing resolution, it is necessary to introduce the concept of correlation maps, which will be a crucial theme throughout the rest of this work. I have already discussed how to use cross correlations to calculate the time delay between two waveforms, which effectively tells you the potential directions the signal could be arriving from. However, you can also think about the reverse problem, where you ask what time delay is associated with a given direction, and calculate the strength of the correlation waveform at that time delay. By combining the information from all channel pairs and all directions, a three-dimensional correlation measure can be introduced as the following:

$$C(R, \phi, \theta) = \frac{\sum_{i=1}^{n_{ant}-1} \sum_{j=i+1}^{n_{ant}} \text{Corr}_{i,j}[\tau(R, \phi, \theta)]}{\binom{n_{ant}}{2}}, \quad (5.3)$$

In practice, this is often calculated using only two dimensions by setting the radius of

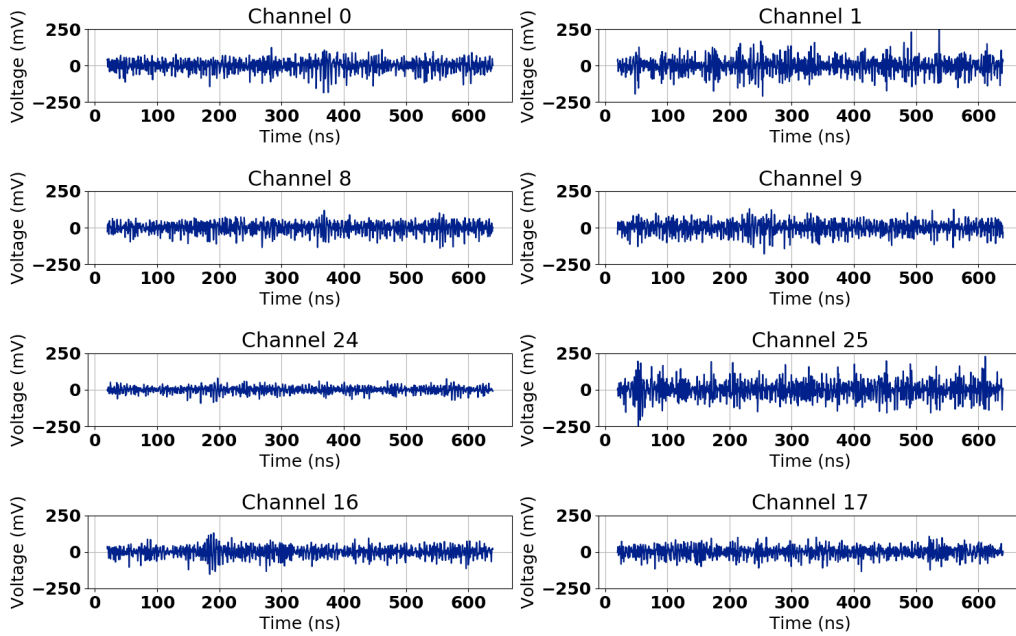


Figure 5.12: An example low SNR event from the local calibration pulser, as seen by the ARA instrument.

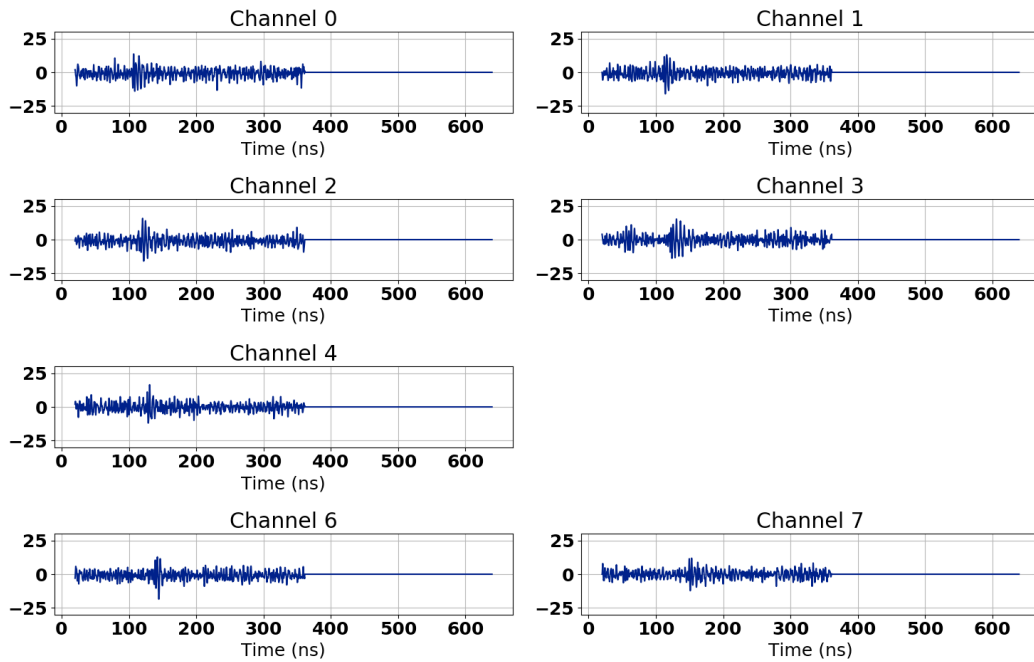


Figure 5.13: An example low SNR event from the local calibration pulser, as seen by the Phased Array instrument.

the map to a constant, usually the radius of the local calibration pulser. This results in a two-dimensional correlation measure referred to as a correlation map.

With a system like ARA Station 5, there are many ways to display the information in a correlation map. Previously, correlation maps for ARA Stations are displayed in terms of azimuth and elevation, and generated for two radii: the radius of the calibration pulser and a far-field radius of 2 km. Additionally, these correlation maps were generated using correlation waveforms that had been Hilbert-transformed, a process that smooths the data to allow for less-precise antenna locations to be used to make less-precise maps. However, for ARA Station 5, the precision of the antenna locations is high enough that the full correlation waveforms can be used, which makes the maps less visually appealing but more precise.

Additionally, for ARA Station 5, correlation maps can be made with or without the Phased Array antennas for events that are triggered by the Phased Array and externally trigger the ARA5 instrument. The most noticeable feature from adding the Phased Array antennas is a horizontal band that brightens at a specific elevation angle, reflecting the fact that the Phased Array is excellent at pointing in zenith. Adding the Phased Array is also visibly helpful when looking at lower SNR signals. Examples of various correlation maps for ARA Station 5 are shown in Figures 5.14 through 5.17.

From this point, it is relatively straightforward to estimate the pointing resolution. By taking the brightest location on the map for multiple calibration pulser events, the spread of the peak can be demonstrated as it is in Figures 5.18, 5.19 and 5.20. The Hilbert envelope resolution is the worst, as expected, as information is lost in the smoothing of the correlation waveforms that cannot be gained back by precise antenna locations. The full waveform correlation maps using only the classic ARA antennas have very high precision, nearly an order of magnitude higher than the previous ARA precision. And finally, adding in the Phased Array waveforms helps the precision even further.




Figure 5.14: Left: a correlation map for a high SNR local calibration pulse, using only the classic ARA antennas and the Hilbert envelope of the correlation waveforms. Right: the same, except for a low SNR local calibration pulse.




Figure 5.15: Left: a correlation map for a high SNR local calibration pulse, using the Phased Array and ARA antennas, generated with the Hilbert envelope of the correlation waveforms. Right: the same, except for a low SNR local calibration pulse.




Figure 5.16: Left: a correlation map for a high SNR local calibration pulse, using only the classic ARA antennas and the full correlation waveforms. Right: the same, except for a low SNR local calibration pulse.




Figure 5.17: Left: a correlation map for a high SNR local calibration pulse, using the Phased Array and ARA antennas, generated with the full correlation waveforms. Right: the same, except for a low SNR local calibration pulse.




Figure 5.18: Left: the reconstructed azimuth and zenith using only classic ARA baselines, and Hilbert enveloped-correlation waveforms. Right: the same, without applying a Hilbert envelope.




Figure 5.19: Left: the reconstructed azimuth and zenith using all available ARA baselines, and Hilbert enveloped-correlation waveforms. Right: the same, without applying a Hilbert envelope.

5.8 Lessons Learned

Because ARA Station 5 was unique to the previous ARA experiments, both in terms of the number of local calibration pulsars and the existence of the Phased Array, a completely new method had to be developed to locate the antennas. While this method has met the necessary requirements for an analysis, there are fairly simple ideas that can be included in future experimental designs that would make calibration easier. These are listed below.

More local calibration sources : having only one local calibration source means that the ice model has a considerable impact on the results. This means that the cable delays, the antenna locations, and the ice model must all be solved somewhat simultaneously. When possible, a second or, better, third local source would allow the antenna locations to be known even more precisely, with less calibration steps. This would allow distant sources like the SPIceCore hole or the IceCube deep pulsar to be used solely as




Figure 5.20: A plot comparing the resolution in both azimuth and zenith for various combinations of baselines.

validation.

Measure cable delays in realistic conditions : while the cable speeds were measured prior to deployment, the lengths were not measured precisely enough to be used in the fit without any additional correction, and the speed of the cable at cold temperatures was different than that measured in the lab. Spending a few days taking these measurements will remove nearly 25% of the unknowns of a given system, allowing the antenna locations to be the focus of the calibration.

Deploy calibration pulsers strategically : The choice of where to deploy the calibration pulser is sometimes up to chance, but when possible, deploying calibration antennas orthogonally from each other relative to the station will likely yield the best results. Otherwise, it can be tough to distinguish the lateral distances between the boreholes, as the time delays most directly measure the difference in distance between the source and the antennas.

Make signals across antennas as identical as possible : While it is possible to cross correlate two signals that do not look similar, it is far easier to do this when the signals are identical. Otherwise, the signals could be aligned incorrectly, and unless the individual antenna response as a function of elevation angle is known, this is not a possible effect to account for. The Phased Array has succeeded at this; the distributions of time delays from the Phased Array antennas are more than two times more narrow, mostly because they are impedance matched in a way that the ARA antennas are not.

Better HPol Data : So far in this chapter the discussion has been nearly completely focused on the VPol antennas, while the HPol antennas have had nearly no focus or investigation. This is partly because the Phased Array is a VPol instrument, but it is mostly because the available calibration data in HPol is extremely limited. For the purposes of this work, this has not been a dealbreaker, but future analyses of

ARA Station 5 may want to double check the assumptions about the location of the HPol antennas relative to the VPol antennas. However, this is not generally a priority, considering the ARA Station 5 HPols were disconnected in December 2019 after a USB failure caused the classic ARA Station DAQ to become unreachable.

Electric Channel Name	x [m]	y [m]	z [m]	Cable Delay [ns]
ARA 0	51.26	43.17	-194.74	141.58
ARA 1	51.24	43.16	-165.09	20.61
ARA 8	23.16	62.13	-196.20	138.64
ARA 9	23.16	62.12	-166.54	23.57
ARA 16	8.65	37.81	-177.74	140.11
ARA 17	8.64	37.81	-147.21	21.91
ARA 24	33.94	14.56	-190.86	130.50
ARA 25	33.95	14.57	-161.02	12.45
PA 0	21.61	46.46	-172.64	0.0
PA 1	21.61	46.46	-173.65	5.55
PA 2	21.61	46.46	-174.66	9.68
PA 3	21.61	46.46	-175.68	15.08
PA 4	21.61	46.46	-176.70	19.65
PA 5	33.94	14.57	-161.02	131.95
PA 6	21.61	46.46	-178.75	30.10
PA 7	21.61	46.46	-180.79	39.41

Table 5.3: The locations of each of the channels in the ARA Station 5 experiment. Note that channels PA 5 and ARA 24 refer to the same physical antenna. The Phased Array cable delays are reported relative to the Phased Array channel 0.

CHAPTER 6

THE PHASED ARRAY ANALYSIS

In this chapter, I will discuss the first analysis conducted with the Phased Array instrument installed as part of ARA Station 5. The analysis consists of the data that was taken during the calendar year 2019 with the Phased Array trigger and its antennas as the primary instrument. Each year, the Phased Array was operated in a slightly different configuration, motivating a year by year approach to the analysis. In particular, the 2019 dataset is the first to include the split ARA channel on the Phased Array side, which allows events between the systems to be easily matched together.

Within 2019, there are approximately 209 days of livetime in which both the Phased Array instrument and the ARA5 instruments were both online and taking data. This number is limited by the Phased Array overheating during the warmest months (usually December-March) and by a failure in a USB port on the ARA instrument that permanently impacted its ability to take data after November 8, 2019. Future Phased Array instruments that are built to consume less power are not expected to have this overheating issue.

The analysis is blinded, using the “pre-scaling” method described by Klein and Roodman [64], in which 10% of data is selected by randomly selecting one in every 10 events. Additionally, a full 24-hour sample was unblinded to better understand cyclical backgrounds such as weather balloons. With this dataset, a deep region is defined, analysis variables are calculated, and a background value is estimated, before unblinding the data and seeing the results. These steps will be described in the next sections.

6.1 Simulation

The simulation used to model this station is a modified version of AraSim, available on GitHub as the branch “PhasedArray”. Because the analysis focuses only on events triggered




Figure 6.1: The in-situ efficiency of the Phased Array instrument, compared to the efficiency of the classic ARA trigger [12].




Figure 6.2: The effective volume of the old Phased Array simulation compared to the new effective volume simulation.

by the Phased Array instrument, the simulation needed only to simulate this trigger. While a full beamforming trigger was outside of the scope of this project, a simple trigger was designed to match the in-situ performance of the phased array. This performance, shown in Figure 6.1, was measured by running the local calibration pulser at various power levels and at a steady rate, so the efficiency of the trigger could be calculated as a function of pulser strength.

In the simulation, the average SNR of the signal is calculated using only the signal seen by the seven VPol channels on the Phased Array string. Additionally, the incoming angle of the signal is calculated to determine whether the signal is on-beam or off-beam, For each event, a random number between 0 and 1 is generated, and if that number is less than the efficiency at that SNR and at that angle, the event passes. This is the same methodology used to write the Phased Array instrument paper [12] and was separately implemented here. The effective volume of this simulation is compared to that of the instrument paper in Figure 6.2. The slightly lower effective volume at lower energies could potentially be caused by the

different ice model used in this simulation, to match the ice model determined in Chapter 5. Overall, however, this effective volume looks generally consistent with what has been previously found. Future experiments with more thorough beamforming triggers can be used to better approximate the performance of this instrument.

For this analysis, the neutrinos are simulated using a flux based on a mixed composition of galactic cosmic rays and an optimistic model of source evolution [18]. A separate neutrino set simulated in half decade energy bins is used to calculate the efficiencies as a function of energy.

6.2 RZ Correlation Maps and Analysis Variables

In a previous chapter, the concept of Correlation Maps has already been introduced for 2D projections in azimuth and elevation. However, as mentioned previously, these maps are 2D projections of a three-dimensional phase space, where the neutrino interaction occurs at a specific (x, y, z) location in the ice. Additionally, the Phased Array itself is only one-dimensional, meaning it is only capable of measuring the incoming zenith angle of each signal at the antenna array. Because the index of refraction changes as a function of time, however, the incoming zenith angle can be extrapolated out to determine possible locations of the signal based on the local ice environment. In this way, data from the Phased Array alone can be used to generate R-Z Correlation maps, which show how well the event correlates for each pair of lateral distance (“R”) and depth (“Z”) coordinates. Each coordinate in the (R, Z) plane is described by the following equation:

$$C(R, Z) = \frac{\sum_{i=1}^{n_{ant}-1} \sum_{j=i+1}^{n_{ant}} \text{Corr}_{i,j}[\tau(R, Z)]}{\binom{n_{ant}}{2}}, \quad (6.1)$$

where R is the lateral distance from the Phased Array instrument, Z is the depth in the ice, n_{ant} is the number of antennas, $\text{Corr}_{i,j}$ is the correlation waveform between the i th and j th

antennas, and $\tau(R, Z)$ is the time delay expected between the two antennas at the location (R, Z) . All correlation maps are normalized so that the maximum possible value is 1.0 (in the case of perfect correlation) and the minimum possible value is -1.0 (in the case of perfect anti-correlation).

Because the index of refraction of ice is approximately exponential as a function of depth, it is possible to have multiple paths that travel through the ice and hit the detectors. Often, these paths are sorted into **direct paths**, which describe the shortest path between the signal creation point and the antennas, and **refracted or reflected paths**, which describe paths that bend above the detector and either bend back down, possible because of the exponential model, or hit the ice-air boundary and are reflected back down. Correlation maps can be generated using either type of map; for this analysis, three maps are used: direct maps, refracted maps, and combination maps that average the two results together. The combination maps are especially sensitive to double pulse events, in which both direct and refracted pulses are visible in the same event.

Example R-Z Correlation Maps are shown in Figure 6.3. For nearby sources like the local calibration pulser, the map is able to resolve distance well, because the signal is too close to behave like a plane wave. For distant sources, like the rooftop pulser mounted on the roof of the Ice Cube Lab building, the distance of the source is not possible to resolve. Instead, the correlation map diverges into two possible solutions at the boundary between the ice and the air. For a known calibration source like the rooftop pulser, the correct solution is known to be the surface, but this ambiguity in the correlation maps can make some neutrino signals difficult to separate from anthropogenic backgrounds occurring at the surface.

Also shown in Figure 6.3 is an example correlation map for a thermal noise event and a simulated neutrino event. In the case of the thermal noise event, the maximum value of the map is much lower than the maximum in the other example maps, with multiple potential source locations all approximately equally likely. The simulated neutrino event is




Figure 6.3: Example correlation maps for four types of events. Upper left: the local calibration pulser, with a blue dot for the actual location. Upper right: the rooftop calibration pulser, located approximately 4 km away at the surface. Lower left: a thermal noise trigger. Lower right: a simulated double pulse neutrino event, with the blue dot indicating the true location.

an example of a “double-pulse event”. When these events are plotted with a correlation map, both solutions are visible on the resulting map. The solutions overlap at the location where the event originated, marked on the map with a blue dot.

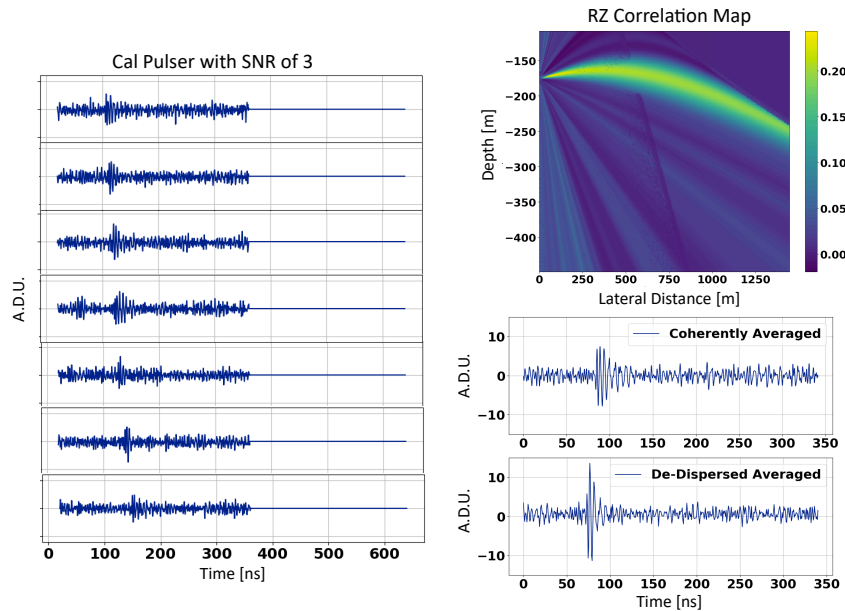


Figure 6.4: Left: An example low SNR calibration pulser. Top right: The correlation map for the low SNR calibration pulser. Bottom right: The coherently averaged waveform and the de-dispersed coherently averaged waveform (CSW).

Each of these categories of events have very distinct characteristics, many of which are clearly visible by just comparing quantities present in the correlation maps themselves: maximum correlation value, correlation near the surface, or the location of the peak. Other quantities require first creating a coherently-summed waveform (CSW), generated by removing the system response from each of the antennas, then adding the waveforms on all channels together using the time delays from the brightest spot on the correlation map. An example CSW for a low-SNR local calibration pulser, is shown in Figure 6.4. The CSW before and after removing the system response is shown for comparison.

The CSW is especially useful for calculating analysis variables related to impulsivity.

Impulsivity can be generally be characterized by the power distribution within a waveform. Impulsive events will have most of their power contained to a short time window, while non-impulsive events will have power evenly distributed throughout the entire waveform. To help describe impulsivity, I will first consider the Hilbert envelope of the CSW by calculating the Hilbert transform $H[x(t)]$ of a signal $x(t)$:

$$H[x(t)] = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{x(\tau)}{t - \tau} d\tau \quad (6.2)$$

and then calculating the analytic signal $z(t)$, defined as:

$$z(t) = x(t) + iH[x(t)]. \quad (6.3)$$

From this, the instantaneous amplitude of the analytic signal $a(t)$, also called the Hilbert envelope, can be calculated:

$$a(t) = \sqrt{x(t)^2 + H[x(t)]^2} \quad (6.4)$$

An example of a Hilbert envelope is shown in Figures 6.5 and 6.6. To describe impulsivity, I locate the peak of the Hilbert envelope and integrate outwards from that peak, creating a CDF that describes power distribution away from the peak. This CDF is referred to below as P_{CDF} . Waveforms that have most of their power concentrated to a small time window will have a highly non-linear CDF, while non-impulsive signals will have a fairly linear CDF, as shown in Figures 6.5 and 6.6. From this CDF, many impulsivity-related analysis variables are calculated, one of which is the average value of P_{CDF} , which naturally ranges from approximately 0.5 to 1.0 and is rescaled to span the range of approximately 0.0 to 1.0. Here, the lower range is approximate because for thermal events, it is possible to have an average that is lower than 0.5, if there is considerable power away from the Hilbert peak. This analysis variable is the same as that used in the ANITA III analysis [50].




Figure 6.5: Left: The CSW of a local calibration pulser in blue, with the Hilbert envelope in green. Right: The CDF of the Hilbert envelope starting from the peak. The average value is marked with a horizontal line.




Figure 6.6: Left: The CSW of a thermal noise event in blue, with the Hilbert envelope in green. Right: The CDF of the Hilbert envelope starting from the peak. The average value is marked with a horizontal line.

Variable Name	Variable Description	Variable Range
Maximum Correlation*	Maximum value on the R - z correlation map	[-1.0, 1.0]
Best R*	Location of the maximum correlation (lateral distance, m)	[10 m, 5500 m]
Best Z*	Location of the maximum correlation (depth, m)	[-1500 m, 0 m]
Hilbert Peak	Magnitude of the peak of the coherently summed waveform (CSW)	[0, 63]
SurfaceCor	Maximum correlation within 10 m of the surface	[-1.0, 1.0]
SurfaceZ	Location of SurfaceCor (depth, m)	[-10 m, 0 m]
SurfaceR	Location of SurfaceCor (lateral distance, m)	[10 m, 5500 m]
CloseSurface	Maximum value of correlation map above the surface	[-1.0, 1.0]
Zenith Angle	Best reconstructed zenith angle, calculated at the antenna array	[0, 180.0]
WindowedZenith	Best reconstructed zenith angle, windowed around the trigger	[0, 180.0]
CoherentSNR	SNR of the CSW	No Constraint
AvgSNR	Average of the SNR of each individual VPol waveform	No Constraint
Impulsivity*	Average of the cumulative distribution function (CDF) around the peak of the CSW	[0, 1.0]
R^{2*}	Correlation Coefficient for the linear fit to Impulsivity CDF	[0, 1.0]
Slope*	Slope of linear fit to impulsivity CDF	No Constraint
Intercept*	Intercept of linear fit to impulsivity CDF	No Constraint
PowerSpot*	Location of peak power along the CSW	[0,2560]
A5 Correlation*	Maximum correlation between) Phased Array CSW and A5 channel	[-1.0, 1.0]
K-S*	A Kolmogorov–Smirnov (K-S) test, with the impulsivity CDF against a linear hypothesis	No Constraint

Table 6.1: Table of analysis variables used in the Fisher Discriminant, along with their definitions and expected ranges where applicable.

The full list of analysis variables is described in Table 6.1. Variables marked with a * were calculated three times: once using the direct map, once using the refracted/reflected map, and once using the combination map.

6.3 Continuous Wave (CW) Background Events

One of the most common sources of radio background are continuous wave (CW) events, which are events that occur at a specific frequency intermittently depending on factors like the time of day, the time of the year, or the station activity at the South Pole. One example of this is the South Pole station communication, which operates on a radio frequency of 450 MHz and is so strong that the hardware is equipped with a notch filter so that this frequency does not dominate our triggers.




Figure 6.7: Before and after applying the Sine Subtraction method to a channel with 405 MHz CW.




Figure 6.8: All of the frequencies removed during the 24-hour unblinded period, centered around 405 MHz.

Another example of CW background is weather balloons, which are launched from the South Pole daily and transmit their data at a frequency of 405 MHz. These balloons are usually airborne for a time period of around 30 minutes to an hour, and only fly near the ARA stations some of the time. During the time periods where the balloon is near the ARA Stations, the CW causes the Phased Array to trigger more often, which in turn causes the beam thresholds on all beams to rise in response. These CW events are limited to a single

frequency, which is very unlike the expected broadband signal of a neutrino event, making these easily identifiable as non-neutrino signals. However, the CW content causes the events to have higher correlation values than a thermal noise event. Because of this, the CW events either need to be removed entirely, or the CW needs to be removed from the events.




Figure 6.9: The Maximum Correlation as a function of time before applying the Sine Subtract method (left) and after applying it (right).

In this analysis, I opt to remove the CW from the events using the Sine Subtraction method developed for the ANITA experiment [50]. In this method, a specific frequency range is identified to search for CW contamination, and a power threshold is set. In the case of the weather balloon source, the frequency range is chosen to be 400-410 MHz, and the power threshold is set to 4% of the total waveform's power. From here, the frequency in the given range with the most power is subtracted from the waveform by fitting the waveform to a sine wave at the given frequency and subtracting it. If the power of the resulting waveform has been reduced by at least the value of the power threshold, then the subtracted waveform is used. If not, the next highest frequency value is subtracted, and the process is repeated.

After repeating a maximum of three times, and no frequency successfully reduces power, the waveform is assumed to have no frequency contamination in the range specified. An example of how this method impacts a waveform in the frequency domain is shown in Figure 6.7.

To measure how successful this method is at removing CW events, the 24-hour burn sample can be used to learn about what happens when both CW-contaminated and non-CW-contaminated events are sent through the Sine Subtraction algorithm. In Figure 6.9 I compare the maximum correlation as a function of time for the raw data and for data that has been sent through the Sine Subtraction algorithm. In the raw data, there is a clear bump on the left hand side of the plot which corresponds to the weather balloon being visible. After processing, this bump has been removed completely, while other events are unaffected. I choose to apply this filter to all events, allowing any event contaminated with CW to stay in the analysis after the CW is removed.

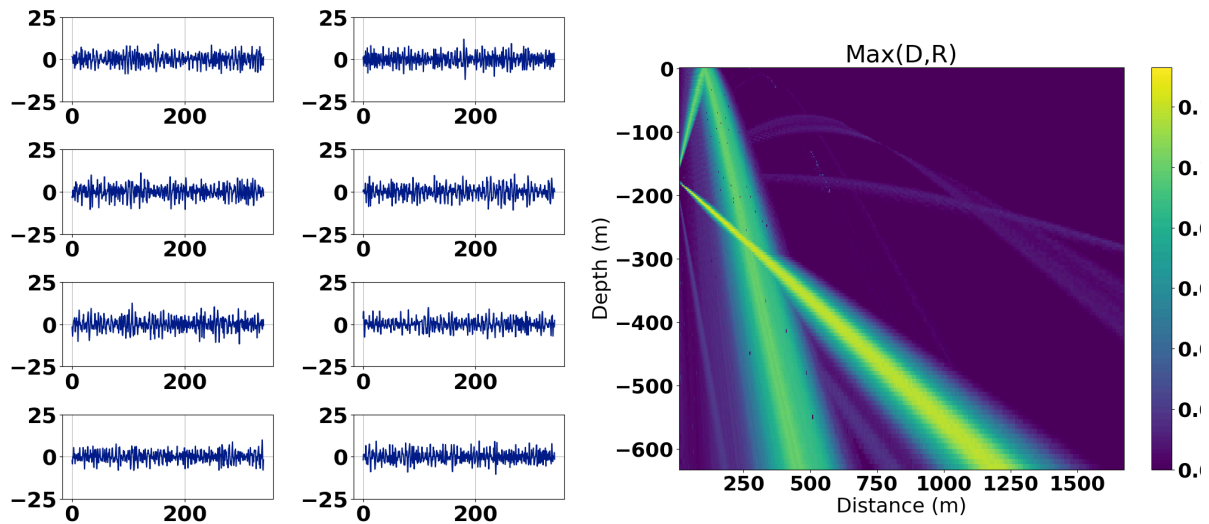


Figure 6.10: An example of an event with CW Contamination from satellite communications. Typically thermal noise events have a maximum correlation value below 0.10; here the maximum is 0.12.

Because this method is waveform-specific, not event-specific, this method is capable of detecting faint CW that impacts only a few channels. This is the case for satellite events, which were only categorized after investigating some of the events in the 10% burn sample with slightly higher maximum correlation. These events are slightly contaminated by a CW signal at 137 MHz, which is the frequency used by some satellites to communicate with the ground. An example satellite event is shown in Figure 6.10 (Run 3001, Event 88765). By adding a Sine Subtraction filter between 135-140 MHz, these satellite frequencies can be filtered from the dataset. This background source also explains a long-standing mystery in which trigger beam 4 of the Phased Array regularly spikes up. The correlation between the satellite event count against the beam 4 trigger threshold is clear, as shown in Figure 6.11. The only remaining mystery is which satellites are contributing to this beam threshold increase; a cursory search through an available catalogue did not point to a single satellite, potentially suggesting that multiple satellites are contributing.




Figure 6.11: A comparison of the events with CW contamination at 137 MHz (Top) and the beam 4 threshold pattern (Bottom).

In addition to the targeted CW filters targeting known CW frequencies, a general filter was applied between 200-700 MHz, for less common but more powerful CW contamination that may not have shown up in the 10% sample. This CW filter removed any event that had more than 10% of its power at a single CW frequency. After unblinding, no new CW sources were found, making this broad filter likely unnecessary for future analyses.

6.4 Livetime Cuts

6.4.1 Anthropogenic Events




Figure 6.12: A plot of outlier events in the sample during a time with a lot of known station activity (October 2019). Four runs account for nearly all of the outlier events.

Anthropogenic backgrounds can be caused by a myriad of sources, from airplanes to deployment activity at the ARA station. For backgrounds that are not CW-based, a more strict cut must be introduced to remove these sources. The most common characteristic of these anthropogenic backgrounds is that they cluster in time to other background events,

often with multiple events occurring within the same run. Considering the expected rate of both neutrinos and cosmic ray events are much lower, having multiple outlier events in a single run is a telltale sign that the run is contaminated with anthropogenic backgrounds.

Surprisingly, this effect is most visible in the data in October 2019, after the classic ARA5 station became unresponsive and multiple attempts to revive it failed. During this time period, plotting the maximum correlation value of the 10% sample over time shows multiple outliers clustered in time, as shown in Figure 6.12. This suggests that the act of logging in to the ARA5 DAQ computer is enough to increase the number of anthropogenic surface events, something that is not well understood at this time.




Figure 6.13: A plot of the surface correlation values for thermal noise events in the 10% sample. The red line indicates an exponential fit.

To remove runs with clusters of anthropogenic events, I consider the value of the analysis variable `SurfaceCor`, a measure of the maximum correlation value with 10 m of the surface. For anthropogenic events, I expect high correlation near the surface, and more than one high correlation event in the same run. This measure is quantified by plotting the values

of SurfaceCor for all events in the 10% sample and fitting the main distribution to an exponential, as shown in Figure 6.13. This exponential fit is then extrapolated to find the expected boundary for the full dataset. Any runs that have more than one event past this boundary in the 10% sample are excluded. This removes the events listed in Table 6.2, corresponding to approximately 18 hours of livetime.




Figure 6.14: A plot of the calibration pulser rate for each run. Most of the drops in event rate correspond to know mode switches or alternative calibration pulser configurations.

6.4.2 Calibration Pulsar Switches

The local calibration pulser at ARA Station 5 is by default sending VPol calibration pulses, although it is possible to switch to either HPol mode, or noise mode. In noise mode, the output is not impulsive or polarized, and the output occurs constantly instead of at a rate of 1 Hz. Over the course of year of operation, the calibration pulser is switched between these modes. The only way to switch to a new calibration mode is to cycle through all the available options while connected to the calibration pulser serial port. This means that a

run documented as an HPol run could have a handful of events that are untagged noise calibration pulser events. To avoid these events ending up in the signal region, I choose to remove runs with either documented calibration pulser switches, or unexplained calibration pulser rate drops. This is illustrated in Figure 6.14. The removed runs are included in Table 6.2.

Run Number	Event Number
3394	Noisy Run (23 events)
3455	Noisy Run (5 events)
3519	Noisy Run (6 events)
3001	Noisy Run (2 events)
3351	Noisy Run (2 events)
3407	Noisy Run (2 events)
1952, 2771, 2777, 2778, 2849, 2857, 2907, 2936, 2104, 3208, 3246, 3327, 3334	Cal Pulser Rate Drop

Table 6.2: The list of runs that were removed.

6.5 Defining a Deep Region

At this point, I have defined a set of analysis variables, identified the range of events to be included in the analysis, and removed periods with high anthropogenic backgrounds or changing calibration pulser modes. The remaining dataset is expected to include mainly thermal noise events (expected to account for $> 99\%$ of all events) and some number of individual, isolated anthropogenic events and cosmic ray events. These single events are the most challenging to remove, which like neutrino events are expected to be broadband, impulsive signals [65]. The best discriminator between these events and neutrino events is the incoming direction: the neutrino events mostly originate in deep ice, while the cosmic ray events and anthropogenic events point back towards the surface.

With this in mind, the remaining 10% sample is separated into two regions, the Deep Region and the Shallow Region. The boundaries between these regions are defined using




Figure 6.15: The population of simulated neutrinos (left) and the 10% sample (right, middle) plotted with the two variables defining the deep region: Zenith Angle and Surface Correlation Ratio. The color axis on the left and middle plots are event counts; the color axis on the right is the maximum correlation in each bin.

the two variables shown in Figure 6.15: Zenith angle reconstructed at the antenna array, and Surface Correlation Ratio, a ratio between the correlation within 10 m of the surface and the global maximum correlation. Surface events are expected to have low zenith angles and high values of Surface Correlation Ratio, because the global maximum correlation will be similar to that near the surface. Meanwhile, most neutrino events will have lower zenith angle values and lower values of Surface Correlation Ratio, with the exception of neutrino events that are refracted near the surface and thus appear to come from the surface when they originate in the deep ice. The exact boundaries were chosen to be 57° in zenith and 0.5 in Surface Correlation Ratio. Approximately 21% of simulated neutrinos, simulated using a flux based on a mixed composition of galactic cosmic rays and an optimistic model of source evolution [18], are removed from the signal region in this step.

In future works, the Surface Region could also be used as a signal region, but would require a more detailed cosmic ray simulation that did not exist at the time of writing this document. Generally, cosmic ray events and anthropogenic backgrounds can be differentiated from each other by their incoming azimuthal direction, a step that requires integrating the antennas from the classic ARA5 station.

6.6 Deep Region Analysis

After defining the signal region for the analysis, the following quality cuts are applied to the dataset:

Calibration Pulser Timing Cut In normal operating mode, the local calibration pulser emits a pulse at the beginning of the second at a rate of 1 Hz. To remove these events, a timing cut is applied to remove these events based on when the trigger occurs within the second. In the 10% sample, this timing cut removes 100% of the calibration pulser events.

Calibration Pulser Geometry Cut In previous ARA analyses, the calibration pulsers were found to misfire about once every 10,000 pulses. In the 10% sample, 0 calibration pulses were identified as misfired, allowing us to place a 90% confidence level upper limit of 21 misfired events in the full sample. To manage these misfired events, I apply a geometry cut to remove events reconstructing within a box defined as between 45 – 55 m radially from the antenna array and between 171 – 175.5 m deep. From this, I expect 99.8% of calibration pulser events to be removed. The background estimate from this is 0.009 events. After unblinding, I found all calibration pulser events were successfully removed with the timing cut, suggesting that future analyses with ARA Station 5 could reasonably remove this cut.

Software Trigger Cut In addition to RF triggers, the Phased Array instrument self triggers at a rate of 1 Hz, sampling the noise environment and allowing the nominal performance of each channel to be monitored over time. These events are expected to be random samples of noise, not neutrino events, and are removed from the analysis.

The remaining events in the 10% sample are shown in Figure 6.16, by comparing the SNR and the maximum correlation, two of the analysis variables. The remaining events in the




Figure 6.16: The remaining 10% sample (left) and the simulated neutrino set (middle, right) plotted with two of the analysis variables: SNR and maximum correlation.

signal region of the 10% sample are expected to have the characteristics of thermal noise and thus are considered background. Comparing the SNR and the maximum correlation of the background events to the simulated neutrinos in Figure 6.16, the differences between the two distributions are already apparent. However, ideally I would utilize more information from more than just two analysis variables, to capitalize on the differences between the simulated set and the 10%. This is the motivation for applying a Fisher Discriminant as the final analysis cut.

The Fisher Discriminant, also called a Linear Discriminant, is a type of basic machine learning algorithm that seeks to find the projection of maximum separation between classes of data [66]. Consider first the simplest case: separating two classes which are each described by two variables, illustrated in Figure 6.17. A single cut in one dimension along the axis of one of the variables would result in a high number of events sorted into the wrong class. Instead, the best cut would be a diagonal projection dependent on each of the two input

variables. The Fisher discriminant determines which projection maximizes the separation between the classes by finding the projection with the highest ratio of between-class variance to within-class variance.




Figure 6.17: An example of a dataset being split by a Fisher Discriminant using the Scikit-learn package. From [14].

Here, I implement the Fisher Discriminant from the Scikit-Learn package, with the two classes defined as the remaining 10% burn sample, assumed to be thermal noise, and the simulated neutrino set. The variables from Table 6.1 are the discriminant variables which are combined into a single discriminant coordinate that maximizing the separation between the classes. The resulting distributions of data and simulation are shown in Figure 6.18.

The final step is to determine the cut value in this coordinate space. The location of the cut will impact both the background estimate and the simulation efficiency in inverse ways. The farther you move the cut value to the right on Figure 6.18, the lower the expected background is, but the efficiency of the cut also decreases as less simulated neutrinos end up in the signal region. The opposite is also true: a cut moved towards the left would have

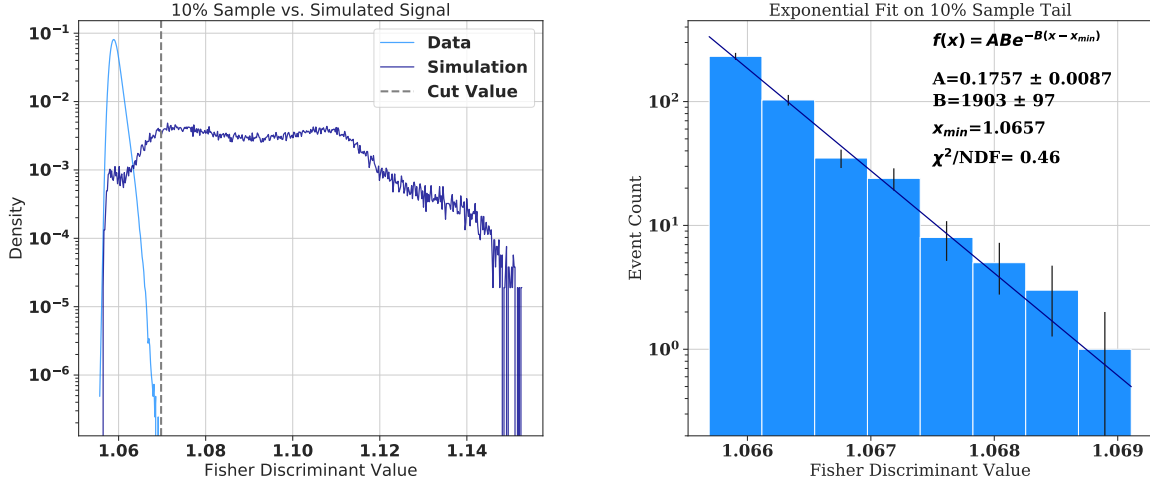


Figure 6.18: Left: A comparison between the 10% sample distribution and the simulated neutrino flux, with the optimized threshold marked with a vertical line. Right: the tail of the 10% sample, with the exponential fit overlaid. The value of X_{min} is defined as the minimum value of the Fisher Discriminant in the tail of the 10% sample.

a higher background and a higher efficiency. The solution to this optimization problem is to optimize for the best expected sensitivity. Specifically, I optimize for smallest value S , defined as:

$$S = \frac{n_i}{\epsilon} \quad (6.5)$$

where n_i is the average Feldman-Cousins 90% upper limit [67] after running 5000 pseudo-experiments with the expected background associated with the chosen cut value, and ϵ is the efficiency on simulated neutrinos.

To model the expected background, the tail of the distribution of the 10% sample is assumed to be exponential and fit to an equation of the form:

$$f(x) = AB e^{-B(x-x_{min})} \quad (6.6)$$

where x is the Fisher Coordinate and x_{min} is set to the minimum value included in the fit.

This parameterization was chosen to help the fit converge. The parameters were fit using ROOT and calculated to be $A = 0.1757 \pm 0.0087$ and $B = 1903 \pm 97$.

For each potential cut value, 100,000 pseudo-experiments are run, with the parameters of $f(x)$ selected from a Gaussian distribution using the reported error on each parameter as the standard deviation. Then, for the selected parameters, the background is estimated to be:




Figure 6.19: A plot of the estimated background distributions as a function of cut value. As the cut value decreases, the background distribution gets wider.

$$bg = \frac{t_{up}}{w} * \int_{cut}^{\infty} f(x)dx + bg_{other} \quad (6.7)$$

where $t_{up} = 9$ is the scale-up factor to account for the additional livetime in the full burn sample, w is the width of the binning used, and bg_{other} is the background from earlier cuts, namely the geometric calibration pulser cut. In this way, 100,000 pseudo-experiments create a distribution of bg values. The median background value for each distribution is selected as

Cut Name	Events Remaining	Background	Signal Efficiency
None	18,651,857	N/A	100%
Deep Region Boundary	6,005,122	N/A	79.03%
Cal Pulsar Gate Flag	4,423,436	N/A	99.98%
Cal Pulsar Geometry Cut	4,411,686	0.009	99.64%
Software Trigger Cut	4,014,776	N/A	100%
Fisher Discriminant	0	$0.09^{+0.06}_{-0.04}$	86.58%
Total	0	$0.10^{+0.06}_{-0.04}$	68.16%

Table 6.3: Table of cuts, background estimates, and analysis efficiencies from the 10% sample in the deep region. For the Fisher Discriminant, the median background is reported, as are the 16th and 84th percentiles.

input to determine the 90% upper limit. Figure 6.19 shows how the background distributions change as a function of cut value.

The cut value that given the minimum value of S is 1.0698, with an estimated background from the Fisher discriminant of $0.09^{+0.06}_{-0.04}$. The total estimated background including the geometric calibration pulsar cut is $0.10^{+0.06}_{-0.04}$.

Table 6.3 summarizes the cuts applied to the 10% sample, the background estimate, and the efficiency of the cut. The results will be discussed in the next section.

6.7 Deep Region Results

After unblinding the deep region, one event passed all cuts, shown in Figure 6.20. This event was triggered independently by both the ARA5 instrument and the Phased Array instrument, occurred in the middle of the Antarctic winter, and does not point back towards the South Pole Station. The incoming zenith angle of this event is approximately 60° , and the approximate maximum depth of the event is 30 m. This event is less impulsive than expected for a neutrino event, suggesting it is potentially a new type of background that was not modeled in the analysis. Looking at the correlation map in Figure 6.20, it seems like a potential source location could be the shadow zone near the surface, indicating that the

event could have been caused by a cosmic ray impacting core event that traveled through ice that was not well described by the exponential model used. Development of a cosmic ray simulation, as well as more ice studies of the shadow zone region, could help explain this type of event.

In future analyses, assuming this background event type is studied and understood, it will be possible to remove events like this with a relatively small impact on neutrino efficiency. For example, a stronger impulsivity cut could be required, or even a more stringent zenith cut. In Figure 6.21, I report the limit from this analysis, including the passing event in the




Figure 6.20: Top: the correlation map of the passing event. Bottom left: the Coherently Summed Waveform of the passing event, compared to a simulated event at the same incoming angle. The noise in the simulated waveform is normalized to have the same root-mean-square as the average from data. Bottom right: a table showing some of the values of analysis variables for the passing event.




Figure 6.21: The limit from this Phased Array analysis, using one station and six months of livetime, along with the projected sensitivity from the entire available livetime of 2.6 years. The expected sensitivity is calculated by directly scaling the livetime and expected backgrounds, selecting the median expected background, and calculating the 90% upper limit. Plotted for comparison are the previous published results from ARA (two stations, each with 4.0 station-years), along with results from ANITA [9], ARIANNA (seven stations, each with 4.5 station-years) [8], Auger [6] and IceCube [4] [5]. Additionally, theoretical production models of cosmogenic neutrinos are plotted [15] [16] [17].

signal region. In the projected limit, using all the data currently available from the Phased Array system, I assume a zenith cut of 90 degrees as a pessimistic lower bound of what is possible.




Figure 6.22: Analysis efficiency of this analysis compared to the most recent ARA Station 2 analysis efficiency [7]. The blue curve is the efficiency on neutrinos in the deep region; the purple curve includes the efficiency loss from removing the surface region which results in a 21% loss in efficiency averaged over all energies. The analysis efficiency shown in the plot on the left is generated for each energy bin separately. The analysis efficiency shown in the plot on the right is generated based on a cosmogenic flux [18]. The trigger efficiency vs. SNR is from [12], while the analysis efficiency from ARA Station 2 is adapted from [7].

The efficiency of this analysis is higher than previous ARA analyses, as shown in Figure 6.22. The Fisher discriminant efficiency, shown in blue, is above 80% at energies above 10^{17} eV, while the total efficiency, shown in purple, is higher at lower energies by nearly a factor of 10. Additionally, the efficiency vs. SNR plot in Figure 6.22 shows the Phased Array analysis is efficient in analysis on low SNR events that the classic ARA instrument could not have triggered on. This is a significant step forward in proving that the designing a system to trigger on low SNR events is a worthwhile endeavor, as the increased effective volume at low energies is something that can continue to be capitalized on in analysis.

6.8 Surface Region Results

For this analysis, the surface region was unblinded without considering it as a signal region. Because of this, the events in the surface region were run through the Fisher Discriminant

that had been trained on the deep events; no additional Fisher Discriminant was trained on the surface sample. If the surface region had been part of the signal region, a separate Fisher Discriminant would have needed to be trained.

In the surface region, 46 events passed all cuts. Seven of these events look like noise by eye; the rest are sorted into the following categories based only on inspection and not from a detailed analysis:

Airplane-like events Three events are clustered in time on January 11, 2019 and point nearly directly above the array, and are classified as potential airplane events. While access to flight arrival and departure times was denied by the NSF, rough flight logs show that four flights landed at the South Pole throughout January 11, making this hypothesis plausible.

Air shower-like events Eight events arrived between the zenith angles of $25 - 32^\circ$, which is too steep of an angle than expected for events that originate from anthropogenic sources a few kilometers away. The hypothesis for these events are air showers occurring directly above the array. These events have a mixture of VPol and HPol content; because these events were triggered by the Phased Array, which has a VPol-only trigger, it is possible that these events are not representative of the polarization generally expected from cosmic ray showers.

Total Internal Reflection events 14 events arrive at a zenith angle between $34 - 37^\circ$ of zenith, which corresponds to the angle in which a horizontally propagating surface signal would hit the antenna array. This likely means that this sample is a mixture of anthropogenic and cosmic ray events, although more study must be done before determining what fraction of each type of event is present.

Impacting Ice Core Events 13 events arrive beyond the total internal reflection angle, meaning that they likely do not come from the surface and instead are likely created

just below the surface of the ice, as a product of a cosmic ray hitting the ice surface and creating a particle shower. These events also have a mixture of VPol and HPol content and are generally impulsive and isolated in time. More simulation work should be done to fully understand these events.

6.9 Future Work

The results from this preliminary analysis are encouraging, suggesting that the design of the Phased Array instrument at the South Pole is worthwhile to include on future neutrino detection projects. The passing event can be used to better understand the backgrounds expected at the South Pole, and future simulation work could recover even more efficiency at the surface. The analysis techniques described in this chapter can be adapted for other experiments and would likely yield similar results.

Already, multiple experiments are being built that implement a phased array trigger technique, namely RNO-G and PUEO, and this design decision is partly due to the success of the phased array instrument on ARA Station 5. Other radio detection experiments, like the planned IceCube-Gen2 radio array, are finalizing their designs soon, and it seems likely that they too could incorporate this trigger type into their station designs.

It seems very likely that soon (maybe in the next decade or so) we could have evidence of multiple UHE neutrino signals detected via the radio technique. When this happens, it will be because of decades of detector development, rigorous analysis efforts, and collaboration between scientists all over the world. I am already looking forward to the coming era of neutrino astronomy.

CHAPTER 7

THE BEAMFORMING ELEVATED ARRAY FOR COSMIC NEUTRINOS (BEACON)

This chapter was adapted from a proceeding written for ICRC 2019 [68].


7.1 Motivation

The Beamforming Elevated Array for Cosmic Neutrinos (BEACON) experiment, like ARA, is built to detect astrophysical neutrinos above 10 PeV, which have yet to be detected by any experiment. BEACON is specifically sensitive to radiation from both cosmic ray air showers and upward going tau lepton-induced showers caused by tau neutrino interactions in the Earth. The main discriminator between the cosmic ray showers and the neutrino signals is the incoming direction: cosmic rays are expected to come from above the horizon, while neutrinos will come from at or just below the horizon.

BEACON is the first experiment to utilize a beamforming trigger for the purpose of detecting tau-lepton induced showers, although the method is planned to be used in PUEO, which will have a tau-specific detection channel. The hypothesis is that the trigger will make radio detection in the presence of anthropogenic signals more possible, and the high duty cycle of BEACON compared to other optical air shower detectors is another benefit. The geometry of a BEACON station is an additional benefit, increasing the field of view compared to experiments like [69].


While the successes of the interferometric trigger has been discussed in the earlier chapters, the early results from the BEACON experiment are promising. In this chapter, the BEACON concept will be discussed and the early analysis work from the first year of the prototype will also be described. Since this first year, multiple rounds of development on the BEACON prototype have been implemented, and the current prototype has been suc-

cessfully running for nearly a year. The work described below does not include any of the most recent developments, as it is outside the scope of my work as a student.



((a)) Spectra from Vertical Polarization. The summit shows a single spectral run, while the observatory shows the averages of 50 spectra.

((b)) Spectra from Horizontal Polarization, averaged over 50 spectra.



((c)) Rates from Vertical Polarization.

((d)) Rates from Horizontal Polarization.

Figure 7.1: Spectra and Rates from White Mountain Site Study.

7.2 White Mountain Site Study

The Barcroft Research Station [70] is located on White Mountain in California, overlooking both Owens Valley to the West and Fish Lake Valley to the East at an elevation of 3.8 km. The prominence over Owens Valley and Fish Lake Valley are 2.6 km and 2.4 km, respectively. This prominence is the major factor in choosing Barcroft for the BEACON prototype. Additionally, Barcroft Station has cabled power which can be used to power the

BEACON prototype, something that other prominent mountains do not have. Development is ongoing for a future BEACON system that can be autonomously powered by wind or solar.

To determine the feasibility of using the Barcroft site for the prototype, the local RFI environment was studied. Electrically-short bicone antennas (ETS-Lindgren 3180 C), sensitive between 30-1000 MHz, were used to measure the local radio environment in two frequency ranges: 30-90 MHz and 200-1200 MHz. In the low frequency range, signals were amplified with a total gain of 79.5 dB in two stages, with a tunable tank filter applied to cut out regular CW pulse packets at 42.5 MHz. In the high frequency range, the amplification ranged linearly from 79.5 dB at 200 MHz down to 64 dB at 1200 MHz. While both frequency ranges were investigated, this study confirmed a previous result that the low frequency range had less contamination from anthropogenic sources. Only the low frequency range is reported from this point forward.

A spectrum analyzer was used to calculate the average continuous wave (CW) spectra present in each direction, in each frequency range. The results of this study are shown in Figure 7.1. Comparing the results in the East vs. the results from the West, the East had less power than the West by more than 10 dB at some frequencies. While the power in the horizontal and vertical polarizations were on average about the same, spikes of high-powered CW were present in the horizontal polarization intermittently between 50-60 MHz in the East and 60-65 MHz in the West. A very small amount of the band above 100 MHz was

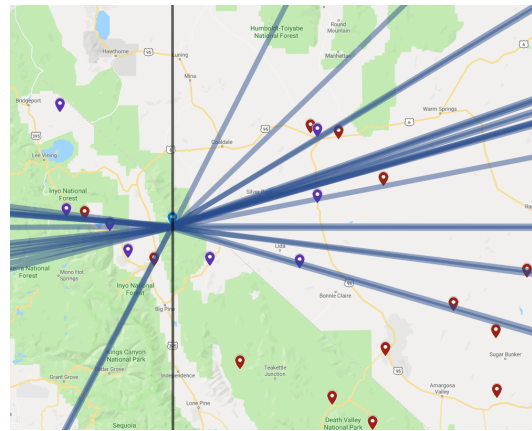


Figure 7.2: A map of the source directions as found by the single baseline interferometer. Lines to the left of center represent data taken while the interferometer was placed overlooking the West Valley, and lines to the right are from the East.

consistent with the expectations of thermal noise, making it this part of the band unusable for an experiment like BEACON [71]. Additional tests were run at the actual summit of White Mountain, which had even higher RFI than either of the Barcroft sites.

As an additional test, a pulse counter was used to measure the trigger rates as a function of threshold, also shown in 7.1. The results are compared to a terminated input, used as a proxy for thermal noise. The results from both the West and the East valley both show a significant amount of background not consistent with thermal noise, suggesting that the transient backgrounds would need to be reduced by at least a factor of 1000.

Although neither of the valleys were clear of RFI, the East Valley was less noisy, motivating the choice for the BEACON prototype to overlook Fish Lake Valley. However, significant vetoing abilities must be a feature of any trigger overlooking either valley in order to reject RFI. The first attempt at such a system is described later in this chapter.

7.3 Interferometry Studies

Two different interferometry studies were conducted on White Mountain: first, a two-day long, single baseline interferometer; and second, a year-long, four-antenna interferometer with a phased array trigger, referred to here as the BEACON prototype. This section describes each in detail.

7.3.1 A Two-Antenna Short Term Interferometer

The bicones used for the site study were reused here to create a short term, single baseline interferometer, with filters selecting the 30-80 MHz range. The interferometer was placed two separate times, once overlooking the West valley and once overlooking the East. Each antenna was at approximately the same height, causing the time difference between each antenna to be directly related to the azimuthal angle at which the signal arrived. Because there are only two antennas, the source could equally be reconstructed on either side of the

antenna pair. However, here the source direction was chosen as the solution that pointed back out towards the valley at the time. This choice was due to the placement of the bicones down the ridge of the mountain, obscuring the view from the opposite valley.

Waveforms were recorded in varying time windows, with the longest being 20 ms, sampled at 0.4 ns. These were then broken into smaller time windows and analyzed for pulses using cross correlation. Pulses that were identified could then be used to calculate the arrival angle. This was done for all pulses for both Owens Valley and Fish Lake Valley.

In Figure 7.2, some source arrival directions do correlate with known locations of local towns and airports. Additionally, expected background rates were extrapolated from this data set; for a thermal event rate of 1-100 Hz, the expected background rate is 250 Hz. This presents a clear challenge to design a experimental system capable of rejecting a high level of RFI contamination.



Figure 7.3: View of BEACON prototype from above the mountain ridge. Antennas were camouflaged to minimize environmental impact.




Figure 7.4: Approximate beam directions of the BEACON prototype. Here due East corresponds to 0 azimuth.

7.3.2 *The BEACON Prototype: A Four-Antenna, Year-Long Interferometer*

Following the White Mountain Site Study, a four-antenna station was set up facing the Fish Lake Valley. The antennas chosen for the BEACON prototype were inverted-V cross dipole

antennas, as shown in Figure 7.3, also used as part of the Long Wavelength Array (LWA) experiment at the Owens Valley Radio Observatory [72]. This antenna was chosen for its sensitivity to 30-80 MHz frequencies, as well as its active balun that includes conversion to a coaxial cable line and 35 dB of amplification [73].

Each antenna was installed on the slope of White Mountain along a 20° gradient. However, the beam pattern greatly favors signals coming from directly above the antenna; though installing the antennas on a slope of 20° improves the response at the horizon slightly, there is still a loss of 5-7 dB. Current development work is being done to improve the antenna design for future iterations.

Additionally, the BEACON prototype was equipped with a phased array trigger system, using the same underlying framework as the interferometric trigger designed and deployed as part of the ARA5 station [12]. The BEACON prototype system has a total of 20 beams, with approximate directions shown in Figure 7.4.

The trigger threshold on each beam changes with time as the RF environment changes, as is evident in Figure 7.6. A few known sources, such as a strong 48 MHz CW signal caused by local towers, are known to turn on and off every 30-60 minutes. Evidence of this and other repeating sources can be seen in Figure 7.5. In an effort to cut down on the high volume of local RFI, a number of RF vetos were added to the software at the trigger level. These vetos enabled the BEACON prototype to ignore events that were overly saturated, had significant CW content, or had high voltage on some antennas but not others. Notably, after these RF vetos were added to the trigger, the thresholds on most beams lowered to between 5-20 σ . The two outlier




Figure 7.5: A spectrogram showing the changing frequency content over the course of 100 minutes.

beams are the ones most directly facing into Fish Lake Valley.

7.4 Understanding Local RFI

7.4.1 *High-Voltage Power Lines*

One of the most prevalent types of triggered events are repeating, poorly correlating events that occur in intervals of approximately 60 Hz spacing, as seen in Figure 7.7. This 60 Hz structure in the trigger time is directly related to the frequency that is carried on power lines in the United States. This hypothesis was confirmed in the field using a log-periodic dipole antenna (LPDA) and an oscilloscope at a power substation along the high-voltage lines in the Fish Lake Valley.

Because at least 10 percent of triggered events appear to be caused by this 60 Hz background, it is imperative that future iterations of BEACON are able to veto this type of event. Likely this can be done on the trigger level; this is an area of active development.

7.4.2 *Galactic Noise*

As the BEACON prototype is most sensitive between 30-80 MHz, it is expected that the radio emission from the galactic center should be visible. One way to find the galactic center would be to construct the average correlation map of many forced trigger events in celestial coordinates; however, this method requires the antenna positions to be well calibrated and farther apart for more precise correlation maps.

An easier solution is to calculate the median RMS of all forced trigger events over the course of many day. If the galactic center is visible, the RMS should increase as the galactic center rises, and fall as it sets. From Figure 7.8, the median RMS is correlated with the rising and setting of the galactic center, which is evidence that it is visible.

7.4.3 Airplane Tracks

The BEACON prototype is also sensitive to RFI from airplanes as they move into and out of its field of view. While airplanes themselves broadcast at frequencies higher than 80 MHz, radio emission from the ground is reflected by the body of the airplanes, making them possible to detect.

Candidate airplane events were found by clustering causal events, i.e. events that returned physically possible time delays based on the physical spacing of the baselines, by their time delays across each baseline of antennas. Events were considered clustered if

all six time delays were within 2 ns. Events were removed if from one of the top 25 largest clusters, each containing more than 1000 events in the analyzed sample.

The remaining events were then sorted based on “impulsivity”, an observable borrowed from ANITA [50] and also discussed as part of the phased array analysis in this thesis. Highly impulsive events contain a high fraction of power around the peak of the coherently summed waveform. A set of high-impulsivity events were found with slowly changing time delays across all baselines over the course of around 30 seconds. Four planes were identified using this method over the course of two days of data. An example set of airplane events is shown in Figure 7.9, providing clear evidence that the BEACON prototype can trigger on impulsive transient events.

Airplane events are easy to spot when multiple events are recorded and can be tracked; it is much more challenging to select airplane events in which only one or two events are




Figure 7.9: Time delays for each baseline as a function of readout time, using vertically-polarized channels. The sloped nature of the lines is evidence that the source is moving over the course of 25 seconds.

detected. To assist with detecting all plane events, a flight tracker was installed in 2019 that detects and records the time, latitude, longitude, and altitude of each airplane, which is broadcast from the airplane every second at 1090 MHz [74]. This setup has made it possible to detect and verify many airplane events.

7.5 Work Since 2018

Since the first year of the BEACON prototype, an upgraded system with custom-built electrically short dipole antennas was installed in 2019, with additional modifications in 2020 and 2021. The upgraded antennas have higher gain at the horizon, and are farther from the ground, making them more sensitive. Additionally, the location of the antennas was calibrated using differential GPS modules, with additional pulsing confirming the results of the GPS measurements. With this upgrading pointing ability, airplane events can be identified not just as airplane events, but by which specific airplane was causing the RFI. This has made it easier to remove airplane events in analysis.

Currently, a cosmic ray search is underway using the data taken during 2021. Future iterations of the BEACON project may include testing the autonomous DAQ, conducting site surveys of other prominent mountain sites, and tuning the phased array trigger concept so it is more efficient in highly RFI-dominated areas.




Figure 7.6: A comparison of the thresholds of each beam formed with the BEACON prototype system. One σ on this plot corresponds to the average power over 16 samples of a typical noise trigger.

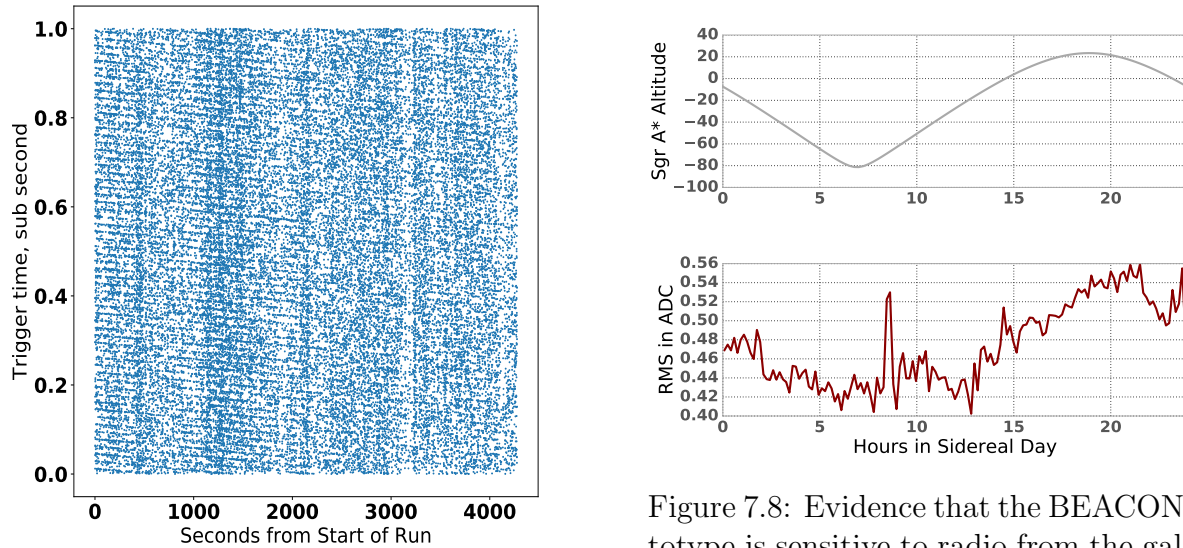


Figure 7.7: Example of the 60 Hz structure that appears during data taking.

Figure 7.8: Evidence that the BEACON prototype is sensitive to radio from the galactic center. Here the altitude of Sgr A* is compared to the median RMS of noise trigger events after undergoing a low pass filter.

REFERENCES

- [1] Imre Bartos and Marek Kowalski. Multimessenger astronomy. In *Multimessenger Astronomy*, 2399-2891, pages 1–1 to 1–18. IOP Publishing, 2017.
- [2] J. A. Aguilar et al. Design and sensitivity of the radio neutrino observatory in greenland (rno-g). *Journal of Instrumentation*, 16(03):P03025, Mar 2021.
- [3] Mauricio Bustamante and Amy Connolly. Extracting the energy-dependent neutrino-nucleon cross section above 10 TeV using IceCube showers. *Physical Review Letters*, 122(4), jan 2019.
- [4] M. G. Aartsen et al. Differential limit on the extremely-high-energy cosmic neutrino flux in the presence of astrophysical background from nine years of icecube data. *Phys. Rev. D*, 98:062003, Sep 2018.
- [5] M. G. Aartsen et al. OBSERVATION AND CHARACTERIZATION OF a COSMIC MUON NEUTRINO FLUX FROM THE NORTHERN HEMISPHERE USING SIX YEARS OF ICECUBE DATA. *The Astrophysical Journal*, 833(1):3, dec 2016.
- [6] A. Aab et al. Probing the origin of ultra-high-energy cosmic rays with neutrinos in the EeV energy range using the pierre auger observatory. *Journal of Cosmology and Astroparticle Physics*, 2019(10):022–022, oct 2019.
- [7] P. Allison et al. Constraints on the diffuse flux of ultrahigh energy neutrinos from four years of askaryan radio array data in two stations. *Physical Review D*, 102(4), Aug 2020.
- [8] A. Anker et al. Targeting ultra-high energy neutrinos with the arianna experiment. *Advances in Space Research*, 64(12):2595–2609, 2019.
- [9] P. W. Gorham et al. Constraints on the ultrahigh-energy cosmic neutrino flux from the fourth flight of anita. *Phys. Rev.*, D99(12):122001, 2019.
- [10] Ara Collaboration, P. Allison, J. Auffenberg, R. Bard, J. J. Beatty, D. Z. Besson, S. Böser, C. Chen, P. Chen, A. Connolly, J. Davies, M. Duvernois, B. Fox, P. W. Gorham, E. W. Grashorn, K. Hanson, J. Haugen, K. Helbing, B. Hill, K. D. Hoffman, E. Hong, M. Huang, M. H. A. Huang, A. Ishihara, A. Karle, D. Kennedy, H. Landsman, T. C. Liu, L. Macchiarulo, K. Mase, T. Meures, R. Meyhandan, C. Miki, R. Morse, M. Newcomb, R. J. Nichol, K. Ratzlaff, M. Richman, L. Ritter, C. Rott, B. Rotter, P. Sandstrom, D. Seckel, J. Touart, G. S. Varner, M. Z. Wang, C. Weaver, A. Wendorff, S. Yoshida, and R. Young. Design and initial performance of the askaryan radio array prototype eev neutrino detector at the south pole. *Astroparticle Physics*, 35(7):457–477, February 2012.
- [11] Brian A Clark. *Optimization of a Search for Ultra-High Energy Neutrinos in Four Years of Data of ARA Station 2*. PhD thesis, 2019.

- [12] P. Allison et al. Design and performance of an interferometric trigger array for radio detection of high-energy neutrinos. *Nuclear Instruments and Methods in Physics Research A*, 930:112–125, Jun 2019.
- [13] L L Ruckman and G S Varner. The pro1 asic for fast wilkinson encoding. *Journal of Instrumentation*, 3(12):P12003–P12003, Dec 2008.
- [14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [15] A. V. Olinto, K. Kotera, and D. Allard. Ultrahigh energy cosmic rays and neutrinos. *Nucl. Phys. Proc. Suppl.*, 217:231–236, 2011.
- [16] Kumiko Kotera and Angela V. Olinto. The astrophysics of ultrahigh energy cosmic rays. *Ann. Rev. Astron. Astrophys.*, 49:119–153, 2011.
- [17] Markus Ahlers and Francis Halzen. Minimal cosmogenic neutrinos. *Phys. Rev. D.*, 86:083010, 2012.
- [18] K. Kotera, D. Allard, and A. V. Olinto. Cosmogenic neutrinos: parameter space and detectability from pev to zev. *JCAP*, 1010(10):013, 2010.
- [19] Bruce T. Cleveland, Timothy Daily, Jr. Raymond Davis, James R. Distel, Kenneth Lande, C. K. Lee, Paul S. Wildenhain, and Jack Ullman. Measurement of the solar electron neutrino flux with the homestake chlorine detector. *The Astrophysical Journal*, 496(1):505–526, mar 1998.
- [20] K. Hirata, T. Kajita, M. Koshiba, M. Nakahata, Y. Oyama, N. Sato, A. Suzuki, M. Takita, Y. Totsuka, T. Kifune, T. Suda, K. Takahashi, T. Tanimori, K. Miyano, M. Yamada, E. W. Beier, L. R. Feldscher, S. B. Kim, A. K. Mann, F. M. Newcomer, R. Van, W. Zhang, and B. G. Cortez. Observation of a neutrino burst from the supernova sn1987a. *Phys. Rev. Lett.*, 58:1490–1493, Apr 1987.
- [21] R. M. Bionta, G. Blewitt, C. B. Bratton, D. Casper, A. Ciocio, R. Claus, B. Cortez, M. Crouch, S. T. Dye, S. Errede, G. W. Foster, W. Gajewski, K. S. Ganezer, M. Goldhaber, T. J. Haines, T. W. Jones, D. Kielczewska, W. R. Kropp, J. G. Learned, J. M. LoSecco, J. Matthews, R. Miller, M. S. Mudan, H. S. Park, L. R. Price, F. Reines, J. Schultz, S. Seidel, E. Shumard, D. Sinclair, H. W. Sobel, J. L. Stone, L. R. Sulak, R. Svoboda, G. Thornton, J. C. van der Velde, and C. Wuest. Observation of a neutrino burst in coincidence with supernova 1987a in the large magellanic cloud. *Phys. Rev. Lett.*, 58:1494–1496, Apr 1987.
- [22] M. G. Aartsen et al. Neutrino emission from the direction of the blazar txs 0506+056 prior to the icecube-170922a alert. *Science*, 361(6398):147–151, 2018.

- [23] M. G. Aartsen et al. Multimessenger observations of a flaring blazar coincident with high-energy neutrino icecube-170922a. *Science*, 361(6398):eaat1378, 2018.
- [24] Yiqian Xu. Tev-pev neutrino-nucleon cross section measurement with 5 years of icecube data. *PoS, DIS2018:019*, 2018.
- [25] Philipp Baerwald, Mauricio Bustamante, and Walter Winter. Neutrino decays over cosmological distances and the implications for neutrino telescopes. *Journal of Cosmology and Astroparticle Physics*, 2012(10):020–020, oct 2012.
- [26] Mauricio Bustamante, John F. Beacom, and Kohta Murase. Testing decay of astrophysical neutrinos with incomplete information. *Physical Review D*, 95(6), mar 2017.
- [27] Carlos. A. Argüelles, Teppei Katori, and Jordi Salvado. Effect of new physics in astrophysical neutrino flavor. *Physical Review Letters*, 115(16), oct 2015.
- [28] Markus Ahlers, Mauricio Bustamante, and Siqiao Mu. Unitarity bounds of astrophysical neutrinos. *Physical Review D*, 98(12), dec 2018.
- [29] L. Oc. Drury. Review article: An introduction to the theory of diffusive shock acceleration of energetic particles in tenuous plasmas. *Reports on Progress in Physics*, 46(8):973–1027, August 1983.
- [30] ENRICO Fermi. On the origin of the cosmic radiation. *Phys. Rev.*, 75:1169–1174, Apr 1949.
- [31] Ke Fang and Kohta Murase. Linking high-energy cosmic particles by black-hole jets embedded in large-scale structures. *Nature Physics*, 14(4):396–398, jan 2018.
- [32] Xavier Rodrigues, Jonas Heinze, Andrea Palladino, Arjen van Vliet, and Walter Winter. Active galactic nuclei jets as the origin of ultrahigh-energy cosmic rays and perspectives for the detection of astrophysical source neutrinos at EeV energies. *Physical Review Letters*, 126(19), may 2021.
- [33] Kohta Murase, Yoshiyuki Inoue, and Charles D. Dermer. Diffuse neutrino intensity from the inner jets of active galactic nuclei: Impacts of external photon fields and the blazar sequence. *Physical Review D*, 90(2), jul 2014.
- [34] Denise Boncioli, Daniel Biehl, and Walter Winter. On the common origin of cosmic rays across the ankle and diffuse neutrinos at the highest energies from low-luminosity gamma-ray bursts. *The Astrophysical Journal*, 872(1):110, feb 2019.
- [35] Kohta Murase. High energy neutrino early afterglows from gamma-ray bursts revisited. *Physical Review D*, 76(12), dec 2007.
- [36] Ke Fang, Kumiko Kotera, Kohta Murase, and Angela V. Olinto. Testing the newborn pulsar origin of ultrahigh energy cosmic rays with eev neutrinos. *Phys. Rev. D*, 90:103005, Nov 2014.

- [37] K. Greisen. End to the cosmic-ray spectrum? *Physical Review Letters*, 16:748–750, April 1966.
- [38] G. T. Zatsepin and V. A. Kuzmin. Upper limit of the spectrum of cosmic rays. *JETP Lett.*, 4:78–80, 1966. [Pisma Zh. Eksp. Teor. Fiz.4,114(1966)].
- [39] 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)].
- [40] P. W. Gorham et al. Observations of the askaryan effect in ice. *Phys. Rev. Lett.*, 99:171101, 2007.
- [41] David Saltzberg et al. Observation of the askaryan effect: Coherent microwave cherenkov emission from charge asymmetry in high-energy particle cascades. *Physical Review Letters*, 86(13):2802–2805, Mar 2001.
- [42] P. W. Gorham et al. Accelerator measurements of the askaryan effect in rock salt: A roadmap toward teraton underground neutrino detectors. *Phys. Rev. D*, 72:023002, Jul 2005.
- [43] K. Belov et al. Accelerator measurements of magnetically-induced radio emission from particle cascades with applications to cosmic-ray air showers. *Phys. Rev. Lett.*, 116(14):141103, 2016.
- [44] A. Aab et al. Probing the radio emission from air showers with polarization measurements. *Phys. Rev. D*, 89:052002, Mar 2014.
- [45] E.C. Andrés, P. Askebjerg, S.W. Barwick, R.C. Bay, L. Bergström, A. Biron, J. Booth, O. Botner, A. Bouchta, S. Carius, M. Carlson, W. Chinowsky, D. Chirkin, J. Conrad, C.G.S. Costa, D. Cowen, E. Dalberg, T. DeYoung, J. Edsjö, P. Ekström, A. Goobar, L. Gray, A. Hallgren, F. Halzen, R. Hardtke, S. Hart, Y. He, C.P. de los Heros, G. Hill, P.O. Hulth, S. Hundertmark, J. Jacobsen, A. Jones, V. Kandhadai, A. Karle, J. Kim, H. Leich, M. Leuthold, P. Lindahl, I. Liubarsky, P. Loaiza, D. Lowder, P. Marciniwski, T.C. Miller, P. Miocinovic, P.C. Mock, R. Morse, M. Newcomer, P. Niessen, D. Nygren, R. Porrata, D. Potter, P.B. Price, G. Przybylski, W. Rhode, S. Richter, J. Rodriguez, P. Romanesko, D. Ross, H. Rubinstein, T. Schmidt, E. Schneider, R. Schwarz, U. Schwendicke, G. Smoot, M. Solarz, V. Sorin, C. Spiering, P. Steffen, R. Stokstad, O. Streicher, I. Taboada, T. Thon, S. Tilav, C. Walck, C.H. Wiebusch, R. Wischnewski, K. Woschnagg, W. Wu, G. Yodh, and S. Young. The AMANDA neutrino telescope. *Nuclear Physics B - Proceedings Supplements*, 77(1-3):474–485, may 1999.
- [46] Matteo Sanguineti. ANTARES and KM3NeT: The Latest Results of the Neutrino Telescopes in the Mediterranean. *Universe*, 5(2):65, 2019.
- [47] A. V. Olinto, J. Krizmanic, J. H. Adams, R. Aloisio, L. A. Anchordoqui, A. Anzalone, M. Bagheri, D. Barghini, M. Battisti, D. R. Bergman, M. E. Bertaina,

- P. F. Bertone, F. Bisconti, M. Bustamante, F. Cafagna, R. Caruso, M. Casolino, K. Černý, M. J. Christl, A. L. Cummings, I. De Mitri, R. Diesing, R. Engel, J. Eser, K. Fang, F. Fenu, G. Filippatos, E. Gazda, C. Guepin, A. Haungs, E. A. Hays, E. G. Judd, P. Klimov, V. Kungel, E. Kuznetsov, Š. Mackovjak, D. Mandát, L. Marcelli, J. McEnery, G. Medina-Tanco, K. D. Merenda, S. S. Meyer, J. W. Mitchell, H. Miyamoto, J. M. Nachtman, A. Neronov, F. Oikonomou, Y. Onel, G. Osteria, A. N. Otte, E. Parizot, T. Paul, M. Pech, J. S. Perkins, P. Picozza, L. W. Piotrowski, Z. Plebaniak, G. Prévôt, P. Reardon, M. H. Reno, M. Ricci, O. Romero Matamala, F. Sarazin, P. Schovánek, V. Scotti, K. Shinozaki, J. F. Soriano, F. Stecker, Y. Takizawa, R. Ulrich, M. Unger, T. M. Venters, L. Wiencke, D. Winn, R. M. Young, and M. Zotov. The poemma (probe of extreme multi-messenger astrophysics) observatory. 2020.
- [48] P. W. Gorham et. al. The antarctic impulsive transient antenna ultra-high energy neutrino detector design, performance, and sensitivity for 2006-2007 balloon flight. *Astroparticle Physics*, 32, 2009.
- [49] 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. D.*, 82, 2010.
- [50] P. W. Gorham et al. Constraints on the diffuse high-energy neutrino flux from the third flight of anita. *Phys. Rev. D*, 98:022001, Jul 2018.
- [51] P. W. Gorham et al. Constraints on the diffuse high-energy neutrino flux from the third flight of anita. *Phys. Rev. D.*, 98, 2018.
- [52] Q. Abarr et al. The payload for ultrahigh energy observations (pueo): a white paper. *Journal of Instrumentation*, 16(08):P08035, Aug 2021.
- [53] R. U. Abbasi, M. Abe, T. Abu-Zayyad, M. Allen, R. Anderson, R. Azuma, E. Barcikowski, J. W. Belz, D. R. Bergman, S. A. Blake, R. Cady, M. J. Chae, B. G. Cheon, J. Chiba, M. Chikawa, W. R. Cho, T. Fujii, M. Fukushima, T. Goto, W. Hanlon, Y. Hayashi, N. Hayashida, K. Hibino, K. Honda, D. Ikeda, N. Inoue, T. Ishii, R. Ishimori, H. Ito, D. Ivanov, C. C. H. Jui, K. Kadota, F. Kakimoto, O. Kalashev, K. Kasahara, H. Kawai, S. Kawakami, S. Kawana, K. Kawata, E. Kido, H. B. Kim, J. H. Kim, J. H. Kim, S. Kitamura, Y. Kitamura, V. Kuzmin, Y. J. Kwon, J. Lan, S. I. Lim, J. P. Lundquist, K. Machida, K. Martens, T. Matsuda, T. Matsuyama, J. N. Matthews, M. Minamino, K. Mukai, I. Myers, K. Nagasawa, S. Nagataki, T. Nakamura, T. Nonaka, A. Nozato, S. Ogio, J. Ogura, M. Ohnishi, H. Ohoka, K. Oki, T. Okuda, M. Ono, A. Oshima, S. Ozawa, I. H. Park, M. S. Pshirkov, D. C. Rodriguez, G. Rubtsov, D. Ryu, H. Sagawa, N. Sakurai, A. L. Sampson, L. M. Scott, P. D. Shah, F. Shibata, T. Shibata, H. Shimodaira, B. K. Shin, J. D. Smith, P. Sokolsky, R. W. Springer, B. T. Stokes, S. R. Stratton, T. A. Stroman, T. Suzawa, M. Takamura, M. Takeda, R. Takeishi, A. Taketa, M. Takita, Y. Tameda, H. Tanaka, K. Tanaka, M. Tanaka, S. B. Thomas, G. B. Thomson, P. Tinyakov, I. Tkachev, H. Tokuno, T. Tomida, S. Troitsky,

- Y. Tsunesada, K. Tsutsumi, Y. Uchihori, S. Udo, F. Urban, G. Vasiloff, T. Wong, R. Yamane, H. Yamaoka, K. Yamazaki, J. Yang, K. Yashiro, Y. Yoneda, S. Yoshida, H. Yoshii, R. Zollinger, and Z. Zundel. INDICATIONS OF INTERMEDIATE-SCALE ANISOTROPY OF COSMIC RAYS WITH ENERGY GREATER THAN 57 EeV IN THE NORTHERN SKY MEASURED WITH THE SURFACE DETECTOR OF THE TELESCOPE ARRAY EXPERIMENT. *The Astrophysical Journal*, 790(2):L21, jul 2014.
- [54] Stephanie Wissel et al. Prospects for high-elevation radio detection of 010 PeV tau neutrinos. *Journal of Cosmology and Astroparticle Physics*, 2020(11):065–065, nov 2020.
- [55] A.G. Vieregg, K. Bechtol, and A. Romero-Wolf. A technique for detection of peV neutrinos using a phased radio array. *Journal of Cosmology and Astroparticle Physics*, 2016(02):005–005, Feb 2016.
- [56] M G Aartsen et al. Icecube-gen2: the window to the extreme universe. *Journal of Physics G: Nuclear and Particle Physics*, 48(6):060501, Apr 2021.
- [57] A. Nelles et al. Measuring a cherenkov ring in the radio emission from air showers at 110–190mhz with lofar. *Astroparticle Physics*, 65:11–21, 2015.
- [58] G. S. Varner, Larry L. Ruckman, Pisin. Chen, Chih-Ching Chen, and C. C. Yao. Specifications for the iceray sampler (irs) asic. 2008.
- [59] Specifications for the iceray sampler (irs) asic. https://www.phys.hawaii.edu/~varner/IRS_spec_v01.pdf, 2009. Accessed: 2020-03-20.
- [60] William R Leo. *Techniques for nuclear and particle physics experiments: a how-to approach; 2nd ed.* Springer, Berlin, 1994.
- [61] Kaeli Hughes. https://github.com/vPhase/PA_Analysis/blob/master/, 2020.
- [62] Jarred M. Roberts, Gary S. Varner, Patrick Allison, Brendan Fox, Eric Oberla, Ben Rotter, and Stefan Spack. Lab4d: A low power, multi-gsa/s, transient digitizer with sampling timebase trimming capabilities. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 925:92–100, 2019.
- [63] P. Allison et al. Long-baseline horizontal radio-frequency transmission through polar ice. *Journal of Cosmology and Astroparticle Physics*, 2020(12):009–009, Dec 2020.
- [64] J. R. Klein and A. Roodman. Blind analysis in nuclear and particle physics. *Ann. Rev. Nucl. Part. Sci.*, 55:141–163, 2005.
- [65] Krijn D. de Vries et al. The cosmic-ray air-shower signal in askaryan radio detectors. *Astroparticle Physics*, 74:96–104, February 2016.

- [66] R. A. Fisher. The statistical utilization of multiple measurements. *Annals of Eugenics*, 8(4):376–386, 1938.
- [67] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, Apr 1998.
- [68] Kaeli Hughes, J. Alvarez-Muñiz, W. Carvalho, A. Cummings, C. Deaconu, G. Hallinan, A. Ludwig, E. Oberla, C. Paciaroni, A. Rodriguez, A. Romero-Wolf, H. Schoorlemmer, D. Southall, B. Strutt, M. Vasquez, A. Viereg, S. Wissel, and E. Zas. Towards interferometric triggering on air showers induced by tau neutrino interactions. *Proceedings of Science*, 358, 2019.
- [69] Jaime Álvarez-Muñiz, Rafael Alves Batista, Aswathi Balagopal V., Julien Bolmont, Mauricio Bustamante, Washington Carvalho, Didier Charrier, Ismaël Cognard, Valentin Decoene, Peter B. Denton, Sijbrand De Jong, Krijn D. De Vries, Ralph Engel, Ke Fang, Chad Finley, Stefano Gabici, QuanBu Gou, JunHua Gu, Claire Guépin, HongBo Hu, Yan Huang, Kumiko Kotera, Sandra Le Coz, Jean-Philippe Lenain, Guo-Liang Lü, Olivier Martineau-Huynh, Miguel Mostafá, Fabrice Mottez, Kohta Murase, Valentin Niess, Foteini Oikonomou, Tanguy Pierog, XiangLi Qian, Bo Qin, Duan Ran, Nicolas Renault-Tinacci, Markus Roth, Frank G. Schröder, Fabian Schüssler, Cyril Tasse, Charles Timmermans, Matías Tüeros, XiangPing Wu, Philippe Zarka, Andreas Zech, B. Theodore Zhang, JianLi Zhang, Yi Zhang, Qian Zheng, and Anne Zilles. The giant radio array for neutrino detection (GRAND): Science and design. *Science China Physics, Mechanics & Astronomy*, 63(1), aug 2019.
- [70] The white mountain research station. <https://www.wmrc.edu/>.
- [71] Stephanie Wissel et al. Concept study for the beamforming elevated array for cosmic neutrinos (beacon). *ICRC Conference Proceedings*, 2019.
- [72] Michael W. Eastwood, Marin M. Anderson, Ryan M. Monroe, Gregg Hallinan, Benjamin R. Barsdell, Stephen A. Bourke, M. A. Clark, Steven W. Ellingson, Jayce Dowell, and Hugh Garsden. The radio sky at meter wavelengths: m-mode analysis imaging with the ovro-lwa. *Astron. J.*, 156(1):32, Jul 2018.
- [73] In-situ testing of frequency labeling and aliasing of dp and measurement of signal cross-coupling at lwa1. 2014.
- [74] S. Cabler. Airworthiness approval of automatic dependents surveillance- broadcast out systems, 2015.
- [75] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [76] Giulia De Bonis. Recent results of the antares neutrino telescope. *Nuclear and Particle Physics Proceedings*, 273-275:419–424, 2016. 37th International Conference on High Energy Physics (ICHEP).

- [77] E Zas. Searches for neutrino fluxes in the eev regime with the pierre auger observatory. *Proc. ICRC*, 301:972, 2017.
- [78] IceCube. Evidence for high-energy extraterrestrial neutrinos at the icecube detector. *Science*, 342(6161):1242856–1242856, 2013.
- [79] Markus Ackermann et al. Astrophysics uniquely enabled by observations of high-energy cosmic neutrinos. *Bulletin of the American Astronomical Society*, 51 (3)(3):185, 2019.
- [80] The ARA Collaboration. <https://github.com/ara-software/AraSim>, 2020.
- [81] M. G. Aartsen et al. Probing the origin of cosmic rays with extremely high energy neutrinos using the icecube observatory. *Phys. Rev.*, D88:112008, 2013.
- [82] Ilya Kravchenko, David Besson, and Josh Meyers. In situ index-of-refraction measurements of the south polar firn with the rice detector. *Journal of Glaciology*, 50(171):522–532, 2004.
- [83] M. G. Aartsen et al. A combined maximum-likelihood analysis of the high-energy astrophysical neutrino flux measured with icecube. *Astrophys. J.*, 809(1):98, 2015.
- [84] P. Allison et al. Measurement of the real dielectric permittivity ϵ_r of glacial ice. *Astropart. Phys.*, 108:63–73, 2019.
- [85] Denis Allard, M. Ave, N. Busca, M. A. Malkan, A. V. Olinto, E. Parizot, F. W. Stecker, and T. Yamamoto. Cosmogenic neutrinos from the propagation of ultrahigh energy nuclei. *JCAP*, 0609(09):005, 2006.
- [86] P. Allison et al. First constraints on the ultra-high energy neutrino flux from a prototype station of the askaryan radio array. *Astropart. Phys.*, 70:62–80, 2015.
- [87] P. Allison et al. Performance of two askaryan radio array stations and first results in the search for ultrahigh energy neutrinos. *Phys. Rev.*, D93(8):082003, 2016.
- [88] P. Allison et al. Constraints on the ultra-high energy neutrino flux from gamma-ray bursts from a prototype station of the askaryan radio array. *Astropart. Phys.*, 88:7–16, 2017.
- [89] P. Allison et al. Observation of reconstructable radio emission coincident with an x-class solar flare in the askaryan radio array prototype station. 2018.
- [90] Jaime Alvarez-Muniz, Andres Romero-Wolf, and Enrique Zas. Practical and accurate calculations of askaryan radiation. *Phys. Rev.*, D84:103003, 2011.
- [91] Jaime Alvarez-Muniz, Andres Romero-Wolf, and Enrique Zas. Time-domain radio pulses from particle showers. *Nucl. Instrum. Meth.*, A662:S32–S35, 2012.
- [92] Jaime Álvarez Muñoz et al. The giant radio array for neutrino detection (grand): Science and design. *Sci. China Phys. Mech. Astron.*, 63(1):219501, 2020.

- [93] A. Anker et al. A search for cosmogenic neutrinos with the arianna test bed using 4.5 years of data. *JCAP*, 03(03):053, 2020.
- [94] J. Avva, J. M. Kovac, C. Miki, D. Saltzberg, and A. G. Vieregg. An in situ measurement of the radio-frequency attenuation in ice at summit station, greenland. *J. Glaciol.*, 61:1005–1011, 2015.
- [95] G. A. Askar’yan. Coherent radio emission from cosmic showers in air and in dense media. *Soviet Journal of Experimental and Theoretical Physics*, 21:658, Sep 1965.
- [96] S. W. Barwick et al. Design and performance of the arianna hra-3 neutrino detector systems. *IEEE Trans. Nucl. Sci.*, 62(5):2202–2215, 2015.
- [97] V. S. Berezinsky and G. T. Zatsepin. Cosmic rays at ultrahigh-energies (neutrino?). *Phys. Lett.*, 28B:423–424, 1969.
- [98] B. L. Berman and S. C. Fultz. Measurements of the giant dipole resonance with monoenergetic photons. *Rev. Mod. Phys.*, 47:713–761, 1975.
- [99] A. Bellétoile, R. Dallier, A. Lecacheux, V. Marin, L. Martin, B. Revenu, and D. Torres. Evidence for the charge-excess contribution in air shower radio emission observed by the codalema experiment. *Astropart. Phys.*, 69:50–60, 2015.
- [100] Amy Connolly, Robert S. Thorne, and David Waters. Calculation of high energy neutrino-nucleon cross sections and uncertainties using the mstw parton distribution functions and implications for future experiments. *Phys. Rev.*, D83:113009, 2011.
- [101] Jan Conrad, O. Botner, A. Hallgren, and Carlos Perez de los Heros. Including systematic uncertainties in confidence interval construction for poisson statistics. *Phys. Rev.*, D67:012002, 2003.
- [102] Amanda Cooper-Sarkar, Philipp Mertsch, and Subir Sarkar. The high energy neutrino cross-section in the standard model and its uncertainty. *JHEP*, 08:042, 2011.
- [103] Brian Dailey. *Analysis of the second flight of the ANtarctic Impulsive Transient Antenna with a focus on filtering techniques*. PhD thesis, The Ohio State University, 2017.
- [104] P. W. Gorham et al. The antarctic impulsive transient antenna ultra-high energy neutrino detector design, performance, and sensitivity for 2006-2007 balloon flight. *Astropart. Phys.*, 32:10–41, 2009.
- [105] Robert J. Gould and Gérard P. Schréder. Opacity of the universe to high-energy photons. *Phys. Rev.*, 155:1408–1411, Mar 1967.
- [106] Christian Haack and Christopher Wiebusch. A measurement of the diffuse astrophysical muon neutrino flux using eight years of icecube data. *PoS, ICRC2017*:1005, 2018.

- [107] Jordan C. Hanson. Ross ice shelf thickness, radio-frequency attenuation and reflectivity: Implications for the arianna uhe neutrino detector. In *Proceedings, ICRC 2011*, volume 4, page 169, 2011.
- [108] Jordan C. Hanson and Amy L. Connolly. Complex analysis of askaryan radiation: A fully analytic treatment including the lpm effect and cascade form factor. *Astropart. Phys.*, 91:75–89, 2017.
- [109] Gary C. Hill. Comment on ‘including systematic uncertainties in confidence interval construction for poisson statistics’. *Phys. Rev.*, D67:118101, 2003.
- [110] Eugene Hong. *Searching for Ultra-high Energy Neutrinos with Data from a Prototype Station of the Askaryan Radio Array*. PhD thesis, The Ohio State University, 2014.
- [111] Dan Hooper, Andrew Taylor, and Subir Sarkar. The impact of heavy nuclei on the cosmogenic neutrino flux. *Astropart. Phys.*, 23:11–17, 2005.
- [112] Ming-Yuan Lu, Carl G Pfendner, and Andrew Shultz. Ultra-high energy neutrino search with the askaryan radio array. *PoS, ICRC2017:966*, 2018.
- [113] Kohta Murase. *Active Galactic Nuclei as High-Energy Neutrino Sources*, pages 15–31. World Scientific, 2017.
- [114] Buitink, S., Scholten, O., Bacelar, J., Braun, R., de Bruyn, A. G., Falcke, H., Singh, K., Stappers, B., Strom, R. G., and al Yahyaoui, R. Constraints on the flux of ultra-high energy neutrinos from westerbork synthesis radio telescope observations. *A&A*, 521:A47, 2010.
- [115] M. Tanabashi and others. Review of particle physics. *Phys. Rev. D*, 98:030001, Aug 2018.
- [116] Jarred M. Roberts, Eric Oberla, Patrick Allison, Gary S. Varner, Stefan Spack, Brendan Fox, and Ben Rotter. Lab4d: A low power, multi-gsa/s, transient digitizer with sampling timebase trimming capabilities. *Nucl. Instrum. Meth.*, A925:92–100, 2019.
- [117] A. Romero-Wolf et al. An interferometric analysis method for radio impulses from ultra-high energy particle showers. *Astropart. Phys.*, 60:72–85, 2015.
- [118] P. Schellart et al. Detecting cosmic rays with the lofar radio telescope. *Astron. Astrophys.*, 560:A98, 2013.
- [119] O. Scholten et al. Measurement of the circular polarization in radio emission from extensive air showers confirms emission mechanisms. *Phys. Rev.*, D94(10):103010, 2016.
- [120] Jessymol K. Thomas, Reetanjali Moharana, and Soebur Razzaque. Ultrahigh energy neutrino afterglows of nearby long duration gamma-ray bursts. *Phys. Rev.*, D96(10):103004, 2017.

- [121] Arjen van Vliet, Jörg R. Hörandel, and Rafael Alves Batista. Cosmogenic gamma-rays and neutrinos constrain uhecr source models. *PoS, ICRC2017:562*, 2018. [35,562(2017)].
- [122] Eli Waxman and John N. Bahcall. Neutrino afterglow from gamma-ray bursts: Similar to 10^{18} -ev. *Astrophys. J.*, 541:707–711, 2000.
- [123] G. S. Varner, L. L. Ruckman, P. W. Gorham, J. W. Nam, R. J. Nichol, J. Cao, and M. Wilcox. The large analog bandwidth recorder and digitizer with ordered readout (labrador) asic. *Nucl. Instrum. Meth.*, A583:447–460, 2007.
- [124] The ANITA Collaboration. Constraints on cosmic neutrino fluxes from the anita experiment. *Phys. Rev. Lett.*, 96, 2005.
- [125] K. Kotera and A. V. Olinto. The astrophysics of ultrahigh-energy cosmic rays. *Annual Review of Astronomy and Astrophysics*, 49:119–153, September 2011.
- [126] G.A. Askaryan. Excess negative charge of an electron-photon shower and its coherent radio emission. *Sov.Phys.JETP*, 1961.
- [127] Steffen Hallmann et al. Sensitivity studies for the icecube-gen2 radio array. *PoS, ICRC2021:1183*, 2021.
- [128] Paramita Dasgupta et al. The calibration of the geometry and antenna delay in askaryan radio array station 4 and 5. *PoS, ICRC2021:1086*, 2021.
- [129] David Besson Ilya Kravchenko and Josh Meyers. In situ index-of-refraction measurements of the south polar firn with the rice detector. *Cambridge University Press*, 2017.
- [130] C. Deaconu et al. Measurements and modeling of near-surface radio propagation in glacial ice and implications for neutrino experiments. *Phys. Rev. D*, 98:043010, Aug 2018.
- [131] Edwin B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927.
- [132] M. Mikhailova et al. Triboelectric backgrounds to radio-based uhe neutrino experiments. 3 2021.
- [133] S. Barwick et al. South polar in situ radio-frequency ice attenuation. *Journal of Glaciology*, 51(173):231–238, 2005.
- [134] K. D. de Vries, A. M. van den Berg, O. Scholten, and K. Werner. Coherent cherenkov radiation from cosmic-ray-induced air showers. *Phys. Rev. Lett.*, 107:061101, Aug 2011.
- [135] J. Quinn et al. Detection of gamma rays with [ITAL]e[/ITAL]] 300 GeV from markarian 501. *The Astrophysical Journal*, 456(2), jan 1996.

- [136] M. Punch, C. W. Akerlof, M. F. Cawley, M. Chantell, D. J. Fegan, S. Fennell, J. A. Gaidos, J. Hagan, A. M. Hillas, Y. Jiang, A. D. Kerrick, R. C. Lamb, M. A. Lawrence, D. A. Lewis, D. I. Meyer, G. Mohanty, K. S. O’Flaherty, P. T. Reynolds, A. C. Rovero, M. S. Schubnell, G. Sembroski, T. C. Weekes, T. Whitaker, and C. Wilson. Detection of tev photons from the active galaxy markarian 421. , 358(6386):477–478, August 1992.
- [137] M. Ageron, J.A. Aguilar, I. Al Samarai, A. Albert, F. Ameli, M. André, M. Anghinolfi, G. Anton, S. Anvar, M. Ardid, K. Arnaud, E. Aslanides, A.C. Assis Jesus, T. Astraatmadja, J.-J. Aubert, R. Auer, E. Barbarito, B. Baret, S. Basa, M. Bazzotti, Y. Becherini, J. Beltramelli, A. Bersani, V. Bertin, S. Beurthey, S. Biagi, C. Bigongiari, M. Billault, R. Blaes, C. Bogazzi, N. de Botton, M. Bou-Cabo, B. Boudahef, M.C. Bouwhuis, A.M. Brown, J. Brunner, J. Busto, L. Caillat, A. Calzas, F. Camarena, A. Capone, L. Caponetto, C. Cârloganu, G. Carminati, E. Carmona, J. Carr, P.H. Carton, B. Cassano, E. Castorina, S. Cecchini, A. Ceres, Th. Chaleil, Ph. Charvis, P. Chauchot, T. Chiarusi, M. Circella, C. Compère, R. Coniglione, X. Coppolani, A. Cosquer, H. Costantini, N. Cottini, P. Coyle, S. Cuneo, C. Curtil, C. D’Amato, G. Damy, R. van Dantzig, G. De Bonis, G. Decock, M.P. Decowski, I. Dekeyser, E. Delagnes, F. Desages-Ardellier, A. Deschamps, J.-J. Destelle, F. Di Maria, B. Dinke-spiler, C. Distefano, J.-L. Dominique, C. Donzaud, D. Dornic, Q. Dorosti, J.-F. Drogou, D. Drouhin, F. Druillole, D. Durand, R. Durand, T. Eberl, U. Emanuele, J.J. Engelen, J.-P. Ernenwein, S. Escoffier, E. Falchini, S. Favard, F. Fehr, F. Feinstein, M. Ferri, S. Ferry, C. Fiorello, V. Flaminio, F. Folger, U. Fritsch, J.-L. Fuda, S. Galatá, S. Galeotti, P. Gay, F. Gensolen, G. Giacomelli, C. Gojak, J.P. Gómez-González, Ph. Goret, K. Graf, G. Guillard, G. Halladjian, G. Hallowell, H. van Haren, B. Hartmann, A.J. Heijboer, E. Heine, Y. Hello, S. Henry, J.J. Hernández-Rey, B. Herold, J. Höbl, J. Hogenbirk, C.C. Hsu, J.R. Hubbard, M. Jaquet, M. Jaspers, M. de Jong, D. Jourde, M. Kadler, N. Kalantar-Nayestanaki, O. Kalekin, A. Kappes, T. Karg, S. Karkar, M. Karolak, U. Katz, P. Keller, P. Kestener, E. Kok, H. Kok, P. Kooijman, C. Kopper, A. Kouchner, W. Kretschmer, A. Kruijjer, S. Kuch, V. Kulikovskiy, D. Lachartre, H. Lafoux, P. Lagier, R. Lahmann, C. Lahonde-Hamdoun, P. Lamare, G. Lambard, J.-C. Languillat, G. Larosa, J. Lavalle, Y. Le Guen, H. Le Provost, A. LeVanSuu, D. Lefèvre, T. Legou, G. Lelaizant, C. Lévêque, G. Lim, D. Lo Presti, H. Loehner, S. Loucatos, F. Louis, F. Lucarelli, V. Lyashuk, P. Magnier, S. Mangano, A. Marcel, M. Marcelin, A. Margiotta, J.A. Martinez-Mora, R. Masullo, F. Mazéas, A. Mazure, A. Meli, M. Melissas, E. Migneco, M. Mongelli, T. Montaruli, M. Morganti, L. Moscoso, H. Motz, M. Musumeci, C. Naumann, M. Naumann-Godo, M. Neff, V. Niess, G.J.L. Nooren, J.E.J. Oberski, C. Olivetto, N. Palanque-Delabrouille, D. Palioselitis, R. Papaleo, G.E. Pāvālaš, K. Payet, P. Payre, H. Peek, J. Petrovic, P. Piattelli, N. Picot-Clemente, C. Picq, Y. Piret, J. Poinignon, V. Popa, T. Pradier, E. Presani, G. Prono, C. Racca, G. Raia, J. van Randwijk, D. Real, C. Reed, F. Réthoré, P. Rewiersma, G. Riccobene, C. Richardt, R. Richter, J.S. Ricol, V. Rigaud, V. Roca, K. Roensch, J.-F. Rolin, A. Rostovtsev, A. Rottura, J. Roux, M. Rujoiu, M. Ruppi, G.V. Russo, F. Salesa, K. Salomon, P. Sapienza, F. Schmitt, F. Schöck, J.-P. Schuller, F. Schüssler,

D. Sciliberto, R. Shanidze, E. Shirokov, F. Simeone, A. Sottoriva, A. Spies, T. Spona, M. Spurio, J.J.M. Steijger, Th. Stolarczyk, K. Streeb, L. Sulak, M. Taiuti, C. Tamburini, C. Tao, L. Tasca, G. Terreni, D. Tezier, S. Toscano, F. Urbano, P. Valdy, B. Valage, V. Van Elewyck, G. Vannoni, M. Vecchi, G. Venekamp, B. Verlaat, P. Vernin, E. Virique, G. de Vries, R. van Wijk, G. Wijnker, G. Wobbe, E. de Wolf, Y. Yakovenko, H. Yepes, D. Zaborov, H. Zacccone, J.D. Zornoza, and J. Zúñiga. ANTARES: The first undersea neutrino telescope. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 656(1):11–38, nov 2011.