Elsevier Editorial System(tm) for Nuclear Inst. and Methods in Physics Research, A
Manuscript Draft

Manuscript Number: NIMA-D-17-00920R1

Title: Dynamic tunable notch filters for the Antarctic Impulsive Transient Antenna (ANITA)

Article Type: Full length article

Section/Category: High Energy and Nuclear Physics Detectors

Keywords: neutrino radio detection,
ultra-high-energy,
notch filtering,
military communications satellites

Corresponding Author: Ms. Oindree Banerjee, M.Sc.

Corresponding Author's Institution: The Ohio State University

First Author: ANITA Collaboration

Order of Authors: ANITA Collaboration; Oindree Banerjee, M.Sc.

Abstract: The Antarctic Impulsive Transient Antenna (ANITA) is a NASA long-duration balloon experiment
with the primary goal of detecting ultra-high-energy
($>10^{18}\,\mbox{eV}$) neutrinos via the Askaryan Effect.
The fourth ANITA mission, ANITA-IV, recently flew from Dec~2 to Dec~29, 2016.
For the first time,
the Tunable Universal Filter Frontend (TUFF) boards were deployed
for mitigation of narrow-band, anthropogenic noise with tunable, switchable notch filters.
The TUFF boards also performed second-stage amplification by approximately
45~dB to boost the $\sim\,\mu\mbox{V-level}$ radio frequency (RF) signals to $\sim$ mV-level for digitization, and
supplied power via bias tees to the first-stage, antenna-mounted amplifiers.
The other major change in signal processing in ANITA-IV is the resurrection of the
90° hybrids deployed previously in ANITA-I, in the trigger system, although in this paper we focus on the TUFF boards.
During the ANITA-IV mission, the TUFF boards were successfully operated throughout the flight.
They contributed to
a factor of 2.8
higher total instrument livetime on average in ANITA-IV compared to ANITA-III
due to reduction of narrow-band, anthropogenic noise before a trigger decision is made.

Dynamic tunable notch filters for the Antarctic Impulsive Transient Antenna (ANITA)

P. Allison[a], O. Banerjee[a,*], J. J. Beatty[a], A. Connolly[a], C. Deaconu[g], J. Gordon[a], P. W. Gorham[h], M. Kovacevich[a], C. Miki[h], E. Oberla[g], J. Roberts[h], B. Rotter[h], S. Stafford[a], K. Tatem[h], L. Batten[b], K. Belov[c], D. Z. Besson[d,1], W. R. Binns[e], V. Bugaev[e], P. Cao[f], C. Chen[i], P. Chen[i], Y. Chen[i], J. M. Clem[f], L. Cremonesi[b], B. Dailey[a], P. F. Dowkontt[j], S.Hsu[i], J. Huang[i], R. Hupe[a], M. H. Israel[e], J. Kowalski[h], J. Lam[j], J. G. Learned[h], K. M. Liewer[c], T. C. Liu[i], A. Ludwig[g], S. Matsuno[h], K. Mulrey[f], J. Nam[i], R. J. Nichol[b], A. Novikov[d,1], S. Prohira[d], B. F. Rauch[e], J. Ripa[i], A. Romero-Wolf[c], J. Russell[h], D. Saltzberg[j], D. Seckel[f], J. Shiao[i], J. Stockham[d], M. Stockham[d], B. Strutt[j], G. S. Varner[h], A. G. Vieregg[g], S. Wang[i], S. A. Wissel[l], F. Wu[j], R. Young[d]

[a] *Dept. of Physics, The Ohio State Univ., Columbus, OH 43210; Center for Cosmology and AstroParticle Physics.*
[b] *Dept. of Physics and Astronomy, University College London, London, United Kingdom.*
[c] *Jet Propulsion Laboratory, Pasadena, CA 91109.*
[d] *Dept. of Physics and Astronomy, Univ. of Kansas, Lawrence, KS 66045.*
[e] *Dept. of Physics, Washington Univ. in St. Louis, MO 63130.*
[f] *Dept. of Physics, Univ. of Delaware, Newark, DE 19716.*
[g] *Dept. of Physics, Enrico Fermi Institute, Kavli Institute for Cosmological Physics, Univ. of Chicago , Chicago IL 60637.*
[h] *Dept. of Physics and Astronomy, Univ. of Hawaii, Manoa, HI 96822.*
[i] *Dept. of Physics, Grad. Inst. of Astrophys., Leung Center for Cosmology and Particle Astrophysics, National Taiwan University, Taipei, Taiwan.*
[j] *Dept. of Physics and Astronomy, Univ. of California, Los Angeles, Los Angeles, CA 90095.*
[k] *National Research Nuclear University, Moscow Engineering Physics Institute, 31 Kashirskoye Highway, Russia 115409*
[l] *Dept. of Physics, California Polytechnic State Univ., San Luis Obispo, CA 93407.*

Abstract

*Corresponding author
Email address: oindreeb@gmail.com (O. Banerjee)

The Antarctic Impulsive Transient Antenna (ANITA) is a NASA long-duration balloon experiment with the primary goal of detecting ultra-high-energy ($> 10^{18}$ eV) neutrinos via the Askaryan Effect. The fourth ANITA mission, ANITA-IV, recently flew from Dec 2 to Dec 29, 2016. For the first time, the Tunable Universal Filter Frontend (TUFF) boards were deployed for mitigation of narrow-band, anthropogenic noise with tunable, switchable notch filters. The TUFF boards also performed second-stage amplification by approximately 45 dB to boost the $\sim$ μV-level radio frequency (RF) signals to $\sim$ mV-level for digitization, and supplied power via bias tees to the first-stage, antenna-mounted amplifiers. The other major change in signal processing in ANITA-IV is the resurrection of the 90° hybrids deployed previously in ANITA-I, in the trigger system, although in this paper we focus on the TUFF boards. During the ANITA-IV mission, the TUFF boards were successfully operated throughout the flight. They contributed to a factor of 2.8 higher total instrument livetime on average in ANITA-IV compared to ANITA-III due to reduction of narrow-band, anthropogenic noise before a trigger decision is made.

Keywords: neutrino radio detection, ultra-high-energy, notch filtering, military communications satellites

1. Introduction

The Antarctic Impulsive Transient Antenna (ANITA) is a NASA long-duration balloon-borne ultra-high-energy (UHE) neutrino detector [1]. ANITA looks for radio impulses produced via the Askaryan Effect by UHE neutrinos interacting in the Antarctic ice. The Askaryan Effect, as formulated by Askaryan *et al.* [2] and observed in ice by the ANITA collaboration in a beam test [3], is the production of coherent Cherenkov radio impulses due to a charged particle shower traveling in a dielectric medium at a speed faster than the speed of light in that medium.

The fourth ANITA flight, ANITA-IV, was launched on Dec 2, 2016 from the NASA Long Duration Balloon (LDB) Facility located 10 km from McMurdo Station in Antarctica. The flight was terminated on Dec 29, 2016 and landed approximately 100 km from the South Pole. The Tunable Universal Filter Frontend (TUFF) boards were deployed for the first time in the ANITA-IV mission, and are the subject of this paper.

1.1. Continuous-wave (CW) interference

The principal challenge of the ANITA experiment is to distinguish neutrino signals from radio frequency (RF) noise. The two main sources of noise are thermal radiation by the Antarctic ice and anthropogenic noise, much of which is modulated continuous-wave (CW) interference.

While Antarctica itself is relatively free of CW transmissions, except for bases of human activity, transmissions from geosynchronous satellites are continuously in view. The average full-width-at-half-maximum (FWHM) beamwidth of the ANITA antennas is approximately 45°. Although the ANITA antennas are canted downward by 10°, the beam of the antennas extends to horizontal from the perspective of the payload and above. The Antarctic science bases, the most prominent being McMurdo and South Pole Station, are more radio-loud than the rest of the continent, producing CW interference, for example, in the $430 - 460\,\mathrm{MHz}$ band.

CW interference due to military satellites has affected all ANITA flights. ANITA-I (Dec. 2006 - Jan. 2007) and ANITA-II (Dec. 2008 - Jan. 2009) observed CW interference primarily in the $240 - 270\,\mathrm{MHz}$ band, peaking at $260\,\mathrm{MHz}$. This frequency range is predominantly used by the aging Fleet Satellite (FLTSAT) Communications System and the Ultra High Frequency Follow-On (UFO) System, both serving the United States Department of Defense since year 1978 and 1993 respectively. In addition to CW interference at $260\,\mathrm{MHz}$, ANITA-III (Dec. 2014 - Jan. 2015) observed CW interference at $375\,\mathrm{MHz}$ which is thought to be due to the newer Mobile User Objective System (MUOS) satellites that were launched during the period from Feb. 2012 - June 2016 [4]. The CW signals generate events with excess power in left circular polarization (as expected) above the horizon, in approximately stationary positions.

The ANITA-III experiment was most affected by CW interference due to military satellites. Due to the design of the ANITA-I and ANITA-II trigger, which required coincidences among different frequency bands, the CW interference did not overwhelm the acquisition system. However, ANITA-III was redesigned for improved sensitivity and based its trigger decisions on full-bandwidth ($200 - 1200\,\mathrm{MHz}$) signals. The modulation present in the CW interference produced trigger rates far in excess of the digitization system's readout capabilities ($\sim 50\,\mathrm{Hz}$) for thresholds comparable to those used in previous flights. Thus, the ANITA-III experiment was susceptible to digitization deadtime (defined in Section 6.1) throughout the flight.

The lesson learned from the ANITA-III flight was that a new method of mitigation of CW signal had to be a priority for the ANITA-IV flight. Before ANITA-IV, the available methods to reduce digitization deadtime were masking and decreasing thresholds (described in Sections 3 and 6.2) when in the presence of higher levels of noise. A decrease in thresholds corresponds to higher power of the incoming signal as explained in Section 3.1. Masking and decreasing thresholds come at the cost of instrument livetime (defined in Section 6.1) and sensitivity to neutrinos, respectively. For about 90% of the time during the ANITA-III flight, masking was used to veto triggers from over half of the payload field-of-view to keep the trigger rate at or below 50 Hz. This significantly lowered the total instrument livetime. For ANITA-IV, the TUFF boards were built with tunable notch filters to restore triggering efficiencies in the presence of CW interference. Additionally, the 90° hybrids, previously deployed in ANITA-I as described in our design paper [1], were added to the ANITA-IV trigger system to require signals to be linearly polarized.

2. ANITA Payload

The ANITA payload is designed to view the ice out to the horizon at 700 km distance with complete azimuthal coverage and good reconstruction capability, while its shape and size is constrained by its NASA launch vehicle "The Boss," pictured in Figure 1. The ANITA-III and ANITA-IV payloads each have 48 antennas. The antennas are arranged in three aligned rings of 16 antennas, termed the top, middle, and bottom rings. The top ring consists of two staggered sub-rings each having eight antennas.

The three rings of antennas and a phi sector of ANITA-IV are pictured in Figure 1. The FWHM beamwidth of the antennas is approximately 45°. The antennas in the top ring are evenly spaced by 45° in azimuth. The two sub-rings in the top ring are offset by 22.5° for uniform coverage. The antennas in the middle ring are evenly spaced by 22.5°. The antennas in the bottom ring are evenly spaced by 22.5°. All the antennas are angled downward by 10° to preferentially observe signals coming from the ice as opposed to from the sky. Each group of three antennas in a vertical column, taking one antenna from each ring, forms a phi sector, viewing a 22.5° region in azimuth.

The ANITA Instrument Box is placed on a deck above the middle ring of antennas, also seen in Figure 1. The Instrument Box contains different units for signal processing, as illustrated in Figure 2. More details on signal

processing are in Section 3. In ANITA-IV, the 12-channel TUFF modules reside inside four Internal Radio Frequency Conditioning Modules (IRFCMs) inside the Instrument Box.

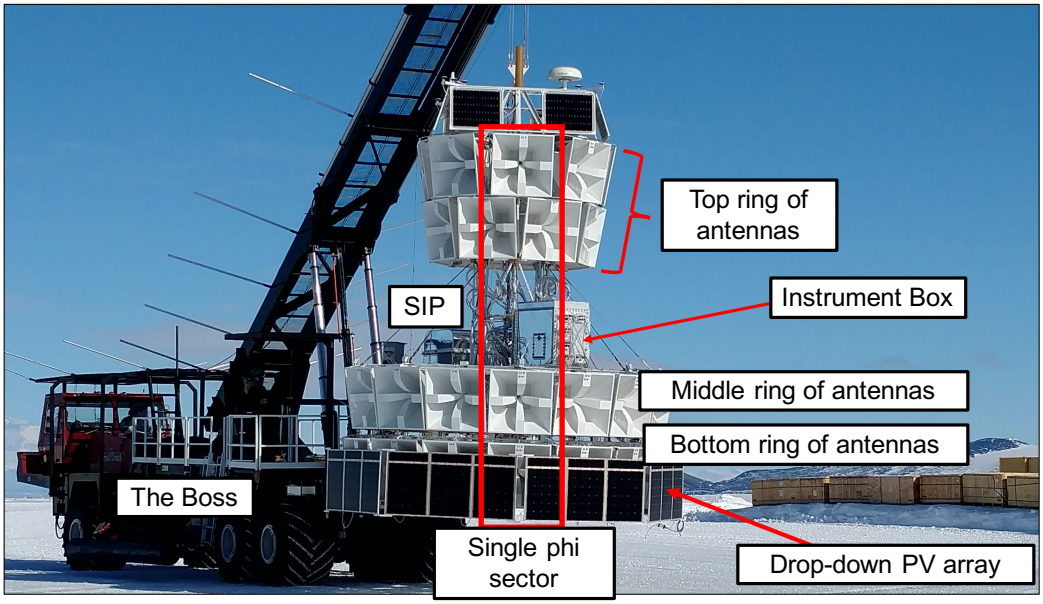

Figure 1: The ANITA-IV payload just prior to launch at the NASA LDB Facility near McMurdo Station, Antarctica. The red box encloses three antennas that make up a single phi sector.

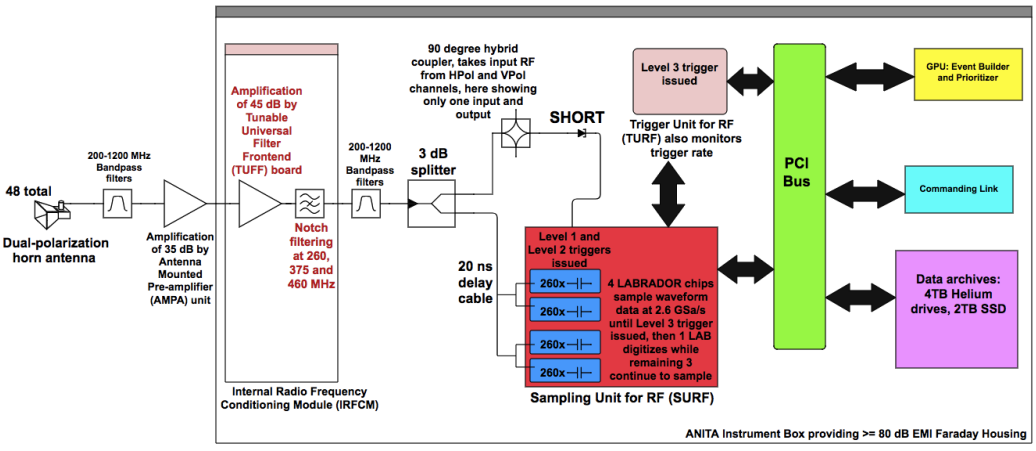

Figure 2: The ANITA-IV signal processing chain for a single RF channel.

The NASA Science Instrument Package (SIP) also sits on the deck. The SIP is powered and controlled by NASA. It is used for flight control such as ballast release and flight termination. The SIP also provides a connection to the ANITA payload during flight through line-of-sight transmission, the Iridium satellites, and the Tracking and Data Satellite System (TDRSS). This allows us to monitor the payload continuously during the flight. A small fraction of data (less than 1%) is transferred from the payload through telemetry. Commands to perform different functions, such as tuning a TUFF notch filter, can be sent to the payload in real time using the SIP connection.

3. ANITA Signal Processing

In this section we describe the signal processing chain for ANITA-IV, and in particular the steps that are relevant to understanding the role of the TUFF boards. We will note when and where the ANITA-III signal processing differed. The RF signal processing chain for ANITA-IV is illustrated in Figure 2. Each ANITA antenna is dual-polarized with feeds for vertically and horizontally polarized (VPol and HPol) signals. Therefore, for 48 antennas there are 96 total full-band ($200 - 1200\,\mathrm{MHz}$) RF signal channels.

Each channel goes through the Antenna-Mounted Pre-amplifier (AMPA) unit before entering the Instrument Box. There is an AMPA unit connected directly to the VPol and HPol outputs of each antenna. The AMPA contains a $200 - 1200\,\mathrm{MHz}$ bandpass filter, followed by an approximately 35 dB Low Noise Amplifier (LNA). Following the AMPA unit, the RF signal travels through 12 m of LMR240 coaxial cable to the Instrument Box. Inside the Instrument Box, the signal first goes through second-stage amplification (performed by a different module in ANITA-III) and notch filtering (unique to ANITA-IV), both performed by the TUFF boards in ANITA-IV. Then it passes through another set of bandpass filters before being split into digitization and triggering paths. The triggering and digitization processes are detailed below.

3.1. Triggering:

In the triggering path, the RF signals from both the VPol and HPol channels of a single antenna are passed through a 90° hybrid (hybrids were absent in ANITA-III). The outputs from the 90° hybrid are the left and right circular polarized (LCP and RCP) components of the combined VPol and HPol

signals from an antenna. The hybrid outputs are input to the SURF (Sampling Unit for RF) high-occupancy RF Trigger (SHORT) unit before being passed to the SURF board. Each SHORT takes four channels as its input. In a SHORT channel, the RF signal passes through a tunnel diode and an amplifier. The output of the SHORT is approximately proportional to the square of the voltages of the input RF signal integrated over approximately 5 ns. It is a measure of the power of the incoming signal and is typically a negative voltage. The SHORT output is routed to a SURF trigger input where it enters a discriminator that compares this negative voltage in Digital-to-Analog Converter (DAC) counts to the output of a software-controlled DAC threshold on the SURF, henceforth referred to as the SURF DAC threshold. The SURF DAC threshold is expressed in arbitrary units of DAC counts corresponding to voltages. Lower thresholds correspond to higher voltages and therefore, higher power of the incoming signal. The SURF DAC threshold can be changed during flight. During the ANITA-III flight, CW interference overwhelmed the digitization system, forcing us to impose frequent and large changes in the SURF DAC thresholds. A comparison of SURF DAC thresholds between ANITA-III and ANITA-IV is presented in Figure 3. Note that the lower overall threshold for ANITA-IV is primarily due to the modified triggering scheme, which requires more overall coincidences between channels. The increased stability of the ANITA-IV thresholds, due to the CW mitigation schemes presented here, is clearly apparent.

Trigger logic: Due to power and bandwidth limitations, ANITA is not able to constantly record data. Digitization of data only occurs when the trigger conditions are satisfied. The ANITA-IV trigger consists of three triggering levels: Level 1, Level 2 and Level 3. The trigger requirements at each of these three levels is described below.

Level 1 trigger: The Sampling Unit for RF (SURF) board issues the Level 1 trigger. To form a Level 1 trigger, the SHORT outputs of the LCP and RCP channels from the same antenna are required to exceed the SURF DAC threshold within 4 ns. This LCP/RCP coincidence requirement was added to the ANITA-IV trigger to mitigate anthropogenic and thermal backgrounds. The signals of interest are known to be linearly polarized, whereas satellite emission is often circularly polarized and thermal noise is unpolarized. In the presence of a continuous source of CW signal such as satellites, the LCP/RCP

coincidence may still allow a combination of circularly polarized satellite noise and the circularly polarized component of thermal noise to satisfy the Level 1 trigger requirement. Therefore, the LCP/RCP coincidence aids in reducing triggers induced by satellites but does not completely mitigate their effect.

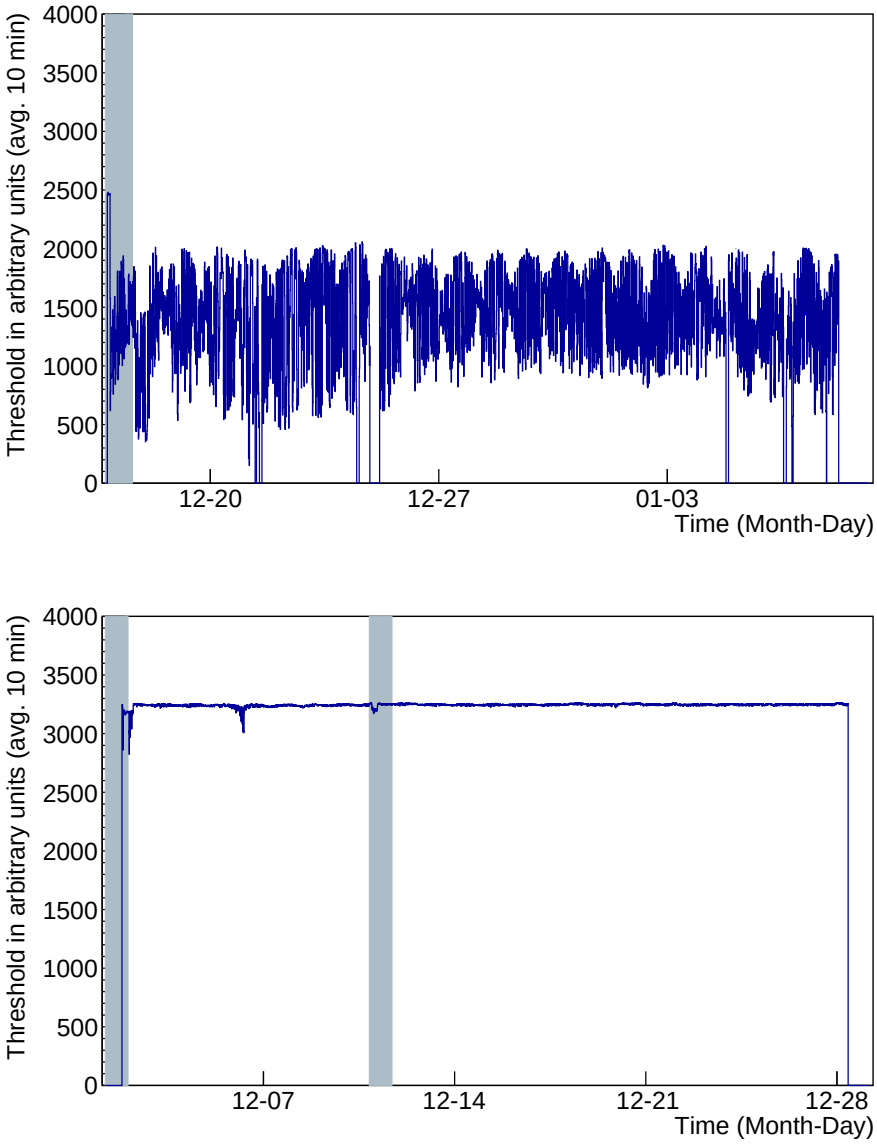

Figure 3: SURF DAC thresholds in arbitrary units of DAC counts for a single channel for the ANITA-III (top) and ANITA-IV (bottom) flights. Changing thresholds is a secondary method of avoiding digitization deadtime due to CW interference. The TUFF boards helped to maintain constant thresholds in ANITA-IV, whereas in ANITA-III, thresholds had to be changed throughout the flight. Note that a lower threshold corresponds to a higher and therefore, stricter requirement on the power of the incoming signal, and so during periods of high anthropogenic noise, the SURF DAC thresholds were lowered. The shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility.

Level 2 trigger: The SURF board issues the Level 2 trigger. A Level 1 trigger opens up a time window. If there are two Level 1 triggers in the same phi sector within the allowable time window, then a Level 2 trigger is issued. The allowable time window depends on which antenna had the first Level 1 trigger. Time windows of 16 ns, 12 ns and 4 ns in duration are opened up when a Level 1 trigger is issued in the bottom, middle and top ring respectively. These time windows were chosen to preferentially select signals coming up from the ice. The Level 2 trigger decisions are passed from the SURF boards to a dedicated triggering board called the Triggering Unit for RF (TURF). The Level 2 trigger timing in ANITA-IV differed from that used in ANITA-III as changes were made to further restrict the allowed timing of the antenna coincidences to better match timing expected from an incoming plane wave.

Level 3 trigger: The TURF board issues the Level 3 trigger. A field programmable gate array (FPGA) on the TURF board monitors Level 2 triggers. A Level 3 trigger is issued by the TURF board when there are Level 2 triggers in two adjacent phi sectors within 10 ns. When there is a Level 3 trigger, the TURF board instructs the SURF board to begin digitization.

3.2. Digitization:

The digitization of the signal is performed by the SURF board. There are twelve SURF boards, each containing four custom-built Application Specific Integrated Circuits called Large Analog Bandwidth Recorder And Digitizer with Ordered Readout (LABRADOR).

LABRADOR chip and digitization deadtime: ANITA-IV uses the third generation of LABRADOR chips that are described by Varner *et al.* [5]. Each LABRADOR chip has a 260-element switched capacitor array (SCA) for each of its 9 input channels, with one channel used for timing synchronization. The RF signal entering a SURF gets split and fed into four parallel LABRADOR chips (forming four "buffers" for digitization). The SCAs sample waveform data at the rate of 2.6 GSa/s. At any moment, the charge stored in an SCA is a 100 ns record of the signal voltage. This 100 ns snapshot of the incoming plane wave is known as an "event." When a Level 3 trigger occurs, a single LABRADOR chip stops sampling and is "held." It then digitizes the stored data, which is then read out by the flight computer, taking approximately $5 - 10$ ms. If all four LABRADOR chips are held,

the trigger is "dead" and the accumulated time when the trigger is dead is recorded as digitization deadtime by the TURF board.

Masking: During ANITA-III, digitization deadtime due to high levels of anthropogenic noise was reduced by excluding certain phi sectors from participating in the Level 3 trigger. This is called phi-masking. Alternatively, specific channels (each antenna has two channels) were excluded from participating in the Level 1 trigger. This is called channel-masking. Together these are referred to as masking. Because of CW interference by military communications satellites, over half of the payload had to be masked during most of the ANITA-III flight. This strongly motivated the creation of the TUFF boards with tunable, switchable notch filters. A comparison of masking between ANITA-III and ANITA-IV is presented in Figure 4.

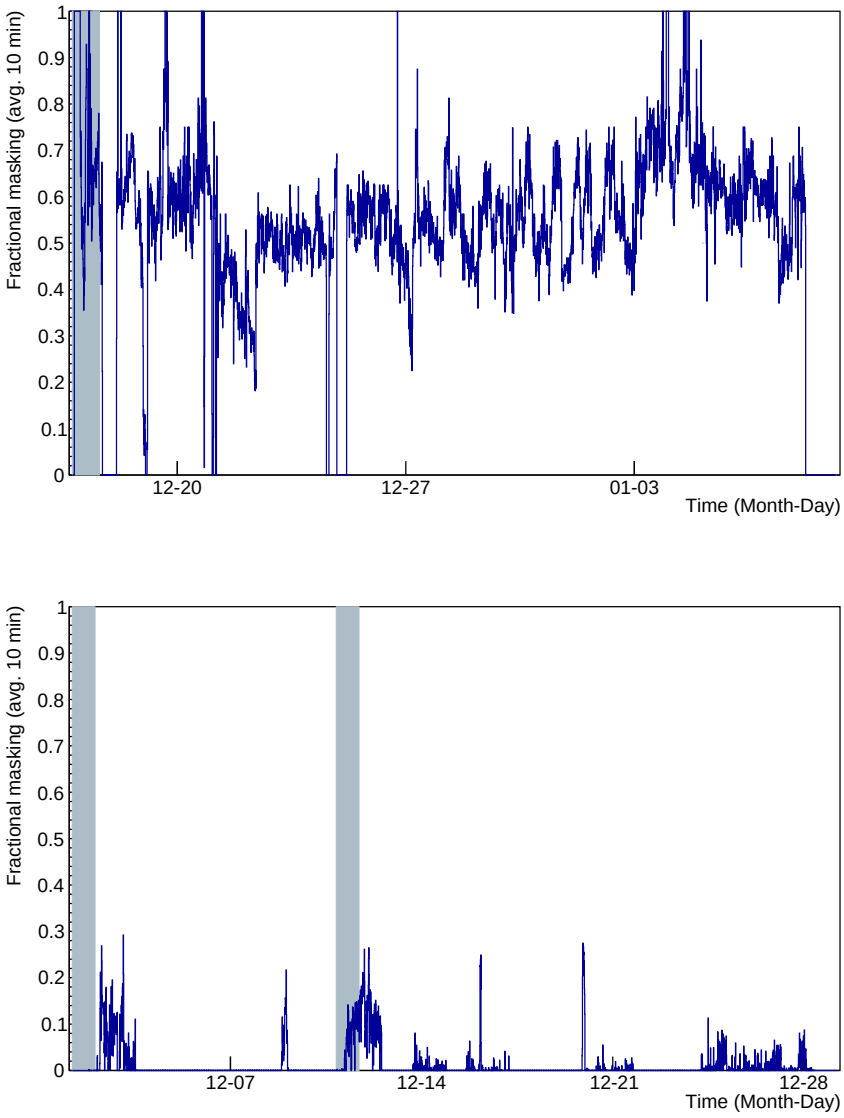

Figure 4: Masking in the ANITA-III (top) and ANITA-IV (bottom) flights. Before ANITA-IV, masking was the primary method of avoiding digitization deadtime due to CW interference. For the majority of the ANITA-III flight, over half of the payload was masked. Due to the mitigation of CW noise in ANITA-IV to acceptable levels by the TUFF notch filters, the need for masking was strikingly reduced. The shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility.

4. TUFF Board Design

For ANITA-IV, we built and deployed 16 TUFF boards (not counting spares) with six channels each for the 96 total full-band RF channels of ANITA. Figure 2 shows, for a single RF channel in ANITA-IV, where the TUFF boards are in the signal processing chain. The main components of each TUFF channel comprise two amplifiers, three notch filters, a microcontroller and a bias tee, as highlighted in Figure 5. In the signal processing chain, the notch filters were included in both the trigger and signal paths for simplicity, to conserve some dynamic range of the digitizer when interference was present, and to ensure that the filters come immediately after a directive element (the final amplifier).

The design of the TUFF board was affected by the low power budget of ANITA as well as the weight and size restrictions of a balloon mission, as described in Section 2. The TUFF boards needed to be low-power and light, and compact to fit into the existing amplifier housing locations along with necessary cabling to match to the existing connectors.

Figure 5 shows a single TUFF channel, each of which is approximately 56 mm in width. Each printed circuit board has four layers of copper with an FR-4 dielectric material. The TUFF boards operate on 3.3 V and 4.7 V power sources provided by a MIC5504 from Microchip Technologies Inc. and a ADM7171 from Analog Devices Inc. Both voltage regulators draw from a 5 V source supplied by the DC/DC unit in the ANITA Instrument Box. A single TUFF channel consumes only 330 mW of power. The total power consumed by the ANITA payload is approximately 800 W.

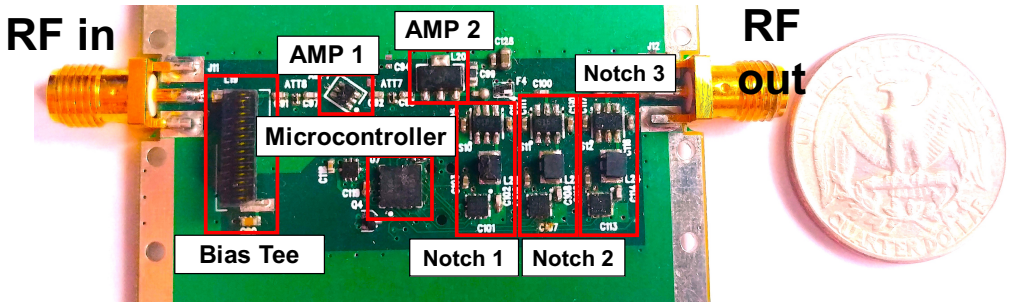

Figure 5: A single TUFF board unit (channel) that powers the first-stage antenna-mounted amplification unit and performs second-stage amplification and notch filtering of a single RF channel (out of 96 total). Each TUFF board has six such channels. The main components of the channel are highlighted here.

Two TUFF boards were assembled into a final 12-channel aluminum housing. This provides heat-sinking, structural support, and RF isolation. Two of these 12-channel modules were placed inside an Internal Radio Frequency Conditioning Module (IRFCM) inside the Instrument Box of ANITA. Figure 6 shows the inside of an IRFCM. The main components of a TUFF channel are described below.

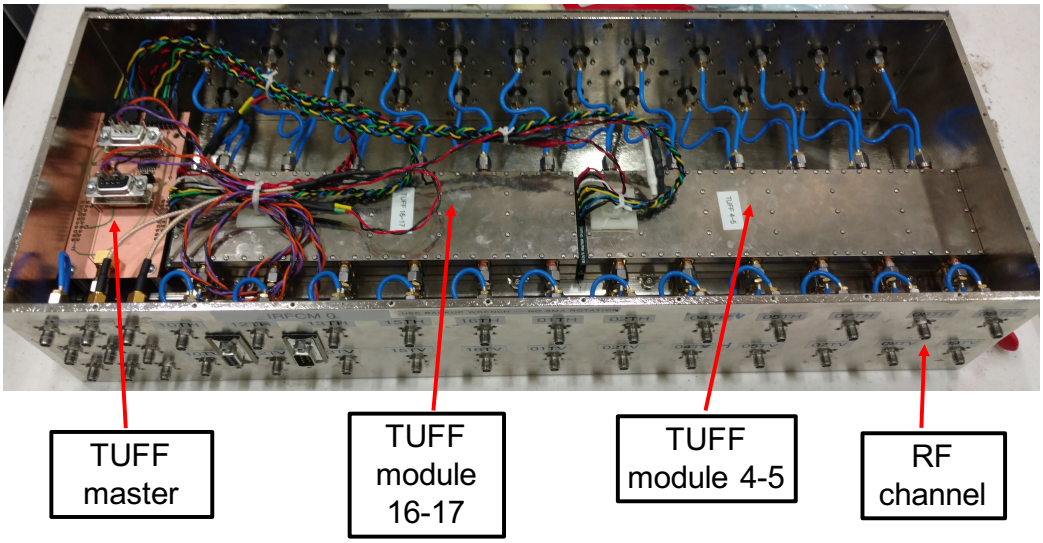

Figure 6: Internal Radio Frequency Conditioning Module (IRFCM) containing two 12 channel TUFF modules serving 24 RF channels total, together with a TUFF Master for sending commands to the TUFF boards from the flight computer.

4.1. Amplifiers and bias tee

There are two amplifiers connected in series that together produce second-stage RF power amplification of approximately 45 dB. The gain of a TUFF channel, as measured in the lab, is shown in Figure 7. In Figure 5, AMP 1 is a BGA2851 by NXP Semiconductors and AMP 2 is an ADL5545 by Analog Devices. There is an attenuator producing 1 dB of attenuation to the RF signal as it leaves AMP 1 and before it enters AMP 2. The BGA2851 provides a gain of 24.8 dB at 950 MHz. It has a noise figure of 3.2 dB at 950 MHz. It consumes 7 mA of current at a supply voltage of 5 V, or 35 mW of power. The ADL5545 provides a gain of 24.1 dB with broadband operation from $30 - 6000$ MHz. Out-of-band power at frequencies above 2 GHz is suppressed by a filter on each TUFF channel. Additionally, there are band-pass filters

immediately after the TUFF boards in the signal processing chain allowing power only in the frequency range $200 - 1200\,\mathrm{MHz}$. The ADL5545 has a noise figure of 2.9 dB at 900 MHz and a 1 dB compression point (P1dB) of 18.1 dBm at 900 MHz. It consumes 56 mA of current at a supply voltage of 5 V, or 300 mW of power. Thus, this amplifier consumes the majority of the power required by a single TUFF channel. The table below summarizes properties of the amplifiers.

Amplifier	Part name	Gain	Power consumed	Noise figure
AMP 1	BGA2851	24.8 dB	35 mW	3.2 dB
AMP 2	ADL5545	24.1 dB	300 mW	2.9 dB

There is a bias tee on each TUFF channel that remotely powers the AMPA (antenna-mounted pre-amplifier) unit at the other end of the coaxial cable connecting an AMPA and that channel. It consists of a 4310LC inductor by Coilcraft in series with a $0.1\,\mu$F capacitor. The inductor delivers DC to the AMPA unit while the capacitor prevents DC from passing through to the signal path of the TUFF channel. The bias tee allows RF signal traveling from the AMPA unit through the coaxial cable to pass through to the rest of the signal path of the TUFF channel.

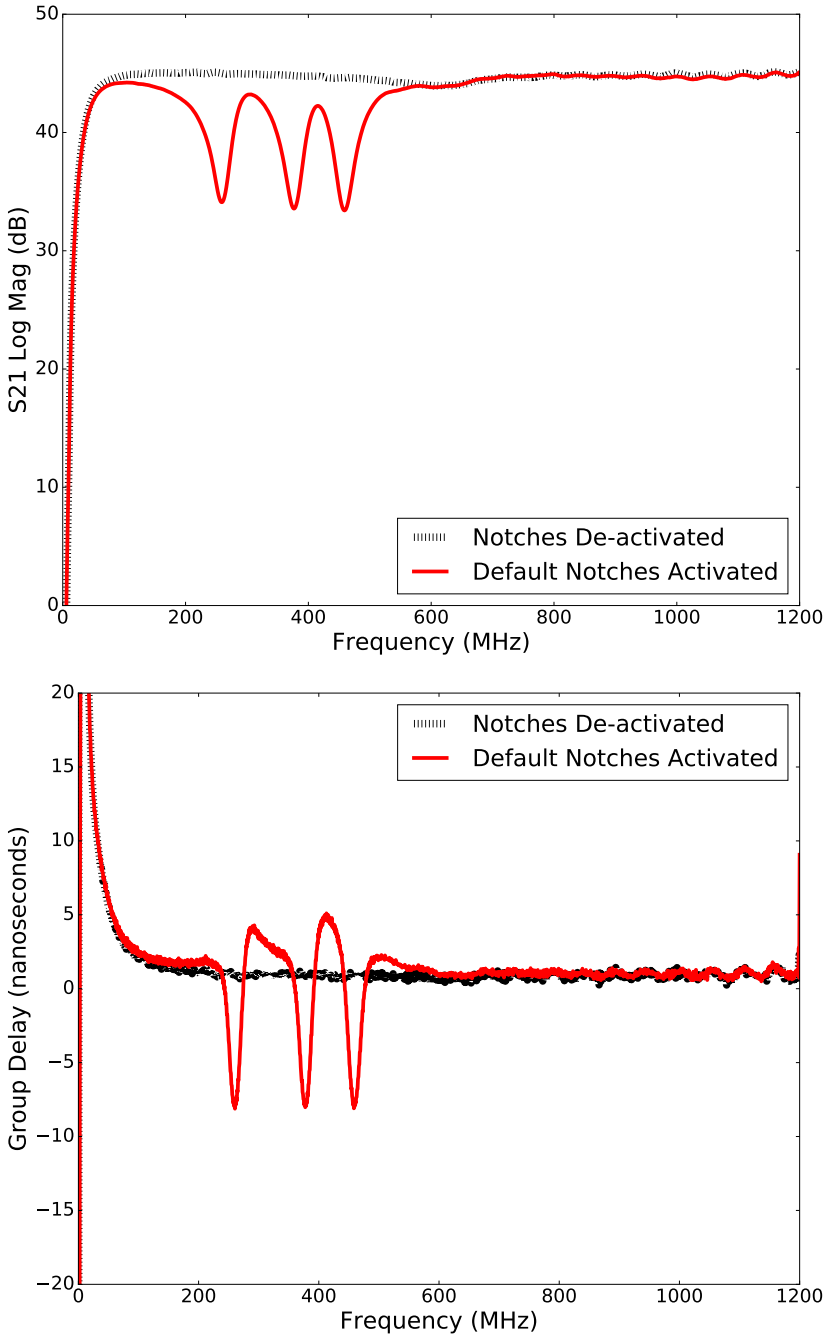

Figure 7: The forward transmission coefficient, S21, or the gain (top) and group delay (bottom) of a TUFF channel as measured in the lab with all notches de-activated (black dashed line) and all notches activated at their default frequencies (red solid line). There is approximately 13 dB of gain attenuation in the notched regions.

4.2. Notch filters

There are three tunable, switchable notch filters for mitigation of CW noise at the default frequencies of 260 MHz (Notch 1), 375 MHz (Notch 2) and 460 MHz (Notch 3). CW noise at the first two frequencies are thought to be caused by military communications satellites, specifically, the FLTSAT and UFO systems and the MUOS system, respectively. The third notch filter is present to curb CW interference seen when the ANITA payload is near Antarctic science bases such as McMurdo and South Pole Station.

The gain and group delay of a TUFF channel, with all notch filters activated as well as de-activated, are shown in Figure 7. The TUFF notches were able to achieve a maximum attenuation of approximately 13 dB, and were implemented as a simple RLC trap (Carr *et al.* [6]). The added group delay in the regions between the notches was below the effective integration time of the SHORT, and so should only have a minor effect on the contribution to the trigger for those frequencies.

In each notch, the resistance R originated from the parasitic on-resistance of a dual-pole, single-throw RF switch and the DC resistance of the remaining components. This is approximately $6 - 7\,\Omega$. The inductance L is fixed at 56 nH. The capacitance C is a combination of a fixed capacitor and a PE64906 variable capacitor from Peregrine Semiconductor. Simulations using the device model of the variable capacitor suggest that the mounting pads of the components contribute $\sim$ 0.6 pF of parasitic capacitance.

With the tuning capability of the variable capacitor, the resonant frequency of the RLC circuit was modified during flight to dynamically mitigate CW interference. The variable capacitor in a notch can be tuned in 32 discrete steps of 119 fF in the range $0.9 - 4.6$ pF and for each notch, is connected in series or parallel with a constant capacitance. For Notch 1, the variable capacitor is in parallel with a 1.8 pF capacitor. For Notches 2 and 3, the variable capacitor is in series with a 12.0 pF (Notch 2) and a 1.5 pF (Notch 3) capacitor for increased tuning capability. Figure 8 shows a simplified circuit diagram.

Previous tunable notch designs were typically implemented as multiple-pole filters (e.g. Brank *et al.* [7]), having both parallel and series connected components. Adding the capability to de-activate these notches would have required multiple switches per notch, a significant increase in circuit size, and resulted in the switches being present in the signal path even when the notch was de-activated (Wong *et al.* [8]). An alternative approach would have been a coupled transmission line (Wu *et al.* [9]), however coupled inductors

over these frequencies are both large and low-performance. The simple notch structure used here, while limited in rejection and bandwidth, results in an extremely compact filter bank and less than 0.1 dB insertion loss when the notch is de-activated.

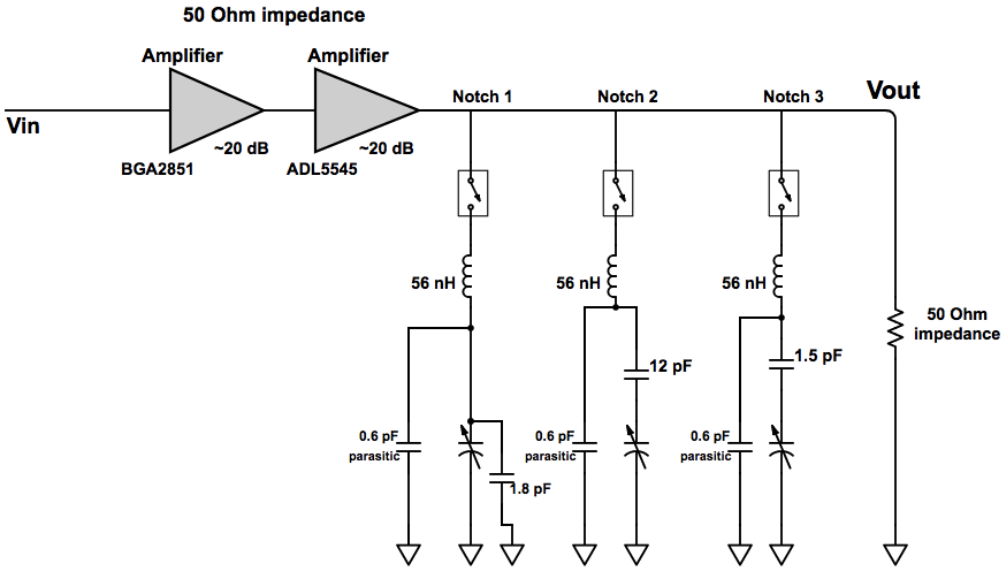

Figure 8: Circuit diagram showing the different components of the TUFF notch filters.

4.3. Microcontroller

We use an ultra-low-power microcontroller, specifically a MSP430G2102 by Texas Instruments. This features a powerful 16-bit Reduced Instruction Set Computing (RISC) central processing unit (CPU). There are five low-power modes optimized for extended battery life. The active mode consumes 220 μA at 1 MHz and 2.2 V. The standby mode consumes only 0.5 μA and the RAM retention-off mode consumes 0.1 μA. The digitally-controlled oscillator allows wake-up from low-power modes to active mode in less than 1 μs.

During the ANITA-IV flight, commands could be sent using the SIP connection to set the state of the variable capacitor of each TUFF notch filter via the microcontroller of that channel. This was done in real time if a re-tune of a notch filter was necessary to mitigate CW interference. Commands could be sent to de-activate or activate a notch filter using the switch associated with each notch. Each microcontroller has the capability to communicate over universal serial communication interface.

5. TUFF notch filter operations during the ANITA-IV flight

Deployed for the first time in ANITA-IV, the TUFF boards were heavily used throughout the flight. Figure 9 summarizes the activation status of each notch as a function of time during the flight.

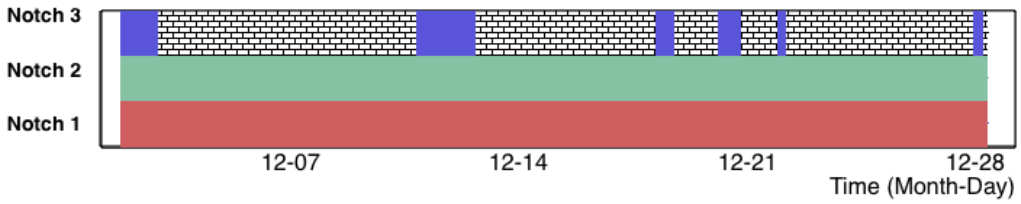

Figure 9: The activated (solid red for Notch 1, solid green for Notch 2, solid blue for Notch 3) or de-activated (hatched) status for each TUFF notch filter during the flight.

Notch 1: 260 MHz During the ANITA-III flight, a CW signal at 260 MHz from military satellite systems (CW peak seen in Figure 10) was present throughout the flight. This CW signal was omnipresent during the ANITA-IV flight as well, and so Notch 1 needed to be active throughout the flight. Notch 1 (usually centered at 260 MHz) was re-tuned on Dec 14 as we saw CW interference at 250 MHz and was tuned back to 260 MHz later that day.

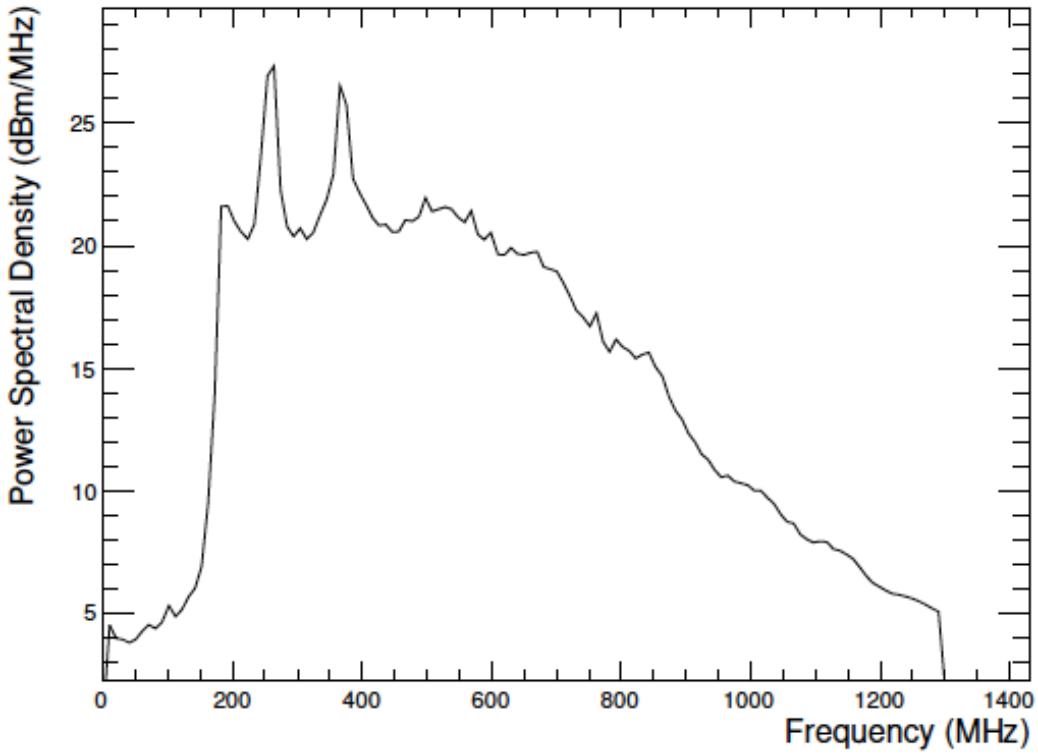

Figure 10: A plot of the average power spectral density over 1 min for one channel from when the ANITA-III payload was near WAIS Divide in Antarctica. The two peaks at 260 MHz and 375 MHz, presumably from military satellites, are visible here. The 260 MHz peak was present throughout the flight and the 375 MHz peak was present during less than half of the flight. These CW peaks motivated the installation of the TUFF notch filters in ANITA-IV. As it turns out, Notch 1 (to curb the left peak) and Notch 2 (to curb the right peak) both needed to be active for essentially the entire flight in ANITA-IV.

Notch 2: 360 – 390 MHz During the ANITA-III flight, a second CW peak at 375 MHz from military satellite systems (CW peak seen in Figure 10) was sometimes present. The MUOS-1 and MUOS-2 satellites are suspected to have caused the second CW peak in ANITA-III. This peak is always present during the ANITA-IV flight. The enhanced second peak in ANITA-IV is likely due to the presence of three additional MUOS satellites, that is, MUOS-3, MUOS-4 and MUOS-5, in orbit during the ANITA-IV flight. During the ANITA-IV flight, Notch 2, although de-activated twice (Dec 2, Dec 19), needed to be activated again within minutes due to this CW noise. This is illustrated in Figure 11 where we show averaged spectra over all

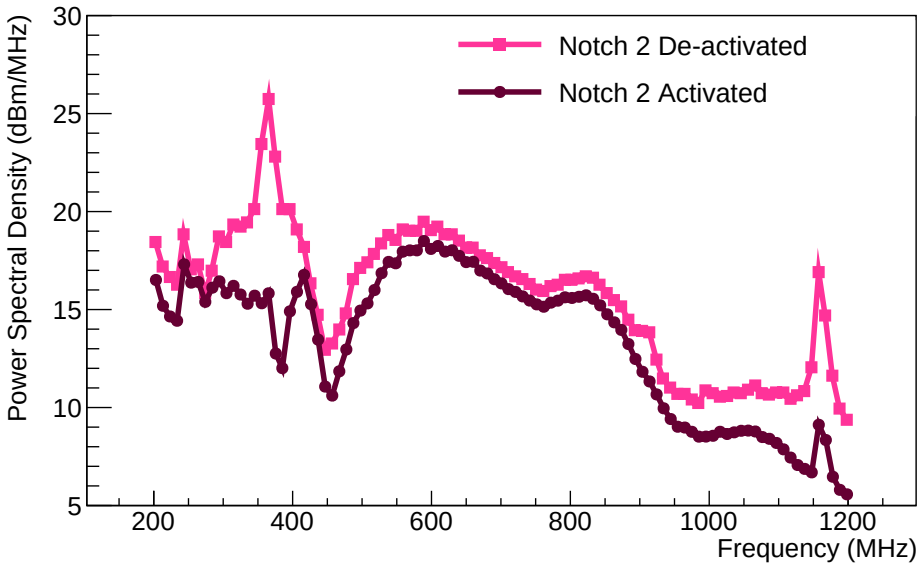

Figure 11: Power spectra with Notch 2 de-activated and Notch 2 activated (spectra averaged over 16.5 minutes) during the ANITA-IV flight. Notch 2 was de-activated on Dec 2 for 16 minutes resulting in a CW peak seen in the spectra. Notch 2 was then activated again, and the CW peak was curbed. Although we show only phi sector 16 here, excess CW noise upon de-activating Notch 2 and the effect of activating Notch 2 again was seen in almost all phi sectors. Note that the lower scale when Notch 2 is activated is not due to signal loss, but rather due to ANITA no longer triggering on the interference.

waveforms from a phi sector with the notch de-activated and then activated (Dec 2), and in Figure 12 where the trigger rate is shown to be nearly double when Notch 2 was de-activated (Dec 19). Notch 2 was re-tuned during flight a few times (Dec 6-8) to dynamically combat CW interference in the range of $360 - 390\,\mathrm{MHz}$. Figure 13 shows the effect of real time tuning of Notch 2 on Dec 7 for mitigation of CW interference at $390\,\mathrm{MHz}$. Tuning the notch brought the CW noise power down.

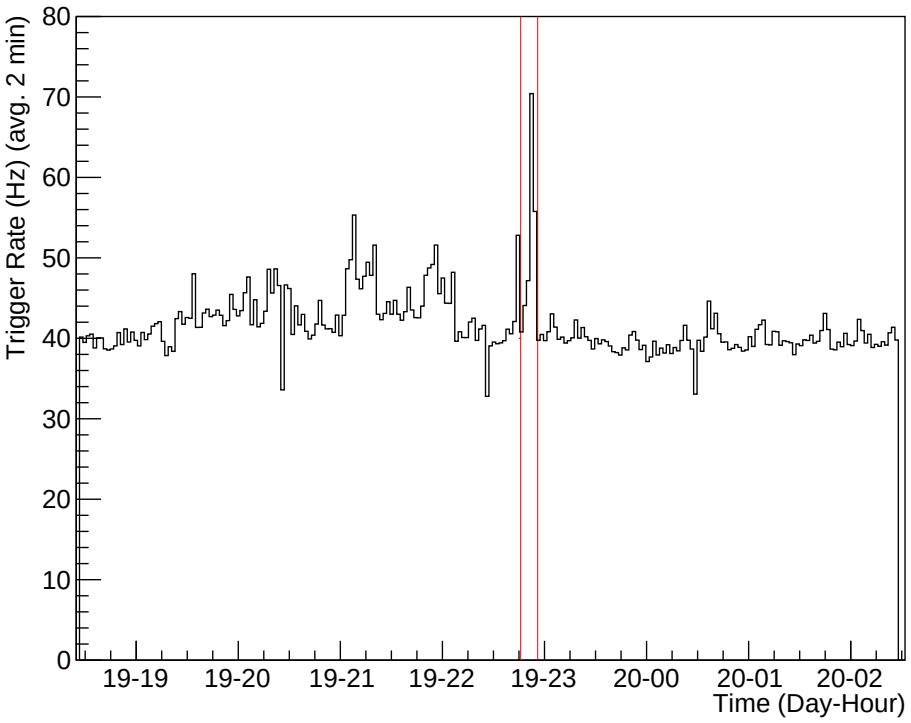

Figure 12: On Dec 19 at approximately 10:46 PM, Notch 2 was de-activated for approximately 10 minutes. The vertical red lines enclose the duration of time during which Notch 2 was de-activated. A trigger rate above $\sim$ 50 Hz incurs digitization deadtime. The spike in event rate shows that Notch 2 was crucial to keeping CW interference in check. Even with the LCP/RCP coincidence required by the ANITA-IV trigger, further mitigation of CW interference by the TUFF boards was necessary to avoid masking.

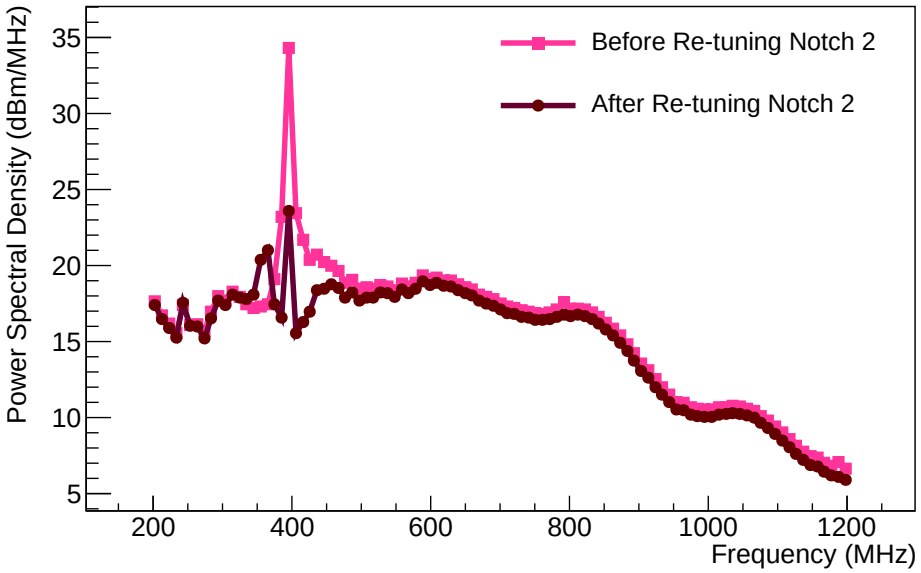

Figure 13: Power spectra (averaged over 2 hours) from before and after dynamically tuning Notch 2 during the ANITA-IV flight. On observing a large CW peak at 390 MHz on Dec 7, Notch 2 was re-tuned. Although we show only phi sector 8 here, similar CW peaks and effects of notch tuning were seen in all phi sectors.

Notch 3: 460 MHz Notch 3 was generally activated when the payload was in view of Antarctic bases and filtered the $450 - 460$ MHz frequency region. Notch 3 was de-activated on Dec 2 for a few minutes but had to be activated again as the payload was close to McMurdo Station at the time.

6. Performance of ANITA-IV compared to ANITA-III

During the ANITA-IV flight, we utilized all features of the TUFF notch filters to achieve decreased masking, increased stability of trigger rate and SURF DAC thresholds, and increased instrument livetime, from 31.6% in ANITA-III to 91.3% in ANITA-IV. These results are summarized in Figures 3, 4, 14 and 15.

6.1. Livetime in ANITA

Increasing livetime was the primary motivation behind building and deploying the TUFF boards in ANITA-IV. There are two types of livetime in ANITA, which are described below.

Digitization livetime In ANITA, deadtime due to digitization by all four LABRADOR chips of the SURF board is recorded by the TURF board, as illustrated in Figure 2. This deadtime is recorded as a fraction of a second. Digitization livetime per second can be obtained by subtracting this from one. Increasing the digitization livetime increases the probability of receiving RF signal due to an UHE neutrino.

Instrument livetime At any given time, the digitization livetime multiplied by the fraction of unmasked phi sectors (after accounting for channel-masking) gives us the instrument livetime per second. In other words, instrument livetime accounts for the fraction of observable ice in azimuth after accounting for masking.

6.2. Methods adopted to reduce digitization deadtime

Masking Before ANITA-IV, the primary method of reducing digitization deadtime due to CW signal was masking, which includes both phi-masking and channel-masking. However, masking leads to instrument deadtime as parts of the payload become unavailable for neutrino detection. Due to the TUFF boards, fractional masking below 0.3 was maintained during the ANITA-IV flight, as seen in Figure 4.

Changing SURF DAC thresholds In addition to masking, adjusting the SURF DAC thresholds is also a method of reducing digitization deadtime. The distribution of SURF DAC thresholds for the ANITA-III and ANITA-IV flights is shown in Figure 3. It is evident that the method of changing thresholds to minimize digitization deadtime was heavily adopted during the ANITA-III flight. As the ANITA-III payload was continuously exposed to CW interference, it was unable to maintain stable SURF DAC thresholds. As the TUFF boards mitigated CW interference to acceptable levels in ANITA-IV, the thresholds are kept nearly constant during this flight.

6.3. Livetime in ANITA-IV compared to ANITA-III

The total digitization livetime for the ANITA-III and ANITA-IV flights was calculated to be 73.7% and 92.3% respectively. The distribution of digitization livetime per second as a function of time is shown for ANITA-III and ANITA-IV in Figure 14. The TUFF boards dynamically notch-filtered CW peaks in the power spectrum of a received signal at an early stage of signal

processing. This brought the rate of triggers due to CW signal to acceptable levels and thereby increased digitization livetime.

Most importantly, the TUFF boards helped to increase the instrument livetime (digitization livetime weighted by the fraction of unmasked phi sectors) in the ANITA-IV flight, mainly by decreasing the need for masking. The distribution for instrument livetime per second as a function of time is shown for ANITA-III and ANITA-IV in Figure 15. The total instrument livetime for ANITA-III and ANITA-IV was calculated to be 31.6% and 91.3% respectively. On average, instrument livetime in ANITA-IV was 2.8 times higher than that in ANITA-III.

6.4. Impact on signal power and acceptance

A full account of the impact of the TUFF notch filters on neutrino sensitivity is under investigation and beyond the scope of this paper. We note that each notch removes approximately 5% of the system bandwidth ($200 - 1200\,\mathrm{MHz}$). Although the impact of increased digitization livetime is straightforward to estimate, the increase in sensitivity due to the reduction in masking will require a full account of the time- and azimuthal-dependent exposure of ANITA to neutrinos.

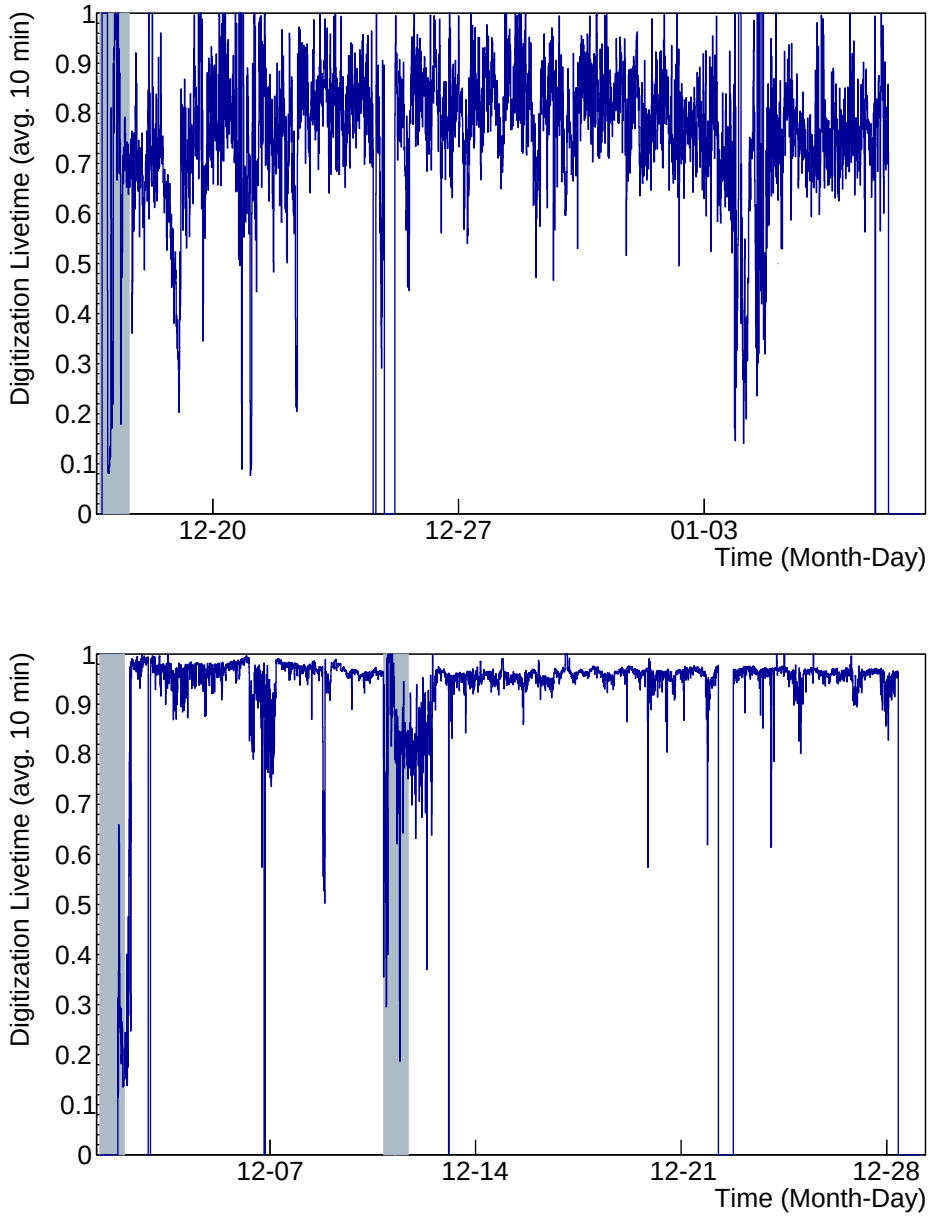

Figure 14: Digitization livetime per second for ANITA-III (top) and ANITA-IV (bottom). As the ANITA-III payload was inundated by CW interference, digitization livetime was reduced. In ANITA-IV, the TUFF boards helped to reduce triggers due to CW signal and therefore, increased the digitization livetime. The total digitization livetime for the ANITA-III and ANITA-IV flights was calculated to be 73.7% and 92.3% respectively. The shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility.

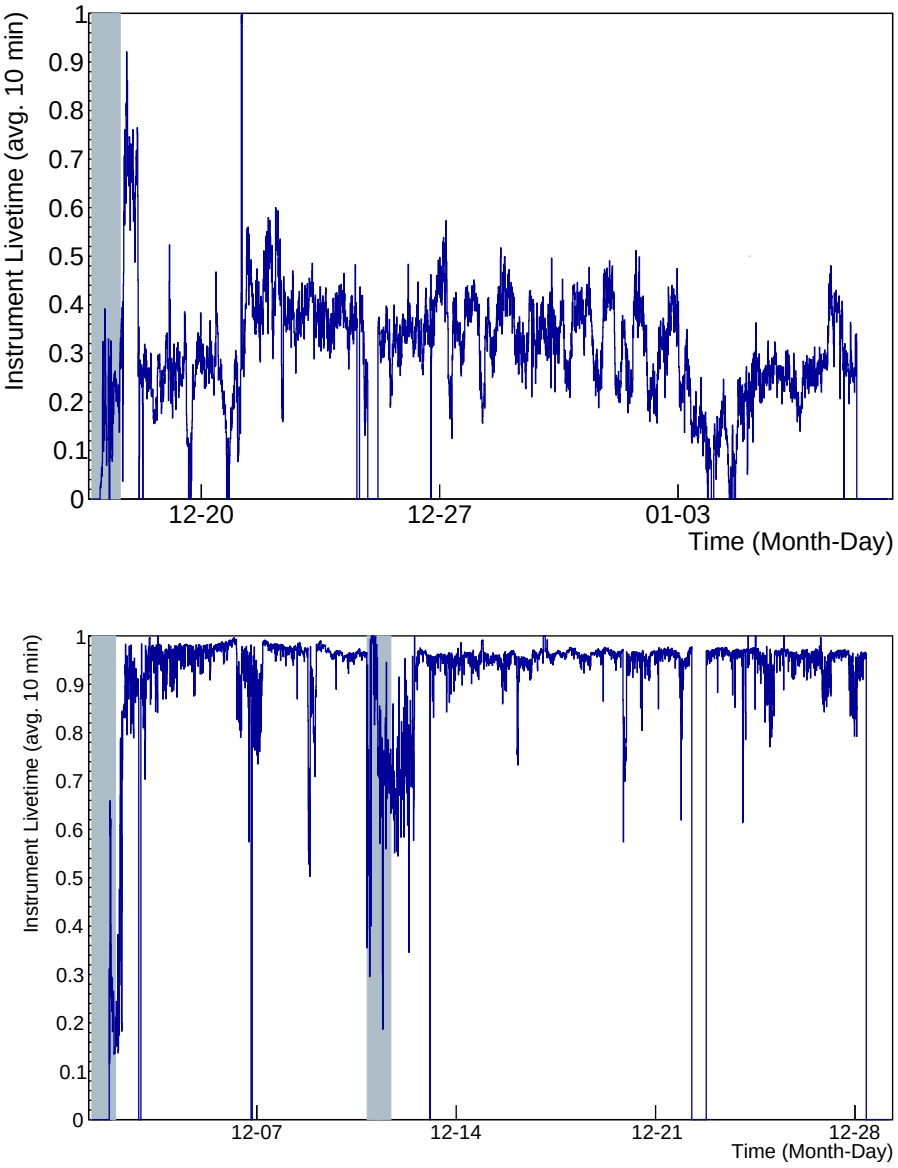

Figure 15: Instrument livetime per second, obtained by weighting digitization livetime by the fraction of unmasked phi sectors, for ANITA-III (top) and ANITA-IV (bottom). In ANITA-III, masking had to be implemented heavily and throughout the flight, which led to a dramatic reduction of instrument livetime. The TUFF boards largely removed the need for masking in ANITA-IV. This helped to increase the instrument livetime of ANITA, with 91.3% total instrument livetime in ANITA-IV, compared to 31.6% in ANITA-III. The shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility.

7. Future plans

A fifth ANITA flight (ANITA-V) is currently proposed. For this flight, we will explore the option of developing software to dynamically activate the notch filters only when satellites come into view of an antenna in order to eliminate any sensitivity loss. Development of new triggering and digitization systems for ANITA-V is currently underway. The main proposed upgrade to the triggering system include the Realtime Independent Three-bit Converter (RITC) as described by Nishimura *et al.* [10]. This triggering system, conceived for ANITA-III and intended for ANITA-IV, was held back until CW mitigation, as shown here, could be demonstrated. The RITC will perform continuous, low-resolution digitization in order to carry out interferometry of all incoming data in realtime. This will be used to generate a system trigger. Once triggered, high-resolution digitization of the data will be performed by new SURF boards.

Proposed upgrades to the SURF board include new LABRADOR chips (LAB4D), as described by Roberts *et al.* [11]. There will be 12 LABRADOR chips per SURF board. Each LABRADOR chip will sample data from one RF channel using 32 blocks of 128-element SCAs. The SCAs will sample waveform data at 3.2 GSa/s, eight blocks at a time (forming four buffers per LABRADOR chip). When a Level 3 trigger is issued, sampling will be frozen for the 8 blocks of SCAs to digitize data, while the remaining 24 blocks continue to sample.

8. Acknowledgments

We are grateful to NASA for their support for ANITA through Grant NNX15AC20G. We thank the U.S. National Science Foundation-Office of Polar Programs. A. Connolly would like to thank the National Science Foundation for their support through CAREER award 1255557. This work was also supported by collaborative visits funded by the Cosmology and Astroparticle Student and Postdoc Exchange Network (CASPEN). Lastly, we thank Brian Clark, Ian Best and Suren Gourapura for their valuable feedback.

9. References

References

[1] 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 (2009) 10–41, doi:10.1016/j.astropartphys.2009.05.003.

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

[3] P. W. Gorham, et al., Observations of the Askaryan effect in ice, Phys. Rev. Lett. 99 (2007) 171101, doi:10.1103/PhysRevLett.99.171101.

[4] C. K. Matassa, Comparing the capabilities and performance of the ultra high frequency follow-on system with the mobile user objective system, Master's thesis, Naval Postgraduate School, URL `https://calhoun.nps.edu/handle/10945/5711`, 2011.

[5] G. S. Varner, L. L. Ruckman, P. W. Gorham, J. W. Nam, R. J. Nichol, J. Cao, M. Wilcox, The large analog bandwidth recorder and digitizer with ordered readout (LABRADOR) ASIC, Nucl. Instrum. Meth. A583 (2007) 447–460, doi:10.1016/j.nima.2007.09.013.

[6] J. Carr, The Technician's EMI Handbook: Clues and Solutions, Newnes, 1 edn., ISBN 0750672331, 2000.

[7] J. Brank, RF MEMS-based tunable filters, Int. J RF Microwave CAE 11 (5) (2001) 276–284.

[8] P. W. Wong, I. Hunter, Electronically Tunable Filters, IEEE Microwave Magazine 10 (6) (2009) 46–54.

[9] Z. Wu, Y. Shim, M. Rais-Zadeh, Switchable wide tuning range bandstop filters for frequency-agile radios, IEEE Xplore .

[10] K. Nishimura, M. Andrew, Z. Cao, M. Cooney, P. Gorham, L. Macchiarulo, L. Ritter, A. Romero-Wolf, G. Varner, A low-resolution, GSa/s streaming digitizer for a correlation-based trigger system, ArXiv e-prints 1203.4178 .

[11] J. Roberts, E. Oberla, P. Allison, G. Varner, S. Spack, B. Fox, B. Rotter, LAB4D: A Low Power, Multi-GSa/s, Transient Digitizer with Sampling Timebase Trimming Capabilities, Nucl. Instrum. Meth. 1 (2017) 12.

On behalf of the ANITA Collaboration I would like to thank the reviewers for their thoughtful comments, which improved the clarity of the paper.

Reviewer #1 -------

The article gives a nice overview of the improvements made in ANITA-IV, in particular the tunable notch filters that helped increase the instrument livetime very significantly. It is well written and clearly presented. I recommend it be accepted for publication after a few minor revisions as indicated below.

Thank you so much!

The paper is sometimes a bit redundant in repeating how well the notch filters worked. It would be good to streamline this a little bit - the message will still be very clear.

Thank you for pointing this out.

We made some changes in Section 5 as follows:

We edited this part “The TUFF boards were deployed for the first time in ANITA-IV and proved to be critical to the success of the mission. The TUFF notch filters were heavily used throughout the flight. Figure 9 summarizes the status of each notch as a function of time during the flight.”

→ “Deployed for the first time in ANITA-IV, the TUFF boards were heavily used throughout the flight. Figure 9 summarizes the activation status of each notch as a function of time during the flight.”

We edited this part “After being de-activated on Dec 2 for ~ 16minutes, Notch 2 needed to be activated again. Excess CW noise upon de-activating Notch 2 was seen in almost all phi sectors. In Figure 11 we show spectra averaged over all waveforms from one phi sector during this period. Notch 2 was de-activated again on Dec 19 for approximately 10 minutes. The trigger rate was nearly doubled almost as soon as the notch was de-activated (see Figure 12) and excess CW noise was seen in several phi sectors.”

→ “This is illustrated in Figure 11 where we show averaged spectra over all waveforms from a phi sector with the notch de-activated and then activated

(Dec~2), and in Figure 12 where the trigger rate is shown to be nearly double when Notch 2 was de-activated (Dec~19).”

The trigger path indicated in figure 2 between the splitter and the SURF board could be made more explicit in indicating the SHORT, as this is heavily referred to in the text.

Thank you, we changed a label in Figure 2 to say “SHORT”.

The second half of page 9 remains a bit unclear to me. Is there a difference between LABRADOR and LAB? How many LAB chips are there? On the one hand the text states "12 SURFS with 4 LAB chips each", which would add up to 48. A few lines further down it sounds as if there are 4 LAB chips per each of the 96 channels. I suggest the authors rephrase this part to clarify.

Thanks. There is no difference between LAB and LABRADOR.

There are four LAB chips per SURF board, so a total of 48 LAB chips.

In the text, we changed LAB to LABRADOR. We say LAB for short, but LABRADOR is more clear.

We changed subsection “Digitization:” to read “There are twelve SURF boards, each containing four custom-built Application Specific Integrated Circuits called Large Analog Bandwidth Recorder And Digitizer with Ordered Readout (LABRADOR).”

We changed Sec 3.2 (LAB chip and…) to say “Each LABRADOR chip has a 260-element switched capacitor array (SCA) for each of its 9 input channels, with one channel used for timing synchronization.”

On page 12 the statement that the TUFF board unit is about "twice the size of the coin" does not really fit the photo, the board seems bigger than that.

Thanks. We changed Sec. 4 to say “Fig. 5 shows a single TUFF channel, each of which is approximately 56 mm in width.” (we deleted the ‘It can be seen’ sentence)

Can the authors include a few words on how the notched parts of the

frequency band are handled in analysis? Is there an end-to-end record of when which notch filter was tuned to which frequency? Does it matter that the effective bandwidth of the experiment changes with time?

Thanks.

Yes, there is a record of all notch configurations during the flight. There is a model associated with each notch configuration. These are included in both the simulation and analysis software. However, a discussion on how the notch filtering affects analysis is beyond the scope of this paper. There is a note in Sec 6.4 addressing this as follows:

“A full account of the impact of the TUFF notch filters on neutrino sensitivity is under investigation and beyond the scope of this paper. We note that each notch removes approximately 5% of the system bandwidth (200 − 1200 MHz). Although the impact of increased digitization livetime is straightforward to estimate, the increase in sensitivity due to the reduction in masking will require a full account of the time- and azimuthal-dependent exposure of ANITA to neutrinos.”

In Figure 7 the transmission coefficients are shown. What about the group delay? Is there any relevant dispersion introduced by the notch filters, which would decrease the signal-to-noise ratio of radio pulses? It would be good to include some information on this.

Thank you.

We added a plot (Figure 7 bottom plot) showing group delay, and a brief mention that the added group delay (dispersion) is small compared to the timescale of the trigger integration.

We changed in Sec 4.2:

“The gain and group delay of a TUFF channel, with all notch filters activated as well as de-activated, are shown in Figure 7. The TUFF notches were able to achieve a maximum attenuation of approximately 13 dB, and were implemented as a simple RLC trap (Carr et al. [6]). The added group delay in the regions between the notches was below the effective integration time of the SHORT, and so should only have a minor effect on the contribution to the trigger for those frequencies.”

A more general question: One could have considered to only notch the signal split off for triggering and keep the signal sampled by the LAB chips in its original form. Why was such an approach not adopted? A short discussion in the paper would be helpful.

Thank you.

The notch filters were included in both the trigger and signal paths for simplicity, to conserve some dynamic range of the digitizer when interference was present, and to ensure that the filters come immediately after a directive element (the final amplifier).

We added in Sec. 4, first paragraph, at the end. “The main components of each TUFF channel comprise two amplifiers, three notch filters, a microcontroller and a bias tee, as highlighted in Figure 5. In the signal processing chain, the notch filters were included in both the trigger and signal paths for simplicity, to conserve some dynamic range of the digitizer when interference was present, and to ensure that the filters come immediately after a directive element (the final amplifier).”

In figure 11, why is the notched signal so different (lower power) than the non-notched signal. Were these data not recorded at approximately the same time? If so, could the authors make this plot with data that were recorded at more similar times to allow a better comparison? Also, what is the peak around 1200 MHz and why does this get suppressed as well when the notch filter at 260 MHz is switched on?

Thank you.

The power spectra are computed from triggered data – when the interference was removed, the average power observed in triggered data is lower because the system no longer triggers on the interference. This is now clarified in the caption of figure 11:

“Note that the lower scale when Notch 2 is activated is not due to signal loss, but rather due to ANITA no longer triggering on the interference.”

The peak at ~1200 MHz is due to a different anthropogenic source (likely Iridium satellite) and the magnitude of the peak varied throughout the flight. It happens to be smaller in the case of Notch 2 being activated but not related.

In section 6.3 and 6.4 it might be worth mentioning that in addition to the greatly increased instrument livetime, also the detection thresholds were on average higher (more sensitive). If there is an easy way to quantify this (e.g., average threshold level), that could be useful.

Thank you. The primary point of Fig. 3 is to show the increased stability of the thresholds over time - while it is true the higher level indicates a lower threshold, the trigger scheme for ANITA-4 requires a higher number of coincident triggers (due to the L/R coincidence), meaning that the lower per-channel threshold would not necessarily translate into a lower overall trigger threshold. A direct comparison of the absolute levels is therefore quite complicated, and we did not mention it in the text for this reason. However, we added a brief note in the text to indicate this.

lines 142-143 (at the end): Note that the lower overall threshold for ANITA-IV is primarily due to the modified triggering scheme, which requires more overall coincidences between channels. The increased stability of the ANITA-IV thresholds, due to the CW mitigation schemes presented here, is clearly apparent.

Reviewer #2: ——

This is interesting hardware work from the ANITA collaboration and has the potential to be useful to other radio neutrino experiments, and possibly to other radio frequency efforts outside of the particle astrophysics community. I strongly endorse the publication of this paper in NIM in a timely manner. I have a few concerns with the references and literature search on these techniques, and some questions of the EMI issues during the ANITA balloon flights that necessitated this solution described in the paper. I think addressing these two issues would enhance the impact of the paper, for the former, it could increase and broaden possible readership, and for the later, would strengthen the background of the paper. There are a few minor comments and questions as well that are below as well.

Thank you!

References/literature:
All of the references are to either the historical Askaryan Effect or internal to ANITA except for a mention of a thesis involving the MUOS satellite systems. Have folks not ever made programmable, tunable RF filters before? Whether custom, or using off-the-shelf parts? A quick search showed dozens of references and resources, including parts which replicate much of this built system as a purchasable item. I'd think some of this was worth referencing and acknowledging. Possibly to answer questions such as: Is this a design from first principles? Is this design based on previous work? (Including the data sheets and white papers from Mini Circuits.) Is this design applicable to other projects and other readers of the paper? How is this superior to other approaches?

Thank you.

This design was essentially from first principles, in that it's a common hobby technique for a simple notch in TV tuning (a parallel trap). It's not commonly used outside of hobby applications due to its limitations (wide bandwidth and limited rejection), which aren't problems in this application.

We changed Sec. 4.2:

"were implemented as a simple RLC trap [ref 1]" and added

Previous tunable notch designs were typically implemented as multiple-pole filters (cite ref 2), having both parallel- and series-connected components. Adding the capability to deactivate these notches would have required multiple switches per notch, a significant increase in circuit size, and resulted in the switches being present in the signal path even when the notch was deactivated (cite ref 3). An alternative approach would have been a coupled transmission line (cite ref 4), however coupled inductors over these frequencies are both large and low-performance. The simple notch structure used here, while limited in rejection and bandwidth, results in an extremely compact filter bank and less than 0.1 dB insertion loss when the notch is deactivated.

Ref 1: J. Carr, "The Technician's EMI Handbook: Clues and Solutions", 2000.

Ref 2: Brank et al., Int. J RF Microwave CAE 11: 276-284, 2011.

Ref 3: P. Wong, I. Hunter. IEEE Micro. Mag., vol 10, no. 6, 46-54, 2009.

Ref 4: Z. Wu, Y. Shim, M. Rais-Zadeh, IEDM, 2051-2054, 2011.

Satellite EMI:
The fundamental rationale for this work is to increase lifetime of ANITA in the presence of unexpected RF backgrounds. Those backgrounds are discussed in section 1.1 but I found this section to feel very incomplete and poorly articulated. ANITA is supposed to observe RF transients (impulsive events) yet the backgrounds mentioned here are called out at continuous wave (CW). An explanation of how CW can come to dominate the triggering, whether it's simply a matter of total RF power or something which could be avoided with a different triggering scheme, should be mentioned. Is this not a similar environment to cell phone signaling? A related, but probably larger, gap is the argument that the CW is satellite-based. The third paragraph in section 1.1 makes a plausible case for the CW interference to be coming from military satellites (including additional sources launched between ANITA flights) but doesn't seem like proof. When the signals were observed, were other hypotheses considered? Tested? Can you point back to the satellites? Can you trigger and observe phase differences across the antennas consistent with a downward signal? I assume all of this was done, and the answers confirm the satellite hypothesis but it seems not completely shown here. 335-399MHz are Federal Government reserved frequencies, and a quick look in the literature shows long-term use of 375MHz in the Transit (commercial) satellites as well as more recent military usage. Is the broadcast modulation determinable from the ANITA data?

Thanks.

The transmitters are not a pure tone, but are modulated for data transmission, which triggers the system. While a different trigger scheme would avoid this problem (as it did in ANITA-2), it comes at the cost of reduced sensitivity, as mentioned in the text.

We changed Sec 1.1 to specify:

“The two main sources of noise are thermal radiation by the Antarctic ice and anthropogenic noise, much of which is modulated continuous-wave (CW) interference.”

and

“that were launched during the period from Feb. 2012 – June 2016. The CW signals generate events with excess power in left circular polarization (as expected) above the horizon, in approximately stationary positions.”

And

“However, ANITA-III was redesigned for improved sensitivity and based its trigger decisions on full-bandwidth (200-1200 MHz) signals. The modulation present in the CW interference produced trigger rates far in excess of the digitization system’s readout capabilities (~50 Hz).”

The exact confirmation of the source of the interference is not really a prime topic of the paper, which is why we stated that the interference was “*thought to be due to the newer Mobile User Objective System.”* All evidence (above horizon, stationary, increase in number of sources from ANITA-3 to ANITA-4) points to them being the source, however we are only interested in eliminating them from the data.

ANITA events are far too short (100 ns) to determine any modulation scheme (microsecond-scale).

Smaller comments:
last paragraph section 1.1: "high price" and "majority" could be quantified perhaps

Thanks.

We changed “During the majority of the ANITA-III flight…” to “For about 90% of the time during the ANITA-III flight, masking was used to veto triggers from over half of the payload field-of-view to keep the trigger rate at below 50Hz.”

We deleted “Both of these methods come at a high price.”

We added “Masking and decreasing thresholds come at the cost of instrument livetime (defined in Section 6.1) and sensitivity to neutrinos, respectively.”

We deleted “Decreasing thresholds led to reduction of sensitivity to neutrinos during noisy periods.”

sections 2 & 3: is there a standard ANITA hardware reference? couldn't a lot of this material be referenced rather than cut and pasted into this paper? there's a lot here which doesn't seem relevant to the tunable notch filters. also, might be worth calling out details on some parts, or references? The LNAs, and the hybrids, by manufacturer perhaps. Ditto the band filters

Thanks.

This is the first paper that describes the trigger systems for ANITA-III and ANITA-IV, which is needed to understand the notch filters. We do reference the ANITA instrument paper.

Level 1 trigger: this calls out the satellite interference as primarily circularly polarized, is that correct? is that consistent with the satellites being named as the EMI sources? does the balloon to satellite geometry and antenna polarizations all make sense?

Thanks.

Yes, the satellites are left-hand circularly polarized (LCP). We clarified this in Sec. 1.1.

LAB chip: much of this could just be a reference to the mentioned paper I think

Thanks.

This paper explains how the digitizer buffer works as it is critical to understanding digitization deadtime. The mentioned paper does not explain the buffer.

section 4: the circuit diagram is more useful than the picture, maybe combine them? also delete the "It can be seen that a single channel…"

Thanks.

We deleted “It can be seen that a single channel..” The picture is present to indicate the extremely compact size of the TUFF, along with the packaging picture which shows the physical constraints the system had to fit under. Some detail added in the text to stress the physical constraints.

We changed the text (lines 221-222) to have:

“The TUFF boards needed to be low-power and light, and compact to fit into the existing amplifier housing locations along with necessary cabling to match to the existing connectors.”

figure 6: a drawing might be more useful to show the packaging than the photo

Thank you.

We would like to keep the photo to show the physical constraints the system had to fit under.

section 4.1: a table could help with readability here

Thank you for the suggestion.

We added a table and the line “The table below summarizes properties of the amplifiers.”

figure 7: so 13dB reduction in those bands is sufficient to fix the problem with the satellite transmissions, this implies than the dynamic range of the experiment is quite limited, this might be worth commenting here. isn't 13dB just about 2 bits of an ideal ADC?

Thanks.

The satellites do not cause a (significant) problem in the dynamic range of the overall system. The problem is essentially due to triggering on the modulation of the CW.

section 4.2: the default notch filters seem to correspond to the always-there-since-ANITA1 satellite, the new satellites, and Antarctic bases UHF

radios, is that right? maybe call it out here. also, the South Pole and McMurdo LMR systems are at 450MHz, what is the frequency resolution of the ANITA digitizers for a CW source?

Thanks.

We added in section 4.2: “CW noise at the first two frequencies are thought to be caused by military communications satellites, specifically, the FLTSAT and UFO systems and the MUOS system, respectively. The third notch filter is present to curb CW interference seen when the ANITA payload is near Antarctic science bases such as McMurdo and South Pole Station.”

There are many transmitters present near McMurdo – the notch was placed at ~460 MHz to roughly cover the observed range of transmitters seen from ANITA-3 rather than targeting any one system.

figure 8: call out the amplifier parts on the figure

Thanks for that suggestion.

We changed the figure to say the amplifier part names.

figure 11: this implies a loss of signal strength across the whole band, not just the notch, can you quantify this? how does this efficiency loss compare to the lifetime gain from knocking off the peak?

Thanks.

The power spectra are computed from triggered data – when the interference was removed, the average power observed in triggered data is lower because the system no longer triggers on the interference. This is now clarified in the caption of figure 11:

“Note that the lower scale when Notch 2 is activated is not due to signal loss, but rather due to ANITA no longer triggering on the interference.”

section 6.3: highlight this, perhaps at the start of section 6 as this is the main goal I believe. if the filters hadn't been programmable, this would be

almost as good (or as good) right? what would be lost if the 450-460MHz was on all of the time?

Thanks.

We changed beginning of Section 6 to read:

“During the ANITA-IV flight, we utilized all features of the TUFF notch filters to achieve decreased masking, increased stability of trigger rate and SURF DAC thresholds, and increased instrument livetime, from 31.6% in ANITA-III to 91.3% in ANITA-IV. These results are summarized in Figures 3, 4, 14 and 15.”

Note that as specified previously in the text and shown in Fig 13, the programmability of the notches was in fact necessary, as the main interference frequency did move slightly (from 380 to approximately 390) over one portion of Antarctica. The notches were in fact made programmable initially simply for manufacturing purposes (to compensate for part-to-part variation), and in-flight retuning was only an “if necessary” possibility – which it was.

section 7: Is ANITA-V a proposed project? Funded? The RITC system was described as being for ANITA-3. Might want to comment on that here.

Thanks.

We changed Section 7 as follows:

7. Future plans

“A fifth ANITA flight (ANITA-V) is currently proposed. For this flight, we will…”

and

“The main proposed upgrade to the triggering system include the Realtime Independent Three-bit Converter (RITC) as described by Nishimura et al. [6] This triggering system, conceived for ANITA-3 and intended for ANITA-4, was held back until CW mitigation, as shown here, could be demonstrated. The RITC will…”