Highly Multiplexed Superconducting Detectors and Readout Electronics for Balloon-Borne and Ground-Based Far-Infrared Imaging and Polarimetry

by

Samuel Gordon

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved May 2019 by the
Graduate Supervisory Committee:

Philip Mauskopf, Chair
Judd Bowman
Christopher Groppi
Daniel Jacobs
Paul Scowen

ARIZONA STATE UNIVERSITY

August 2019

ABSTRACT

This dissertation details the development of an open source, frequency domain multiplexed (FDM) readout for large-format arrays of superconducting lumped-element kinetic inductance detectors (LEKIDs). The system architecture is designed to meet the requirements of current and next generation balloon-borne and ground-based submillimeter (sub-mm), far-infrared (FIR) and millimeter-wave (mm-wave) astronomical cameras, whose science goals will soon drive the pixel counts of sub-mm detector arrays from the kilopixel to the megapixel regime. The in-flight performance of the readout system was verified during the summer, 2018 flight of ASI's OLIMPO balloon-borne telescope, from Svalbard, Norway. This was the first flight for both LEKID detectors and their associated readout electronics. In winter 2019/2020, the system will fly on NASA's long-duration Balloon Borne Large Aperture Submillimeter Telescope (BLAST-TNG), a sub-mm polarimeter which will map the polarized thermal emission from cosmic dust at 250, 350 and 500 microns (spatial resolution of 30″, 41″ and 59″). It is also a core system in several upcoming ground based mm-wave instruments which will soon observe at the 50 m Large Millimeter Telescope (e.g., TolTEC, SuperSpec, MUSCAT), at Sierra Negra, Mexico.

The design and verification of the FPGA firmware, software and electronics which make up the system are described in detail. Primary system requirements are derived from the science objectives of BLAST-TNG, and discussed in the context of relevant size, weight, power and cost (SWaP-C) considerations for balloon platforms. The system was used to characterize the instrumental performance of the BLAST-TNG receiver and detector arrays in the lead-up to the 2019/2020 flight attempt from McMurdo Station, Antarctica. The results of this characterization are interpreted by applying a parametric software model of a LEKID detector to the measured data

in order to estimate important system parameters, including the optical efficiency, optical passbands and sensitivity.

The role that magnetic fields (B-fields) play in shaping structures on various scales in the interstellar medium is one of the central areas of research which is carried out by sub-mm/FIR observatories. The Davis-Chandrasekhar-Fermi Method (DCFM) is applied to a BLASTPol 2012 map (smoothed to 5′) of the inner $\sim$1.25 deg^2 of the Carina Nebula Complex (CNC, NGC 3372) in order to estimate the strength of the B-field in the plane-of-the-sky ($B_{\rm POS}$). The resulting map contains estimates of $B_{\rm POS}$ along several thousand sightlines through the CNC. This data analysis pipeline will be used to process maps of the CNC and other science targets which will be produced during the upcoming BLAST-TNG flight. A target selection survey of five nearby external galaxies which will be mapped during the flight is also presented.

ACKNOWLEDGMENTS

There are people in this world to whom I owe an immeasurable debt of gratitude. That I somehow become mixed up with all of them in the first place is one of the greatest fortunes of my life.

To Phil Mauskopf, Gena Pilyavsky, Adrian Sinclair, Caleb Wheeler, Natalie Hinkel, Chris Groppi, Hamdi Mani, Paul Scowen, Judd Bowman, Danny Jacobs, Boom Kittiwisit, Tom Mozden, Barb Mozden, Rick Sarmento, Sean Bryan, Paul Horton, Jose M. Chavez-Garcia, Emilie Dunham, Nari Miller, Rubble, George Che, Kay Davis, Marko Neric, Becky Jackson, Emily Lunde, Justin Mathewson, Matt Underhill, Eric Weeks, Cassie Whitton, Jonathan Hoh, Andrea Aparicio, Mark Devlin, Giles Novak, Jeff Klein, Laura Fissel, Federico Nati, Nate Lourie, Ian Lowe, Paul Williams, Brad Dober, Tyr Galitzki, Javier Romualdez, Jordan Wheeler, Pete Barry, Sam Rowe, Matt Strader, Glenn Jones, Peter Day, Jiansong Gao, Chris McKenney, Grant Wilson, Kamal Souccar, Silvia Masi, Paolo De Bernardis, Francesco Piacentini, Ellia Battistelli, David Schiminovich, Erika Hamden, Szabolcs Marka, Petr Kout, Heather McCarrick, Luke Goetzke, Elena Aprile, Adam Greenberg, Guy Tower, Jules Halpern, Sharon Bender (Mom), Bob Gordon (Dad), Josh Gordon, Karl Biemuller, Merlin, Pieces:

Thank you, thank you, thank you.

This work was funded in part by a NASA Earth and Space Science Fellowship in Astrophysics (NESSF NNX16AO91H).

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

GLOSSARY

A_V Total visual extinction [mag].

$B_{\rm LOS}$ B-field along the line-of-sight.

$B_{\rm POS}$ B-field in the plane-of-the-sky.

I_0 Modified Bessel function of order zero.

I_c Critical current of a superconductor.

K_0 Modified Bessel function of the second kind of order zero.

L_m Magnetic inductance.

$N(\rm H_2)$ Column density of molecular hydrogen.

N_0 Single-spin density of states at the Fermi Energy.

N_{qp} Total number of quasiparticles.

$P_{\rm abs}$ Total power which is absorbed by the detectors, accounting for the system optical and detector efficiencies.

$P_{\rm opt}$ Optical power which in incident on the telescope optics or cryostat window.

$P_{\rm ro}$ Microwave readout power.

Q_c Couping quality factor.

Q_i Internal quality factor.

Q_r Resonator quality factor.

R_P Fractional frequency responsivity to absorbed optical power $(df/f_0)/pW$.

R_T Fractional frequency responsivity to changes in base temperature $(df/f_0)/mK$.

R_{sq} Sheet resistance in the normal state.

S_{21} Forward-transmission scattering parameter.

S_Φ Polarization-angle disperion.

$S_{\delta f/f}$ Power spectral density of fractional frequency flucuations. Expressed in units of Hz^{-1} or dB/Hz.

S_{XX} Power spectral density of fractional frequency fluctuations, in the phase direction [Hz^{-1}].

S_{YY} Power spectral density of fractional frequency fluctuations, in the dissipation direction [Hz^{-1}].

S_{ϕ} Power spectral density of phase fluctuations between I and Q. Expressed in units of $\mathrm{rad}^2/\mathrm{Hz}$ or dBc/Hz.

T_{base} Base temperature of the focal-plane-array.

T_c Critical temperature.

Δ_0 One half of the zero-temperature superconducting band gap energy.

Δf Frequency shift relative to a channel's resonant frequency.

Γ_{gen} Quasiparticle generation rate.

Γ_{rec} Cooper pair recombination rate.

Φ Inferred angle of $\mathrm{B}_{\mathrm{POS}}$ ($\Psi \pm 90^{\circ}$).

Ψ Polarization-angle of the dust emission.

α Kinetic inductance fraction.

$\boldsymbol{v}_i$ Polarization vector for the *ith* pixel.

ϵ_{Xpol} Cross-polarization efficiency.

ϵ_{a} Resonator asymmetry parameter.

ϵ_{opt} Optical efficiency.

ϵ_{pol} Polarization efficiency.

η_{det} Detector efficiency.

κ_{ν} Dust opacity, [cm^2/g].

$\mathbf{S}(\theta)$ Polarization response as function of grid rotation angle θ.

$\mathcal{L}_k$ Kinetic inductance per unit length.

$\mathscr{L}_{k,sq}$ Kinetic inductance per square.

$\nabla_{I,Q}$ I/Q gradient: $(dI/df, dQ/df)$.

$\phi_{I/Q}$ Phase of S_{21}.

σ_1 Conductivity of the quasiparticle fluid.

σ_2 Conductivity of the superfluid.

σ_d Drude conductivity.

σ_n Normal state conductivity.

σ_v Line-of-sight velocity dispersion.

τ_{qp} Quasiparticle recombination time.

τ_s Scattering time.

f_c The 1/f corner-frequency.

$m^\star$ Effective electron mass.

m_e Rest mass of an electron.

$n(\mathrm{H}_2)$ Number density of molecular hydrogen.

n_{occ} Photon occupation number.

n_{qp} Quasiparticle number density.

v_A Alfvén velocity.

Chapter 1

INTRODUCTION

This dissertation describes the development and implementation of a readout system for a novel detector technology known as the microwave kinetic inductance detector (MKID, or KID) (Day et al., 2003; Mazin, 2005). MKIDs are superconducting resonators which change their resonant frequency in response to the absorption of incident photons that have enough energy to break apart superconducting Cooper pairs. The change in resonant frequency is proportional to the amount of optical power which has been absorbed by the inductor of the resonator.

When combined, the detector arrays and readout system form the core of a powerful astronomical camera. Although many MKID architectures exist, the one which is primarily discussed in this document is the lumped-element kinetic inductance detector (LEKID) (Doyle et al., 2008). The tunable resonant frequencies of MKIDs allow them to be frequency domain multiplexed (FDM) on a single feedline. Current multiplexing factors are in the kilopixel regime. However, the demands of next generation submillimeter (sub-mm), far-infrared (FIR) and millimeter-wave (mm-wave) observatories (e.g., CMB-S4 (Abitbol et al., 2017)) are rapidly pushing the pixel count requirement into the megapixel regime. Fortunately, recent advances in FPGA-based electronics have enabled the use of high speed digital signal processing (DSP) algorithms which are required to match the high multiplexing factors of the detector arrays.

The ASU LEKID readout incorporates firmware, software and electronics which are specifically designed to work with astronomical MKID/LEKID instruments ob-

serving in the sub-mm/FIR/mm-wave (see Section 1.1). These instruments are installed on ground-based and stratospheric balloon-based observatories. Although the ASU detector readout has been integrated into several current and planned instruments, it was primarily designed for NASA's Next Generation Balloon-Borne Large Aperture Submillimeter Telescope (BLAST-TNG) (Dober et al., 2014). Its design architecture and initial performance are described in Gordon et al. (2016).

BLAST-TNG (see Section 1.4) is scheduled for a 28 day flight from NASA's Long Duration Balloon Facility (LDB) near McMurdo Station, Antarctica, during winter 2019/2020. If it succeeds, this flight will be the second for an MKID-based camera system. In July, 2018, the readout system flew on the Italian Space Agency's (ISA) OLIMPO balloon, which was the first test of these emerging technologies in a space-like environment (Masi et al., 2019). The lessons learned from each successive application of the system lead to improvements in system performance and reliability. Successful balloon flights can advance the NASA technology-readiness-level (TRL) of such systems to the point at which they can be used on a space-based observatory.

1.1 Submillimeter and Millimeter Wave Astronomy

The primary purpose of the detector technology which is the subject of this dissertation is to enable sub-mm/FIR and mm-wave astronomy. The terminology used to refer to these regions of the electromagnetic spectrum is sometimes the source of confusion. In this work, the terms mean the following:

Submillimeter: Wavelengths of 0.1–1 mm (frequencies of 0.3–3 THz, which some astronomers refer to as the terahertz band).

Millimeter: Wavelengths of 1–10 mm (frequencies of 30–300 GHz). The microwave band, here defined as 0.3–300 GHz, overlaps the entire mm-wave band.

Far-infrared: Wavelengths of 25–350 μm (frequencies of 0.857–12 THz). FIR shares some overlap with the sub-mm band.

1.1.1 Atmospheric Transmission

The sub-mm/FIR/mm-wave radiation that we detect from space is heat emitted by microscopic particles which fill the spiral arms of galaxies. These particles are collectively referred to as cosmic dust. Since the discovery of cosmic dust in the 1930s-40s, its study has advanced at a constant but relatively slow pace when compared to other areas of research within astronomy and astrophysics. This lag in development can be attributed to two main factors.

The first factor is that the photons which are emitted by the dust, which is mostly at temperatures $\lesssim$100 K, are at very low energies. Their successful detection therefore requires extremely low backgrounds, along with the use of very sensitive devices, such as superconductors. The second factor is that the troposphere is mostly opaque to the dust emission, due to absorption by water vapor, oxygen and carbon dioxide. Observations of the sub-mm emission from space must therefore be conducted from sites that are as high in altitude and as dry (in terms of precipitable water vapor) as possible. At present, the two best sites in the world from which to conduct sub-mm astronomy are the Atacama desert in Chile, and Antarctica.

1.2 Cosmic Dust

In the 19th century, the question of why the night sky is mostly dark, given a presumably infinite number of stars, came to be known as Olbers' Paradox. The commonly known solutions, which have become more nuanced over the course of time, usually cite the confluence of the vastness of space and the finite speed of light[1]. While these solutions suffice, they typically leave out an often overlooked contributor to the darkness of the night sky. This is cosmic dust.

Cosmic dust fills the spiral arms of the Milky Way and other galaxies which host active star formation. On a very dark night, the dust lanes of the Milky Way are noticeably darker than surrounding starlit regions of the sky (see Figure 1).

Figure 1. The Milky Way galaxy, from Cerro Paranal, Chile. The dust, which is mostly confined inside the spiral arms, obscures the glow of billions of background stars. Image from Bruno Gilli/ESO: https://www.eso.org

It has been almost a century since astronomers first became interested in interstellar dust and started to study its properties in earnest. These earlier astronomers thought of the dust as being somewhat akin to smoke. In *Dust in the Galactic Environment* (Whittet, 2002), the author writes:

[1]This speed limit applies to all forms of information transfer.

"The term *smoke* was often used to describe these particles in the early literature. *Smoke* implies the product of combustion, whereas *dust* implies finely powdered matter resulting from the abrasion of solids. The former is arguably more appropriate as a description of the particles condensing in a stellar atmosphere, now regarded as an important source of interstellar grains."[2]

1.2.1 Observational Inferences

Several key observational inferences about the dust point to its role as an important catalyst in both small and large-scale processes related to star formation in galaxies. The grains exhibit differential, or wavelength-dependent extinction, absorbing photons in the optical and near-infrared (O/NIR) while scattering higher energy photons out of the line-of-sight (LOS). The radiation which is absorbed by the grains is re-emitted in the sub-mm/FIR/mm-wave bands. This thermal emission is linearly polarized (see Section 1.2.3).

The extinction curve of the interstellar dust indicates a size distribution of $\sim$0.05–0.5 μm, where the smallest grains are individual molecules (e.g., polycyclic aromatic hydrocarbons (PAHs)). The exact composition and material properties of the grains are the subject of ongoing study (see, e.g., Andersson et al. (2015); Draine (2003)). The O/NIR dust SED indicates that the continuum emission is due to a population of silicate grains, while narrow emission features correspond to PAHs and other carbonaceous grains. The sub-mm/FIR polarization spectrum of the dust, as measured in nearby molecular clouds (MCs), contains a minimum at $\sim$350 μm (Figure 2). This feature indicates that the observed SED is the result of two distinct dust populations.

[2]Italics added.

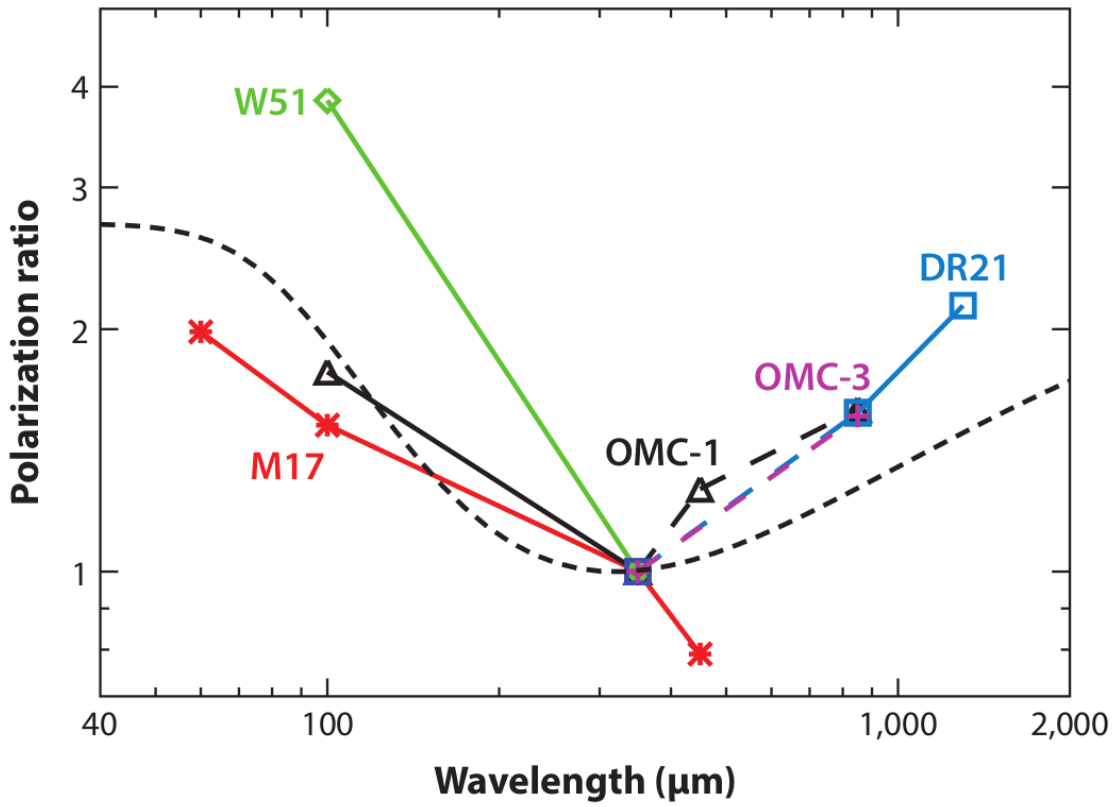

Figure 2. The sub-mm polarization spectrum (normalized to the polarized emission at ~350 μm) for several nearby MCs (Andersson et al., 2015). The minimum in the spectrum at ~350 μm points to a two-component dust population.

Measurements of the spectral energy distribution (SED) of the extragalactic background light (EBL) show that roughly the same amount of energy is contained in the sub-mm/FIR part of the spectrum as in the O/NIR (see Figure 3). These two distinct regions of the SED are known as the cosmic-optical-background (COB) and cosmic-infrared-background (CIB). The COB and CIB peak at ~1 μm and ~100 μm. By comparing measurements of the GOODS-S region of the sky taken by BLAST 2006 at 250, 350 and 500 μm with those taken in other wavebands, it was discovered that approximately half of the CIB can be associated with individual sub-mm galaxies (SMGs) at cosmological redshifts of $1 \leq z \leq 4$ (Devlin et al. (2009a); Marsden

et al. (2009); Pascale et al. (2009)). This redshift range corresponds to a period in the universe's history when star formation occurred at rates which are many times higher than that which is observed in the present day universe.

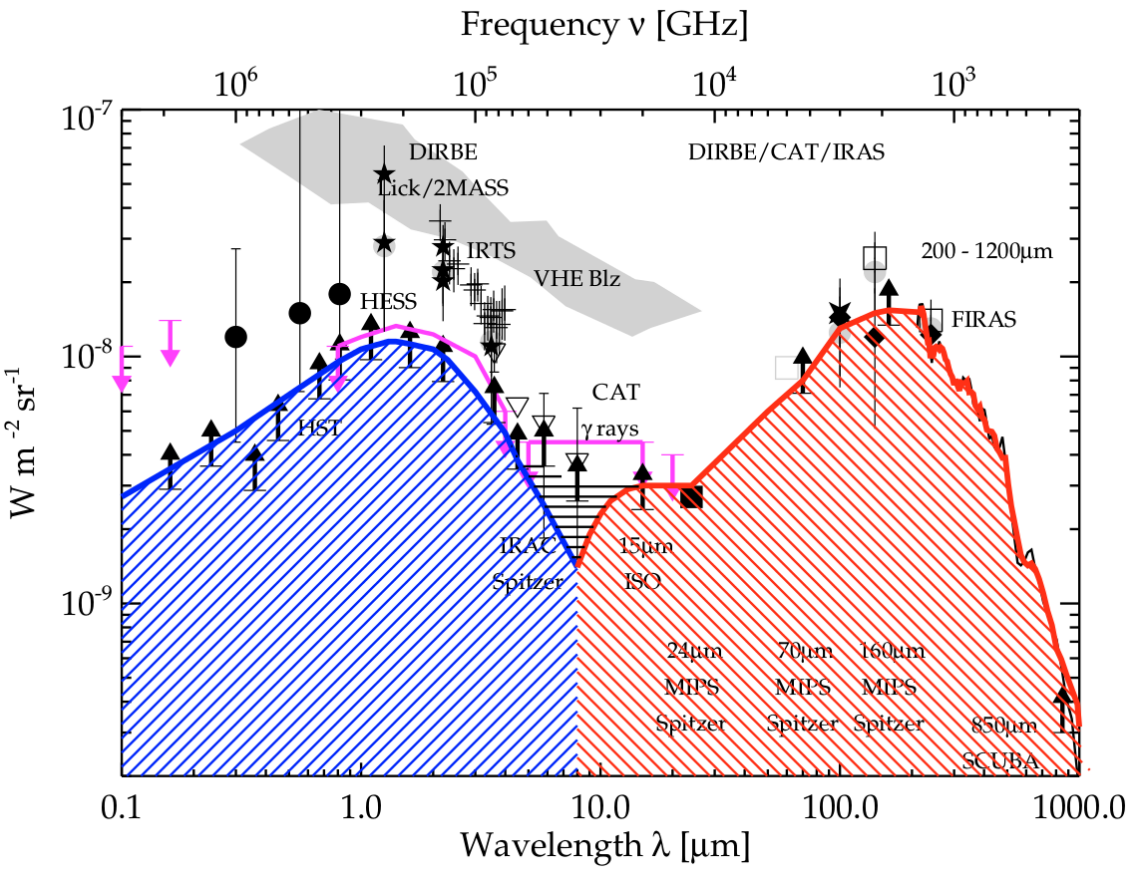

Figure 3. The SED of the EBL, produced from Spitzer observations (Dole et al., 2006). The amount of energy contained in the COB (blue) and CIB (red) is roughly equal.

1.2.2 The Lifecycle of Cosmic Dust

Interstellar dust is perpetually recycled in the ISM of galaxies. The raw elements of which it consists are created by main-sequence stars. These refractory elements

condense out of the cool photospheres of AGB stars and planetary nebulae, and coalesce into grains. More dust is created by supernovae.

After it is created, the dust is dispersed into the ISM surrounding its parent star by stellar winds produced by OB and AGB stars, and by the violent explosions of supernovae. Once in the diffuse ISM, it becomes incorporated into MCs, the birth places of stars. A strong correlation has been observed between the star formation rate (SFR) of galaxies and their dust mass (da Cunha et al., 2010). As a rule of thumb, the gas-to-dust mass ratio of the ISM of the Milky Way and other spiral galaxies is commonly assumed to be $\sim$100.

Inside MCs, The charged surfaces of the grains serve as the sites of molecular hydrogen formation. In the densest regions of MCs, the dust is collisionally coupled to the gas, and shields it from the ionizing radiation from nearby stars. These two effects allow the gas and dust to cool together, further collapsing under gravity. In prestellar cores, the dust temperatures can range from 10–100 K.

Following star formation, some of the dust grains are recycled by the star, and others become incorporated into newly formed planets, asteroids and other solar system bodies. In regions of a galaxy which contain little to no star formation, the remaining dust grains are eventually destroyed by cosmic rays.

1.2.3 The Role of Dust as a Tracer of Cosmic Magnetic Fields

Observed SFRs in Galactic MCs are far lower than what is expected from a theoretical model of cloud collapse which considers only gravity. Theorists seeking an explanation for what could be slowing the collapse of prestellar cores have settled on two likely culprits. These are supersonic turbulence and magnetic fields (B-fields).

Turbulence prevents the effective dissipation of gravitational energy, while B-fields provide an additional source of pressure support that slows, or even halts, the gravitational collapse of MCs.

In the 1940s, astronomers discovered that a small fraction of the O/NIR light which is received from stars is linearly polarized (Hall (1949); Hiltner (1949)). The correlation between the polarization of starlight and the previously observed phenomenon of interstellar reddening pointed to cosmic dust as the cause of the polarization. Since the O/NIR starlight is linearly polarized, so is the portion of the starlight which was absorbed by interstellar dust and re-emitted at sub-mm/FIR wavelengths.

The cause of the linear polarization of thermal dust emission is rooted in the aspherical nature of the grains. Through a complex alignment mechanism, the aspherical dust grains, which are spinning due to interactions with the gas and radiation field, end up having their long axes orthogonal to the local B-field direction. Consequently, the electric field (E-field) of the thermal radiation emitted by the dust is preferentially oriented along the long axis of the grains. By rotating the measured polarization angle of the E-field by 90 degrees, the direction of the B-field in the plane-of-the-sky ($B_{\rm POS}$) can be inferred. Because the polarity of $B_{\rm POS}$ is unknown, it is referred to as a pseudovector.

The physics of the alignment mechanism which is responsible for aligning the dust grains with the local B-field has been the subject of long debate (Andersson et al., 2015). The most commonly accepted mechanism is described by the theory of radiative alignment torques (RAT) (Lazarian and Hoang, 2007). In RAT theory, the alignment is facilitated by both the anisotropy of the radiation field and the paramagnetic nature of the dust grains.

Without additional knowledge of the B-field, the magnitude of $B_{\rm POS}$ is difficult

to estimate. However, indirect methods, such as the Davis-Chandrasekhar-Fermi method (DCFM) (Chandrasekhar and Fermi, 1953), can sometimes be used to place upper constraints on the field strength (see Chapter 5).

1.2.4 Combining Total Intensity and Polarization Measurements

By combining total intensity and polarization measurements of the same region, it is possible to ascertain the relative orientations between $B_{\rm POS}$ and matter structures composed of gas and dust. The data yielded by such observations can be used to model the relationship between the B-field and several environmental variables, including temperature, hydrogen and dust column density and polarization fraction.

The relationship between the above parameters continues to be examined across a wide range of spatial scales. An example of one such analysis is shown in Figure 4 (Fissel et al., 2016). The color scale is the BLASTPol 2012 500 μm dust intensity, and the drapery pattern is a line-integral-convolution (LIC) of $B_{\rm POS}$ (see Section 5.2). The spatial resolution in the image is ∼2.5 pc, which is sufficient to resolve the large scale filamentary structure of MCs, but not structures within individual filaments (this requires resolution of $\lesssim$0.1 pc).

1.2.5 The Role of Cosmic Dust as a CMB Foreground

By virtue of being polarized, thermal emission from interstellar dust grains poses a problem for polarization measurements of the Cosmic Microwave Background (CMB) radiation. It is now known that the polarized emission from diffuse Galactic dust is the dominant CMB foreground at frequencies above 100 GHz (Adam et al.,

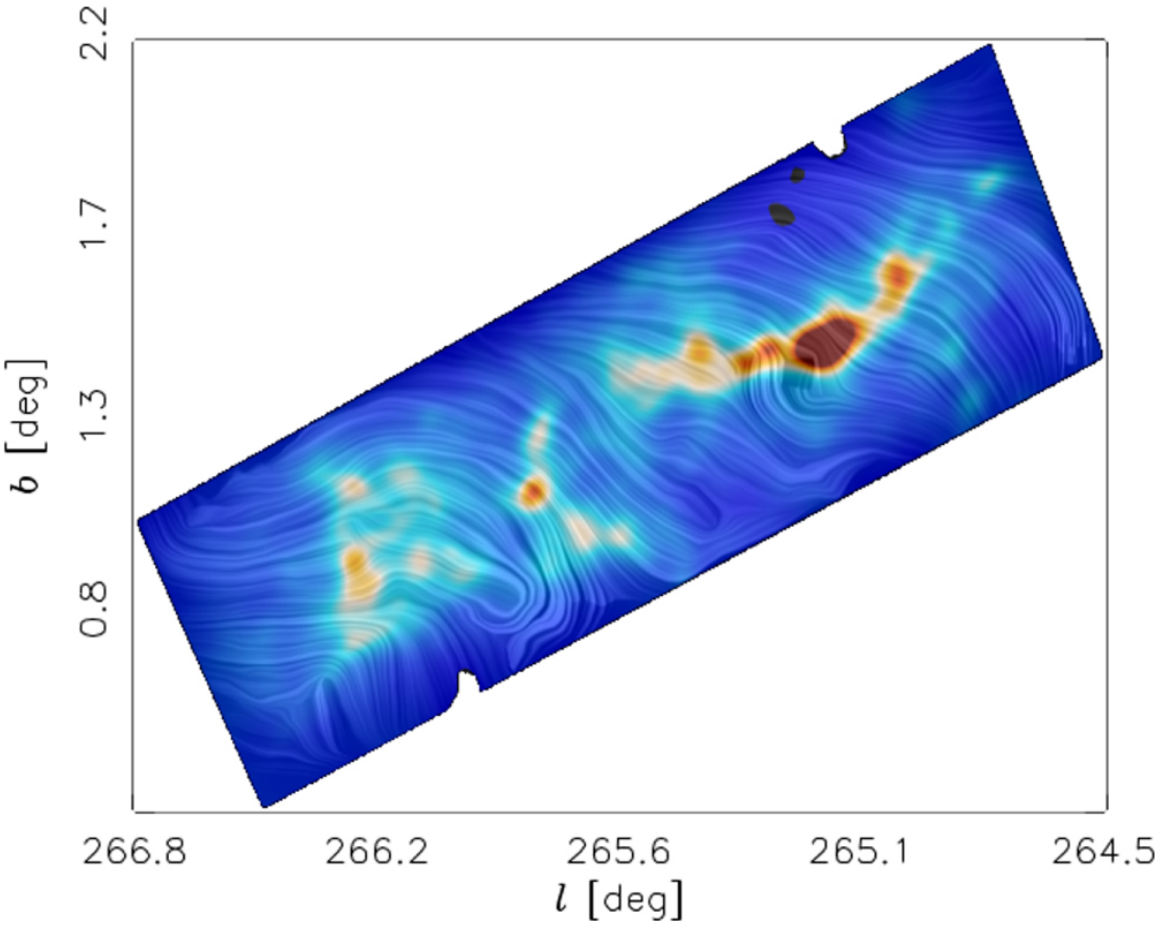

Figure 4. A LIC map of the Vela-C MC, produced from BLASTPol 2012 data (Fissel et al., 2016). The color scale is 500 μm dust intensity.

2016). Any attempt to detect primordial B-mode polarization in the CMB will be hindered by the polarized dust foreground emission. The solution to this problem is to carefully map the dust polarization in the same regions which are observed by CMB telescopes, over as wide a waveband as possible. The foreground signal can then be cross-correlated with the CMB measurements, and removed.

1.3 Technologies for Submillimeter and Millimeter-Wave Astronomy

MKIDs are incoherent detectors, measuring the total power in the incident electromagnetic radiation (photons), but not the phase. In a semiconductor sensor such as a charge-coupled device (CCD), the energy band gap of ∼1 eV restricts their optical response to wavelengths between ∼0.3–1 μm. The superconducting band gap is $\mathcal{O}(10^{-}3)$ smaller than that of a semiconductor. Consequently, MKIDs (and other superconducting detector technologies) can be engineered to detect radiation which spans the energy range between mm-waves and Gamma-rays.

The narrow energy band gap of MKIDs and other superconducting detectors gives them better sensitivity than semiconductor-based detectors. However, they must be operated at a temperature which minimizes the thermal device noise (generation-recombination (GR) noise) to below the level of the photon (shot) noise in the passband. As a rule of thumb, this temperature should be $\sim T_c/8$ (Mazin et al., 2013), where T_c is the critical temperature of the device. For a typical MKID T_c of ∼1.5 K, this corresponds to fridge temperatures of 100–300 mK.

To achieve and maintain such low temperatures throughout an astronomical observation is an engineering feat in its own right. In order to shield the detectors from the thermal radiation of the elements within the optical system, such as mirrors, filters and optomechanics, it is necessary to cool as much of the optical system as possible (sometimes to as low as ∼4 K). Cooling of the optics and detectors (also some of the readout electronics, in certain cases) is facilitated using combinations of open and closed-cycle cryostats which use liquid and gaseous Helium-4 (LH4) and Helium-3 (LH3) as coolant.

1.3.1 KIDs and TESs

In many aspects, the KID may be considered a successor to the transition-edge sensor (TES) bolometer. TESs, which take a variety forms, have been used by many successful ground, balloon and spaced-based sub-mm/FIR/mm-wave observatories. They are also the chosen sensor technology of several current and planned astronomical observatories (e.g., The Simons Observatory (Ade et al., 2019)).

TESs consist of a superconducting absorber which is part of a voltage-biased circuit that has been tuned to just above the edge of its superconducting transition-edge. The TES is inductively coupled to a superconducting quantum interference device (SQUID) followed by a low-noise amplifier (LNA). Incident photons cause the absorber to heat up, increasing the resistance of the TES. Because the TES is voltage-biased, the rise in resistance causes the current through the TES to decrease. The change in current is read out using the SQUID. The voltage-biasing of TESs establishes negative electrothermal feedback, which restores them to their superconducting state after each absorption of a photon.

As astronomical detectors, KIDs and TESs have relative advantages and disadvantages, particularly in regards to their readout systems. A detailed comparison of the two technologies can be found in Mauskopf (2018). While more recent TES cameras use FDM readouts (μMUX, see e.g, Stanchfield et al. (2016)), TESs have been traditionally time-domain multiplexed (TDM). In the TDM readout, $N \times M$ detectors are read out in sequence by N-columns of SQUID-switches. In a given data frame, each TES is readout $1/N$ of the time, where N is the multiplexing factor. At present, the highest multiplexing factors which have been achieved are N of 64–128 (Henderson et al. (2016); Mates et al. (2017)).

In addition to their limited multiplexing factor, the cryogenic part of the TES readout requires many cryogenic wires/wirebonds ($N + M$ wires for $N \times M$ TESs (Mauskopf, 2018)). The thermal noise in the readout chain can be higher than in KID systems due to the dissipative readout. KIDs, by virtue of being FDM, require only one cryogenic feedline and LNA for N detectors. With the exception of the LNA, their readout electronics are entirely room-temperature.

Despite the relative complexity of their readouts, TESs do have advantages over KIDs. The dissipative readout gives the advantage of self-regulation through electrothermal equilibrium, which increases their response time after each photon hit. KIDs have no equivalent self-regulating mechanism, which makes biasing them somewhat challenging. Another advantage of the TES readout is that the sum of the absorbed optical power and readout power is constant. With KID readout systems, signal processing must be performed on the raw data outputs in order to make it proportional to the absorbed optical power (see Chapter 4).

As for their noise properties, TESs are typically stable at frequencies down to a few millihertz. Many KIDs are observed to have excess noise at low frequencies, which has implications for the mapping capabilities of the experiment (see Section 3.1.3).

1.4 BLAST-TNG

The Next Generation Balloon-borne Large Aperture Submillimeter Telescope (BLAST-TNG) is a sub-mm imaging polarimeter which will map the polarized thermal emission from interstellar dust, revealing B-field structures in nearby MCs, giant molecular clouds (GMCs), the diffuse interstellar medium and in nearby external

	BLAST (2006)	BLASTPol (2010)	BLASTPol (2012)	BLAST-TNG (2020)
Bands (μm)	250, 350, 500	250, 350, 500	250, 350, 500	250, 350, 500
Det. type	SiN SWB[1]	SiN SWB	SiN SWB	LEKID
$N_{\rm det}$	149, 88, 43	149, 88, 43	139, 88, and 43	∼1500, 700, 300
Pol.	N	Y	Y	Y
Aperture (m)	1.8	1.8	1.8	2.5
Beam FWHM (″)	36, 42, 60	∼150[2]	∼150[3]	30, 41, 50
Flight time (days)	12	9	13	28 (planned)

(1) Silicon Nitride (SiN) spider web bolometers (SWB). These were flown on the Herschel Space Observatory's Spectral and Photometric Imaging Receiver (SPIRE) 2009–2013 (Griffin et al., 2003).
(2) A melted IR blocking filter reduced the overall resolution to ∼2.5′, after correcting the maps (Matthews et al., 2014).
(3) The point-spread function (PSF) of the optical beam was warped, resulting in an overall resolution ∼2.5′ after correcting the maps (Fissel et al., 2016).

Table 1. Comparison of camera parameters between BLAST-TNG, BLASTPol 2012, BLASTPol 2010 and BLAST 2006.

galaxies. It will observe in three 30% bands centered at 250, 350 and 500 microns (1200, 857 and 600 GHz), with spatial resolution of 30″, 41″ and 50″. These physical scales bridge the gap between the 5′ all-sky polarization maps produced by the ESA's Planck space observatory, and ALMA's sub-arcsecond resolution.

The experiment is scheduled for a ∼28 day flight from NASA's Long Duration Balloon Facility, near McMurdo Station, Antarctica, in winter 2019/2020[3]. BLAST-TNG is the latest incarnation of a series of BLAST experiments which date back to 2005. Previous BLAST payloads had successful flights in 2005, 2006, 2010 and 2012. Table 1 lists key instrumental parameters for each of these experiments. BLAST-TNG is 'next-generation' for several reasons. Chief among these are:

Detectors and readout electronics: BLAST-TNG is the first NASA-funded balloon experiment to feature large-format (kilopixel scale) arrays of LEKIDs, along with the highly multiplexed FPGA-based electronics which are required to read them out[4]. Its camera features ∼3,000 dual-polarization sensitive, horn-coupled LEKIDs, which are fabricated from a trilayer substrate consisting of titanium nitride/titanium/titanium nitride (TiN/Ti/TiN) (Hubmayr et al., 2014). In this capacity, BLAST-TNG serves as a pathfinder for upcoming balloon experiments, including EXCLAIM (Switzer, 2017) and TIM/STARFIRE (Aguirre et al., 2018).

Primary mirror: BLAST-TNG's 2.5 m diameter primary mirror is based on a composite carbon fiber reinforced polymer (CFRP) design developed by Alliance

[3]Several launch attempts were made in winter 2018/2019, but the flight was ultimately delayed until the following season.

[4]In summer, 2018, the ASI's OLIMPO balloon succeeded in becoming the first flight for LEKID detectors. Their readout system was based on the firmware and software from the ASU detector system, with a modified set of intermediate frequency (IF) electronics (see Section 1.5).

Spacesystems[5]. The primary mirror will be the largest to-date which has flown on a long-duration stratospheric balloon. The optical design of BLAST-TNG is described in Lourie et al. (2018).

Mapping speed: With its large aperture primary mirror and ∼3,00 LEKID detectors, the mapping speed of BLAST-TNG is expected to be much greater than that of its predecessor, BLASTPol 2012. The resulting maps will contain hundreds of thousands of B-field pseudo-vectors– an order of magnitude improvement over BLAST-Pol.

1.4.1 Science with BLAST-TNG

BLAST-TNG will perform five science surveys, covering the following categories:

Nearby Molecular Clouds: These are the closest sites of active star formation (visible from McMurdo Station, Antarctica). BLAST-TNG will map $B_{\rm POS}$ down to the scale of the filaments of gas and dust which thread MCs. Many of the dense cores which form the nodes of these filaments will form into stars. These maps will enable the study of correlations between physical and environmental parameters in the MC, including the dust polarization fraction and spectrum, gas and dust temperatures, column densities and B-field direction and strength (see, e.g., Galitzki et al. (2014a); Fissel et al. (2016, 2018); Shariff et al. (2019); Gandilo et al. (2016)).

Giant Molecular Clouds: GMCs are among the densest regions of the ISM, and

[5]https://alliancespacesystems.com/

host large amounts of massive star formation. The clusters of massive O and B stars which form inside GMCs inject energy into the surrounding ISM via ionizing radiation, stellar winds and supernovae.

Diffuse Galatic ISM: In order to constrain physical models of different dust populations and their alignment mechanisms, it is necessary to probe structures within the ISM across a wide range of densities. This includes mapping the lower density, diffuse regions of the ISM. These observations will provide input to numerical models which test the correlation between the measured dust polarization properties and the orientation of the B-field.

CMB Foregrounds: Mm-wave measurements of polarization in the CMB are contaminated by the polarized dust emission. High resolution maps of the polarized dust emission in the regions observed by CMB observatories are required to clean the CMB maps of this foreground. Around 100 hours of the BLAST-TNG flight is designated to mapping polarized dust emission on degree scales in regions of the sky that overlap with those previously mapped by CMB observatories.

Nearby External Galaxies B-fields constitute a significant part of the energy density in galactic ISMs. Through flux freezing and channeling of cosmic ray electrons, they can influence the evolution of structure over a wide range of physical scales within a galaxy. At radio wavelengths, synchrotron emission and Faraday rotation has been used to map the large scale fields in a variety of galaxy types, including spirals, ellipticals, dwarf irregulars and interacting pairs. BLAST-TNG will complement these radio observations with the first comprehensive survey of $B_{\rm POS}$ as traced by interstellar dust (see Section 6.2).

1.5 Current KID-based Instruments

BLAST-TNG and the ASU LEKID readout build on the legacy of past sub-mm/FIR/mm-wave cameras (e.g., MUSIC (Golwala et al., 2012), MAKO (Swenson et al., 2012), NIKA (Bourrion et al., 2013), NIKA2 (Bourrion et al., 2016), AMKID (van Rantwijk et al., 2016), DESHIMA (Endo et al., 2012)), as well as MKID-based O/NIR cameras (e.g., ARCONS (Mazin et al., 2013), DARKNESS (Meeker et al., 2018; Strader, 2016)). While much of the electronics, firmware and software which is required to readout MKIDs is similar for sub-mm/mm-wave and O/NIR cameras, the mode in which the detectors are operated is different.

1.5.1 Considerations on Sub-mm/mm-wave and O/NIR MKID Readout Systems

In the sub-mm/mm-wave, MKIDs are used in bolometer-mode, as power detectors. In the O/NIR, they are used in Geiger-mode, to count single photon arrival times and measure their energy. When KIDs are operated in bolometer (or background-limited infrared photon, BLIP) mode, slow variations in the background loading cause them to drift off-resonance. To track the detectors' responsivities during an observation requires 'tuning' the probe tones which are generated by the readout system in amplitude and frequency. Tuning can be implemented in either firmware or software, and depending on the number of detectors and detector parameters which are taken into account, can be very resource and time intensive (see, e.g., Dodkins et al. (2018)).

While bolometric MKID systems must address the tuning requirement, they have

the advantage of being able to read out data at lower-rates than O/NIR systems, which must be able to register the arrival times of individual photons.

1.5.2 KID Instruments which Incorporate Elements of the ASU LEKID Readout

At the time of writing, the ASU LEKID readout is a core system of several current and planned sub-mm/mm-wave instruments. Because each instrument has different system requirements, the elements of the readout (IF electronics, firmware, software) are used to varying extents in each system. Key instrumental parameters for each camera are listed in in Table 2, and the instruments are summarized below.

OLIMPO: OLIMPO (Masi et al. (2008); Paiella et al. (2019)) is a large aperture (2.6 m) balloon–borne mm-wave telescope which is designed to measure the Sunyaev-Zeldovich effect (SZ). It has four bands, centered at 150, 250, 350, and 460 GHz (2, 1.2, 0.86, 0.67 mm), which utilize 19, 37, 23 and 41 aluminum (Al) horn-coupled LEKIDs. The instrument performs spectrophotometry using a differential Fourier transform spectrometer (DFTS) which can be inserted into the optical path using a movable relay mirror.

In summer, 2018, OLIMPO had a successful flight from Svalbard, Norway. This was the first flight for KID detectors and their associated FPGA-based readout electronics. The camera incorporated the ASU ROACH2 electronics, firmware and software (two ROACH2 slices). In addition to the milestone which it represents for this detector technology, the data which was obtained during the flight regarding the performance of both the detectors and the readout system is invaluable to upcoming KID-based balloon-borne cameras, including BLAST-TNG.

An analysis of the camera's in-flight performance is presented in Masi et al. (2019).

	TolTEC	BLAST-TNG	MUSCAT	OLIMPO	SuperSpec
$N_{Det.}$	900, 1800, 3600	300, 700, 1500	1600	19, 37, 23, 41	300
N_{R2}	13	5	4	2	1
f_{res} (MHz)	∼500–1000	∼500–1000	∼500–1000	∼0–150	∼100–600
Bands (GHz)	150, 214.14, 272.5	600, 856, 1200	272.5	150, 250, 350, 460	185–315
Bands (mm)	2, 1.4, 1.1	0.500, 0.350, 0.250	1.1	2, 1.2, 0.86, 0.67	1.62–0.95
Mode[1]	BP	BP	BP	SP	SP
Couplng[2]	HC	HC	HC	HC	AC
Material	TiN/Ti/TiN	TiN/Ti/TiN	Al	Al	TiN
T_{base} (mK)	0.1	0.28	0.1	0.3	0.25
Site	LMT	LDB	LMT	LMT	Svalbard
Altitute (km)	4.6	∼35	4.6	∼38	4.6
Aperture (m)	50	2.5	50	2.6	50
Beam FWHM (″)	9.5, 6.3, 5	50, 41, 30	5.5	194, 116, 83, 63	NA
R ($\lambda/\Delta\lambda$)	NA	NA	NA	∼100	∼300
Platform	Ground	Balloon	Ground	Balloon	Ground
Total NEP W/$\sqrt{\text{Hz}}$ $\times 10^{-17}$	∼1[a]	∼5[b]	∼1[c]	∼1[d]	∼0.3–0.4[e]

(1) Horn coupled (HC) or antenna coupled (AC).
(2) Bolometric polarimetry (BP) or spectro-photometry (SP).
(a) http://toltec.astro.umass.edu/about
(b) See Chapter 4.
(c) Assuming comparable to TolTEC.
(d) Paiella et al. (2019)
(e) McGeehan et al. (2018)

Table 2. Camera parameters for systems currently using elements of the ASU LEKID readout.

The analysis examines the detector and readout system's noise performance, response to cosmic ray hits and thermal behavior during the flight. It determined that in all categories, the camera performed as expected– a finding that bodes very well for future KID cameras.

Figure 5 (Masi et al., 2019) shows a detector phase timestream containing a calibration lamp (cal lamp) chop (left) and detector noise PSDs (right) for the same channel, taken on the ground (red) and during the flight (blue). The cal lamp chop is a diagnostic used to measure the detector responsivities. The noise is lower in-flight by a factor of ∼2.5, due to the decreased thermal loading on the detectors from the colder environment at float altitude (∼38 km).

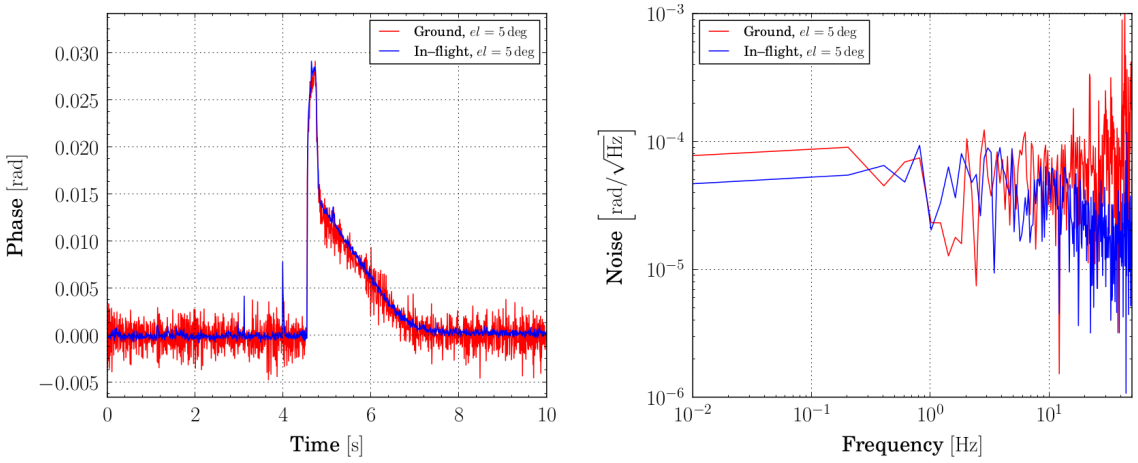

Figure 5. Data from a single detector taken during the 2018 OLIMPO balloon flight. The left frame shows the timestream of a cal lamp chop, taken on the ground (red) and in flight (blue). The right frame is the noise PSD for the same channel, on the ground (red) and in-flight (blue). The noise is lower in-flight by a factor of ∼2.5, due to lower thermal loading on the detectors. Image is from Masi et al. (2019).

TolTEC: TolTEC (Velazquez et al., 2016) is a mm-wave camera which will be a facility instrument on the LMT, with first-light planned for 2019. Its camera is based on ∼7,000 dual-polarization sensitive trilayer TiN/Ti/TiN horn-coupled LEKIDs,

similar to those used in BLAST-TNG, distributed between three bands centered at 2, 1.4, and 1.11 mm (50, 220, and 280 GHz) (Austermann et al., 2018). The readout uses the ROACH2 firmware, software and IF electronics which have been developed for BLAST-TNG (∼13 readout slices, versus the 5 used for BLAST-TNG). When coupled to the 50 m aperture of the LMT, TolTEC will have unprecedented spatial resolution in the mm-wave band (9.5″, 6.3″ and 5″). This will enable it to produce several unique data sets.

Soon after first-light, the TolTEC collaboration will conduct Four Public Legacy Surveys which will be made freely accessible to the public after ∼1–2 years. Some of first science surveys that are planned for TolTEC include[6]:

- 100 deg^2 Large Scale Structure Survey
- 1 deg^2 confusion limited Ultra-Deep Galaxy Survey
- 90 deg^2 Clouds-to-Cores Legacy Survey
- Legacy Surveys of Nearby Galaxies
- The Cluster Cosmology Survey
- Magnetic Fields and Star Formation
- The Interrogation of Comets

Mexico-UK Submillimeter Camera for Astronomy (MUSCAT): MUSCAT (Brien et al., 2018) is a LEKID camera which will be installed at the LMT along with TolTEC. It is based on four sub-arrays of horn-coupled LEKIDs (1,600 in total) which are cooled to 100 mK, and will observe in a single band centered at 1.1 mm. MUSCAT utilizes the BLAST-TNG ROACH2 firmware and IF electronics.

SuperSpec: SuperSpec (Shirokoff et al. (2012);Wheeler et al. (2018)) is an on-chip

[6]http://toltec.astro.umass.edu/science

filter-bank spectrometer based on antenna-coupled TiN KIDs which will have its first on-sky engineering run at the LMT in 2019. It will perform moderate resolution (R $\sim$ 300) spectroscopy over 190–310 GHz (1.57–0.98 mm) using 300 channels. SuperSpec will measure redshifted line emission (e.g., [CII]) in submillimeter galaxies (SMGs) at cosmological redshift z of 5–9, as well as CO emission in galaxies at lower z.

The SuperSpec camera incorporates the ROACH2 electronics, firmware and software developed for the BLAST-TNG readout. Figure 6 (McGeehan et al., 2018) shows a comparison between the fractional frequency noise S_{XX} of the system measured at a base temperature of 210 mK using the SuperSpec ROACH2 system and a Vector Network Analyzer (VNA) for tone combs consisting of 1 (blue and green) and 50 (red and pink) channels. The noise power spectral densities (PSDs) are shown for timestreams which correspond to the on and off-resonance states (see Chapter 4). Horizontal lines represent the expected photon-noise at the LMT sight. The system noise (including the LNA, readout, and detectors) which is measured with the ROACH2 system is comparable to that which is measured using VNA, for both the single and multitone combs.

Future Instruments: In addition to the experiments listed above, two recently funded balloon-borne cameras will incorporate elements of the ASU LEKID readout, with added improvements to the firmware and electronics which are made possible by the advent of ultra-high bandwidth FPGA boards (see Chapter 7). These are: The Experiment for Cryogenic Large-aperture Intensity Mapping (EXCLAIM) (Switzer, 2017), and the Terahertz Imager (TIM/STARFIRE) (Aguirre et al., 2018).

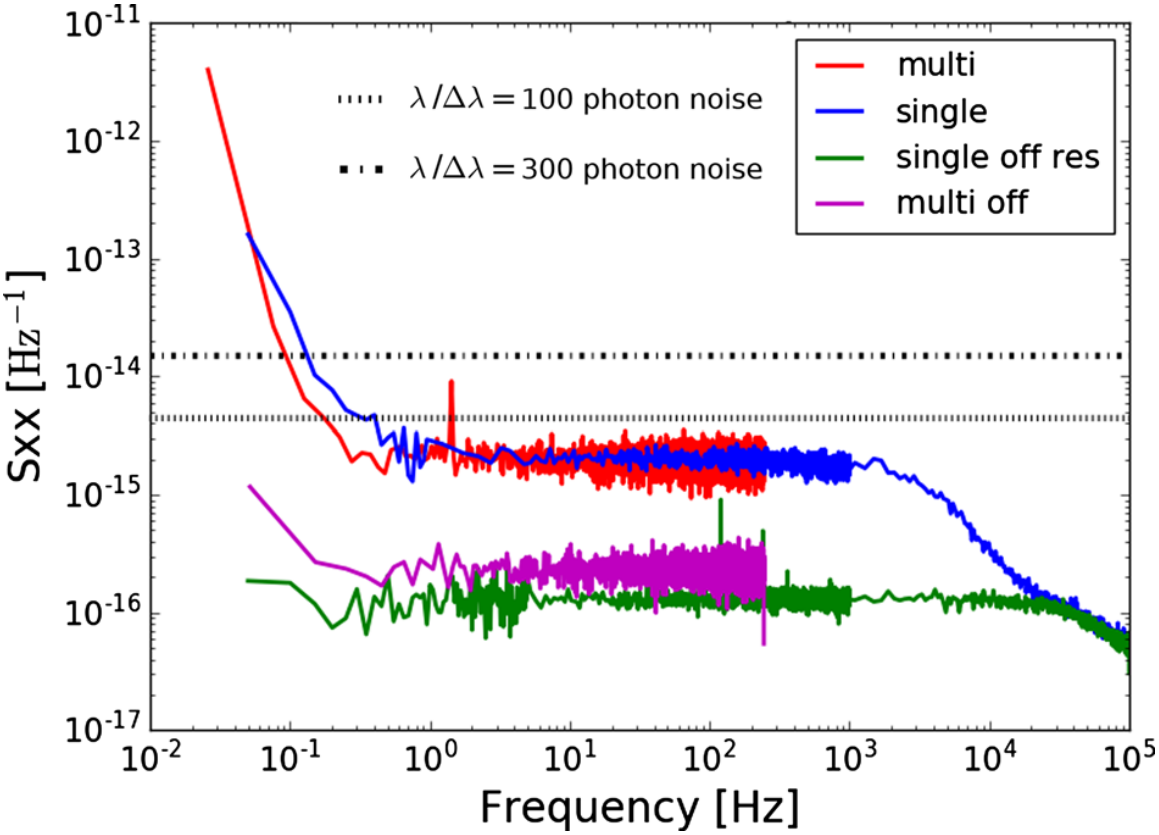

Figure 6. A comparison of the fractional frequency noise S_{XX} measured using the SuperSpec ROACH2 system and a VNA for single and multitone combs. The horizontal lines represent the expected photon-noise at the LMT sight. The ROACH2 system noise is comparable to that of the VNA, for both the single and multitone combs. Image taken from McGeehan et al. (2018).

1.6 Dissertation Outline

This dissertation is organized as follows:

- Chapter 2 describes a parametric model of a LEKID which can be used to simulate the behavior of both dark and optically loaded pixels, as well as to fit measurements of actual LEKIDs.
- Chapter 3 introduces the ASU LEKID readout, and details the system require-

ments, the architecture of the firmware, software and electronics, and the verification of the DSP pipeline and noise performance.

- Chapter 4 describes the instrumental characterization of the BLAST-TNG receiver and detector arrays which was conducted during the instrument integration at NASA's Columbia Scientific Balloon Facility (CSBF) in summer, 2018, as well as during pre-flight testing at NASA's Long Duration Balloon site (LDB) near McMurdo Station, Antarctica, during November through January, 2018/2019. The parametric LEKID model presented in Chapter 2 is applied to measured detector data in order to understand various characteristics of the detector arrays and the instrument as a whole.
- Chapter 5 presents an original data analysis of the B-field morphology in the Carina Nebula Complex (CNC, NGC 3372), using a map of the region taken during the BLASTPol 2012 flight. In that analysis, we present what are possibly the first estimates of the strength of B_{POS} over the inner $\sim$1.25 deg^2 of the CNC.
- Chapter 6 presents a target selection survey of five nearby external galaxies which will be mapped during the 2019/2020 BLAST-TNG flight.
- Chapter A contains an operator's manual for the *C*-based BLAST-TNG detector readout flight software, which includes a description of the in-flight detector readout strategy.
- Chapter 7 presents conclusions, with a discussion of how the technological landscape has evolved since this work began, and how it can benefit from recent developments in the area of high speed DSP.

Chapter 2

PARAMETRIC MODEL OF A LEKID

In the following, we develop a parametric model of a kinetic inductance detector (KID) which can be used to interpret or fit actual measurements produced by KID cameras or experiments. The model simulates both intrinsic circuit parameters and measured experimental values as a functions of base temperature $T_{\rm base}$ and absorbed optical power $P_{\rm abs}$. To be able to make accurate predictions, the model must incorporate the underlying device physics of KIDs. These physics are described within the framework of superconductivity.

The KID model contains several device properties which require initial values. These parameters, and the initial values which were used to produce the figures presented in the following sections, are listed in Table 3. The definition of each parameter can be found in the glossary section at the beginning of this document, or in the text below.

This chapter is organized as follows:

- Section 2.1 provides a brief overview of relevant concepts from superconductivity.
- Section 2.2 introduces the kinetic inductance (KI).
- Section 2.3 describes the temperature and power dependence of key instrinsic and empirically determined values.
- Section 2.4 defines the KID responsivity to changes in T_c (R_T) and $P_{\rm abs}$ (R_P).
- Section 2.5 defines the instrumental sensitivity, and describes the instrumental

noise budget for a KID instrument, detailing how each of the dominant noise sources contributes to the total noise-equivalent power (NEP).

Parameter	250 μm	350 μm	500 μm
$w_{\rm L}$ (μm)	8	10	4
$l_{\rm L}$ (μm)	600	900	670
$t_{\rm L}$ (μm)	18	18	20
N_0 ($\rm ev^{-1}\mu m^{-3}$)	9.8×10^{10}	9.5×10^{10}	3×10^{10}
T_c (K)	1.46	1.55	1.30
$\tau_{\rm rec,0}$ (s)	3×10^{-6}	3×10^{-6}	3×10^{-6}
$T_{\rm LNA}$ (K)	6	6	6
$Q_{\rm loss}$	10^7	10^7	10^7
Q_c	4×10^4	4×10^4	4×10^4
σ_n $(\Omega$ cm$)^{-1}$	7000	7000	7000
$T_{\rm base}$ (mK)	290	290	290
$P_{\rm opt}$ (dBm)	-90	-90	-90
$\epsilon_{\rm opt}$	1	1	1
$\eta_{\rm det}$	1	1	1

Table 3. Initial values of device parameters used in the KID model.

2.1 Cooper Pairs and Quasiparticles

A superconductor is a conductor whose DC resistance R disappears when it is cooled below a material-specific temperature known as the critical temperature T_c. The disappearance of R is accompanied by the expulsion of magnetic field (B-field) lines from inside the material, which increases the magnetic flux at the surface. The latter phenomenon is known as the Meissner effect.

That R goes to zero when a superconductor is cooled below T_c implies (by Ohm's Law) that an infinite DC current can be generated by the application of a finite voltage. However, superconductors have non-zero impedance, because the reactance

(capacitive and inductive) does not disappear along with R. The inductance of the superconductor prevents the instantaneous formation of an infinite current. Beyond a maximum current known as the critical current I_c, the superconductor reverts to its normal state. Similarly, the superconductivity can be destroyed by applying a an external B-field whose strength exceeds the material's critical field B_c.

The underlying physics of superconductivity was first described by Bardeen-Cooper-Schrieffer (BCS) theory (Bardeen et al., 1957). In BCS, when a superconductor is cooled below T_c a supercurrent is created that consists of paired normal charge carriers (holes and electrons) called Cooper pairs (CPs). The remaining charge carriers which have not formed pairs are referred to as quasiparticles (QP). QP have the same spin and charge as normal electrons, but have an effective mass $m^\star$ which differs from the rest mass of an electron m_e.

CPs are modeled as quasi-bosons with mass of $2m_e$ and charge of $2e$. The pairs arise from electron-phonon interactions in the crystal lattice of the conductor. Cooling the material decreases the thermal vibrations in the cation lattice. A single QP moving through the lattice attracts several cations, creating a local overdensity of positive charge, or a phonon, which propagates along with the QP. A second QP, having opposite spin to the first, becomes attracted to the phonon, forming the second member of the CP. The pairs have coherence lengths which are typically several orders of magnitude greater than the lattice spacing.

From an energy standpoint, the formation of supercurrent is the consequence of the creation of an energy band gap which forms above and below the material's Fermi energy:

$$E_g(T=0) = \Delta_0 = \frac{3.52}{2} k T_c \tag{2.1}$$

where k is Boltzmann's constant, and Δ_0 is the value of the band gap at absolute-zero. The Fermi energy, E_F, is the difference in the energy between the highest and lowest states which are occupied by the fermions (electrons and holes) at $T = 0$ K (Kittel and Kroemer, 1998):

$$E_F = \frac{\hbar^2}{2m^\star}\left(\frac{3\pi^2 N_0}{\Sigma}\right) \tag{2.2}$$

where Σ is the volume of the conductor, and N_0 is the single-spin density of states at the Fermi energy. During cooling, the rate of the decline in R with T is material dependent. For example, the decline is much faster in aluminum (Al) devices than in titanium nitride (TiN) devices (Mauskopf, 2018). The band gap energy Δ_0 is equivalent to the CP binding energy, and is typically of $\mathcal{O}(10^{-3})$ eV. This band gap is around three orders of magnitude smaller than that of semiconductors, including CCDs, which illustrates the potential sensitivity advantages of using superconductors as detectors.

Any energy greater than $2\Delta_0$ which is absorbed by a superconductor will unbind, or break apart a CP into its constituent QP. This breaking apart of CPs can be measured to estimate the total power which is absorbed by a device. In the case of a KID, the splitting of CPs is measured by proxy, using the corresponding change in KI (see Section 2.2). In Section 2.1.1 we will derive an equation which contains the band gap energy's dependence on T_c and $P_{\rm abs}$.

2.1.1 Quasiparticle Generation and Recombination

In a KID, QP are generated and recombine (CPs dissociate and reform) due to the absorption of thermally generated phonons in the material. These phonons may

be produced by temperature fluctuations in $T_{\rm base}$, or from power dissipated by the microwave carrier tone which is used to readout the device. QP are also generated by the absorption of photons with frequency $\nu \geq 2\Delta_0/h$. In the superconducting materials used to fabricate sub-mm/mm-wave KID devices, typical QP lifetimes (recombination times) τ_{qp}, are $\mathcal{O}(10)$ µs. An empirical expression for the QP lifetime which has been found to apply to several materials is (Zmuidzinas, 2012):

$$\tau_{qp} = \frac{\tau_{max}}{1 + n_{qp}(T)/n^{\star}} \tag{2.3}$$

where $n_{qp}(T)$ is the number density of thermally generated QP, and $n^{\star} \sim 100$ µs is the cross-over number density. The QP which are generated by thermal phonon absorption and by pair-breaking photons can be modeled as two distinct populations.

The total number density of QP is (Mauskopf, 2018):

$$\begin{aligned} n_{qp,\rm tot} &= n_{qp}(T) + n_{qp}(P) \\ &= 4N_0 \int_0^{\infty} \frac{dE}{1 + e^{E/kT_c}} \\ &= 4N_0 kT_c \ln(2) \end{aligned} \tag{2.4}$$

To separate out the effects of varying $T_{\rm base}$ and $P_{\rm abs}$, the thermal and optical QP number densities must be written explicitly. In the low temperature limit ($T \ll T_c$) the number density of thermal QP is

$$\begin{aligned} n_{qp}(T) &\simeq 2N_0\sqrt{2\pi kT\Delta_0}e^{-\Delta_0/kT} \\ &= \frac{N_{qp}(T)}{\Sigma} \end{aligned} \tag{2.5}$$

where Σ is the volume of the inductor, and $N_{qp}(T)$ is the total number of thermally generated QP. Figure 7 shows $N_{qp}(T)$ for a range of base temperatures.

The number of QP produced per absorbed photon of frequency ν can be found by taking the ratio of the photon energy to the band gap energy:

$$N_{qp}(\nu)_{,\mathrm{per}} = \frac{\eta_{\mathrm{det}} h\nu}{\Delta_0} \tag{2.6}$$

where η_{det} is the detector quantum efficiency (typically ~80%). The number of broken CPs per absorbed photon would be half of Equation 2.6. The corresponding number of optically generated QP is:

$$\begin{aligned} N_{qp}(P) &= \frac{\eta_{\mathrm{det}} \epsilon_{\mathrm{opt}} N_{qp}(\nu)_{,\mathrm{per}} P_{\mathrm{abs}} \tau_{qp}}{h\nu} \\ &= \frac{N_{qp}(\nu)_{,\mathrm{per}} P_{\mathrm{opt}} \tau_{qp}}{h\nu} \end{aligned} \tag{2.7}$$

The ratio $N_{qp}(P)/N_{qp}(T = 290)$ is shown in Figure 8 for a range of P_{abs}, with $T_{\mathrm{base}} = 290$ mK. Using Equations 2.5 and 2.7 to express the total number of QP $N_{qp,,\mathrm{tot}}$, the band gap energy can now be expressed as a function of both T_c and P_{abs}:

$$\Delta(T, P) = \frac{\Delta_0}{2}\left(1 + \sqrt{1 - \frac{2N_{qp,\mathrm{tot}}}{\Sigma N_0 \Delta_0}}\right) \tag{2.8}$$

2.1.2 The Two-Fluid Model

Several of the KID parameters discussed in the following sections depend on the conductivity σ of the superconductor. Here we describe the two-fluid model of conductivity (Glover and Tinkham (1957); Mattis and Bardeen (1958)), which relates the Drude conductivity σ_d in a normal conductor to the CPs and QP of a superconductor.

In a conductor, the equation of motion for a charge carrier is:

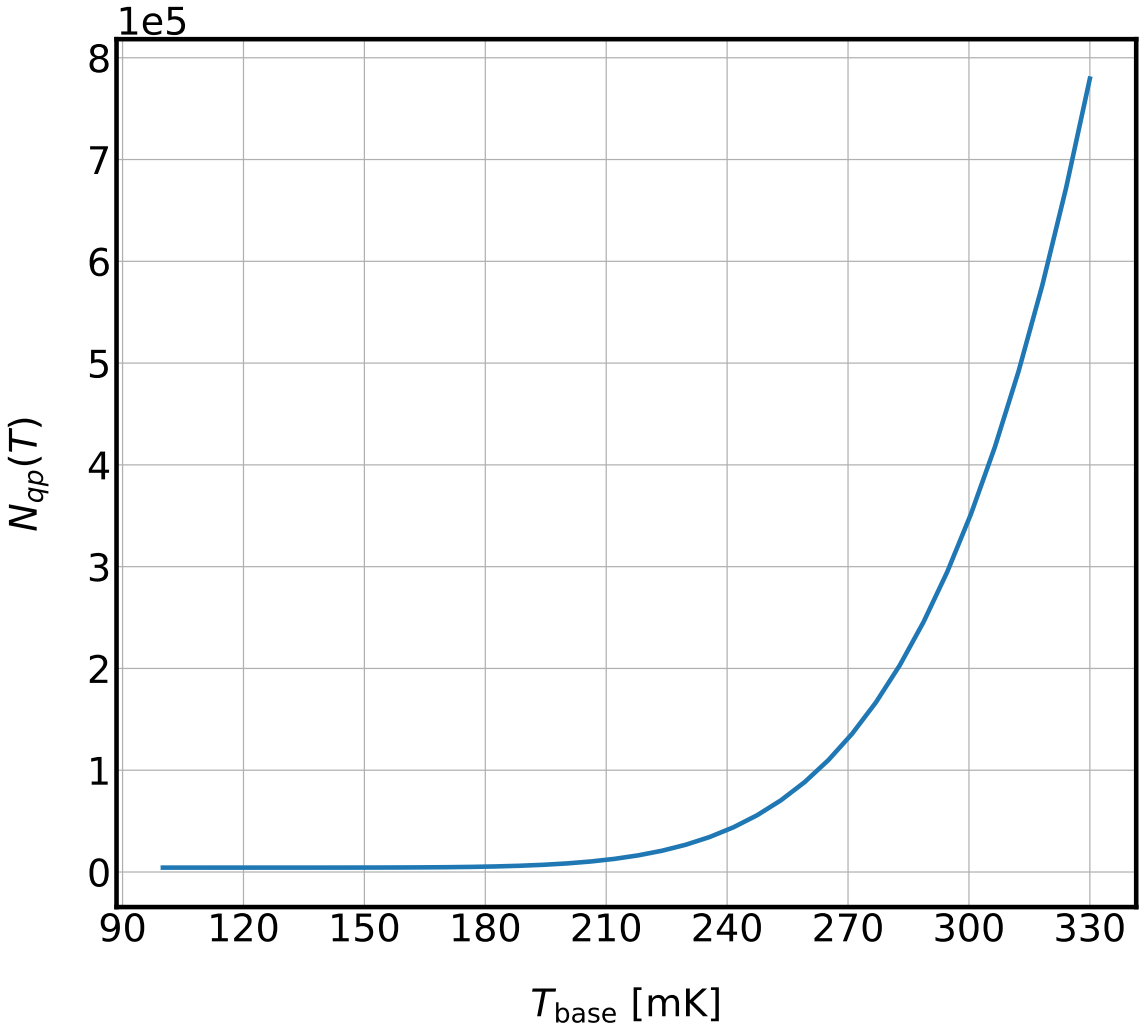

Figure 7. The total number of thermally generated QP as a function of base temperature.

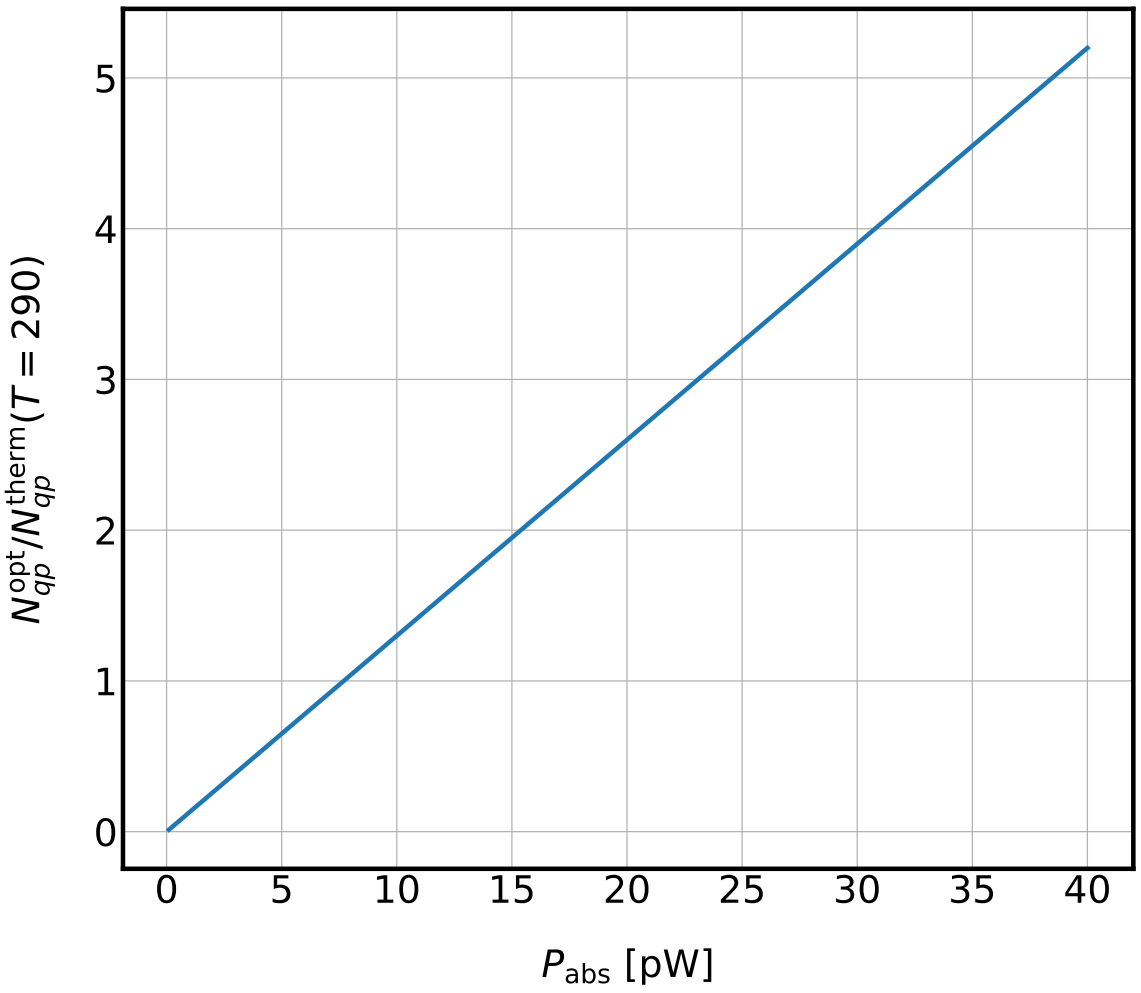

Figure 8. The ratio of the number of optically generated QP to thermally generated QP, $N_{qp,\mathrm{opt}}/N_{qp,\mathrm{therm}}$, as a function of P_{abs}. T_{base} is fixed at 290 mK.

$$m^{\star}\dot{v} = -eE - \frac{m^{\star}v}{\tau_s} \tag{2.9}$$

where ω is the frequency of a time-varying electromotive force (EMF) being applied to the circuit, $m^{\star}$ is the effective mass of the charge carrier, and $\boldsymbol{v}$ is their drift velocity ($\propto e^{j\omega\tau_s}$). The second term on the right-hand side of the equation is a damping term which accounts for charges scattering off of the lattice, with τ_s being the

characteristic time between collisions. Ohm's Law provides the relation between conductivity and current density: $J = \sigma_d E$, with $J = n_e e v$.

Inserting these relations into the equation of motion, and solving for σ_d yields the Drude conductivity:

$$\begin{aligned}\sigma_d(\omega) &= \frac{n_e e^2 \tau_s}{m_e} \frac{1}{1 + j\omega\tau} \\ &= \frac{\sigma_n}{1 + j\omega\tau_s} \\ &= \frac{\sigma_n}{1 + \omega^2 {\tau_s}^2} - j\frac{\sigma_n \tau_s}{1 + \omega^2 {\tau_s}^2} \\ &= \sigma_1 - j\sigma_2\end{aligned} \tag{2.10}$$

where σ_n is the normal conductivity, $\sigma_n = n_e e^2 \tau_s / m^\star$. In a normal metal, τ_s is small enough so that $\sigma_d \approx \sigma_n$ is negligible at frequencies ω below the optical range. For $kT \ll \Delta_0$ and $\hbar\omega \ll \Delta_0$, the real and imaginary conductivities can be written as (Gao, 2008):

$$\sigma_1(n_{qp}, T)_{\text{therm}} = \frac{4\sigma_n \Delta_0}{\hbar\omega} e^{-\frac{\Delta_0 - \mu^*}{kT}} \sinh(\xi) K_0(\xi) \tag{2.11}$$

$$\sigma_2(n_{qp}, T)_{\text{therm}} = \frac{\sigma_n \Delta_0}{\hbar\omega} \left[1 - \sqrt{\frac{2\pi kT}{\Delta_0}} e^{-\frac{\Delta_0}{kT}} - 2e^{\frac{\Delta_0}{kT}} e^{-\xi} I_0 \right] \tag{2.12}$$

$$\sigma_1(n_{qp}, T)_{\text{opt}} = \frac{2\Delta_0}{\hbar\omega} \frac{n_{qp}}{2N_0\sqrt{\pi kT\Delta_0}} \sinh(\xi) K_0(\xi) \tag{2.13}$$

$$\sigma_2(n_{qp}, T)_{\text{opt}} = \frac{\pi\Delta_0}{\hbar\omega} \left[1 - \frac{n_{qp}}{2N_0\Delta_0} \left(1 + \sqrt{\frac{2\Delta_0}{\pi kT}} e^{-\xi} I_0 \right) \right] \tag{2.14}$$

Figure 9 shows the ratio of σ_1 to σ_2 as a function of T_{base} ($P_{\text{abs}} = 0$). The ratio asymptotically decays to zero for $T \gtrsim 250$ mK. The ratio of σ_1 to σ_2 as a function of

P_{abs} (T_{base} = 290 mK) is shown in Figure 10. The ratio falls below one for $P_{\mathrm{abs}} \gtrsim$ 12 pW.

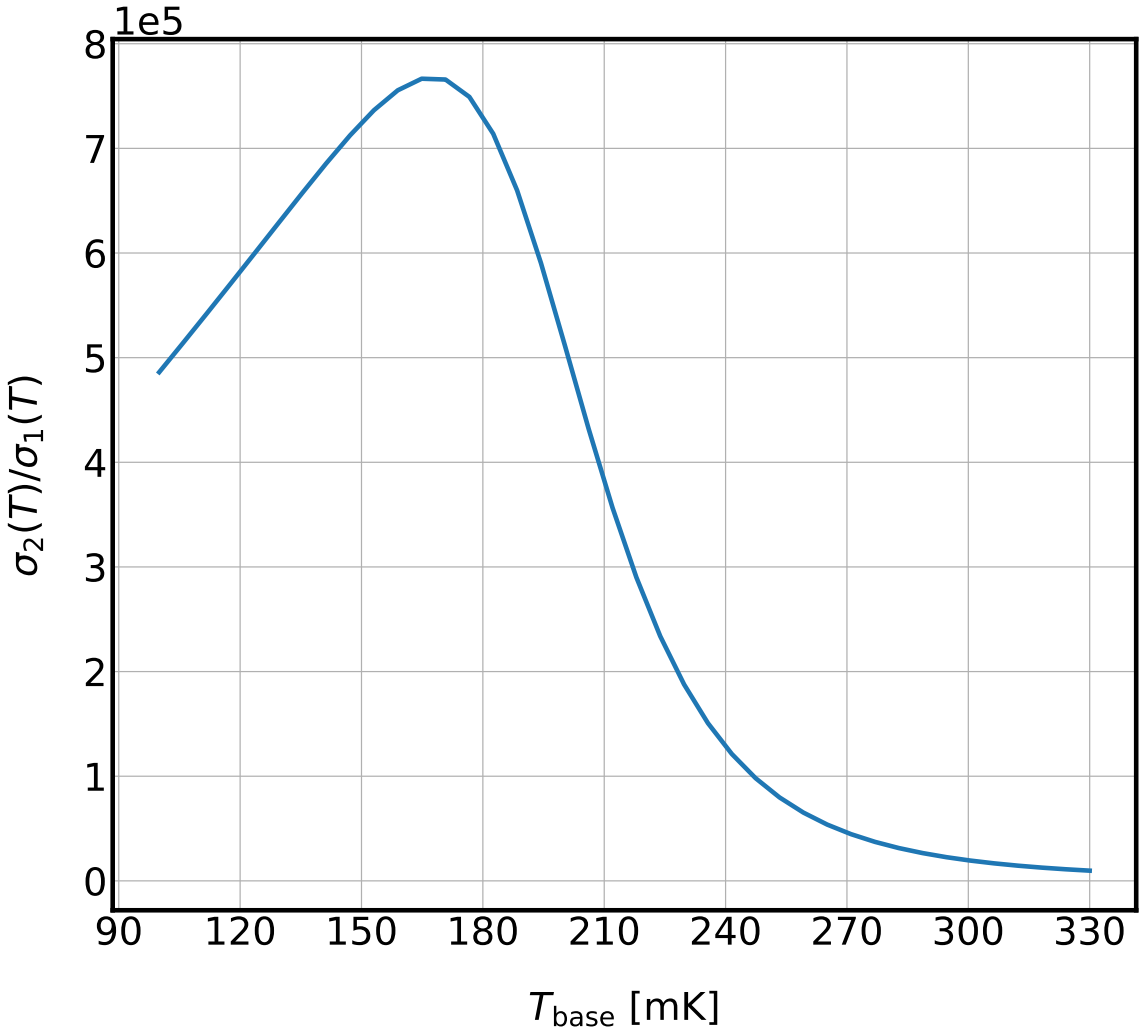

Figure 9. The ratio of σ_1 to σ_2 as a function of base temperature.

where μ^* is an effective chemical potential which is added to the Fermi distribution function to account for QP generation due to pair breaking (for a derivation, see Owen and Scalapino (1972); Gao (2008)).

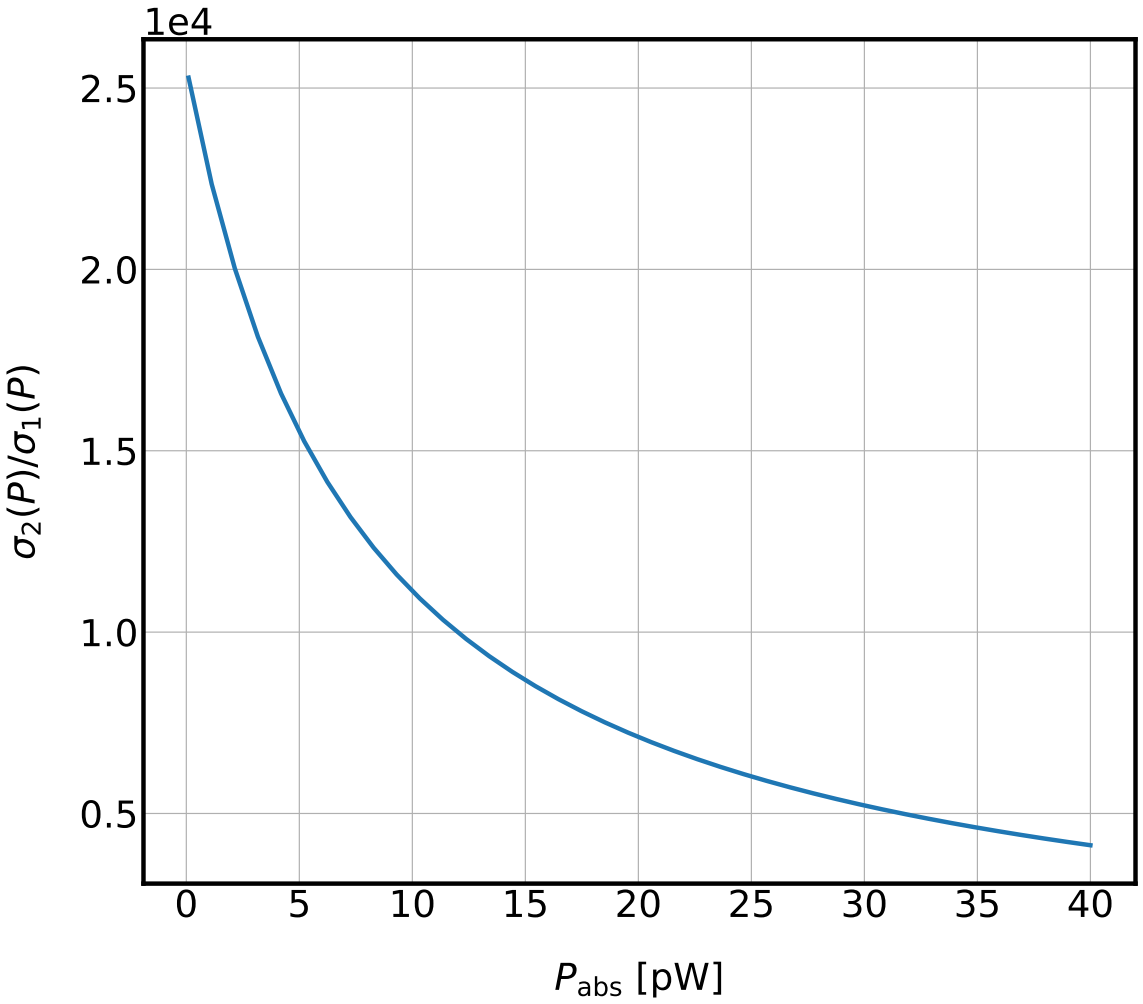

Figure 10. The ratio of σ_1 to σ_2 as a function of absorbed optical power, at a base temperature of 290 mK.

$$\mu^*_{\mathrm{therm}} = kT \ln\left(\frac{n_{qp}}{2N_0\sqrt{2}\pi kT\Delta_0}\right) + \Delta_0 \tag{2.15}$$

$$\mu^*_{\mathrm{opt}} = kT \ln\left(\frac{{N_{qp}}^{tot}}{2VN_0\sqrt{2\pi kT\Delta}}\right) + \Delta \tag{2.16}$$

where ω is the KID resonant frequency, σ_n is the normal conductivity of the material, $\xi = \hbar\omega/2kT$, K_0 is the Modified Bessel function of the second kind of order zero, and I_0 is the Modified Bessel function of order zero.

2.2 Kinetic Inductance

When a voltage is applied to a closed circuit, current does not start to flow instantaneously. The reason for this is that some energy from the applied voltage must first be stored in a B-field, via the inductance L. Inductance is defined as the ratio of the applied voltage to the rate of change of the current: $L = \frac{V}{dI/dt}$ [Ω/Hz]. The energy stored in the inductance is $E_L = \frac{1}{2}LI^2$. The magnetic inductance L_m of a superconducting film depends solely on its geometric parameters: Length l, width w, thickness t.

In all conductors there is an additional form of inductance, known as the kinetic inductance (KI) L_k. The total inductance is the sum of the magnetic (or geometric) and kinetic inductances:

$$L_{\text{tot}} = L_m + L_k \tag{2.17}$$

It's useful to define a KI fraction, as:

$$\alpha = \frac{L_k}{L_{\text{tot}}} \tag{2.18}$$

The KI arises from the inertia of the CPs. As charge carriers scatter (or break apart, in the case of CPs), the kinetic energy (velocity) of the remaining charge carriers must increase in order to keep the current constant. This increase in velocity cannot happen instantaneously. Therefore, if an AC current is circulated through a superconducting circuit, and the KI is increased, the current undergoes a phase shift. There is a maximum velocity imposed upon the charge carriers, corresponding to the critical, or pair-breaking current of the material, which is described in the Ginzberg-Landau theory (GL) of superconductivity (see Section 2.4.3) (Tinkham, 2004).

From an energy standpoint, an expression for the KI can be derived by equating the energy in the KI to the kinetic energy of a CP:

$$\begin{aligned}\frac{1}{2}L_k I^2 &= \frac{1}{2}(2m_e)v_{cp}^2 N_{cp} \\ &= m_e^2 v_{cp}^2 n_{cp} l A\end{aligned} \tag{2.19}$$

where $A = wlt$ is the cross-sectional area of the transmission line (TL), I_s is the supercurrent, v_{cp} is the CP velocity and N_{cp} is the total number of CPs. Defining the supercurrent as

$$I_s = \frac{1}{2}n_{cp}(2e)v_{cp}A \tag{2.20}$$

the KI per unit length is

$$\mathcal{L}_k = \frac{m_e}{n_{cp}e^2 A} \tag{2.21}$$

The KI can also be understood by considering the impedance. Because CPs are collisionless, their scattering time τ_s is infinite. Taking the limit of Equation 2.10 as τ_s goes to infinity:

$$\begin{aligned}\lim_{\tau_s \to \infty} \sigma_d(\omega) &= -j\frac{n_e e^2}{m_e \omega} \\ &= \sigma_2\end{aligned} \tag{2.22}$$

Therefore, the KI per unit length is

$$\mathcal{L}_k(\omega) = \frac{1}{\omega \sigma_2 A} \tag{2.23}$$

The impedance of the superconducting film can be written as a function of the conductivity:

$$
\begin{aligned}
Z(\omega) &= \frac{1}{t\sigma(\omega)} \\
&= \frac{1}{t}\left(\frac{\sigma_1}{{\sigma_2}^2} - j{\sigma_2}^{-1}\right) \\
&= \mathcal{R} + j\omega\mathscr{L}_{k,sq}
\end{aligned}
\tag{2.24}
$$

where $\mathcal{R}$ is the resistance per square.

Then,

$$
\mathscr{L}_{k,sq}(\omega) = \frac{1}{\sigma_2 \omega t} \tag{2.25}
$$

Using the fact that in the low-frequency limit ($\hbar\omega \ll kT$), Mattis-Bardeen (MB) theory allows the complex conductivity to be expressed as (Mauskopf (2018); Annunziata et al. (2010)):

$$
\sigma_2 = \frac{\sigma_n \pi \Delta}{\hbar\omega} \tanh\left(\frac{\Delta_0}{2kT}\right) \tag{2.26}
$$

and the sheet resistance of the film in the normal state is $R_{sq} = 1/\sigma_n$t, the KI per square becomes:

$$
\mathscr{L}_{k,sq}(T,P) = \frac{\hbar R_{sq}}{\pi\Delta(T,P)} \tag{2.27}
$$

Equation 2.27 illustrates the fact that materials with higher normal state sheet resistivity, and/or smaller band gaps, have higher intrinsic KI fractions α. Figures 11 and 12 show α as a function of T_{base} and P_{abs}. The KI fraction is expontial in T_{base}, but linear in P_{abs}.

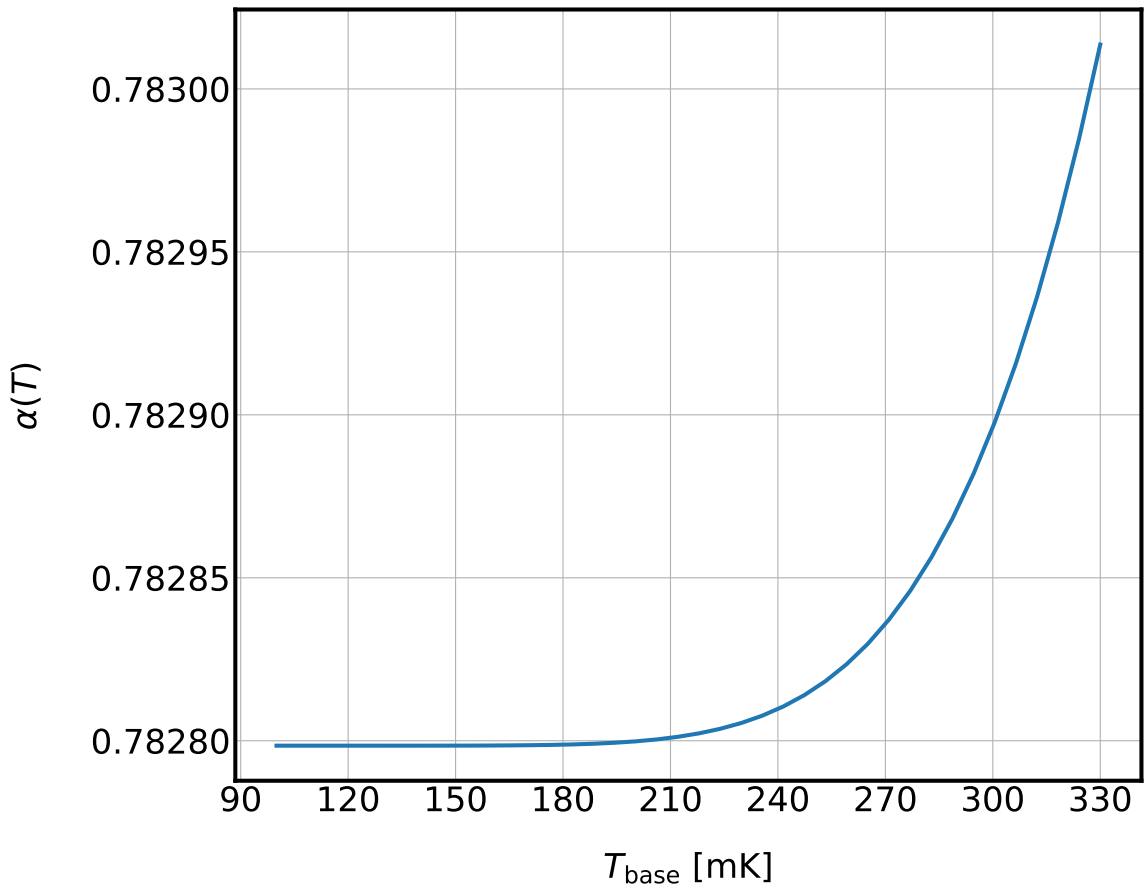

Figure 11. The kinetic inductance ratio α as a function of base temperature.

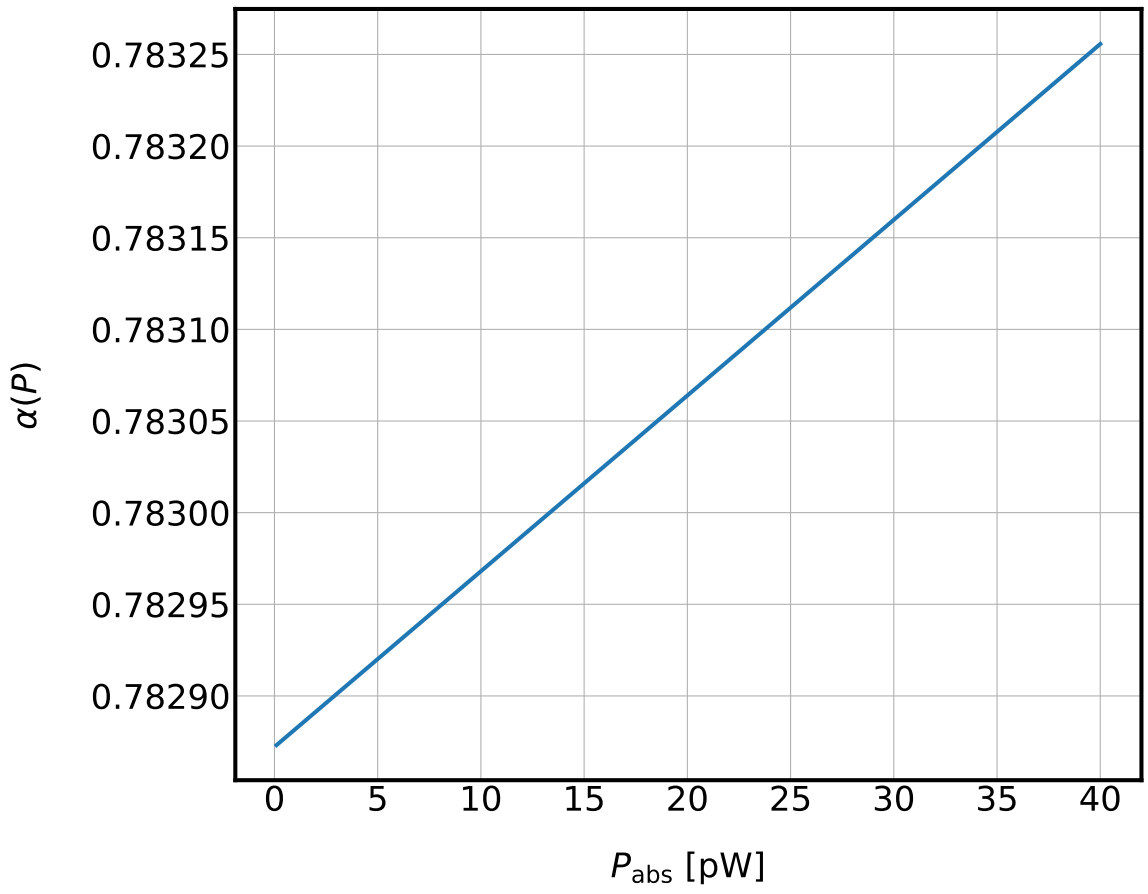

Figure 12. The kinetic inductance fraction α as a function of optical power.

2.3 Principle LEKID Parameters

In the previous sections, several superconducting device properties were derived as functions of base temperature $T_{\rm base}$ and absorbed optical power $P_{\rm abs}$. Now, we introduce a circuit model for a lumped-element KID (LEKID), and derive the temperature and power dependence of several intrinsic and empirical (determined from measurements) parameters.

2.3.1 Intrinsic and Empirical Parameters

A LEKID is a capacitively-coupled lumped-element resonant circuit (tank circuit). It has five intrinsic circuit parameters. These are:

$L_{\rm tot}$ The total resonator inductance

C_r The resonator capacitance

C_c The coupling capacitance

$R_{\rm eff}$ An effective resistance, in series with $L_{\rm tot}$, which accounts for the real part of the resonator impedance

Z_0 The feedline impedance

Figure 13 shows a schematic which illustrates how a LEKID is constructed from the basic elements listed above.

There are four principle empirical parameters:

Q_r Resonator quality factor

Q_c Coupling quality factor

ω_0 Resonant frequency

ϵ_{a} A parameter which accounts for a frequency dependent asymmetry in resonator impedance

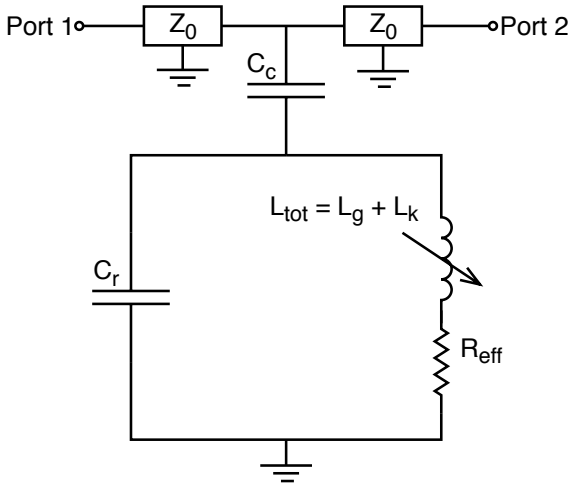

Figure 13. A schematic of a single LEKID.

The resonant frequency is set by the total inductance and resonator capacitance:

$$\omega_0 = \frac{1}{\sqrt{\mathcal{L}_{\mathrm{tot}}\mathcal{C}_r}} \tag{2.28}$$

From the perspective of a microwave carrier tone which is input to the feedline at frequency $\omega_c = 2\pi f_c$, the LEKID acts as a frequency dependent shunt impedance to ground. The resonant frequency of the LEKID is changed by modulating the KI, with a combination of changes in absorbed optical power, base temperature and microwave probe tone power. Taking the derivative of Equation 2.28 with respect to $\mathcal{L}_k$, and defining the resonator capacitance as $C_r = \frac{1}{L_{\mathrm{tot}}\omega_0^2}$,

$$\frac{df_0}{d\mathcal{L}_k} = \frac{-\alpha f_0}{2\mathcal{L}_k} \tag{2.29}$$

The asymmetry parameter is defined as $\epsilon_{\mathrm{a}} = \frac{C_r}{C_c Q_i}$ where Q_i is the internal quality factor (Mauskopf, 2018). The internal quality factor is itself defined as:

$$Q_i = \frac{\Im(Z_r)}{\Re(Z_r)} = \frac{\omega_c L_{\rm tot}}{R_{\rm eff}} = \frac{\omega_c E_r}{P_{\rm diss}} \tag{2.30}$$

where E_r is the stored (internal) energy of the resonator, and $P_{\rm diss}$ is the dissipated power, or energy dissipated per cycle of the carrier tone at frequency ω_c. The coupling quality factor is (see Barry (2014)):

$$Q_c = \frac{2(C_r + C_c)}{\omega_0 C_c^2 Z_0} \tag{2.31}$$

The measured resonator quality factor, Q_r, is related to the internal and coupling quality factors as:

$$\frac{1}{Q_r} = \frac{1}{Q_i} + \frac{1}{Q_c} + \frac{1}{Q_{\rm loss}} \tag{2.32}$$

where $Q_{\rm loss}$ is an external loss factor, typically of $\mathcal{O}(10^6)$, which accounts for parasitic loss, including parasitic capacitance and inductance.

Figures 14 and 15 show Q_r as a function of $T_{\rm base}$ and $P_{\rm abs}$. While $Q_r(P)$ decreases monotonically with increasing absorbed optical power, $Q_r(T)$ does not. Starting at $T_{\rm base} = 90$ mK, $Q_r(T)$ first increases with base temperature, reaching a maximum at $T_{\rm base} \sim 180$ mK, before decreasing almost linearly.

The empirically determined resonator parameters are determined by measuring the phase and amplitude of the carrier tone after it passes through the LEKID circuit. The resonator impedance Z_r has a frequency dependence which can be expressed as:

$$Z_r(\omega) \simeq Z_0 \left[\frac{Q_c}{2Q_i} + jQ_c \delta x \right] \frac{1}{1 + j\epsilon_{\rm a}} \tag{2.33}$$

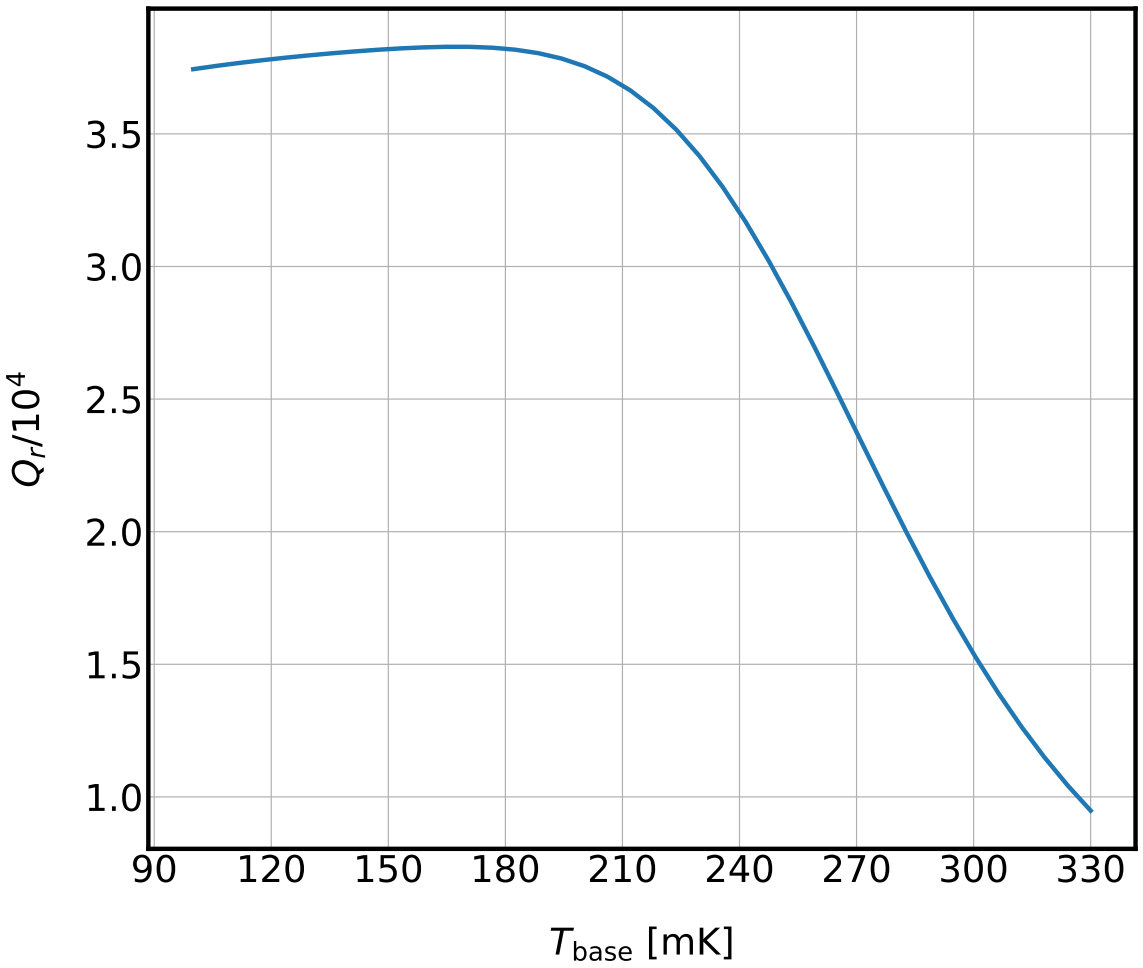

Figure 14. Simulated Q_r as a function of base temperature.

where $\delta x = \frac{\omega - \omega_0}{\omega_0}$. The forward transmission coefficient, S_{21}, of the carrier tone expressed as a function of frequency, base temperature and optical power (the dependence of the latter two parameters is implicit) is:

$$\begin{aligned} S_{21}(\omega, T, P) &= 1 - \frac{1}{1 + 2Z_r/Z_0} \\ &= 1 - \frac{1 + j\epsilon_{\mathrm{a}}}{1 + j\epsilon_{\mathrm{a}}\frac{Q_r}{Q_c}} \frac{Q_r}{Q_c} \left[\frac{1}{1 + 2jQ_r\delta x/(1 + j\epsilon_{\mathrm{a}}\frac{Q_r}{Q_c})} \right] \end{aligned} \tag{2.34}$$

for $\epsilon_{\mathrm{a}} \ll 1$, this reduces to:

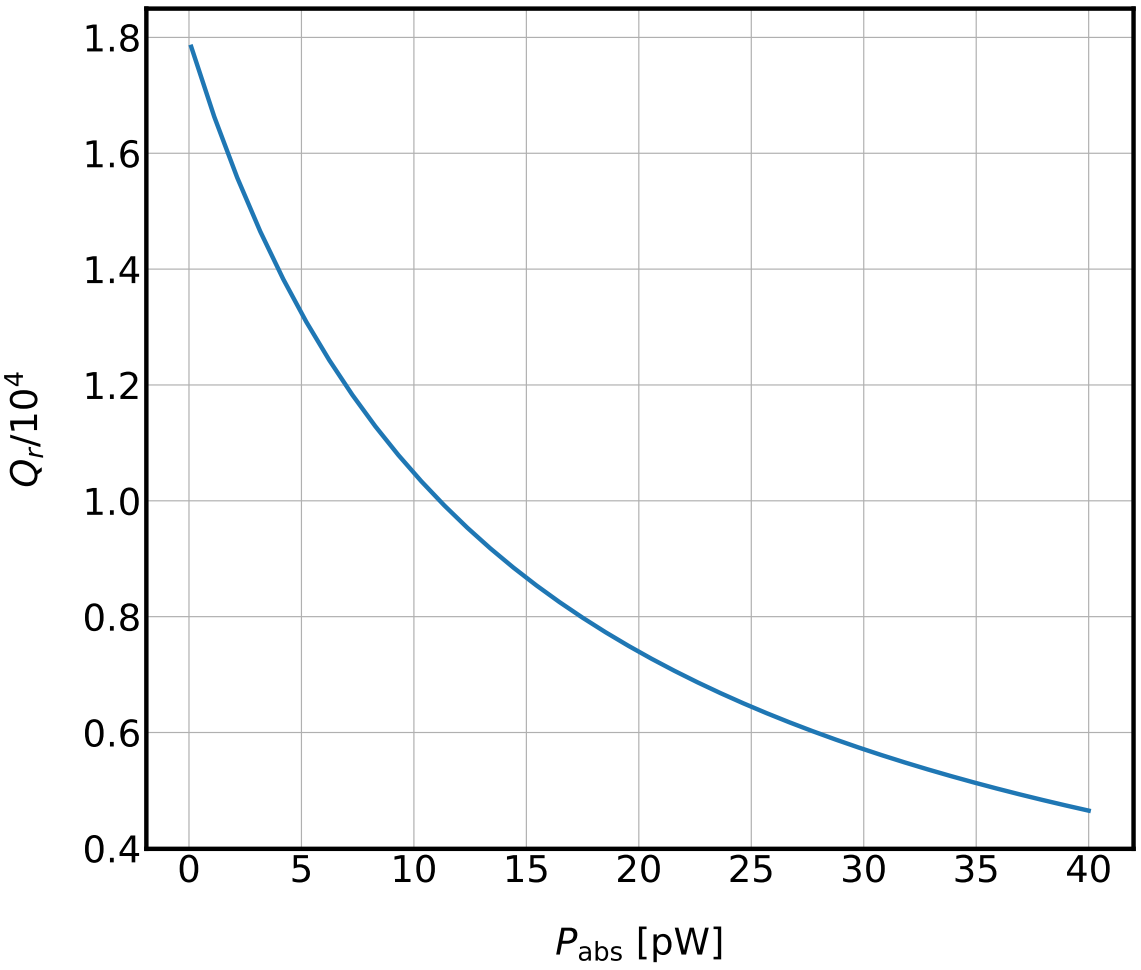

Figure 15. Simulated Q_r as a function of optical power.

$$S_{21}(\omega, T, P) \simeq 1 - \frac{Q_r}{Q_c} \frac{1}{1 + 2jQ_r\delta x} \tag{2.35}$$

Figure 16 shows an example $|S_{21}|(\omega)$ for a wide range (-3–3) of ϵ_a values. The temperature and power dependence of S_{21} are shown in Figures 17 and 18, for a

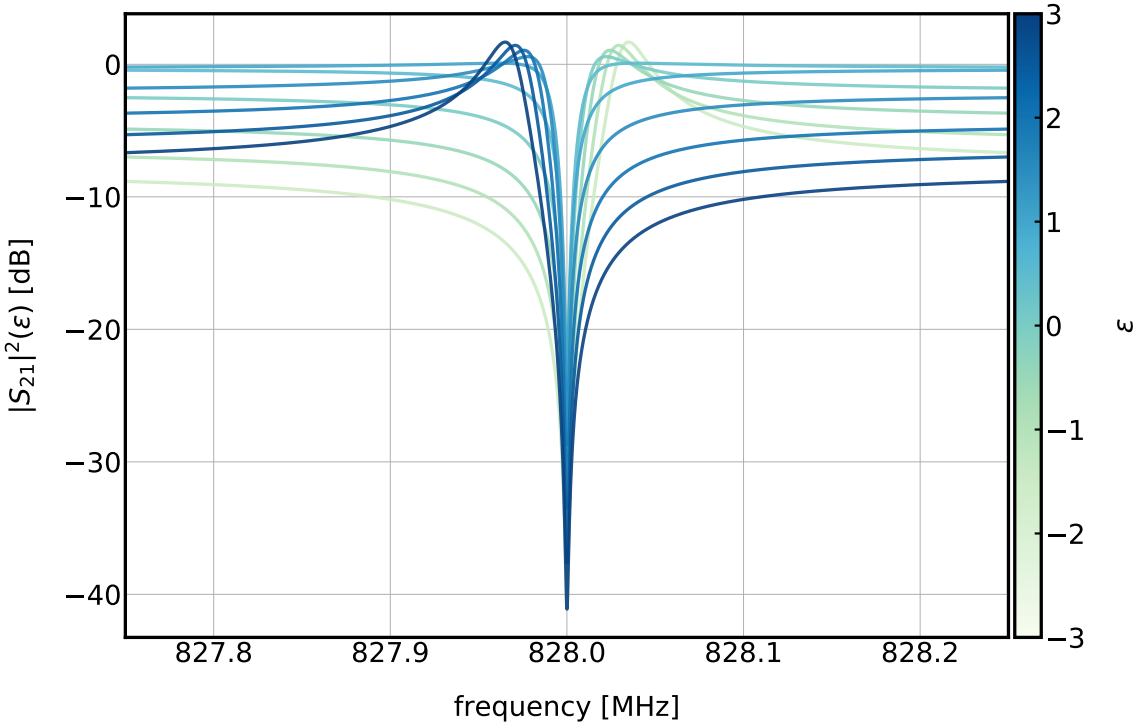

Figure 16. $|S_{21}|(\omega)$ shown for a range of ϵ_a.

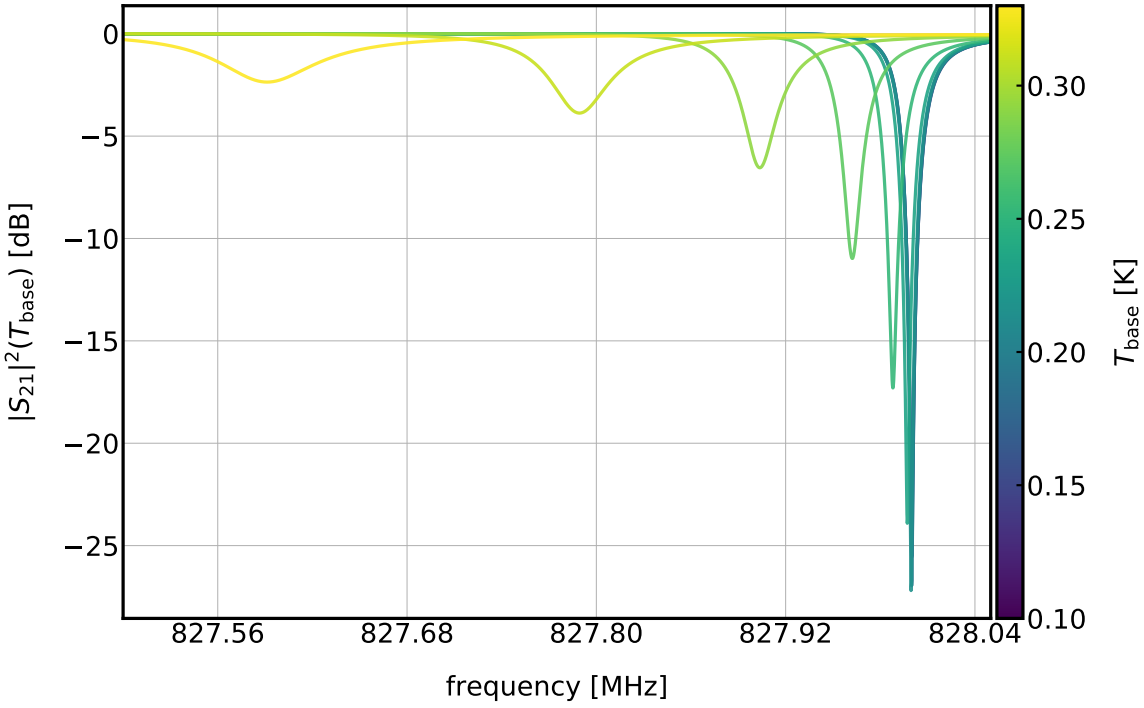

Figure 17. Simulated $|S_{21}|(\omega)$ as a function of base temperature.

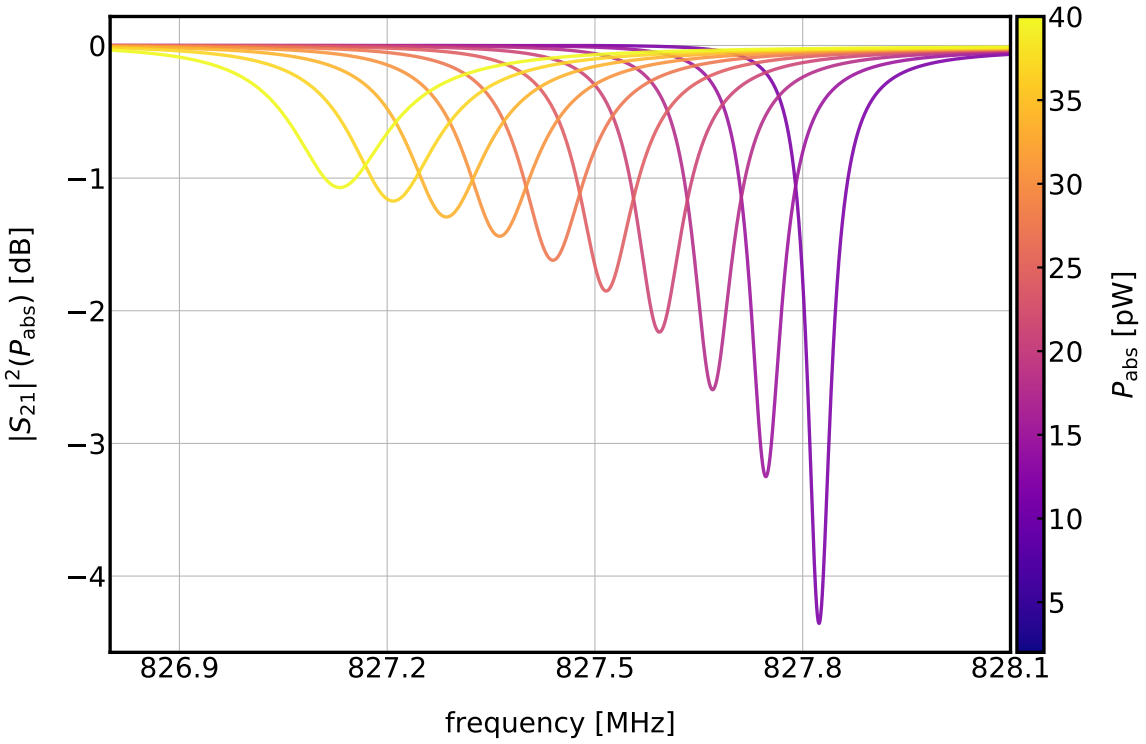

Figure 18. Simulated $|S_{21}|(\omega)$ as a function of optical power.

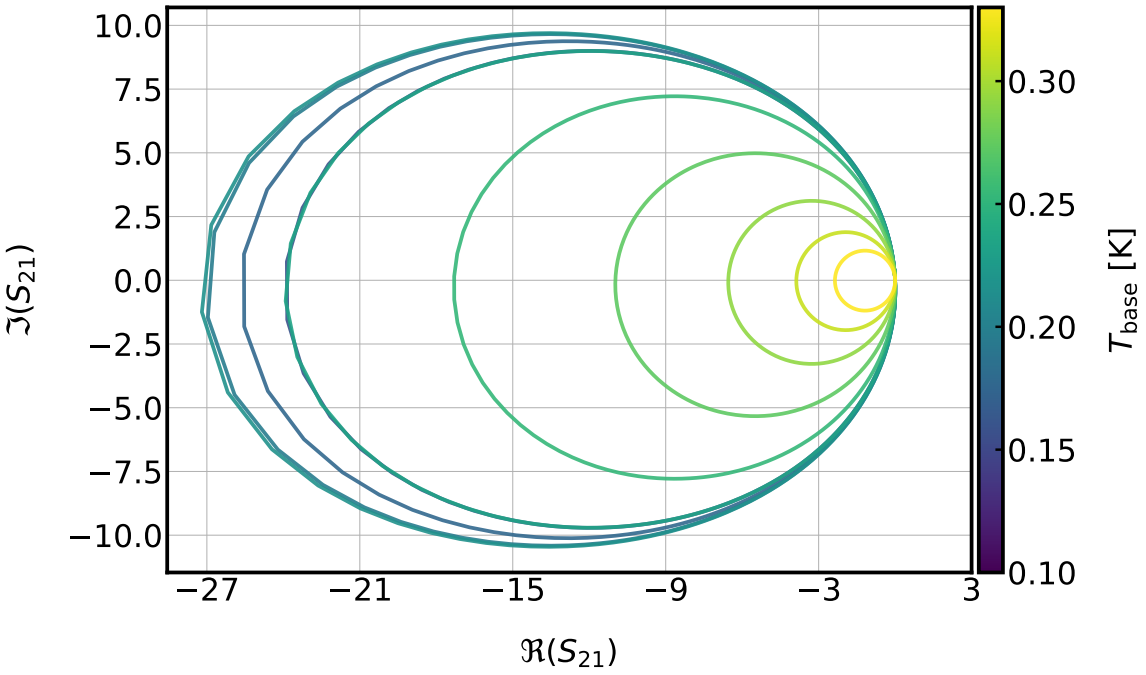

Figure 19. The simulated I/Q loop as a function of base temperature.

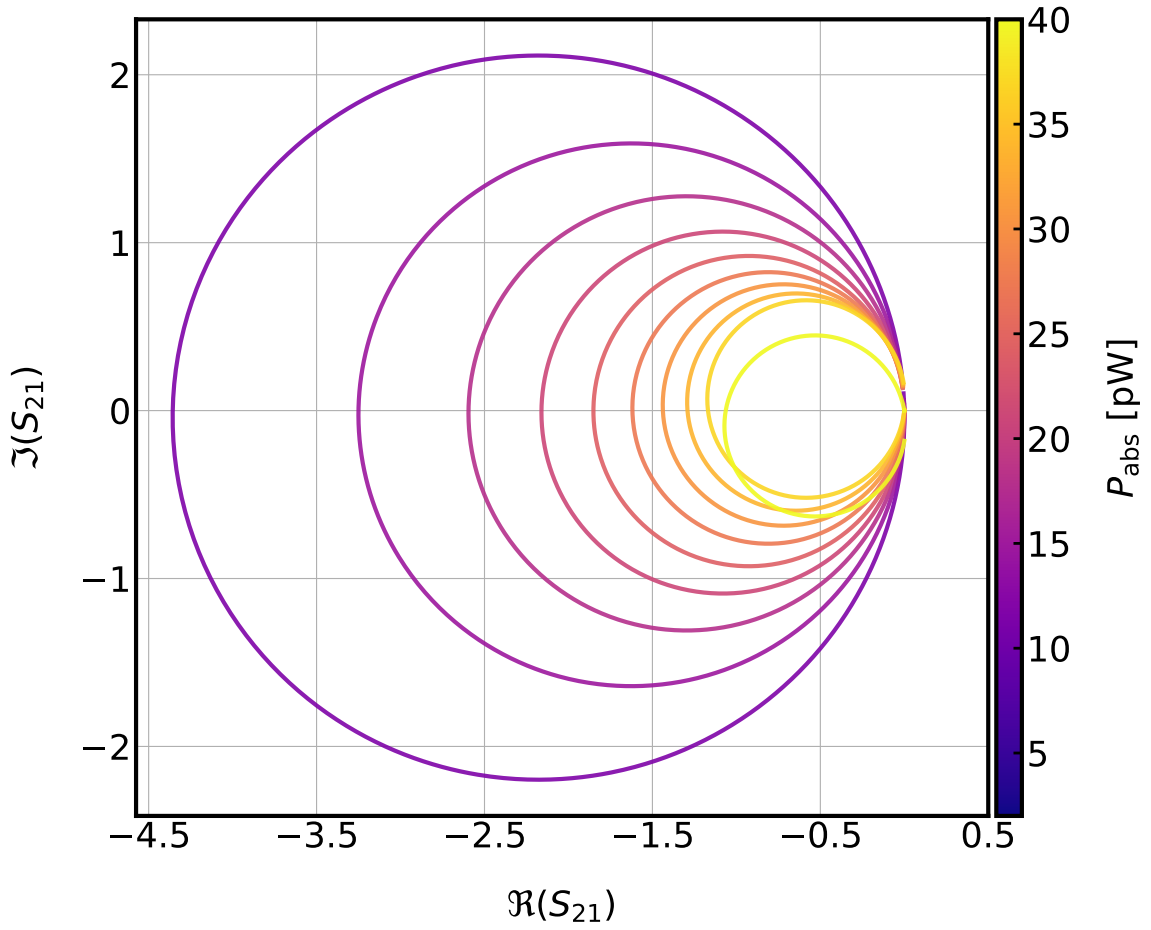

Figure 20. The I/Q loop as a function of optical power.

simulated BLAST-TNG 250 µm resonator with $\omega_0 = 828$ MHz. The 'dip-depth' of the S_{21} traces is a function of the quality factors:

$$D = 20 \log_{10} \left(1 - \frac{Q_r}{Q_c} \right) \tag{2.36}$$

The 'I/Q' loops corresponding to Figures 17 and 18 are shown in Figures 19 and 20, where $I = \Re(S_{21})$ and $Q = \Im(S_{21})$.

2.4 Responsivities

In this section, we derive the LEKID responsivity to fluctuations in base temperature T_{base} and absorbed optical power P_{abs}.

2.4.1 Temperature Responsivity

The change in the LEKID's resonant frequency in response to a change in base temperature can be written as:

$$\begin{aligned} R_T &= \frac{df_0}{dT} \\ &= \frac{df_0}{d\sigma_2}\frac{d\sigma_2}{dn_{qp}}\frac{dn_{qp}}{dT} \end{aligned} \tag{2.37}$$

where:

$$\frac{df_0}{d\sigma_2} = \frac{\alpha f_0}{2\sigma_2} \tag{2.38}$$

and for $T_{\text{base}} \ll T_c$,

$$\frac{dn_{qp}}{dT} \simeq \frac{n_{qp}}{T}\left(\frac{1}{2} + \frac{\Delta_0}{kT}\right) \tag{2.39}$$

Derivations of Equations 2.38 and 2.39 can be found in Mauskopf (2018).

Equation 2.37 is shown graphically in Figure 21. It can be seen that R_T decreases monotonically with increasing T_{base}.

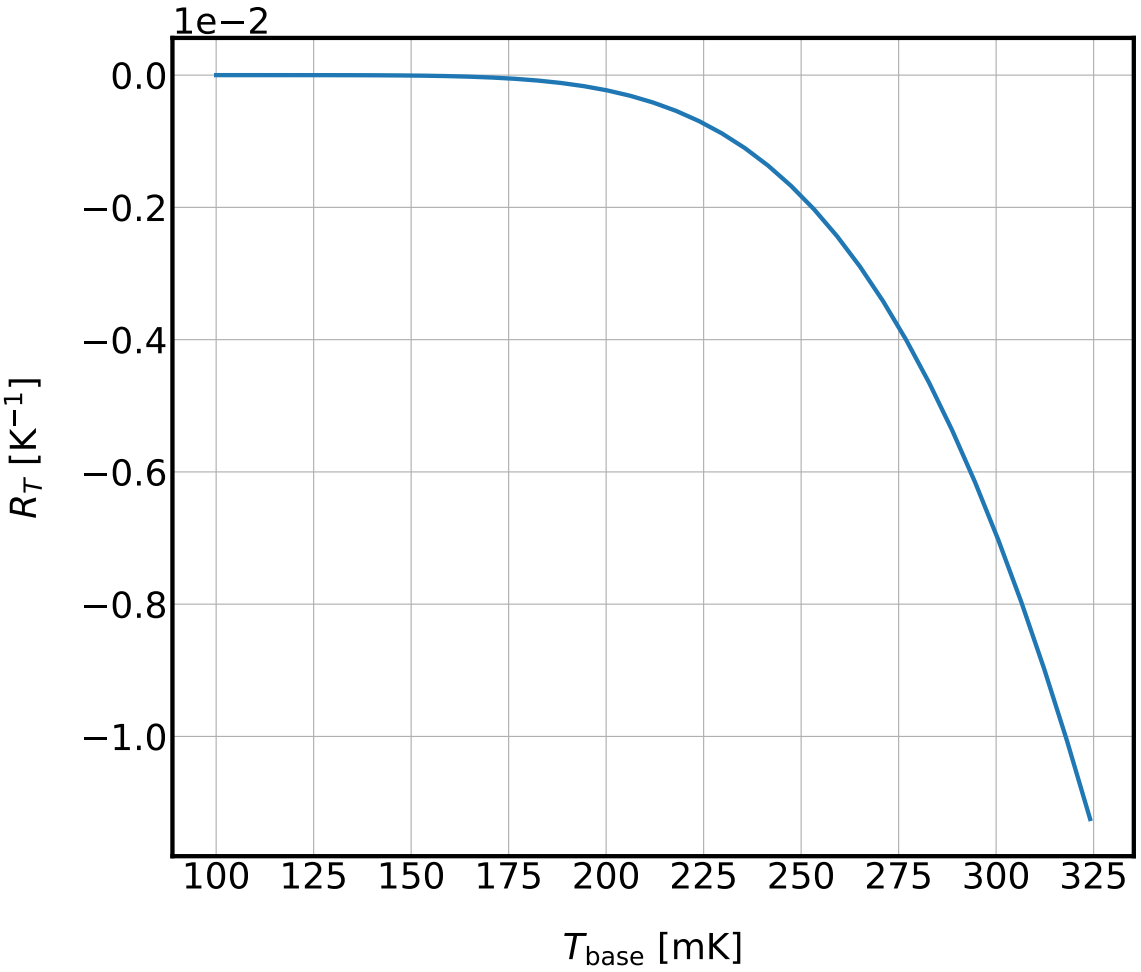

Figure 21. Simulated LEKID temperature responsivity R_T.

2.4.2 Optical Responsivity

At a fixed base temperature, the change in the LEKID's resonant frequency in response to a change in absorbed optical power can be written as:

$$R_P = \frac{df_0}{dP_{\text{abs}}} = \frac{df_0}{d\sigma_2}\frac{d\sigma_2}{dn_{qp}}\frac{dn_{qp}}{dP_{\text{abs}}} \tag{2.40}$$

where:

$$\frac{df_0}{d\sigma_2}\frac{d\sigma_2}{dn_{qp}} \simeq -\frac{\alpha f_0}{4N_0\Delta_0}\left(1+\sqrt{\frac{2\Delta_0}{\pi KT}}\right) \tag{2.41}$$

The expression for $\frac{dn_{qp}}{dP_{\rm abs}}$ is case dependent. Here, we provide a brief overview of three cases, as found in Mauskopf (2018). We first consider the rate of change in the total number density of QP:

$$\begin{aligned}\frac{n_{qp}}{dt} &= \Gamma_{\rm opt} + \Gamma_{\rm ro} + \Gamma_{\rm therm} + \Gamma_{\rm rec} \\ &= \frac{\eta_{\rm det}P_{\rm abs}}{\Delta\Sigma} + \frac{\epsilon_{\rm ro}P_{\rm ro}}{\Delta\Sigma} + \gamma {N_0}^2 8\pi kT\Delta e^{-2\Delta/kT} - \frac{n_{qp}}{\tau_{qp}}\end{aligned} \tag{2.42}$$

where:

$\Gamma_{\rm opt}$ is the QP generation rate due to absorbed optical power.

$\Gamma_{\rm ro}$ is the QP generation rate due to microwave readout power.

$\Gamma_{\rm therm}$ is the QP generation rate due to thermal phonons.

$\Gamma_{\rm rec}$ is the QP recombination rate into CPs.

$\eta_{\rm det}$ is the efficiency of QP generation from absorbed optical power.

$\epsilon_{\rm ro}$ is the internal QP generation efficiency for absorbed readout power.

$\gamma = 1/(n_{qp}\tau_{qp})$ is a constant relating the number density of QP to the QP recombination time.

In the first case (*Case 1*), we assume that QP generation is dominated by the absorption of photons. Equation 2.42 then becomes

$$\frac{n_{qp}}{dt} \simeq \frac{\eta_{\rm det}P_{\rm abs}}{\Delta\Sigma} - \gamma {n_{qp}}^2 \tag{2.43}$$

Then, $\frac{dn_{qp}}{dPP_{\rm abs}}$ can be shown to be:

$$\frac{dn_{qp}}{dP_{\rm abs}} = \frac{1}{2}\sqrt{\frac{\eta}{\gamma P_{\rm abs}\Delta\Sigma}} = \frac{n_0}{2P_{\rm abs}}\frac{1}{1+j\omega\tau_{qp}/2} \tag{2.44}$$

Behavior corresponding to *Case 1* has been observed in Al KIDs (e.g., De Visser (2014); Flanigan et al. (2016); Mauskopf et al. (2014)).

The second case to consider *Case 2* is if the QP generation rate is dominated by thermal phonon generation, or has a recombination time that does not depend on $P_{\rm abs}$. In this case, Equation 2.42 becomes:

$$\frac{n_{qp}}{dt} \simeq \frac{\eta_{\rm det} P_{\rm abs}}{\Delta\Sigma} + \Gamma_{\rm therm} - \frac{n_{qp}}{\tau_{\rm eff}} \tag{2.45}$$

and

$$\frac{dn_{qp}}{dP_{\rm abs}} \simeq \frac{\eta\tau_{\rm eff}}{\Delta\Sigma}\frac{1}{1+j\omega\tau_{\rm eff}} \tag{2.46}$$

Behavior corresponding to *Case 2* has been observed in TiN KIDs (e.g., Catalano et al. (2014); Hubmayr et al. (2015); Hailey-Dunsheath et al. (2016)).

In the intermediate case, there is an effective 'dark' power loading on the device, $P_{\rm dark}$, which generates a constant background of QP whose rate of creation is independent of absorbed optical power. In this case,

$$\frac{\epsilon_{\rm ro} P_{\rm ro}}{\Delta\Sigma} + \gamma {N_0}^2 8\pi k T \Delta e^{-2\Delta/kT} - \gamma {n_{qp}}^2 \tag{2.47}$$

and

$$\frac{dn_{qp}}{dP_{\rm abs}} = \sqrt{\frac{\eta_{\rm det}}{\Delta\Sigma\gamma}}\frac{1}{\sqrt{P_{\rm dark}+P_{\rm abs}}} \tag{2.48}$$

A model of the *Case 1* and *Case 2* responsivities is shown in Figure 22. In *Case 1*, the frequency shift due to absorbed power, in fractional frequency units (df_0/f_0) is linear. In *Case 2*, the frequency shift is approximately linear for low amounts of

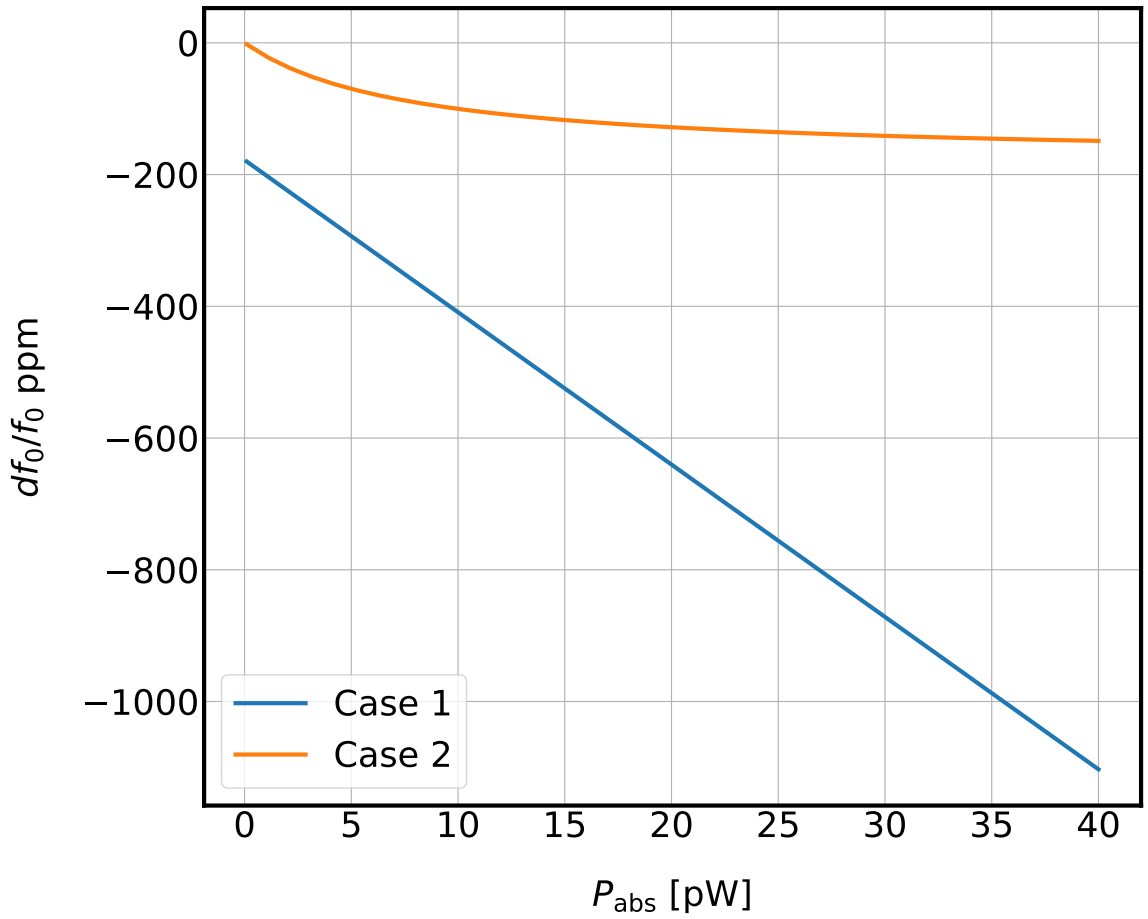

Figure 22. A model of the *Case 1* and *Case 2* optical responsivities.

absorbed power, then asymptotically approaches a constant value at higher powers. Figure 23 shows the *Case 2* responsivity for a range of base temperatures. The slope of the low-power, linear part of the curve becomes steeper at lower T_{base}.

2.4.3 The Nonlinear Kinetic Inductance

The KI has a second-order term whose effects on various LEKID parameters must be considered during the biasing stage of detector calibration. The nonlinear KI can be written as a function of internal current, I_{int}:

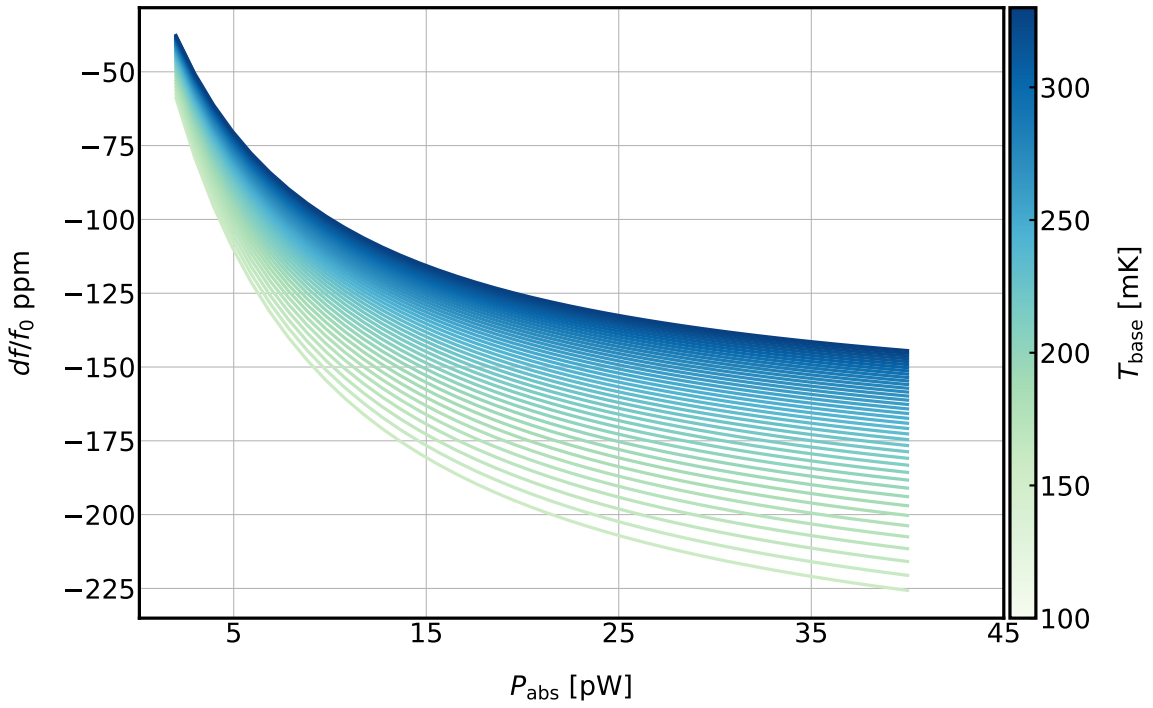

Figure 23. The *Case 2* responsivity shown for a range of base temperatures.

$$\mathcal{L}_k(I_{\rm int}) = \mathcal{L}_k(0)\left[1 + \left(\frac{I_{\rm int}}{I_\star}\right)\right] \tag{2.49}$$

where $I_\star$ is the critical current corresponding to the maximum charge-carrier velocity before pair-breaking occurs (Tinkham (2004); Anlage et al. (1989); Annunziata et al. (2010)). In a LEKID camera, the source of this internal current is the microwave probe tone generated by the readout system. For LEKIDs, the readout tone power, $P_{\rm ro}$, at the device typically ranges from ∼-60 to -120 dBm. The nonlinearity adds an additional frequency detuning δx, of $\mathcal{O}(I^2/I_\star^2)$:

$$\delta x = \frac{\delta\omega}{\omega_0} = -\frac{1\delta L}{2L} = \frac{\alpha}{2}\left(\frac{I}{I_\star}\right)^2 = -\frac{E_r}{E_\star} \tag{2.50}$$

where E_r is the internal resonator energy, and $E_\star = LI_\star^2/2\alpha = N_0\Delta^2\Sigma/2$ is the condensation energy of the inductor (Mauskopf (2018); Tinkham (2004)). As in

Section 2.2, the resonator current is related to the internal resonator energy by the inductance (here assuming that $L_m \ll L_K$). The resonator energy can be written as a function of tone power (Swenson et al., 2013), as:

$$\begin{aligned} E_r &= \frac{1}{2} L I_{\text{int}}^2 \\ &= \frac{2{Q_r}^2}{Q_c} \frac{1}{1 + 4{Q_r}^2 x^2} \frac{P_{\text{ro}}}{\omega_r} \end{aligned} \tag{2.51}$$

where ω_r is the shifted resonant frequency.

The total fractional detuning, x, is an implicit equation:

$$\begin{aligned} x &= x_0 + \delta x \\ &= x_0 - \frac{E_r(x)}{E_\star} \end{aligned} \tag{2.52}$$

Using the relations above, it's possible to estimate the degree of nonlinearity from a frequency sweep of S_{21}. In standard practice, it's preferable to bias each resonator so that the frequency detuning contributed by the readout tone power is less than one resonator line-width (FWHM). This ensures that the resonator doesn't bifurcate. The width of the resonator is $\delta\omega \simeq \omega_0/Q_r$. Therefore, the fractional detuning due to the nonlinearity should be: $|\delta x| = \frac{E_r}{E_\star} < \frac{1}{Q_r}$. Using the equations above, satisfying this condition requires that an upper limit to the internal resonator energy is:

$$E_r < \frac{N_0 {\Delta_0}^2 \Sigma}{2Q_r} \tag{2.53}$$

and the maximum readout tone power is:

$$P_{\text{ro}} < \frac{Q_c}{2{Q_r}^3} N_0 \Delta^3 \Sigma \omega \left(1 + 4{Q_r}^2 x^2\right) \tag{2.54}$$

At a fixed base temperature, larger optical loading permits the use of higher probe tone powers. This behavior is illustratd in Figure 24, which shows the maximum probe tone power for a range of optical powers.

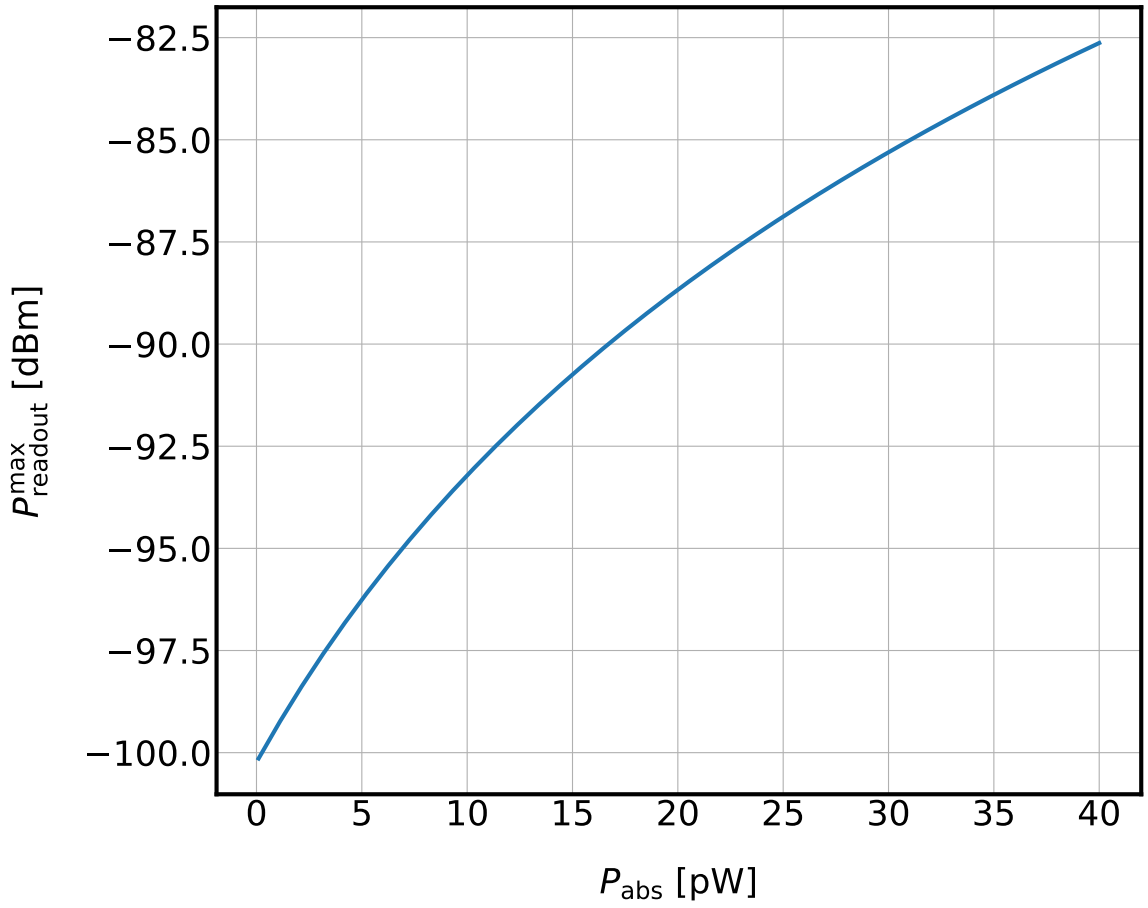

Figure 24. The maximum probe tone power shown for a range of optical power loadings.

2.4.4 Bifurcation

At a particular probe tone power, the resonance enters an unstable equilibrium between two states, which is known as bifurcation. This phenomenon can be under-

stood from the perspective of soft-spring Duffing effects (see, e.g., Swenson et al. (2013); Duffing (1918)). Bifurcation can be induced during an S_{21} frequency sweep, by sweeping the resonance either from below or from above the resonant frequency. These two cases correspond to positive and negative feedback, respectively.

Upward sweep, positive feedback: As the probe tone is swept upward in frequency from below the resonant frequency, the resonant frequency is detuned, and drawn toward the probe tone frequency. This positive feedback increases until the resonator bifurcates between the energized state and its rest state.

Downward sweep, negative feedback: As the probe tone is swept downward in frequency from below the resonant frequency, the resonator is detuned farther below the probe tone frequency. When the probe tone eventually passes the resonant frequency, the resonator bifurcates between its energized and rest states.

The bifurcation power can be estimated using Equation 2.54. If the resonator quality factor is dominated by the coupling quality factor, $Q_r \approx Q_c$, and $x \ll 1$,

$$P_{\text{bifurc}} = \frac{1}{2{Q_c}^2} N_0 \Sigma \omega_0 \Delta^2 \tag{2.55}$$

2.5 Sensitivity

The sensitivity of a sub-mm detector is typically reported as either a noise-equivalent power (NEP) or noise-equivalent temperature (NET). The NEP (NET) describes the detector's sensitivity to absorbed optical power (temperature). NEP is commonly written in units of $\text{W}/\sqrt{\text{Hz}}$. This represents the noise power inside a measurement bandwidth of 1 Hz. Equivalently, it is the noise power which is measured after a 0.5 second integration (by the Shannon-Nyquist sampling theorem). At the

time of writing, typical KID NEPs (under optical loading) are of $\mathcal{O}(10^{-17})$ W/$\sqrt{\text{Hz}}$. With that NEP, one can measure a signal level of 10^{-17} W with a signal-to-noise-ratio (SNR) of 1, over a 0.5 s integration.

KIDs, like most other astronomical detectors, are designed to be photon noise limited. Photon noise follows Poisson, or shot noise statistics. In a measurement dominated by shot noise, the SNR increases as $\sqrt{N}$, where N is the number of measurements. In a typical sub-mm observation, the detector noise is limited by the shot noise of the background of in-band photons emanating from the receiver, optical system, atmosphere and astronomical source. In this scenario, the detectors are said to be operating in the background-limited IR photodetector (BLIP) limit.

During an on-sky KID measurement, it is preferable to view the detector timestreams in units of frequency shift (or fractional frequency) Δf, relative to the resonant frequency of each pixel. If a responsivity (e.g., df/dT, df/dP) is known, then the NEP (or NET) can be calculated by dividing the square-root of the power spectral density (PSD) of the frequency fluctuations (units of Hz/$\sqrt{\text{Hz}}$) by the corresponding responsivity.

2.5.1 Instrumental Noise Hierarchy

The NEP which is calculated during a measurement is the sum of contributions from several different noise sources. The order in which each noise source should dominant forms the instrumental noise hierarchy. In a LEKID camera, the noise hierarchy (from most to least dominant) is:

- Photon noise

- Amplifier noise (dominated by the first-stage cryogenic low noise amplifier (LNA))
- Detector noise (generation-recombination (GR))
- Readout noise

The readout noise is discussed in Chapter 3. The other three noise sources are treated individually in the following sections.

2.5.2 Generation-Recombination Noise

In the absence of readout and amplifier noise, the intrinsic noise floor of a KID measurement is set by GR noise. GR noise originates from the instantaneous uncertainty in the rate of CPs recombining into QPs (or, equivalently, the rate at which CPs are broken into QP pairs). A simple expression for the GR contribution to the fractional noise spectral density follows from a consideration of the generation and recombination rates ($\Gamma_{\rm gen}$ and $\Gamma_{\rm rec}$), assuming a steady state ($\Gamma_{\rm gen} = \Gamma_{\rm rec}$). The uncertainty in the number of CPs is:

$$\begin{aligned} \delta N_{cp} &= \Gamma_{\rm gen}\delta t - \Gamma_{\rm rec}\delta t \\ &= \sqrt{(\Gamma_{\rm gen} + \Gamma_{\rm rec})\delta t} \\ &= \sqrt{2\Gamma_{\rm rec}\delta t} \end{aligned} \tag{2.56}$$

where $\delta t = 1/2\delta f$ is the measurement integration time (with measurement bandwidth δf). Using the fact that $\Gamma_{\rm rec} = \frac{N_{qp}}{2\tau_{qp}}$, the uncertainty in the recombination rate is:

$$\begin{aligned}\delta\Gamma_{\mathrm{rec}} &= \frac{\delta N_{cp}}{\delta t} \\ &= \sqrt{\frac{2\Gamma_{\mathrm{rec}}}{\delta t}} \\ &= \sqrt{\frac{2N_{qp}}{2\tau_{qp}\delta t}}\end{aligned} \tag{2.57}$$

Equation 2.56 becomes:

$$\begin{aligned}\delta N_{cp} &= \delta\Gamma_{\mathrm{rec}}\tau_{qp} \\ &= \sqrt{\frac{N_{qp}\tau_{qp}}{\delta t}} \\ &= \sqrt{2N_{qp}\tau_{qp}\delta f}\end{aligned} \tag{2.58}$$

Dividing Equation 2.58 by the square-root of the measurement bandwidth, and multiplying by the frequency responsivity to fluctuations in the number of QP yields an expression for the frequency noise PSD:

$$e_{f,GR} = \sqrt{2N_{qp,\mathrm{tot}}\tau_{qp}}\frac{df_0}{dN_{qp}} \qquad \left[\frac{\mathrm{Hz}}{\sqrt{\mathrm{Hz}}}\right] \tag{2.59}$$

From Equation 2.59, the NEP contribution from GR can be calculated as:

$$NEP_{GR} = e_{f,GR}\Big/\left(\frac{df_0}{dN_{qp}}\frac{dN_{qp}}{dP_{\mathrm{opt}}}\right) \qquad \left[\frac{\mathrm{W}}{\sqrt{\mathrm{Hz}}}\right] \tag{2.60}$$

2.5.3 Two-Level System Noise

Nonthermal two-level system (TLS) fluctuations occur throughout the amorphous dielectric layer of the superconducting substrate. The TLS systems have electric and magnetic dipoles which can interact with external fields around the device. These interactions manifest as fluctuations in the complex permittivity and/or per-

meability of the material. In a KID, random fluctuations in C_r are indistinguishable from fluctuations in KI, because both introduce jitter to the resonant frequency. The TLS noise is generally modeled as being capacitive in origin (see, e.g., Gao (2008); Zmuidzinas (2012)). It is thought to be the underlying cause of an often observed anomalous increase in resonant frequency with increasing base temperature (a violation of predictions made by MB theory), however other explanations have been offered (see, e.g., the Kondo effect (Noguchi et al., 2018)).

The TLS loss tangent, δ_{TLS}, can limit the Q_i of a KID, and has an empirical dependence on the internal power stored in the resonator (proportional to microwave probe tone power), expressed as (Barry (2014); Martinis et al. (2005)):

$$\delta_{TLS} = {Q_i}^{-1} = F_{TLS}\delta_0 \left(\frac{1}{1+|P_{\rm int}/P_c|}\right)^{1/2} \tag{2.61}$$

where:

F_{TLS} is the ratio of the electric field stored in the TLS systems to the total electric field

δ_0 is the intrinsic loss tangent

$\chi = \hbar\omega/2kT$

$P_{\rm int}$ is the internal resonator microwave power

$P_{\rm sat}$ is the TLS saturation power

where δ_{TLS} is known to saturate above a characteristic power, $P_{\rm sat}$. Because the maximum microwave readout power which can be used before nonlinearity sets in (Equation 2.54) is proportional to inductor volume ($P_{ro,\max} \propto \Sigma$). The TLS noise therefore increases as $\Sigma^{1/2}$. The number of TLS systems scales with the volume of the amorphous dielectric, or Σ, and the way in which the KID responsivity varies with Σ depends on the substrate material (see Section 2.4).

For Al KIDs, the responsivity does not depend on Σ (Equation 2.44). The TLS noise therefore increases relative to GR and other noise sources as $\Sigma^{1/2}$, and the inductor volume can be made large. In TiN KIDs, $df_0/dP_{\rm abs} \propto \Sigma^{-1}$, and these devices typically use smaller Σ (Mauskopf, 2018). TLS noise mitigation techniques which involve the use of crystalline, rather than amorphous dielectrics, are presently an active area of research (see, e.g., Weber et al. (2011)).

A semi-empirical model for the TLS noise spectral density is presented in Gao et al. (2008). The noise PSD has a colored spectrum ($S_{\rm XX} \propto f^{-\alpha}$) with spectral index $\alpha \approx 0.5$, and depends on microwave readout power and $T_{\rm base}$ as:

$$S_{{\rm XX},TLS} \propto T^{-t} P_{\rm ro} \tag{2.62}$$

where $t = 1.5$–2. NEP_{TLS} can be found by dividing Equation 2.62 by a responsivity (e.g., R_T or R_P).

2.5.4 Amplifier Noise

The LNA which is used to amplify the probe tone comb at the output of the KID array adds thermal, or Johnson-Nyquist noise to the signal (Johnson (1928); Nyquist (1928)). Johnson-Nyquist noise in a resistor originates from thermal motions of the electrons which manifest as voltage fluctuations across its terminals. Following Kittel and Kroemer (1998), an expression for the RMS voltage contribution from Johnson-Nyquist noise is found by considering a simple circuit model. The circuit consists of a white noise generator with resistance R which transfers power to a resistive load with resistance R_L. The power transferred from the noise generator to the load is:

$$P = \left\langle I^2 \right\rangle R_L = \frac{\langle V^2 \rangle R}{(R + R_L)^2} \tag{2.63}$$

The maximum power transfer occurs when the source and load are perfectly matched ($R = R_L$). In this case,

$$P = \frac{\langle V^2 \rangle R}{4R} \tag{2.64}$$

and the RMS voltage is:

$$\left\langle V^2 \right\rangle = 4PR = 4kT \int_{f_0}^{f_1} R df = 4kTRB \tag{2.65}$$

where T is the temperature of the Thevenin equivalent of the circuit, and B is the bandwidth over which the voltage is being measured. The resistance is assumed to be constant with frequency. Returning to the case of the LNA, the one-sided PSD of the voltage fluctuations is:

$$e_{V,\mathrm{amp}} = 4kT_{\mathrm{amp}}Z \qquad \left[\frac{\mathrm{V}^2}{\mathrm{Hz}}\right] \tag{2.66}$$

where $Z = R$ is the input impedance of the amplifier. The PSD is converted to frequency units by dividing by the voltage responsivity:

$$e^2_{f,\mathrm{amp}} = \frac{e^2_{v,\mathrm{amp}}}{|dV/df_0|^2} \tag{2.67}$$

The voltage responsivity can be approximated using Equation 2.35:

$$\begin{aligned} \frac{dV_{\mathrm{out}}}{df_0} &= V_{\mathrm{in}} \frac{dS_{21}}{df_{\mathrm{out}}} \\ &\simeq -2jV_{\mathrm{in}} \frac{{Q_r}^2 f_{\mathrm{probe}}}{Q_c f_0^2} \frac{1}{(1 + 2jQ_r x)^2} \end{aligned} \tag{2.68}$$

Near resonance, $x \ll 1$, and

$$\frac{dV_{\rm out}}{df_0} \simeq 2jV_{\rm in}\frac{{Q_r}^2}{Q_c f_0} \tag{2.69}$$

Equation 2.67 can then be written as:

$$\begin{aligned} e^2_{f,\rm amp} &= 4kT_{\rm amp}Z_0\left(\frac{{Q_c}^2 f_0^2}{4{Q_r}^4}\right) \\ &= kT_{\rm amp}\left(\frac{{Q_c}^2 f_0^2}{{Q_r}^4 P_{\rm ro}}\right) \\ &= kT_{\rm amp}\left(\frac{f_0^2}{{Q_r}^2 P_{\rm ro}}\right) \end{aligned} \tag{2.70}$$

where the probe tone power $P_{\rm ro} = V^2_{\rm in}/Z_0$. For a resonator limited by the coupling quality factor ($Q_r \approx Q_c$),

$$e^2_{f,\rm amp} = kT_{\rm amp}\left(\frac{f_0^2}{{Q_r}^2 P_{\rm ro}}\right) \qquad \left[\frac{\rm Hz^2}{\rm Hz}\right] \tag{2.71}$$

Therefore, for a probe tone near the resonant frequency, the frequency noise decreases with increasing readout power. As the tone power approaches the bifurcation power (Equation 2.55), the frequency noise becomes independent of tone power. The LNA contribution to the NEP can therefore be written as:

$$NEP_{\rm amp} = e_{f,\rm amp}\Big/\left(\frac{df_0}{dN_{qp}}\frac{dN_{qp}}{dP_{\rm opt}}\right) \qquad \left[\frac{\rm W}{\sqrt{\rm Hz}}\right] \tag{2.72}$$

The $NEP_{\rm amp} \propto 1/\sqrt{P_{\rm ro}}$ dependence makes it possible to suppress the amplifer noise contribution by using higher microwave readout powers.

2.5.5 Photon Noise

The optical power input into the system is assumed to be incoherent emission from a thermal blackbody source. To write an expression for the total optical power

P_{opt}, we first consider the number of modes contained in the incident light beam, which has a central frequency ν and optical bandwidth $\Delta\nu$. The total number of modes is the product of the number of spatial, temporal and polarization modes:

$$\begin{aligned} N_{\text{modes}} &= (N_{\text{spatial}})\,(N_{\text{temporal}})\,(N_{\text{pol}}) \\ &= \left(\frac{A\Omega\nu^2}{c^2}\right)(\tau_{\text{int}}\Delta\nu)\,(m) \end{aligned} \tag{2.73}$$

where:

$A\Omega$ is the *é*ntendue

$\nu^2/c^2 = \lambda^2$ is the *é*ntendue of coherence

τ_{int} is the integration time

m is the number of polarization modes [1,2]

The photon arrival times follow Bose-Einstein statistics, with an occupation number, or mean number of photons per mode, of:

$$n_{\text{occ}} = \frac{\langle N_{\text{photons}}\rangle}{\text{mode}} = \frac{1}{e^{h\nu/kT}-1} \tag{2.74}$$

where T is the blackbody temperature of the source. With $m = 2$, the total optical power can be expressed as:

$$\begin{aligned} P_{\text{opt}} &= \left(N_{\text{modes}} \times \frac{\langle N_{\text{phot}}\rangle}{\text{mode}} \times \text{Energy per photon}\right) \Big/ \tau_{\text{int}} \\ &= \left(\frac{mA\Omega\nu^2}{c^2}\right)(\tau_{\text{int}}\Delta\nu)\, n_{\text{occ}}\frac{h\nu}{\tau_{\text{int}}} \\ &= A\Omega\Delta\nu\left(\frac{2h\nu^3}{c^2}\frac{1}{e^{h\nu/kT}-1}\right) \\ &= A\Omega\Delta\nu B_\nu(\nu,T) \qquad [\text{W}] \end{aligned} \tag{2.75}$$

where $B_\nu(\nu,T)$ $\left[\text{Wm}^{-2}\text{str}^{-1}\text{Hz}^{-1}\right]$ is the Planck function for spectral radiance.

The incident power which is absorbed by the detectors is the total optical power multiplied by the optical efficiency $\epsilon_{\rm opt}$ and quantum efficiency of the detectors $\eta_{\rm det}$:

$$\begin{aligned} P_{\rm abs} &= \epsilon_{\rm opt}\eta_{\rm det}P_{\rm opt} \\ &= \epsilon_{\rm opt}\eta_{\rm det}A\Omega\Delta\nu B_\nu(\nu,T) \qquad [\rm W] \end{aligned} \tag{2.76}$$

The mean-squared noise power measured by the detectors originates from fluctuations in $n_{\rm occ}$. Because its probability distribution function is the Bose-Einstein distribution, its variance is:

$$\sigma^2_{n_{occ}} = n_{\rm occ} + \frac{{n_{\rm occ}}^2}{N_{\rm modes}} \tag{2.77}$$

where a factor corresponding to the case of $N_{\rm modes} \ll 1$ has been omitted (see, e.g., Fox (2006); Rowe (2015)). The mean-squared noise power input to the optical system is therefore

$$\begin{aligned} \sigma^2_P &= \sigma^2_{n_{occ}} N_{\rm modes}\left(\frac{h\nu}{\tau_{\rm int}}\right)^2 \\ &= \frac{h\nu P_{\rm opt}}{\tau_{\rm int}} + \frac{{P_{\rm opt}}^2}{N_{\rm modes}} \qquad \left[\rm W^2\right] \end{aligned} \tag{2.78}$$

The noise spectral density (square-root of the NEP) is found by dividing Equation 2.78 by the measurement bandwidth, $B = f_s/2 = \tau_{\rm int}$:

$$S_{xx} = NEP^2_{\rm phot} = h\nu P_{\rm opt} + \frac{{P_{\rm opt}}^2}{N_{\rm modes}B} \qquad \left[\frac{\rm W^2}{\rm Hz}\right] \tag{2.79}$$

When the NEP is expressed in terms of $P_{\rm opt}$, it is referred to as the *electrical* NEP. When the NEP is expressed in terms of $P_{\rm abs}$, it is referred to as the *optical* NEP. The first term in Equation 2.79 dominates for low $n_{\rm occ}$ (shot noise limit, $h\nu/kT \gg 1$), and the second term dominates for high occupation number (wave noise limit, $h\nu/kT \ll 1$):

$$NEP^2_{\rm phot} = NEP^2_{\rm shot} + NEP^2_{\rm wave} \tag{2.80}$$

In each of BLAST-TNG's three observation bands (250, 350 and 500 μm) the shot noise term dominates the NEP. In frequency units, the NEP can be calculated as:

$$e_{f,\rm phot} = NEP_{\rm phot} \frac{df_0}{dN_{qp}} \frac{dN_{qp}}{dP_{\rm opt}} \qquad \left[\frac{\rm Hz}{\sqrt{\rm Hz}}\right] \tag{2.81}$$

2.5.6 Total NEP

In frequency units, the total noise spectral density measured by the detectors is the quadrature sum of the noise contributions from each of the sources discussed in the above sections:

$$e_f = \left(e^2_{f,\rm phot} + e^2_{f,\rm amp} + e^2_{f,TLS}\right)^{1/2} \qquad \left[\frac{\rm Hz}{\sqrt{\rm Hz}}\right] \tag{2.82}$$

Figure 25 shows Equation 2.82 as a function of absorbed optical power (omitting the contribution from TLS). Photon noise is dominant for absorbed power between ∼1–18 pW, above which amplifier noise dominates due to the decrease in resonator quality factor Q_r. The physical parameters in this example simulation (which are representative of the BLAST-TNG 250 μm array) predict white noise levels, in fractional frequency units, of a few $\times\ 10^{-17}$ [Hz^{-1}]. The ratio of the photon noise to the amplifier and GR noise is illustrated in Figure 26. The photon noise dominates over the GR noise at absorbed powers of a few picowatts.

The total optical NEP is:

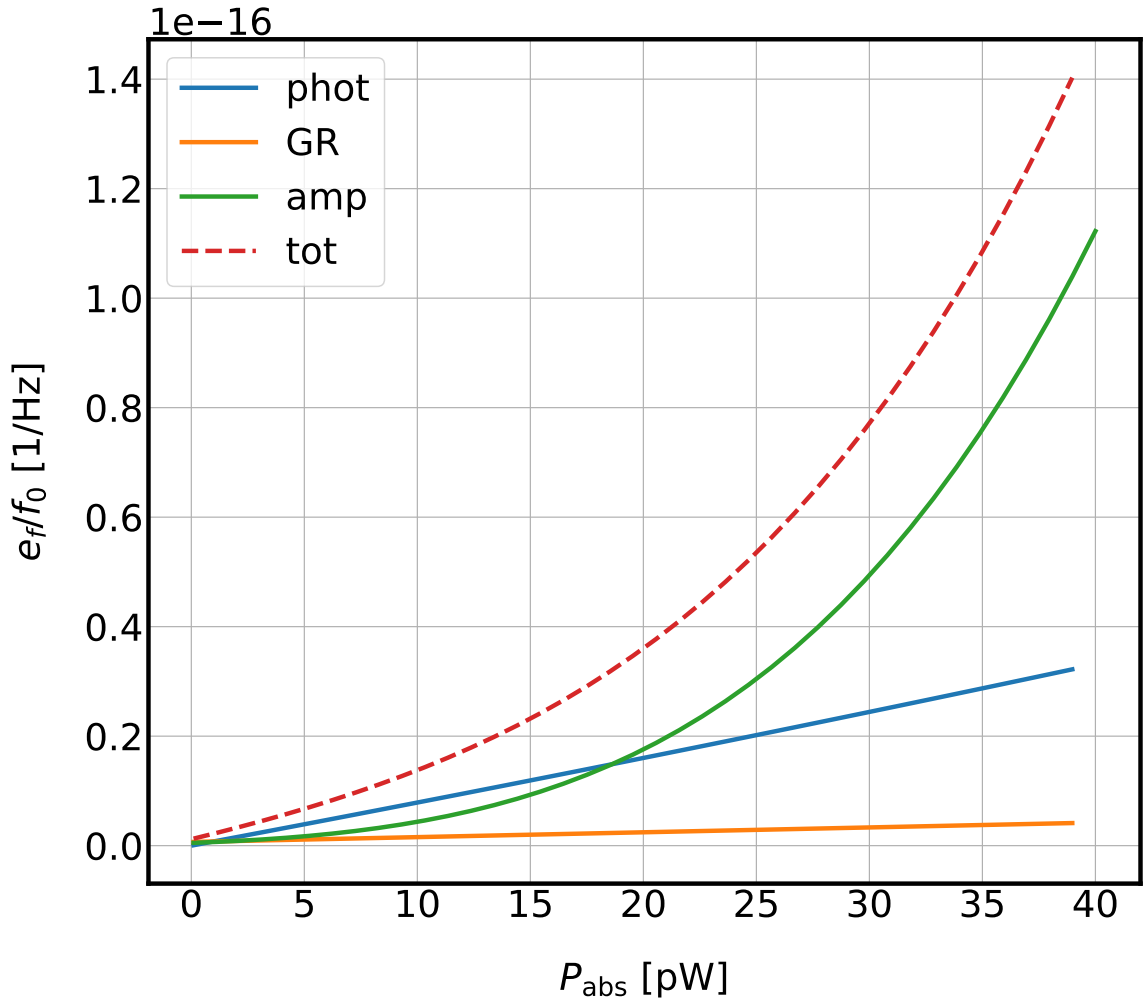

Figure 25. The spectral density, in fractional frequency units, as a function of optical power, for the major contributors of instrumental noise.

$$NEP_{\mathrm{tot}} = \left(NEP^2_{\mathrm{amp}} + NEP^2_{GR} + NEP^2_{TLS} + NEP^2_{\mathrm{phot}}\right)^{1/2} \qquad \left[\frac{\mathrm{W}}{\sqrt{\mathrm{Hz}}}\right] \qquad (2.83)$$

The NEP for each noise source (except for TLS) is shown in Figure 27. At absorbed powers above a few picowatts, the simulated NEP is $\propto \sqrt{P_{\mathrm{opt}}}$.

and the total NET can be calculated as:

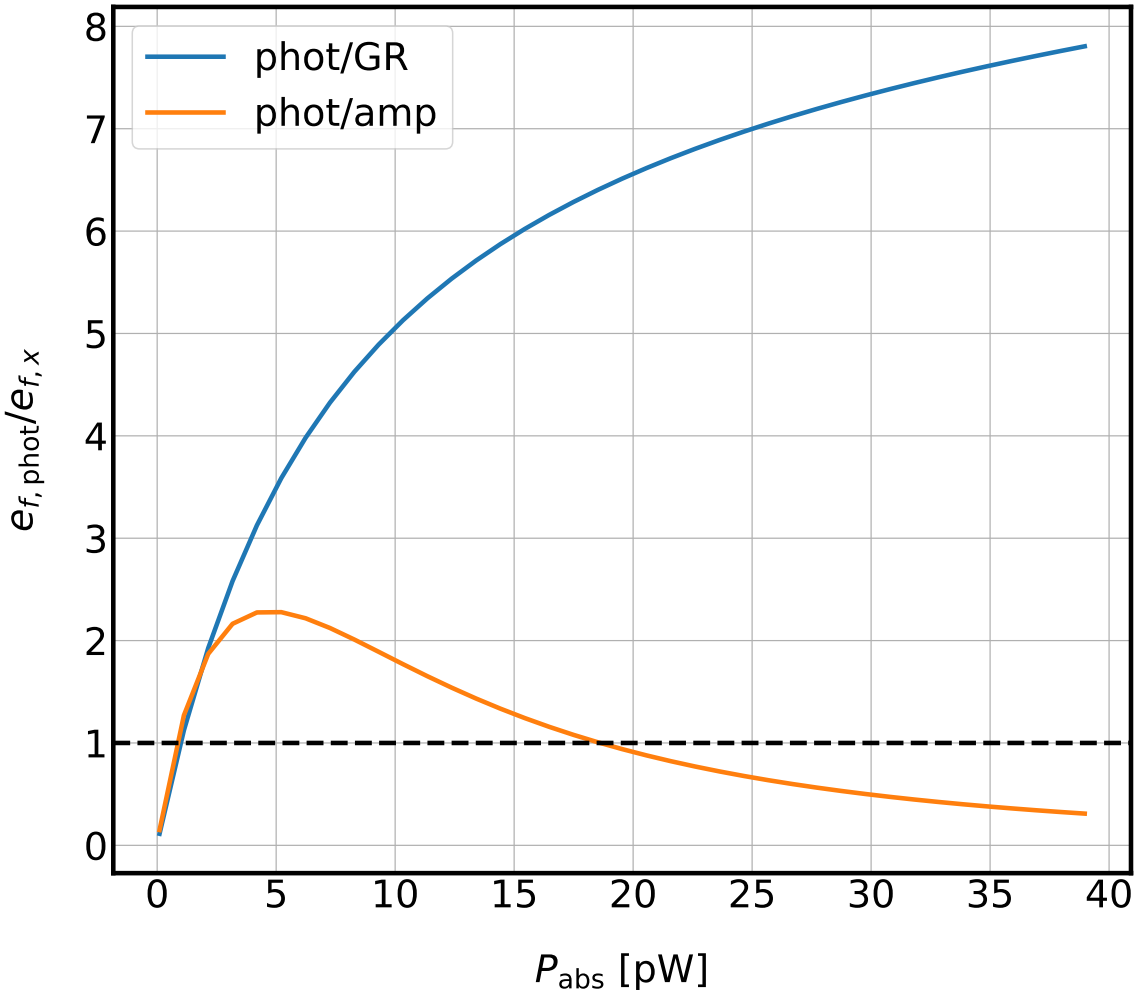

Figure 26. The ratio of $e_{f,phot}$ to $e_{f,GR}$ and $e_{f,\mathrm{amp}}$.

$$NET_{\mathrm{tot}} = \frac{NEP_{\mathrm{tot}}}{k\epsilon_{\mathrm{opt}}\Delta\nu} \qquad \left[\frac{\mathrm{K}}{\sqrt{\mathrm{Hz}}}\right] \tag{2.84}$$

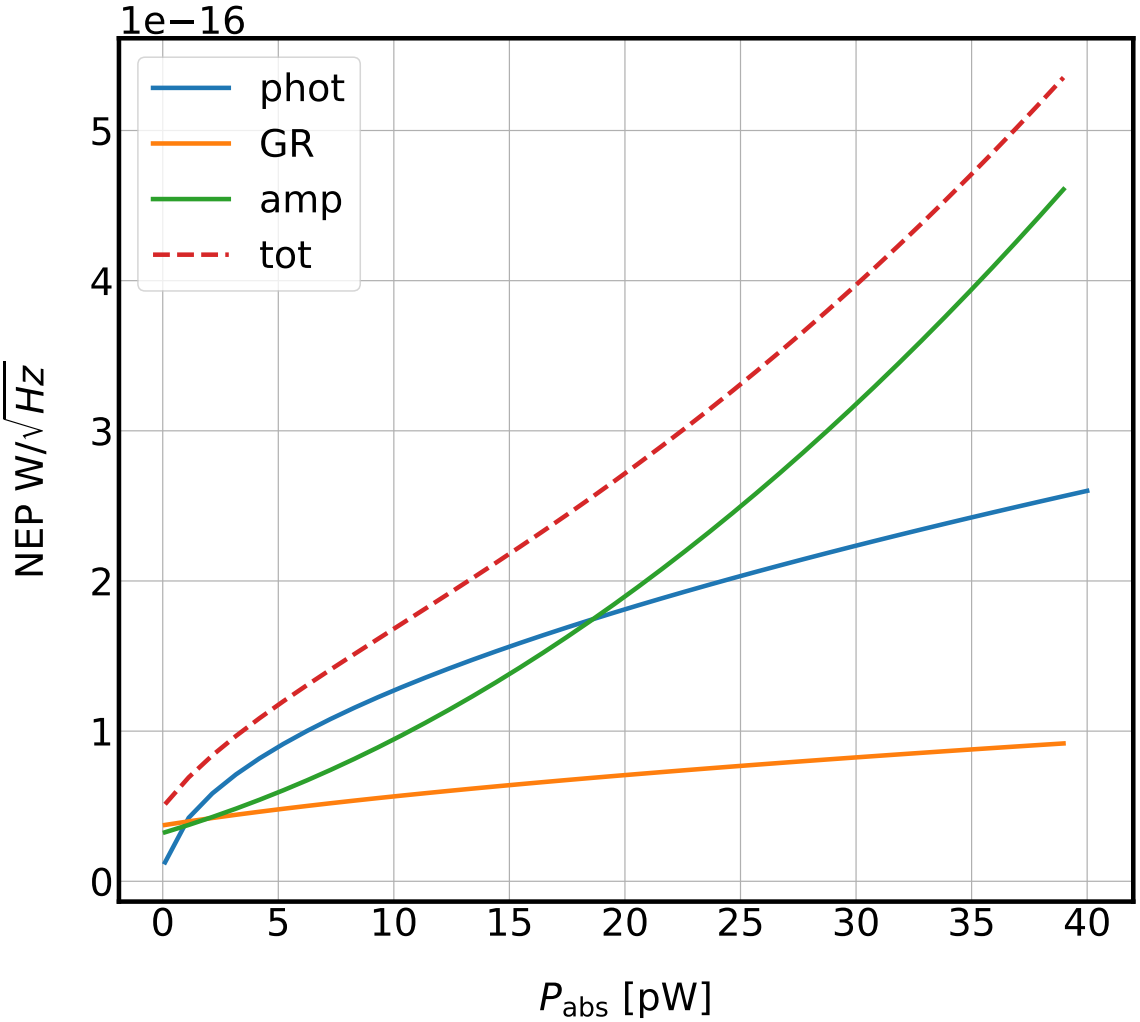

Figure 27. The optical NEP as a function of optical power, for the major contributors of instrumental noise.

Chapter 3

DESIGN OF AN FPGA-BASED LEKID SPECTROMETER

The ASU open-source LEKID readout is a 1,000-channel transceiver, consisting of digital and analog electronics, firmware and software. This technology builds on the legacy of previous MKID demonstrator instruments which have used Field Programmable Gate Array (FGPA) platforms (e.g. MUSIC (Golwala et al., 2012), ARCONS (McHugh et al., 2012), NIKA (Monfardini et al., 2014), MUSTANG-2 (Dicker et al., 2014), MAKO (Swenson et al., 2012), SPACEKIDS (van Rantwijk et al., 2016), DARKNESS (Strader, 2016)). FPGAs enable the high speed digital signal processing (DSP) techniques which are required for both time and frequency-domain-multiplexed readout (FDM) of superconducting detectors.

The FPGA board which serves as the foundation for the LEKID readout described in this work is the second generation Reconfigurable Open Architecture Computing Hardware (ROACH2), an open-source DSP board developed by the Collaboration for Astronomical Signal Processing and Electronics Research (Werthimer (2011); Hickish et al. (2016)). The ROACH2 is a general purpose DSP board whose primary use is as a digital backend for radio astronomy observatories. The ASU LEKID readout, which is the subject of this chapter, is one of several MKID readouts which have used the ROACH1/ROACH2 platform, including MUSIC (ROACH1), MAKO (ROACH2), ARCONS (ROACH1/2) and DARKNESS (ROACH2/DARKNESS board).

The design and creation of the readout system was driven by a philosophy based on adaptability and scalability. Consequently, the system now serves as a core part of

several current and upcoming astronomical instruments. Among these are balloon-borne (e.g., OLIMPO (Masi et al., 2019), BLAST-TNG (Gordon et al., 2016)) and ground-based cameras (e.g., TolTEC (Austermann et al., 2018), SuperSpec (Wheeler et al., 2018), MUSCAT (Brien et al., 2018)).

In the following sections, we describe the design and verification of the LEKID readout system, with a primary focus on the system which has been developed for the BLAST-TNG stratospheric balloon platform. The chapter is organized as follows:

- Section 3.1.1 provides a brief overview of the principal readout functions.
- Section 3.1 describes the primary system requirements.
- Section 3.3 details the design and verification of the ROACH2 firmware and DSP algorithms.
- Section 3.2 describes the design and construction of the BLAST-TNG readout hardware and electronics.
- Section 3.4 presents the results of noise verification measurements.

3.1 System Requirements

The core system requirements for a KID readout system fall under categories of multiplexing factor, data rate, noise performance, size, weight, power and cost (SWaP-C) and timing synchronization. The readout system must be able to read out the required number of LEKID detectors at a rate determined by the telescope scan speed while contributing less noise than the detectors and cryogenic low-noise amplifiers (LNAs). To facilitate map making, the detector data packets must be timestamped in such a way that makes it possible to later align them with the timing information provided by the pointing system.

It is important to note that the requirements which are discussed here pertain primarily to sub-mm/FIR/mm-wave KID cameras. O/NIR systems have slightly different requirements (see Section 1.5.1). The following sections discuss the motivation behind each system requirement. The requirements are listed in Table 4. We begin with an overview of the system.

Parameter	Requirement
RF bandwidth (MHz)	512
S_ϕ / chan (dBc/Hz)	$\lesssim$ -95
1/f corner f_c (Hz)	$\lesssim 0.5$
Data rate (Hz)	200–500
Power dissipation (W)	$\lesssim 65$
Time stamp precision (ms)	$\lesssim 1$
DAC Tone Resolution (Hz)	≤ 1000
Multiplexing Factor	$\gtrsim 500$

Table 4. LEKID readout system requirements.

3.1.1 System Overview

The LEKID readout processes 512 MHz of instantaneous complex baseband and RF (radio frequency) bandwidth using quadrature (I/Q) signal processing. The total bandwidth is divided between 1,000 channels, each with readout bandwidth of ∼30–500 Hz. For each channel, the system outputs a 64-bit I/Q sample at a configurable data rate of ∼60–1,000 Hz. The I/Q pairs represent the total transmission of each resonator ($I + jQ = V_{\text{in}}S_{21}$), and contain the AM (amplitude modulation) and PM (phase modulation) of the detectors. If the detectors are properly biased, the PM is

proportional to the absorbed optical power. For KID readout, typically only the PM is used.

In order to accurately probe the resonant frequency of each detector, the probe (carrier) tones synthesized by the DAC must have a frequency resolution $f_{\mathrm{res}} \lesssim 1,000$ Hz (this requirement scales with the resonator quality factor Q_r).

To readout the detectors, the system must do the following:

1. Synthesize a baseband probe tone ($I_{\mathrm{DAC}} + jQ_{\mathrm{DAC}}$) in software and firmware with a unique amplitude, frequency and phase for each resonator.
2. Generate the tone comb with the digital-to-analog converters (DACs).
3. Upconvert the tone comb to the RF band of the detectors and send to the detector arrays.
4. Receive the detector-modulated tone comb.
5. Downconvert the modulated tone comb from RF to baseband, amplify/attenuate/filter as needed.
6. Digitize the tone comb and analyze with a polyphase filterbank/Fourier transform (PFB-FFT, see Section 3.3.3).
7. Discard unwanted FFT bins (remaining bins are labeled as channels).
8. Digitally demodulate each channel (digital downconverter, see Section 3.3.5).
9. Optionally apply a finite impulse response filter (FIR) to the demodulated signals.
10. Coherently accumulate each demodulated signal to filter and downsample to the desired output rate (see Section 3.3.6).
11. Packetize I and Q for each channel into UDP frames, add timing and checksum information.
12. Stream packets to the data network.

13. In software, perform all necessary steps to calculate the FM of each resonator, and convert to absorbed optical power.

3.1.2 White Noise

The readout noise is characterized by its power spectral density (PSD), S_{XX} [x^2/Hz], which consists of additive white Gaussian noise (AWGN, hereafter WN) and 1/f ('one over f') components. The latter is hereafter referred to as flicker noise (FN). The 1/f knee, or corner-frequency f_c of the FN sets the lower limit on the telescope scan speed ω_{scan} and readout rate f_{ro} which will permit for adequate mapping of the science targets.

Array	250 μm	350 μm	500 μm
Readout Modules	3	1	1
Number of Tones per Module	∼500	∼700	∼300

Table 5. Detector counts for the three BLAST-TNG wavebands. A detailed description of the detector counts is provided in Chapter 4 of this document.

The readout must meet its noise requirements at the highest multiplexing factor set by the pixel count of each detector array. BLAST-TNG's ∼2,500 detectors are divided between five independent readout slices: One each for the 350 and 500 μm arrays, and three for the 250 μm array. Each slice contains a ROACH2 board, DAC/ADC board and a set of both commercial-off-the-shelf (COTS) and custom-made RF electronics (see Section 3.2). Although the 250 μm array could have been distributed between two readout slices, the array was designed as one chip

which contains a set of three identical rhombuses. The ∼700 pixels of the 350 μm array sets the most stringent multiplexing requirement on a single slice. Approximate detector counts for each BLAST-TNG band are listed in Table 5.

The requirement on the WN level of each readout channel is that it must be lower than that of both the LNA and detectors, at all multiplexing factors (see Chapter 2 for a detailed description of the instrumental noise hierarchy). The FM of the resonators is in the phase quadrature of the signal, while the amplifier noise is in the dissipation quadrature. The phase of I/Q is:

$$\phi_{I/Q} = \arctan2(Q, I) \tag{3.1}$$

and its PSD S_ϕ has units of $\mathrm{rad}^2/\mathrm{Hz}$. It is useful to express S_ϕ in units of dBc/Hz (decibels relative-to-the-carrier, $10\log_{10}(S_\phi)$). The amplifier and detector noise are commonly expressed using the PSD of their fractional frequency fluctuations $S_{\delta f/f}$ [1/Hz] (see Chapter 2). Following Gao (2008), S_Φ can be converted to $S_{\delta f/f}$ using a known or estimated resonator quality factor, Q_r:

$$S_{\delta f/f} = \frac{S_\phi}{16{Q_r}^2} \qquad \left[\mathrm{Hz}^{-1}\right] \tag{3.2}$$

by: $\frac{\delta f}{f} = \frac{\delta \phi}{4Q_r}$. For LEKID detectors like the ones used in BLAST-TNG, the optically loaded detector noise is $S_{\delta f/f} \sim 10^{-17}$ Hz^{-1}.

The frequency noise of the LNA $S_{\delta f/f,\mathrm{amp}}$ can be estimated as (Barry, 2014):

$$S_{\delta f/f,\mathrm{amp}} \simeq \frac{kT_N}{P_{\mathrm{ro}}} \frac{{Q_c}^2}{{Q_r}^4} \tag{3.3}$$

where k is Boltzmann's constant, T_N is the noise temperature of the LNA and P_{ro} is the power in the microwave probe tone. For BLAST-TNG, $T_N \approx 6$ K. P_{ro} ranges from -80 to -90 dBm at the detectors, and Q_r and Q_c range from ∼1–5 × 10^4. Using

these ranges in Equation 3.3 suggests that a lower limit on $S_{\delta f/f,\mathrm{amp}}$ is $\sim 5 \times 10^{-20}$ Hz^{-1}. Converting this value to a phase noise using Equation 3.2 gives a required phase noise of $\lesssim$ -95 dBc/Hz.

The WN floor for a single probe tone is set by the bit-width of the analog-to-digital converters (ADCs). Given an effective number of bits (ENOB) and ADC sampling frequency f_s, this floor is:

$$-20\log_{10}(2^{ENOB}) - 10\log_{10}(f_s/2) \qquad \left[\frac{\mathrm{dBc}}{\mathrm{Hz}}\right] \tag{3.4}$$

For an ENOB of 10-bits and $f_s = 512$ MHz (see Section 3.2), this noise floor is $\approx$-144 dBc/Hz. For a multitone comb of N tones, the WN level of each tone can be estimated as:

$$-144 \text{ dBc/Hz} + 10\log_{10}(N) + CF \tag{3.5}$$

where CF is the crest factor of the waveform:

$$CF = 20\log_{10}\frac{|V_{\mathrm{peak}}|}{V_{\mathrm{rms}}} \qquad [\mathrm{dB}] \tag{3.6}$$

Adding the contribution of the warm IF electronics to Equation 3.5, the total WN contribution from the readout system is:

$$WN_{\mathrm{read}} = -144 \text{ dBc/Hz} + 10\log_{10}(N) + CF + WN_{IF} \tag{3.7}$$

3.1.3 Flicker Noise: Readout Rate and Map Size

The FN is characterized by a PSD which is proportional to $1/f_c^{\alpha}$, where f_c is the 1/f corner, or knee-frequency, and α is the FN spectral index. Typically, $0 \leq \alpha \leq 2$.

In the following analysis, we set $\alpha = 1$. Because FN is colored noise, a correlation can be assumed to exist between some component of the FN of each readout channel. The correlated component of the noise PSDs can be subtracted using the technique of common-mode subtraction, as in van Rantwijk et al. (2016). In this work, we limit our scope to determining how the measured FN level of the readout will impact the scanning strategy of the telescope with which it will be used.

Because the telescope scan speed $\omega_{\rm scan}$ has a fixed upper limit, f_c establishes a lower limit for the readout rate, and an upper limit on map size. In the following, we take $\omega_{\rm scan} = 0.5$ deg/s, which is the approximate upper limit of the telescope scan speed for BLAST-TNG. To critically (Nyquist) sample features on the scale of the detector beams, the readout must sample I and Q at a frequency of at least

$$f_{\rm ro} = 2\omega_{\rm scan}/\theta_{\rm beam} \tag{3.8}$$

where $\theta_{\rm beam}$ is the beam diameter ($\approx 30''$ for the BLAST-TNG 250 μm band). It is preferable to more than critically sample each beam. For $\omega_{\rm scan} = 1$ deg/s, a twice-critically sampled beam requires that $f_{\rm ro} \geq 480$ Hz.

The largest spatial scale that can be mapped before the statistics of the noise deviates too far from WN corresponds to f_c. For example, with $f_c = 0.5$ Hz, the largest spatial scale that can be mapped is $w_{\rm scan}/f_c = 1$ deg (corresponding to a map area of 1 deg^2). Consequently, smaller maps are better for imaging fainter sources (e.g., diffuse Galactic emission and CMB foregrounds).

The upper limit on map size set by f_c in turn constrains the amount of time which is required to produce maps at various spatial scales with 1-σ error bars on either the polarization fraction p or total intensity of a specified percentage. To see this, we write $NEP_{\rm phot}$ (W/$\sqrt{\rm Hz}$) as a function of scan frequency, $f_{\rm scan}$:

$$
\begin{aligned}
NEP^2_{\rm phot}(f_{\rm scan}) &= 2h\nu P_{\rm opt}\left(\frac{1+m\eta_{\rm det}}{\eta_{\rm det}}\right)(1+f_c/f_{\rm scan}) \\
&= \frac{2h\nu}{\eta_{\rm det}}P_{\rm opt}\,(1+f_c/f_{\rm scan}) \qquad \left[\frac{\rm W}{\sqrt{\rm Hz}}\right]
\end{aligned}
\tag{3.9}
$$

where $f_{\rm scan}$ is determind by the telescope scan speed and a spatial scale $L_{\rm spatial}$ of interest, $f_{\rm scan} = L_{\rm spatial}/\omega_{\rm scan}$, m is a photon-bunching (wave noise) term that can be omitted at the BLAST-TNG observation bands (see Chapter 2), and $\eta_{\rm det}$ is the detector efficiency ($\sim$80%). The optical power $P_{\rm opt}$ corresponds to a source of flux density (intensity) $I_{\rm source}$ [MJy/sr]:

$$
P_{\rm source} = I_{\rm source}A_{\rm prim}B_{\rm opt}\epsilon_{\rm opt}p \tag{3.10}
$$

where $A_{\rm prime}$ is the collecting area of the primary telescope ($\approx$ 4.5 m^2, for BLAST-TNG), $B_{\rm opt}$ is the optical bandwidth ($\approx 0.3\nu_{\rm center}$), $\epsilon_{\rm opt}$ is the system optical efficiency ($\approx$ 30%), and p is the polarization fraction of the source ($\approx$ 2%). Using these approximate values as examples, a source with $I = 200$ MJy/sr imaged in the BLAST-TNG 250 μm band corresponds to $P_{\rm opt} \approx 2.58 \times 10^{-16}$ W.

Assuming the detectors are photon noise limited, the time required to map an area of $A_{\rm map}$ with a root-mean-square (RMS) uncertainty on the polarization fraction σ_p is (for a single pixel):

$$
t_{\rm req,FN} = 2\frac{A_{\rm map}}{A_{\rm beam}}\frac{NEP_{\rm phot}}{\sigma_p P_{\rm opt}} \tag{3.11}
$$

where $A_{\rm beam}$ is the beam area, and the factor of two comes from the fact that $\sigma_p = 2\sigma_I/I$, in the case of photon noise limited operation. With $fc = 0.5$ Hz, and the values given above, Equation 3.11 shows that to map scales of $L_{\rm spatial} = 1$ deg over a 1 deg^2 map with $\sigma_p = 0.5\%$ in the 250 μm band will take $\sim$1.35 hr.

3.1.4 Multiplexing Factor

Ideally, the readout system would have a multiplexing factor which is as high as desired. However, the readout electronics have a limited signal processing bandwidth, which is set by the digital electronics (FPGA/DAC/ADC). For the ROACH2 system, the DAC/ADC board is clocked at $f_{\mathrm{DAC}} = 512$ MHz, and the FPGA is clocked at $f_{\mathrm{FPGA}} = 256$ MHz. By virtue of using quadrature signals, the system bandwidth is equal to f_{DAC}.

The number of detectors which can be readout on a single slice is a function of the signal processing bandwidth and resonator quality factors. Assuming that the resonators are equally spaced, a spacing of 500 kHz would allow for ∼1,000 channels per slice. In practice, the LEKID resonant frequencies are not evenly spaced, and are sometimes within 100 kHz of each other. The readout architecture is designed to allow for this scenario (see Section 3.3.5).

3.1.5 Data Rate

Both the rate at which data is saved to disk and the total amount of data which must be saved over the duration of a science observation are important considerations. The packet routing network on the telescope must be able to handle the throughput of each readout channel, and there must be sufficient disk space to store all of the data. The ROACH2 data packets are in User Datagram Protocol (UDP) format. Regardless of the number of active channels, the data packet size is fixed at 8234-B. In BLAST-TNG there are five ROACH2 slices, and packets are output at $f_{\mathrm{packet}} = 488.28125$ Hz. The total data rate is therefore:

$$N_{\mathrm{slices}} \times f_{\mathrm{packet}} \times 8234 = 20.1 \qquad \mathrm{MB/s} \tag{3.12}$$

or ≈ 4 MB/s per readout slice. Over a 24 hr period, the total data payload for BLAST-TNG is ≈ 1.7 TB. Therefore, a 28 day balloon flight requires a storage capacity of at least 50 TB (with no backup storage).

3.1.6 Timing and Packet Synchronization

During a telescope observation, each readout data packet must be timestamped in a way that allows for the detector signals to later be synchronized with the timing information of the other telescope subsystems. Most critically, in order to be able to make maps, it must be possible to align the data packet timestamps with those of the pointing system.

In BLAST-TNG, a GPS card provides the flight computers with the absolute time, which is accurate to within a few microseconds. The BLAST-TNG data frames which are constructed by the flight software contain subsystem data which is sampled at various rates (e.g., 1 s, 2 s, 5 s), with the highest rate being that of the ROACH2 packets ($f_s = 488.28125$ Hz). To be able to make accurate science maps, the information from each frame must be synchronized to within a few milliseconds.

To timestamp the ROACH2 data packets, the GPS-disciplined PPS is input into the FPGA and used to flag packets which correspond to the rising edge of each pulse (within one FPGA clock cycle; a period of ≈ 2 ns). All packets which occur between PPS boundaries are stamped with the number of FPGA clock cycles which have elapsed since the previous pulse. The FPGA clock itself is disciplined by an external 10 MHz reference, which is split between each ROACH2 slice. In addition

to the relative timestamp information, an absolute timestamp (containing the flight computer time) can periodically be added to a packet on a PPS boundary.

3.1.7 Power

Given an unlimited power supply, the primary requirement regarding power dissipation is that the readout system must be able to operate without melting itself, or any other subsystem. However, from a SWaP-C standpoint it is desirable to minimize the power-per pixel to as low a level as possible. On a space-based or balloon-borne platform, the electronics must be able to passively cool.

By virtue of flying during the Antarctic summer, the BLAST-TNG gondola is provided with 24 hr sunlight. Its total power budget is $\sim$1 kW. Using a thermal model of the BLAST-TNG gondola (see Gordon et al. (2016)) it was determined that each readout slice should dissipate $\lesssim$ 65 W. This places the power-per pixel at $\sim$65 mW/pixel.

3.2 Readout Electronics

The readout electronics for a single readout slice consist of digital and both active and passive analog components. Depending on the system architecture (e.g., its space constraints and thermal considerations), the digital and intermediate frequency (IF) components can be housed in either the same or separate enclosures. For the BLAST-TNG balloon platform, the decision was made to house all five digital and electronics

slices in the same enclosure, which is known to the collaboration as the 'ROACH2 Motel'[7].

3.2.1 Digital Electronics

The digital electronics contain a ROACH2 Virtex-6 FPGA board coupled to a MUSIC DAC/ADC board (Duan et al., 2010). An embedded processor (AMCC PowerPC 440EPx, hereafter PPC) acts as an interface between the PPC/FPGA and data acquisition (DAQ) or flight computer (FC). The PPC runs a daemonized Karoo Array Telescope Protocol (KATCP) server, which facilitates communications between the DAQ and FPGA[8]. The MUSIC board includes two 12-b 550 Msps ADC chips[9], and two 16-bit 1,000 Msps DAC chips[10]. Each DAC and ADC is designated to either the I or Q signal quadratures.

Communication between the PPC/FPGA and DAQ/FC is facilitated by one-gigabit (1GbE) Ethernet transceivers (one for the PPC and one for the FPGA). The PPC link is half-duplex. It is used for programming the FPGA firmware, commanding the firmware during readout operation, and receiving small quantities of diagnostic data from various points in the DSP chain using CASPER 'snap' blocks which are

[7]BLAST-TNG is not the first experiment to use the name 'ROACH Motel' for its collection of ROACH boards. See: https://casper.ssl.berkeley.edu/wiki/ROACH_Motel

[8]KATCP has been developed by the Square Kilometer Array South Africa (SKA SA) collaboration for use on their CASPER hardware-based correlators and beam formers. See https://casper.berkeley.edu/wiki/KATCP

[9]ADS54RF63, Texas Instruments Inc.

[10]DAC5681, Texas Instruments Inc.

built into the firmware (see Section 3.3). The FPGA Ethernet link is one-way, and is used stream the User Datagram Protocol (UDP) data packets to the DAQ.

The clock for the DAC/ADC board and FPGA is driven by one channel of a dual-channel voltage-controlled oscillator (VCO) board (Valon 5009 Synthesizer[11]). The DAC/ADC board is driven at 512 MHz, and the FPGA divides this frequency by two to produce a 256 MHz clock signal. The second channel of the Valon 5009 is used as a tunable local oscillator (LO) for the quadrature modulator and demodulator (see below). The Valon 5009 has an internal frequency reference, but also accepts an external reference. In order to synchronize the clock phase of each of the readout slices used in the BLAST-TNG ROACH2 Motel, each Valon is driven by an external reference which is generated by an OctoClock CDA-2990[12]. The Octoclock also provides a GPS-disciplined PPS signal, which is input to the ROACH2 'sync-in port' via SMA, and used to timestamp each data pacekt (see Section 3.3.6.1).

To allow for control of the Valon synthesizer and programmable attenuators (see Section 3.2.2), the BLAST-TNG electronics incorporate a Raspberry Pi 3[13] single-board computer in each readout slice. The programmable attenuators are also powered via USB through the Pis. Having a Pi in each readout slice reduces the amount of cable routing which is required between the electronics and DAQ/FC.

Table 6 lists the measured power dissipation of major components on the ROACH2 and MUSIC board with the firmware loaded and running (values taken

[11] 5009 Dual-Frequency Synthesizer, Valon Technology Inc.:https://www.valonrf.com

[12] Ettus Research: https://www.ettus.com

[13] Raspberry Pi:https://www.raspberrypi.org

from Gordon et al. (2016)). At 30 W, the FPGA is the dominant power consumer in the readout system. The total power dissipation for the digital electronics is ≈47.8 W.

Component	Quantity	Power Dissipation (W)
FPGA	1	30
Power PC	1	5.0
RAM	1	4.4
ADCs	2	2.6
QDR	4	1.8
PHY	2	1.0
DACs	2	0.5
Raspberry Pi 3	1	2.5
Total		47.8

Table 6. The power budget for the digital electronics used in each BLAST-TNG readout slice.

3.2.2 Intermediate Frequency Electronics

The primary purpose of the IF electronics is to move the signals which are produced by the ROACH2 system between complex baseband (-256 to +256 MHz) and RF (∼500–1100 MHz). Additionally, the signals require amplification after passing through the cryostat. The IF electronics in each readout slice include the following components[14]:

- Four 1–256 MHz anti-aliasing filters (custom-made at ASU, one each for the *I* and *Q* DACs and ADCs)

[14]Not included in the list: Several meters of coaxial cable; many handfuls of SMA adaptors.

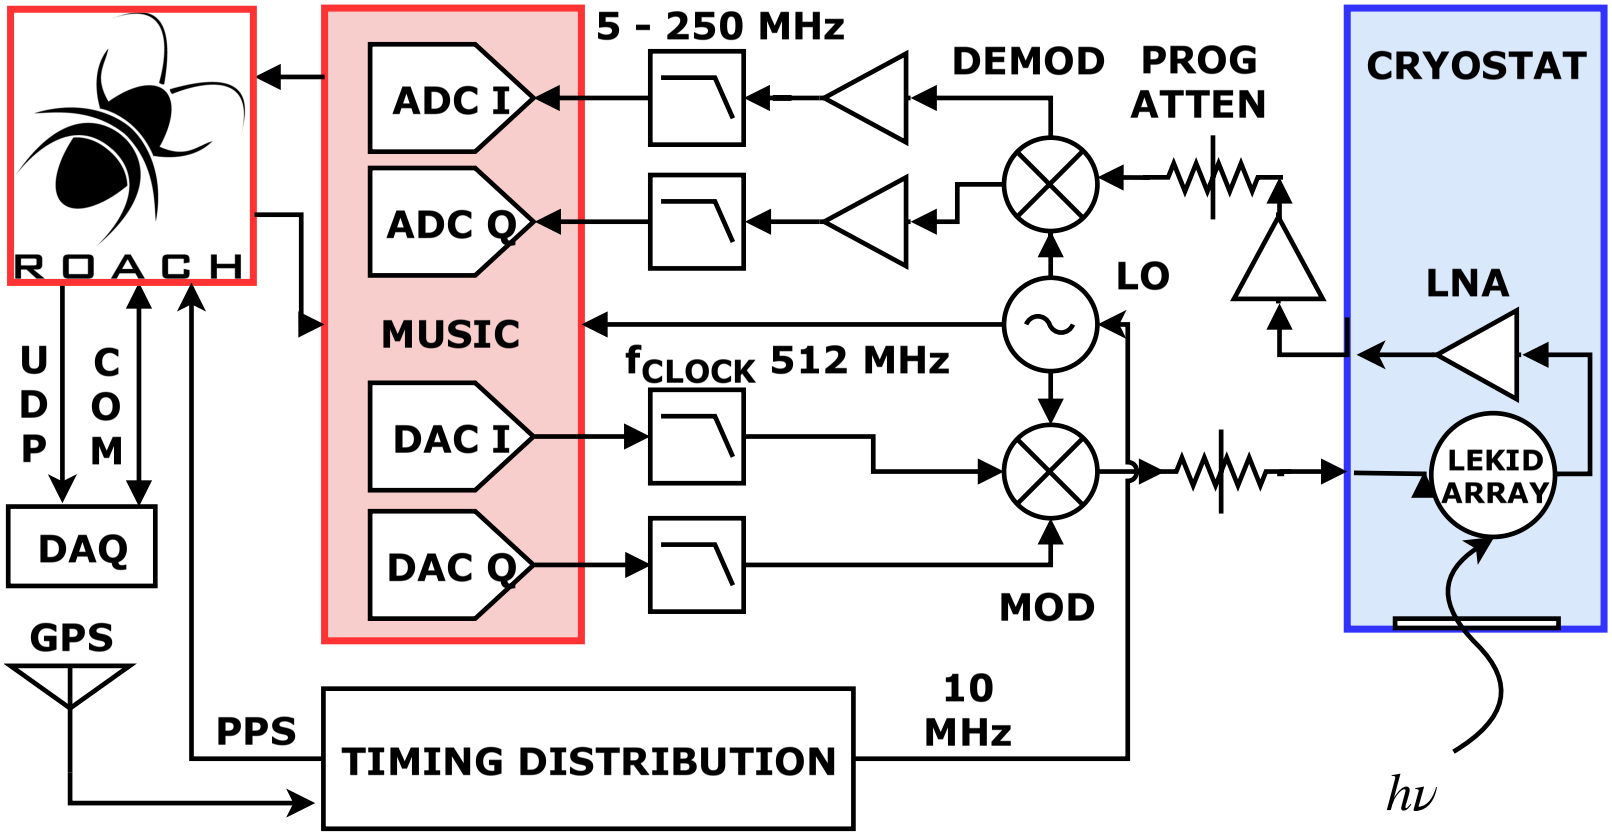

Figure 28. A schematic of the BLAST-TNG readout electronics.

- One quadrature modulator (Polyphase Microwave AM0350A[15])
- One quadrature demodulator (Polyphase Microwave AD0105B)
- Two Mini-Circuits[16] 6000–30 programmable attenuators (steppable between 0–30 dB in 0.25 dB increments)
- One room-temperature (RT) LNA, (custom-made at ASU)
- Two RT baseband amplifiers (custom-made at ASU, one each for I and Q on the downconversion side.)
- One dual-channel synthesizer (the Valon 5009)
- Two 1 GHz low-pass filters (Mini-Circuits VLF-1000+, used to filter out LO harmonics and at the RF input of the readout system).
- Two baluns (Mini-Circuits ADT2-1T+, used to convert the single-ended I/Q modulator inputs from single-ended to differential.

A schematic of the digital and IF electronics signal processing chain is shown in Figure 28. For explanation purposes, the signal chain can be split into two chains which merge on one end at the ROACH2 system and at the cryostat (or detector arrays) at the other.

Upconversion chain: Starting with the ROACH-2 board, the firmware generates a baseband probe tone comb containing the resonant frequencies of each detector in the LEKID array. The signal is output in quadrature (I and Q) by the DACs, and low-pass filtered to remove higher-order Nyquist zones (see Section 3.3.1). The signal is then upconverted to RF in the quadrature modulator, which takes the I/Q quadratures as differential inputs and outputs a single RF signal. To be input into

[15]Polyphase Microwave:https://polyphasemicrowave.com/

[16]Mini-Circuits:https://www.minicircuits.com/

the modulator, the *I* and *Q* signals are converted to differential using baluns[17]. The upconverted signal is then attenuated by the output-side (relative to the ROACH2) programmable attenuator before being fed into the cryostat on a single coaxial cable. The cryostat contains several temperature stages between 300 K and 300 mK, where the detector arrays are kept. Each stage contains a fixed amount of attenuation whose purpose is to reduce the thermal load on the following (colder) stages. The total gain of the cryostat depends on fixed attenuation, cable losses and LNA gain. For BLAST-TNG, the cryogenic gain is ∼-20 dB. After being modulated by the LEKID detectors the signal is amplified by a cryogenic 6 K LNA[18] with 30–35 dB of gain, and sent back out of the cryostat.

Downconversion chain: The downconversion chain begins when the signal exits the cryostat. It is first low-pass filtered (cutoff frequency at ∼1 GHz) to remove any signal above ∼1 GHz which was amplified by the high gain of the cryogenic LNA. The signal then enters a second-stage amplifier, which for BLAST-TNG is a RT LNA which has been modified to operate on 5 VDC. The second-stage amplifier is followed by the programmable input attenuator, which feeds the demodulator. The demodulator splits the signal back into quadratures, which are passed through a third-stage of amplification that provides ∼20 dB of gain. The purpose of the third-stage amplifiers is to ensure that enough signal power enters the ADC when the probe tones are on resonance. The amplified quadratures are then low-pass filtered through the second pair of 1–256 MHz anti-aliasing filters, and input to the *I* and *Q* ADCs.

[17]The baluns are not required for this particular modulator. However, it is desirable to have as little LO leakage as possible at the RF output. Using baluns reduces the LO leakage to ∼4 dB above the probe tone level. When driven single-ended, the LO is ∼20 dB above the probe tones.

[18]A SiGe LNA from Groppi Labs at ASU: http://thz.asu.edu/products.html

The power dissipation for each of the active IF components is listed in Table 7. Their total power dissipation is ≈ 26.7 W. Adding this value to the ≈ 48.7 W which are dissipated by the digital electronics (Section 3.2.1) gives a total power dissipation of ≈ 74.5 W. This per-readout slice power dissipation is within acceptable range of the BLAST-TNG power requirement of $\lesssim$65 W per slice.

Part	Quantity	Volts (V)	Current (mA)	Total Power (W)
Valon 5009	1	6	600	3.6
RUDAT 6000-30	2	5	60	6
Polyphase AM0350A	1	+5/-5	250/30	2.75
Polyphase AD0540B	1	+5/-5	290/50	3.95
RT LNA	1	5	30	1.5
Baseband RT Amps	2	5	145	1.45
Total				26.7

Table 7. The power dissipation for the IF electronics used in each BLAST-TNG readout slice

3.2.3 Readout Electronics for the Balloon Platform: The BLAST-TNG ROACH2 Motel

The ROACH2 Motel, shown fully assembled in Figure 29, is a custom aluminum (Al) enclosure that houses the set of five ROACH2 readout slices, including their digital and IF electronic components. The electronics for each readout slice are mounted to 1/4″ backing plates, which are stacked side by side between two 5/8″ thick Al side panels. To allow for continuous operation at float altitude (∼35 km), where fans have little effect due to the lack of atmosphere, the enclosure itself must provide a thermal link to the inner frame of the balloon gondola to enable the heat generated

by the electronics to dissipate and radiate to space. Heat from each slice's backing plate flows through the two Al side panels and into two 8″ × 5″ × 1/4″ right angle brackets which secure the enclosure to the inner frame of the gondola (see Figure 30). All contact joints between metal components within the enclosure are coated with non-conductive thermal joint compound to increase their thermal conductivity.

Figure 29. The BLAST-TNG ROACH2 Motel, fully assembled and taking in the sunset in Palestine, Texas, USA.

Each of the electronics used in the BLAST-TNG detector readout are rated for operation to at least 85°C. In order to stay within their operational temperature limits, the FPGA, PPC and ADCs must be well heat sunk to the inner plates of the enclosure, which in turn must have their own path for dissipating heat. The ROACH2's Xilinx

Virtex-6 FPGA is by far the largest producer of heat in each electronics slice. To address this, a copper assembly containing two 5 mm diameter water filled, sintered copper heat pipes[19] are mounted directly on top of the FPGA chips. The heat pipes are soldered into the heat sink assemblies using bismuth tin (BiSn) solder paste. BiSn is used because its relatively low melting point of 138˚C allows it to liquify during fabrication without damaging the heat pipes.

Besides the FPGA, the two most sensitive components are the PPC and ADCs. The PPC is heat-sunk to the FPGA's heat pipe assembly using two smaller copper heat pipes, and the the ADCs are heat strapped directly to the 1/4″ Al backing plates using 10 AWG copper wires which are secured to the chips using non-conductive thermal epoxy. The DACs, which run cooler than the ADCs, are connected to the backing plate by a single 14 AWG copper wire. The Raspberry Pi 3 boards are mounted to Al plates which are secured to the backing plate of each slice. The IF electronics are mounted to the same backing plate as the digital electronics.

During operation of the system, the temperature of the FPGAs and PPCs is logged via temperature sensors which are secured to the heat sinking assemblies of each device. The CPU temperatures of each Raspberry Pi 3 board are also logged.

Power suppy and input and output signals are routed through the front and back panels of the ROACH2 Motel. Each slice's front panel includes SMA input ports for: A 10 MHz reference, a PPS signal, an (optional) external LO, RF input and output, and a spare input. The back panel of each slice contains a 4-pin military connector which receives 28 VDC from the balloon gondola's power distribution system. This

[19]Enertron Inc:https://www.enertron-inc.com/

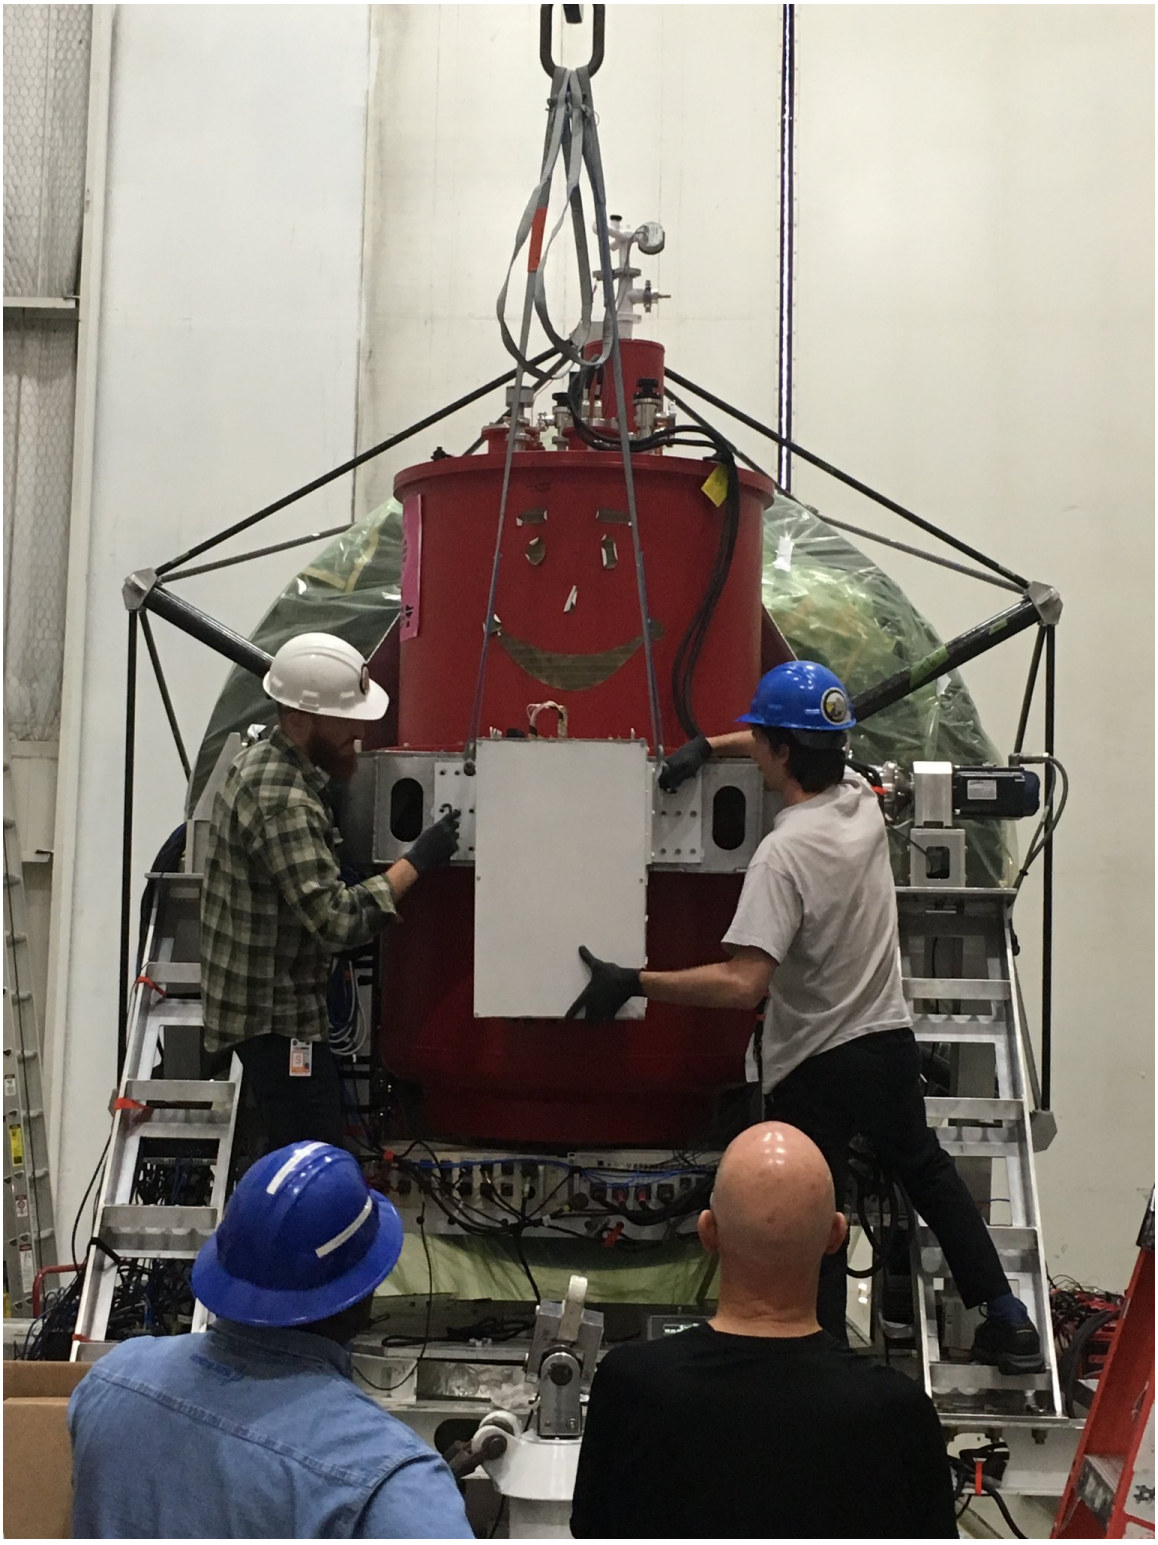

Figure 30. The 'Roach Boys' installing the Roach2 Motel on the BLAST-TNG gondola in Palestine, Texas, USA (July, 2018).

voltage is distributed to several Vicor[20] DC-DC converters which are mounted to the enclosure's backing plate. The DC converters supply ±5, 6 and 12 VDC to the IF electronics. The back panels also contain Ethernet port for communicating with the PPC, FPGA, and Raspberry Pi 3, as well as a USB port for interfacing directly with the Linux OS (BORPH) which runs on the PPC.

Figure 31 is an annotated photograph of one of the five BLAST-TNG ROACH2 slices, taken during the 2018/2019 Antarctic campaign. Critical components are labeled with numbers. These are: 1) ROACH2 board 2) FPGA heat pipe assembly 3) Raspbery Pi 3 4) MUSIC DAC/ADC board 5) anti-aliasing filters 6) programmable attenuators 7) demodulator 8) modulator 9) second-stage LNA (not visible) 10) Valon synthesizer (not visible). A view of the fully assembled motel is shown in Figure 32. The outside of the enclosure is painted white to increase its albedo. Ethernet ports for the FPGA and PPC can be seen forming two columns that run along the back panel. The power distribution for the exposed ROACH2 slice is routed through a D-sub connector at the lower left corner of the image.

[20]Vicor:http://www.vicorpower.com

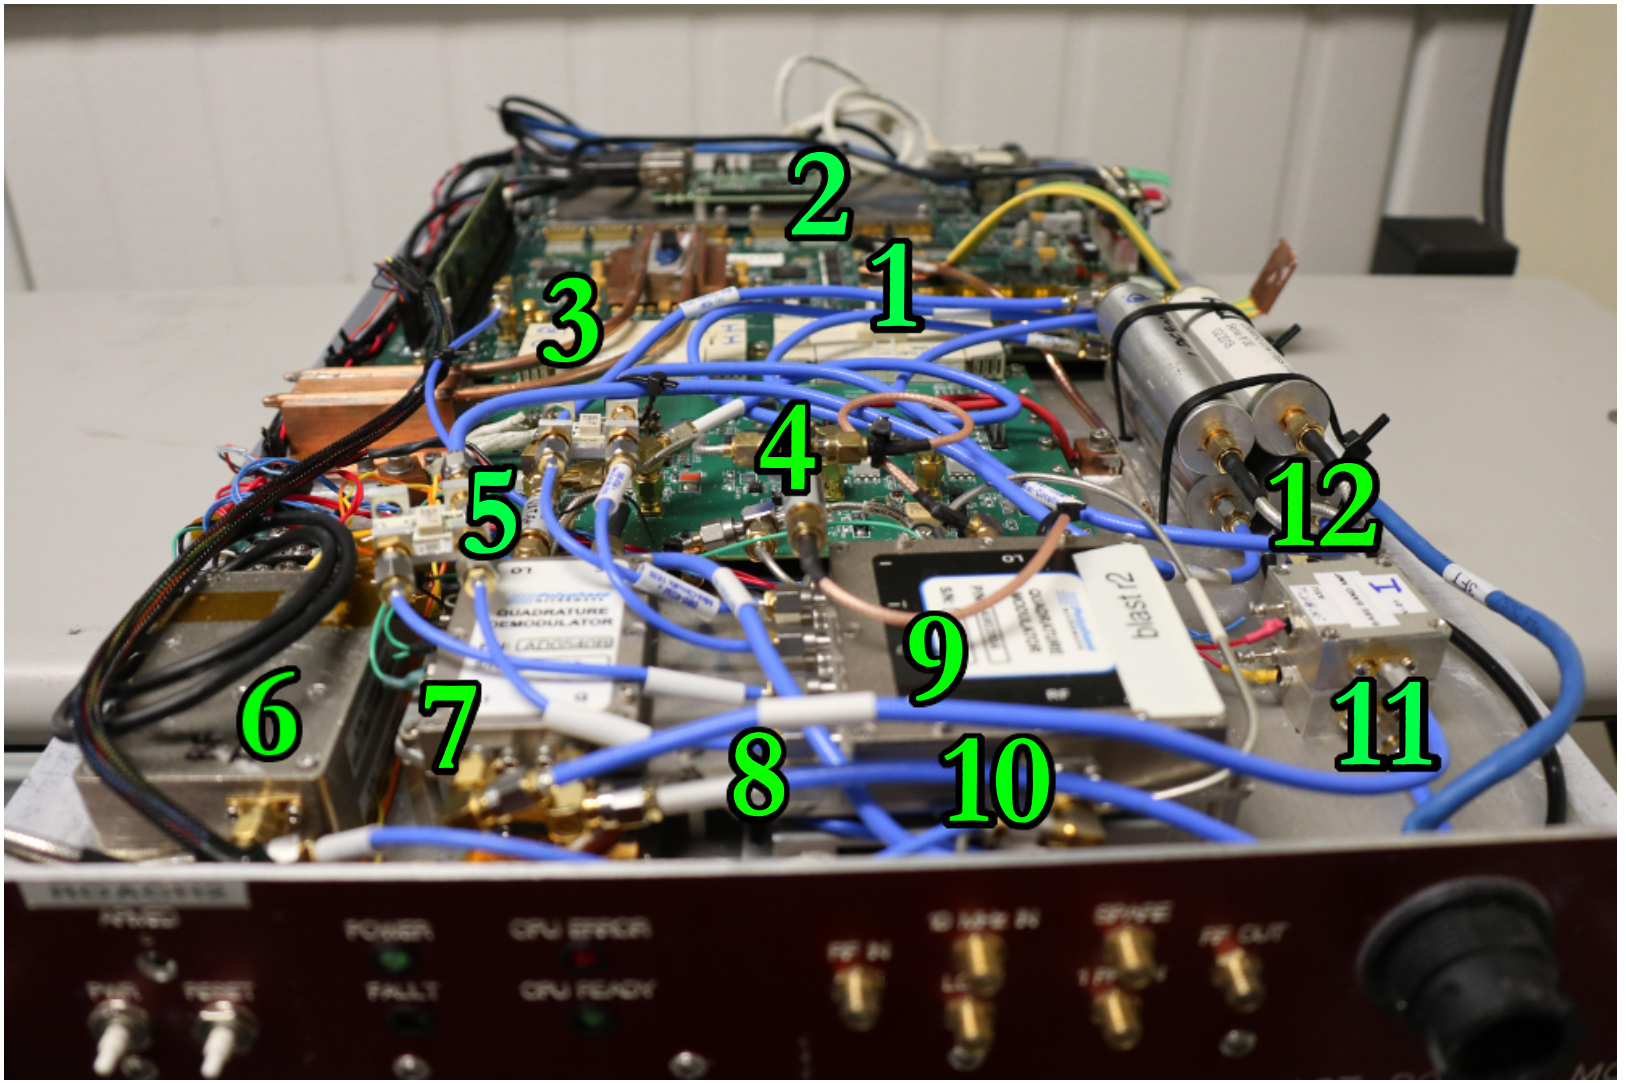

Figure 31. An annotated photograph of one of the five BLAST-TNG ROACH2 electronics slices. Labeled components are: 1) ROACH2 board 2) Raspberry Pi 3 3) FPGA heat pipe assembly 4) MUSIC DAC/ADC board 5) baluns 6) programmable attenuators 7) demodulator 8) RT LNA (not visible) 9) modulator 10) Valon synthesizer (not visible) 11) third-stage amplifiers 12) anti-aliasing filter stack.

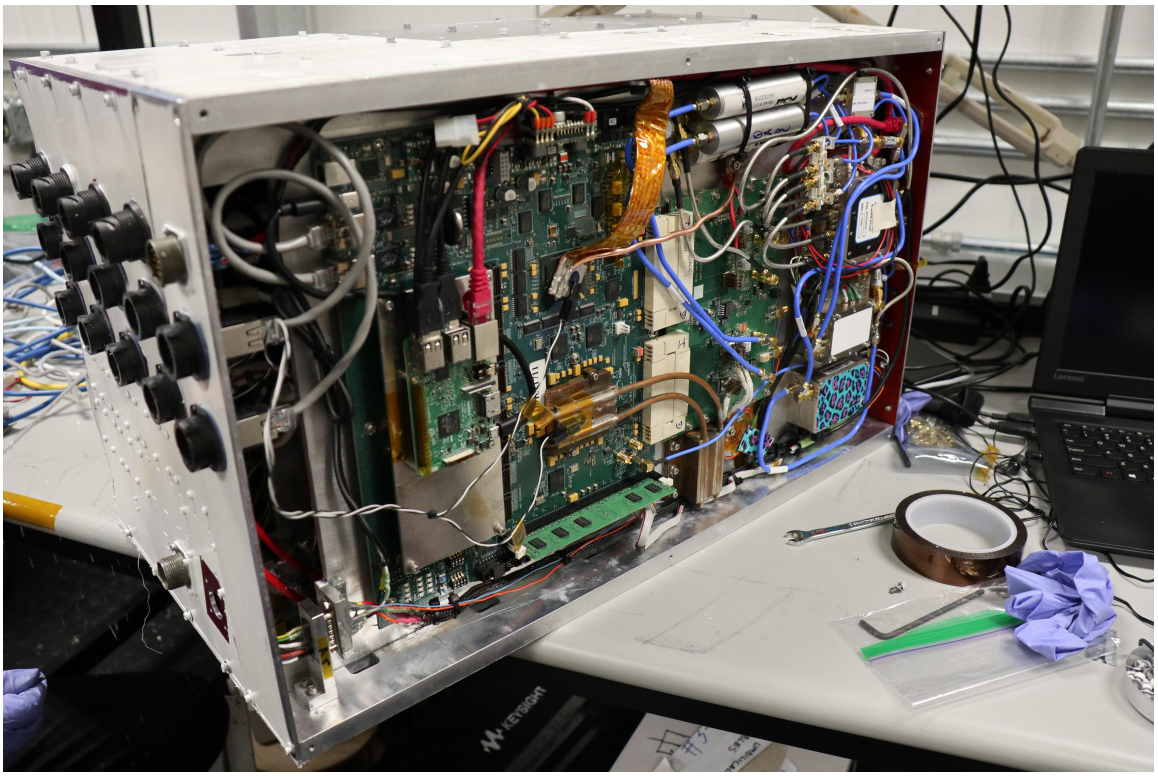

Figure 32. A top view of the BLAST-TNG ROACH2 Motel with the lid removed.

3.3 ROACH2 Firmware Overview

In this section we describe in detail the DSP which is performed in the FPGA firmware. The firmware is written using the MATLAB[21] / Simulink[22] / System Generator[23] / EDK[24] (MSSGE) Toolflow developed by the CASPER collaboration. CASPER 'snap' blocks allow for pre-specified amounts of data from the data stream

[21] Mathworks:https://www.mathworks.com/

[22] Mathworks

[23] Xilinx ISE 14.7 Design Suite

[24] Xilinx Embedded Development Kit

to be saved in block RAM (BRAM) on the FPGA at key points in the DSP chain. This data can then be converted into figures of merit to be used for making real-time adjustments to either the RF electronics or baseband frequency combs. Critical parameters, including the readout bandwidth, may be adjusted during operation using software inputs.

The DSP chain performs several synthesis and analysis functions, which are described sequentially in the sections below. A block diagram of the DSP chain is shown in Figure 33.

3.3.1 Tone Comb Synthesis

The DSP chain begins with the synthesis of the DAC probe tone and DDC tone look-up-tables (LUTs). The DDC tones are used to digitally downconvert the output timestreams of the PFB-FFT (see Section 3.3.5), and the DAC probe tones circulate between the readout system and the LEKID arrays throughout an observation. If the resonator frequencies are not known, they can be discovered using the channel identification technique described in Section 4.1.3. In the following, we assume that the LEKID resonant frequencies (f_{res}) are already known.

The time domain waveform of the DAC probe tone is synthesized in software (`Python/C/C++`) by assembling a dummy spectrum in the frequency domain, and then taking its inverse-Fourier transform (IFT). In the dummy spectrum, each probe tone is represented by an amplitude A and unique phase ϕ. During an observation, the probe tone amplitudes are set to unique values corresponding to the specific tone power requirements of the detectors (including a global transfer function correction to correct for the ROACH2-output side IF electronics, described in Section 4.1.2). In

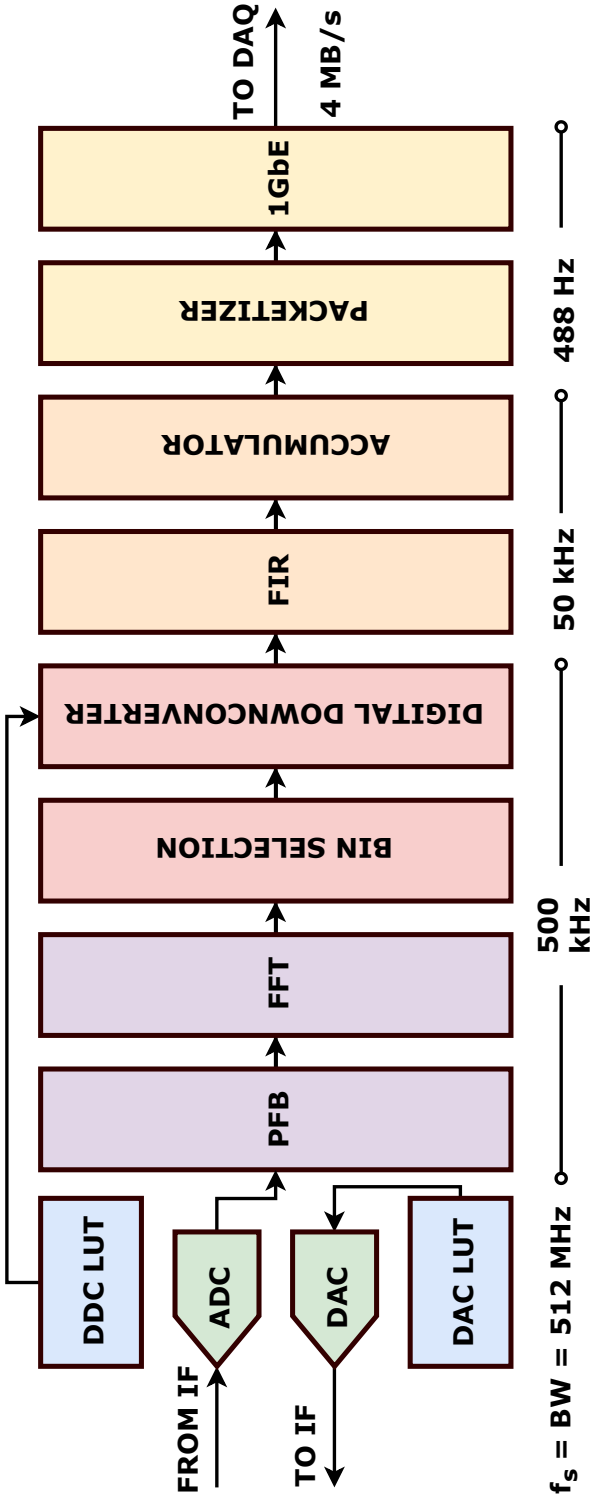

Figure 33. A schematic of the firmware DSP chain.

the examples which are provided in the following, the tone amplitudes are normalized to their maximum allowable value, which is 32767 analog-digital units (ADU). At the DAC output, this level corresponds to $V_{pp} \simeq 1.4$ V.

The phases of the probe tones are chosen from a normal distribution. This produces a near-optimal CF for a multitone comb with an arbitrary number of tones (Boyd (1986)). The tone frequencies occupy the complex baseband bandwidth of the ROACH2/MUSIC board system, which spans from -256 to +256 MHz (for a good explanation of negative frequency, see Lyons (2004)). The tone frequencies are rounded to increments of the final channelizer bandwidth of 488.28125 Hz.

The DDC tones are computed in a similar fashion as the DAC tones, although their frequencies correspond to the difference between the DAC probe tone frequencies and the center frequencies f_k of their firmware FFT bins. For a given FFT bin index k, $f_k = kf_s/N_{\mathrm{FFT}}$. The frequency of the DDC tone corresponding to a probe tone at frequency f_{pr} which is located in bin k is therefore $f_{k,\mathrm{DDC}} = f_{\mathrm{pr}} - f_k$. The sampling frequency for the DDC LUT is the PFB-FFT output frequency, $f_{\mathrm{FFT}} = 500$ kHz.

Once synthesized in software, the DAC and DDC comb LUTs are uploaded to the ROACH2 board, where they occupy two of the ROACH2's four quad-data-rate (QDR) SRAM chips[25], hereafter designated QDR_I and QDR_Q. Each LUT contains 2^{21} signed 16-bit time ordered samples.

The QDR is logically accessed as 2^{19} addresses $\times$ 16-B. The KATCP protocol is used to upload the LUT data to QDR RAM. To facilitate uploading the two LUTs to each QDR, the *I* and *Q* components are interwoven into two separate LUTs (LUT_I, LUT_Q) of 2^{22} time ordered samples each (this is based on the method used

[25]Cypress, CY7C2565KV18

by McHugh et al. (2012)). For example, the order of samples contained within LUT_I is:

$$I^1_{\mathrm{DAC}}, I^0_{\mathrm{DAC}}, I^1_{\mathrm{DDC}}, I^0_{\mathrm{DDC}}, \ldots, I^n_{\mathrm{DAC}}, I^{n-1}_{\mathrm{DAC}}, I^n_{\mathrm{DDC}}, I^{n-1}_{\mathrm{DDC}} \tag{3.13}$$

where superscripts refer to even and odd numbered samples, and each value is 16-bit wide.

After the LUTs have been uploaded to the QDRs, they are read from the QDR buffers, sliced into their original 16-bit values and re-cast as fixed-point 16.15 values[26]. On each clock cycle, four consecutive samples are read out from QDR_I (QDR_Q): $I^1_{\mathrm{DAC}}, I^0_{\mathrm{DAC}}, I^1_{\mathrm{DDC}}, I^0_{\mathrm{DDC}}$ (same for Q).

The DDC LUT samples are sent directly to the DDC section of the firmware, while the DAC LUT samples are input to the MUSIC board DAC gateware. Since the DACs are clocked at twice the rate of the FPGAs, two consecutive samples (e.g., $I^1_{\mathrm{DAC}}, I^0_{\mathrm{DAC}}$) are processed on each FPGA clock cycle. To ensure proper synchronization of each quadrature component, the DAC is synchronized by a pulse which simultaneously resets the QDR address counter.

Because the DAC waveforms are reconstructed using a zero-order hold (ZOH), the frequency response of the waveform is a sinc function: $\mathcal{X}(f) = |\mathrm{sinc}(\pi f / f_s)|$. The sinc response produces two effects which must be considered individually. The first effect is a roll-off of $\sim$6 dB between the center of the band and the edge of the first Nyquist zone, at $f_{\mathrm{FPGA}} = 256$ MHz. This roll-off is dealt with by applying an inverse transfer function correction to the probe tone amplitudes.

[26] In this fixed-point notation, the first number represents the total bit width, with the second number being the radix point. All numbers in this notation are assumed to be signed unless otherwise noted.

The second consequence of the ZOH is that several higher order Nyquist zones are output by the DAC, which must be removed using high-order low-pass filters. Figure 34 shows the measured output of one DAC for a tone comb comprised of 1,000 evenly spaced tones, ranging from -245–245 MHz. The blue trace is the unfiltered frequency response out to 1.5 GHz. The power in the second Nyquist zone (f_s—$\frac{3}{2}f_s$) is ∼10 dB below the peak at $f \simeq 250$ MHz. The orange trace shows the simulated sinc response. The simulation is a good match the first Nyquist zone, but does not taper off at higher zones (no frequency dependent attenuation factor was included in the simulation). The red trace shows the DAC output after the application of the anti-aliasing filters, which cut off at ∼245 MHz.

3.3.2 Digitization

After the tone comb has been upconverted by the IF electronics, passed through the detector arrays, and once again downconverted to complex baseband, the I and Q signals are digitized by two independent ADCs which share the same clock source as the DACs. Each ADC outputs two consecutive time samples per FPGA clock cycle: $\mathrm{I}^1_{\mathrm{ADC}}$, $\mathrm{I}^0_{\mathrm{ADC}}$ ($\mathrm{Q}^1_{\mathrm{ADC}}$, $\mathrm{Q}^0_{\mathrm{ADC}}$). The time ordered I/Q pairs are concatenated and sent to the first stage of channelization, which is the PFB-FFT. CASPER snap blocks are utilized at this stage to store some of the ADC timestream in BRAM, which can be downloaded to the DAQ at any time in order to calculate V_{rms} at the ADC input for both I and Q. The ADCs have a full-scale input range of 1.1 V, which corresponds to a maximum allowable input power of 10.8 dBm.

Figure 35 shows 2,000 samples of the digitized timestream of a single tone at $f_{\mathrm{pr}} = 50.0125$ MHz, as well as for a tone comb with $N_{\mathrm{tones}} = 1000$ with evenly

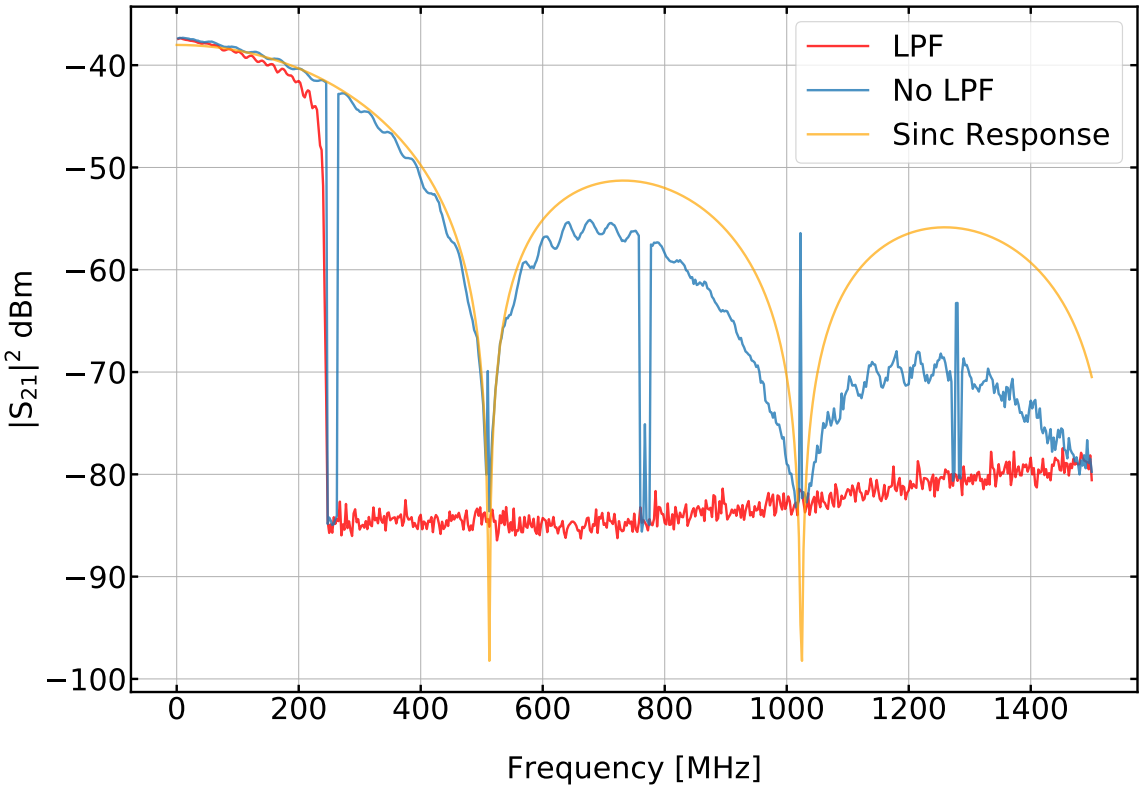

Figure 34. The DAC frequency response for a multitone probe comb with $N = 1000$, showing the unfiltered (blue), simulated (orange) and measured filtered (red) response out to 1.5 GHz.

spaced frequencies. The power in the single tone has been attenuated using the output attenuator inside the IF electronics, to avoid saturating the ADCs. Because the phases of the multitone comb are randomized, the peak power in the comb scales as $P_{\max} \propto \sqrt{N_{\text{tones}}}$. During an observation, the IF attenuators are adjusted to keep the RMS voltage of the tone comb to below ≈ 100 mV. This ensures that the ADC will not saturate as the LEKID resonant frequencies drift from the initial locations of the probe tones, increasing their transmission through the system.

The CF of the waveform can also be measured from a small number of ADC samples, for both I and for Q. For the data shown in Figure 35, the CF is measured

as 3 dB for the single tone, and 11.3 dB for the 1,000 tone comb. The real-time CF measurements are useful for making noise estimates (see Section 3.4).

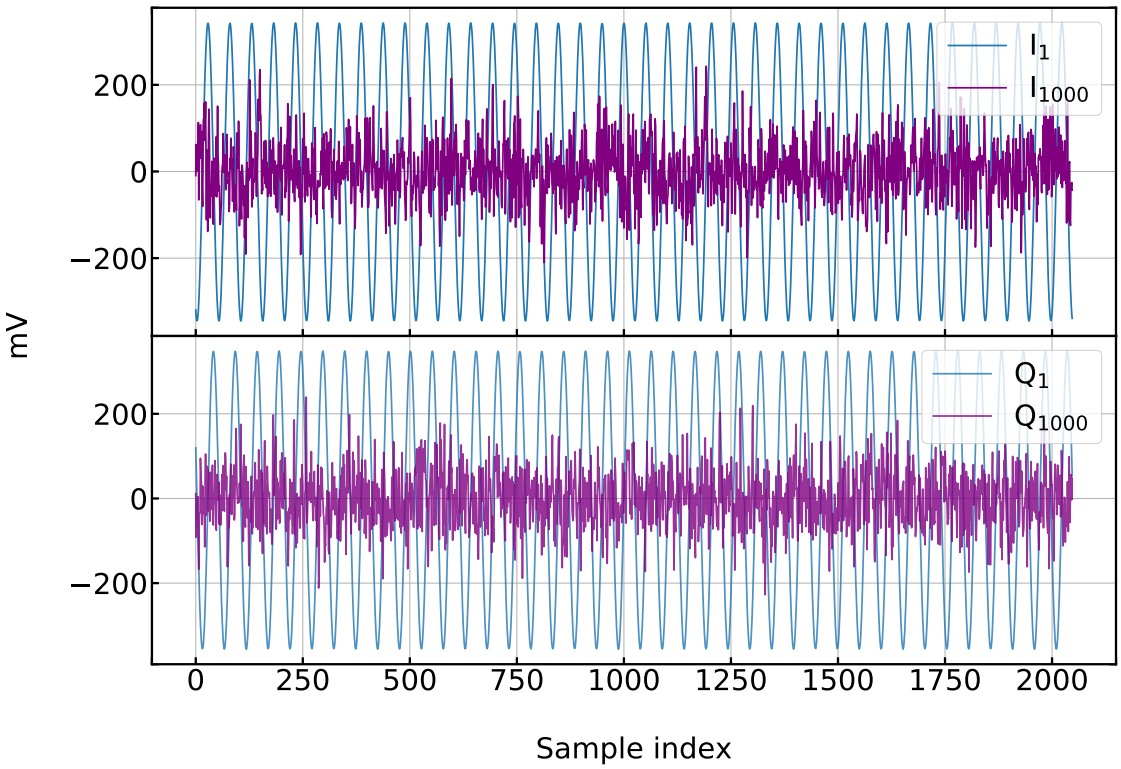

Figure 35. 2,000 I/Q samples of a digitized waveform at $f = 50.0125$ Hz (blue), and for a tone comb consisting of 1,000 evenly spaced tones.

3.3.3 Coarse Channelization: The PFB-FFT

Inside the firmware, the digitized tone comb is processed using two consecutive stages of channelization. The first stage of channelization is a PFB-FFT (AKA filterbank). A detailed description of PFBs is found in Price (2016). Utilizing a PFB before an FFT allows for frequency response of the FFT bins to be shaped according

to the chosen PFB parameters. Specifically, the PFB is intended to minimize the two most deleterious side-effects of FFTs, which are spectral leakage and scalloping loss.

A PFB accomplishes this by multiplying a segment of timestream (for either I or Q) with a window function, dividing the windowed segment into P 'phases' of length N_{FFT} (the poly-phases), and summing each of the P phases. The summed segment of length N_{FFT} is then fed into a FFT, also of length N_{FFT} (the PFB-FFT is the combination of the two). The PFB acts as a set of sub-filters which condition the FFT bin transfer functions so that they are more rectangular than they would be otherwise. Lower sidelobe levels can be achieved by widening the the FFT bins.

The PFB-FFT is implemented using the CASPER `pfb_fir` and `biplex_fft` blocks. For the PFB, we use a $P = 4$ taps, a bin scaling factor of 2, and a Hamming window for the filter coefficients. The bin scaling factor of 2 widens the FFT bins so that they overlap at -6 dB. Their width at FWHM is 450 kHz. On each clock cycle, the biplex FFT receives two consecutive complex time ordered samples I_0, I_1, Q_0, Q_1, and outputs the complex amplitudes $\widetilde{I}_0, \widetilde{I}_1, \widetilde{Q}_0, \widetilde{Q}_1$ of two consecutive (k, $k+1$) frequency bins. One $N = 1024$ FFT is processed every 512 clock cycles, corresponding to an FFT-rate of 500 kHz. A synchronization pulse which is emitted on the last clock cycle before the first valid data of each consecutive FFT is used to synchronize all following stages of the firmware. Since the average individual detector bandwidth is $\sim$50 kHz, several detectors may safely fall within a single FFT bin. Each bin pair output by the FFT is concatenated into a single 72-bit word (4×18-bit) before being stored in block RAM (BRAM) in the FPGA for channel selection.

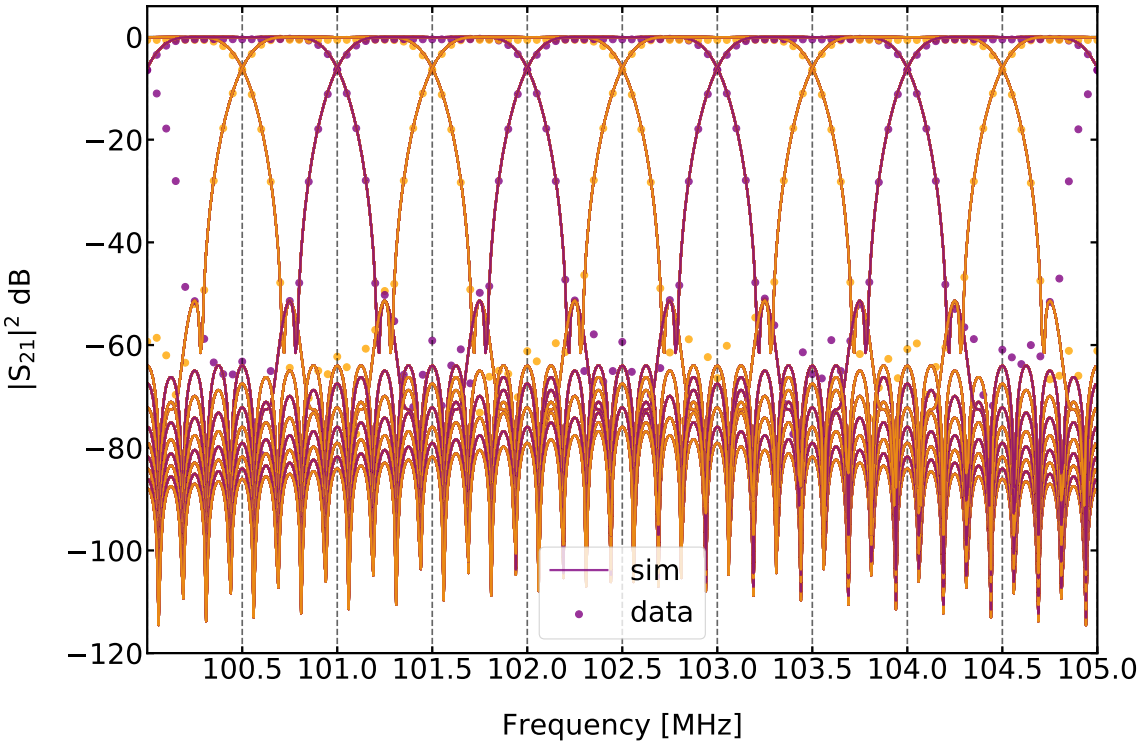

Figure 36. A 5 MHz section of the transfer function of the firmware PFB-FFT, showing measured (dots) and simulated values (lines). The PFB has widened the FFT bins so that they overlap at -6 dB. Their FHWM width is 450 kHz.

3.3.4 FFT Bin Selection

When the LEKID resonant frequencies are irregularly spaced (as is the case for the BLAST-TNG arrays), some PFB-FFT bins contain multiple probe tones, while others remain empty (containing only noise). Only the former set of bins requires further channelization. The channel selection logic requires that up to 1024 channels from the FFT bin stream be selected within 512 clock cycles. To manage this while continuously streaming data, a buffered switch is constructed using Xilinx dual-port BRAM blocks. During software synthesis of the I/Q waveform buffers, a list of up to 1024 bins is pre-calculated based on known resonator frequencies and loaded into

a dual-port 'bin select' BRAM. The list may consist of any combination of the 1024 available bin indices, including a single bin index repeated 1024 times. If only a subset of bins is required, any unused RAM addresses are initialized to zero. Once the bin list is loaded into RAM, the bins are referred to as channels, with the channel order corresponding to the order of the original list. During operation, the bin indices for two consecutive channels are read out in parallel from the dual-port RAM. Each bin index is halved to represent the clock cycle ('clock address') corresponding to its offset in cycles from the zeroth FFT bin, and these clock addresses are to be used as read addresses for another dual-port RAM containing the bin data. In the data RAM, the contents of two consecutive bins are stored in each address slot.

In read mode, the contents of two bins are presented at each output port of the dual-port data RAM, the addresses of which are chosen by the clock addresses of the desired bin indices. Out of the four available bins to choose from on each cycle, only one member of each pair is desired. To determine which member of each bin pair to use, the least significant bit of the desired bin index is used to operate a switch that slices the proper bin from each pair. The new pair of bins is then passed through a MUX selector and sent to the first stage of the DDC. To facilitate continuous readout, the bin selector is duplicated into a read branch and a write branch, which together form a buffered switch.

The bins which are selected for further processing are hereafter referred to as channels.

3.3.5 Fine Channelization: Digital Down Conversion

As was discussed in Section 3.3.1 The PFB-FFT operates on the digitized ADC timestream once every 512 clock cycles, and therefore any probe tone waveforms that have a period longer than the filterbank length of 1024 samples will exhibit unwanted phase rotation over the course of several FFTs. This results in amplitude modulation of $\widetilde{I}$ and $\widetilde{Q}$, where the AM frequency $f_k - f_{\rm pr}$. One approach to circumventing this issue would be to use a longer FFT, with bins so narrow that each $f_{\rm pr}$ falls very near to a bin center. The MUSIC firmware employed this approach, with a 2^{16}-point FFT, resulting in bin width of $\sim$7.5 kHz (Duan et al. (2010)). However, this method relies on the LEKID frequency spacing being somewhat uniform. Instead, we use digital down conversion (DDC) to demodulate the residual AM. The technique, which was previously implemented in the ARCONS firmware (McHugh et al., 2012), has the advantage of using fewer FPGA resources then the long PFB-FFT approach.

To downconvert each channel, the $\widetilde{I}/\widetilde{Q}$ output timestreams from each FFT bin are multiplied with their corresponding $I_{\rm DDC}/Q_{\rm DDC}$ entries from the DDC LUT. With each clock cycle, two consecutive channels are operated on in parallel. A single cycle of the operation involves performing the calculation:

$$(\widetilde{I} + j\widetilde{Q})(Q_{\rm DDC} + jI_{\rm DDC}) \tag{3.14}$$

where $\widetilde{I}/\widetilde{Q}$ are of data type 18.17, $I_{\rm DDC}/Q_{\rm DDC}$ are 16.15, and the resulting output I/Q values are 19.17. During this process, FFT bins which contain multiple channels are downconverted once per channel. For successful down conversion to occur, each channel of the DDC LUT must be synchronized with its corresponding $\widetilde{I}/\widetilde{Q}$ timestream at the output of the FFT. Otherwies, on system start, the first channel

arriving at the downconverter will be out of sync channel-wise with its corresponding DDC tone by some number of clock cycles between 0 and 512. This 'DDC shift' is constant for a given bitstream file, but due to differences in how the System Generator places the logic during each compilation, the DDC shift varies by a small number of clock cycles between different bitstreams.

For each firmware file, the DDC shift can be discovered on system startup by using an algorithm built into the readout software (see Section 3.5). The software steps through each possible DDC shift using the variable delay block while monitoring the snap block data from the DDC for a single channel. At each shift, a software FFT of $\widetilde{I}/\widetilde{Q}$ is compared to that of $I_{\mathrm{DDC}}/Q_{\mathrm{DDC}}$. When the delay has been set properly, the FFTs will match. Once the shift is known, it is programmed into the variable delay block using a software register, and the value is thereafter left unchanged.

As stated above, FFT bins which contain N probe tones are downconverted N times. Each of the N channels which result from this process contain the both the downconverted probe tone, centered at DC, as well as non-zero signal contributions from the other probe tones which were present in the FFT bin prior to downconversion. These adjacent tones must be filtered out. Two methods were explored in order to achieve this. The first method is to add a low-pass FIR filter after the DDC, which is narrow enough to filter out any adjacent tones. The second method, which is what is used in the BLAST-TNG firmware, is to accumulate the I/Q values for each channel (see Section 3.3.6). Accumulation both filters out the adjacent tones (it is a boxcar filter) and reduces the sample rate.

Relative to downsampling by accumulation, the FIR has the disadvantage of requiring many FPGA resources (delays, adders and multipliers). Because the I/Q samples at the output of the DDC are time-multiplexed (equivalently, channel-

multiplexed) ($I_{\mathrm{chan0}}/Q_{\mathrm{chan0}}$, $I_{\mathrm{chan1}}/Q_{\mathrm{chan1}}$, …, $I_{\mathrm{chanN}}/Q_{\mathrm{chanN}}$) the FIR must also act in a time-multiplexed fashion. At the output of the DDC, the samples of any given channel appear once every 512 clock cycles. Therefore, the FIR must include a matching latency of 512 clock cycles between every addition. Following Strader (2016), the nth I value to appear at the output of the FIR is therefore:

$$I_{\mathrm{out}}[n] = \sum_{m=0}^{M-1} h[m] I_{\mathrm{in}}[n - 512m] \tag{3.15}$$

where m is the channel index and h is a filter coefficient. The filter coefficients are calculated in software, using a Hamming or Hanning window, and then programmed into software registers by the user. During system operation they can be reprogrammed at any time.

The need for long latencies make it difficult for System Generator to meet the design's timing constraints during the placement and routing stage of firmware compilation. In addition, to narrow the signal bandwidth from 500 kHz to ∼1 kHz requires the use of many filter taps ($\mathcal{O}(100)$). By simply co-adding (accumulating) consecutive samples from each channel, both filtering and downsampling can be achieved simultaneously, with less resource utilization. Ultimately, this was the option that was taken for the BLAST-TNG firmware.

Figure 37 and 38 show frequency domain examples of the downconversion process for one channel from a tone comb consisting of one and 1,000 tones. In both figures, the top frame shows the input signal to the DDC, $X = \sqrt{\widetilde{I}^2 + \widetilde{Q}^2}$, the middle frame shows the DDC tone X_{DDC}, and the bottom frame shows the downconverted output signal Y. The spurs which are visible at either side of the central tone are due to the effects of quantization errors in the software FFT. In the channel from the multitone comb shown in Figure 38, the group of spurs centered at -50 kHz relative

to the main tone is due to sideband leakage in the IF modulator. The sideband tones are down ∼-20 dB relative to the main tone's peak, and are averaged out during the accumulation stage.

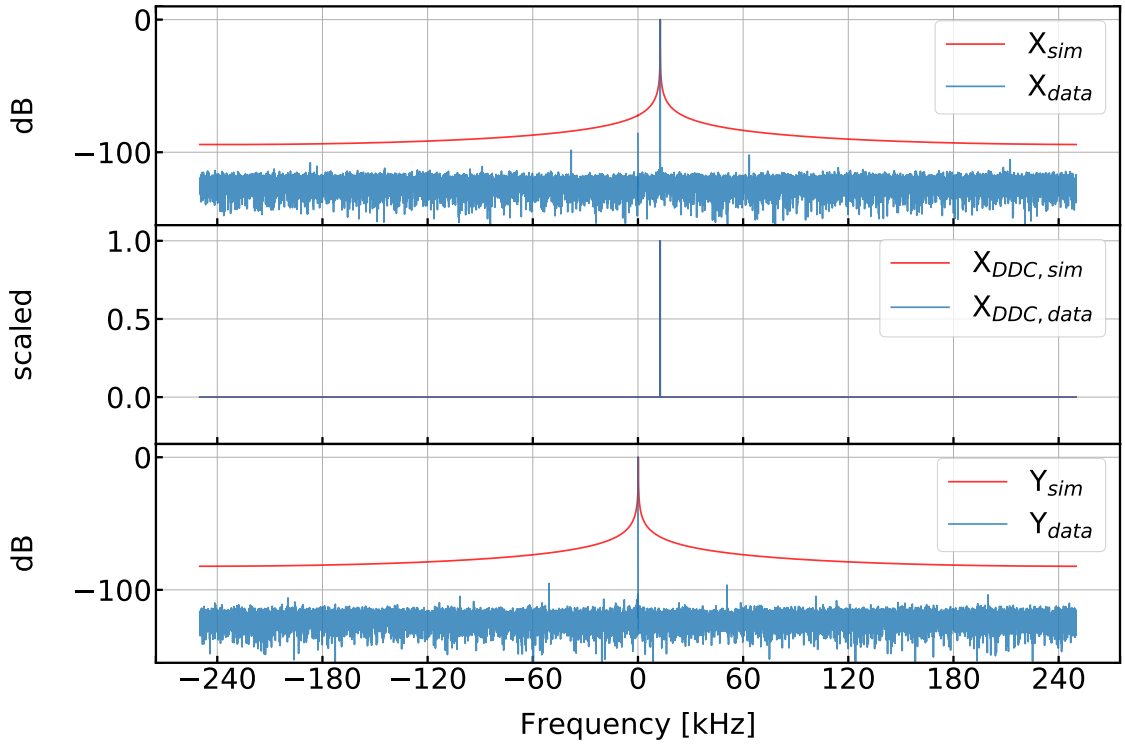

Figure 37. An example of the frequency domain DDC products (measured and simulated) for one channel of a tone comb consisting of a single tone.

3.3.6 Accumulation and Downsampling

The downconverter outputs I/Q samples at a rate of $f_{\mathrm{DDC}} = 500$ kHz. The I/Q timestreams at the output must be downsampled in order to output data at the desired rate. In the version of the firmware which does not contain the post-DDC FIR filter, the downconverted channels still contain signal from probe tones which were within

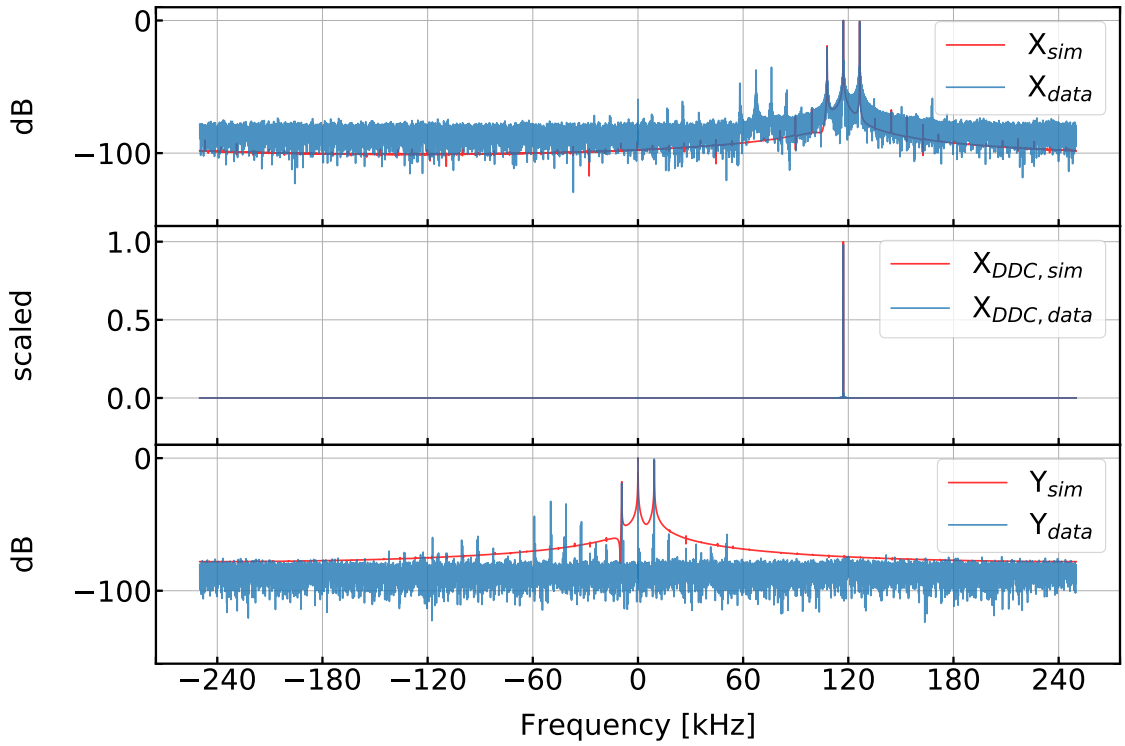

Figure 38. The frequency domain DDC products for one channel of a tone comb consisting of 1,000 probe tones, with measured (blue) and simulated data (red).

the same PFB-FFT bin. While downsampling by decimation would require further filtering, downsampling by accumulation both decreases the data rate and filters out any signal which is non-periodic within the length of the accumulation.

Both downsampling and filtration are achieved by coherent channel-wise accumulation of the I/Q values. The length of the accumulation is configured using a software register. The total number of accumulations used determines the bandwidth of the channels which are output by the system. For BLAST-TNG, the accumulation length is set to 2^{19} clock cycles, which corresponds to $N_{\mathrm{accum}} = 1024$. The resulting readout frequency $f_{\mathrm{ro}} = 488.2815$ Hz, and the channel bandwidth $B_{\mathrm{chan}} = f_{\mathrm{ro}}/2 = 244.14$ Hz. Because the number of accumulations can be adjusted

by factors of 2, $f_{\rm ro}$ may be set to values of $\approx$ 30, 61, 122, 244, 488 or 976 Hz. The firmware architecture does not support $f_{\rm ro} > 976$ Hz.

Accumulation is performed using CASPER vector accumulator blocks of length 512. During the additions, the I/Q values are permitted to grow to 32-bits. The averaging function of the accumulator is effectively a box-car filter, which provides low-pass filtering of the channel timestreams. It is not necessary to complete the average with division by the total number of accumulations. Following accumulation, the channel values are sent to the packetization stage. Figure 39 shows the uncorrected transfer function (frequency response) for the 512 MHz readout band for a tone comb comprised of 1,000 evenly spaced channels (the 'VNA' comb). The orange trace was measured in digital loopback. In digital loopback, the DAC output is filtered and then input directly into the ADC. The blue trace show the full RF loopback (containing all IF electronics). The ripples in the RF loopback trace are a product of the frequency response from each IF component. This transfer function can be corrected for using software, by calculating and applying the inverse transfer function for each readout slice (see Section 4.1.2).

3.3.6.1 UDP Packetization and Time Stamping

UDP packetization is performed using the CASPER 1GbE block. Because the block accepts only one byte at a time, each 4-B I or Q word must be sliced byte-wise before being input to the block.

The data packet size is fixed at 8234-B, regardless of the number of channels which is used. Figure 40 shows the structure of a single UDP frame. Each packet contains a 42-B header, and 8192-B of data, which includes the I/Q values for each

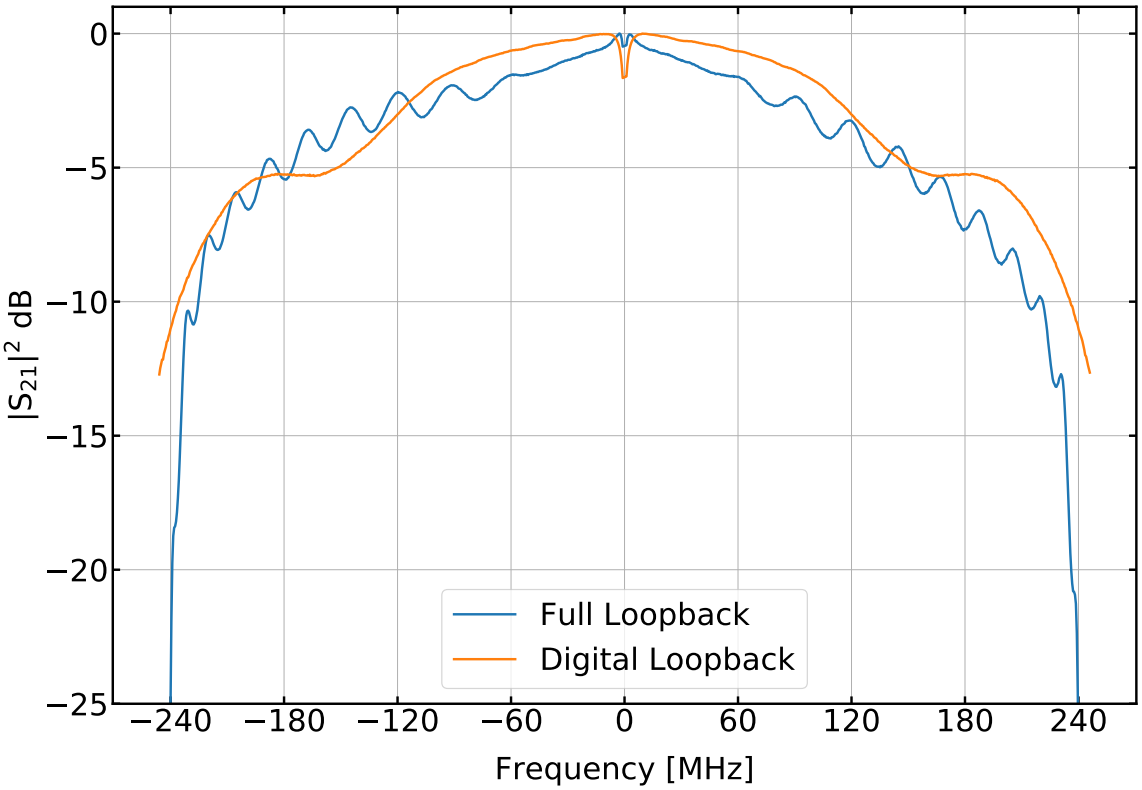

Figure 39. The transfer function of the accumulator ($|S_{21}|^2$), for a tone comb consisting of 1,000 tones, in digital and RF loopback.

channel, as well as time stamp and checksum information. The UDP header contains the source and destination MAC addresses, IP addresses and UDP ports, where the source values correspond to the FPGA Ethernet device on the ROACH2 board, and the destination values correspond to the data acquisition computer (DAQ). Each of these parameters is chosen by the user and programmed into the firmware using software registers during system startup.

After the header, there are 8192-B allocated for data storage. The first 8218-B contain the I/Q values, which are stored as signed, 4-B words in little-endian:

- Words 1–512: I, channels 0–512
- Words 513–1024: I, channels 0–512

- Words 1025–1536: I, channels 513–1024
- Words 1537–2032: I, channels 513–1016

The last 8 Q slots are used to store timestamp values, which limits the maximum channel count to 1016. The final 64-B of each packet contain timetamp information, a data checksum, and a free 4-B slot which can contain extra diagnostic information chosen by the user. The timing information is contained in 4 separate words, which are:

- ctime: The absolute Linux time, taken from the flight computer, and truncated to μs precision.
- PPS count: The number of elapsed PPS pulses since the PPS reset register was toggled.
- clock count: The number FPGA clock cycles which have elapsed between PPS counts.
- packet count: The number of packets which have been generated since the PPS reset register was toggled.

As described in Section 3.1.6, the combination of these values allows for the acquistion time of each data packet to be sychronized with the other telescope subsystems (most crucially, the pointing system) to within the required precision of $\sim$1 ms.

42 Byte header		
DST MAC	SRC MAC	SRC IP
DST IP	SRC PORT	DST PORT
8192 Data Bytes I and Q are 4 B words, signed little endian		
1 - 512: I, CHAN 0 - 512 513 - 1024: Q, CHAN 0 - 512 1025 - 1536: I, CHAN 513 - 1024 1537 - 2032: Q, CHAN 513 - 1016		
Words 2032 - 2048 4 B unsigned big endian		
CTIME	PPS COUNT	CLOCK COUNT
PACKET COUNT	FREE BYTE	FREE BYTE

Figure 40. The structure of a ROACH2 UDP packet

3.4 Noise Verification

In this section, we describe the verification of the WN and FN system requirements described in Sections 3.1.2 and 3.1.3. To characterize these noise properties, we use the one-sided PSD, $S_{xx}(f)$:

$$S_{xx}[f] = \frac{1}{Nf_s}\left|\sum_{n=1}^{N} x_n e^{-j2\pi fn/f_s}\right|^2 \tag{3.16}$$

To obtain a better estimate of the PSD than Equation 3.16 can provide, we use the Welch Periodogram (WP) (Welch, 1967). The WP is an asymptotically unbiased estimator of the PSD. In Welch's method, a timestream is divided into K windowed segments of equal length. The periodogram is calculated for each segment, and then

averaged. The segments may also be overlapped in time. The variance of the WP is reduced by a factor of $1/K$ relative to a regular periodogram, at the expense of lower frequency resolution. The WP can be written as:

$$S_{xx}[f] = \frac{1}{K}\sum_{n=0}^{K-1}\frac{1}{f_s M}\left|\sum_{n=0}^{N-1} x_{n,w} e^{-j2\pi f n/f_s}\right|^2 \tag{3.17}$$

where $x_{n,w}$ is a windowed version of the time series, and M is the length of each segment. To fit the noise PSDs, we use a modified version of Leeson's equation, which empirically models the phase noise spectrum produced by a resonator (Lesson, 1966). The model is a three-parameter fit to S_ϕ:

$$S_{\phi,fit}(f) = W\left[1 + f_c/f + \left(\frac{B_{\rm res}}{f}\right)^2 (1 + f_c/f)\right] \qquad \left[\frac{\rm rad^2}{\rm Hz}\right] \tag{3.18}$$

where W is the WN level, the resonator bandwidth $B_{\rm res} = (f_0/2Q_r)$, and f_c the FN corner-frequency. The PSDs are measured in digital and RF loopback. In digital loopback, the MUSIC board DAC I/Q outputs are passed through anti-aliasing filters and then fed directly into the ADCs. RF loopback includes the IF electronics. To simulate the gain of the BLAST-TNG cryostat, a 20 dB attenuator is placed between the RF out and RF in ports of the readout slice.

Figure 41 shows an example fit (red) using Equation 3.18 to the median PSD of a tone comb containing 1,000 evenly spaced tones (blue), in digital loopback mode. In this particular example, the FN corner frequency $f_c \approx 0.5$ Hz, and the WN level is $\sim$-98 dBc/Hz.

In digital loopback mode, the noise was measured for tone combs containing 1, 100, 500, 618 and 1,000 tones. In the tone combs containing 100, 500 and 1,000 tones, the frequencies were evenly spaced by $\approx$ 490 kHz. The tone comb with length

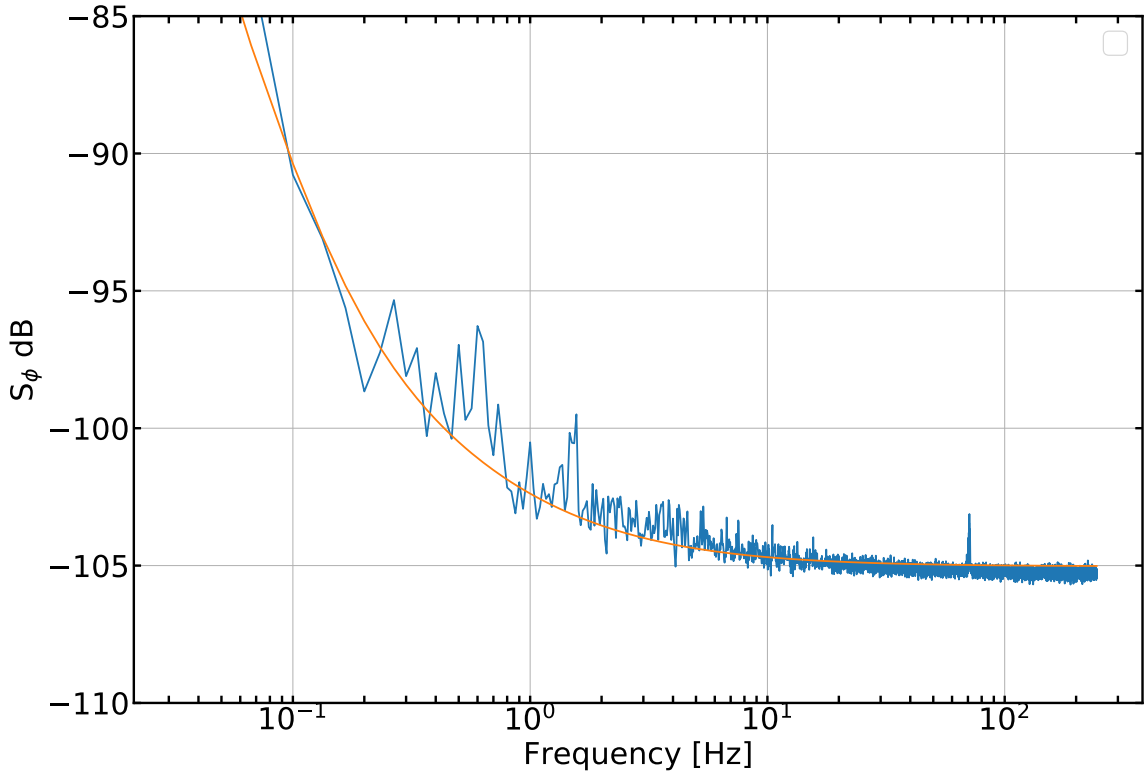

Figure 41. A fit (red) of a modified Leeson's equation to the median PSD (blue) for a tone comb containing 1,000 evenly spaced tones.

618 uses resonator frequencies which were identified for the 350 μm array during detector testing at NASA's Columbia Scientific Balloon Facility in July, 2018 ($618_{350\mu m}$). These tones are not evenly spaced (see the discussion about detector yields in Chapter 4). The median PSDs calculated from each tone comb are shown in Figure 42. Table 8 lists the WN level, in dBc/Hz and dB/Hz (converted using the method described in Section 3.1.2, as well as the estimated CF, using Equation 3.5).

For the single tone comb, the WN level is measured to be ≈-143 dBc/Hz, which is within ∼1 dB of the theoretical minimum predicted by Equation 3.5. The WN levels for the evenly spaced tone combs are also close to their expected values. The $618_{350\mu m}$ tone comb has a median WN level of ≈-105 dBc/Hz, which is close to the level measured for 1,000 evenly spaced tones. This small amount of excess noise is

attributed to the uneven spacing of the probe tones, which creates intermodulation products in the analog modulator/demodulators. For each tone comb, the FN corner frequency is ≈0.5 Hz.

N tones	S_ϕ (dBc/Hz)	$S_{\delta f/f}$ (dB/Hz)	Crest Factor est. (dB)
1	-143	-244	3
100	-113	-215	11
500	-108	-209	10
$618_{350\mu m}$	-105	-206	12.5
1000	-105	-206	10

Table 8. The digital loopback WN levels as a function of 1, 100, 500, 618 and 1,000 tones. $618_{350\mu m}$ is the tone comb for the 350 μm band tones used during the 2018 Palestine instrument integration for BLAST-TNG. $S_{\delta f/f}$ is calculated according to Equation 3.2, with $Q_r = 2.85 \times 10^4$. The PSDs shown are the median of the PSDs calculated for each N value, from 60 s timestreams. CF is the estimated crest factor, in dB.

The same median-PSD analysis was performed in RF loopback, using tone combs containing a single tone, the $618_{350\mu m}$ tones and an evenly spaced comb of 1,000 tones (25 of which were removed from the data due to unusually high noise). The results of the analysis are shown in Figure 43. Table 9 lists the WN levels as in Table 8, except the third column shows the estimated IF noise, in dB. The IF noise was estimated using Equation 3.5 with the CF values in Table 8, and is found to be ∼5 dB in each measurement.

The WN level of the single tone is measured as -136 dBc/Hz, and those of the 1,000 and $618_{350\mu m}$ combs are ∼ -98 dBc/Hz. These values are safely below the WN system requirement derived in Section 3.1.2. However, the margin of error is only a few dB. As in digital loopback mode, the FN corner frequencies are ≈ 0.5 Hz. While this value also satisfies the system requirement, there is little margin for error. It is

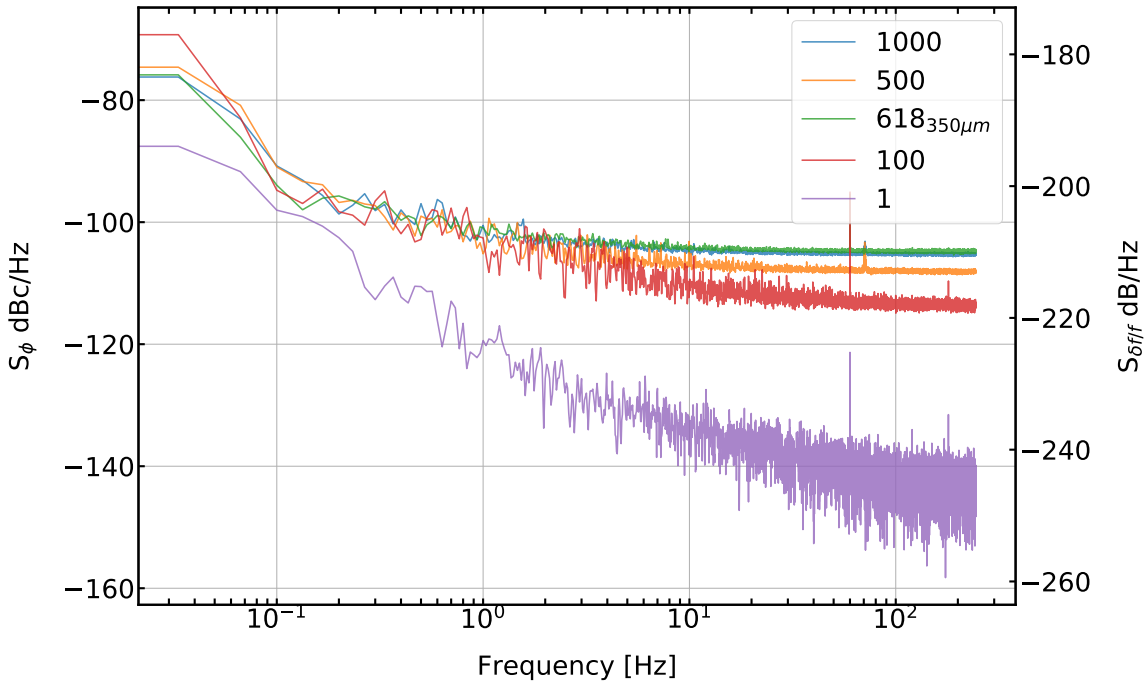

Figure 42. The digital loopback WN levels as a function of 1, 100, 500, 618 and 1,000 tones. $618_{350\mu m}$ is the tone comb for the 350 μm band tones used during the 2018 Palestine instrument integration for BLAST-TNG. $S_{\delta f/f}$ is calculated according to Equation 3.2, with $Q_r = 2.85 \times 10^4$. The PSDs shown are the median of the PSDs calculated for each N value, from 60 s timestreams.

therefore important to keep the noise contribution of the IF electronics as low as possible.

The WN and FN noise estimates presented in this section are in close agreement with those presented in Gordon et al. (2016), which have been measured using similar firmware and IF electronics.

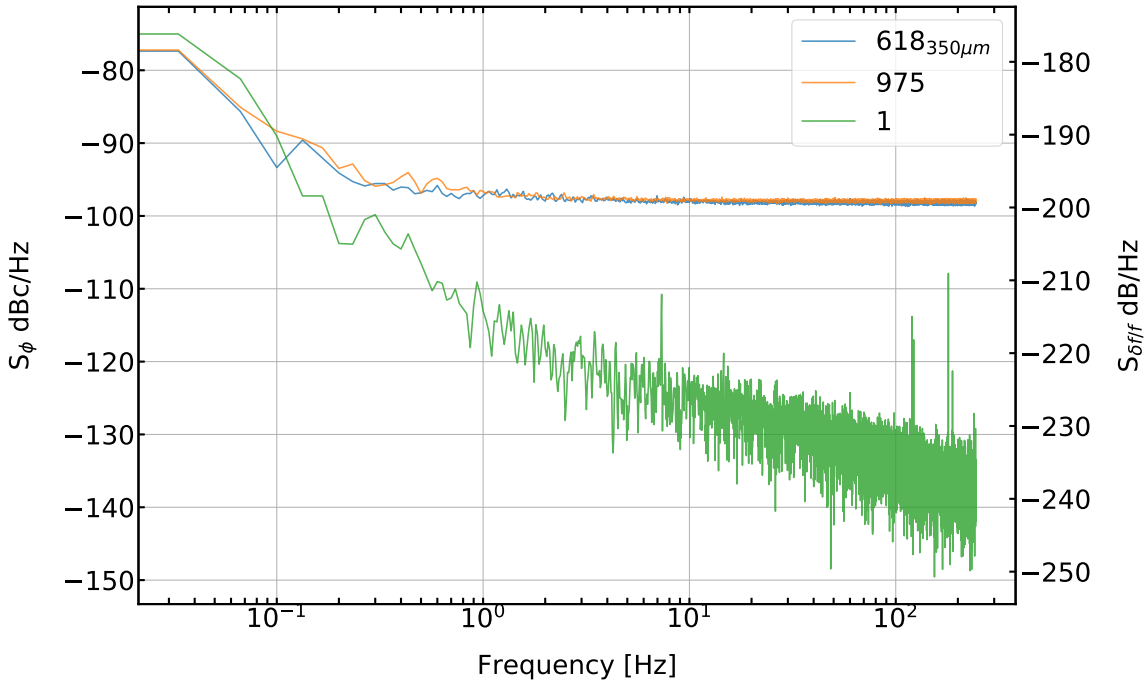

Figure 43. The RF loopback WN levels as a function of 1, 618 and 975 tones. $618_{350\mu m}$ is the tone comb for the 350 µm band tones used during the 2018 Palestine instrument integration for BLAST-TNG. $S_{\delta f/f}$ is calculated according to Equation 3.2, withQ_r = 2.85 × 10^4. The PSDs shown are the median of the PSDs calculated for each N value, from 60 s timestreams.

N tones	S_ϕ (dBc/Hz)	$S_{\delta f/f}$ (dB/Hz)	IF Noise (dB)
1	-136	-237	4
$618_{350\mu m}$	-98	-199	5
1000	-98	-199	5.5

Table 9. The RF loopback WN levels as a function of 1, 618 and 975 tones. $618_{350\mu m}$ is the tone comb for the 350 µm band tones used during the 2018 Palestine instrument integration for BLAST-TNG. $S_{\delta f/f}$ is calculated according to Equation 3.2, with Q_r = 2.85 × 10^4. The PSDs shown are the median of the PSDs calculated for each N value, from 60 s timestreams. The last column is the estimated IF noise, in dB.

3.5 Readout Software

A custom software interface is required to operate the ROACH2 firmware and IF electronics, both in a lab setting and during instrument operation. The extent to which the software is automated depends on the context in which the system is being used. For example, if the system is to be used on a balloon platform (e.g., BLAST-TNG, OLIMPO), most if not all of the software functionality should be automated. At a ground/mountain-based observatory (e.g., TolTEC, MUSCAT, SuperSpec), user input on the observer's part should be incorporated when the outcome of a decision can affect the quality of the data (e.g., during the channel selection stage of operation). Basic housekeeping functions, which include uploading the FPGA firmware, initializing software registers such as the DDC delay and writing and programming probe tone combs tones frequency combs, should be automated no matter the context. Other functions, such as the acquisition of diagnostic snap-block data and the setting of RF power levels via the programmable attenuators, should be easily accessible to the operator.

For lab-based system operation, an open-source `Python`-based software was developed, known as kidPy[27]. KidPy is a terminal user interface (TUI) which enables the manual operation of every step which is required to run the readout and acquire data. For the BLAST-TNG flight code (MCP), kidPY was ported to `C` and fully automated. The BLAST-TNG readout software and in-flight operational strategy is detailed in Appendix A.

[27] https://github.com/sbg2133/kidPy

Chapter 4

INSTRUMENTAL CHARACTERIZATION OF BLAST-TNG

The following sections describe the characterization of the instrumental performance of BLAST-TNG in the lead-up to the 2018/2019 launch attempt from NASA's Long Duration Balloon site (LDB) in Antarctica. Because the camera contains over 2,500 polarization sensitive pixels distributed over three wavebands, it would not be feasible to describe the individual performance of each pixel within this document. Instead, we present a detailed data reduction for a single pixel, which is then extended to statistical results which summarize the overall performance of each BLAST-TNG band. The instrumental performance is characterized by numerous parameters, many of which are interdependent and can not be directly measured. Instead, they must be inferred from the results of empirical measurements, which in some cases can be checked against a model, such as the one presented in Chapter 2.

Out of all of the instrumental parameters that could be examined, a small number of them directly determine how well the camera will perform in terms of its mapping speed and on-sky sensitivity to sub-mm intensity and polarization. These parameters are the detector yields and NEP or NET (sensitivity), optical efficiency, passband shapes and polarization efficiencies. The in-flight performance is also directly related to the operation of the detector readout. For example, to utilize the full dynamic range of each detector during flight will require that the readout probe tones be periodically retuned in frequency and amplitude. Notes on readout operation are interspersed with the following description of how each of the aforementioned parameters has been estimated.

The chapter is organized as follows:

- Section 4.1 introduces detector frequency sweeps, frequency-domain channel identification, and detector yields.
- Section 4.2 introduces the raw I/Q timestreams which are generated by the readout, and how they are calibrated into frequency units which are proportional to absorbed optical power.
- Section 4.3 discusses the responsivities to absorbed optical power and changes in base temperature.
- Section 4.4 presents sensitivity estimates for each band (NEP) and describes how these estimates are made using a combination of sweeps, timestreams and optical tests.
- In Section 4.5 we compare the measured data for a single channel from the 350 μm and 250W arrays to results from the parametric LEKID model presented in Chapter 2.
- In Section 4.2.3 we present histograms showing the distribution of important detector parameters for each of the five BLAST-TNG detector arrays.
- Section 4.6 presents the data reduction and results of passband mapping (Section 4.6.1) and polarization efficiency (Section 4.6.2).

The data used in the following analysis was mostly acquired in May, 2018 (Philadelphia, PA), July, 2018 (NASA's Columbia Scientific Ballooning Facility (CSBF) Palestine, Texas), and winter 2018/2019 (NASA Long Duration Balloon facility (LDB), Antarctica).

4.1 Detector Sweeps

The BLAST-TNG detectors are distributed between five LEKID arrays, with one for the 350 and 500 μm bands, and three for the 250 μm. The three 250 μm arrays (hereafter referred to as the 250U, 250V and 250W) reside on the same chip, but are read out independently. The 250V RF-chain malfunctioned during the pre-flight instrument integration at LDB, and we therefore do not present results for this channel from the Antarctic ('ice') campaign.

Every test of the detectors begins with a frequency sweep of the full 512 MHz RF-bandwidth occupied by each detector array. If the resonator frequencies are not known ahead of time, or have drifted by more than ∼1 resonator linewidth, the wide sweep is used to identify the channel frequencies (see Section 4.1.3). A wide sweep (sometimes referred to as a 'VNA' sweep) is performed by sweeping the 1,000 tone evenly-spaced search-comb over the 500 kHz gap between tones. The multitone comb is swept by stepping the local oscillator (LO) from one end of the 500 kHz gap to the other in frequency steps of ≈ 1 kHz. At each step, I/Q data is saved to disk and co-added, with the number of data points taken at each frequency typically being ∼20. The wide-sweeps take ∼30 seconds to complete. Plotted as a function of RF-frequency, the data shows $S_{21}V_{\rm in}$ of the detector readout chain (hereafter simply referred to as S_{21}), where $V_{\rm in}$ is the amplitude of the probe-tone comb at the output of the ROACH2 digital-to-analog converters (DACs).

4.1.1 The Wide-sweep Frequency Response

Figures 44 and Figures 45 show an overlay of two example sweeps of the 350 μm array. The blue sweep was taken in May, 2018, and the orange trace was taken in July, 2018. The blue sweep shows the raw (uncorrected) frequency response of the readout system and detector chain, whereas the orange trace was produced by applying several corrections to the readout probe tone amplitudes. The transfer function correction is described in Section 4.1.2. Both sweeps were taken with the cryostat window exposed to the ∼300 K thermal radiation of the highbay. A (2.85%) 4% neutral-density filter (NDF) inside the receiver's cold-optics limits the optical power which reaches the detectors (after accounting for an instrumental optical efficiency $\epsilon_{\rm opt}$ of ∼30%, and a detector efficiency $\eta_{\rm det}$ of ∼80%) to ∼10 pW.

The frequency response (or transfer function) of the system contains information about the frequency response of every element in the detector chain. The primary contributors are the ROACH2 electronics, input and output side intermediate-frequency (IF) electronics, room-temperature as well as cryogenic cabling, the cryogenic low-noise amplifier (LNA) and the detector arrays themselves. In the frequency domain, the total system response is the product of the transfer function for each of the individual elements.

The envelope of the uncorrected S_{21} sweep contains several notable features:

- A gradual roll-off in power toward the band-edges.
- A steep roll-off in power at the band-edges.
- A deep null in the center of the band.
- A standing wave (ringing) pattern with a frequency of ∼30 MHz.

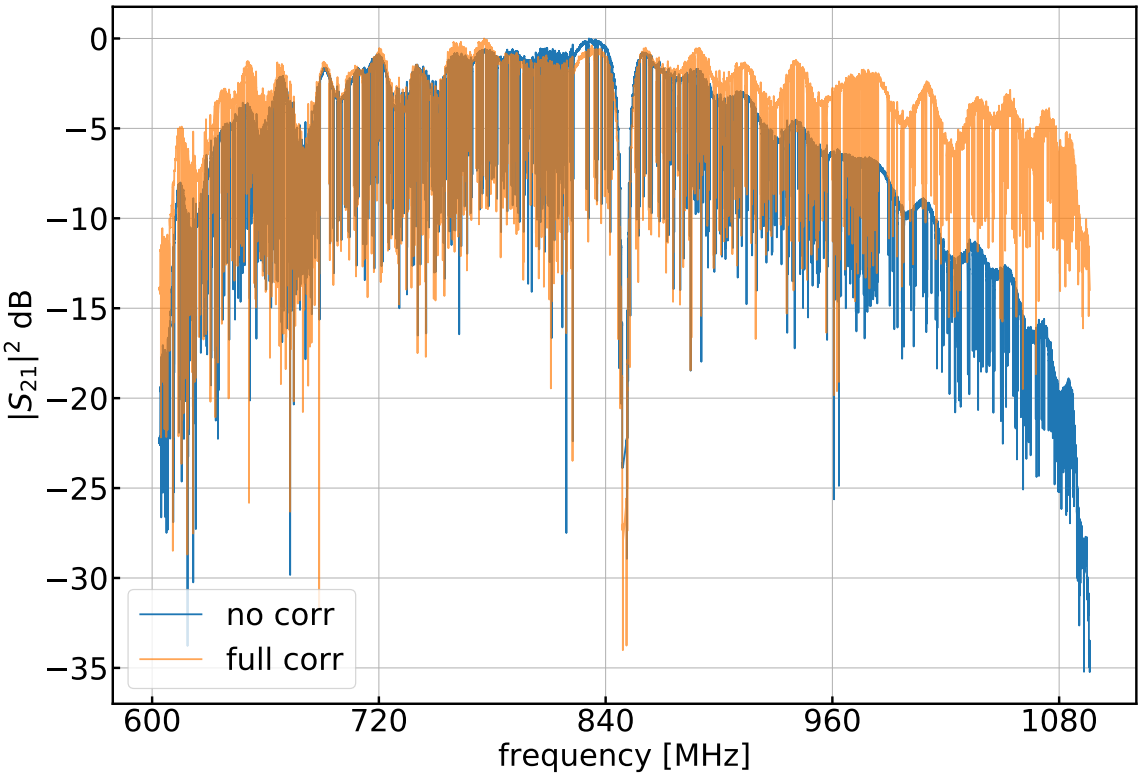

Figure 44. Uncorrected (blue) and corrected (orange) ROACH2 wide-sweeps of the BLAST-TNG 350 μm array.

The gradual roll-off in power toward the band-edges can be attributed to the MUSIC board's DAC frequency response, the LNA frequency response at low frequency (∼500 MHz) as well as to the frequency response of modulator and demodulator (the demodulator contains 300 MHz low-pass filters). The steeper roll-off that begins at ∼15 MHz from the lower and upper band-edges is from the baseband anti-aliasing filters. The ringing pattern in the S_{21} envelope is due to a combination of the anti-aliasing filters and warm and cold cabling. The null in the center of the band is where the LO frequency is located. Typically, no probe tones are generated inside a 20 MHz band centered at the LO frequency.

The remaining features in the S_{21} trace shown in Figure 44 are the individual resonator transfer functions. The 350 μm array contains ∼750 resonators (within

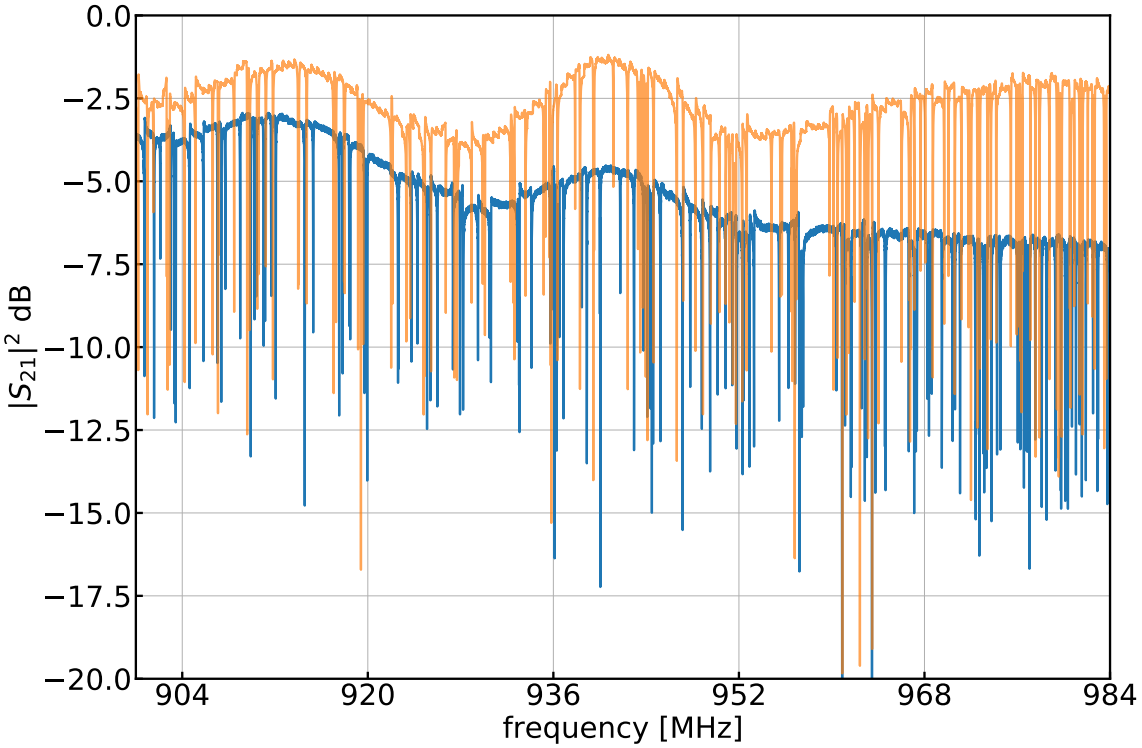

Figure 45. A zoomed-in region of an unprocessed (blue) and processed (orange) wide-sweep of the 350 μm array.

the readout band– see Section 4.1.4). A small number of resonators with their horn antennas blocked off, known as *darks*, have noticeably deeper resonances than the other detectors. The distribution of resonator dip-depths varies widely across the band. This spread is due to variations in several parameters, including intrinsic resonator quality factors, readout powers and optical loading. Much of the large scale variations in readout power can be corrected for by applying a global inverse transfer function to the probe comb. In addition to this, the tone power for each channel may be individually adjusted. The process of determining and applying the inverse transfer function correction is described in the following section.

4.1.2 Transfer function Correction

As discussed in Section 4.1.1, the uncorrected S_{21} frequency response revealed by the ROACH2 wide sweeps contain large scale features that may contribute to uneven detector performance across the band. Some of these features can be compensated for by applying an inverse transfer function to the amplitudes of the DAC probe comb. The goal of applying this correction is to correct for, as much as possible, any transfer function contributions from the input side of the detector chain (relative to the cryostat). All output side transfer functions (after and including the detectors) are left untouched. The components whose transfer functions will be corrected include the ROACH2 output-side electronics (both digital and analog) and any coaxial cables between the ROACH2 electronics and the detector arrays.

The output side system transfer function is measured by connecting the RF output of the ROACH2 IF electronics to a spectrum analyzer. The spectrum is then saved to disk, and interpolated in frequency as needed. Examples of the intermediate products used in the transfer function calculation are shown in Figure 46. The blue trace is the raw output of the RF modulator. A $\sim$10 dB roll-off is present between the center of the band and either band-edge. In addition to correcting for the envelope of the ROACH2 comb, a power ramp (shown as a green line) is added to the correction to compensate for frequency dependent loss in the cryostat cables. The inverse transfer function which is then applied to the DAC comb amplitudes is shown in orange.

The frequency-dependent cable loss which occurs inside the cryostat is visible in Figures 47 to 51, which show S_{21} sweeps taken with a VNA (purple) compared to corrected ROACH2 S_{21} sweeps (orange) for each BLAST-TNG array. The cable

loss is ∼10 dB across the readout-band. With the application of the inverse transfer function, much of the roll-off in the envelope of the frequency response is compensated for. For each array, a slightly different slope for the power ramp is used. An illustration of the effect of using different slope corrections for the 350 µm array is shown in Figure 52, for slopes of 5, 8 and 10 dB (across the entire readout band).

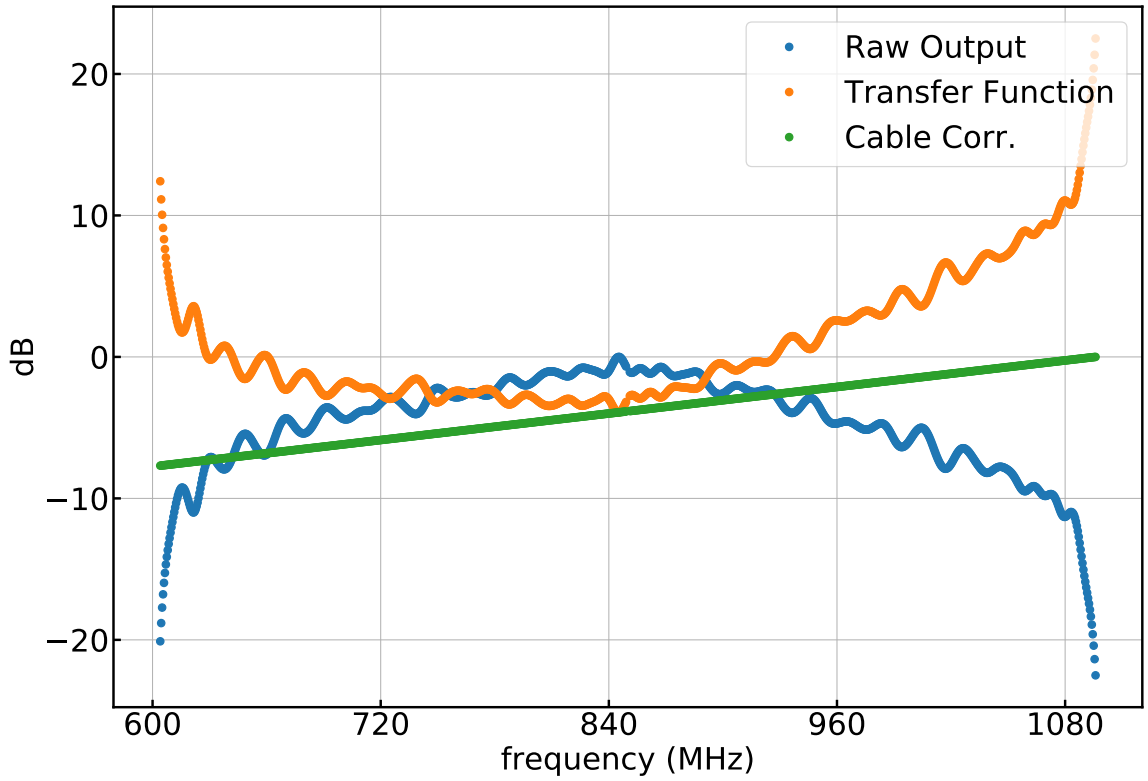

Figure 46. Intermediate products used in the transfer function calculation for the 350 µm array.

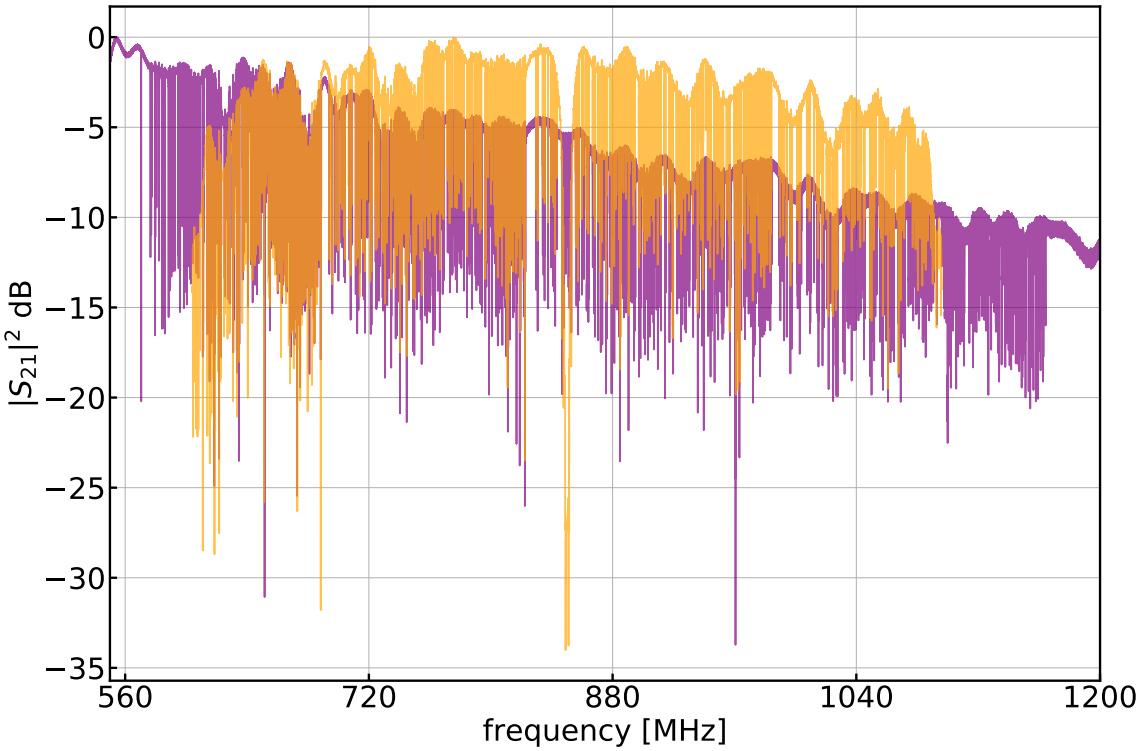

Figure 47. A 350 μm ROACH2 sweep (orange) from the Palestine integration, with a VNA sweep from May, 2018 (purple).

4.1.3 Channel Identification

The channel frequencies for each array can be determined from a wide ROACH2 sweep (or VNA trace of S_{21}). For BLAST-TNG, the steps of the KID-finding algorithm are:

1. Despike the S_{21} trace in linear-space.
2. Apply a high-pass filter (HPF) to flatten the trace.
3. Put the trace into log-space, and apply a second HPF.
4. Apply a low-pass filter (LPF) to remove any remaining spurs.
5. Identify all points in trace below the dip-depth threshold.
6. Step through points and enforce frequency spacing condition.

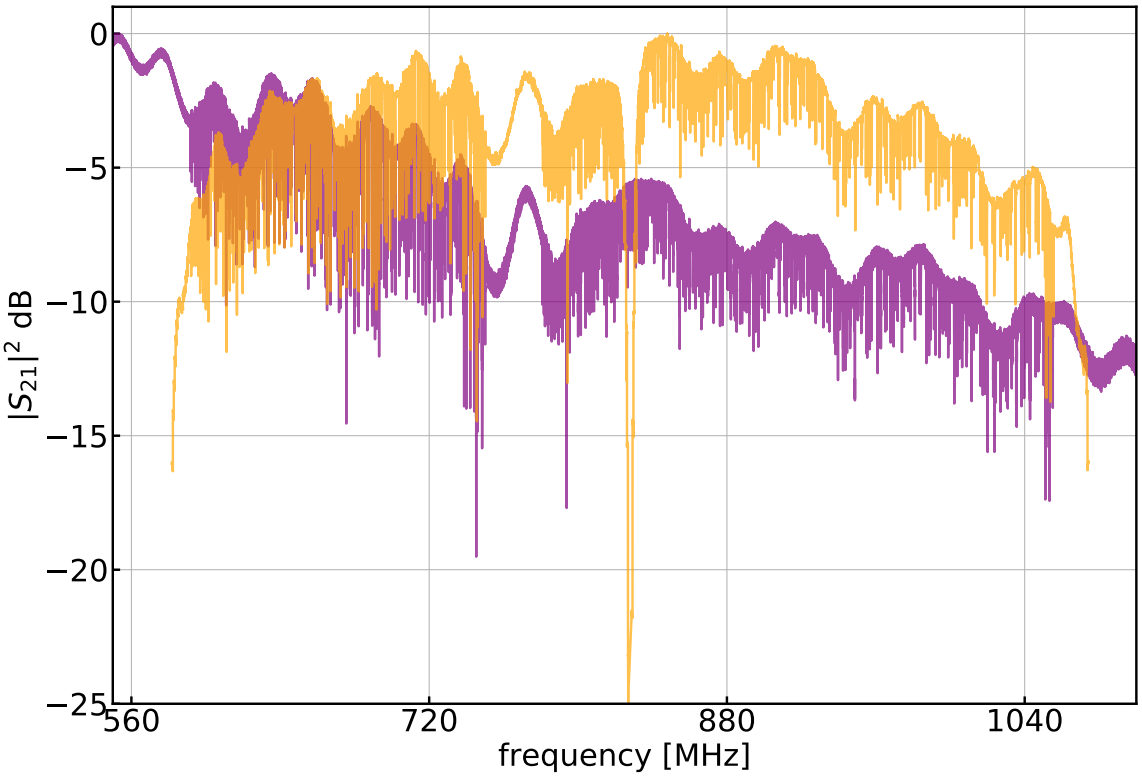

Figure 48. A 250W ROACH2 sweep (orange) from the Palestine integration, with a VNA sweep from May, 2018 (purple).

7. Eliminate frequencies in masked regions.
8. Calculate channel amplitudes from interpolated output TRF.
9. Save channel frequencies and amplitudes to disk.

This sequence is intended to be performed either at the start of a science observation, or when environmental conditions (fridge temperature drift, optical loading) have shifted the resonators far off of resonance. If the frequency shifts are less than ∼2 resonator linewidths, a target sweep may be used to recalibrate the readout frequencies in fewer steps (see Section 4.1.5).

The preflight parameters used in the KID-finding process for each BLAST-TNG

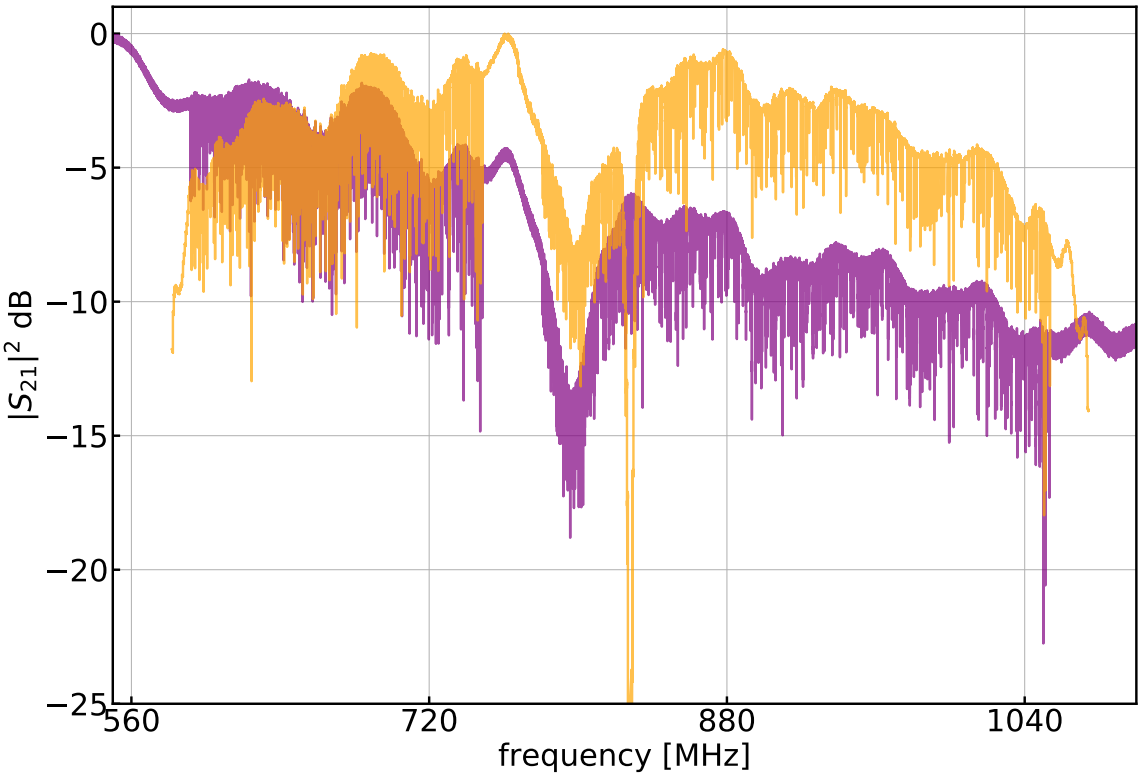

Figure 49. A 250U ROACH2 sweep (orange) from the Palestine integration, with a VNA sweep from May, 2018 (purple).

detector array are listed in Table 10. The S_{21} sweep is high-pass filtered in (2) and (3) to flatten out the trace envelope as much as possible before identifying points below the dip-depth threshold. The high-pass filtering in (2) is achieved by subtracting a LPF from the data in linear-space. An example of this is shown in Figure 53, for the 350 μm array. The close spacing of the channels makes it challenging to flatten out the inter-channel frequency space below a level of ∼0.1–0.5 dB. Consequently, a small number of 'false' channels (∼1–2% of the total number) are identified by the KID-finding algorithm. The false channels can be manually flagged for removal, or added to the masked regions in (7). The dip-depth threshhold for the 500 μm array is set to -0.5 dB due to the intrinsically low Q_r of the resonators. For the other arrays, the threshold is set to -1 dB.

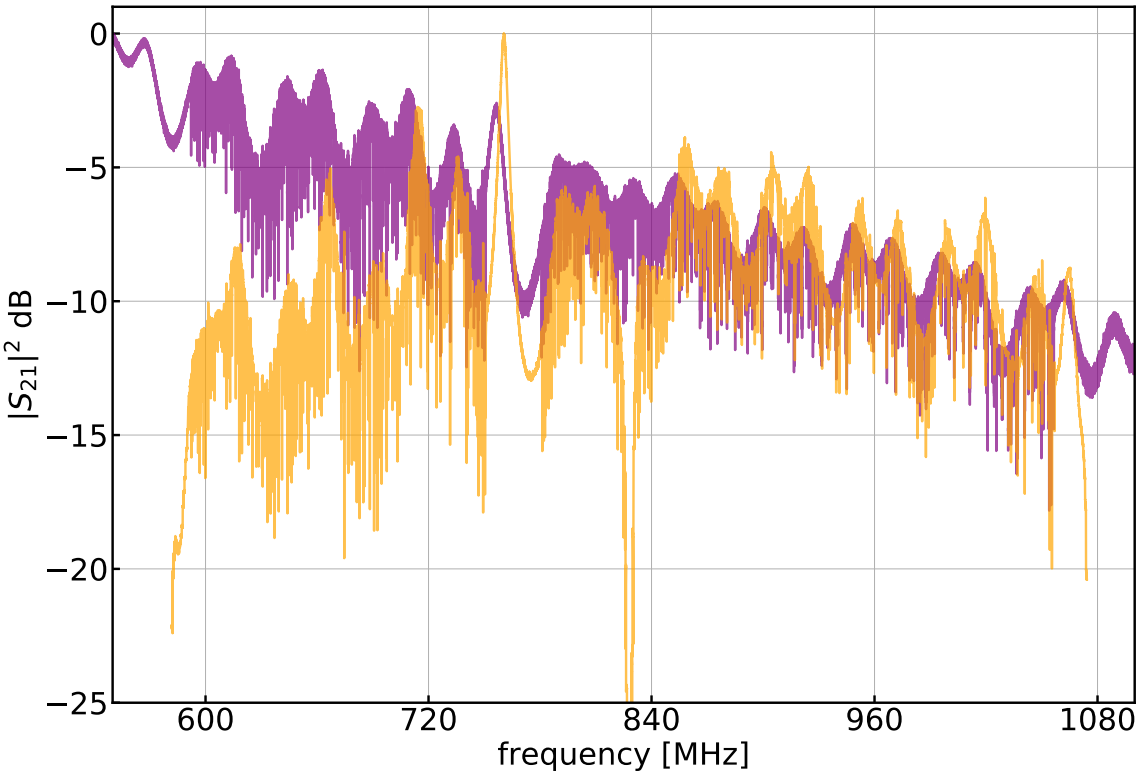

Figure 50. A 250V ROACH2 sweep (orange) from the Palestine integration, with a VNA sweep from May, 2018 (purple).

During channel identification, a frequency spacing threshold of 80–100 kHz is enforced. If two channels are spaced by less than this threshold, the channel with the shallow dip-depth is omitted from the final frequency list. The masked regions listed in Table 10 are ignored during the KID-finding process. These regions include a 10 MHz window above and below the center of the band (f_{LO}), and narrow bands which sometimes contain spurious noise thought to be associated with RF interference (RFI).

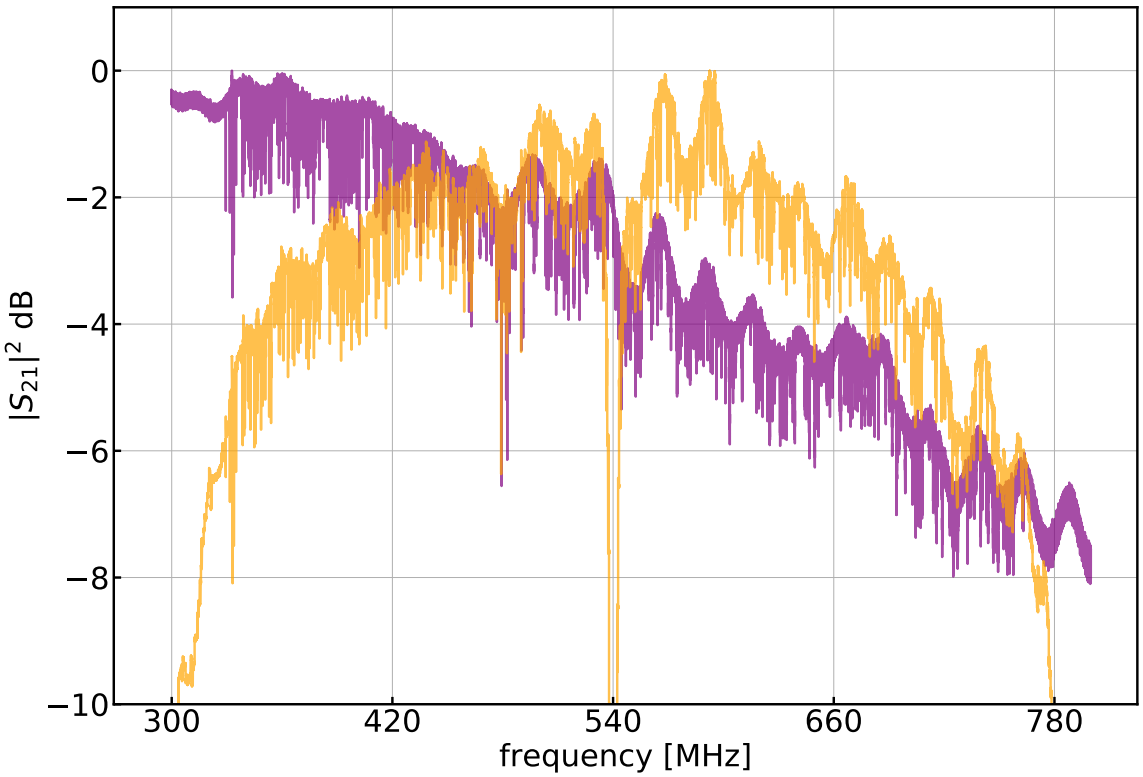

Figure 51. A 500 µm ROACH2 sweep (orange) from the Palestine integration, with a VNA sweep from May, 2018 (purple).

Table 10. Parameters used in KID-finding process for each BLAST-TNG detector array.

	250U µm	250V µm	250W µm	350 µm	500 µm
LPF f_c (MHz)	10	5	5	10	10
Dip-depth thresh. (dB)	-1	-1	-1	-1	-0.5
Spacing thresh. (kHz)	100	100	100	100	100
f_{LO} (MHz)[1]	827	829	828	850	540
$f_{low-cut}$ (MHz)	590	590	590	604.5	326
$f_{high-cut}$ (MHz)	1060	1060	1060	1089	762
$f_{mask,low}$ (MHz)[2]		820	751, 818	845	450.8, 535
$f_{mask,high}$ (MHz)[3]		835	770, 838	855	451.8, 545

[1] The LO center-frequency.
[2] The start frequencies of masked regions.
[3] The stop frequencies of masked regions.

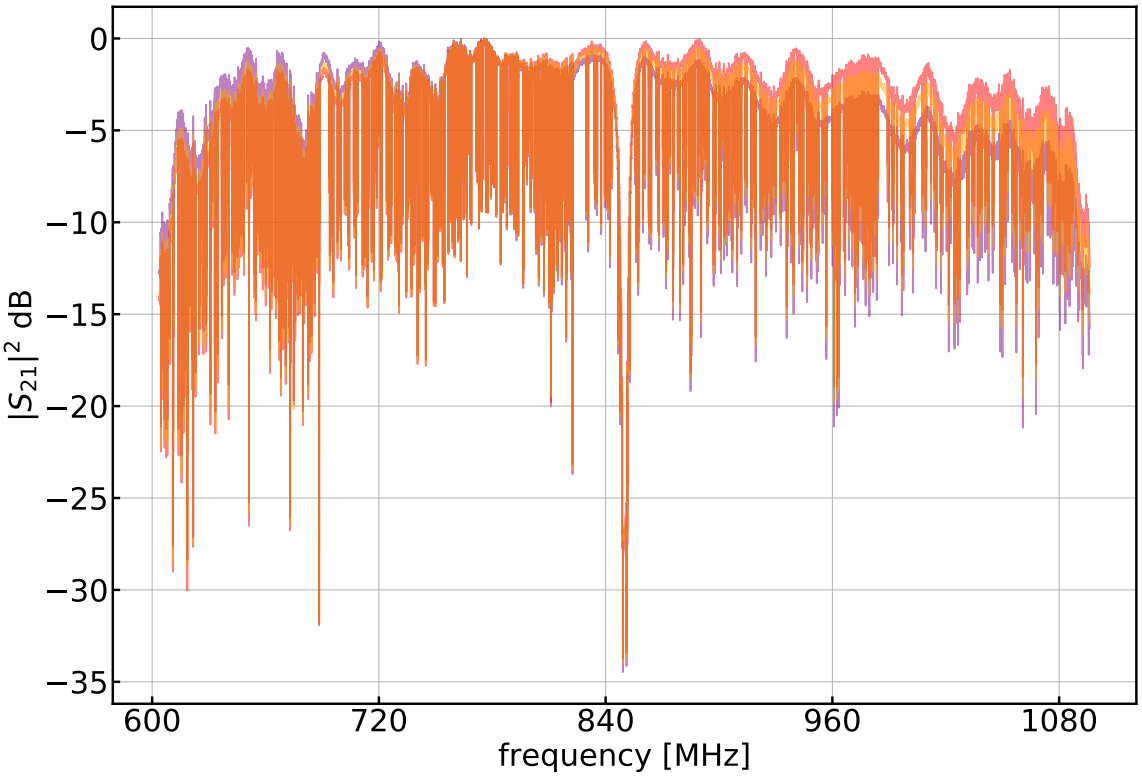

Figure 52. 350 µm output transfer function corrections with 5, 8 and 10 dB slopes accounting for frequency dependent cable loss in the croystat. A VNA trace is shown for comparison.

4.1.4 Detector Yields

Detector yields for each BLAST-TNG array determined using the ROACH2 readout during the Palestine and Antarctica integrations are listed in Table 11. In the table, N_{design} refers to the number of channels on each array which were reported by NIST. With the exception of the 350 µm array, this value is not necessarily the actual number of channels within the 512 MHz ROACH2 readout bandwidth. As a baseline channel count, we adopt the number of channels which were found by applying the KID-finding algorithm to dark VNA S_{21} sweeps taken during May, 2018

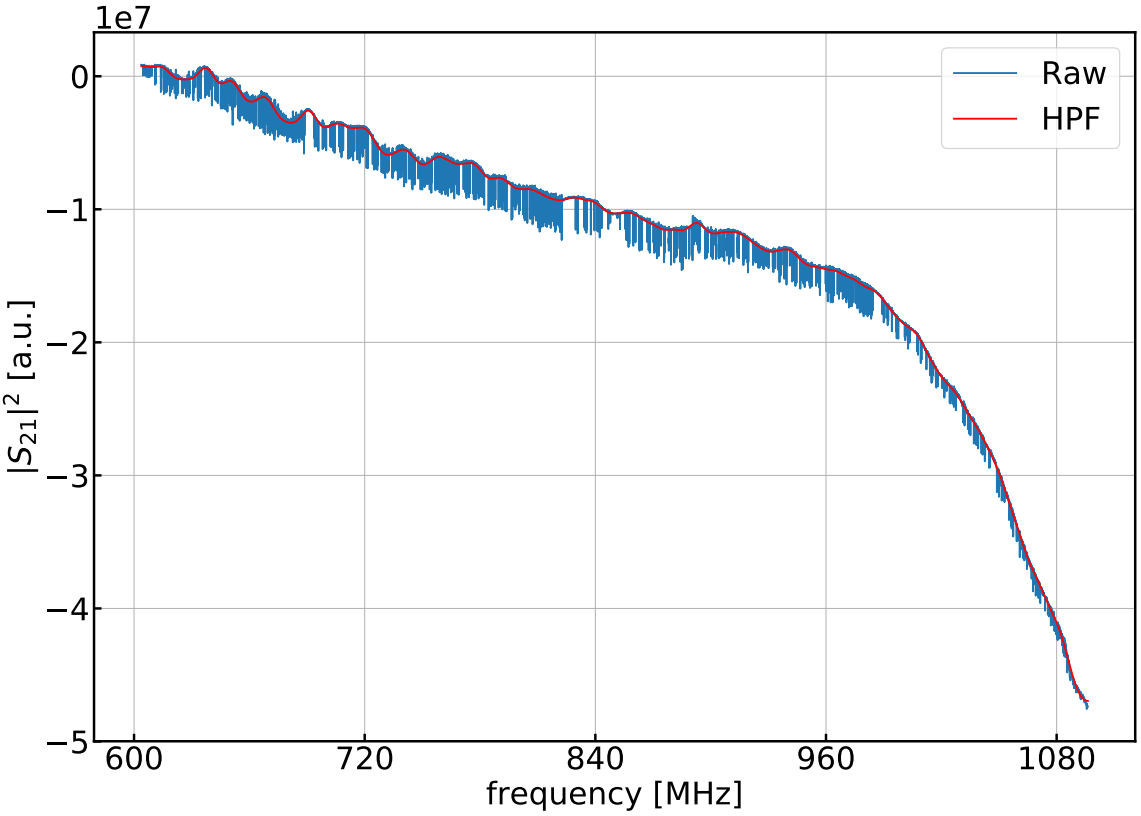

Figure 53. The channel alignment and HPF step of the KID-finding algorithm. The LPF (red) is subtracted from the S_{21} trace (blue) prior to frequency identification.

($N_{May18,VNA}$). The yields are then calculated by comparing the channel counts found during the Palestine and Antarctica integrations to $N_{May18,VNA}$. The Palestine yields ($Y_{PAL/VNA}$) are around 85% for each array. During the Palestine integration, a ~4% neutral-density filter (NDF) was installed in the array. The NDF was not present during the Antarctica integration, and the increased optical loading on the arrays made it challenging to identify all of the channels (particularly in the case of the 500 μm array). The 250V array was not used during ice testing due to a damaged component in the cold RF-chain.

Figures 54 to 58 show the results of the KID-finding algorithm applied to the

Table 11. Detector yields for each BLAST-TNG detector array.

	250U	250V [1]	250W	350 μm	500 μm
N_{design} [3]	612	612	612	775[2]	468
$N_{May18,VNA}$ [4]	500	521	495	683	381
N_{PAL} [5]	390	237	412	626	326
N_{ICE} [6]	508		466	619	260
$Y_{VNA/design}$ [7]	0.82	0.85	0.81	0.88	0.81
$Y_{PAL/VNA}$	0.78	0.45	0.83	0.92	0.86
$Y_{ICE/VNA}$	1.01		0.94	0.91	0.68

[1] The 250V cold RF chain was damaged on the ice, and the array was not used during testing.
[2] Number of detectors within the readout band of 512 MHz.
[3] Number of detectors on each array, reported by NIST.
[4] Number of detectors identified in a May, 2018 VNA sweep.
[5] Number of detectors identified in ROACH sweeps during Palestine integration.
[6] Number of detectors identified in ROACH sweeps during preflight ice integration.

VNA, Palestine and Antarctica sweeps of each array. Selected channels are marked with red stars. The dip-depth threshold is shown as a dashed red line, and masked regions are highlighted in green.

Figure 54. 350 μm array KID-finding results for a VNA sweep (top) and ROACH2 sweeps from Palestine (middle) and the ice (bottom).

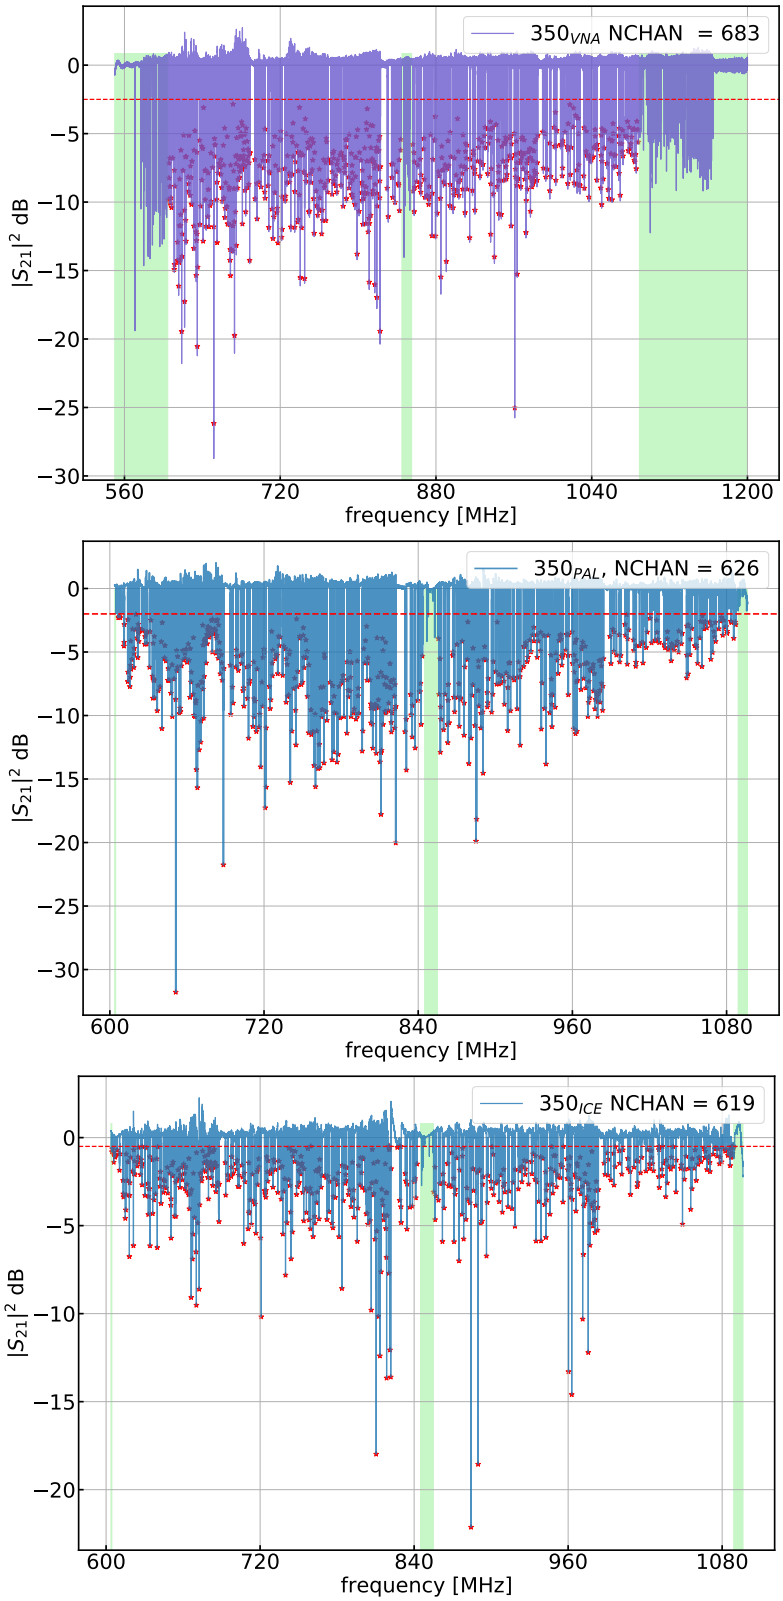

Figure 55. 250W array KID-finding results for a VNA sweep (top) and ROACH2 sweeps from Palestine (middle) and the ice (bottom).

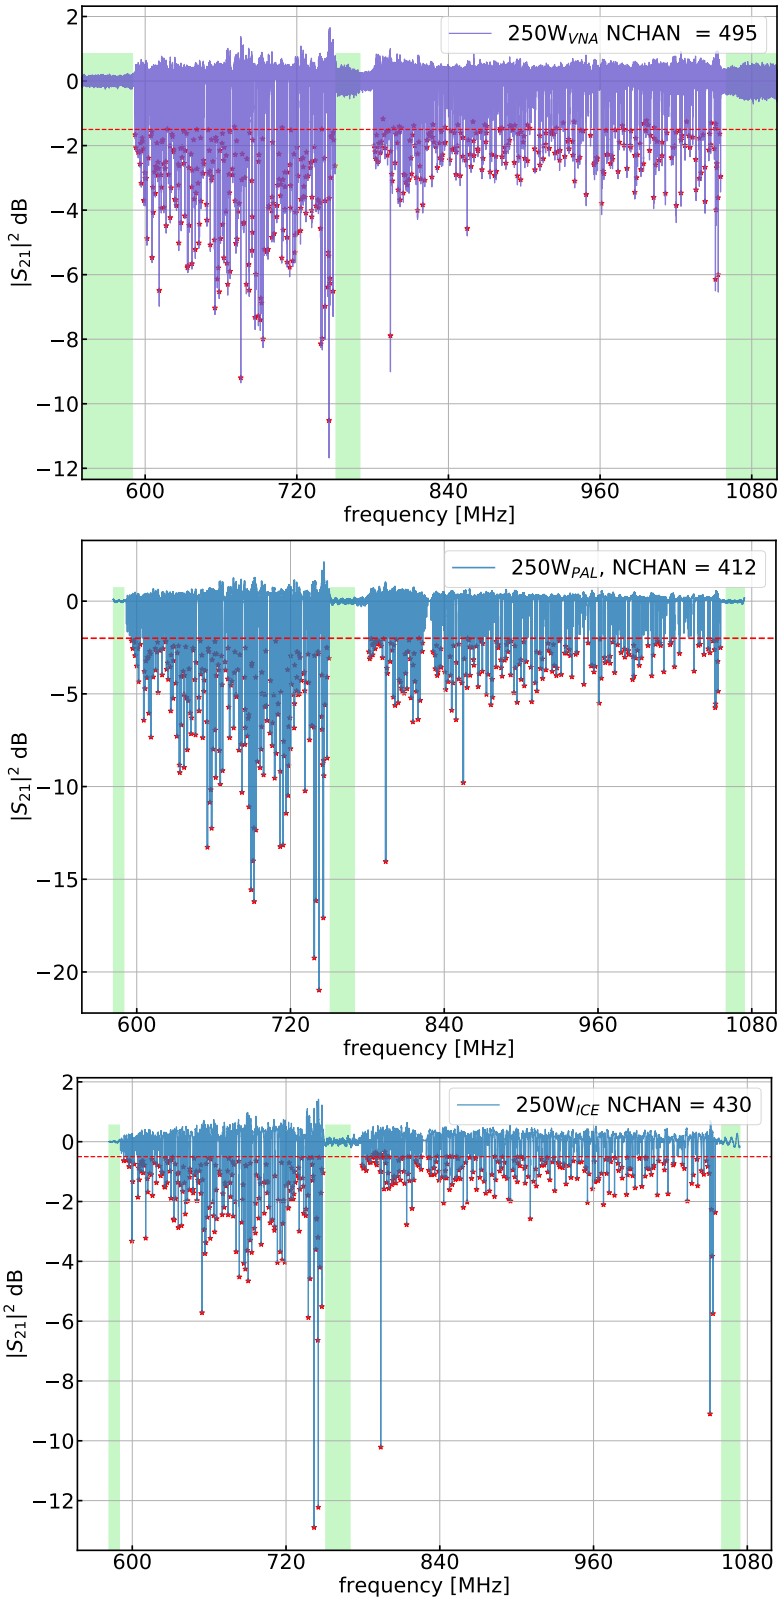

Figure 56. 250U array KID-finding results for a VNA sweep (top) and ROACH2 sweeps from Palestine (middle) and the ice (bottom).

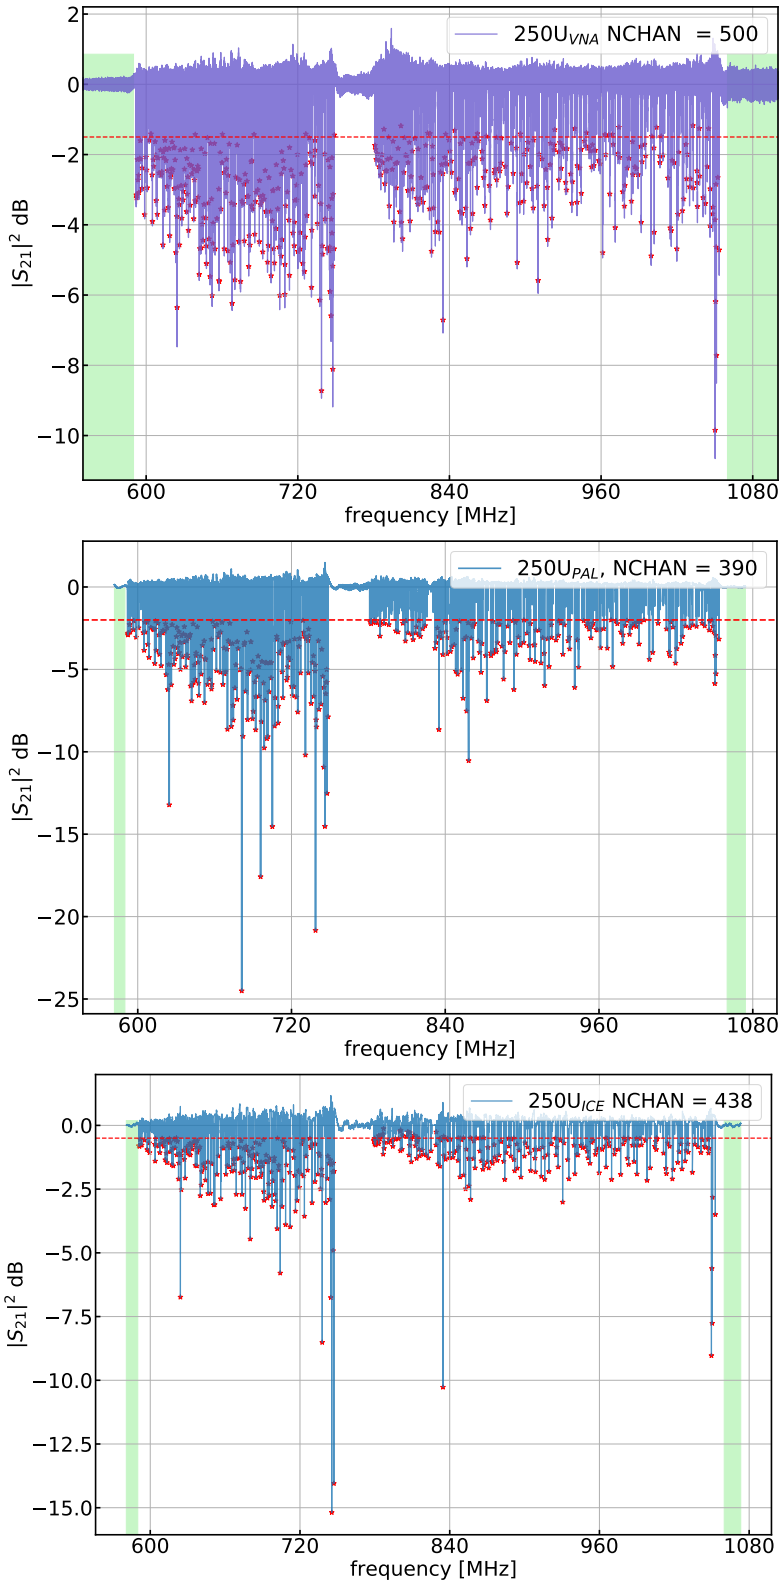

Figure 57. 250V array KID-finding results for a VNA sweep (top) and a ROACH2 sweep from Palestine (bottom).

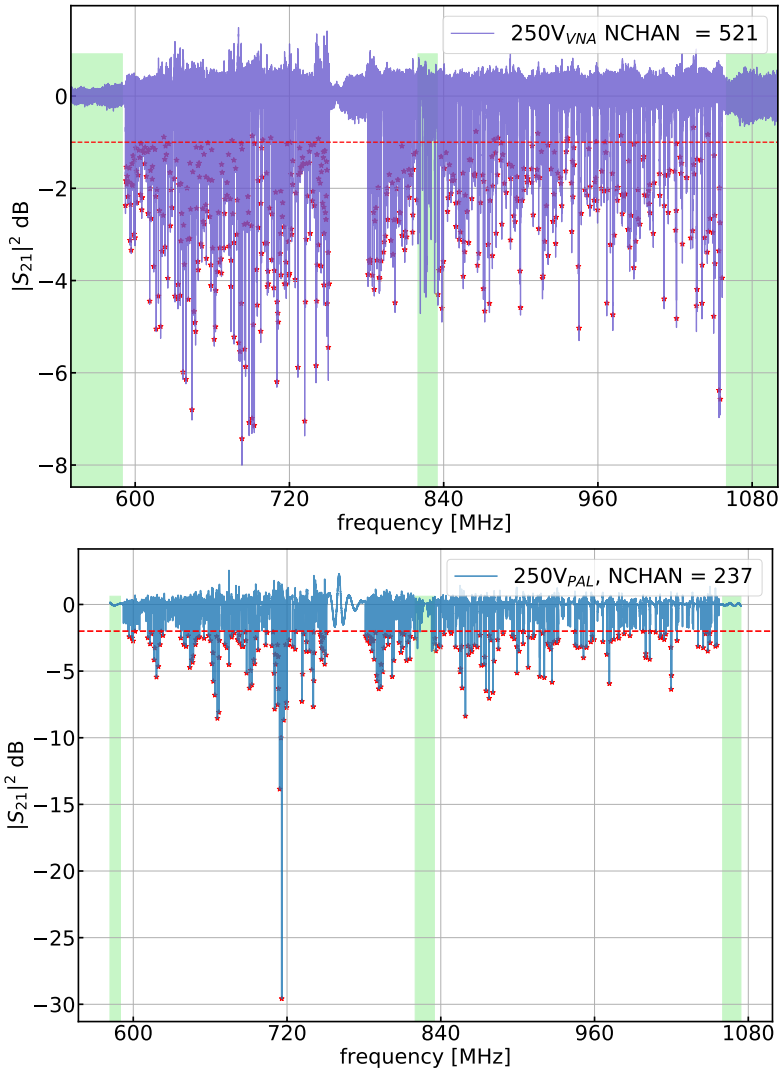

Figure 58. 500 μm array KID-finding results for a VNA sweep (top) and ROACH2 sweeps from Palestine (middle) and the ice (bottom).

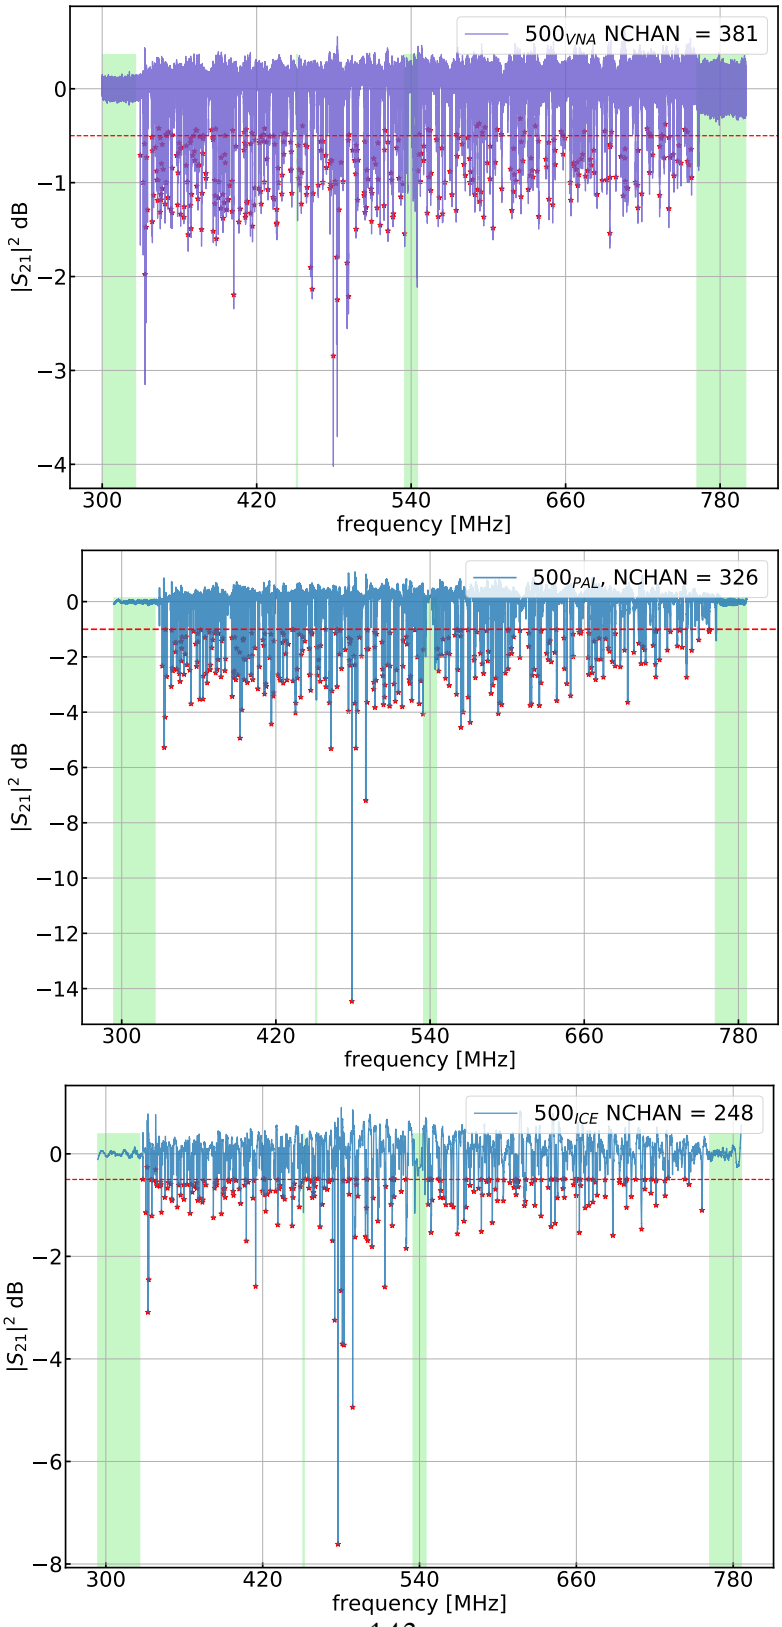

4.1.5 Target Sweeps

Once the channel frequencies for a given detector array have been identified, shorter frequency sweeps, hereafter referred to as target sweeps, may be performed in place of the wide-sweeps. Only tones which correspond to the frequency locations of resonators are used in the target sweep. The tones in the target frequency comb are referred to as channels. The initial channel amplitudes are chosen by interpolating the channel frequency onto the inverse transfer function template calculated for each array (see Section 4.1.2). During instrument operation, target sweeps are done on regular intervals to calibrate the absolute frequency response of each channel. The frequency span of the sweep can be varied in the readout software, but is typically between ∼100–250 kHz, depending on the quality factor of the resonators.

An example of a 350 μm target sweep, taken at CSBF, is shown in Figure 59. Each channel is labeled with a number (the channel index) which ranges from 0 (to the right of the LO) to N (to the left of the LO), where N is the number of channels identified in the KID-finding stage. Figure 60 shows an overlay of two target sweeps taken during the LDB integration. The purple sweep was taken with the shutter closed, and the orange with the shutter open. Several parameters which are used to calculate the absolute frequency shift of each resonator can be extracted from the target sweeps. This calculation is described in Section 4.2.1.

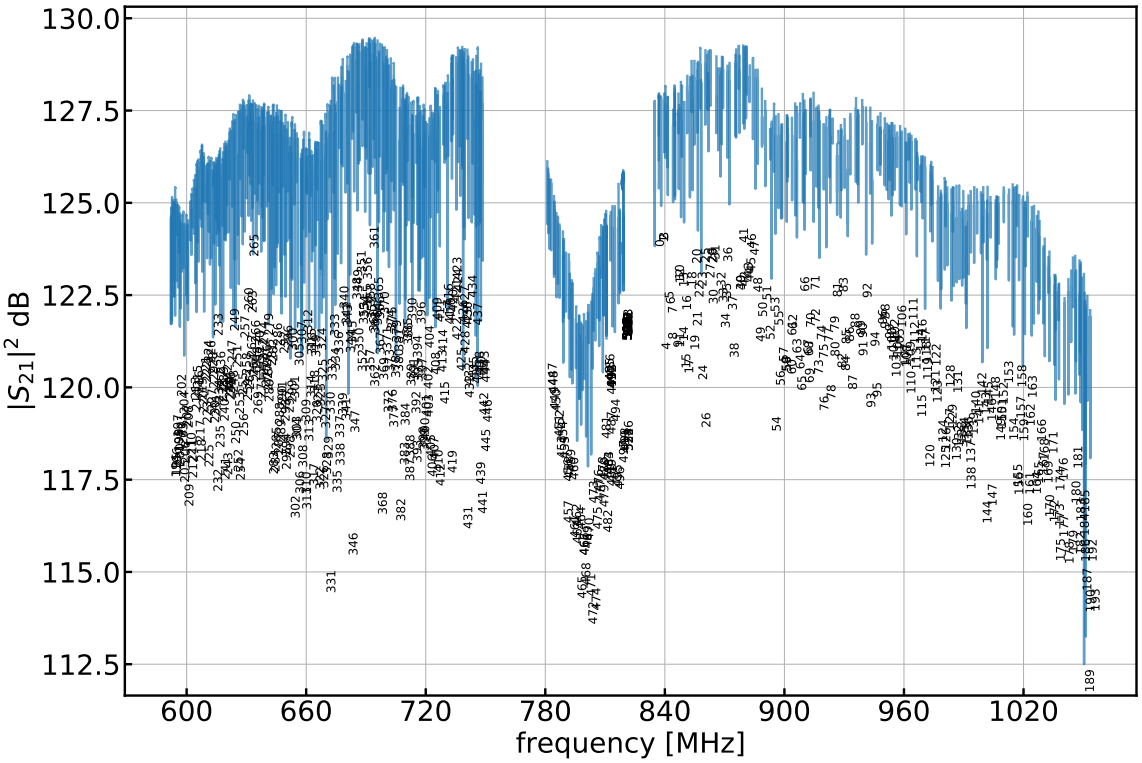

Figure 59. An example 350 μm target sweep.

4.1.6 Readout Power Optimization

The readout probe tones have an effect on resonator frequencies and quality factors. These effects, and their relation to the nonlinear kinetic inductance, are discussed in Chapter 2. Figures 61 and 62 show the effects of varying the readout power for two different resonators on one of the 250 μm arrays. In this example, the readout power was increased in six increments of 3 dB by decreasing the warm output attenuation in the ROACH2 IF electronics, producing a range of tone powers from ~ -90 to -72 dBm. The output attenuator settings, in dB, are listed in the figure legends.

The data shown in Figures 61 and 62 are overlays of target sweeps taken at each

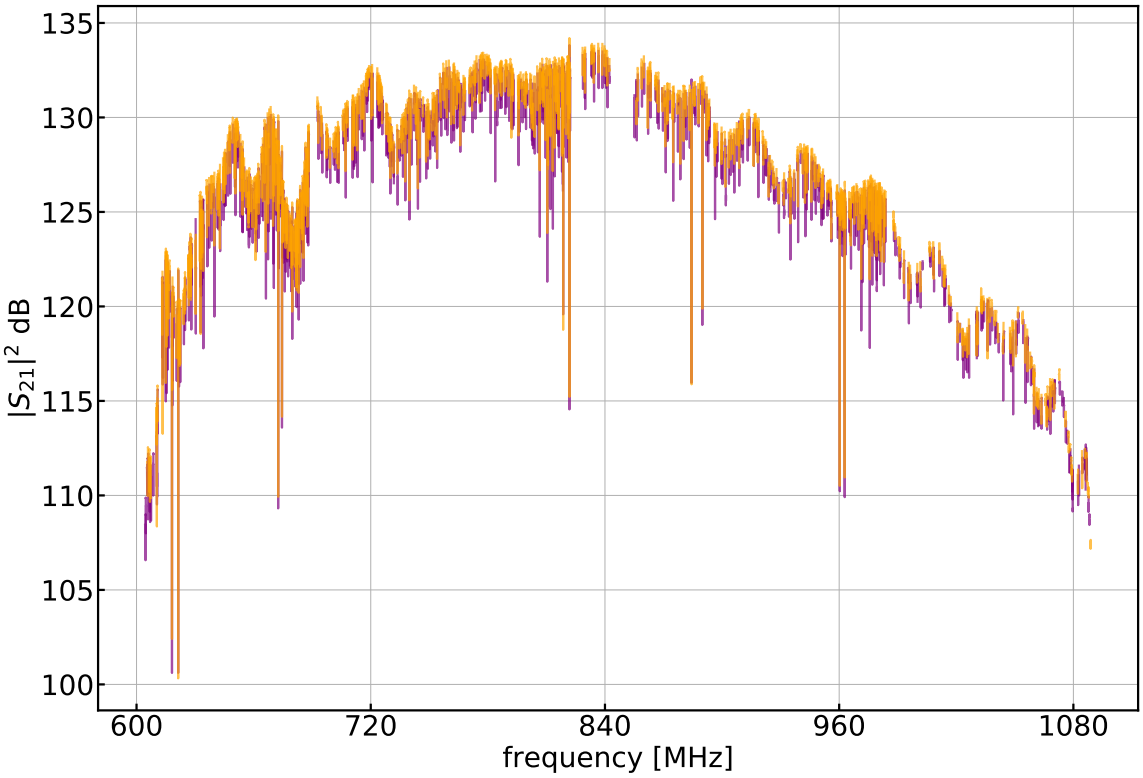

Figure 60. 350 μm target sweeps from the ice, with shutter open (orange) and closed purple.

readout power. In each figure, the top panel shows $|S_{21}|^2$, the middle panel shows the phase of S_{21}, $\phi_{I/Q}$, and the bottom panel shows the I/Q circle in the complex plane, in scaled raw units, for $S_{21}V_{\text{readout}}$. It can be seen in both figures that above a certain readout power, the resonator frequency detuning and Q degradation becomes nonlinear (due to the nonlinear kinetic inductance). For both channels, the resonator detuning reaches ~1 linewidth before the nonlinearity sets in. For the resonator in Figure 62, bifurcation sets in before the last 6 dB of tone power stepping. The bifurcation is visible as a sharp discontinuity in each panel.

At lower readout powers, the effects seen in $|S_{21}|^2$, $\phi_{I/Q}$ and in the I/Q circle are similar to what would be observed from either optical or thermal loading. As tone

power increases, the slope of the phase decreases, and the I/Q circles become more elliptical.

In practice, the optical responsivity of each resonator is maximized at a slightly different readout power. As a rule of thumb, the optimal tone power can be set for each channel by first identifying that channel's bifurcation power, and then backing it off by ∼6 dB. The tone powers can also be optimized in a global sense, by adjusting the ROACH2 output attenuator while measuring the frequency noise on and off-resonance (f_0+ ∼100 kHz). The on-resonance noise level is the sum of the LNA, detector and photon noise, whereas the off-resonance noise level reflects the LNA noise. Therefore, if the noise level on-resonance is greater than the off-resonance noise for a given channel, then the channel is assumed to be at least detector-noise limited. Above some attenuator setting, the vast majority of resonators should be biased in the detector-noise limited regime.

Figure 61. $|S_{21}|^2$, $\phi_{I/Q}$ and I/Q loops for target sweeps of a single 250 µm channel for a range of warm output attenuator settings.

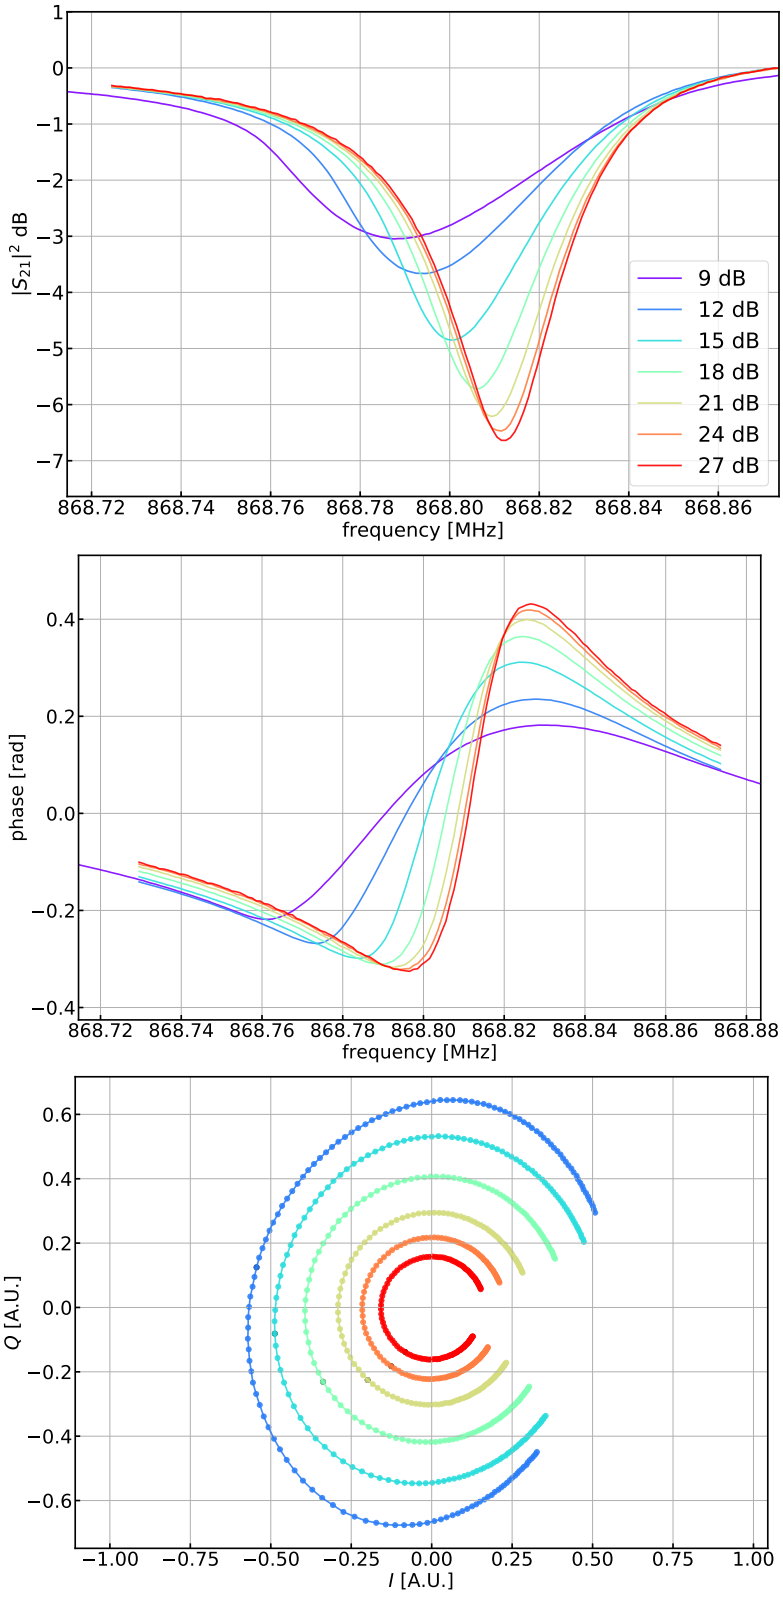

Figure 62. $|S_{21}|^2$, $\phi_{I/Q}$ and I/Q loops for a range of warm output attenuator settings, showing bifurcation.

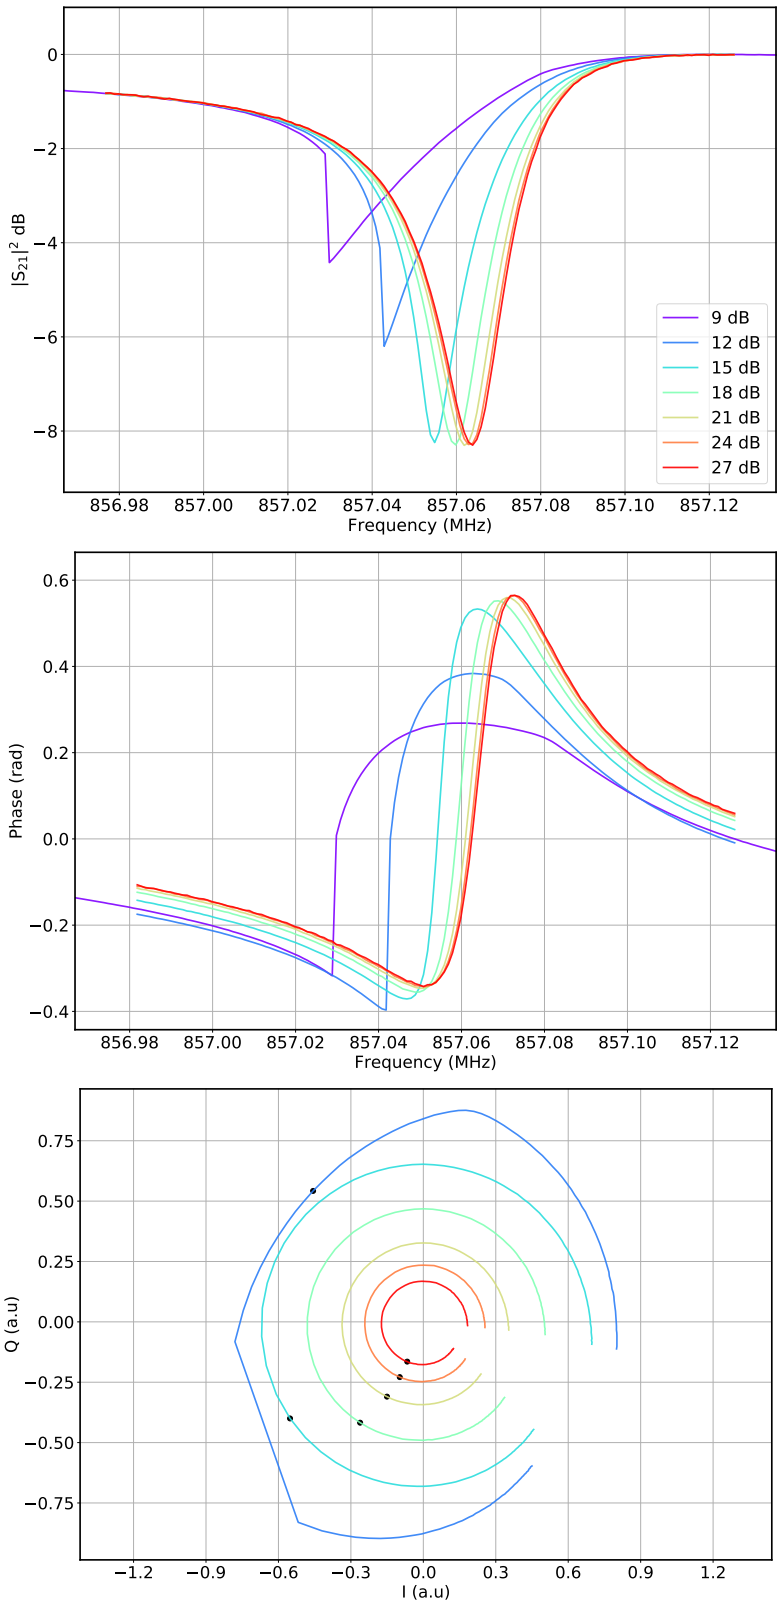

4.2 Detector Timestreams

The readout outputs two 32-bit samples, $I[n]$ and $Q[n]$, for each channel at a rate of 488.28125 Hz (once every $\approx$ 2 ms). I and Q represent the demodulated channel amplitude: $S_{21}n = I[n] + jQ[n]$. With knowledge of the readout processing gain, it would be possible to convert I and Q into absolute units of volts. However, because it is $\phi_{I/Q}$ which is proportional to the absorbed optical power, it is not necessary to perform this calibration. I and Q are therefore saved to disk in raw units, and then used to calculate the frequency shift of the resonator for each sample, Δf (which is proportional to the phase shift of the probe tone). This process is described in the following section.

4.2.1 Converting I/Q to Δf

In the following, we describe the steps of the data reduction which are taken to convert the I/Q timestreams to units of frequency shift, Δf, relative to the resonant frequency of each channel. If the frequency responsivity is known, the system NEP can be calculated by measuring the frequency noise $(e_f/f_0)^2 = S_{\mathrm{XX}} = S_{\delta f/f}$ [Hz^{-1}] (see Section 4.4).

To convert I/Q to units of frequency requires either a recent target (or wide) sweep of the resonator. Hereafter, this sweep is referred to as the reference sweep. Two methods of converting S_{21} to frequency units are commonly used. These are the I/Q gradient ($\nabla_{I,Q}$) and the phase-fitting method. The relative advantages of each method are described in detail in Barry (2014). Choosing between the two methods entails a trade-off between computational efficiency and accurate representation of

the frequency shift. The gradient method is more computationally efficient than the phase method at the expense of not accounting for the change in the shape of the resonant circle when f_0 is shifted far from resonance (more below). Ultimately, the gradient method underestimates the frequency noise by a factor of ∼1.5 (Barry, 2014).

For real-time estimates of Δf during a measurement, we find that the I/Q gradient method is preferred over the phase-fitting method. The gradient method is described in detail in D'Addabbo et al. (2013) (for a description of the phase method, see Gao (2008)). Here, we summarize its key elements. Using the reference frequency sweep, Δf (in the continuous time domain) is calculated from S_{21} as:

$$\delta f(t) = \frac{\delta S_{21}}{\nabla_f S_{21}} = \frac{\delta I(t) + j\delta Q(t)}{\nabla_{I,Q}} \tag{4.1}$$

Rationalizing Equation 4.1 yields:

$$\delta f(t) = \frac{\delta I(t) \cdot dI/df + \delta Q(t) \cdot dQ/df}{(df/dQ)^2 + (dQ/df)^2} - j\frac{\delta I(t) \cdot dQ/df - \delta Q(t) \cdot dI/df}{(df/dQ)^2 + (dQ/df)^2} \tag{4.2}$$

where all of the I/Q values are with respect to the resonant frequency, $\delta I(t) = I(t) - I_{\mathrm{ref}}$ and I_{ref} (Q_{ref}) are the I/Q values on resonance at the time at which the reference sweep was taken. The Δf values produced by this method are valid only if the resonator has shifted by less than a linewidth (see below). Figure 64 shows an example $\nabla_{I,Q}$ calculation for one channel (channel 448) of the 350 μm array. The value of $\nabla_{I,Q}$ which is used in Equations 4.1 and 4.2 is taken where $|\nabla_{I,Q}|$ at its maximum. This point is marked with the dotted black line in Figure 64.

In practice, SNR variations in the target sweeps translate into errors in the $\nabla_{I,Q}$ calculation. During a measurement, it is possible to calibrate $\nabla_{I,Q}$ by chopping the

LO in frequency by a fixed amount both above and below the center frequency. If the probe tone is perfectly centered on the resonant frequency, and Q_r is moderately high, the Δf step observed in the timestreams for each channel should equal the chop Δf. If the probe tone is within a resonator linewidth of f_0, but off-center relative to f_0, the Δf observed in the timestreams will be asymmetric. If the probe tone is too far out of range of f_0, the LO chop will have little to no observable effect in the channel timestreams.

If the probe tone is known to be on or very near resonance, the LO chop provides a good indicator of the accuracy of the $\nabla_{I,Q}$ values. During the BLAST-TNG flight the LO chop will performed during each azimuth turnaround (the data acquired during these turnaround will not be used to make science maps). The Δf chops which are imprinted into each channel's timestream can be used to calibrate the I/Q data post-flight.

Figure 63 shows an example LO chop for one channel from the 250U, 350 and 500 μm arrays, taken during pre-flight preparations at LDB. The frequency step-size is 2.5 kHz for the 250 and 350 μm arrays, and 10 kHz for the 500 μm. At the time when these LO chops were recorded, the probe tones were very close to f_0, and the channel Δf values match the LO chop. The spikes in the data located around the frequency jumps are due to drop-outs in the LO signal when the synthesizer is stepped.

4.2.2 Phase and Dissipation Quadratures

In Equation 4.2, the real part of Δf, $\Re(\Delta f)$, represents the total frequency fluctuation in the phase quadrature (tangent to the I/Q circle in the complex-plane),

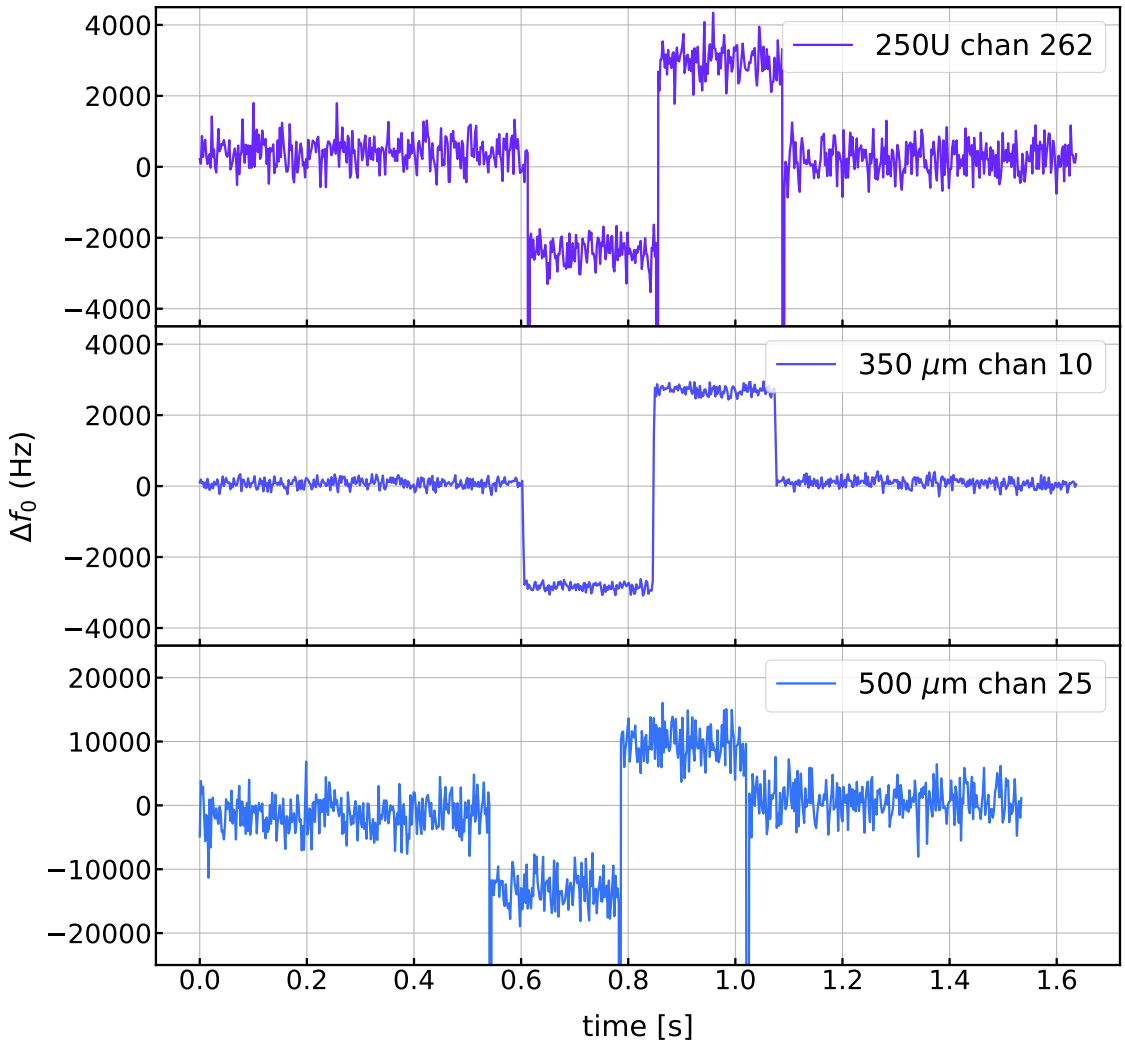

Figure 63. LO chops for the 250U, 350 and 500 μm arrays.

and the imaginary part of Δf, $\Im(\Delta f)$, represents the fluctuation in the dissipation quadrature (orthogonal to the I/Q circle). The noise in the dissipation quadrature is primarily from the LNA ($(e_{f,\mathrm{diss}}/f_0)^2 = S_{\mathrm{YY}}$). The noise in the phase quadrature which is not due to the amplifiers can be estimated as $S_{\mathrm{XX}} - S_{\mathrm{YY}}$. If the detectors are photon noise limited, then $S_{\mathrm{phot}} \simeq S_{\mathrm{XX}} - S_{\mathrm{YY}}$. In the photon noise limited regime, S_{XX} should be greater than S_{YY}.

Figure 65 shows a 10 s timestream for channel 448 of the 350 μm array, recorded at CSBF. The timestream is shown as I (top, raw units), Q (second from top, raw units), $\Re(\Delta f)$ (second from bottom, Hz) and $\Im(\Delta f)$ (bottom, Hz). Figure 66 shows $|S_{21}|^2$ (top), $\phi_{I/Q}$ (middle) and the I/Q circle for the example channel. The dashed red line in $|S_{21}|^2$ is a fit to the data which has been used to estimate the resonator parameters listed in Table 12. The estimated parameters are: f_0, dip-depth, Q_r, Q_c, Q_i, ϵ_{a} and $S_{\delta f/f}$. 10 seconds of timestream values are shown as a scatter ball (orange) in the phase and I/Q circle panels. To plot $\phi_{I/Q}$ and the I/Q circle, the I/Q values have been rotated so that the minimum value of I coincides with the resonance point. This rotation prevents phase wrapping when plotting $\phi_{I/Q}$, and is achieved by multiplying S_{21} by a complex exponential, where

$$S_{21,rot} = S_{21} e^{j\phi_{rot}} \tag{4.3}$$

and $\phi_{rot} = \arctan 2(\langle Q(t) \rangle, \langle I(t) \rangle)$.

Using the resonator parameters which were fit from the S_{21} sweep and I/Q timestream for channel 30, the frequency noise can be estimated without calculating the I/Q gradient $\nabla_{I,Q}$ as:

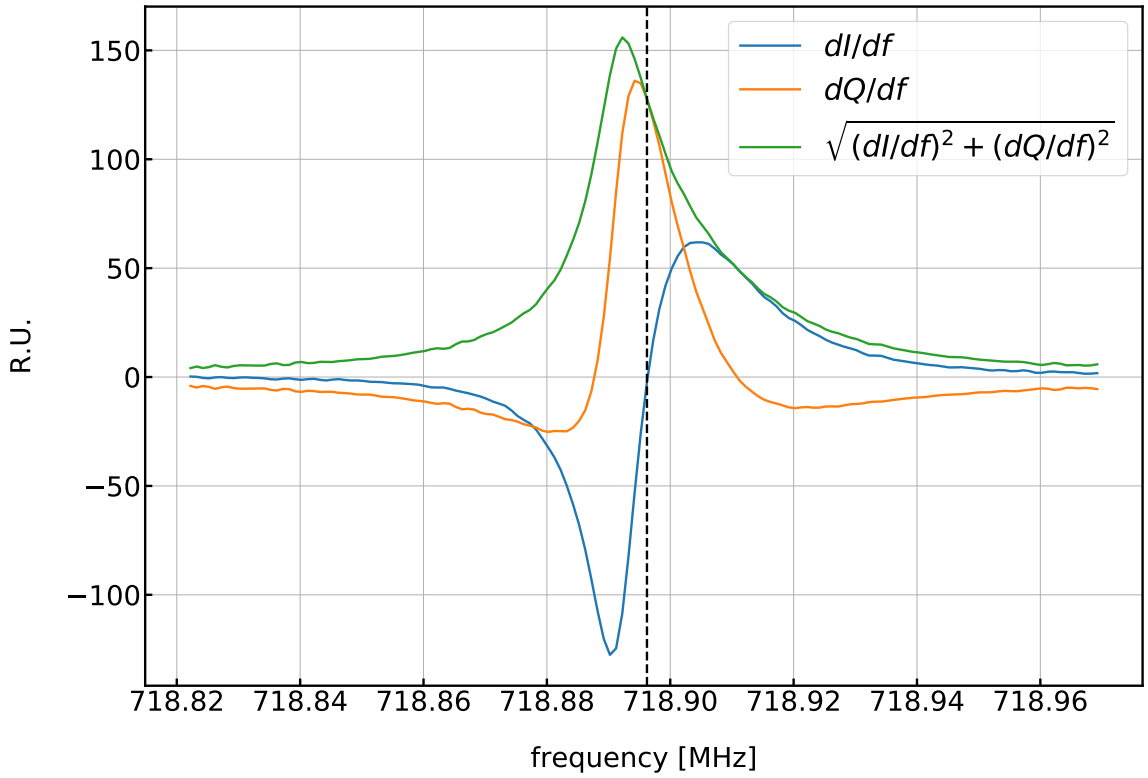

Figure 64. An example of the I/Q gradient $\nabla_{I,Q}$ calculation for a single 350 µm channel.

f_0 (MHz)	Depth (dB)	Q_r	Q_c	Q_i	ϵ_a	$(dI/df, dQ/df)$ (raw)	$S_{\delta f/f}$ (Hz^{-1})
718.90	-8.04	32,826	54,460	83,3325	0.17	-2.28, 127.68	2.78×10^{-18}

Table 12. Resonator parameters for channel 448 of the 350 µm array.

$$\begin{aligned} e_f &= e_v \frac{f_0 Q_c}{2{Q_r}^2} \\ &= \sqrt{\frac{(\sigma_I^2 + \sigma_Q^2)/(f_s/2)}{I^2 + Q^2}} \frac{f_0 Q_c}{2{Q_r}^2} \qquad \left[\frac{\mathrm{Hz}}{\sqrt{\mathrm{Hz}}}\right] \end{aligned} \tag{4.4}$$

where Q_r and Q_c are fit using the sweep, f_s is the readout sampling frequency (e.g., 488.28125 Hz), and the voltage noise e_v ($\mathrm{V}/\sqrt{\mathrm{Hz}}$) is calculated using I/Q at their maximum off-resonance values for each channel.

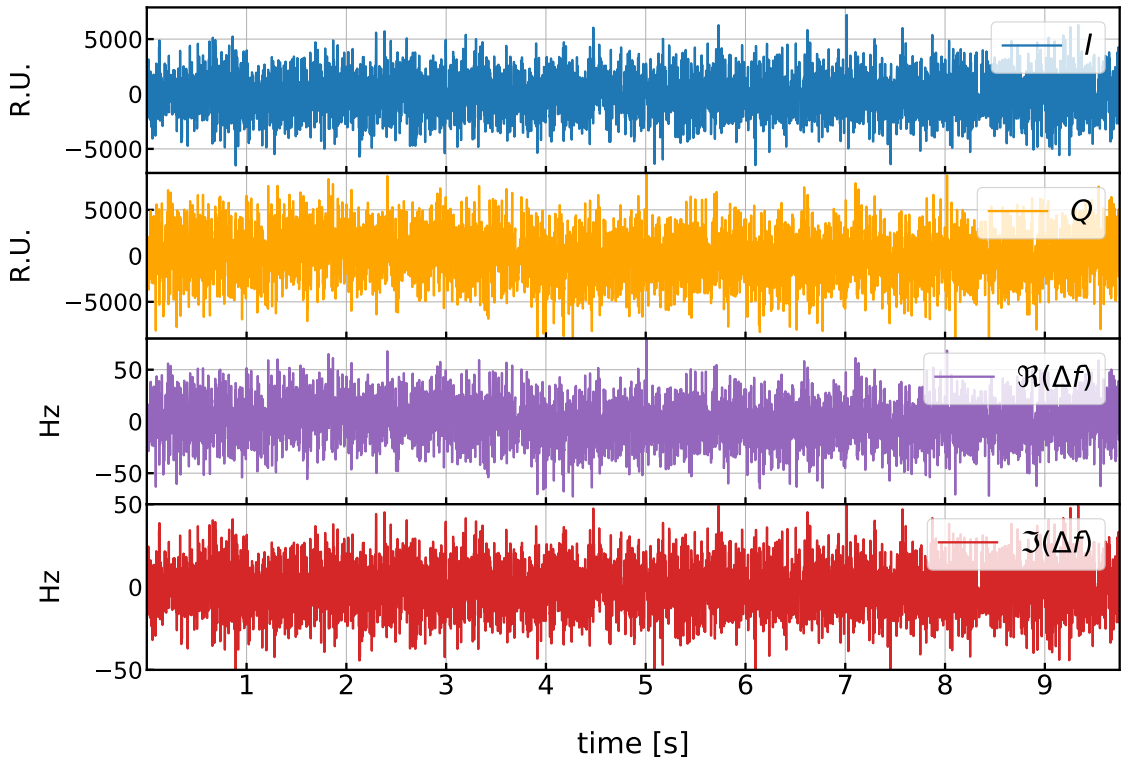

Figure 65. 10 second timestreams for the Palestine 350 µm example channel. The quantities shown are (top to bottom) I (raw units), Q (raw units), $\Re(\Delta f_{\rm res})$ [Hz] and $\Im(\Delta f_{\rm res})$ [Hz].

Using the example channel values for f_0, Q_r and Q_c from Table 12 in Equation 4.4 gives $e_f \approx 1.5$ Hz/$\sqrt{\rm Hz}$. This is equal to the frequency noise calculated for this channel using the gradient method.

4.2.3 Palestine Integration Detector Histograms

Because there are over 2,500 BLAST-TNG detectors, it would be impractical to present results for each one. Figures 69 to 76 show histograms for the measured dip-depths, Q_r, Q_r/Q_c, and electrical NEP. The I/Q timestreams used to produce these

Figure 66. $|S_{21}|^2$ (top), $\phi_{I/Q}$ (middle) and the I/Q loop (bottom) for channel 448 of the 350 µm array, measured at CSBF. 10 seconds of timestream values are shown as a scatter ball (orange) in the phase and I/Q circle panels.

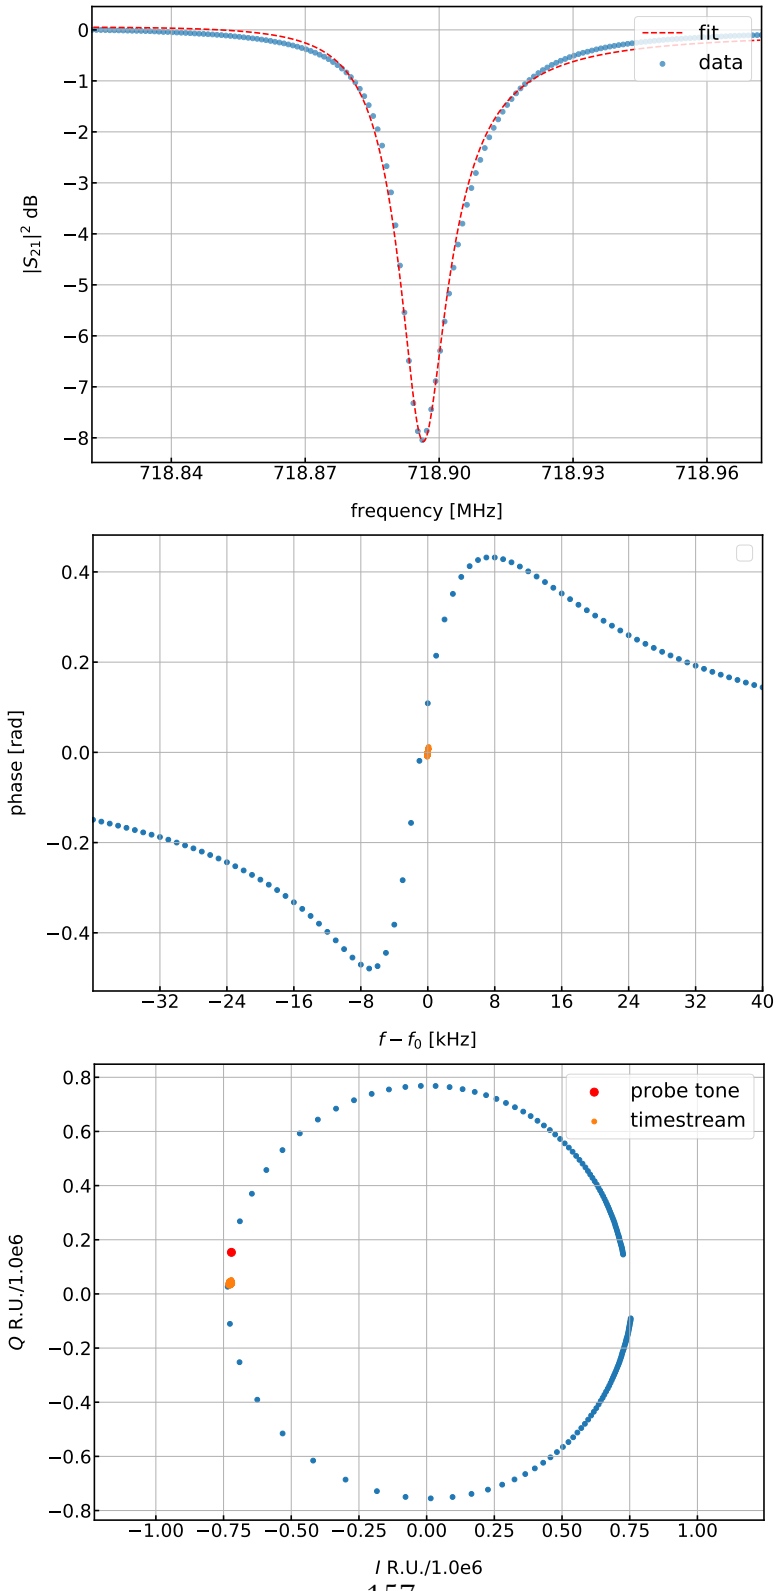

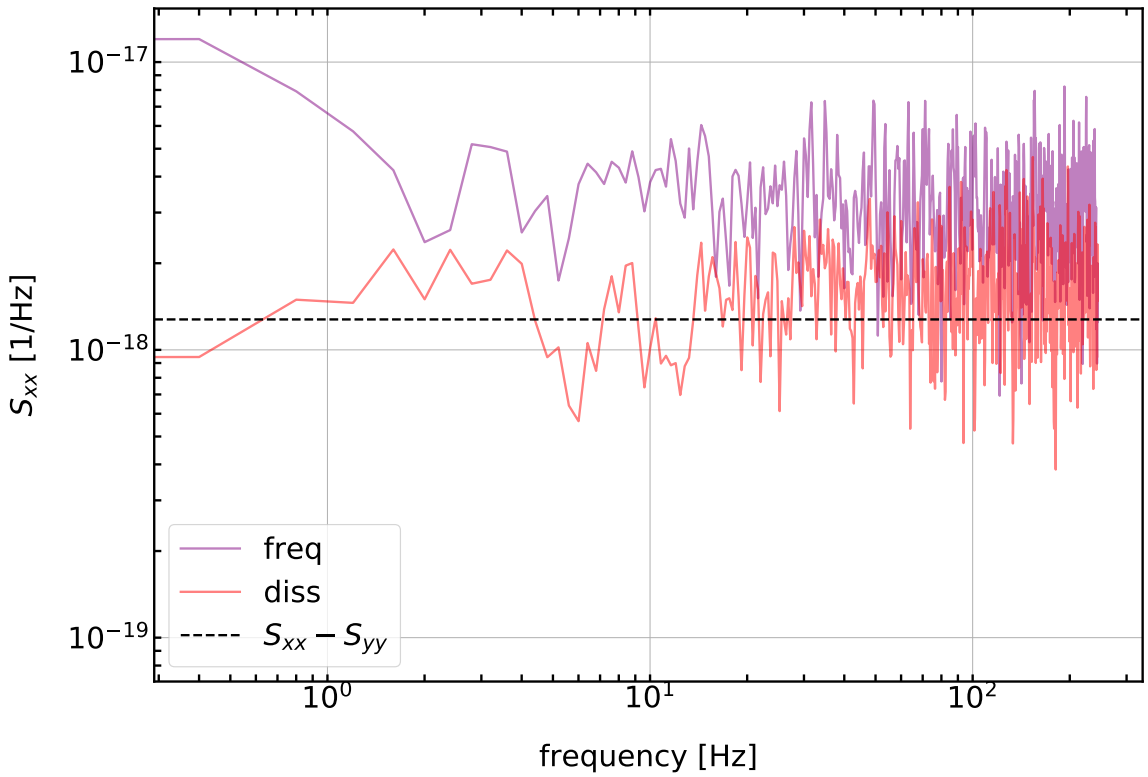

Figure 67. A 350 μm channel PSD, showing S_{XX} and S_{YY}.

estimates was acquired during passband mapping, while the FTS was radiating into the cryostat window (see Section 4.6.1). No data containing interferograms was used in this analysis, although the level of optical loading on the detectors is assumed to be slightly higher than the ambient 290 K of the CSBF highbay. Throughout the measurement, A 4% cryogenic NDF was present in the optical chain. The median values for each of the measured quantities is listed in Table 13. The values for the electrical NEP will be used to convert to noise-equivalent flux density in Section 4.4.2.

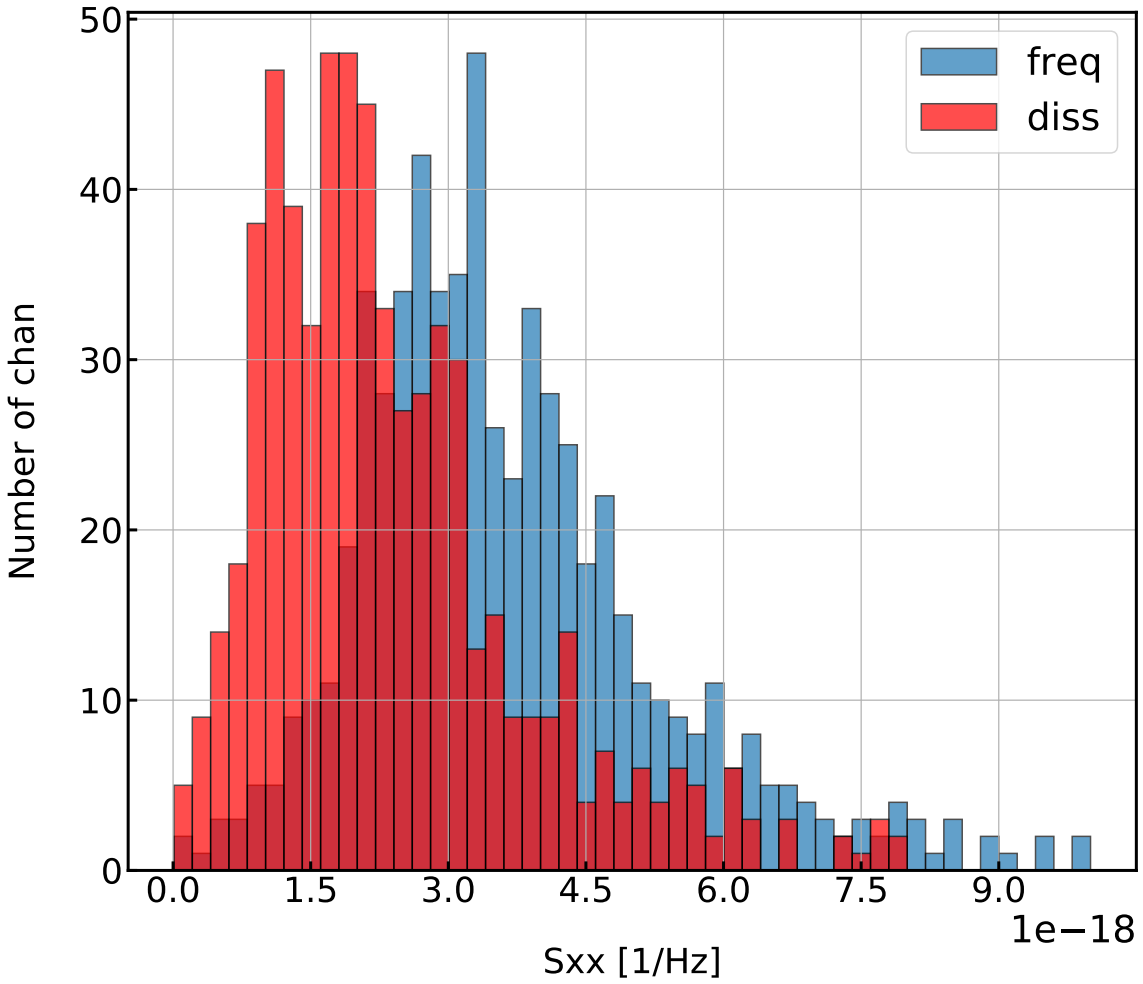

Figure 68. A histogram of S_{XX} and S_{YY} for the 350 µm array, as measured in Palestine, TX.

Table 13. Median values of important detector parameters for all five BLAST-TNG arrays, measured at CSBF.

	500 µm	250V	350 µm	250U	250W
Q_r	5361	12,759	12,759	17,301	16,942
Q_r/Q_c	0.073	0.28	0.56	0.24	0.22
Depth (dB)	-0.74	-3.22	-7.81	-2.52	-2.28
S_{XX} (1/Hz)	28×10^{-17}	30×10^{-17}	0.34×10^{-17}	3.5×10^{-17}	4×10^{-17}
S_{YY} (1/Hz)	25×10^{-17}	27×10^{-17}	0.21×10^{-17}	2.4×10^{-17}	2.5×10^{-17}
$\mathrm{NEP_{freq}}$ W/$\sqrt{\mathrm{Hz}}$	10×10^{-17}	100×10^{-17}	35×10^{-17}	35×10^{-17}	37×10^{-17}

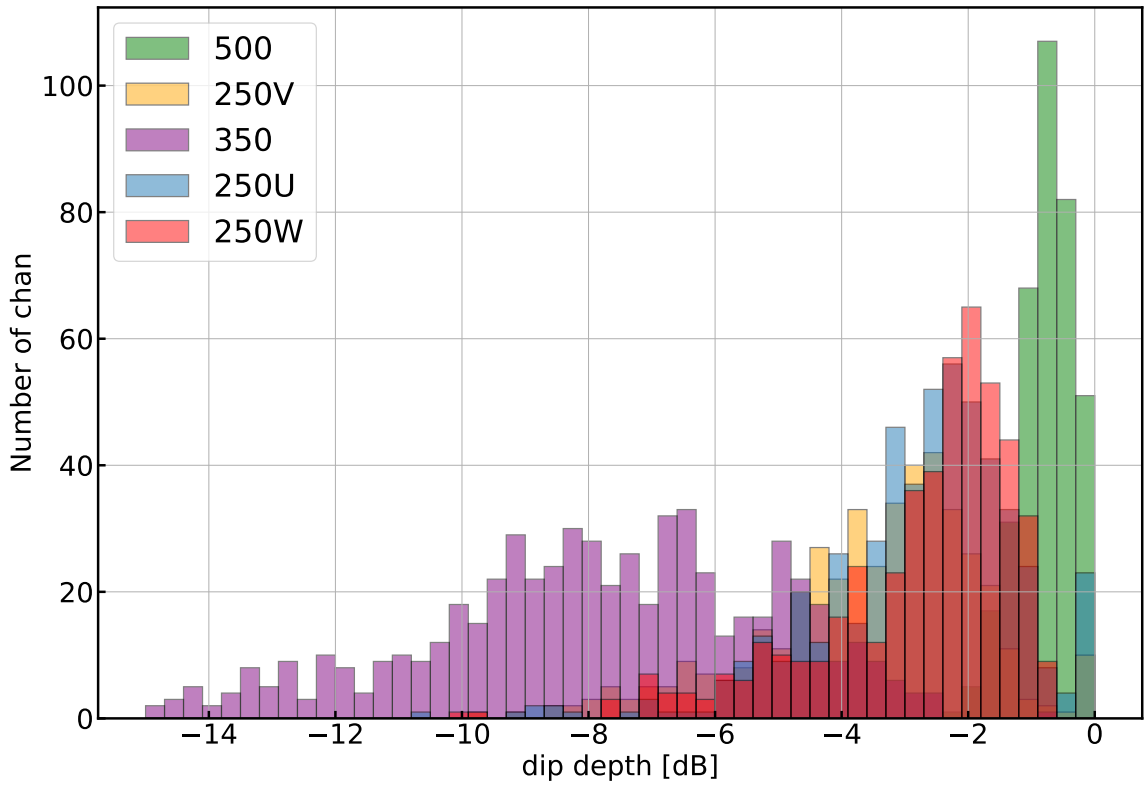

Figure 69. Histograms of resonator dip-depth for all five BLAST-TNG arrays, measured at CSBF.

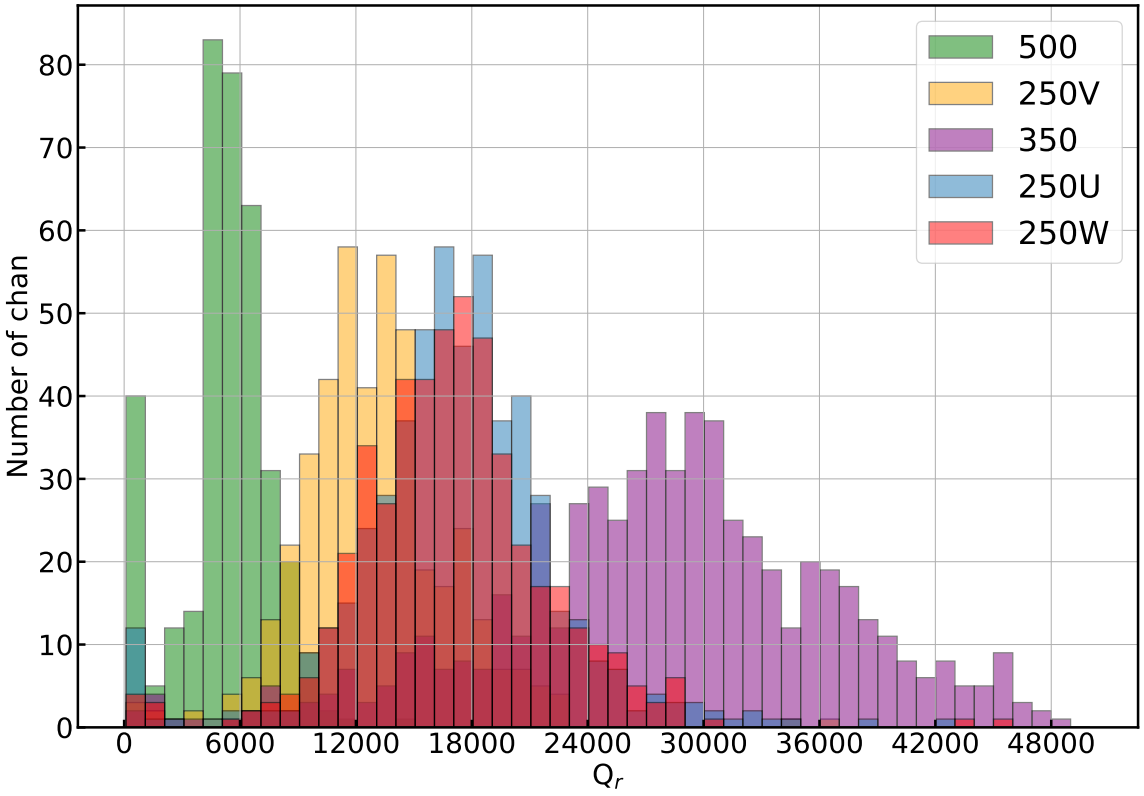

Figure 70. Histograms of Q_r for all five BLAST-TNG arrays, measured at CSBF.

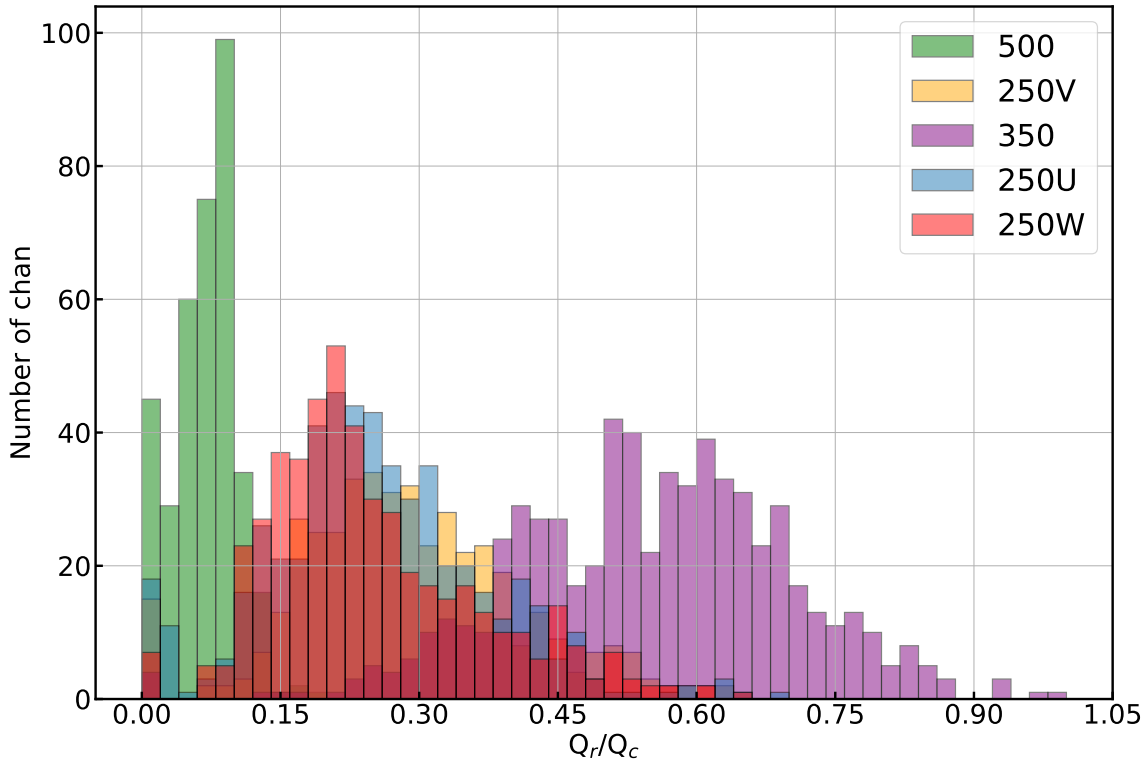

Figure 71. Histograms of Q_r/Q_c for all five BLAST-TNG arrays, measured at CSBF.

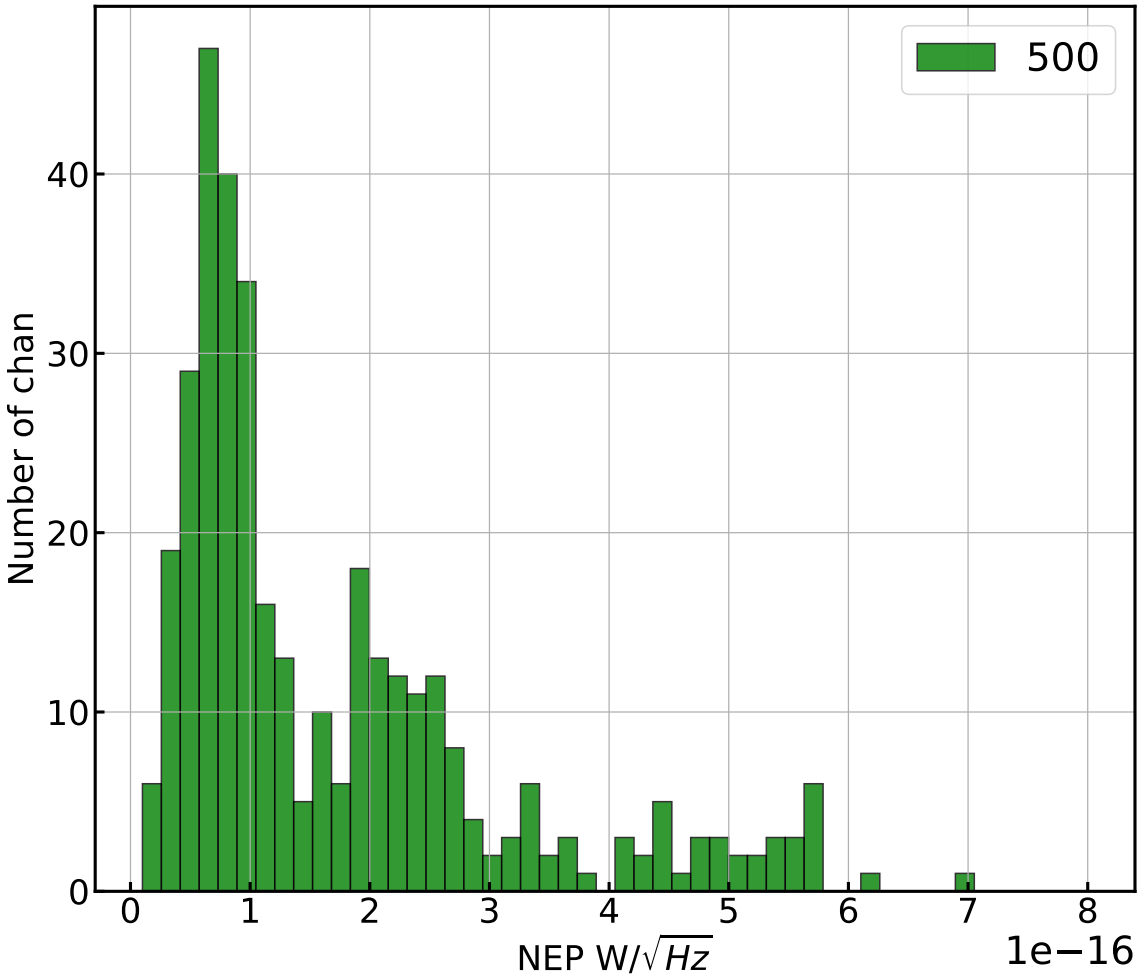

Figure 72. Histogram of NEP_{freq} for the 500 μm array, measured at CSBF.

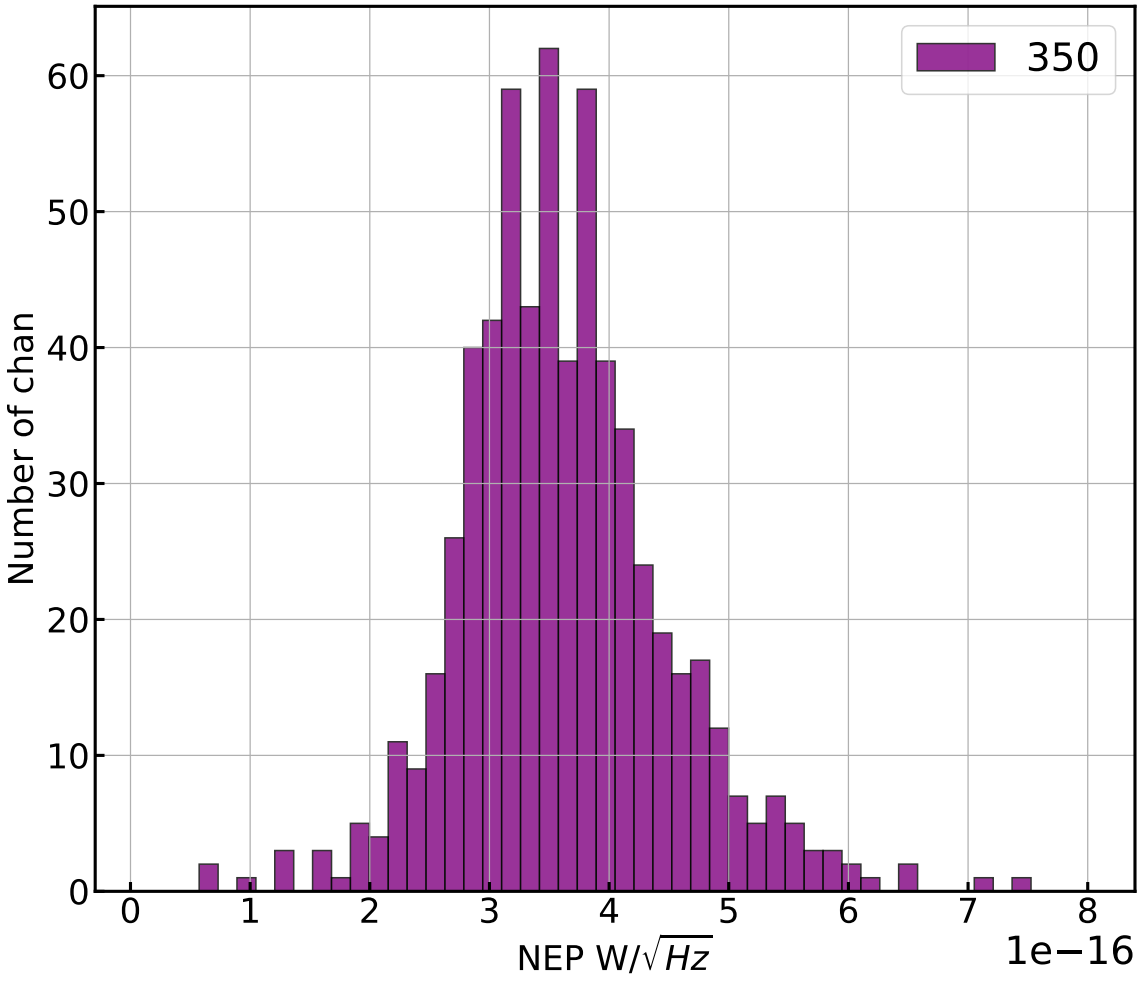

Figure 73. Histogram of NEP_{freq} for the 350 μm array, measured at CSBF.

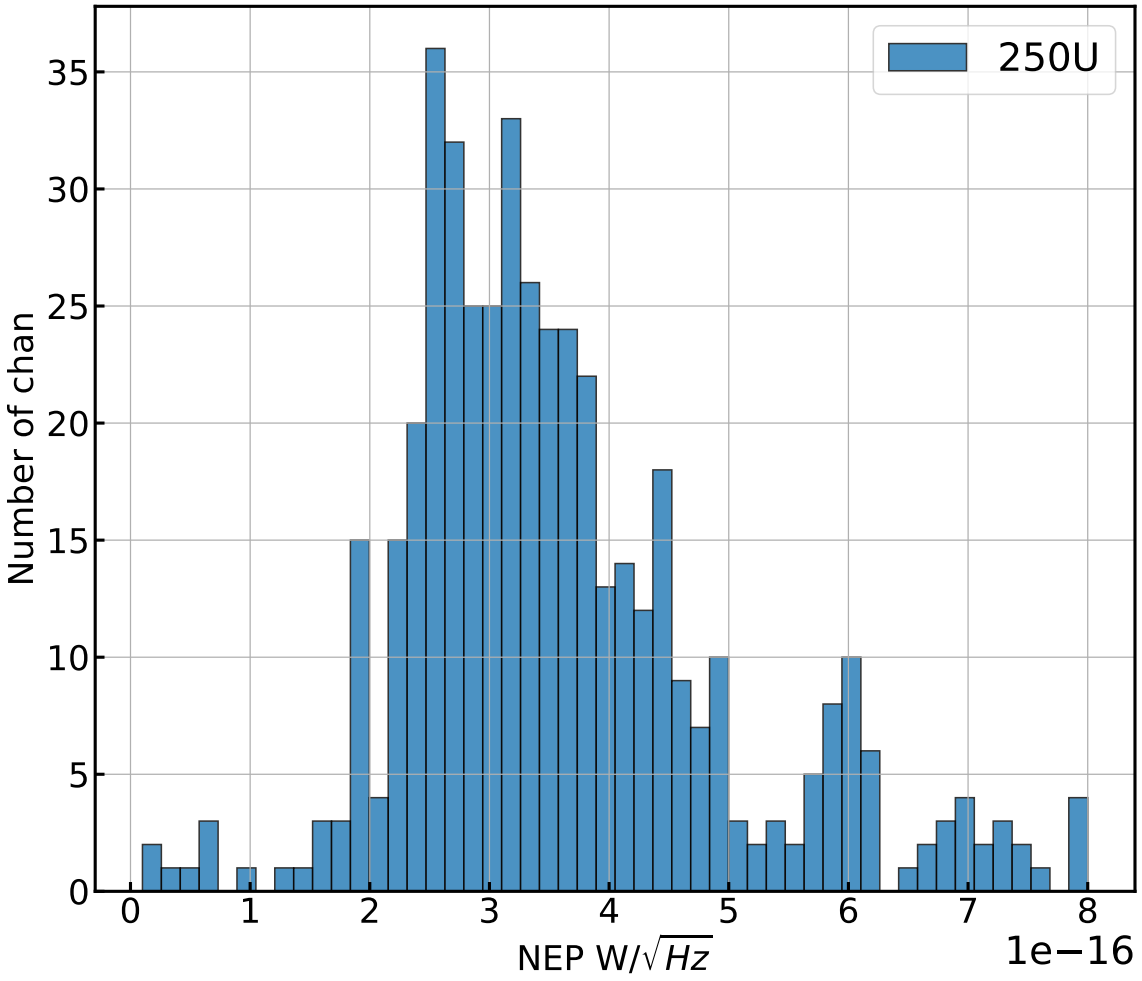

Figure 74. Histogram of NEP_{freq} for the 250U array, measured at CSBF.

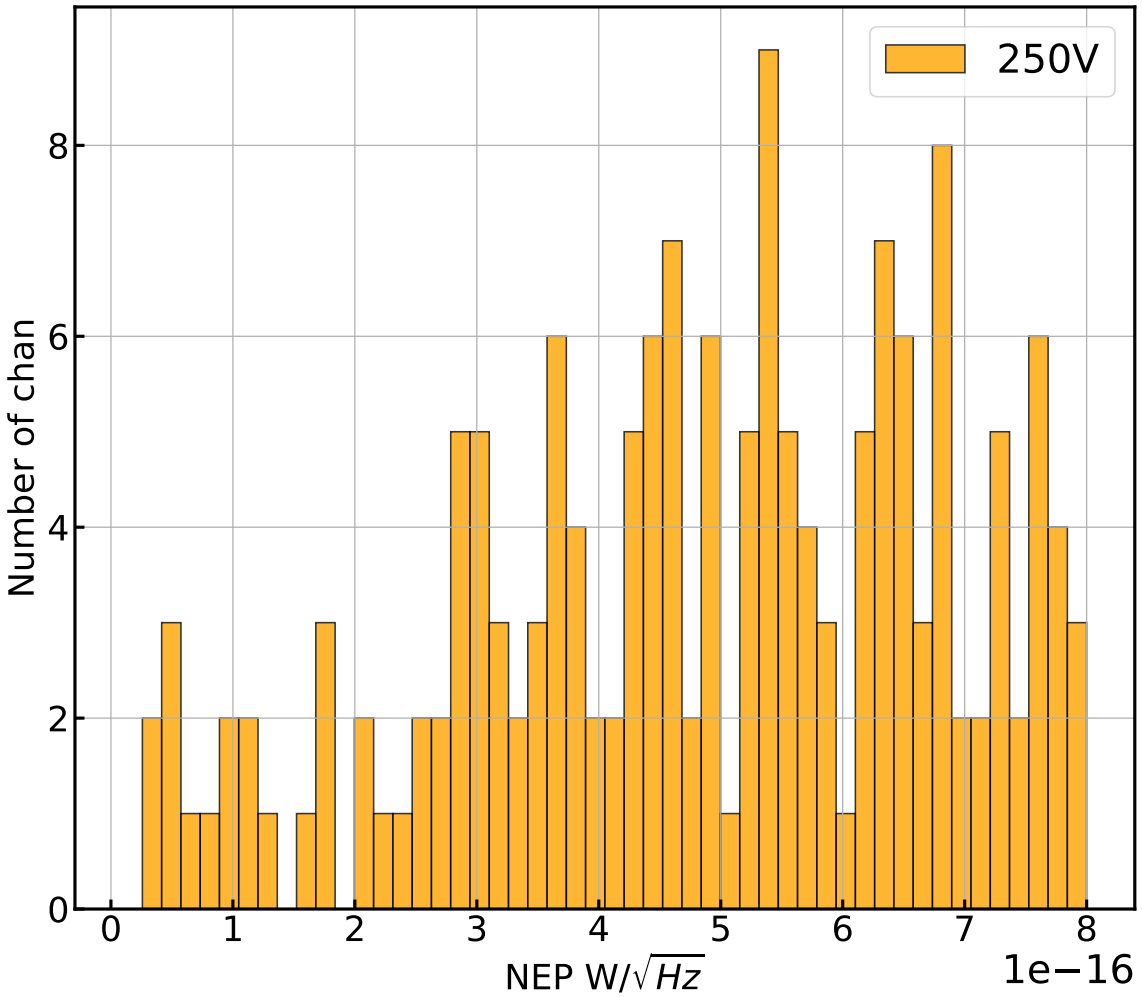

Figure 75. Histogram of NEP_{freq} for the 250V array, measured at CSBF.

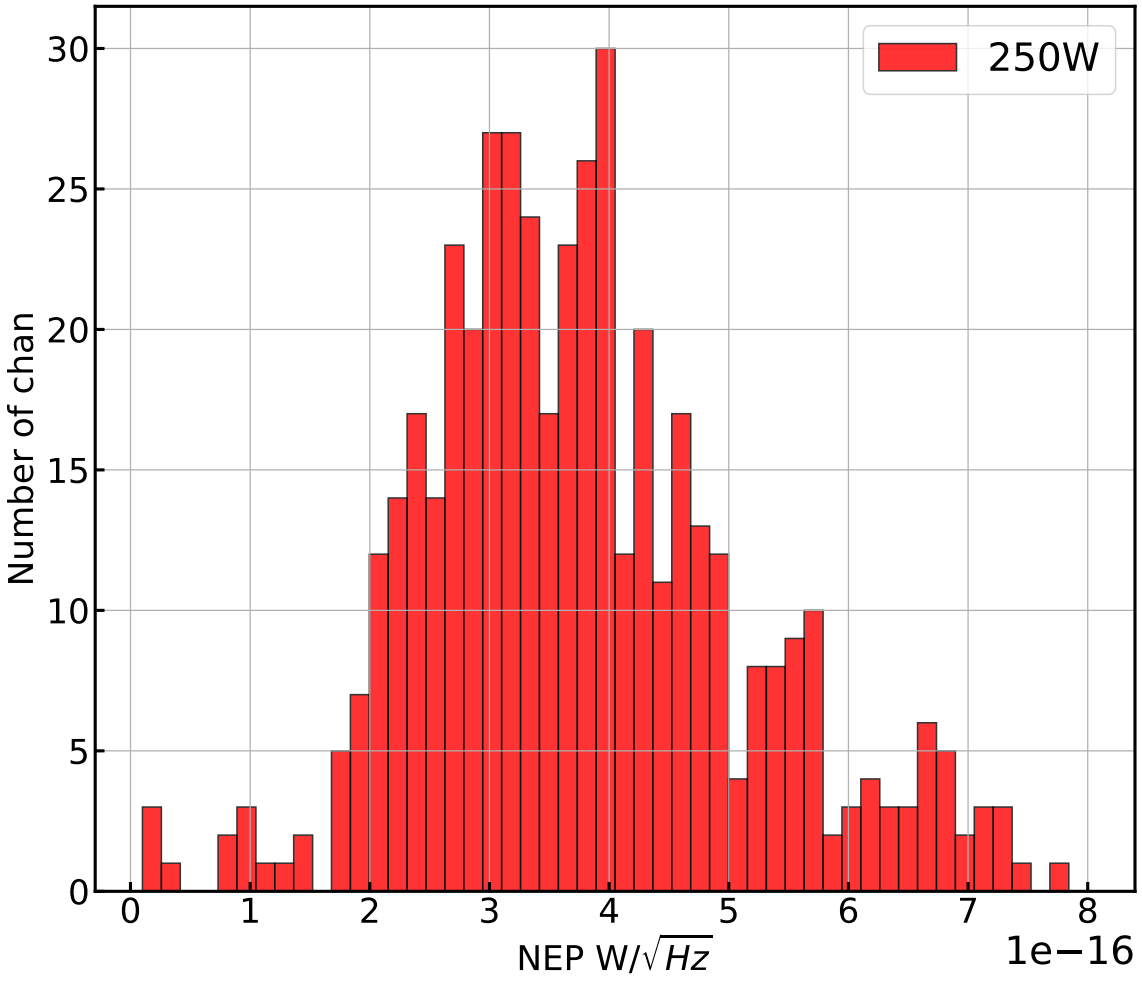

Figure 76. Histogram of NEP_{freq} for the 250W array, measured at CSBF.

4.3 Responsivities

Average responsivities to absorbed optical power, $R_P = (df/f_0)/\mathrm{pW}$, and to changes in base temperature, $R_T = (df/f_0)/\mathrm{mK}$, are listed in Table 14 for the 250, 350 and 500 μm detector arrays. The optical responsivity R_P is difficult to measure, because it requires the use of a well calibrated black-body source, as well as knowledge of the total optical efficiency of the detector test-bed. For the horn-coupled, polarization sensitive LEKIDs used in BLAST-TNG, the total optical efficiency includes the horn-efficiency, polarization efficiency, detector quantum efficiency and the efficiency of any optical elements, such as band-defining filters, which are present in the optical chain. Furthermore, R_P and R_T are different for each detector. In general, either the optical efficiency of the system or the optical responsivity may be estimated by assuming a value for one or the other.

4.3.1 Optical Responsivity

In this work, we assume average values of R_P for each BLAST-TNG detector which were measured at NIST (shared in a communication with collaboration members). In Section 4.5, we compare these values to those produced by the parametric LEKID model presented in Chapter 2. The BLAST-TNG receiver contains a calibration lamp which is of a similar design to the one which was used in Herschel SPIRE (Hargrave et al., 2006). Because the absolute power output from the lamp is unknown, it is used primarily as a relative response calibrator. Examples of 1 second calibration lamp chops from CSBF and LDB are shown in Figures 77 and Figure 78. The CSBF chops are shown in units of normalized $\phi_{I/Q}$, and represent an average

over many channels. The rising edge of the pulse is visible. The chops shown in Figure 78 are for single channels, and have been converted to Δf. A low-pass filtered trace for each chop is overlayed on the raw data. Due to the lower SNR, the rising edge of the lamp pulse is not as visible.

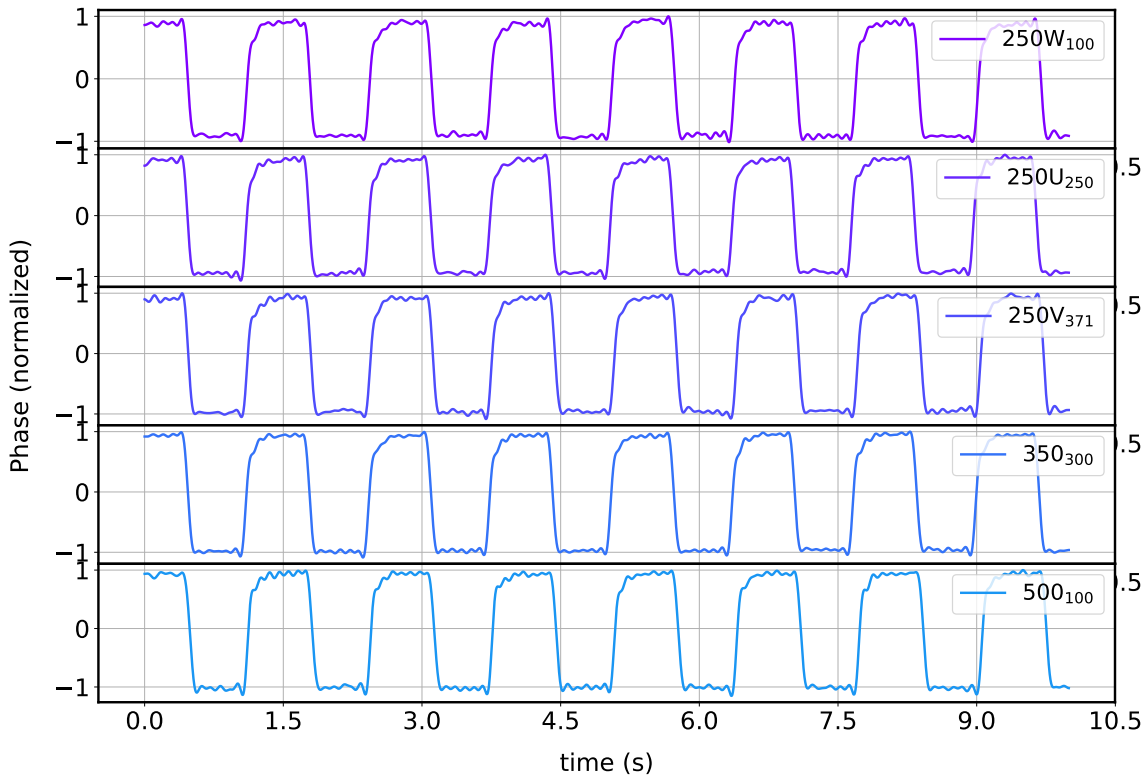

Figure 77. Calibration lamp chops for the 250U, 250W, 250V, 350 and 500 µm arrays, taken at CSBF.

Figure 79 shows a chop for a single channel of the 350 µm array which was produced by chopping a thermal source in front of the cryostat window. The thermal source was at temperature of ~15 K above room temperature. During this measurement, which took place in May, 2018, a 2.85% NDF was installed in the optical chain. The top panel of Figure 79 shows the phase quadrature of Δf, and the bottom shows

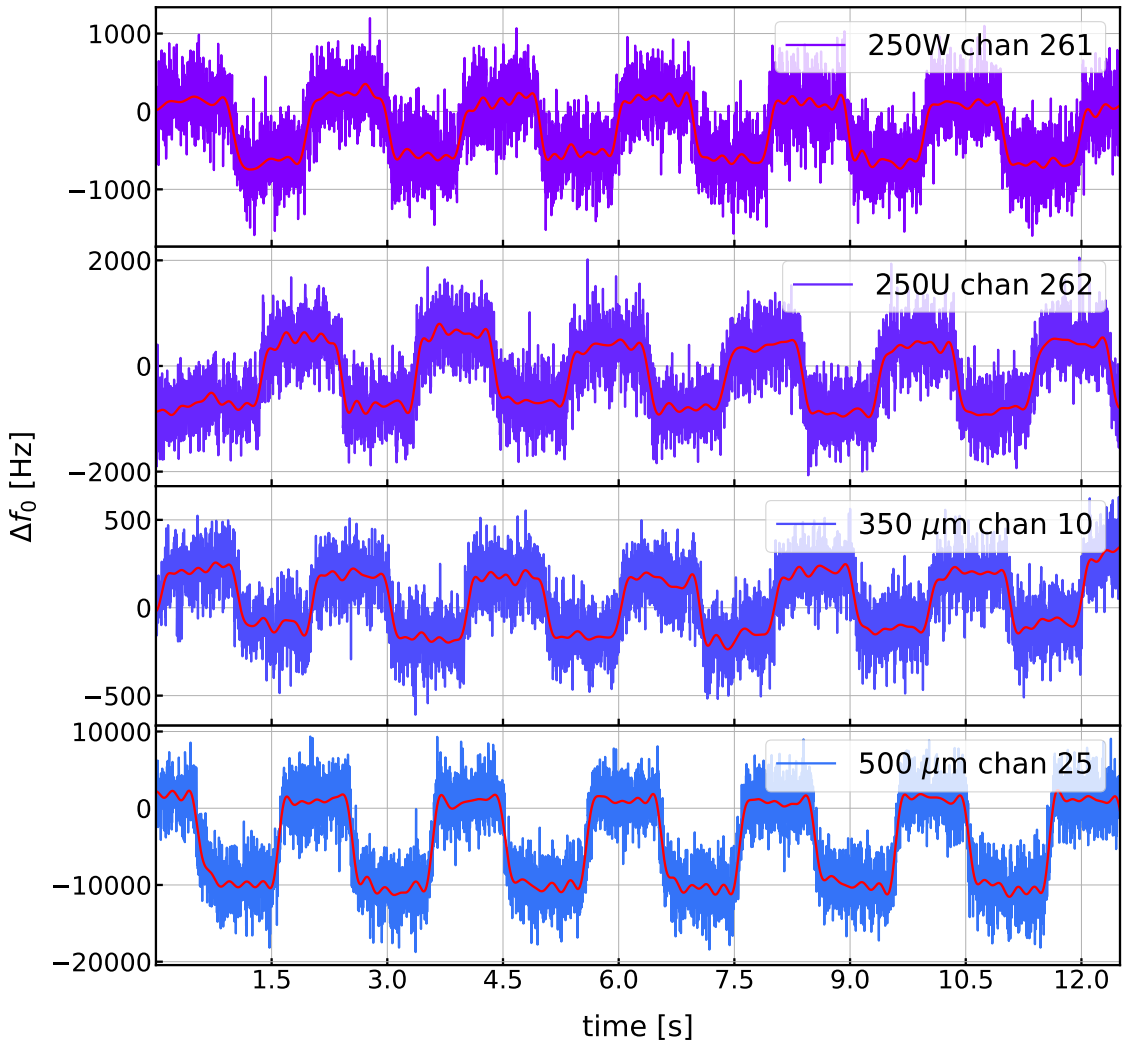

Figure 78. Calibration lamp chops for the 250U, 250W, 350 and 500 μm arrays, taken at LDB.

the dissipation quadrature. If we assume a value for either R_P or $\epsilon_{\rm opt}$, we can estimate either an NEP or an NET for this channel. This is done in Section 4.4.

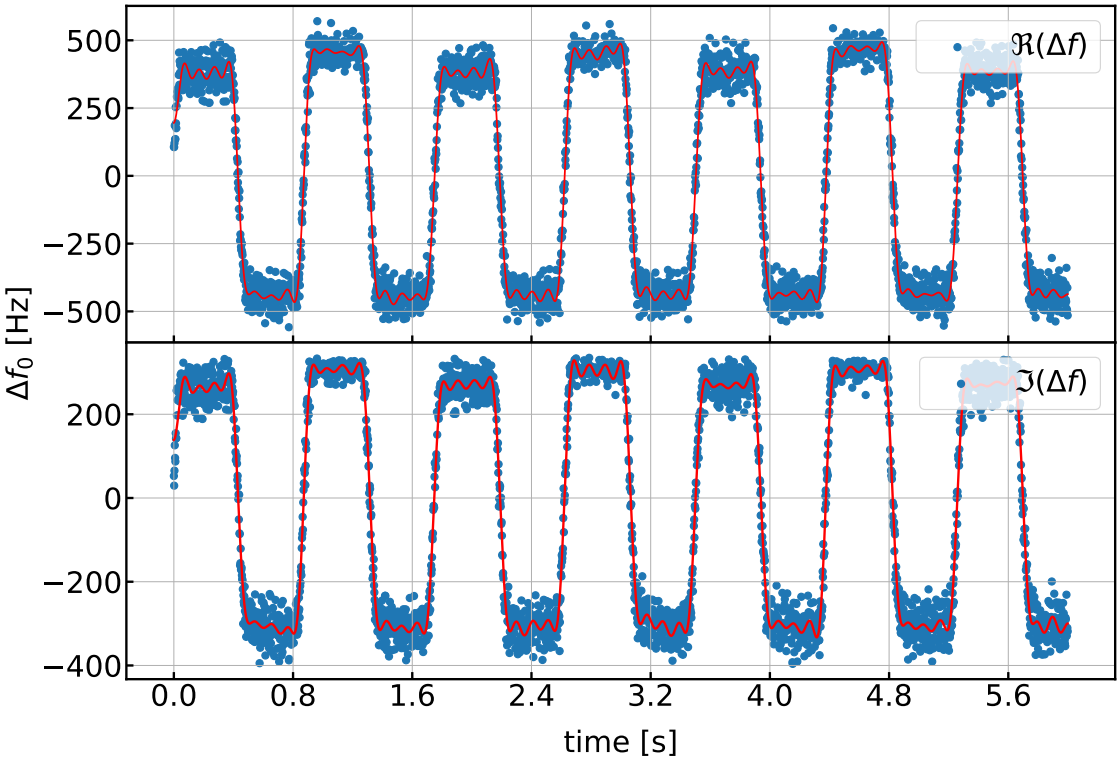

Figure 79. Chops made using a ~15 K thermal source for a single channel of the 350 μm array.

4.3.2 Base Temperature Responsivity

At CSBF, R_T was measured by taking VNA sweeps of S_{21} at different base temperatures. The temperatures used in this calculation are 277, 282 and 286 mK. The df/dT from the sweep was checked against values predicted by the Mattis-Bardeen theory. Figure 80 shows S_{21} traces for the three base temperatures are shown for

a single 350 µm resonator. The R_T predicted by Mattis-Bardeen are shown in Figure 80.

Band	R_P (df/f_0)/pW	R_T (df/f_0)/mK
250 µm	1.7×10^{-5}	9.15×10^{-7}
350 µm	5.2×10^{-6}	1.51×10^{-6}
500 µm	1.6×10^{-4}	3.91×10^{-6}

Table 14. BLAST-TNG optical and temperature responsivities (optical responsivities provided by NIST).

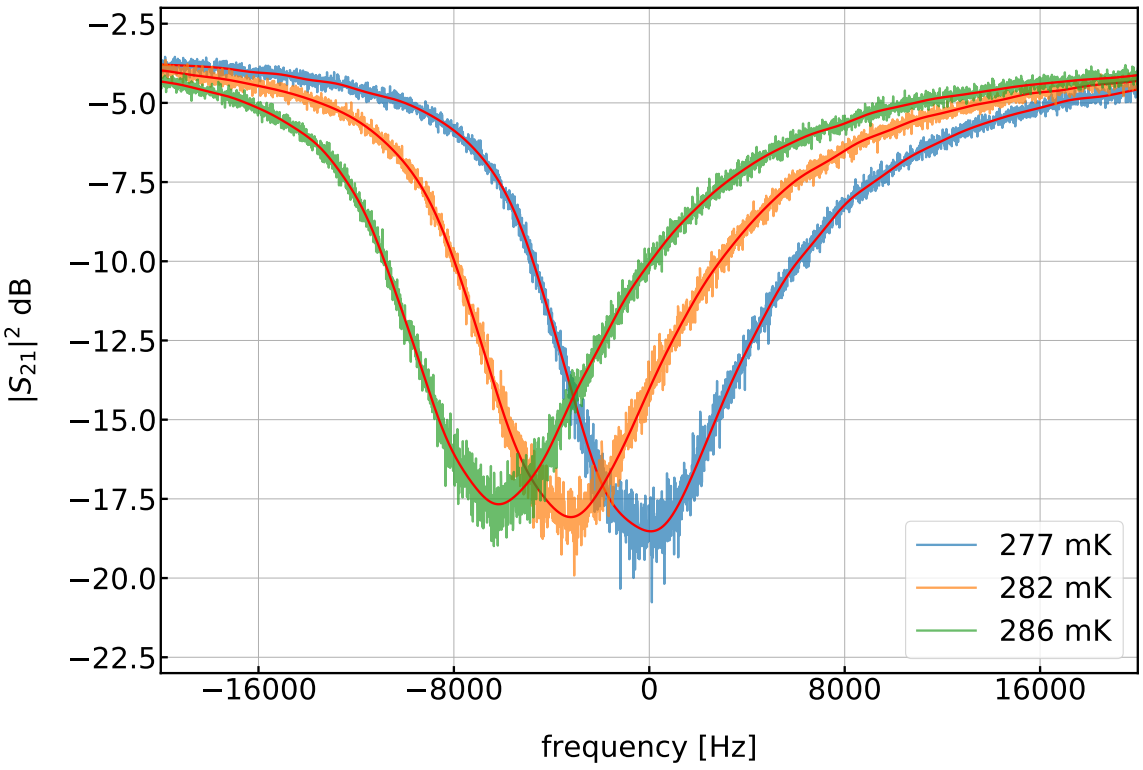

Figure 80. 350 µm VNA $|S_{21}|$ sweeps at different FPA temperatures. The red trace is low-pass filtered data (Figure courtesy of Adrian Sinclair)

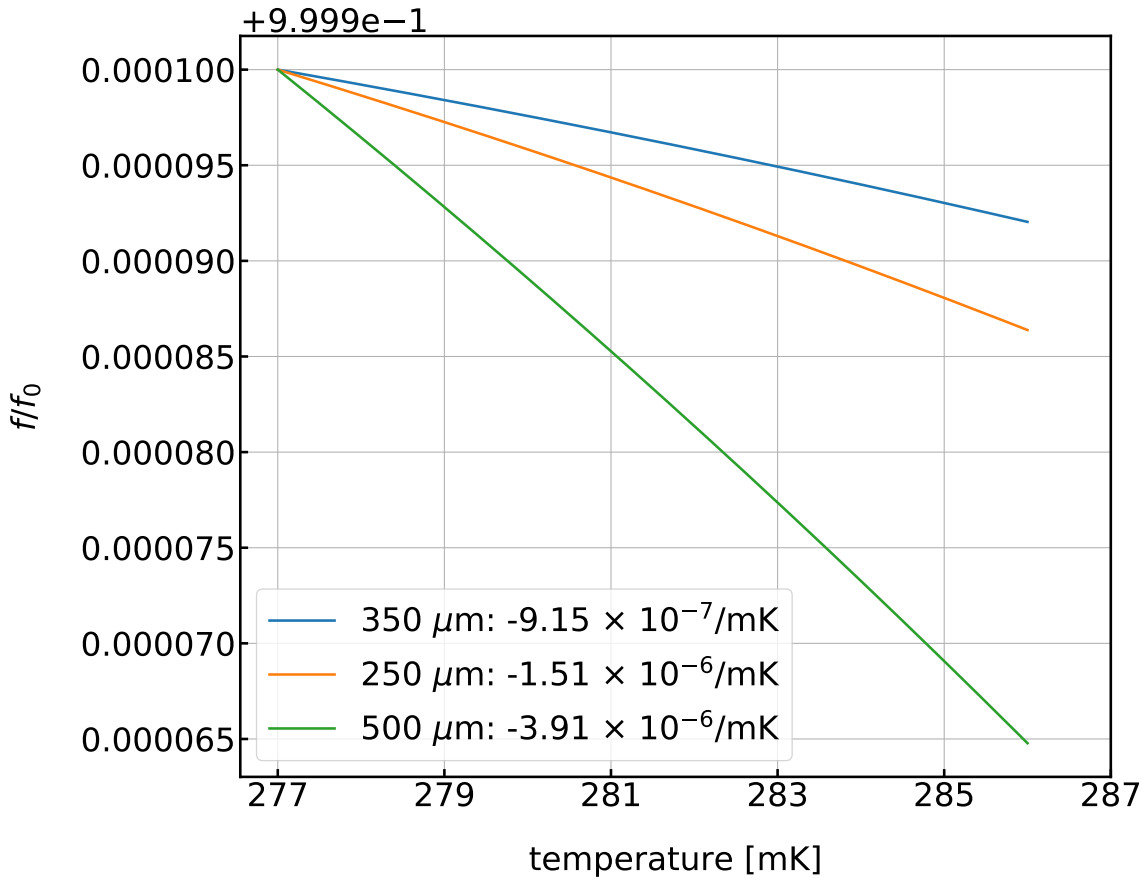

Figure 81. Mattis-Bardeen temperature responsivity curves for the 250, 350 and 500 μm arrays (Figure courtesy of Adrian Sinclair).

4.4 Estimating Sensitivity

The sensitivity of each detector is characterized by its NET, in units of $K/\sqrt{Hz}$, or NEP, in units of $W/\sqrt{Hz}$. These quantities are defined in Chapter 2. In this section, we describe how they can be estimated using the quantities introduced in previous sections.

4.4.1 NET and NEP

NET (or NEP) are estimated by modulating the detectors with a thermal source whose temperature relative to the background is known with some certainty. An I/Q timestream of the chop is recorded for each readout channel. Without doing a frequency sweep (e.g., target or wide sweep), the NET can be calculated from the ΔT and I/Q timestreams as:

$$\begin{aligned} NET &= \Delta T \frac{e_V}{V_{pp}} \\ &= \Delta T \sqrt{\frac{(\sigma_I^2 + \sigma_Q^2)/(f_s/2)}{I^2 + Q^2}} \qquad \left[\frac{\mathrm{K}}{\sqrt{\mathrm{Hz}}}\right] \end{aligned} \tag{4.5}$$

If the I/Q timestreams are converted into units of Δf, then the NET can be calculated as:

$$NET = \Delta T \frac{e_f}{\Delta f / f_0} \qquad \left[\frac{\mathrm{K}}{\sqrt{\mathrm{Hz}}}\right] \tag{4.6}$$

where Δf is the size of the frequency chop, in units of Hz. The frequency noise e_f can be calculated using either the I/Q gradient method, the phase fitting method, or by fitting the resonator parameters from the sweep, as in Equation 4.4.

If the optical responsivity is known or assumed, the NEP can be calculated as:

$$\begin{aligned} NEP &= \frac{e_f/f_0}{R_P} \\ &= \frac{e_f/f_0}{(df/f_0)/dP} \qquad \left[\frac{\mathrm{W}}{\sqrt{\mathrm{Hz}}}\right] \end{aligned} \tag{4.7}$$

Because e_f is complex, Equations 4.6 and 4.7 allow for the NET and NEP in both the phase and dissipation quadratures to be calculated.

The thermal source chop in Figure 79 was produced with $\Delta T \simeq 15$ K. At the time of the measurement there was a 2.85% NDF inside the cryostat. Using Equation 4.6,

and multiplying by the loss from the NDF gives $NET_{\rm tot} = 825$ μK/$\sqrt{\rm Hz}$. Assuming a 30% passband width, this value corresponds to $NEP_{\rm freq,tot} = 5.55 \times 10^{-16}$ W/$\sqrt{\rm Hz}$. Assuming the R_P value for the 350 μm array in Table 14, the optical efficiency for this channel is found to ∼20%. For this particular channel, the frequency noise in the dissipation quadrature was found to be slightly less than half of the noise in the phase quadrature. This indicates that the channel is close to being photon noise limited ($NEP_{\rm phot} > NEP_{\rm diss}$, where $NEP_{\rm diss}$ is calculated using the noise in the dissipation quadrature and the frequency responsivity). Values for the NET and NEP in the two quadratures are listed in Table 15.

Table 15. Estimates of NET and NEP for a single 350 μm channel. This data was recorded in May, 2018. There was a 2.85% NDF in the optical path.

	Channel 618
f_0 MHz	668.44
$S_{\rm XX}$ Hz^{-1}	8.34×10^{-18}
$S_{\rm YY}$ Hz^{-1}	3.99×10^{-18}
NEP$_{\rm freq}$ W/$\sqrt{\rm Hz}$	5.55×10^{-16}
NEP$_{\rm diss}$ W/$\sqrt{\rm Hz}$	3.84×10^{-16}
NET$_{\rm freq}$ μK/$\sqrt{\rm Hz}$	825
NET$_{\rm diss}$ μK/$\sqrt{\rm Hz}$	571
NDF	0.0285
$\epsilon_{\rm opt}$ (estimated)	0.19

4.4.2 NEFD and Mapping Speed

To relate NEP to mapping speed, it is useful to convert it into a Noise-equivalent flux density (NEFD). For sub-mm astronomical instruments, NEFD is typically reported in units of MJy/$\sqrt{\rm Hz}$ or MJy$\sqrt{\rm s}$ (1 Jy = 10^{-26} W/m^{2}Hz). Here, we convert

the median eletrical NEP values measured for each BLAST-TNG detector array (see Section 4.2.3) into NEFD, and then into a per-beam RMS noise level for each waveband, $\sigma_{\rm map}$, in units of mJy/beam. The conversion from NEP in W/$\sqrt{\rm Hz}$ to NEFD in MJy$\sqrt{\rm s}$ is done as:

$$NEFD = \sqrt{0.5}10^{-20}\frac{NEP}{A_{\rm beam}B_{\rm opt}} \qquad \left[\frac{\rm MJy\sqrt{s}}{\rm sr}\right] \tag{4.8}$$

where the factor of $\sqrt{0.5}$ accounts for the conversion from $1/\sqrt{\rm Hz}$ to $\sqrt{\rm s}$, $A_{\rm beam}$ is the beam area in steradians and $B_{\rm opt}$ is the optical bandwidth. Table 16 compares the NEPs and NEFD measured in Palestine for each array to the total (phase quadrature) and photon NEPs listed in the original BLAST-TNG proposal. $I_{\rm min}$ is the source intensity which would be needed to measure 1-σ error bars on the polarization fraction of 0.5% over a 1 deg^2 area after a 5 hour integration (this quantity is discussed in Chapter 5). Values for the map noise $\sigma_{\rm map}$ is shown per-detector, as well as accounting for the total number of detectors in each band. The NEFD which is given accounts for the total number of detectors.

Table 16. NEP and NEFD estimates estimated from Palestine tests, compared to proposal values.

	250 μm	350 μm	500 μm
NEP$_{\rm freq,prop}$[1] W/$\sqrt{\rm Hz}$	6.5×10^{-17}	5.5×10^{-17}	4.7×10^{-17}
NEP$_{\rm phot,prop}$ W/$\sqrt{\rm Hz}$	17×10^{-17}	12×10^{-17}	8.7×10^{-17}
NEP$_{\rm freq,Pal}$ W/$\sqrt{\rm Hz}$	35×10^{-17}	35×10^{-17}	10.0×10^{-17}
NEFD$_{\rm PAL}$ $\frac{MJy\sqrt{s}}{sr}$	0.13	0.12	0.034
$\sigma_{\rm map}$ mJy/beam per detector	68.75	96.25	39.28
$\sigma_{\rm map}$ mJy/beam	2.17	3.77	2.27
$I_{\rm min}$ MJy/sr	78.16	49.46	10.19

[1] Values from the original BLAST-TNG proposal.

The measured electrical NEP values from Palestine are slightly higher than the estimated values in the proposal, but comparable to those given for the photon noise NEPs in the proposal. During the measurement when the electrical NEPs were calculated, it is possible that many of the detectors were photon noise limited. However, they were not optimally biased in terms of readout power, which causes the NEP in the dissipation direction to contribute more to the total NEP. The values for $\mathrm{NEP}_{\mathrm{freq}}$ given in the proposal assume that the dissipation quadrature contributes ∼33% to the total. In the Palestine measurements, the dissipation quadrature constitutes ∼65% of the total.

Despite the slightly higher NEPs, the NEFDs are lower than in BLAST 2006 and the corresponding mapping speed is an order of magnitude higher (Marsden et al., 2009). This bodes well for the planned 2019/2020 flight of BLAST-TNG.

4.5 Applying the LEKID Model to Measured Data

In this section we apply the parametric LEKID model which was developed in Chapter 2 to measured data from the 2018–2019 Antarctic campaign to determine whether or not it can provide reasonable estimates of absorbed optical power and optical efficiency for a single resonator. During pre-flight preparations at LDB, no NDF was installed inside the receiver, and the detectors were exposed to the ambient ∼290 K loading from the highbay. Despite their being heavily loaded, the majority of the detectors still responded to calibration lamp pulses.

At the end of the ice campaign, data was taken with the aluminum camera shutter in the open and closed states. The shutter blocks the optical path between the primary mirror and the cryostat window, and is used during ascent when there is a

risk of exposing the detectors to direct sunlight due to uncontrolled pointing. The WiFi and star-camera video transmitter were powered down during this test, since they had been previously been discovered to interfere with the detector timestreams.

As a test of the model, we choose a single channel from the 350 μm array. The model takes $P_{\rm ro}$, $T_{\rm base}$, $P_{\rm opt}$, f_0, Q_c, $Q_{\rm loss}$ and ϵ_{assym} as inputs. To simulate the test channel, we begin by estimating these and other resonator parameters using the target sweep and I/Q timestreams taken in the shutter-closed and shutter-open states. The estimated parameters, listed in Table 17, include: f_0, Q_r, Q_c, dip-depth, $\epsilon_{\rm a}$, $S_{\rm XX}$, $NEP_{\rm freq}$ and $NEP_{\rm diss}$. For $T_{\rm base}$ we assume 287 mK, which was the temperature of the 350 μm focal-plane-array at the time that the data was recorded. For $P_{\rm opt}$ we use the range between 1–200 pW, divided into 500 intervals. The readout power $P_{\rm ro}$ is -83 dBm.

4.5.1 Optical Response

Given the inputs listed above, the model returns $S_{21}(P_{\rm opt})$ for each optical power in the input range. A fit of S_{21} to both the shutter-closed and shutter-open data is shown in Figure 82. The measured data is shown as blue points, the initial fit is shown as a dashed black line, and the model fit is shown as a dashed red line. To produce the model traces shown in Figure 82, we select the two $S_{21}(P_{\rm opt})$ arrays which correspond to the $f_{0,\rm closed}$ and $f_{0,\rm open}$ values determined from the initial fit. To create a good fit to S_{21}, we find that we must add a frequency offset to the value of f_0 which is input to the model. For this test, we find that a frequency offset of 110 kHz produces good results. The resonator quality factors produced by the parametric model (listed in Table 17) match those of the initial fit to within a few percent.

The frequency shift between the two resonator states is -95.47 kHz. Using this shift, the model can be used to estimate the optical responsivity R_P. The responsivity is sensitive to the value which is chosen for N_0, the single-spin density of states at the Fermi Energy. Using $N_0 = 10^{10}$ eV^{-1} μm^3, the model reports R_P = 5.35 pW^{-1}. This value is close to the average R_P reported by NIST, which is listed in Table 14. Given the frequency shift and responsivity, the power absorbed by the resonator between the two states is $\Delta P_{\rm abs} \approx 22$ pW.

Table 17. 350 μm model parameters.

	Shutter closed (est. from data)	Shutter open (est. from data)	Shutter closed (model)	Shutter open (model)
f_0 (MHz)	813.330	813.235	813.330	813.235
Q_r	15,829	12,106	15,600	11,724
Q_c	25,328	24216		
Q_i	42,316	24,246	40,615	21,829
depth (dB)	-8.42	-5.77	-8.31	-5.40
$\epsilon_{\rm a}$	-0.20	-0.28		
$R_{P,\rm model}$ (1/pW)	5.35×10^{-6}			
NEP$_{\rm freq}$ W/$\sqrt{\rm Hz}$	4.87×10^{-16}	7.11×10^{-16}	4.88×10^{-16}	7.11×10^{-16}
NEP$_{\rm diss}$ W/$\sqrt{\rm Hz}$	4.06×10^{-16}	6.40×10^{-16}		
NEP$_{\rm phot}$ W/$\sqrt{\rm Hz}$			3.99×10^{-16}	5.33×10^{-16}
NEP$_{\rm amp}$ W/$\sqrt{\rm Hz}$			2.34×10^{-16}	4.32×10^{-16}
$S_{\rm XX}$ (1/Hz)	6.80×10^{-18}	1.45×10^{-17}	6.83×10^{-18}	1.45×10^{-17}
S_{YY} (1/Hz)	4.72×10^{-18}	1.17×10^{-17}		
$\Delta P_{\rm abs}$ (pW)		22		22 (58)
$P_{\rm abs}$(pW)			103	161
T$_{\rm shutter}$ (K)				250
$\epsilon_{\rm opt}$				0.28
$Q_{\rm loss}$	3×10^7			
N_0 eV^{-1} μm^3	9.5×10^{10}			

Next, we compare the measured values for fractional frequency noise $S_{\rm XX}$, and NEP, to the values produced by the model. Because of the frequency offset which

was applied to the f_0 values in order to obtain a good S_{21} fit, we do not use the values for S_{XX} and NEP which correspond to those traces. Instead, we locate the two $S_{21}(P_{\mathrm{opt}})$ traces that have a fractional frequency noise S_{XX} (1/Hz) which is closest to the values of S_{XX} which are estimated from the data. The two S_{21} traces which are found to match the closed and open states have f_0 values which are offset by those of the original S_{21} traces by -413 and -566 kHz, respectively. At the time of writing, the source of this frequency offset is not understood. Apart from the frequency offset, the frequency noise and NEP values produced by the model are close matches to the measured values.

Table 17 lists the measured values for S_{XX}, S_{YY}, $\mathrm{NEP}_{\mathrm{freq}}$ and $\mathrm{NEP}_{\mathrm{diss}}$ in the shutter-open and shutter-closed states. Figure 84 shows the fractional frequency power spectral density for the measured S_{XX} and S_{YY}. The black line indicates the white noise level which corresponds to their difference. Because the measured value for S_{XX} was used to select the appropriate S_{21} array, the measured and model values for S_{XX} and $\mathrm{NEP}_{\mathrm{freq}}$ values in Table 17 agree to within a fraction of a percent.

The model S_{21} traces with total frequency noise that matches the measured noise correspond to optical powers of 161 and 103 pW, for the open and closed states. The ratio of these powers is 1.55, which is close to the measured ratio of $\mathrm{NEP}_{\mathrm{freq}}$ between the closed and open states (1.50). Figure 83 shows the predicted ratio of $\mathrm{NEP}_{\mathrm{phot}}$ to $\mathrm{NEP}_{\mathrm{amp}}$ as a function of optical power. At the optical powers corresponding to the open and closed states, these ratios equal 1.23 and 1.7, which indicate that the detector is photon noise limited. However, the value of the optical power given by the model does not necessarily equate to the true optical power (despite being close to what would be expected for an optical efficiency of $\sim$20%.)

To estimate the system optical efficiency ϵ_{opt}, we start by assuming that with the

shutter open and without an NDF inside the cryostat, the detectors see an optical load of ∼290 K. Assuming a detector efficiency $\eta_{\rm det}$ of 80%, a horn efficiency $\gamma_{\rm horn}$ of 70% and a 30% optical passband centered on 350 μm, the total power incident on the detectors is:

$$P_{\rm opt} \simeq kTB_{\rm opt}\eta_{\rm det}\eta_{\rm horn} \tag{4.9}$$

Assuming the above values, $P_{\rm opt} \simeq 576$ pW. Assuming the model value of $\Delta P_{\rm abs} = 161$ pW which is based on the measured frequency noise, $\epsilon_{\rm opt} \sim 28\%$. Given the design target value of 30%, this estimate is reasonable. However, there is a discrepancy between the model's prediction of $\Delta P_{\rm abs} \approx 22$ pW and the difference in absorbed power between the shutter-open and closed states of ∼58 pW which is based on the measured frequency noise. It is unclear whether this discrepancy originates in the model or in the measured data. Because the model responsivity is very close to the responsivity reported by NIST, for the purposes of this test we assume that the discrepancy originates in the model, and disregard the value it reports for $P_{\rm abs}$ in the closed state. Assuming the model values for $\Delta P_{\rm abs}$ and $\epsilon_{\rm opt}$, the shutter temperature is estimated as ∼250 K.

4.5.2 Temperature Response

The parametric model can also be used to estimate the responsivity to changes in base-temperature, R_T. After recording data with the shutter open and closed, it was left closed while the base-temperature of the cryostat was raised. Wide and target

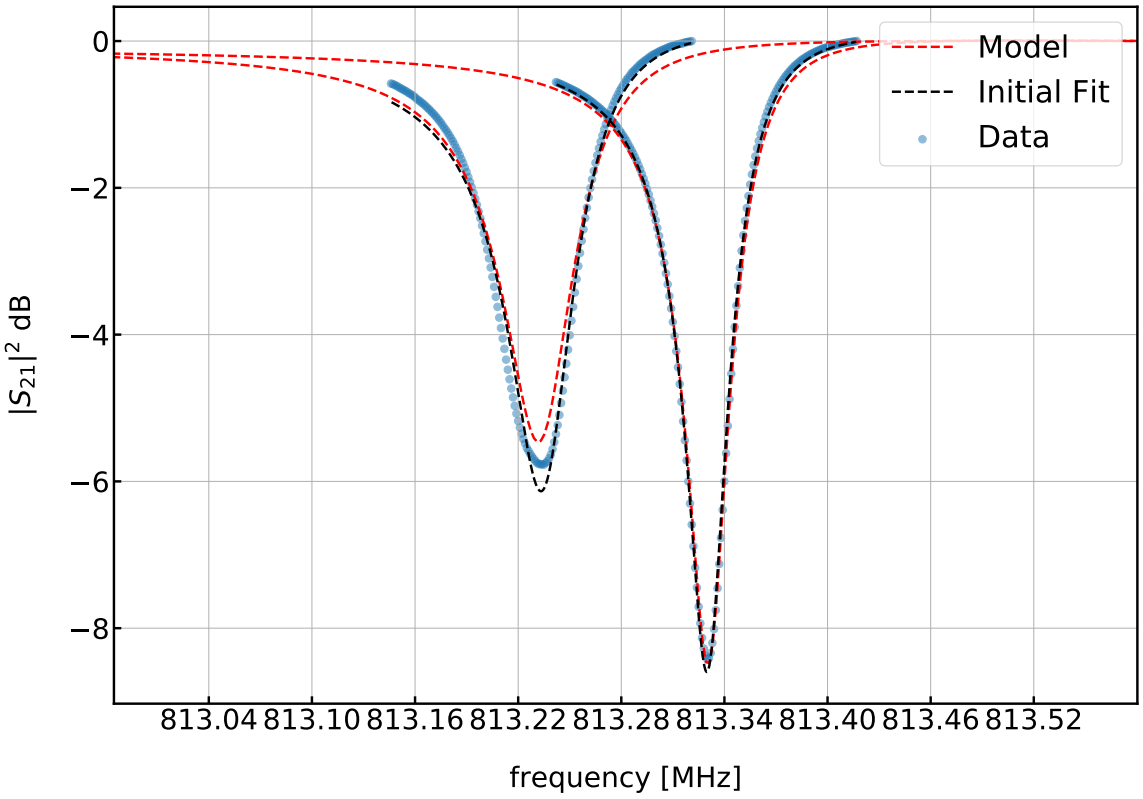

Figure 82. A fit to S_{21} for an optically shifted 350 µm channel using the parametric LEKID model presented in Chapter 2 (red). The blue points are measured data.

sweeps were taken at $T_{\text{base}} \approx 287$ mK and 316 mK. The model fit to $S_{21}(T_{\text{base}})$ at both temperatures is shown in Figure 85. As with the optical responsivity model fit, we use S_{XX} to estimate the T_{base}, and the frequency shift between states to estimate the responsivity. The results of the fit are listed in Table 18.

The fits to the base-temperatures are within 4 K of the temperatures which were measured using the cryogenic thermometry. The R_T value measured from the frequency shift of 32 kHz is 1.4×10^{-7}. This value is lower than the R_T of 9.15×10^{-7} which was measured during the Palestine integration, but the difference is reasonable given that we are comparing two different channels under different conditions.

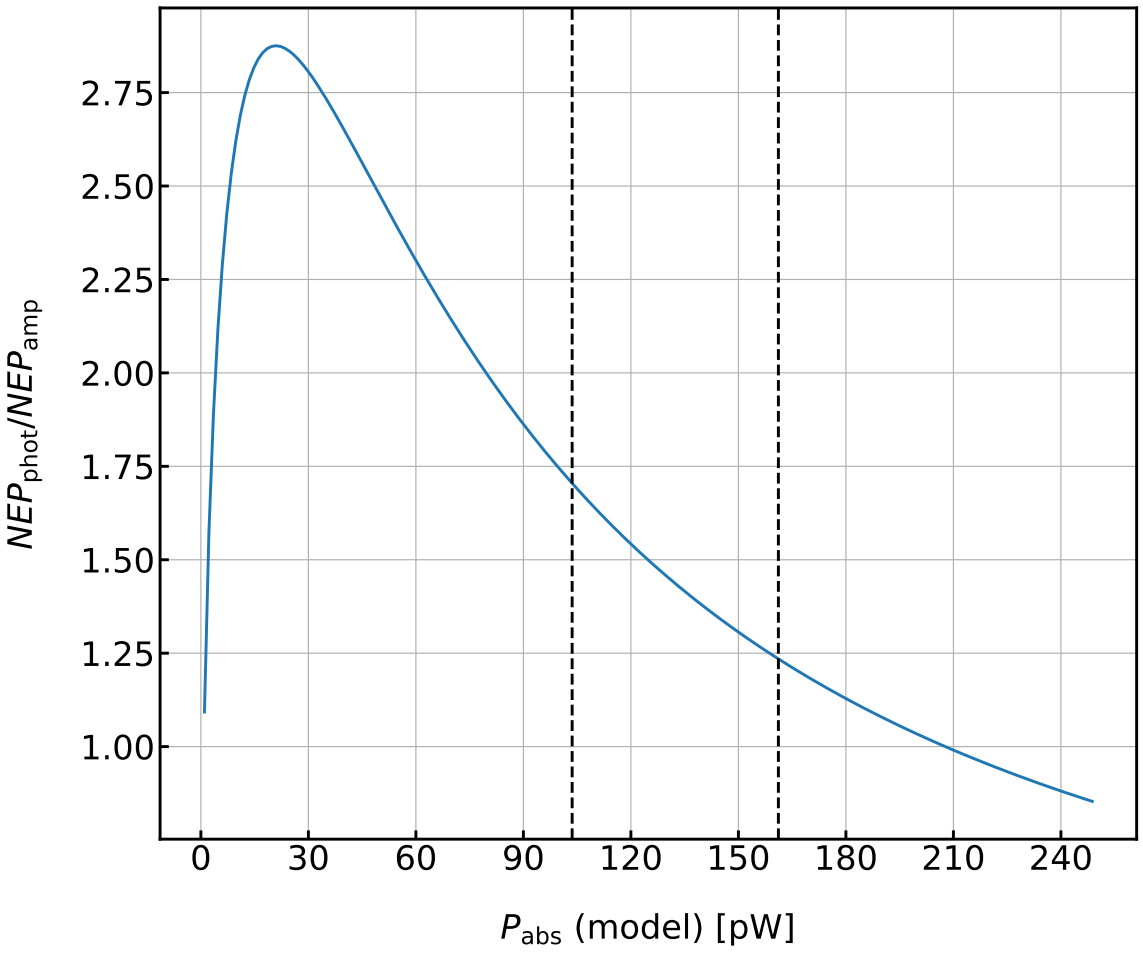

Figure 83. The ratio of $NEP_{\rm phot}$ to $NEP_{\rm amp}$ for a 350 μm channel.

$f_{0,\rm cold}$ (MHz)	$f_{0,\rm warm}$ (MHz)	$T_{\rm warm}$ (meas,mod) mK	$T_{\rm cold}$ (meas,mod) mK	$(df/f_0)/dT$ meas. ($\rm K^{-1}$)	$(df/f_0)/dT$ mod. ($\rm K^{-1}$)
813.328	813.298	316, 311	287, 283	1.35×10^{-7}	1.40×10^{-7}

Table 18. Model and measured parameters for the 350 μm base temperature response.

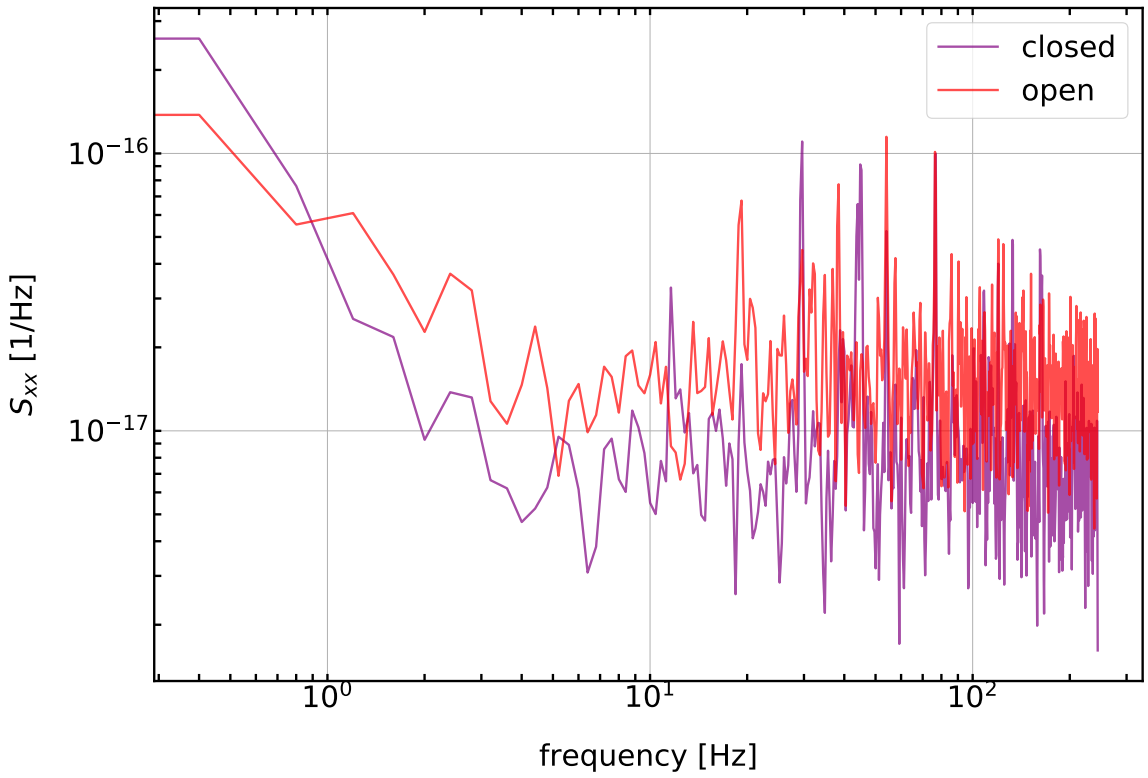

Figure 84. The measured fractional frequency noise S_{XX} for the 350 μm channel, with the camera shutter open and closed.

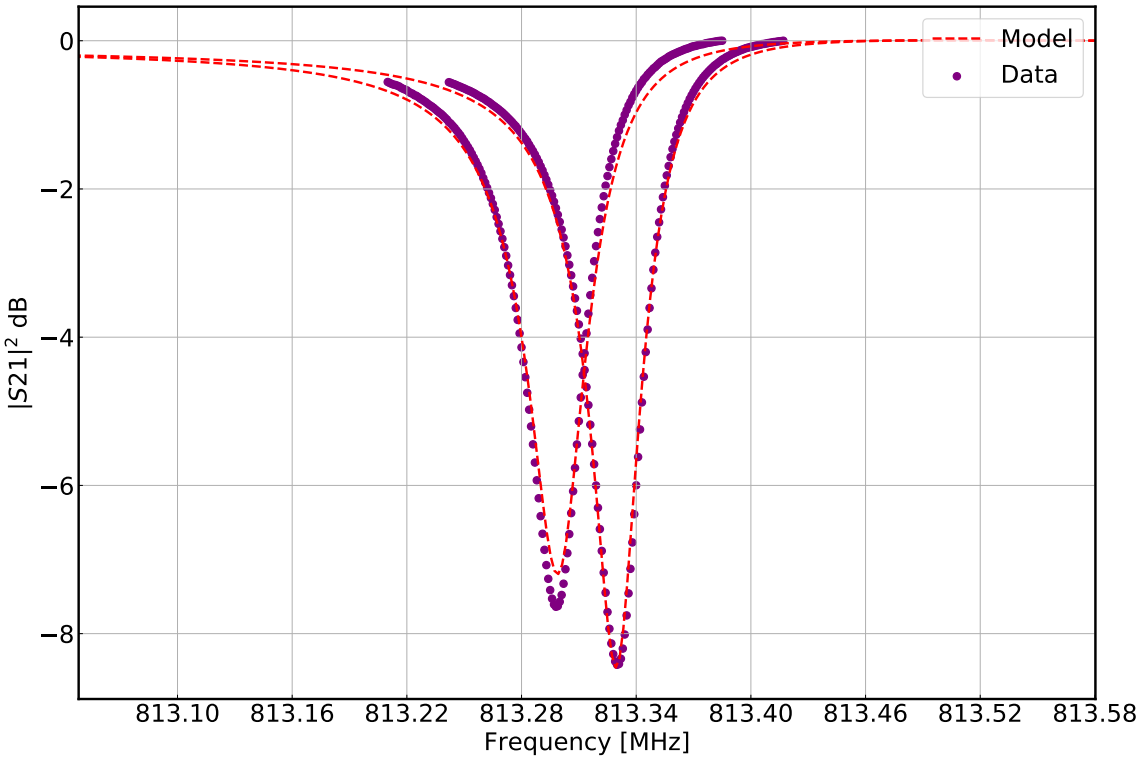

Figure 85. A fit to $S_{21}(T_{\mathrm{base}})$ for a temperature shifted 350 μm channel using the parametric LEKID model presented in Chapter 2 (red). The blue points are measured data.

4.6 Optical Tests

The following sections describe the passband mapping and polarization efficiency tests that were performed in the lead-up to the planned 2018/2019 Antarctic flight of BLAST-TNG.

4.6.1 Passband Mapping

In this section we describe the characterization of the camera's spectral passbands. The low-frequency edge of each BLAST-TNG waveband is defined by the horn array waveguide cutoff frequencies, and the high frequency edges are defined by low-pass edge (LPE) and dichroic filters. Ideally, the passbands should be three boxcar filters centered at 250 μm, 350 μm and 500 μm. Their widths should be ∼30% of each center frequency (360 GHz, 257 GHz and 180 GHz). In practice, the filter shapes will not be rectangular. Their shape will be the product of the filter responses of each component in the optical path of the instrument.

A Fourier-transform spectrometer (FTS) containing a water-cooled mercury arc lamp was used to produce in-band emission over the three wavebands. A low-pass filter inside the FTS prevents high frequency photons from exiting the output aperture. During the measurement, the FTS aperture was positioned a few inches in front of the cryostat window, to ensure that the in-band source filled the entire beam. The setup in the CSBF highbay is shown in Figure 86.

A FTS is a modified Michelson interferometer. A single beam of light is split into two paths by a beamsplitter. The two beams then reflect off of mirrors which send them back through the beamsplitter to recombine. In an FTS, one of the mirrors

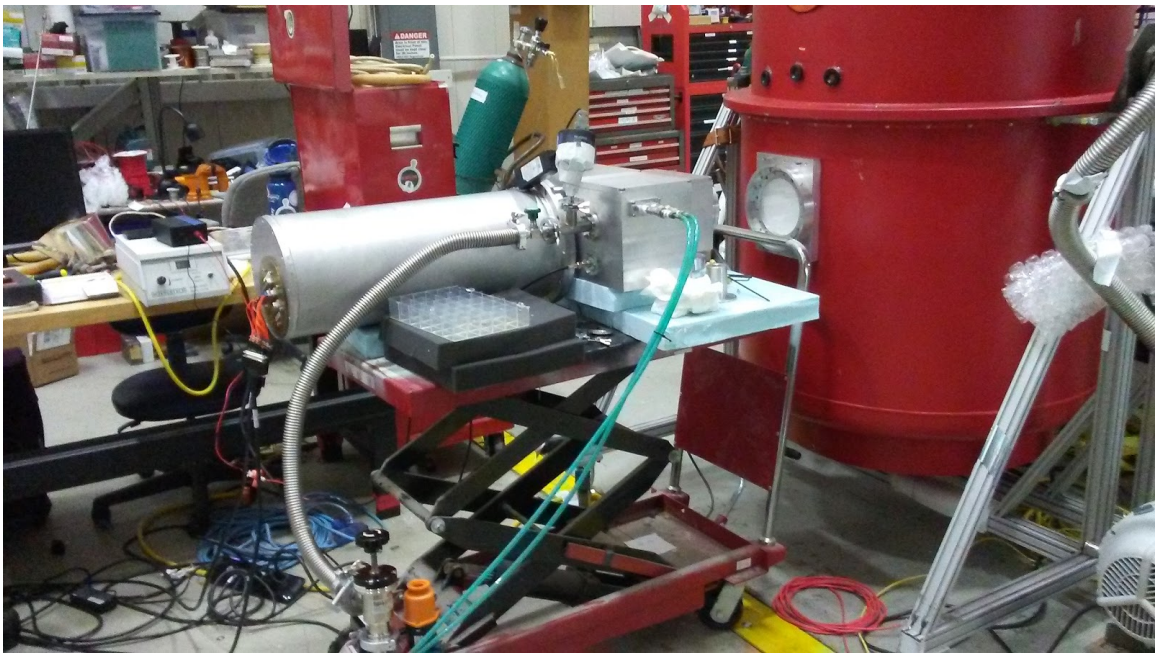

Figure 86. The FTS measurement setup at CSBF.

can be translated back and forth through a distance ΔL to create a variable optical path-length delay (OPD) between the two optical paths. At the output of the FTS, the intensity at optical frequency ν for a given OPD is:

$$I(\nu, \Delta L) = I(\nu, 0)\left[1 + \cos(2\pi\nu(2\Delta L))\right] \tag{4.10}$$

where $I(\nu)_{\mathrm{ZPD}}$ is the intensity corresponding to the zero-path length difference (ZPD). Summing the intensities over all frequencies within the waveband gives:

$$I(\Delta L) = \int_{\nu_0}^{\nu_1} I(\nu, \Delta L) d\nu \tag{4.11}$$

Equation 4.11 is the Fourier cosine transform (real Fourier-transform). As the movable mirror inside the FTS is stepped through ΔL, the recorded timestream for each channel is an interference pattern called an interferogram. Each point in the interferogram corresponds to Equation 4.11 for a given path-length difference ΔL, and

therefore contains information about the intensity of every optical frequency in the waveband. The desired passband spectrum is produced by taking the inverse Fourier-transform of Equation 4.11, $\mathcal{F}^{-1}(I(\Delta L))$, and then shifting the Fourier frequencies into the optical band (see Equation 4.13). The resulting frequency resolution of the spectrum $\nu_{\rm res}$ depends only on the OPD: $\nu_{\rm res} = c/2\Delta L$.

The FTS measurement parameters are shown in Table 19.

Table 19. FTS measurement parameters.

Parameter	Value
ΔL (mm)	100
$\rm v_{mirror}$ (mm/s)	0.1
$\rm f_s$ (Hz)	488.28125
$\rm f_c$[1](Hz)	3
$\rm N_{samp}$	4×10^4
$\nu_{\rm res}$ (GHz)	1.5
$\rm f_{scale}$ (GHz)	1.22×10^{14}

[1] cutoff frequency of the low-pass filter used to process interferograms.

In the following, we describe each step of the FTS data reduction for the 250W, 350 μm and 500 μm detector arrays. The measurement lasted ∼5.5 hours, during which ∼40 min of data was recorded for each BLAST-TNG detector with the FTS positioned in five different orientations relative to the cryostat window– centered, above-center, below-center, right-of-center and left-of-center. These positions were used to maximize the likelihood that the beam for each detector channel would be well illuminated in at least one of the positions. Figure 87 shows a phase timestream of the entire measurement for a single channel of the 350 μm array. Due to slow drifts in the fridge temperature, the readout tones had to be recalibrated three times

during the measurement (visible as large spikes). The features between ∼1–1.1 × 10^7 samples correspond to the cryogenic pumped-pot running out of He4, and then refilling.

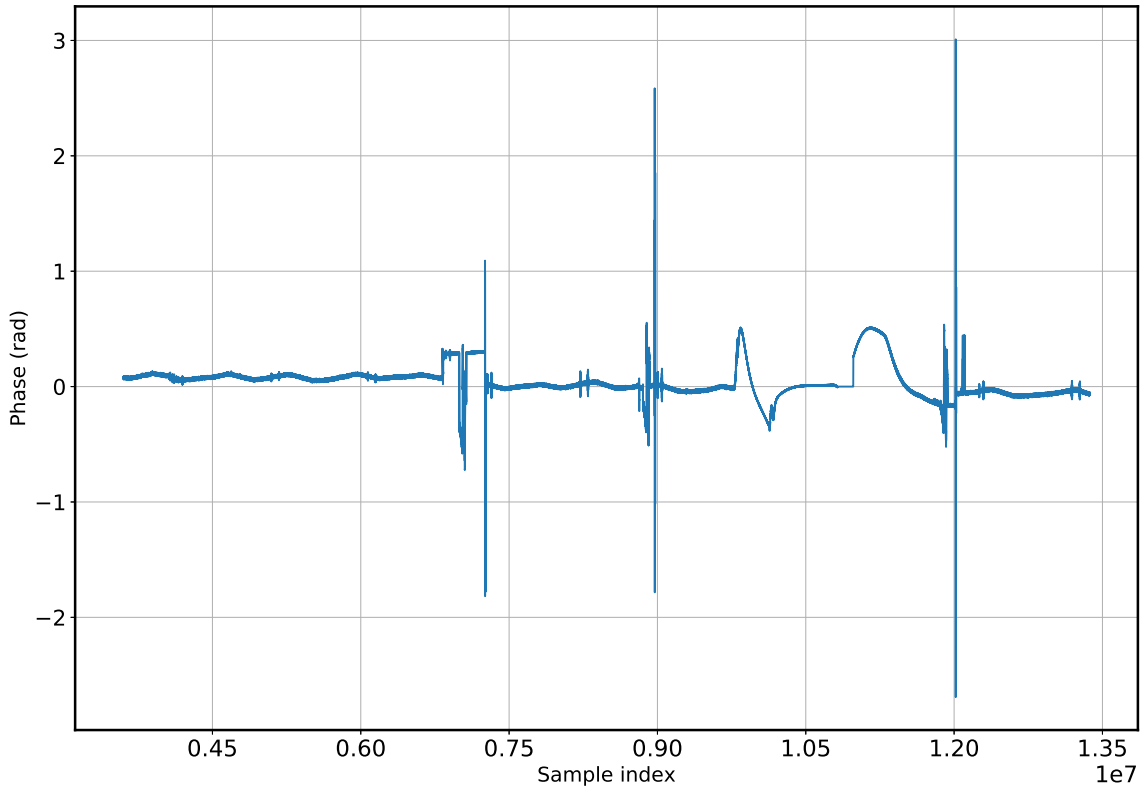

Figure 87. The timestream of the FTS measurement, for a single channel of the 350 μm array. The measurement duration is ∼5.5 hours.

The data reduction steps used to create the passband spectrum spectrum for each band (shown in Figure 91) are:

1. For each channel, identify high SNR interferograms I_c in the total measurement timestream.
2. Align the interferograms according to the location of $I_c(0)$.
3. Normalize each interferogram according to the peak-to-peak amplitude of $I(\nu)_{\mathrm{ZPD}}$ and apply a low-pass filter.

4. Correct each I_c for nonlinearities introduced by the detectors.
5. Co-add each I_c to create an interferogram template I_T.
6. Take the inverse Fourier-transform of the interferogram template and shift the Fourier frequencies to the optical band.

To choose the interferograms to be used in the analysis (1), the approximate locations of $I_c(0)$ for ~10 interferograms were chosen from a high SNR measurement timestream for each band. The exact center locations of each $I_c(0)$ were then calculated to facilitate their alignment in time in (2). The indices corresponding to the locations of each $I_c(0)$ were then used to select the interferograms from each of the other channels. Two selected interferograms from a raw timestream are shown in Figure 88, with their approximate center positions marked with red dots.

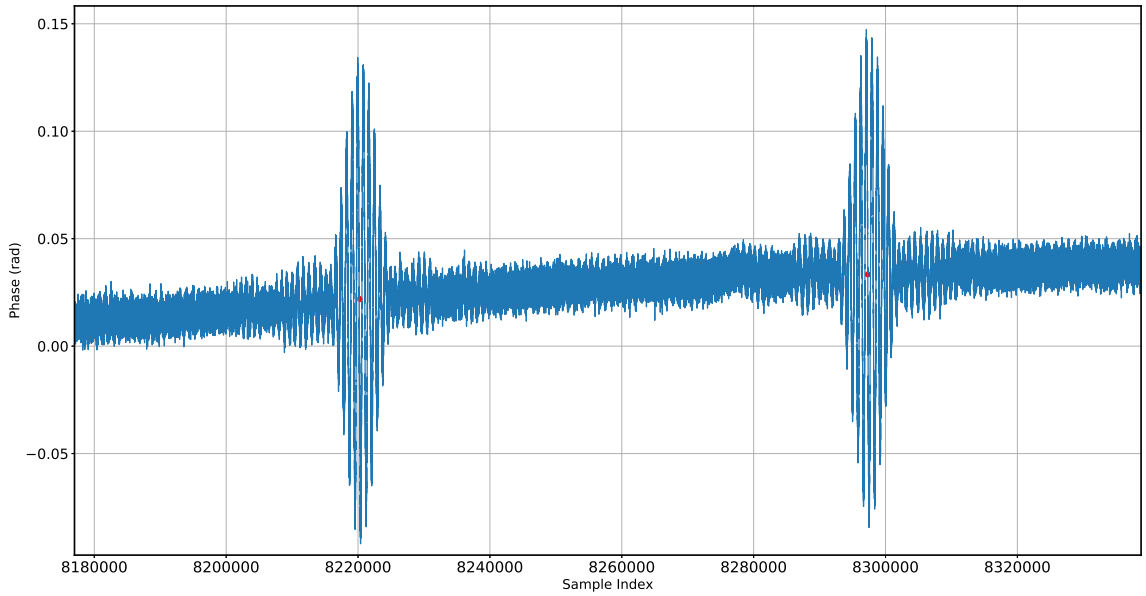

Figure 88. Two interferograms from the raw FTS measurement timestream, with their approximate center locations marked with red dots.

Several aligned and low-pass filtered interferograms for a single channel (prior to scaling) are shown in Figure 89. The variation in peak-to-peak amplitudes is the

result of the different beam-filling factor of each FTS position. After normalizing each I_c, they must be corrected for nonlinearities (4). The nonlinearities appear as amplitude (or phase, or Δf) asymmetries, which are particularly visible in the central portion of each I_c. If the nonlinearity is left uncorrected, sidelobes are produced in the final passband spectrum. Part of the nonlinearities seen in these interferograms is attributable to imperfect readout calibration during the measurement. When the FTS beam was near peak intensity, it shifted the detectors off of resonance, resulting in a loss of response. The nonlinearity is corrected by fitting a second order polynomial to I_c:

$$I_{c,\mathrm{corr}} = I_c \left[1 + a(I_{c,\mathrm{norm}}) + b(I_{c,\mathrm{norm}})^2\right] \tag{4.12}$$

where $I_{c,\mathrm{norm}} = \frac{I_c - \langle I_c \rangle}{\max(I_c) - \langle I_c \rangle}$.

The coefficients which were used for each of the three detector arrays used in this analysis are shown in Table 20. The template interferograms for each band, I_T, are shown in Figure 90, where the solid blue trace is the corrected template and the dashed red trace shows the template created from uncorrected channel interferograms. The nonlinearity is most pronounced in the 500 μm data, which is reasonable given its higher responsivity.

Band	a	b
250W μm	0.01	0.01
350 μm	0.05	0.01
500 μm	0.2	0.1

Table 20. Coefficients used in nonlinearity correction of interferograms.

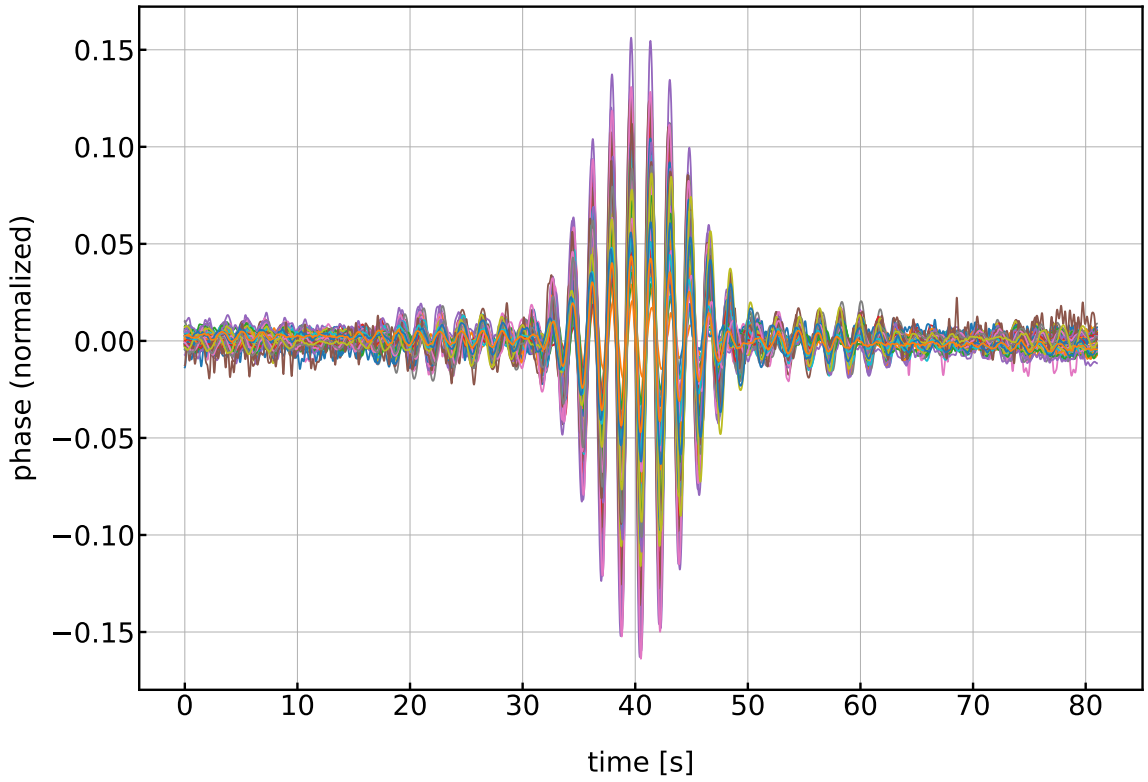

Figure 89. Several aligned interferograms for the same channel, prior to applying weights.

The passband spectra are produced by taking the inverse-Fourier transform of each I_T, and shifting the Fourier-frequencies into the optical band (6). The scale factor used to shift the frequencies is:

$$f_{\text{scale}} = \frac{f_s/2}{f_c}\frac{c}{2v_{\text{mirror}}}$$
$$\nu_{\text{opt}} = f_{\text{scale}} \times f_{\text{Fourier}} \tag{4.13}$$

where f_c is the cutoff frequency of the low-pass filter used to process each I_c. The passbands are shown in Figure 91, with their amplitudes normalized to the maximum and expressed in dB. The solid and dashed traces correspond to the corrected and uncorrected interferograms. As expected, the frequency response of the measured passbands are far from that of ideal band-pass filters. However, their approxi-

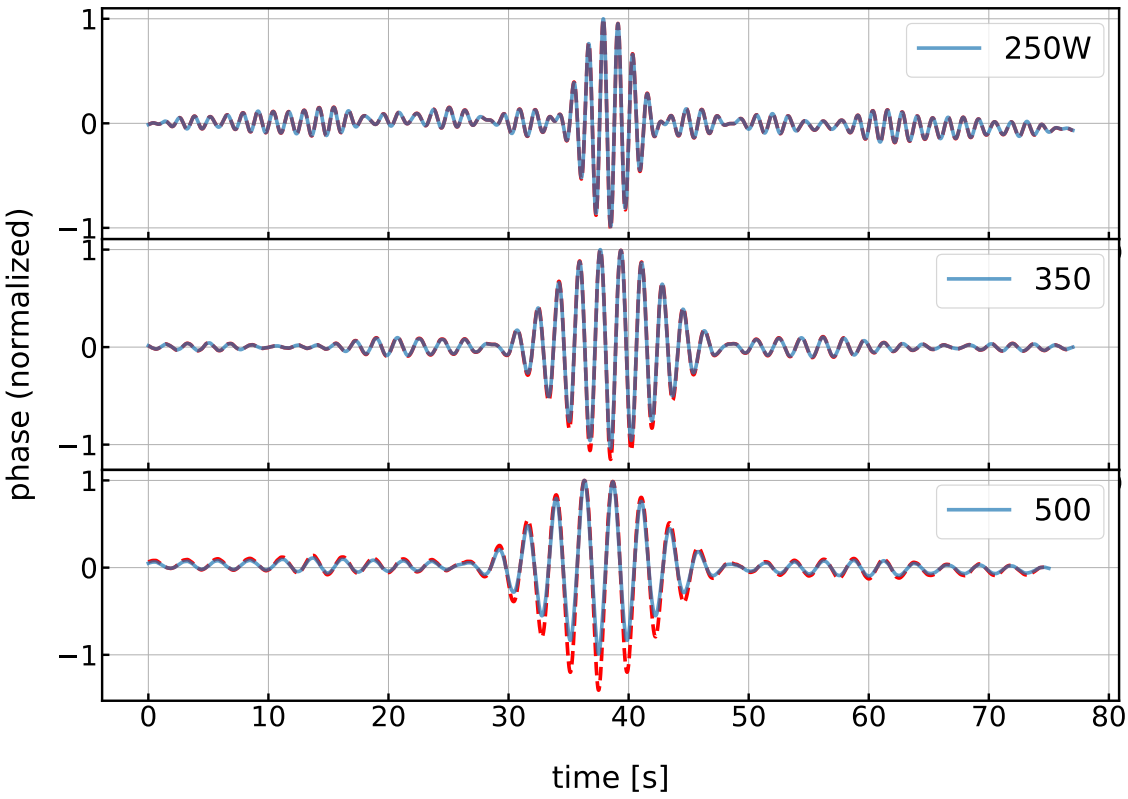

Figure 90. Averaged interferograms for the 250W, 350 and 500 μm arrays. The solid blue trace is corrected for nonlinearity, the red trace is uncorrected.

mate widths are close to the 30% design bandwidth, and their isolation is ∼ - 15 dB. This result is comparable to that measured for previous BLASTs (e.g., Galitzki et al. (2014b)).

4.6.2 Polarization Efficiency

Because BLAST-TNG is a polarimeter, its camera must be able to distinguish signal in one linear polarization from signal in the orthogonal polarization with high accuracy. The polarization efficiency $\epsilon_{\rm pol}$ (or cross-polarization efficiency $\epsilon_{\rm Xpol}$) is

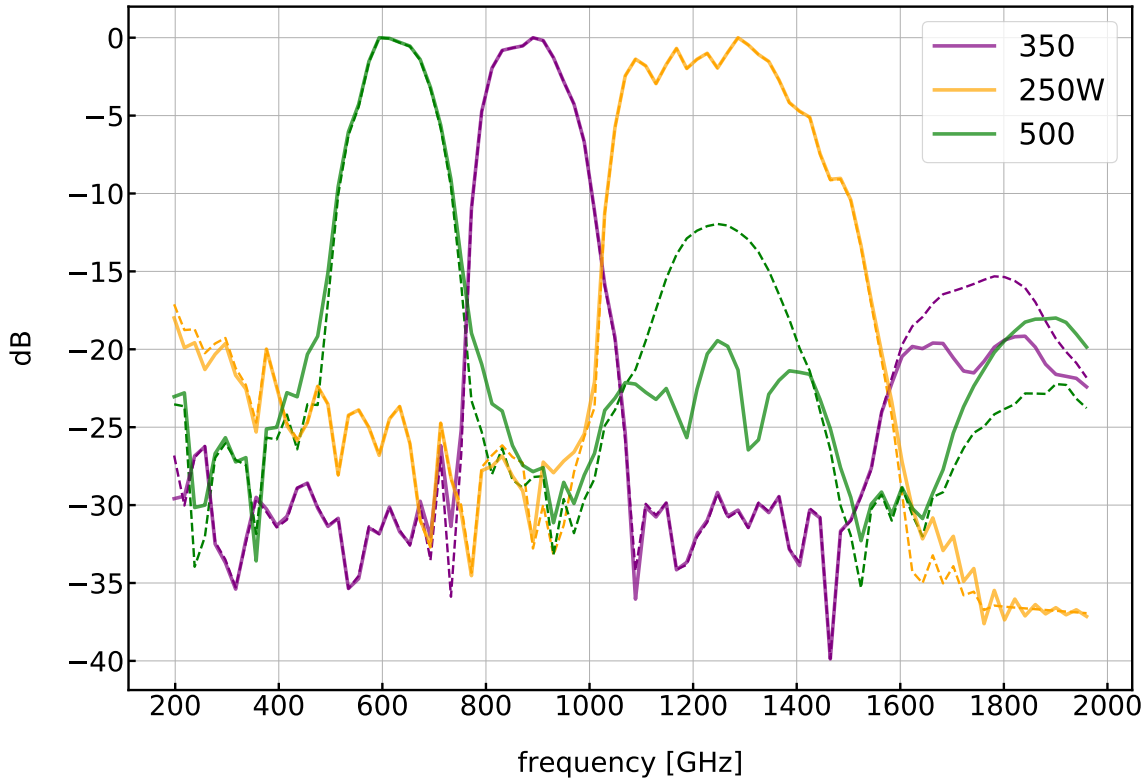

Figure 91. BLAST-TNG optical passbands for the 250W, 350 and 500 μm arrays. The solid traces are corrected for nonlinearity, and the dashed traces are uncorrected.

generally reported as a percentage ϵ_{pol}. It is a measurement of the extent to which detectors with sensitivity to signal with one linear polarization are isolated from signal in the orthogonal polarization (*X* or *Y*, here *pol-1* and *pol-2*). The BLAST-TNG LEKID detectors are dual-polarization sensitive, where each detector is comprised of two pixels with sensitivity to orthogonal linear polarizations (Stokes *Q* and *U*).

Previous lab measurements of a prototype BLAST-TNG 250 μm array showed that ϵ_{Xpol} was at most 0.026 and 0.028 for the *X* and *Y* polarizations (Dober et al., 2016). The instrumental polarization efficiency for the BLASTPol 2012 instrument is reported to be 0.81 (250 μm), 0.79 (350 μm) and 0.82 (500 μm) (ϵ_{Xpol} of 0.19, 0.21, 0.18) (Shariff, 2015).

Here we present initial estimates of the instrumental polarization for three of the BLAST-TNG flight arrays: The 250U, 250W and 350 μm. The data for the 250V and 500 μm arrays was not available at the time of writing. It is important to note that when these measurements were taken the achromatic half wave plate (AHWP) was not installed in the instrument. Additional measurements will be required to estimate any cross-polarization contribution from the AHWP.

A photograph of the measurement setup is shown in Figure 92. A rotatable polarizing grid was attached to the outside of the cryostat window. The plane of the grid was tilted 45° away from the plane of the window to reduce back reflections. A chopper equipped with a heated blackbody source was placed in front of the window. The polarizing grid was then rotated by 360° in 10° intervals. At each rotation angle, channel timestreams were recorded with the chopper in the *on* ($T \simeq 330$) and *off* ($T \simeq 300$) positions. In the following data reduction, we use data corresponding to 22 consecutive rotation angles.

Typically, $\epsilon_{\rm pol}$ is calculated by fitting a channel's polarization response $\mathbf{S}(\theta)$ to a sine wave:

$$\mathbf{S}(\theta) = A \sin(\theta + \Phi) + B \tag{4.14}$$

where θ is the grid angle, and $\epsilon_{\rm pol} = \frac{A-B}{A+B}$. Here, we take a slightly different approach. At any θ, a channel timestream in either the chopper *on* or *off* positions can be written as the sum of a noiseless timestream template T and additive Gaussian white noise (AGWN):

$$\mathbf{x}_c = \alpha_c \mathbf{T} + \mathbf{n} \tag{4.15}$$

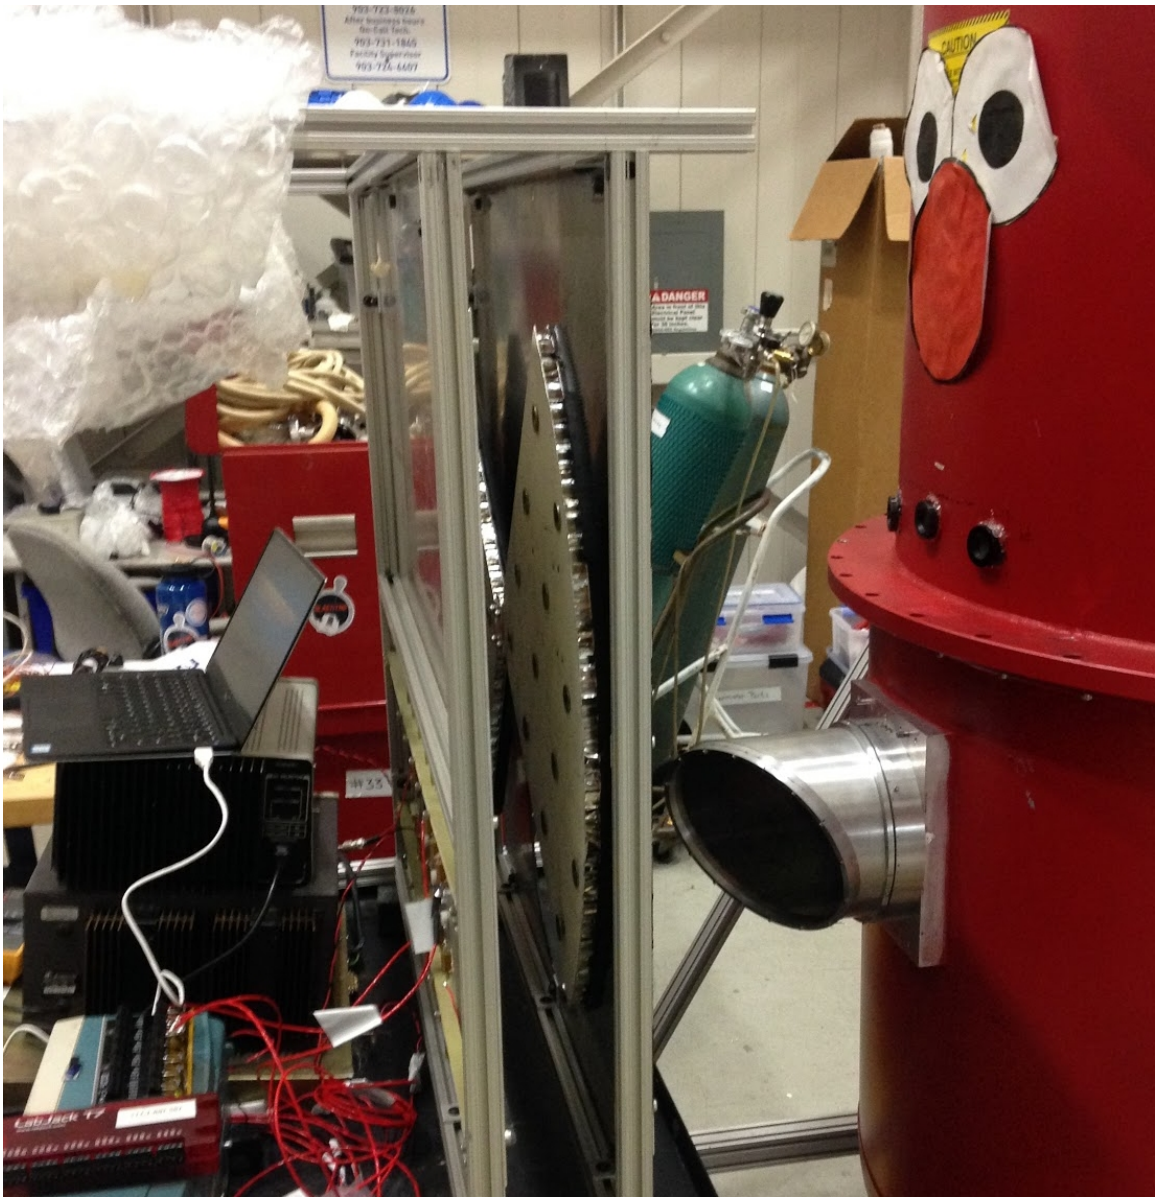

Figure 92. The polarization measurement setup in the CSBF highbay, July, 2018.

where α_c is a channel dependent scaling factor, and $\mathbf{x}_c$ is the phase timestream for channel c. The channel template is calculated by averaging $\mathbf{x}$ over all channels, and normalizing by the peak-to-peak of $\overline{\mathbf{x}}$:

$$\mathbf{T} = 2\frac{\overline{\mathbf{x}}}{max(\overline{\mathbf{x}}) - min(\overline{\mathbf{x}})} \tag{4.16}$$

Then, the scale factor for channel c can be calculated as:

$$\alpha_c = \frac{\mathbf{T}^T\mathbf{x}_c}{\mathbf{T}^T\mathbf{T}} \tag{4.17}$$

The channel's raw phase response ($\mathrm{d}\phi/\mathrm{d}\theta$) is:

$$\mathbf{R}_c(\theta) = \left[\arctan 2(\overline{\mathbf{Q}_c}(\theta)_{\mathrm{on}}, \overline{\mathbf{I}_c}(\theta)_{\mathrm{on}}) - \arctan 2(\overline{\mathbf{Q}_c}(\theta)_{\mathrm{off}}, \overline{\mathbf{I}_c}(\theta)_{\mathrm{off}})\right] \Big/ \Delta\theta \tag{4.18}$$

where $\mathbf{I}_c$ and $\mathbf{Q}_c$ are the I/Q timestreams for channel c. A systematic DC offset can by removed to produce $\mathbf{S}(\theta)$:

$$\mathbf{S}_c(\theta) = \mathbf{R}_c(\theta) - \overline{\mathbf{R}_c}(\theta) + \frac{1}{2}\left[\max\left(\mathbf{R}_c(\theta)\right) - \min\left(\mathbf{R}_c(\theta)\right)\right] \tag{4.19}$$

Finally, the channel's polarization and cross-polarization efficiency are calculated as:

$$\begin{aligned} \epsilon_{\mathrm{pol}} &= \frac{\max(\mathbf{S}_c(\theta)) - \min(\mathbf{S}_c(\theta))}{\max(\mathbf{S}_c(\theta)) + min(\mathbf{S}_c(\theta))} \\ \epsilon_{\mathrm{Xpol}} &= 1 - \epsilon_{\mathrm{pol}} \end{aligned} \tag{4.20}$$

The values for ϵ_{pol} and ϵ_{Xpol} calculated for the 250U, 250W and 350 μm detector arrays are listed in Table 21. Figures 93 to 94 show $\mathbf{S}(\theta)$ for the two orthoganol polarizations (*pol-1* and *pol-2*).

	ϵ_{pol-1}	ϵ_{Xpol-1}	ϵ_{Xpol-1}	ϵ_{Xpol-2}
250U µm	0.818	0.182	0.872	0.128
250W µm	0.884	0.116	0.869	0.131
350 µm	0.844	0.156	0.948	0.052

Table 21. Polarization efficiency $\epsilon_{\rm pol}$ and cross-pol efficiency $\epsilon_{\rm Xpol}$ for three BLAST-TNG detector arrays: 250U, 250W and 350 µm.

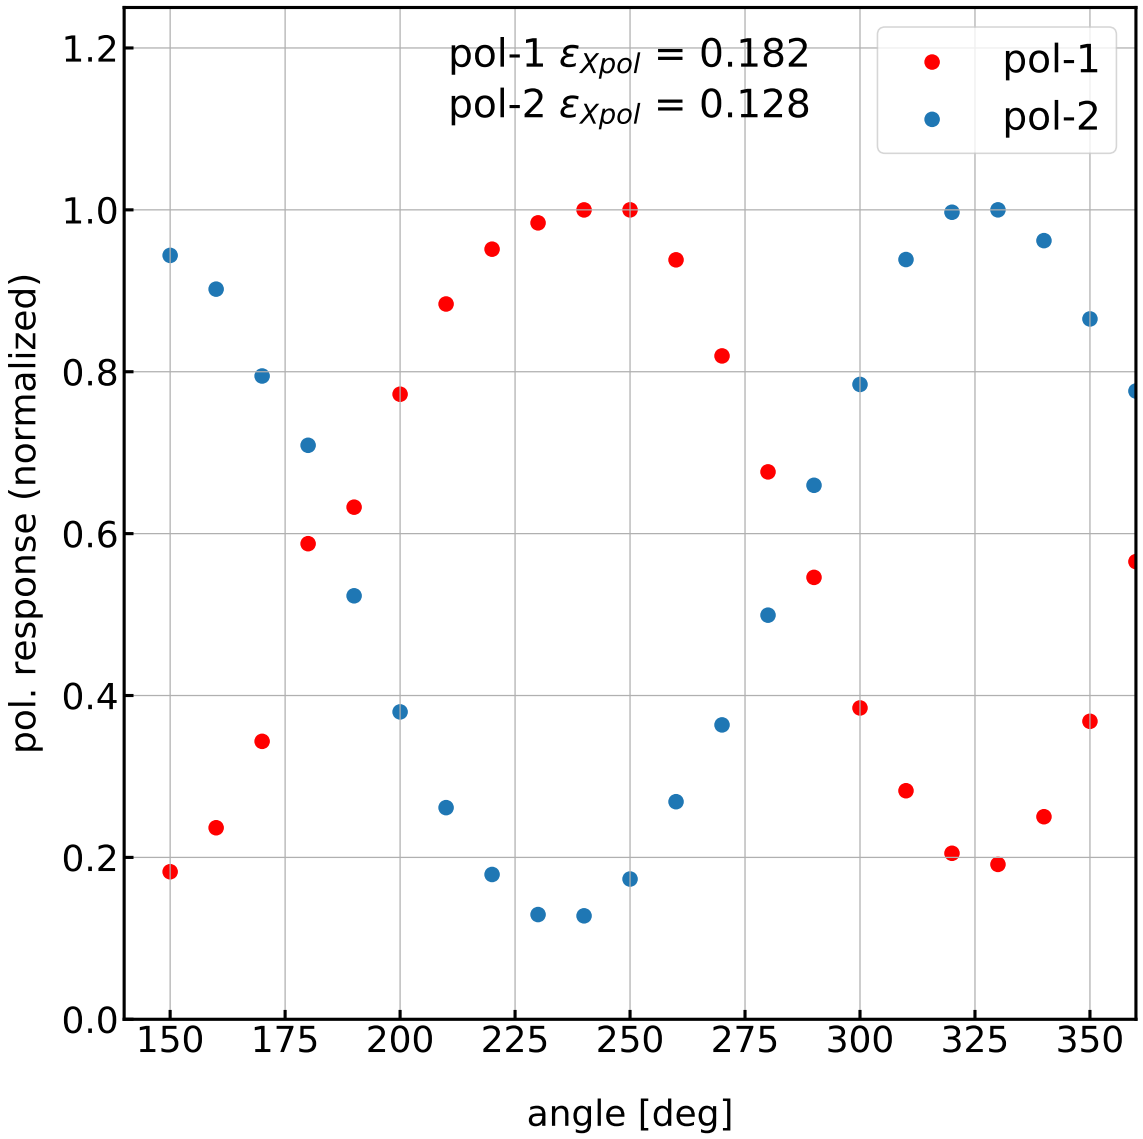

Figure 93. Polarization response $\mathbf{S}(\theta)$ for the BLAST-TNG 250U array.

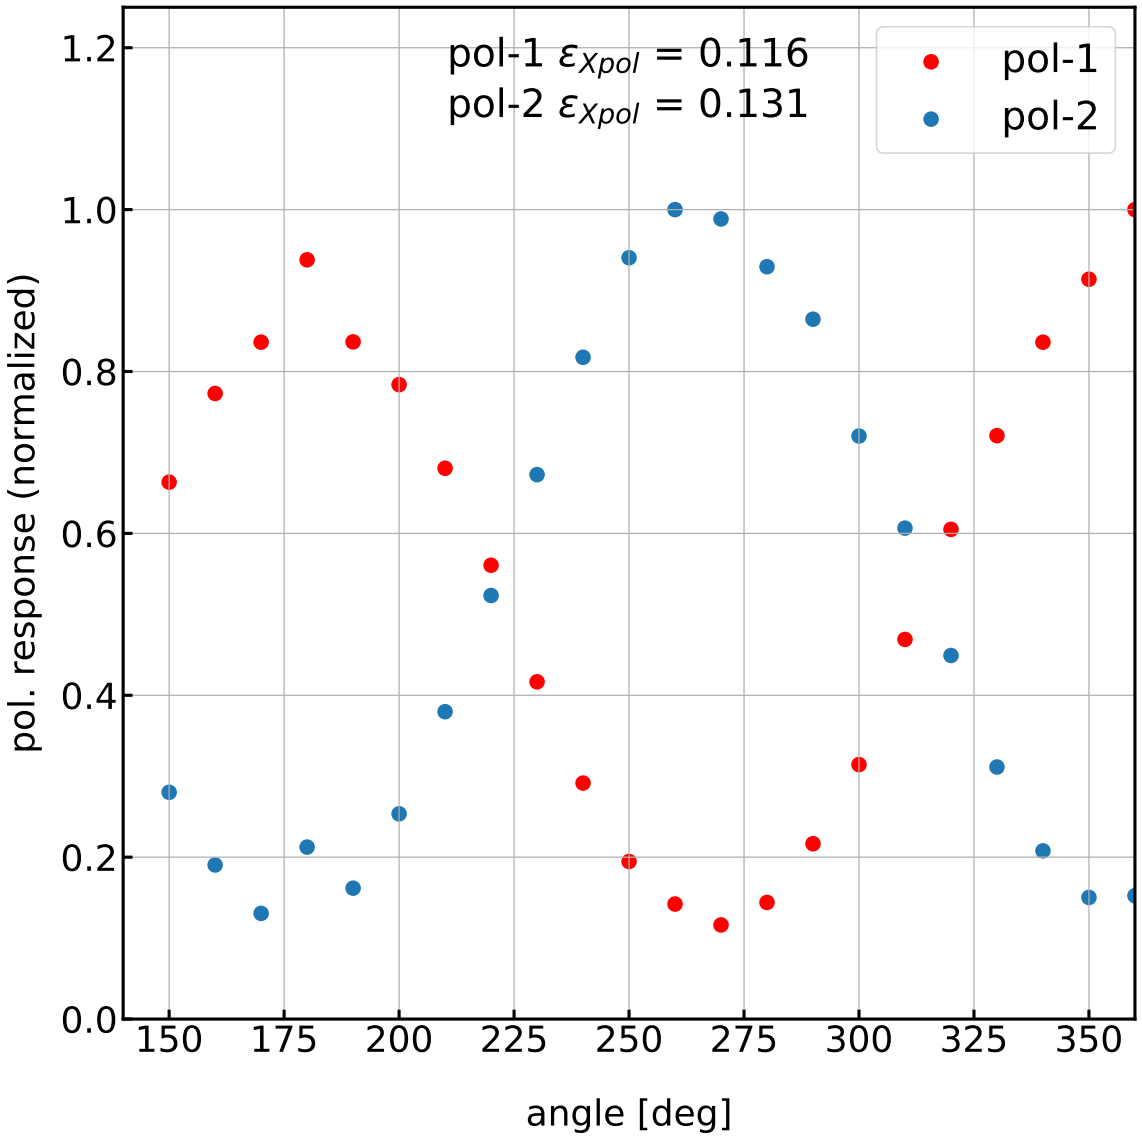

Figure 94. Polarization response $\mathbf{S}(\theta)$ for the BLAST-TNG 250W array.

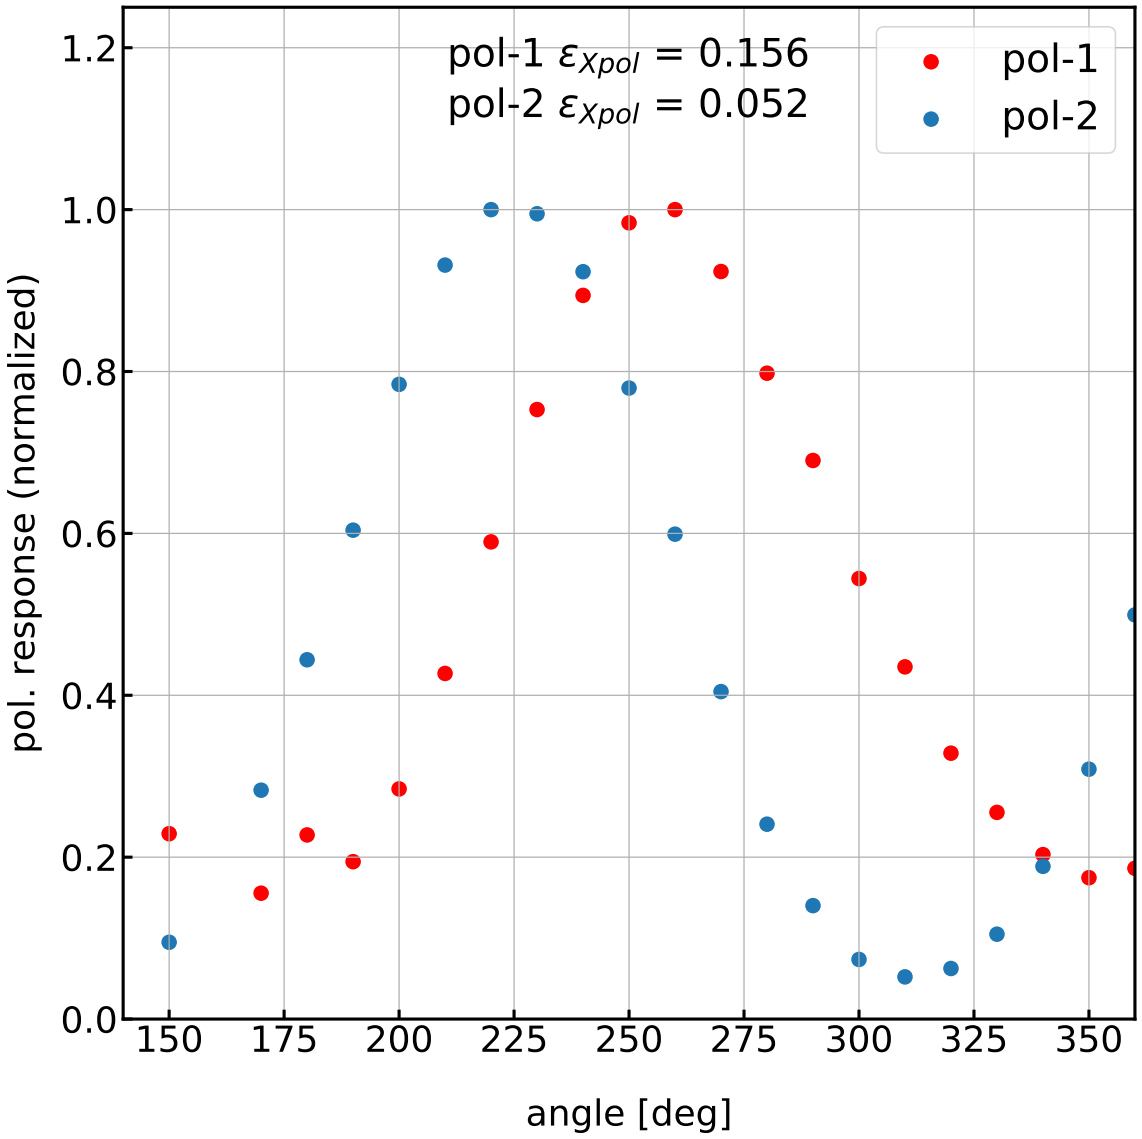

Figure 95. Polarization response $\mathbf{S}(\theta)$ for the BLAST-TNG 350 μm array.

Although these measurements of the instrumental $\epsilon_{\rm pol}$ for BLAST-TNG are preliminary and do not take contributions from the AHWP into account, they are comparable to the instrumental polarizations measured for BLASTPol (Shariff (2015)).

4.6.3 Additional Optical Tests

The results presented in the previous two sections are preliminary, and will be expanded upon in the lead-up to the planned 2020 flight. Additional testing that is underway at the time of writing includes beam mapping, spatial pixel identification and characterization of the HWP.

Chapter 5

MAGNETIC FIELD MORPHOLOGY OF THE CARINA NEBULA COMPLEX

In the following sections, we describe the preliminary data reduction and analysis of BLASTPol 2012 (Galitzki et al., 2014a) observations of the Carina Nebula Complex (NGC 3372, hereafter CNC). These BLASTPol maps have the highest spatial resolution of any polarized sub-mm map of the CNC to date ($\sim$2.5′). Recently, Shariff (2015) used the same data set to produce the CNC's dust polarization spectrum. The authors calculated the polarization ratio $p_\lambda/p_{350\mu m}$ along 314 sightlines, and found it to be flat within $\pm 15\%$ over the three BLASTPol wavebands (250, 350 and 500 μm). While the flatness of the spectrum agrees with other BLASTPol observations of molecular clouds (MCs) (e.g., Ashton et al. (2018); Gandilo et al. (2016)), it is in stark contrast to existing sub-mm/FIR polarization spectra of MCs (see, e.g., Vaillancourt and Matthews (2012)), which show V-shaped spectra with a negative slope towards the FIR, a positive slope towards the millimeter, and minima near 350 μm.

The goal of the data analysis presented in this work is to describe the morphology of the magnetic field (B-field) in the plane-of-the-sky (POS) $B_{\rm POS}$ in the CNC as revealed by the BLASTPol data, compare it to previous observations, and to provide what are possibly the first estimates of the magnitude of $B_{\rm POS}$ along various sightlines through the cloud. To aid in visualizing the field, we apply a vector field visualization technique called line-integral-convolution (LIC). This chapter is organized as follows:

- Section 5.1 introduces the BLASTPol CNC map and describes its acquisition and calibration.
- Section 5.2 presents an original software implementation of the LIC algorithm.
- In Section 5.3, we compare the BLASTPol CNC to maps made in different wavebands in order to investigate how the dust component of the cloud relates to other sources of emission.
- In Section 5.4, we apply the Davis-Chandrasekhar-Fermi Method (DCFM) to the CNC map and present estimates of $B_{\rm POS}$ along several sightlines of interest.

5.1 BLASTPol Observations of Carina

The BLASTPol telescope (Galitzki et al. (2014a)) launched from McMurdo Station, Antarctica, during the summer 2012/2013 season, and observed several targets over the course of ∼12.5 days. Its camera incorporated 139, 88 and 43 pixels in 30% bands centered at 250 μm, 350 μm and 500 μm. The detectors were superconducting spiderweb bolometers (SWBs) based on those used in Herschel SPIRE (Griffin et al., 2003). The BLASTPol telescope had a 1.8 m primary mirror, which provided nominal diffraction limited resolutions in the three bands of 30″, 42″ and 60″. However, deformations in the telescope optics introduced a non-Gaussianity to the point spread function (PSF) (see Fissel et al. (2016)). After correcting for the warped PSF, the resulting resolution of the maps is ∼2.5′.

BLASTPol observed the CNC for 4.2 hours, and the resulting map is ∼2.5 deg^2 in area. The target coverage was chosen to overlap with previous observations taken at 450 μm by the Submillimeter Polarimeter for Antarctic Remote Observations (SPARO, (Li et al., 2006)). As discussed in Benton (2015), the detector timestreams

taken during the CNC mapping were contaminated by interference from the Tracking and Data Relay Satellite System (TDRSS) traffic. The loss of data made it difficult for the map maker to converge to a solution, and some artifacts are present in the resulting maps. The maps were created with the Time-Ordered Astrophysics Scalable Tools (TOAST)[28] package. TOAST produces a (3×3) I, Q, U covariance matrix for each pixel (where I, Q, and U are the Stokes parameters). Raw analog-to-digital-converter (ADC) detector counts were calibrated into beam-averaged flux density (MJy/sr) using the Planck all-sky thermal dust model (Ade et al., 2015).

5.1.1 Data parameterization

The map for each BLASTPol band are parameterized in terms of the first three Stokes Parameters, I, Q and U. The Stokes parameters describe the polarization state of measured light. In this work, they have been calibrated into units of MJy/sr. Each Stokes parameter can be written as either the sum or the difference between the squared magnitudes of electric field (E-field) vectors which are oriented along different axes of a Cartesian coordinate system (for a detailed description of the Stokes parameters, see Hecht (2002)). The coordinate system can be parameterized in terms of Jones Vectors, which form three orthogonal bases: x and y, a and b (rotated 45° with respect to the $x-y$ plane), l and r ($\hat{l} = (\hat{x}+j\hat{y})/\sqrt{2}, \hat{r} = (\hat{x}-j\hat{y})/\sqrt{2}$).

Using these bases, the Stokes parameters can be written as:

[28] https://github.com/tskisner/TOAST

$$
\begin{aligned}
I &= \langle |E_x|^2 \rangle + \langle |E_y|^2 \rangle = \langle |E_a|^2 \rangle + \langle |E_b|^2 \rangle = \langle |E_l|^2 \rangle + \langle |E_r|^2 \rangle \\
Q &= \langle |E_x|^2 \rangle - \langle |E_y|^2 \rangle \\
U &= 2\langle \Re \left(E_x E_y^\star \right) \rangle = \langle |E_a|^2 \rangle - \langle |E_b|^2 \rangle \\
V &= -2\langle \Im \left(E_x E_y^\star \right) \rangle = \langle |E_l|^2 \rangle - \langle |E_r|^2 \rangle
\end{aligned}
\tag{5.1}
$$

where $V = 0$ in the case of linearly polarized light, and the inequality becomes an equality in the case of circularly polarized light. The parameters are related as:

$$
|I|^2 \geq |Q|^2 + |U|^2 + |V|^2 \tag{5.2}
$$

with:

$$
Q + jU = Pe^{j2\Psi} = pIe^{j2\Psi} \tag{5.3}
$$

where $P = \sqrt{Q + U}$ is the total intensity of linearly polarized light, and Ψ is the polarization angle

$$
\Psi = \frac{1}{2} \arctan2(U, Q) \tag{5.4}
$$

where $-90^\circ \leq \Psi \leq 90^\circ$. In this work we use the astronomical polarization angle convention, where Φ increases counter-clockwise from the north (Shariff, 2015). The inferred angle of $B_{\rm POS}$ is $\Phi = \Psi + 90^\circ$.

The polarization fraction is defined as:

$$
p = \frac{\sqrt{Q + U}}{I} = \frac{P}{I} \tag{5.5}
$$

5.1.2 Data Calibration

In this work, the resolution of the CNC Stokes maps are smoothed to 5′ full-width-half-max (FWHM) to match that of the 353 GHz (850 μm) Planck maps. This is done so that the Planck and Carina data can be combined to produce estimates of $B_{\rm POS}$ using the DCFM (see Section 5.4). The smoothing is performed using Lucy-Richardson iterative deconvolution (Lucy (1974); Richardson (1972)). After smoothing, the steps taken to further clean the data are similar to those described in Shariff (2015). First, a background subtraction is applied to the map data. This is done to remove the brightness contribution from the diffuse Galactic background. The background subtraction is performed by first selecting a region near the edge of the map, where the average flux density is close to the minimum of the map. The average flux density in this region is then subtracted from each map pixel.

After background subtraction, the polarization fraction values p are debiased. The debiased polarization fraction p_{db} is calculated as:

$$p_{db} = \sqrt{p^2 - \sigma_p^2} \tag{5.6}$$

where the variance in p is calculated using the covariance maps and standard error propagation techniques. Following the debias, a 2σ cutoff is applied to p, and all pixels containing $p > 0.5$ are eliminated from the map. Finally, the polarization fractions for each of the three observation bands are divided by their respective polarization efficiencies (0.81, 0.79 and 0.82, for 250, 350 and 500 μm). Figure 96 shows the background-subtracted I_{250} map, with regions of interest (ROI) outlined in green. The regions are listed in Table 22, and discussed in Section 5.3. A detailed analysis of

the polarization fraction over the inner region of the CNC is found in Shariff et al. (2019).

In the following analysis, all calculations use the debiased polarization fraction and polarization angle (see Section 5.2).

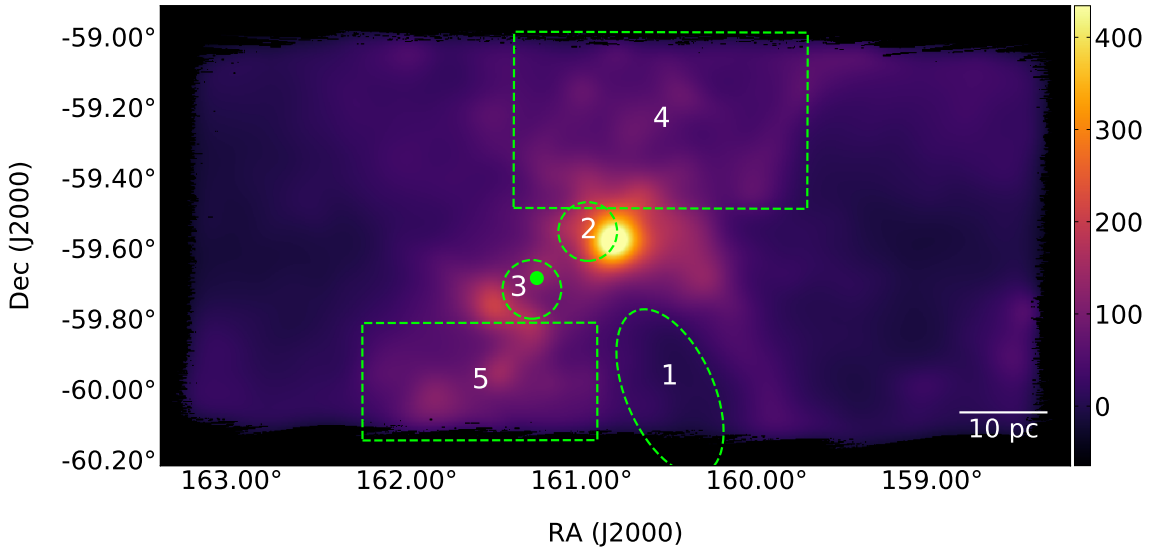

Figure 96. The background-subtracted BLASTPol 2012 I_{250} map of the CNC. The color scale has units of MJy/sr.

5.2 Line-Integral-Convolution

LIC is a data visualization technique which is useful for representing dense vector fields, in either two or three dimensions (Cabral and Leedom, 1993). The visual effect that it produces is akin to that of dropping ink into flowing water. The resulting images provide a visual sense of the direction and uniformity of the vector field,

allowing the eye to notice features such as sources, sinks and loops. The LIC does not, however, indicate the magnitude of the vector field.

The LIC technique has previously been applied to maps of $B_{\rm POS}$ (see, e.g., Ade et al. (2016)). In the following, we present an original software implementation of LIC which has been written in `C/Python`. Although fast, computationally efficient algorithms for LIC exist, the main priorities for this particular implementation were simplicity and flexibility. To that end, the basic skeleton of the code is based on a bare bones LIC algorithm, which is described in Ma (1996).

To create a LIC image requires two basic ingredients: A vector (or pseudovector) field $\boldsymbol{\mathcal{V}}$, and a texture map $\mathcal{T}$, each of size $N \times N$ pixels. In the following example, we use the polarization vector field for Carina, and a texture map consisting of Gaussian white noise. The value of the *ith* pixel in the polarization vector map is a pseudovector whose magnitude is the polarization fraction p_i:

$$\boldsymbol{v}_i = \langle \cos(\Psi_i), \sin(\Psi_i) \rangle \tag{5.7}$$

where the polarization angle Ψ_i is calculated as:

$$\Psi_i = \frac{1}{2} \arctan2(U_i, Q_i) \tag{5.8}$$

The method of LIC is to advect a streamline from each pixel which follows the direction indicated by that pixel's pseudovector (both positive and negative directions) into the adjacent pixels. This process continues until the streamline is terminated by one of several end conditions. The resulting streamlines, when downsampled by some factor, are also useful for map analysis. After the streamlines are created, the values of $\mathcal{T}$ which underly the streamline segments in each pixel are convolved with a filter kernel. These steps are described in more detail in the the following sections.

5.2.1 LIC Algorithm

The first step in creating an LIC image is to launch, or advect, a streamline $\sigma(s)$ from each pixel in the map image, $\mathcal{M}$. At any point in $\mathcal{M}$, the direction of the streamline is tangent to the vector at that point: $\frac{d\boldsymbol{\sigma(s)}}{ds} = \frac{\boldsymbol{v(\sigma(s))}}{|\boldsymbol{v(\sigma(s))}|}$. The entire vector field for the CNC is shown in Figures 97 and 98, where the vectors have been rotated by 90° to correspond to the direction of B_{POS}, and decimated by a factors of 5× and 10×, respectively. The streamlines extend forwards and backwards along their starting vectors, into each adjacent pixel. This process continues until it is terminated by a break condition.

One common break condition is when a streamline enters into a pixel containing a null vector. In this case, the streamline may simply be terminated, or some non-zero value may be substituted for the null vector in order to attempt to continue the streamline. Another common break condition is when a streamline encounters the edge of the map. Discontinuities occur when a streamline enters an adjacent pixel whose vector forms an angle with the first one which exceeds some critical value (e.g., in the range of of 120–180°). There are several ways of dealing with each case, each of which results in a slightly different LIC image.

In the CNC code, streamline advection is achieved using Eulerian advection, as in Cabral and Leedom (1993). In this method, each streamline consists of a series of points $P[i]$ which are calculated recursively as:

$$P[i] = P[i-1] + \frac{\vec{v}_{i-1}}{\|\vec{v}_{i-1}\|}\Delta s_{i-1} \tag{5.9}$$

where:

$P[i]$ (for $i > 0$) is the end point of the streamline segment starting at $P[i-1]$.

$P[0] = (x + 0.5, y + 0.5)$

Δs_{i-1} is the length of the streamline segment starting at $P[i-1]$ and ending at $P[i]$.

$\vec{v}_{i-1}$ is the polarization vector defined in Equation 5.7 which begins at the center of pixel $i-1$.

Δs_{i-1} is the length of the segment which connects $P[i-1]$ to $P[i]$.

The length of each streamline segment, Δs_{i-1}, is the smallest positive distance to any of the adjacent pixel walls. Following Ma (1996), by defining the pixel walls themselves as rays, the distance to each wall can be found using the ray-to-ray intersection formulas:

$$\Delta s_{top} = [(y+1) - P_{i-1,y}] \frac{\|\vec{v}_{i-1}\|}{\vec{v}_{i-1,y}}$$

$$\Delta s_{bot} = [y - P_{i-1,y}] \frac{\|\vec{v}_{i-1}\|}{\vec{v}_{i-1,y}}$$

$$\Delta s_{right} = [(x+1) - P_{i-1,x}] \frac{\|\vec{v}_{i-1}\|}{\vec{v}_{i-1,x}}$$

$$\Delta s_{left} = [x - P_{i-1,x}] \frac{\|\vec{v}_{i-1}\|}{\vec{v}_{i-1,x}}$$

The kernel value h_i to be used in the LIC for each pixel is calculated by summing over the number of steps in the streamline segment between s_{i-1} and s_i:

$$h_i = \sum_{s_i}^{s_{i+1}} k_s \tag{5.10}$$

where:

$$s_0 = 0$$

$$s_i = s_{i-1} + \Delta s_{i-1}$$

$$k[s] = 1 \text{ , for a Boxcar filter}$$

$$k[s] = \frac{\cos(s\pi/L) + 1}{2} \text{ , for a Hanning filter}$$

The length of the kernel L_k can be varied within the software. Larger values of L_k result in there being more correlated values between adjacent pixels appearing in the final LIC image (i.e., it results in a more smeared appearance). The images in this work were produced using L_k between 50–60.

Finally, the LIC value for the output pixel located at coordinate $[x, y]$, is:

$$I[x,y] = \frac{\sum_{i=0}^{l} N[P_i]h[i] + \sum_{i=0}^{l^-} N[P_i^-]h[i^-]}{\sum_{i=0}^{l^-} h[i] + \sum_{i=0}^{l'} h[i']} \tag{5.11}$$

where:

$N[P_i]$ is the value of the noise texture image at vector position (P_x, P_y).

l is a a value of i which satisfies: $s_i \leq L < s_{i+1}$.

The superscripts in Equation 5.11 indicate backwards ($-$) advection.

5.2.2 Visualizing Magnetic Fields with LIC

The LIC method described in the previous sections can produce a wide variety of images, which differ according to the parameters which are used to generate the streamlines and convolution values. The resulting images are most sensitive to the

length of streamlines which are used L_{sl}, and the length of the convolution kernel L_k. Larger values of these parameters produce a smearing effect, where parallel streamlines which are close together appear to combine into larger streamlines which are stretched out in the direction of the bulk flow (they are more correlated). Smaller values of L_{sl} and L_k produce images in which individual streamlines are visible (they are more independent). In this work, we choose a moderate value for L_{sl} and L_{k} of 50 (with a Hanning kernel). Streamlines which could not be advected through 50 pixels are truncated, but still included in the final images.

Out of many possible ways of displaying the LIC images, we find that two representations are particularly useful. These are the intensity overlay and the intensity-weighted LIC. An example of the intensity overlay is shown in Figure 101. To create this type of image, the LIC map is superimposed onto the I intensity map as a transparent overlay. The intensity-weighted LIC is created by multiplying the LIC map with the intensity map, and then stretching the image to highlight different intensity regions. An intensity-weighted LIC with an aggressive stretch is shown in Figure 102. A less aggressive stretch is shown in Figure 103. The numbered regions in Figure 103 are key features of the CNC which are labeled in Table 22 and discussed in Section 5.3.

The structures in B_{POS} which are revealed by the LIC maps are discussed in the following sections.

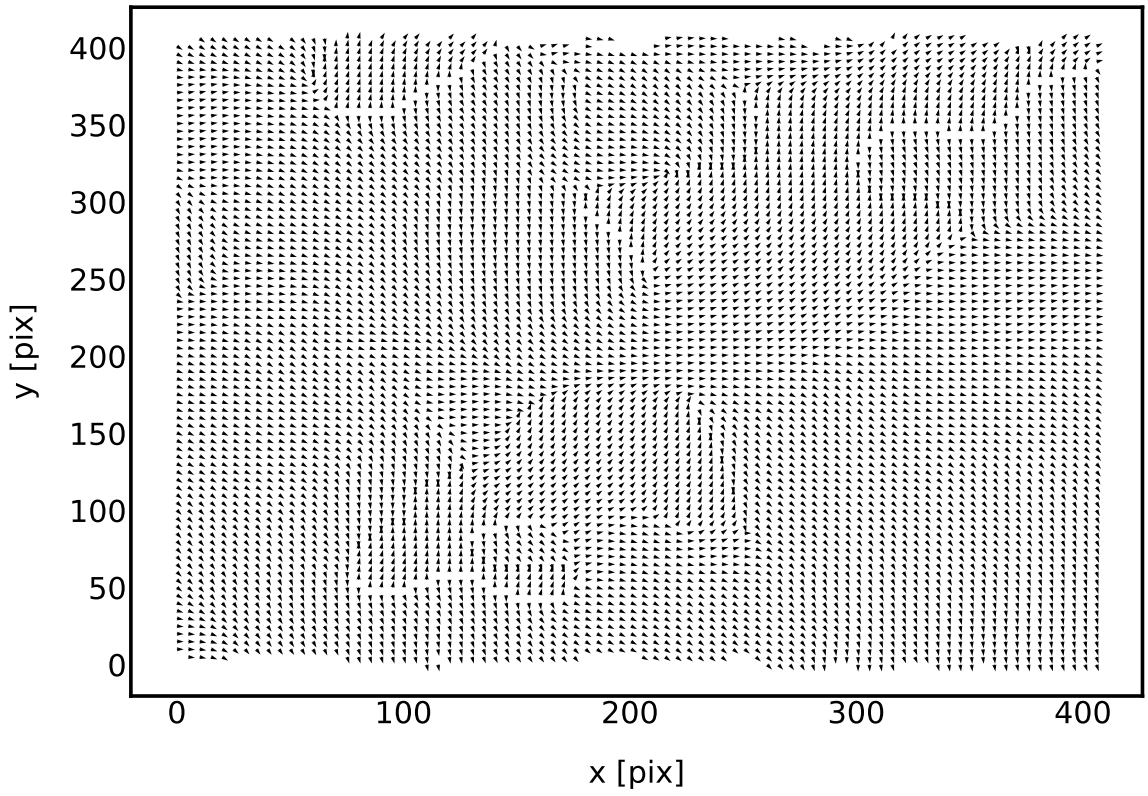

Figure 97. The CNC polarization vector map, with $\boldsymbol{v}_i$ following B_{POS} and decimated by 5×.

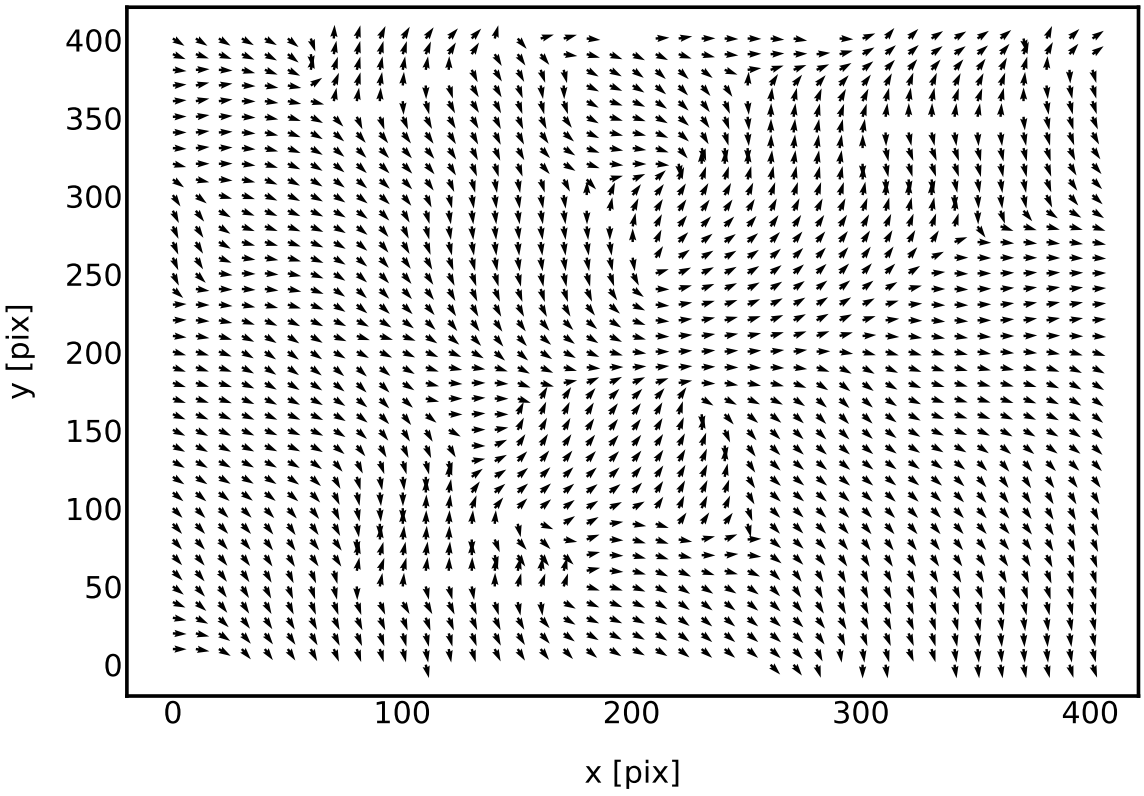

Figure 98. The CNC polarization vector map, with $\boldsymbol{v}_i$ following B_{POS} and decimated by 10×.

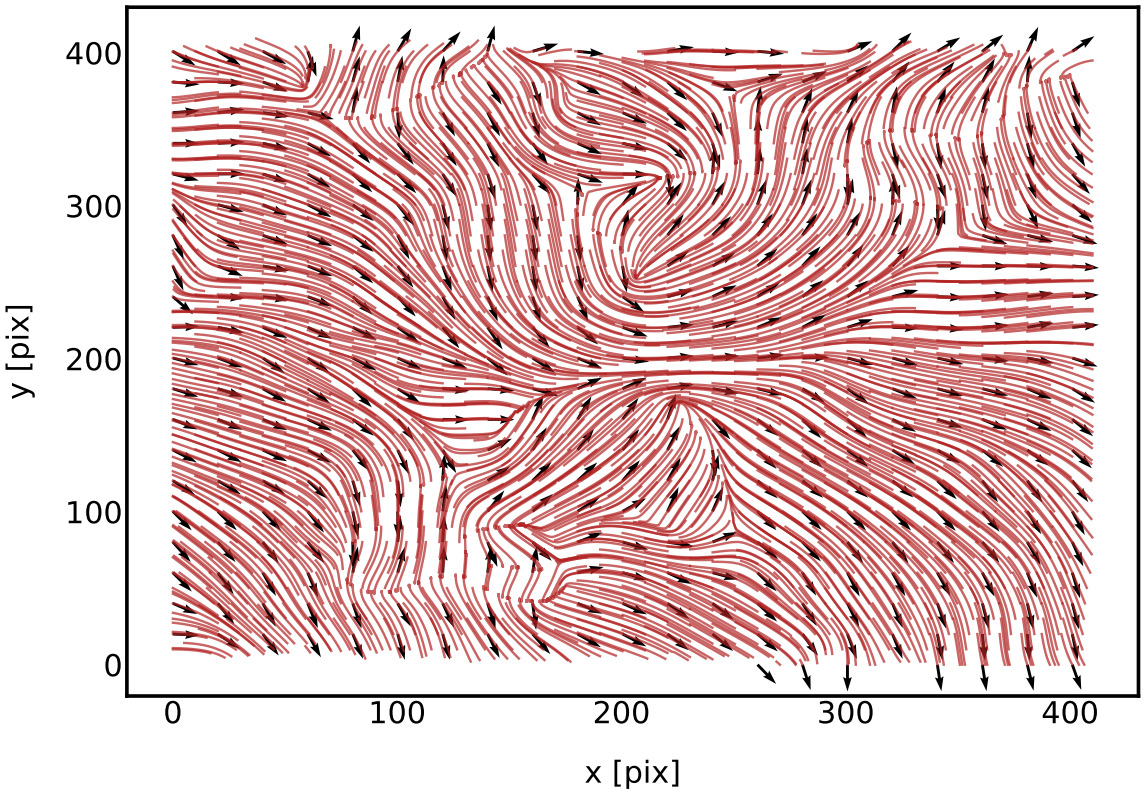

Figure 99. A CNC polarization streamline map, with the streamlines following B_{POS} and decimated by 10×.

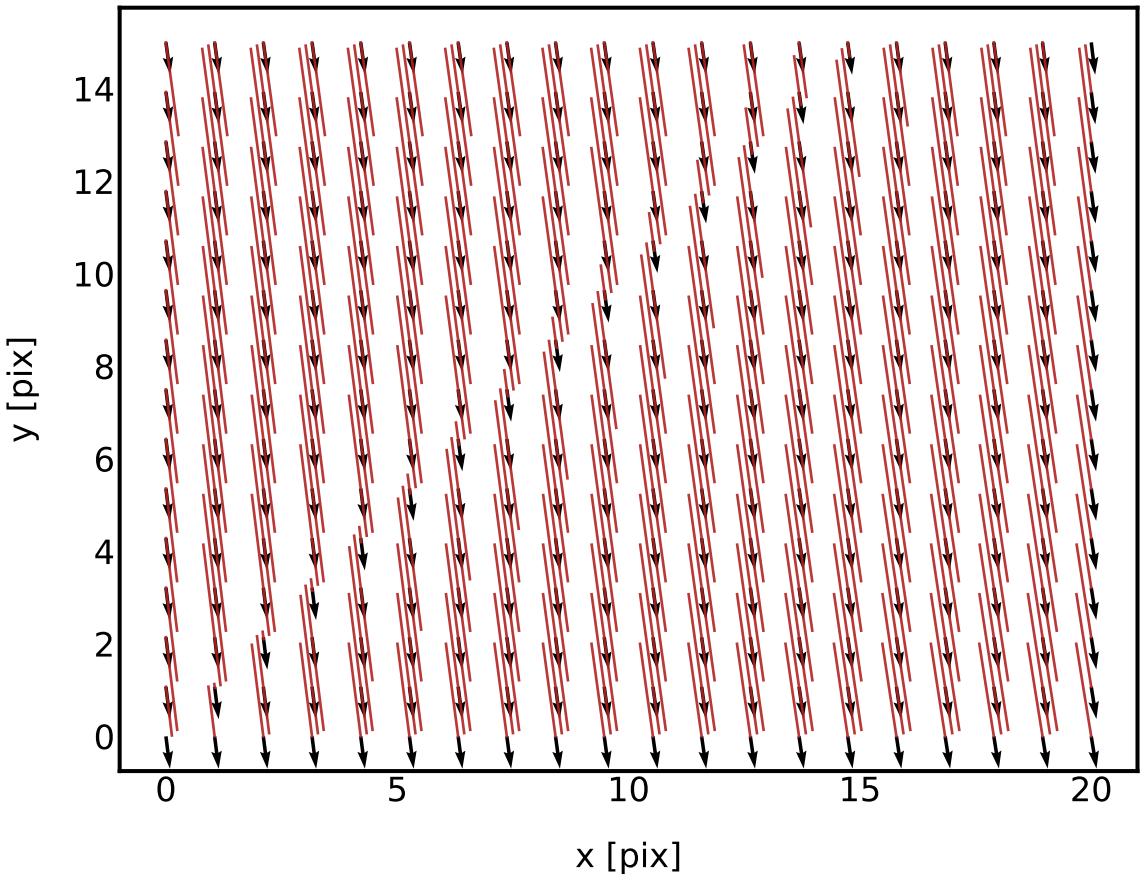

Figure 100. A zoom-in of a region of the vector-streamline map where the polarization angles are very uniform.

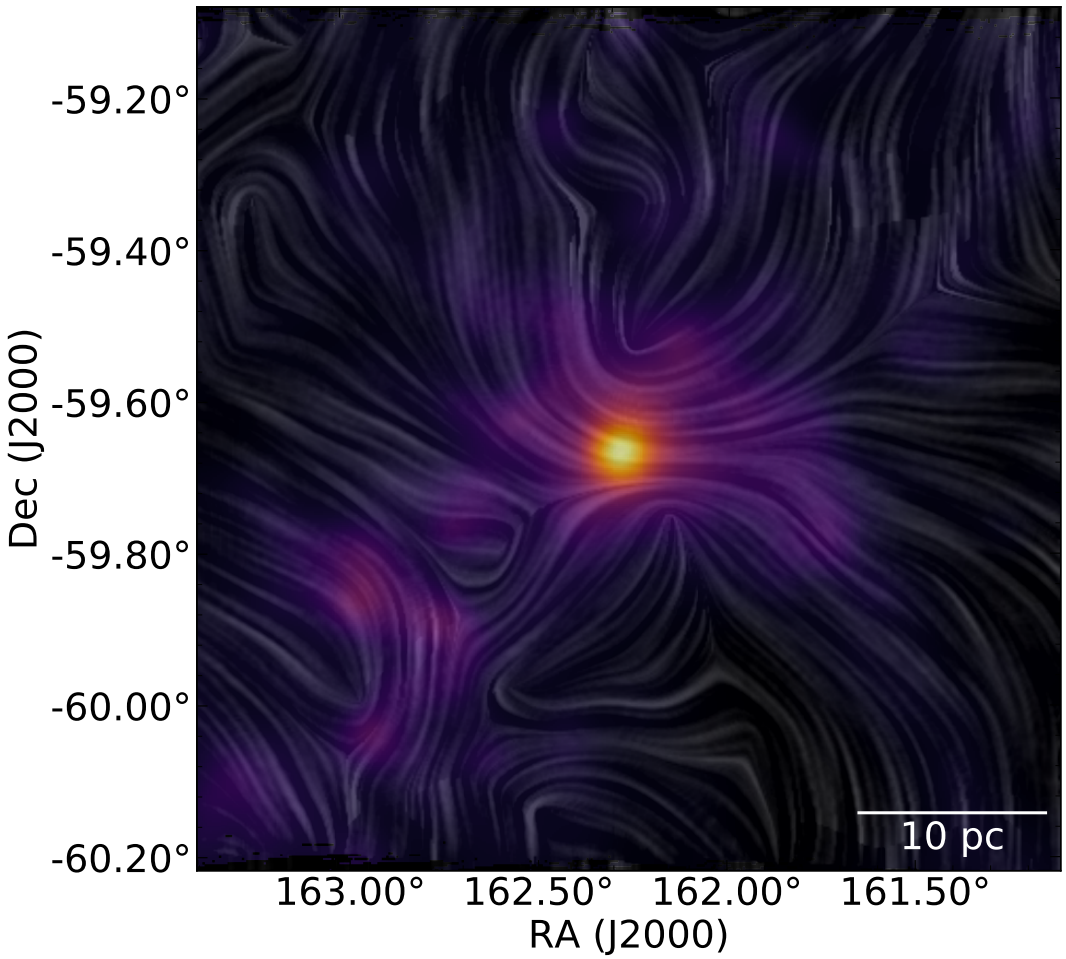

Figure 101. The LIC_{500} map, displayed as a transparent overlay on the I_{500} map.

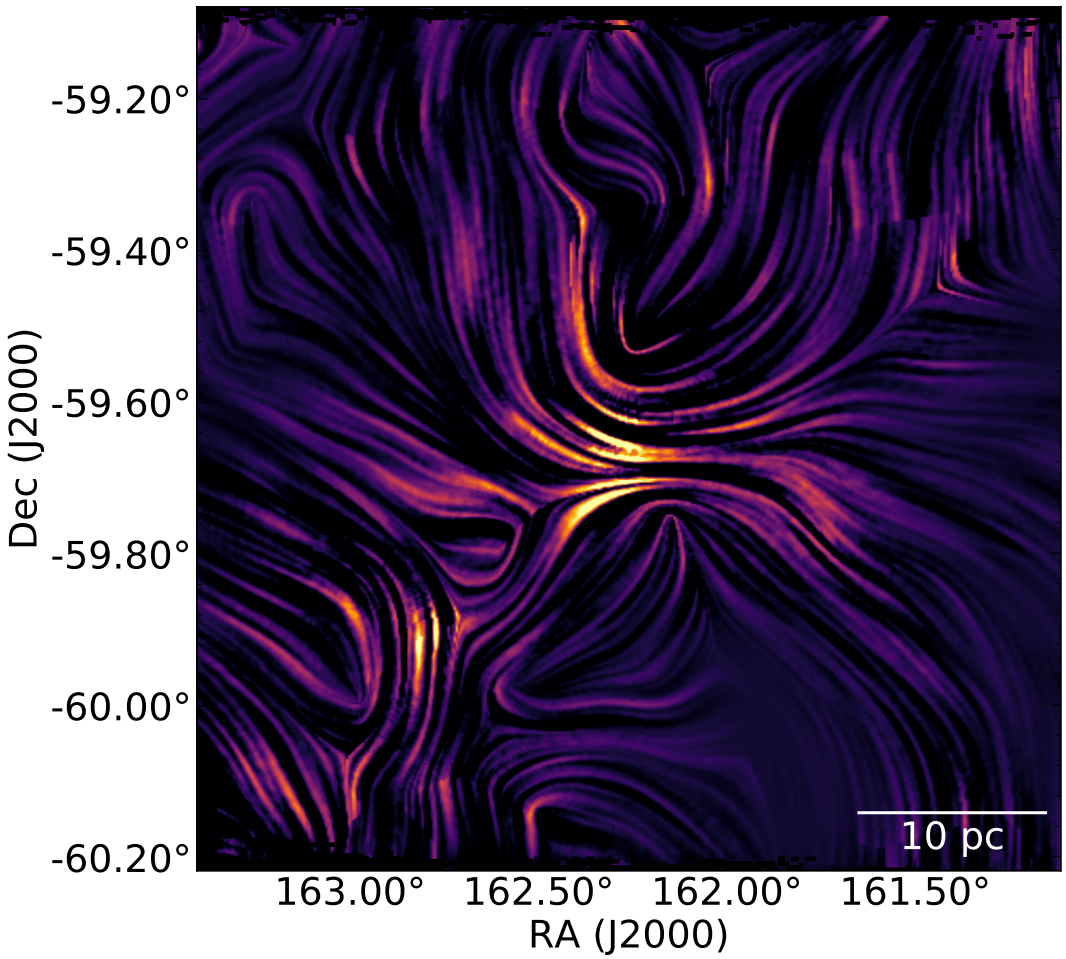

Figure 102. An intensity-weighted LIC_{500} map, displayed with an aggressive intensity stretch.

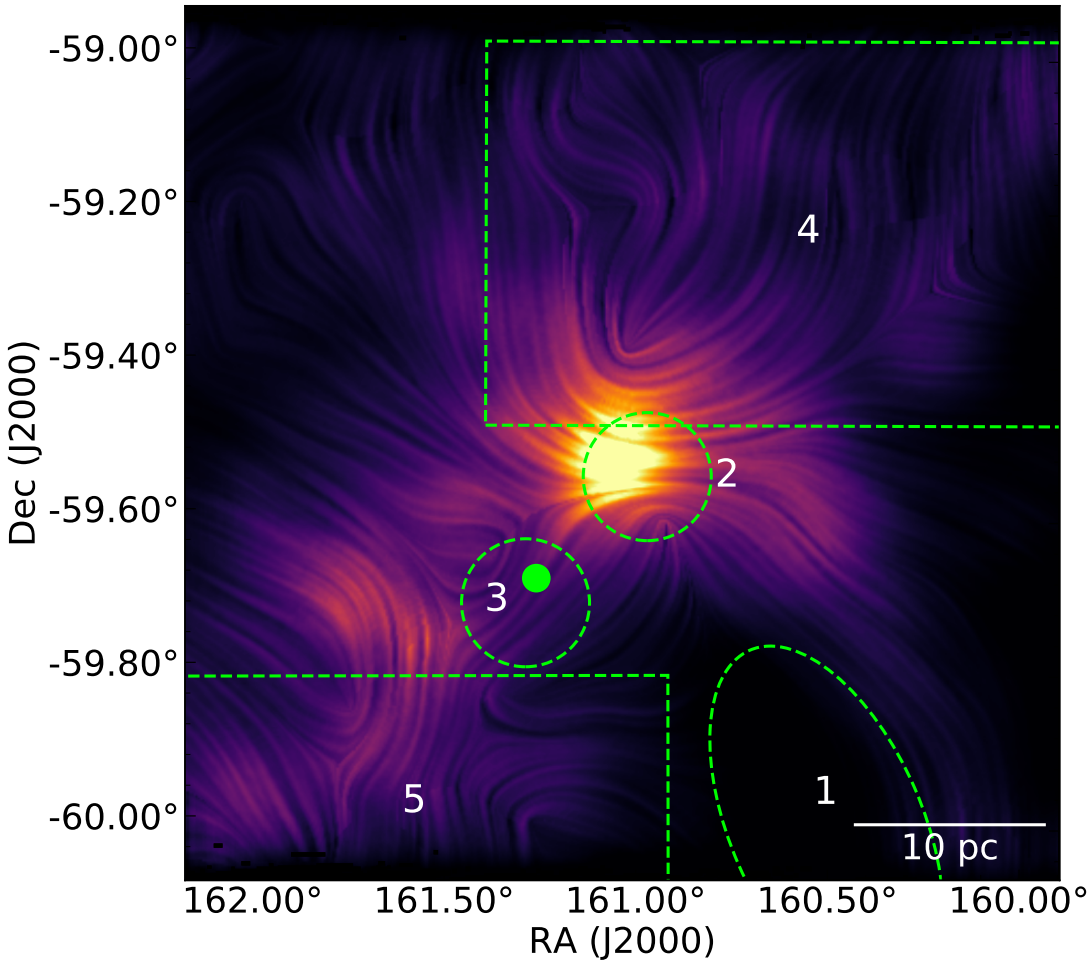

Figure 103. A less aggressive stretch of the intensity-weighted LIC_{500}. The numbered regions are listed in Table 22, and discussed in Section 5.3

.

5.3 Multiwavelength Observations of the CNC

The CNC (10h 45m 08.5s -59° 52′ 04″) is a bright Galactic giant molecular cloud (GMC) which is host to several HII regions and active massive star formation. Its distance is estimated at ∼2.3 kpc (Allen and Hillier (1993); Smith (2006)). However, recent observations by the Gaia space observatory have found that the distance might be closer to 2.6 kpc (Davidson et al., 2018). In this analysis, we assume the more commonly used distance of 2.3 kpc. Although the CNC occupies an area of ∼5 deg^2, in this work we focus on the inner ∼1.25 deg^2.

First discovered in 1752, the CNC has been well studied across the electromagnetic spectrum (see e.g., Smith and Brooks (2008)). It is estimated to contain ∼25,000 $M_\odot$ of material, which includes part of the Carina OB1 association. The inner region of the CNC contains the Homunculus and Keyhole nebulas, as well as two massive open clusters, Trumpler 16 (Tr 16) and Trumpler 14 (Tr 14). Tr 16, the more evolved of the two clusters, is host to two of the brightest star systems in the Milky Way: η Car, a binary system which is known for its strong wind-wind interactions, and WR 25. Tr 14 is thought to be just 0.5 Myr old (Preibisch et al., 2011a).

Ionizing stellar winds which are driven by large OB associations have carved out several conspicuous bubbles throughout the CNC. The most prominent of these is the GUM 31 nebula, to the northwest of Tr 16. Smaller groups of bubbles are visible to the north and south of Tr 16. The variety of high energy processes occurring within the CNC make it an ideal laboratory for the study of stellar feedback and triggered star formation.

The CNC has been observed in the sub-mm/FIR/mm-wave by ground-based (LABOCA, (Preibisch et al., 2011b), SPARO (Li et al., 2006)), balloon-borne

(BLASTPol, (Shariff et al., 2019) and space-based telescopes (e.g., Planck, (Abergel et al., 2014), Herschel, (Preibisch et al. (2012); Gaczkowski et al. (2013); Roccatagliata et al. (2013))). Spectral data in the submillimeter has also been obtained (Oberst et al., 2006). The authors of Li et al. (2006) presented the first maps of $B_{\rm POS}$ pseudovectors in the inner region of the CNC. In their analysis of ∼30 pseudovectors within the inner 1 deg^2 of the CNC, they found that the mean position angle of $B_{\rm POS}$ is within ≈ 15 deg of the GP. They also observed that the pseudovectors to the north and south of the central region appear to run parallel to a section of the perimeters of two large HII regions. Figure 104 illustrates the overlap between the SPARO and BLASTPol map coverages.

Besides BLASTPol and SPARO, the multi-band Planck data constitutes the only polarized sub-mm observations of the CNC to date. In this work, only the Planck 353 GHz (850 μm) maps are used as part of the quantitative analysis.

The Herschel maps constitute the highest resolution FIR maps of the CNC, revealing physical structures in the dust on spatial scales of 0.1–0.4 pc. This spatial scale is thought to correspond to the upper size limit of the filamentary structures which thread the CNC, and which are observed in other MCs. In this analysis, the smoothed resolution of 5′ corresponds to a physical scale of ∼3.5 pc, which is well above the scale of individual filaments. However, this physical resolution is sufficient to gain insight into the larger scale structures seen in the CNC.

The most visible structures revealed in the Herschel maps are discussed in detail in Preibisch et al. (2012). The main features which occur in the inner ∼1 deg^2 of the CNC (and which are visible in the BLASTPol maps) are briefly described below. These features are listed in Table 22. Figures 105 through 107 show semi-

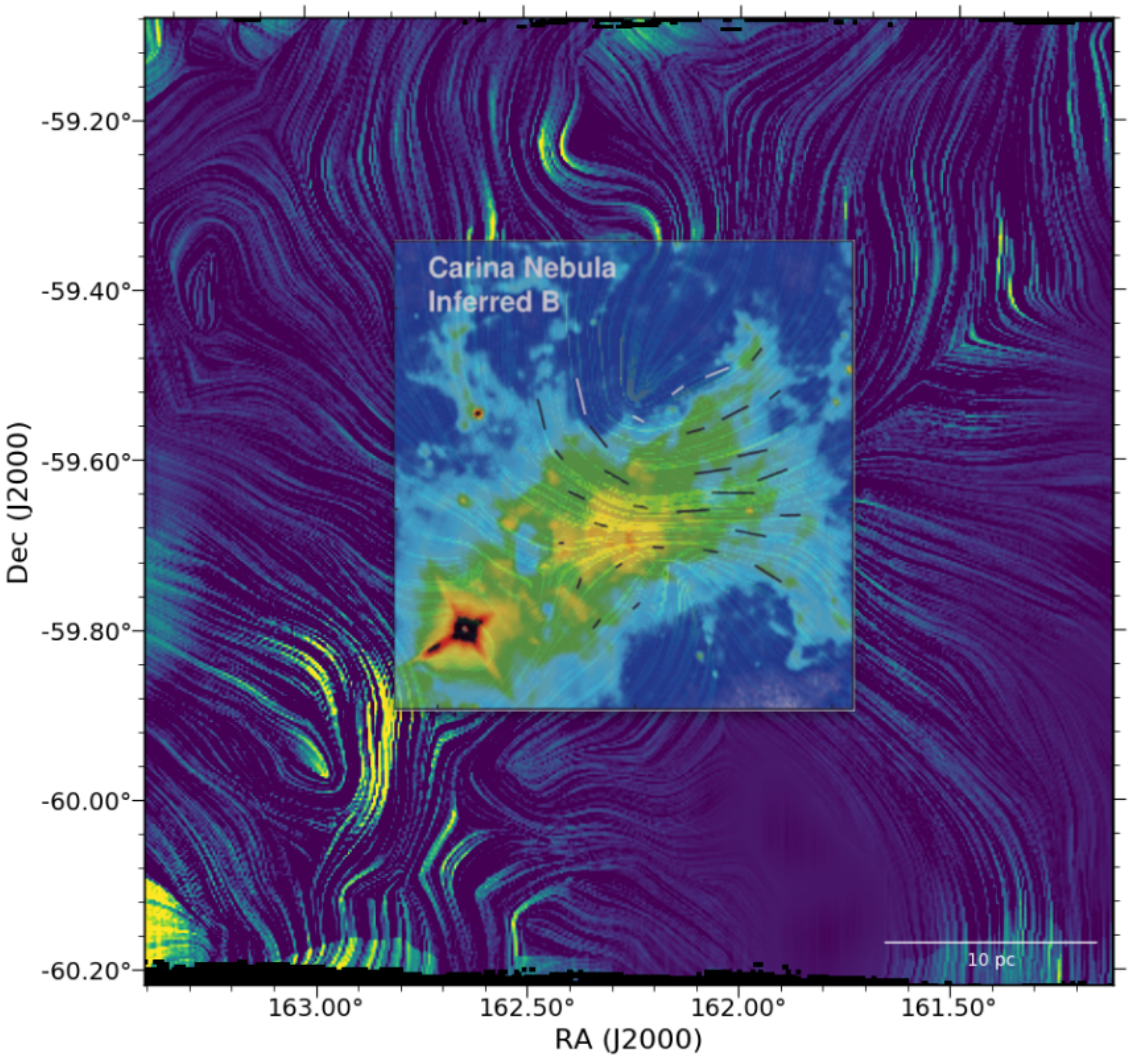

Figure 104. The CNC intensity map taken by SPARO (Li et al., 2006), with polarization vectors (inset), overlayed on the ~1.25 deg^2 of the BLASTPol CNC map.

Region #	Name
1	Southern Bubble (SB)
2	Tr 14
3	Tr 16
4	Northern Bubbles (NB)
5	Southern Pillars (SP)

Table 22. Regions of interest inside the inner ∼1.25 deg^2 of the BLASTPol CNC map. These regions correspond to those used in Preibisch et al. (2012).

transparent overlays of the BLASTPol I_{250} dust maps (in red) with maps in the near-infrared (NIR) (8 μm, Midcourse Space Experiment (MSX, Smith et al. (2000)), R-band (Guide Star Catalog II (GSC2), Palomar and UK Schmidt telescopes (Lasker et al., 2008)) and soft X-ray (Chandra Carina Complex Project (Townsley et al., 2011)). The B-field streamlines produced using Eulerian advection (see Section 5.2) are over-plotted in yellow.

The most prominent features in the sub-mm/FIR maps of the CNC are the voids, or bubbles, to the north and south of Tr 14 and Tr 16. Following Preibisch et al. (2012), we denote these as the Northern Bubbles (NB) and the Southern Bubble (SB). The SB appears as an elongated void to the southwest of Tr 16, with its long axis spanning ∼30 pc. Although the southern extent of the bubble is not visible in the BLASTPol map, a clear, rounded boundary can be seen in the Herschel map. The authors of Preibisch et al. (2012) identify three early-type stars which are potentially located inside the bubble. They point out that the three stars alone most likely do not constitute an energy source which is sufficient enough to have created the SB by way of stellar winds and ionizing radiation. They offer two possible explanations for its origin: 1) That the bubble could be the result of gas which has been forced out

of the central part of the CNC, and 2) that the bubble might be the lower lobe of a bipolar jet formed by stars in the Tr 16 cluster.

The multiband overlay maps support the hypothesis that the Southern Bubble is an enclosed, three-dimensional void. Its edges are well defined by ionized gas which is visible in the R-band map. The Chandra map reveals that the bubble is in fact filled with hot, diffuse gas ($T \geq 10^4$ K). At 8 μm, the emission is from polycyclic aromatic hydrocarbons (PAHs) which have been vibrationally energized by the absorption of far-UV photons. As discussed in Li et al. (2006), the intensity of the 3–11 μm PAH spectrum depends on the ratio of the intensity of the far-UV radiation field to the electron density. Whereas the electron density changes abruptly across the edges of HII regions, the intensity of the illuminating radiation field does not. Therefore, PAH emission is an excellent tracer of the boundaries of HII regions. This effect is visible in the 8 μm overlay, shown in Figure 105.

The direction of the $B_{\rm POS}$ streamlines shown in the overlay images appear to be well correlated with the edges of the SB. Along the boundaries of the SB, $B_{\rm POS}$ is largely parallel to the bubble edges (particularly along the western boundary), which is the expected behavior under the conditions of B-field flux-freezing (Li et al., 2006) (see below).

Unlike the SB, the NB does not appear to be a single HII region. The Herschel maps show it to be a combination of several smaller HII regions which have merged, forming dust clouds and pillars at their intersections. In the multi-band images, then boundaries of individual voids within the NB region are less defined than those of the SB, with the Chandra map showing a more diffuse concentration of ionized gas. However, as with the SB, the $B_{\rm POS}$ streamlines in the NB region do appear to trace out

the shape of a large, semi-spherical bubble. This is especially true for the southern boundary of the NB.

The brightest region of the BLASTPol maps is the central region of the CNC, located between Tr 16 and Tr 14. The eastern edge of Tr 14 is a well-studied photon dominated region (PDR) (see, e.g., Kramer et al. (2008)). The edge of the PDR is loosely defined by the B_{POS} streamlines. The region around Tr 16 is also well defined, with the streamlines appearing to trace out a bubble. The bubble's diameter is ∼3 pc, which is small in comparison to the SB and NB, with a diameter of ∼3 pc. Below this bubble is the region which Preibisch et al. (2012) denotes as the Southern Pillars (SP). This region of pillars is a site of active star formation, and has been shown to contain many compact IR sources (Smith et al., 2010).

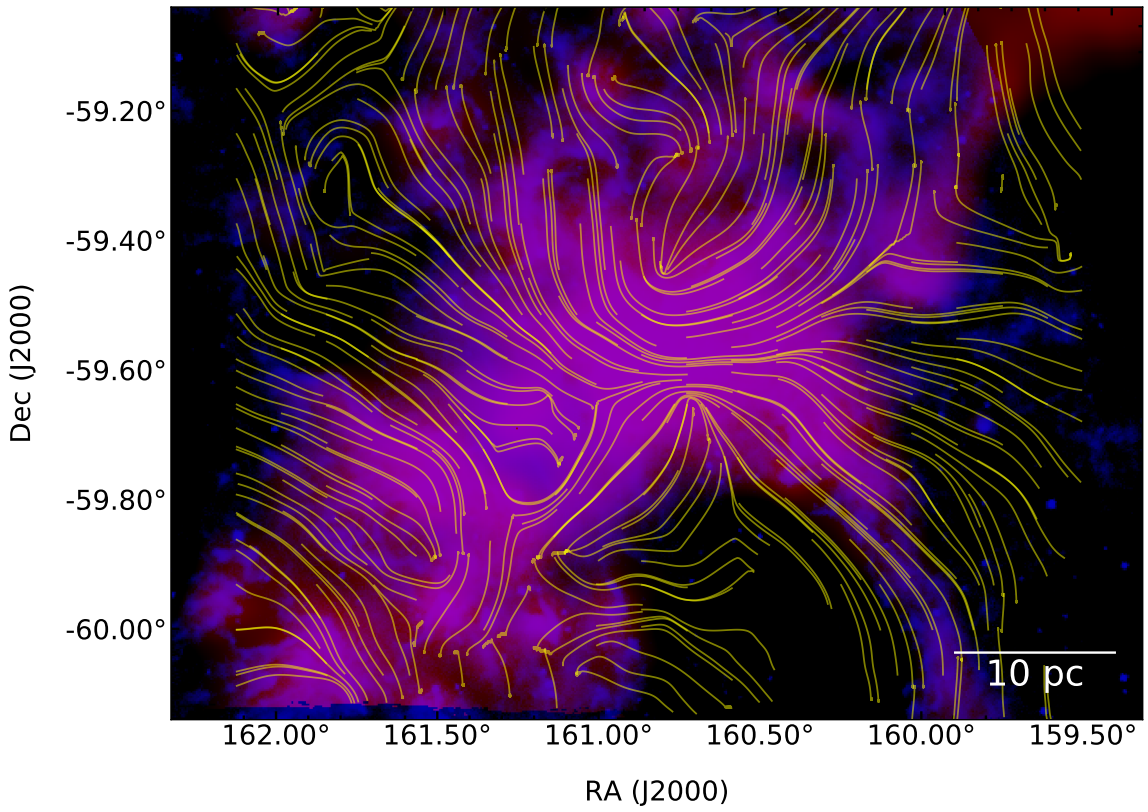

Figure 105. An overlay of the MSX 8 μm CNC map (green) (Smith et al., 2000) with a BLASTPol intensity map (red). A subset of polarization vector streamlines are overplotted in yellow.

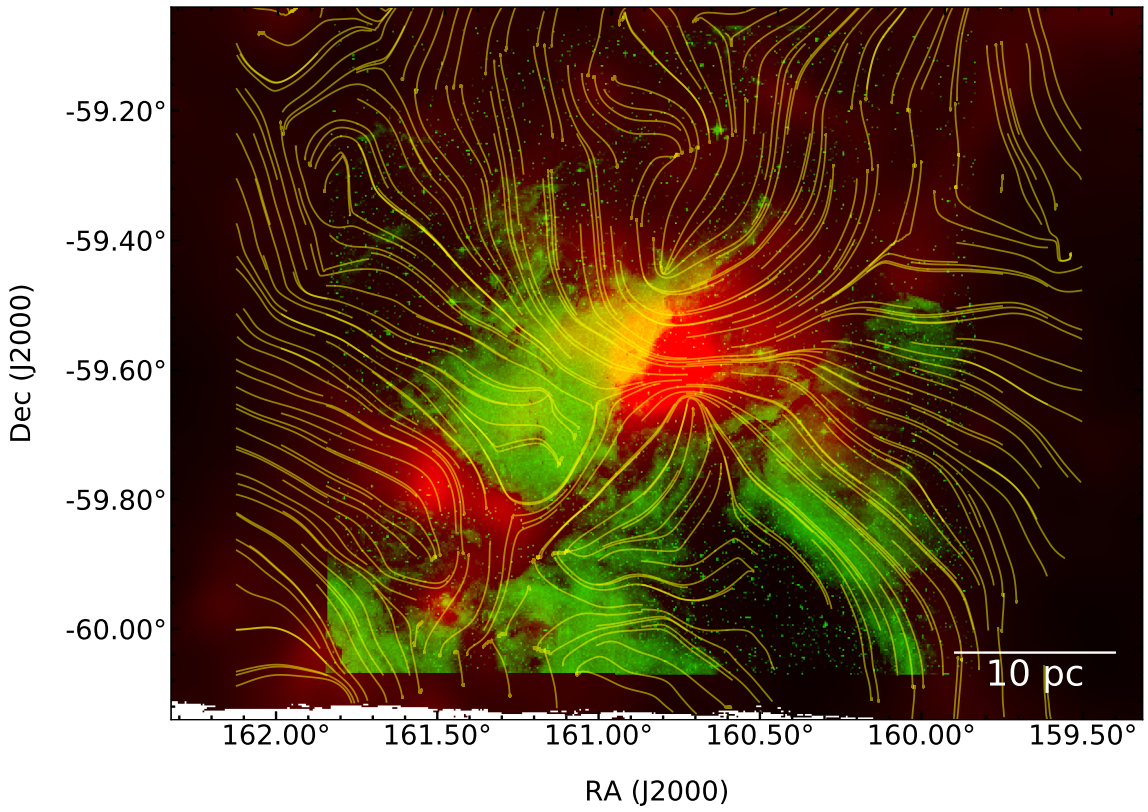

Figure 106. An overlay of the GSC2 R-band CNC map (green) (Lasker et al., 2008) with a BLASTPol intensity map (red). A subset of polarization vector streamlines are overplotted in yellow.

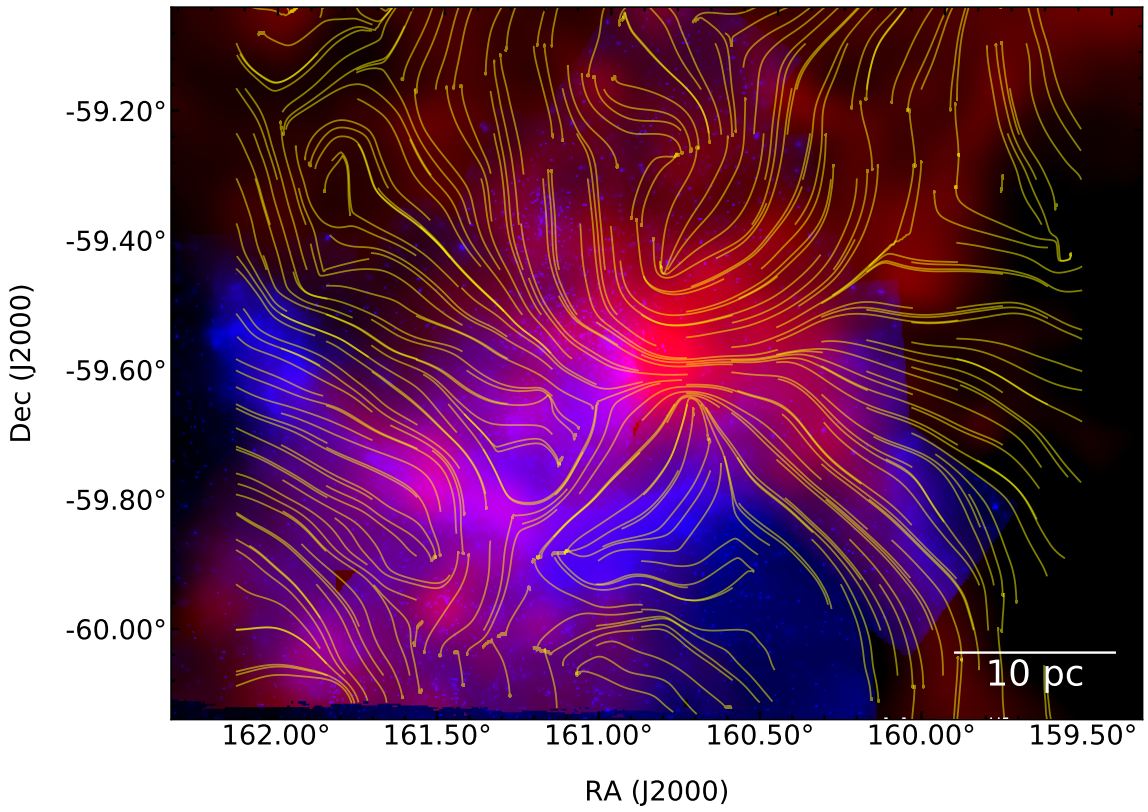

Figure 107. An overlay of the Chandra soft-X-ray CNC map (blue) (Townsley et al., 2011) and BLASTPol intensity map (red). A subset of polarization vector streamlines are overplotted in yellow.

5.4 Estimating $B_{\rm POS}$ using the DCFM

In this section we apply a modified version of the Davis-Chandrasekhar-Fermi Method (DCFM) to the BLASTPol CNC map in order to estimate $B_{\rm POS}$ along various sightlines through the complex. The DCFM is a technique which was first used by Chandrasekhar and Fermi to estimate $B_{\rm POS}$ in the Orion spiral arm of the Milky Way (Davis Jr (1951); Chandrasekhar and Fermi (1953)). Using their method, they accurately estimated $B_{\rm POS}$ at a level of a few μG. Although the estimates of $B_{\rm POS}$ which are produced by the DCFM are known to carry inherent uncertainties of at least a factor of 2 (see Section 5.4.5), it represents the most practical method of gaining insight into $B_{\rm POS}$, as it requires only a map of polarization vectors, and a map of the LOS velocity dispersion in an optically thin emission line.

5.4.1 Description of the DCFM

A magnetized plasma, such as that found throughout the ISM, supports transerve magnetohydrodynamic (MHD) waves known as Alfvén waves. Alfvén waves are oscillations in the plasma which is frozen into the B-field. They're driven by the restorative force associated with the magnetic tension, $B^2/4\pi$ (B^2/μ_0, in SI units), and propagate with a phase velocity (the Alfvén velocity) of:

$$v_A = \frac{B}{\sqrt{4\pi\rho}} \tag{5.12}$$

where ρ is the mass density of ionized material, and Equation 5.12 is in CGS units. If ρ can be estimated along a sightline of interest where there is a uniform B-field of magnitude B_0, then all that is needed to estimate B_0 is knowledge of v_A.

To estimate v_A using the DCFM relies on the assumption that the gas (ionized and neutral) and dust along the sightline are flux-frozen into B_0, so that any transverse spatial deviation of the material from the direction of B_0 can be attributed to an Alfvén wave. Ultimately, the spatial deviation is quantified using the dispersion of E-field polarization angles (see below). For B_0 oriented along the z-axis, and the LOS being the x-direction, the spatial deviation along the y-axis, is:

$$y(x,t) \propto \cos(k(x - v_A t)) \tag{5.13}$$

This assumption is not completely valid in cases where there are significant dynamical forces at play besides the B-field, such as supernovae shocks and ionization fronts associated with HII regions, or strong turbulence. Here, we must assume that the turbulence is sub-Alfvénic. Despite this shortcoming of the DCFM, it is still useful as a method of obtaining rough upper limits on $B_{\rm POS}$.

Taking Equation 5.13 as valid, v_A and $y(x,t)$ can be written into a wave equation as:

$$v_A^2 \left\langle \frac{dy}{dx} \right\rangle^2 = \left\langle \frac{dy}{dt} \right\rangle^2 \tag{5.14}$$

It then remains to estimate dy/dx and dy/dt. The DCFM approach to doing so is to assume that the total B-field in the region is the sum of both the ordered component B_0 and a turbulent component B_t ($B_{\rm tot} = B_0 + B_t$). The turbulent field represents disorder which is caused by turbulent motions in the gas and dust. If the gas pressure is sub-dominant to magnetic pressure, then following Hildebrand et al. (2009), $B_t \ll B_0$, and $\frac{\langle B_t^2 \rangle^{1/2}}{B_0} \approx \frac{\delta B}{B_0}$.

The ratio $\delta B / B_0$ is proportional to both dy/dx and dy/dt. In the DCFM model, the former term can be empirically determined by measuring the dispersion in E-

field polarization angles S_Φ. The second term is determined by measuring the LOS velocity dispersion σ_v in an optically thin emission line. Substituting S_Φ and σ_v for dy/dx and dy/dt in Equation 5.14, and then solving for B yields the principle DCFM equation:

$$B_{\rm POS} = Q\sqrt{4\pi\rho}\frac{S_\Phi}{\sigma_v} \qquad [G] \tag{5.15}$$

where Q is a scale factor which accounts for uncertainties in the measured quantities, as well as the assumptions upon which the method is based. Q is generally set equal to 0.5, although in this work, we choose $Q = 0.1$ (see Section 5.4.5).

More recent users of the DCFM have made slight modifications to Equation 5.15. In this work we adopt the expression derived by Hildebrand et al. (2009), which is also applied by Crutcher et al. (2004) and Franco and Alves (2015). Parameterizing the polarization angle dispersion as $b \approx \sqrt{2S_\Phi^2}$ (Houde et al., 2009), and calculating the FWHM of the LOS velocity line as $\Delta V_{LOS} = \sqrt{8\ln 2}\sigma_V$, the DCFM equation can be expressed as:

$$B_{\rm POS} = 9.3\left(\frac{2n({\rm H_2})}{\rm cm^{-3}}\right)^{1/2}\left(\frac{\Delta V_{\rm LOS}}{\rm km\ s^{-1}}\right)\left(\frac{b}{1^\circ}\right)^{-1} \qquad [\mu G] \tag{5.16}$$

5.4.2 Determining S_Φ

The dispersion in polarization angles S_Φ is calculated according to the method described in Ade et al. (2015). This method has been applied in a recent study of BLASTPol 2012 map of the Vela-C MC (Fissel et al., 2016). The calculation of S_Φ is as follows.

For every N pixels in the polarization angle map the Φ value at that pixel is com-

pared to that of P pixels which surround it in an annular radius of R pixels. For this work, $N = 3$, $P = 12$ and $R = 15$ (corresponding to a diameter of 5′). For the ith pixel, the first step in the calculation is to compare its polarization angle, $\Phi[i]$, with those of the pixels which surround it at a radius of R pixels. The polarization dispersion between each pair of pixels is:

$$S^2_\Phi[i, R] = (\Phi[i] - \Phi[i + R])^2 \tag{5.17}$$

After calculating $S^2_\Phi[i, R]$ for each pair of pixels, the final polarization dispersion value $S^2_\Phi[i]$ is the average of the pairs:

$$S^2_\Phi[i] = \frac{1}{N} \sum_{i=0}^{N} S^2_\Phi[i] \tag{5.18}$$

Following Fissel et al. (2016), the values of S which are used in the final DCFM calculations are debiased by subtracting the variance of $S^2_\Phi[i]$:

$$S_{\Phi,\mathrm{db}} = \sqrt{S^2_\Phi - \sigma^2_{S_\Phi}} \tag{5.19}$$

A histogram of Φ calculated along ∼165,000 sightlines corresponding to the inner ∼1.25 deg^2 of the CNC is shown in Figure 108, and a histogram of S_Φ for ∼33,000 sightlines within the same region is shown in Figure 108. First and second moments for both quantities are reported in Table 23. We find that in the inner region of the CNC, $\langle S_\Phi \rangle \approx 20$ deg.

In Figure 109, the polarization angles are reported relative to Galactic coordinates, where 90 deg corresponds to B_{POS} parallel to the GP. We find that $\langle \Phi \rangle = 96$ deg, and that 50% of the B-field polarization angles are within ± 23 deg of the GP. This finding agrees with Li et al. (2006), which reports $\langle \Phi \rangle \leq 15$ deg of the GP for three GMCs, including the CNC.

	N_{sl}	Mean (deg)	σ (deg)
Φ	165,000	96	33
S_ϕ	33,000	20	26

Table 23. First and second moments of the Φ and S_Φ distribution over the inner ~1.25 deg^2 of the CNC. N_{sl} is the number of sightlines reported.

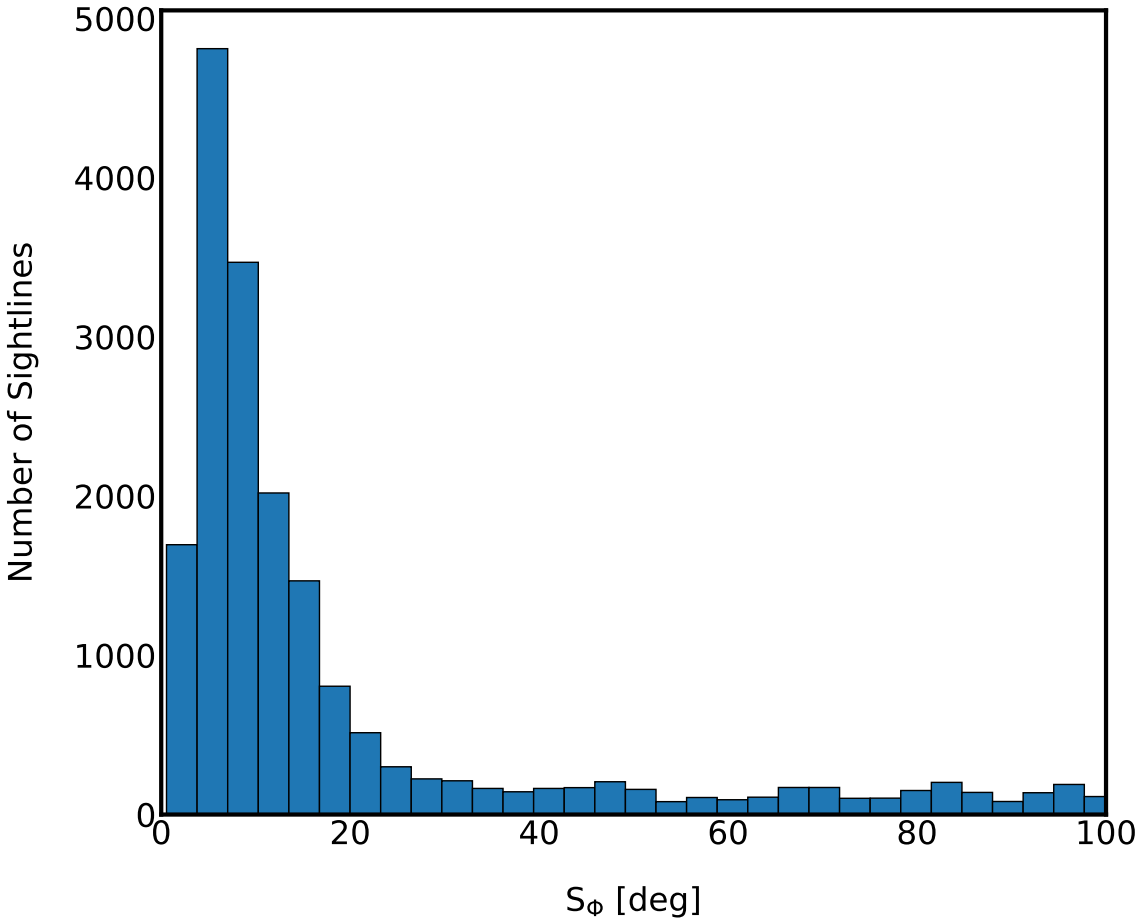

Figure 108. A histogram of S_Φ, the polarization angle dispersion (deg).

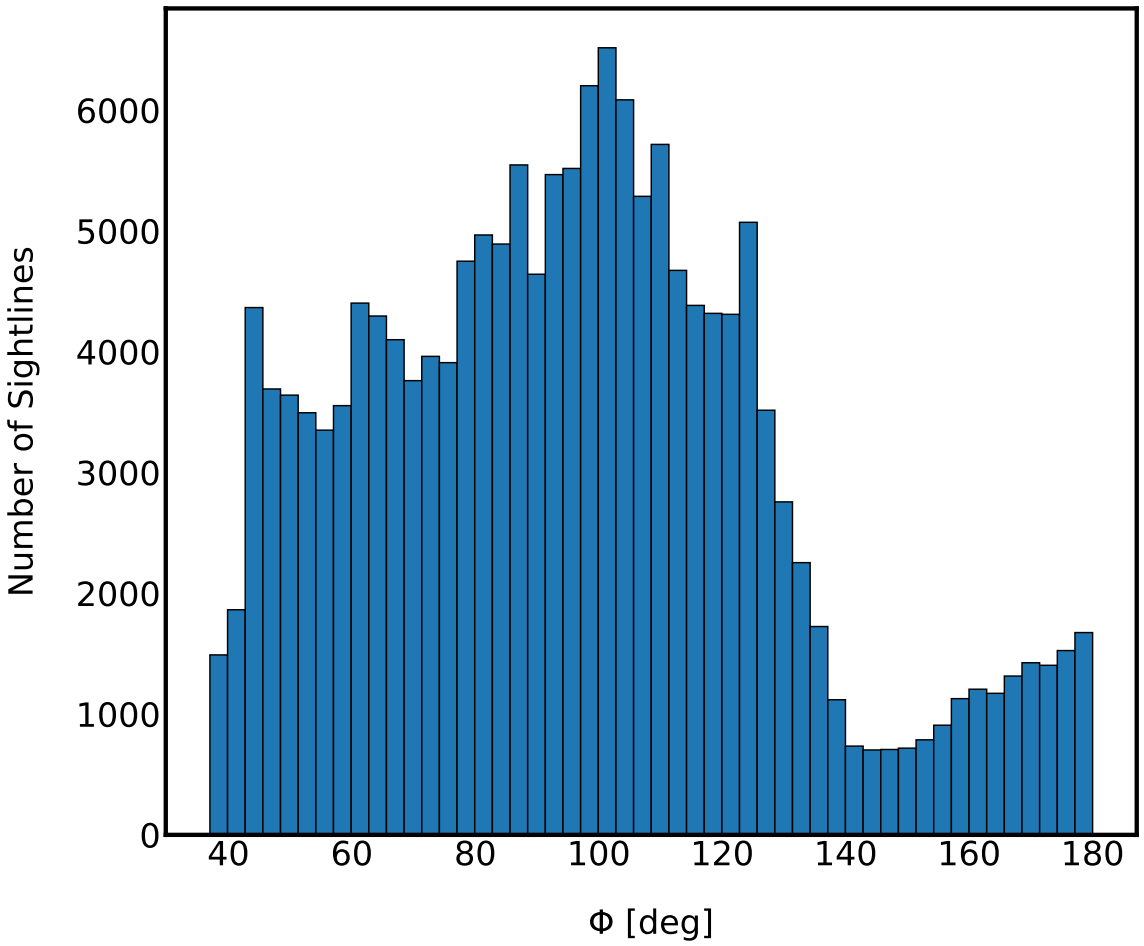

Figure 109. A histogram of Φ (deg), where 90 deg is parallel to the GP.

5.4.3 Determining $n(H_2)$

To determine the number density of molecular hydrogen $n(H_2)$ which is needed to compute B_{POS} using Equation 5.16, we first calculate the column density of molecular hydrogen $N(H_2)$. The column density can be computed from the total intensity maps at each of the three BLASTPol bands as:

$$N(H_2) = 2\frac{\tau_\nu R}{\kappa_\nu m_{H_2}} \qquad \left[\mathrm{cm}^{-2}\right] \tag{5.20}$$

where:

τ_ν is the optical depth at frequency ν

$R = 100$ is the gas-to-dust mass ratio

κ_ν is the dust opacity at frequency ν [cm^2/g]

$m_{H_2} = 3.32 \times 10^{-24}$ is the mass of molecular hydrogen [g]

The FIR dust opacity κ_ν is calculated using the empirical relation found in Mathis (1990):

$$\kappa_\nu = \kappa_\lambda = 13.6(\lambda/250\mu\mathrm{m})^{-2} \qquad \left[\mathrm{cm}^2/\mathrm{g}\right] \tag{5.21}$$

and the optical depth at frequency ν is:

$$\tau_\nu = \frac{F_\nu \Omega_{\mathrm{beam}}}{B_\nu(T_d)} \tag{5.22}$$

$$= \frac{I_\nu}{B_\nu(T_d)} \tag{5.23}$$

$$\tag{5.24}$$

where $B_\nu(T_d)$ is the Planck function, and T_d is the dust temperature, taken from the Planck 353 GHz map of the CNC, and re-gridded to the BLASTPol map coordinates.

Once $N(\mathrm{H_2})$ is known, the total visual extinction A_V can be calculated using the canonical relation described in Bohlin (1978):

$$A_V = \frac{9.4 \times 10^{20}}{N(\mathrm{H_2})} \qquad [\mathrm{mag}] \tag{5.25}$$

Mean values for optical depth τ_ν, dust opacity κ_ν, molecular hydrogen column density $N(\mathrm{H_2})$, and total visual extinction A_V are listed in Table 24 for the inner

$\sim$1.25 deg^2 of the CNC. The quantities are calculated for each of the three BLASTPol bands. A histogram of $N(H_2)$, taken over the same region, is shown in Figure 110. The values for $N(H_2)$ and A_V agree with those presented in Preibisch et al. (2012). It should be noted that the dust opacities (Equation 5.21) carry an uncertainty of a factor of $\sim$2, which translates into an uncertainty in $N(H_2)$.

	250 μm	350 μm	500 μm
$\langle\tau_\nu\rangle$	2.09×10^{-3}	2.06×10^{-4}	3.1×10^{-3}
κ_ν (cm^2/g)	13.16	4.06	3.29
$\langle N(H_2)\rangle$ (cm^{-2})	2.09×10^{21}	4.5×10^{21}	2.88×10^{21}
$\langle A_V\rangle$ (mag)	3.03	2.06	1.03

Table 24. The mean values of optical depth τ_ν, dust opacity κ_ν, molecular hydrogen column density $N(H_2)$ and total visual extinction A_V calculated with the I_{250}, I_{350} and I_{500} maps for the inner $\sim$1.25 deg^2 of the CNC.

To calculate $n(H_2)$, we must assume a column depth through the CNC. The complex morphology of the region makes it infeasible to assign a single column depth to the entire map. Instead, we examine a range of $n(H_2)$, which extends over an order of magnitude: $0.5 \text{ pc} \leq L_{\text{CNC}} \leq 5 \text{ pc}$. The mean and maximum values of $n(H_2)$ corresponding to these two limits are listed in Table 25, for the each of the three BLASTPol bands. A histogram of the $n(H_2)$ values corresponding to $L_{\text{CNC}} = 5$ pc is shown in Figure 111.

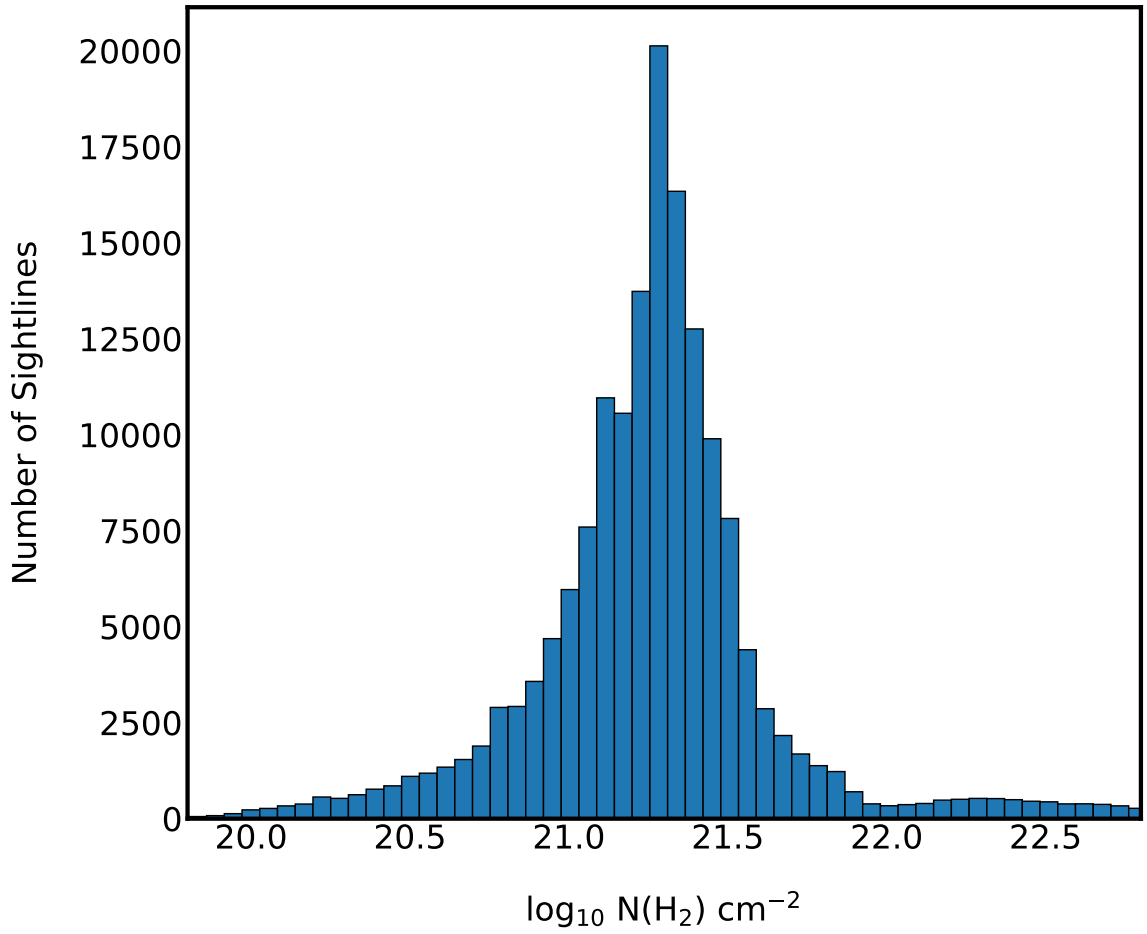

Figure 110. A histogram of $N(H_2)$ over the inner ∼1.25 deg^2 of the CNC.

	250 μm	350 μm	500 μm
$\langle n(H_2)\rangle$ (cm^{-3})	93–930	94–941	133–1330
max $n(H_2)$ (cm^{-3})	4,837–48,370	3645–36,450	4,008–40,080

Table 25. Mean and maximum values of $n(H_2)$ shown for each BLASTPol band. The lower and upper ranges correspond to column depths L_{CNC} of 0.5 pc and 5 pc.

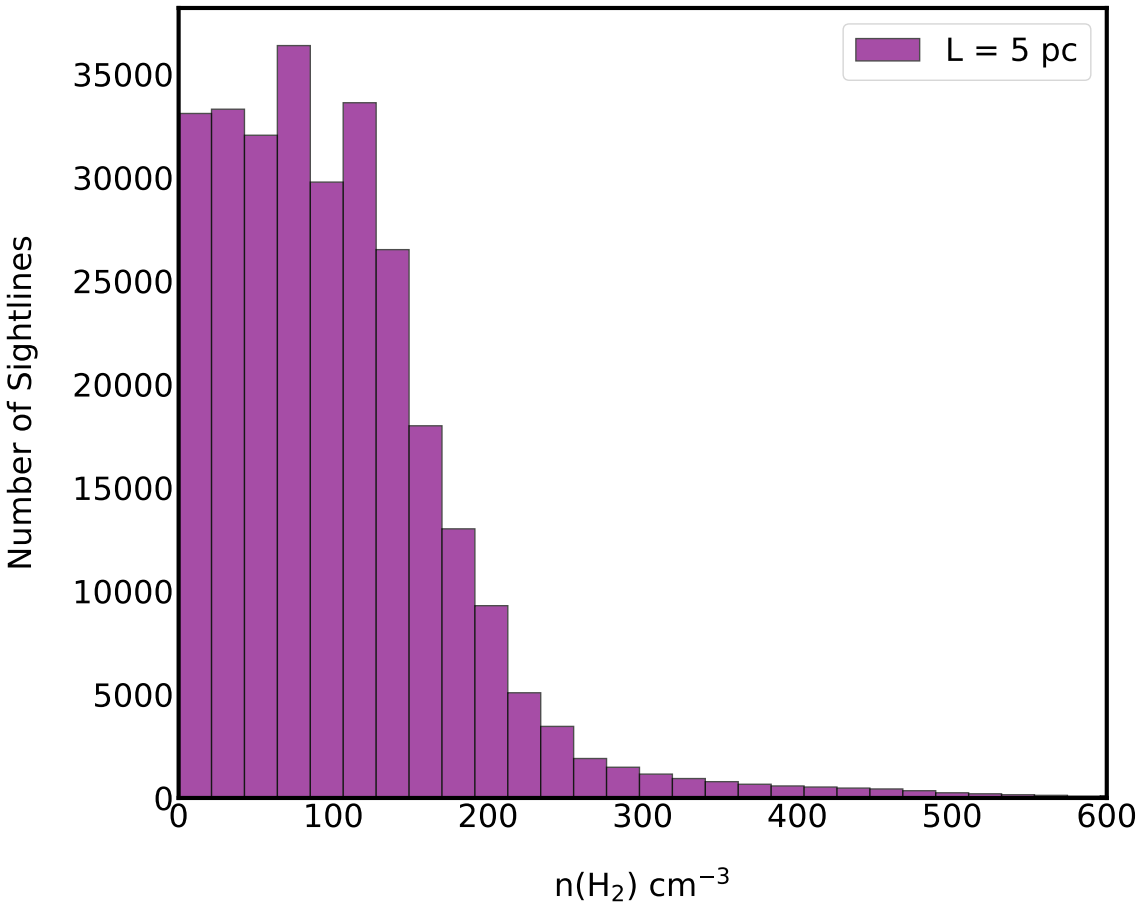

Figure 111. A histogram of $n(H_2)$ in the inner ~1.25 deg^2 of the CNC, shown for an assumed column depth of 5 pc.

5.4.4 Determining σ_v

To estimate $B_{\rm POS}$ using Equation 5.16 requires velocity dispersion values from an optically thin line which are sampled over the BLASTPol CNC field. It is also important that the spatial resolution of the velocity map is at least as high as the BLASTPol map. In this work, we use the HI 21 cm data cube of the CNC region acquired with the Australia Telescope Compact Array (ATCA) (Rebolledo et al., 2017). The spatial coverage of the map contains the entire BLASTPol CNC region, and is taken at a

resolution of $\sim$35″. In the following analysis, the HI data cube is smoothed to 5′ to match the resolution of the Planck maps. The ATCA observations probe a velocity range of $\sim$200 km/s, with a velocity resolution of $\sim$500 m/s.

After smoothing each slice of the ATCA data cube to 5′, the ROI is aligned with the BLASTPol map by converting its coordinates from the galactic to the equatorial frame. The resulting overlay is shown in Figure 112. Once aligned, the average HI line profile along the velocity axis of the cube (i.e., into the plane of Figure 112) is averaged around each map location which was used in the calculation of S_Φ (Section 5.4.2). A subset of these regions is shown in Figure 112, highlighted in yellow.

Before calculating the velocity dispersion σ_v from each average line-profile, the profiles are low-pass filtered, and profiles containing spurious noise are removed from the set. The dispersion of each line profile is estimated by applying a curve fitting algorithm which produces Maximum Likelihood Estimates (MLEs) for the FWHM of a Gaussian profile. A histogram showing the distribution of $v_{\rm FWHM}$ for $\sim$30,000 lines of sight is shown in Figure 113. The distribution of $v_{\rm FWHM}$ is roughly Gaussian, with a mean and standard deviation of $\sim$38 km/s and $\sim$5 km/s. A map of $v_{\rm FWHM}$ over the CNC region used in the DCFM analysis is shown in Figure 114, with $B_{\rm POS}$ pseudovectors (Downsampled by 20x) overplotted in white.

5.4.5 Uncertainties Associated with the DCFM

For any object, the values of $B_{\rm POS}$ which are produced using Equation 5.16 carry uncertainties which are important to address. Firstly, the DCFM relies on the as-

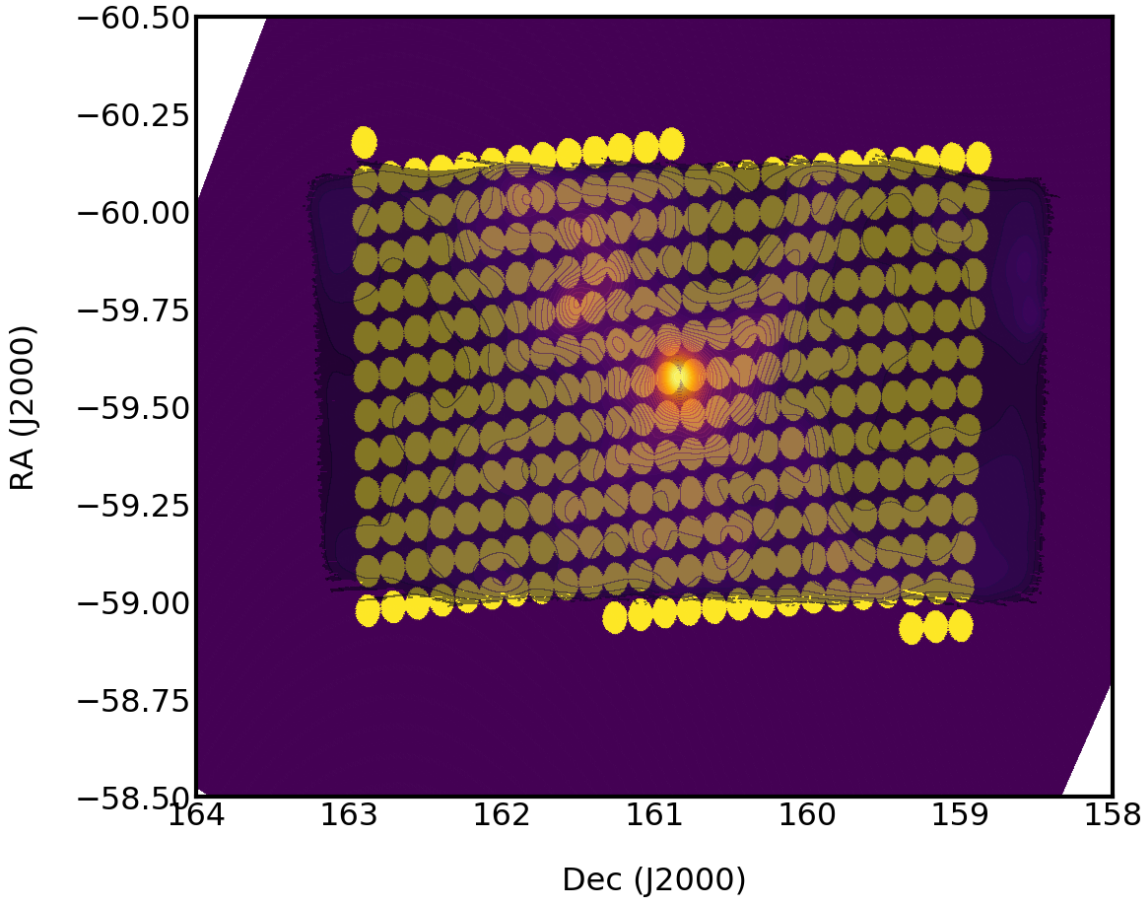

Figure 112. An overlay of one slice of the ATCA HI velocity cube with the BLASTPol intensity map. A subset of the regions used to calculate average HI line profiles are highlighted in yellow. The diameter of the regions is ∼5′.

sumption that flux-freezing applies in the ROI, and that gravity, turbulence and the B-field together account for the dominant energy densities. For the case of a MC, where the DCFM is typically applied, this assumption may hold true to the extent that any uncertainties it introduces into $B_{\rm POS}$ are sub-dominant to other sources of error. Relative to a single MC, the CNC is complex, containing many HII regions and powerful sources of energy injection such as the ongoing star formation in Tr 16 and Tr 14. These dynamical processes, along with the aforementioned ones, have conspired together to produce the complex morphology of the CNC that we now

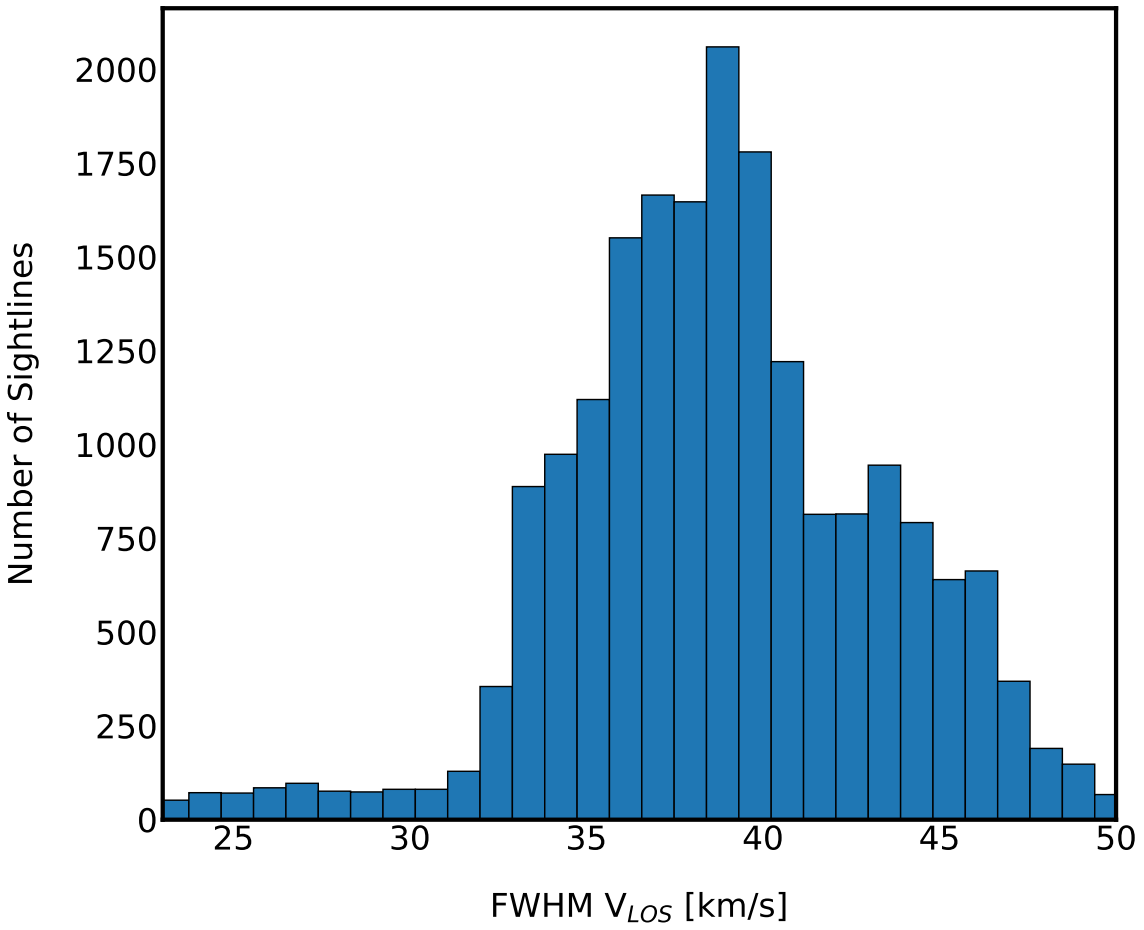

Figure 113. A histogram of the FWHM velocity line widths (in km/s) computed over the inner ∼1.25 deg^2 of the BLASTPol CNC map.

observe. Given this disclaimer, the estimates of $B_{\rm POS}$ provided in this work should be taken as preliminary.

The second type of uncertainty carried by the DCFM pertains to S_Φ and σ_v. Because the velocity dispersion is only calculated along the LOS, the method implicitly assumes that the turbulence is isotropic, which is not necessarily the case. Additional uncertainty, represented by the scale factor Q in Equations 5.15 and 5.16, pertain to the relation of the spatial resolution of the polarization and velocity maps to the driving scale of the turbulence in each ROI. Several studies have performed MHD sim-

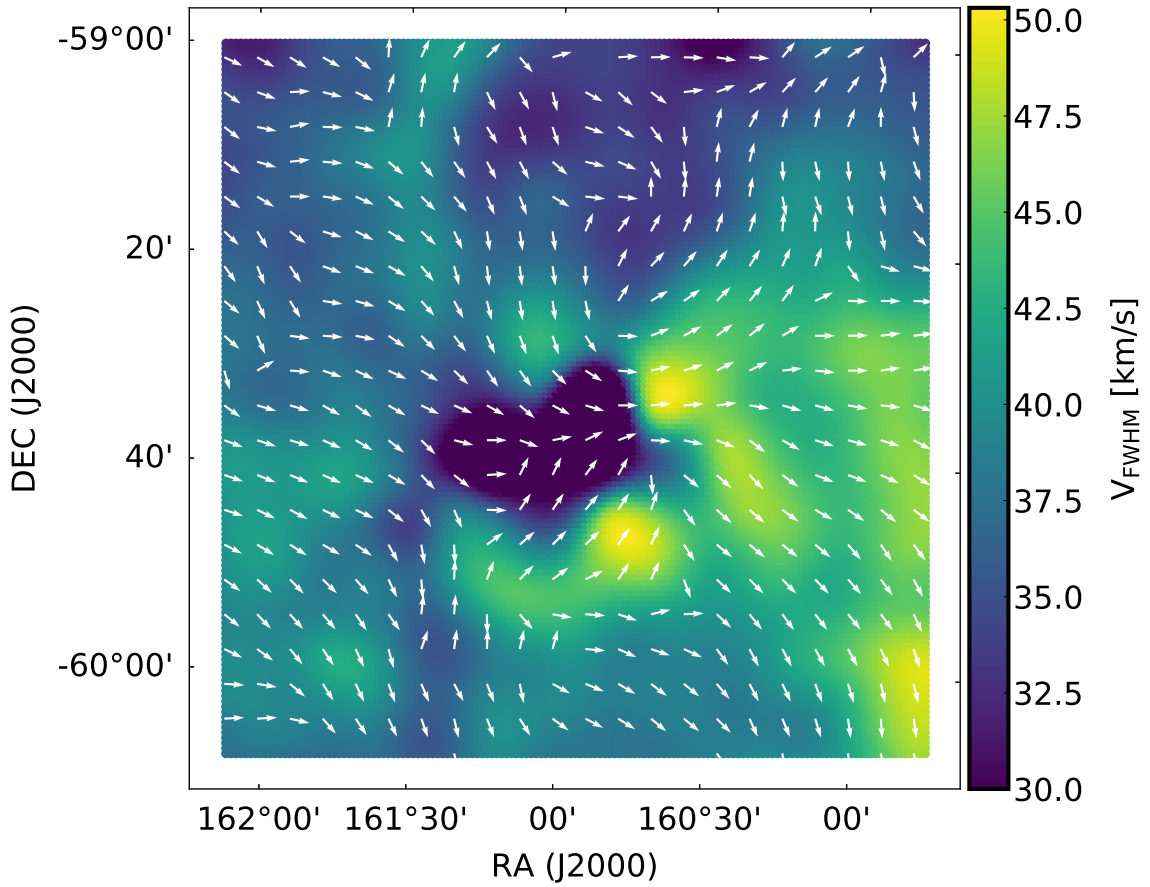

Figure 114. A map of the FWHM velocity line widths (in km/s) computed from the ATCA HI data cube. A subset of the B_{POS} pseudovectors (downsampled by 20x) is overplotted in white.

ulations of the turbulent environment commonly found in MCs in order to estimate Q, and have shown that the B-field estimates produced by the DCFM are sensitive to the level of spatial smoothing which is applied to the maps (see, e.g., Ostriker et al. (2001); Padoan et al. (2001); Heitsch et al. (2001); Kudoh and Basu (2003)).

If S_Φ is produced by averaging over a region which is larger than the driving scale of turbulence, then S_Φ will be underestimated, and B_{POS} will be be overestimated (Cho and Yoo (2016)). The studies referenced above have shown that Q can range from 0.1–0.7, where the lower limit applies to heavy smoothing, and in the case of a

strong B-field with minimal smoothing, Q is usually taken to be 0.5 (Houde, 2004). Further uncertainties in S_Φ and σ_v are added due to smoothing over inhomogeneities in the gas and dust structures, and confusion between LOS and POS features. In this work, the maps have been smoothed to a physical scale of $\sim$3.5 pc, which is likely to be at least an order of magnitude above the driving scale of turbulence for filaments within the CNC, which are of $\mathcal{O}(0.1)$ pc wide (Preibisch et al., 2012). We therefore adopt a conservative value of 0.1 for Q, which reduces the Q value of 9.3 in Equation 5.16 to $Q = 1.86$.

In this work, uncertainty in $n(\mathrm{H_2})$ is likely to be the dominant source of error. $N(\mathrm{H_2})$, which is calculated using the Planck A_V map and canonical relation found in Bohlin (1978), carries an uncertainty which is sub-dominant to the other measured parameters which factor into the DCFM estimates. However, in order to calculate $n(\mathrm{H_2})$, we have had to assume a column depth for the cloud. The variety of physical configurations present in the ROIs discussed in Section 5.3 require the use of a range of column depths, L_{CNC}. For simplicity's sake, the results shown in the following are produced using $L_{\mathrm{CNC}} = 5$ pc.

5.4.6 Results of DCFM Applied to the CNC

The B_{POS} maps of the CNC which have been produced using the DCFM include 19,000 sightlines along the inner $\sim$1.25 deg^2 of the BLASTPol map (they are decimated by a factor of 3 relative to the Stokes maps). Table 26 lists the estimated values for B_{POS}, $\mathrm{v_{FWHM}}$, $n(\mathrm{H_2})$ and S_Φ for 8 sightlines chosen from the regions defined in Table 22. All results presented in this section were produced using the BLASTPol 500 μm Stokes maps, and assuming a uniform column depth $L_{\mathrm{CNC}} = 5$ pc.

The results are displayed in several figures included in this section: Figure 115 shows the interpolated $B_{\rm POS}$ map with polarization streamlines overplotted in yellow. The ROIs which are defined in Table 22 are outlined in white, and the reference sightlines from Table 26 are shown as red dots (the location of η Car is shown as a green dot). Figure 117 shows the $B_{\rm POS}$-weighted LIC ($\mathrm{LIC}_{500} \times B_{\rm POS}$), with overlays as in Figure 115. Figure 116 shows the raw $B_{\rm POS}$ map with a subset of $B_{\rm POS}$ polarization vectors overplotted in white. The Galactic coordinate grid is overplotted in red to illustrate the alignment between $B_{\rm POS}$ and the GP. Figure 118 shows a histogram of $B_{\rm POS}$ over the inner region of the CNC map.

We find that $\langle B_{\rm POS} \rangle = 96$ μG and $\sigma_B = 100$ μG. The maximum value of $B_{\rm POS}$ is 1.2 mG, and the minimum value is close to zero. An idea of the relationship between cloud morphology and $B_{\rm POS}$ can be gleaned from examining Figures 115, 117 and 116. The highest values of $B_{\rm POS}$ are found in the PDR along the western edge of Tr 14, the lower portion of the NB region (slightly north of Tr 14), around the SP region, and along the eastern edge of the SB. There is a concentrated region of high field strength along the southwestern edge of the SB. The outlines of several voids other than the SB can clearly be seen in the $B_{\rm POS}$ maps. These are particurly visible in Tr 16, where there is an apparent bubble centered at the approximate location of η Car. Others are visible in the NB, and in the region between the SB and SP.

Figure 116 reveals the extent to which various features within the CNC are aligned with the GP. As observed by Li et al. (2006), the central region of the CNC is in close alignment. The western edge of the SB is largely orthogonal to the GP, as is the region of the SP which contains the highest values of $B_{\rm POS}$. In Section 5.4.7, we discuss these findings in the context of other studies of B-field morphology inside GMCs and HII regions.

RA, DEC (deg J2000)	$B_{\rm POS}$ (μG)	$v_{\rm FWHM}$ (km/s)	$n(H_2)$ (cm^{-3})	S_Φ (deg)	Region
160.37, -59.78	563.85	44.68	94.59	1.43	1
160.56, -59.42	199.25	43.36	271.16	6.66	4
160.61, -60.08	487.80	38.53	37.35	0.90	1
160.70, -59.59	889.48	44.74	2596.09	4.77	CTR[1]
160.87, -59.75	181.29	43.66	111.54	4.73	1
160.95, -59.55	149.98	24.87	1550.56	12.15	2
161.40, -59.73	121.11	35.87	314.25	9.77	3
161.76, -60.01	297.11	41.74	401.03	5.23	5

[1] A sightline near the center of the map.

Table 26. Estimates of $B_{\rm POS}$, $v_{\rm FWHM}$, $n(H_2)$ and S_Φ for 8 sight-lines along the CNC, produced using the DCFM. The regions in column 6 correspond to those in Table 22, and are outlined in Figure 117. The sightlines are shown as red dots in Figure 117.

5.4.7 Interpretation of $B_{\rm POS}$

The magnitude of the estimates for $B_{\rm POS}$ in the CNC produced using the DCFM is comparable to estimates which have been made in other GMCs using the same method. Franco and Alves (2015) applied the DCFM to the Lupus I MC, and found that at large scales $B_{\rm POS}$ is of $\mathcal{O}(10^2)$ μG. In Soler et al. (2018), the authors applied the DCFM to Planck maps of the Orion-Eridanus superbubble, providing estimates of $B_{\rm POS}$ in the southern edge of the bubble of tens of μG, which were compared to estimates of $B_{\rm LOS}$ obtained from Zeeman splitting of HI. The relationship between $B_{\rm POS}$ and the material at sub-filament scales ($\lesssim$0.1 pc) is of particular interest. MHD simulations of magnetized filaments have predict that there is a critical density, such that $B_{\rm POS}$ is perpendicular to regions of the filament which are above it, and parallel to regions which are below it (Soler et al., 2013). The predictions of these simulations

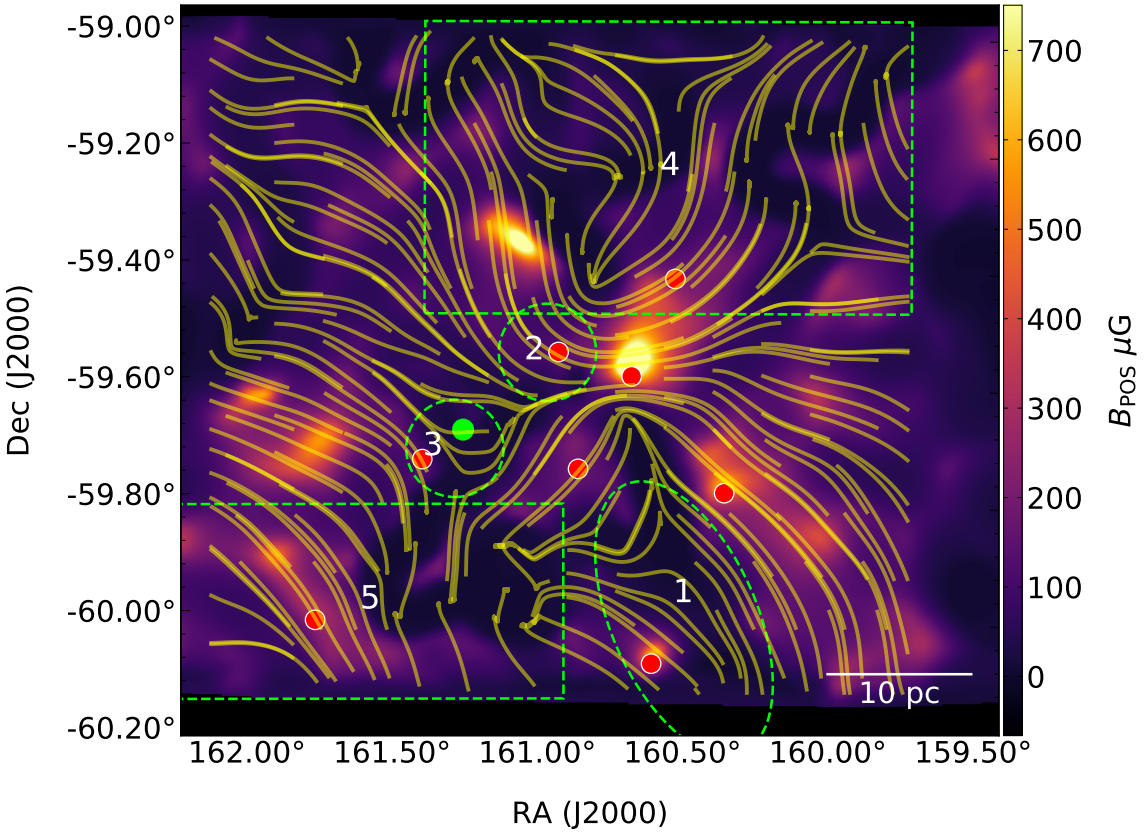

Figure 115. The $B_{POS,500}$ map (every 3rd pixel interpolated) with polarization streamlines overplotted in yellow. The outlined ROIs are defined in Table 22. Reference sightlines listed in Table 26 are shown as red dots.

have since been verified observationally (see, e.g., Ade et al. (2016); Soler et al. (2017); Fissel et al. (2018)).

The maps of the CNC used in this analysis do not possess the requisite spatial resolution needed to probe B_{POS} on the scale of sub-parsec filaments (especially after they have been smoothed to the Planck 353 GHz resolution of ∼5′). However, their resolution is sufficient enough to allow for an examination of the B-field morphology on the scale of the prominent features listed in Table 22.

The question of whether or not B-fields play a significant role in determining the morphologies of Galactic HII regions– or if the structure of the B-field is shaped by

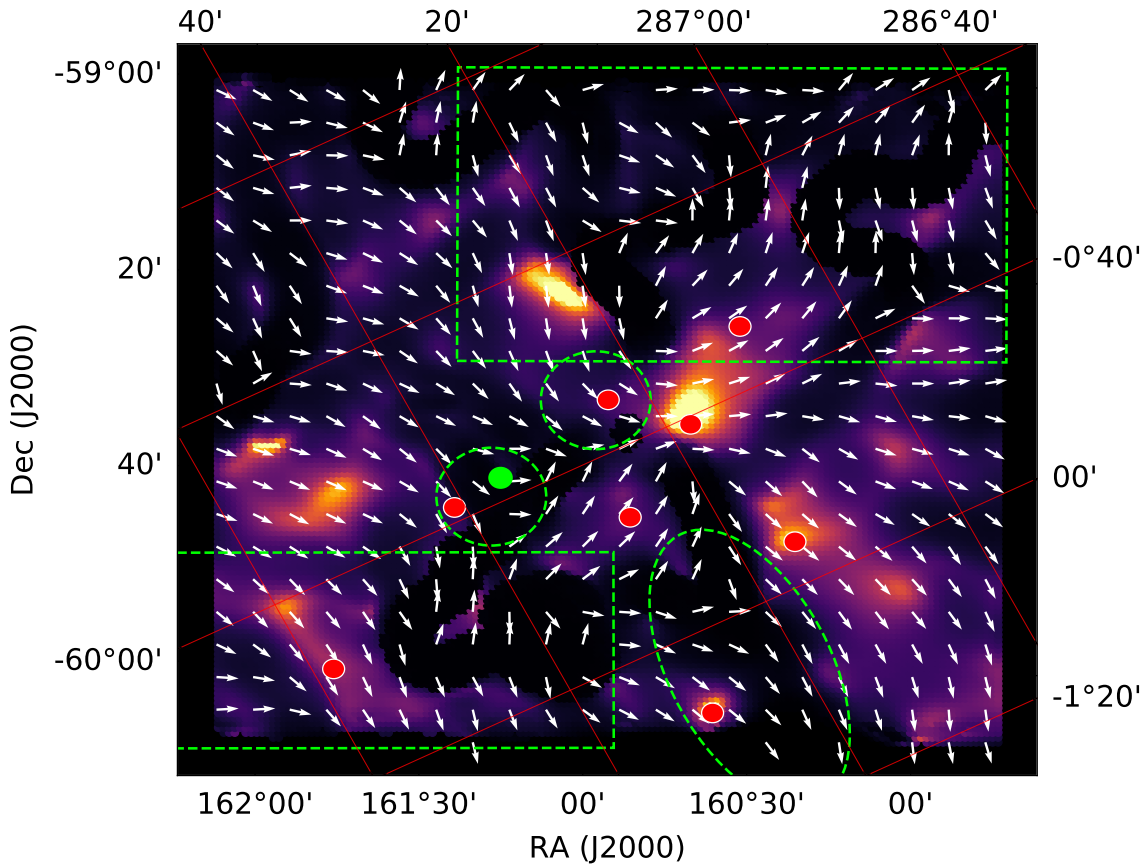

Figure 116. The $B_{\mathrm{POS},500}$ map for $L_{\mathrm{CNC}} = 5$ pc with polarization vectors overplotted in white. The color scale, in units of μG, is identical to that of Figure 115. Both Galactic and Celestial coordinate frames are labeled, with the Galactic coordinate grid shown in red.

the expanding ionization front, is an active area of research (see, e.g., Churchwell et al. (2006); Gendelev and Krumholz (2012); Anderson et al. (2012); Pavel and Clemens (2012); Tremblin et al. (2014); Pellegrini et al. (2007)). As a young HII region expands into the surrounding ISM, the ram pressure of the ionization front encounters resistance from dust and gas which is flux-frozen into the large-scale Galactic B-field. Rather than expand spherically, the shell of the HII region becomes elliptical, with the semi-major axis being parallel to the large-scale direction of the B-field. The B-

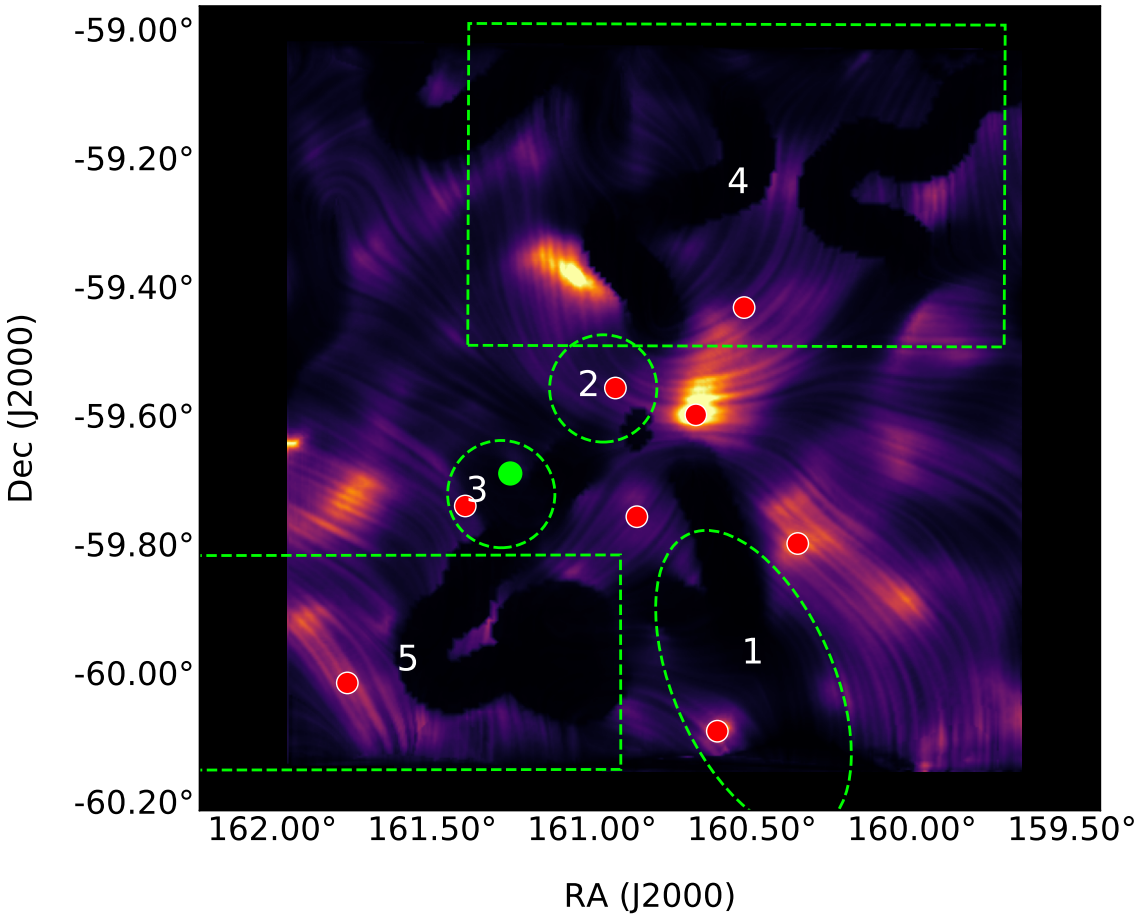

Figure 117. The $B_{\rm POS}$-weighted LIC of the 500 μm map. The outlined ROIs are defined in Table 22. Reference sightlines listed in Table 26 are shown as red dots.

field then becomes amplified along the long axis of the shell, as parallel lines of flux which are tangential to the expansion direction are forced closer together.

Eventually, the B-field dissipates through magnetic diffusion and mass loading, and the ram pressure of the expanding shell overcomes the resistance from the magnetic pressure (Pavel and Clemens, 2012). If measurements of $B_{\rm LOS}$ can be obtained via Zeeman splitting or rotation measures, they can be combined with estimates of $B_{\rm POS}$ to formulate a 3D model of the HII region (see, e.g., Aghanim et al. (2016)).

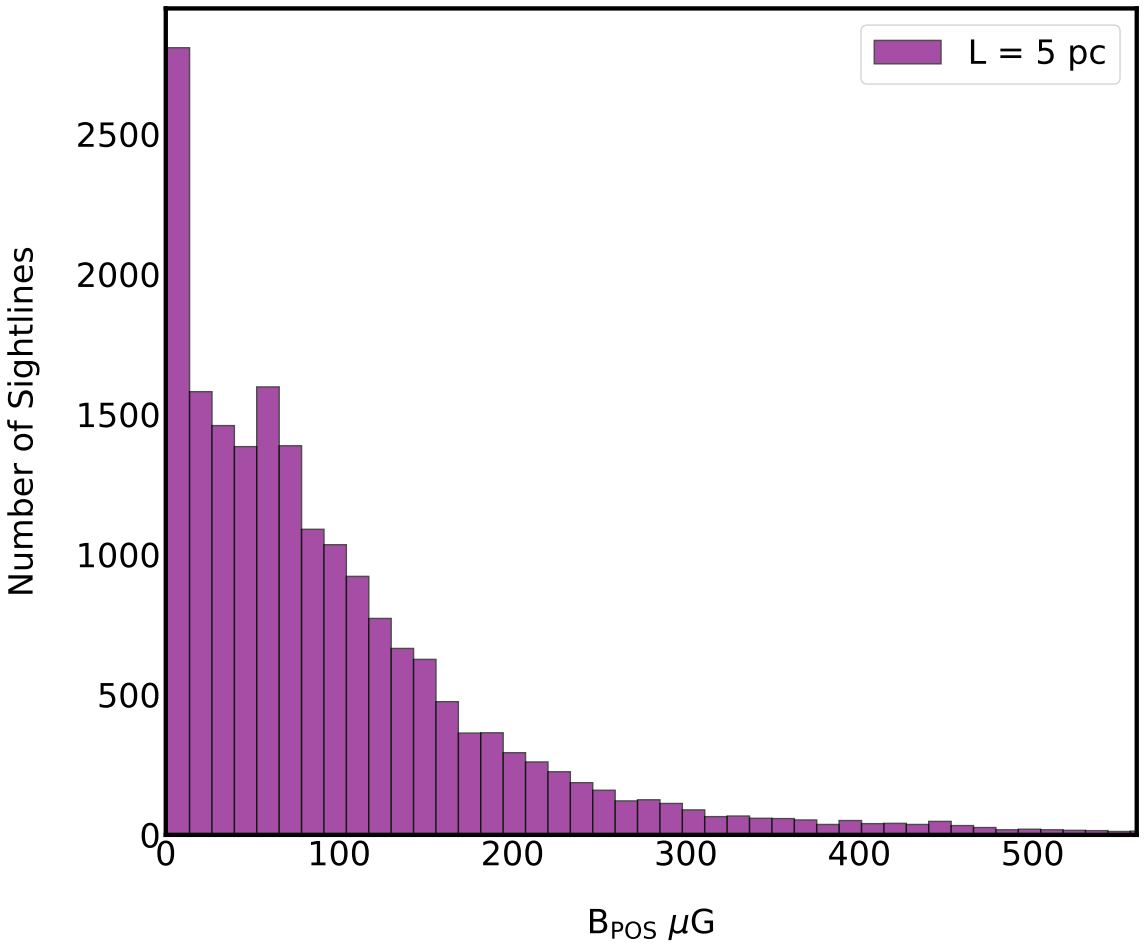

Figure 118. A histogram of $B_{\mathrm{POS},500}$ within the inner $\sim$1.25 deg^2 of the BLASTPol CNC field, for an assumed column depth of 5 pc.

Barring any knowledge of B_{LOS}, a rough lower limit can be assigned to B_{tot} using a simple geometric argument.

Given an HII region with some ellipticity, if all of the ellipticity is assumed to be due to the balance between ram pressure and magnetic pressure, a rough estimate of a lower limit of B_{tot} along the edges of the expansion shell. In doing this analysis, bias effects which are due to the inclination angle of the bubble must be considered, as viewing angles which increase the ellipticity of the bubble will produce higher estimates of B_{tot}. This estimation technique has been applied to ensembles of Galactic

HII regions (Churchwell et al. (2006); Pavel and Clemens (2012)). Following Pavel and Clemens (2012), we examine the SB region of the CNC (region 1 in Figures 115 and 117).

Using Figures 115 and 117 to roughly estimate the spatial extent of the SB, we assume a shell thickness of ∼1 pc. The bubble's fractional thickness $\langle T \rangle$ is calculated as:

$$\langle T \rangle = \frac{1}{2} \frac{\sqrt{R_{\rm out} r_{\rm out}} - \sqrt{R_{\rm in} r_{\rm in}}}{\sqrt{R_{\rm out} r_{\rm out}} + \sqrt{R_{\rm in} r_{\rm in}}} \tag{5.26}$$

We find $\langle T \rangle = 0.44$, which is comparable to values reported in Pavel and Clemens (2012) for other HII regions. The fractional thickness is then used to calculate the ratio between the mass density of the shell and mass density of the ISM ($\rho_{\rm shell}$, $\rho_{\rm ISM}$) into which it is expanding:

$$\frac{\rho_{\rm shell}}{\rho_{\rm ISM}} = \frac{1}{1 - (1 - \langle T \rangle)^3} \tag{5.27}$$

Using Equation 5.27, we find $\rho_{\rm shell} = 1.22 \rho_{\rm ISM}$. For $\rho_{\rm ISM}$, we assume a value corresponding to $n(\mathrm{H_2}) \sim 100$, which is roughly the average number density of molecular hydrogen calculated for the region to the west of the shell. The magnitude of $B_{\rm tot}$ is estimated by equating the ram pressure of the expanding shell to the magnetic pressure of the large-scale field:

$$\begin{aligned} P_{\rm ram} &= P_B \\ \frac{1}{2} \eta \rho_{\rm shell} u_s^2 &= \frac{B^2}{8\pi} \end{aligned} \tag{5.28}$$

where $u_s \approx 10$ km/s is the assumed expansion velocity of the shell, and $\eta = 0.29$ is an efficiency factor calculated in Pavel and Clemens (2012) which accounts for the inclination angle bias mentioned above. Using Equation 5.28 to calculate $B_{\rm tot}$, we

find $B_{\rm tot} = 22$ μG for $\eta = 1$, and $B_{\rm tot} = 12$ μG for $\eta = 0.29$. These values are roughly an order of magnitude below the largest values for $B_{\rm POS}$ estimated for the western edge of the shell using the DCFM (first row of Table 26). However, they are comparable to the mean value of $B_{\rm POS}$, $\langle B_{\rm POS} \rangle = 96$ μG.

At present, the literature on this subject contains evidence which boths supports and refutes the idea that B-fields may play a significant role in the evolution of HII regions. As noted in Tremblin et al. (2014), B-fields of $\mathcal{O}(10^2$–$10^3)$ may have sufficient influence over dense gas at small scales, but will likely not influence the morphology of diffuse nebulae on large scales. Additionally, recent MHD simulations of HII region evolution have found that although magnetic pressure may dominate in the atomic gas component of the region, at large scales the radiation pressure plays a more significant role in shaping the bubble (Rahner et al., 2017). In contrast to the aforementioned findings, in a recent examination of a magnetized PDR in the Galactic HII region M17, Pellegrini et al. (2007) present evidence that the PDR and MC are supported by a B-field with $\langle B \rangle \sim 100$–$600$ μG (inferred from $B_{\rm LOS}$ measurements).

These preliminary estimates of $B_{\rm POS}$ in the innermost regions of the CNC are perhaps calculated at too low of a spatial resolution to draw solid conclusions on the significance, if any, which the B-field plays in shaping the various structures within the cloud. Certainly, at large scales, it is likely that they are subdominant to turbulence and ionizing radiation pressure from the many nascent star clusters embedded within the nebula. Future, higher resolution sub-mm observations of the CNC, such as those which are planned to be undertaken by BLAST-TNG in 2019/2020, will hopefully shed more light on the issue.

Chapter 6

BLAST-TNG TARGET SELECTION: B-FIELD MORPHOLOGY IN NEARBY EXTERNAL GALAXIES

Our understanding of the Milky Way's ISM is greatly augmented by observations of nearby galaxies. Results from the BLAST 2006 balloon flight (2006) revealed that over half of the FIR background is associated with galaxies at $z \geq 1.2$ (Devlin et al., 2009b). These galaxies are spatially unresolved, although estimates of their global properties may be improved by factoring in inferences made from observations of resolved galaxies. Wiebe et al. (2009) attempted to constrain estimates of the star formation rate (SFR) in the aforementioned population of unresolved galaxies by placing an upper limit on the amount of active galactic nuclei (AGN) driven dust heating in a sample of six nearby ($d \leq 20$ Mpc) AGN. This 'AGN fraction' is the observed ratio of core-flux to extended-flux. Using total intensity measurements for six nearby galaxies surveyed by BLAST 2006, the authors were able to derive the dust mass absorption coefficient, and found that the sample of nearby galaxies contains a higher dust mass than would be predicted based on observations of unresolved sub-mm galaxies.

As a sub-mm imaging polarimeter with spatial resolution of $\approx 30''$ at 250 μm, BLAST-TNG will be capable of measuring dozens of $B_{\rm POS}$ pseudovectors for galaxies at distances of around 10 Mpc (a physical scale of $\sim$1 kpc). This mapping resolution is comparable to that which is commonly achieved by single-dish radio observatories operating at wavelengths of a few centimeters. The resulting maps of $B_{\rm POS}$ can be compared to existing observations of $B_{\rm POS}$ in these targets in order to gain

a better understanding of the three-dimensional structure of the large-scale B-field. The total intensity maps will be useful for further analysis of global dust properties within external galaxies.

In the following sections, we describe the selection of five nearby galaxies which will be observed during the 2019/2020 BLAST-TNG flight from McMurdo, Station, Antarctica. They are organized as follows:

- Section 6.1 describes some of the key observational inferences which have been made concerning B-fields in external galaxies.
- Section 6.2 summarizes the galaxy targets which have been selected for observation by BLAST-TNG.
- In Section 6.2.2 we present estimates of the mapping time required for BLAST-TNG to measure sufficient SNRs on the polarization fraction of each target galaxy.

6.1 Summary of Previous Observations

Observations of extragalactic B-fields aim to address the fundamental questions of how the original fields were seeded, amplified, and shaped into their presently observed morphologies. In addition, open questions remain concerning the degree to which large-scale B-fields influence the evolution of galaxies through the role they play in energetic processes, and in governing SFRs. On galaxy scales, the energy density of B-fields is generally assumed to be in equipartition with that of cosmic rays. In most galaxies, including the Milky Way, the average equipartition B-field strength is of $\mathcal{O}(10)$ μG. The field strength is observed to be weaker in radio-quiet galaxies, and stronger in galaxies containing active star formation (Beck, 2016). B-fields are phys-

ically important on a wide range of scales, from the intergalactic to the interstellar medium (IGM, ISM). They are responsible for channeling cosmic rays throughout (and between) galaxies, and play a role in star formation within molecular clouds (MCs). However, astronomical B-fields are difficult to measure. To date, the best tracer of B-fields is polarized emission in wavebands ranging from the radio to the optical. Observations of Zeeman splitting, radio synchrotron emission, Faraday rotation and differential dust extinction (in the O/NIR or sub-mm/mm-wave bands) can reveal the strength of one or more B-field components. However, Zeeman splitting is the only known measurement technique which can directly probe the strength of the field (for $B_{\rm LOS}$).

6.1.1 Radio Measurements of External Galaxies

Over the past few decades, several surveys of polarized synchrotron emission in nearby galaxies have been undertaken in order to map the large-scale fields in a variety of galaxy types, including spirals, ellipticals, dwarf irregulars and interacting pairs (Van Eck et al., 2015). Synchrotron measurements yield estimates of $B_{\rm POS}$. These measurements can be combined with Faraday rotation measures (RM), which trace $B_{\rm POS}$, in order to glean information about the three-dimensional structure of the fields. The large-scale B-field geometries of a few dozen galaxies have been studied in detail (see, e.g, Sofue et al. (1985); Beck et al. (2005); Beck (2016)). While many of the observed extragalactic B-field configurations can be understood as being the result of small and mean-scale dynamo mechanisms (see Section 6.1.2), some morphologies defy classification by existing theories (e.g, M81 Beck (2006)).

These radio observations have led to several key inferences. Perhaps the most

wide-reaching inference is the discovery that the integrated radio continuum emission at centimeter wavelengths is strongly correlated with the degree of sub-mm/FIR emission in galaxies with active star formation (the radio-FIR correlation) (de Jong et al. (1985); Beck (2008)). Because most of the radio emission at centimeter wavelengths is from non-thermal synchrotron radiation, the synchrotron emission (which traces B_{POS}) can be used to estimate the SFR.

The polarized synchrotron emission, which traces the ordered component of galactic B-fields, has been observed to be strongest in the inter-arm regions of spiral galaxies (Beck, 2016). In these galaxies the toroidal field (in the plane of the disk) often forms arms which fill the inter-arm regions. These magnetic arms follow the optically bright spiral arms, but occasionally deviate from the large-scale spiral pattern formed by the gas and dust. Within the material arms, the B-field is turbulent and irregular. The poloidal field (the field which occupies the halo), which is most easily observed in galaxies with high inclination angles (close to edge-on), frequently forms an X shape, which is thought to be generated by bipolar outflows of ionized material (Ferriere and Terral, 2014). Figure 119 shows a map of B_{POS} pseudovectors in M51, obtained from VLA and Effelsberg measurements at 6 cm, overplotted with an optical image from the Hubble Space Telescope (HST) (Fletcher et al., 2011). The inter-arm B-field tracks the spiral pattern of the optical arms, but extends well beyond the visible extent of the gas and dust.

6.1.2 The Modal Structure of Large-Scale Magnetic Fields

Although the mechanism by which the primordial seed fields were created (magnetogenesis) is not well understood (see, e.g., Kandus et al. (2011)), the theory which

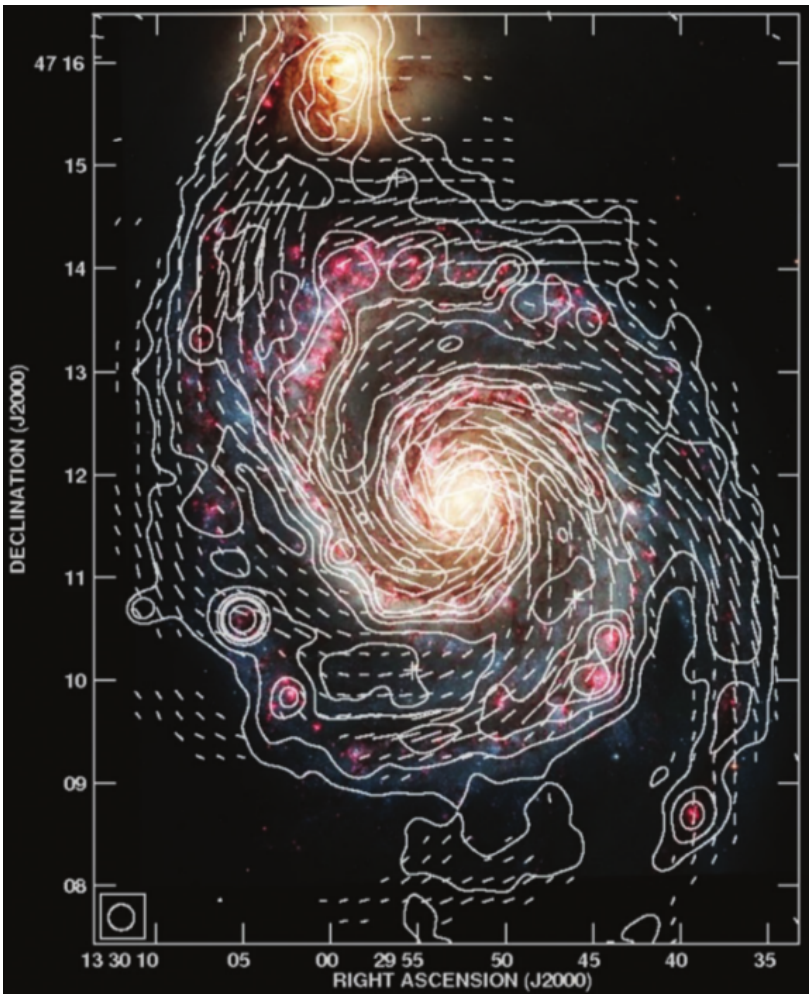

Figure 119. The large-scale B-field in M51 (Φ pseudovectors), revealed by VLA and Effelsberg measurements taken at 6 cm, with a spatial resolution of 6″. The background optical image is from HST (Fletcher et al., 2011).

describes how galactic seed fields are amplified to their observed equipartition values of $\mathcal{O}(10)$ μG has for the most part been well matched by observations (Beck, 2006). The seed fields are thought to have been amplified over several million years by a small-scale dynamo. This dynamo is generated by turbulence created by supernovae and density wave shocks within spiral arms, which shear and compress the B-field. These processes result in a turbulent field geometry, having no regular direction on small scales. The field is then ordered and sustained on large-scales over several billion years by the $\alpha - \Omega$ mechanism, where α refers to the Coriolis Effect, and Ω refers to differential rotation. The mechanism is described by the mean-field dynamo equation:

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (v \times \mathbf{B}) + \nabla \times \alpha \mathbf{B} + \eta \nabla^2 \mathbf{B} \tag{6.1}$$

where $\nabla \times (v \times \mathbf{B})$ represents field amplification, $\nabla \times \alpha \mathbf{B}$ is a gain term from the Coriolis effect, and $\eta \nabla^2 \mathbf{B}$ is a loss term which accounts for the friction of diffusion between ionized and neutral gas. For the dynamo to work, gas flows associated with supernova remnants in a disk expand and rise, causing B-field loops to form via the Coriolis effect. In turn, the field loops generate ion flows that produce new, regular fields loops in a plane perpendicular to the disk. The solutions to the mean-field dynamo equation are azimuthal Fourier modes: $\boldsymbol{B} = \sum_m \beta_m \cos(m\phi - \beta)$, where β is a phase and ϕ is the azimuthal angle in the plane of the disk (Sofue et al., 1985).

The two lowest toroidal modes, axisymmetric ($m = 0$) and bisymmetric ($m = 1$), are accompanied by even (quadrupolar) and odd (dipolar) parity poloidal modes. The lower order modes are easier to excite and amplify than the higher order modes, and most well studied spiral galaxies exhibit a superposition of the two, with axisymmetric being the more dominant mode. Combined radio measurements of polarized syn-

chrotron emission and Faraday RMs can be used to partially decompose the modes (Beck, 2008).

6.1.3 Measurements of Polarized Submillimeter Emission in External Galaxies

Measurements of the sub-mm polarization fraction p in external galaxies can advance theories of magnetic dust-grain alignment. They have also been used to characterize the role that external galaxies may play as a CMB foreground (Seiffert et al., 2006). In contrast to the growing catalog of radio observations of external galaxies, there is a paucity of complementary observations of polarized emission at sub-mm wavelengths. Part of this discrepancy is explained by the relatively small number of sub-mm telescopes which possess the spatial resolution of large, single-dish radio observatories. Consequently, more attempts have been made to map $B_{\rm POS}$ on large-scales using differential dust extinction in the O/NIR (see, e.g., Fendt et al. (1998)).

The first sub-mm mapping of $B_{\rm POS}$ in another galaxy was reported for M82, a nearby starbust galaxy (Greaves et al., 2000). This observation was taken at 850 μm, with the SCUBA instrument (Submillimeter Common User Bolometer Array) on the James Clerk Maxwell Telescope (JCMT). The spatial resolution of the map is $\sim$15″ ($\sim$280 pc). These observations revealed that M82 possesses a large-scale toroidal field which is ordered on scales of at least that of the map resolution. A poloidal field surrounding the nucleus was also noted. Surprisingly, the orientation of the fields as traced by the polarized sub-mm dust emission were found to be unaligned with the galactic plane.

A more recent sub-mm observation of M82 was taken at 53 μm and 154 μm with the HAWC+ instrument (High-resolution Airborne Wideband Camera-plus)

on the Stratospheric Observatory for Infrared Astronomy (SOFIA) (Jones et al., 2019). These observations resolved poalrized emission from thermal dust on scales of ~90 pc. A LIC map of B_{POS} at 154 μm is shown in Figure 120. The LIC map of M82 shows that B_{POS} has a vertical orientation in the inner region of the galaxy, and a planar orientation in the outer regions of the disk. The inner, vertically oriented field corresponds to the region where a starburst is occurring. The authors of the study note that the high amount of beam-averaging prohibits the application of DCFM to these maps.

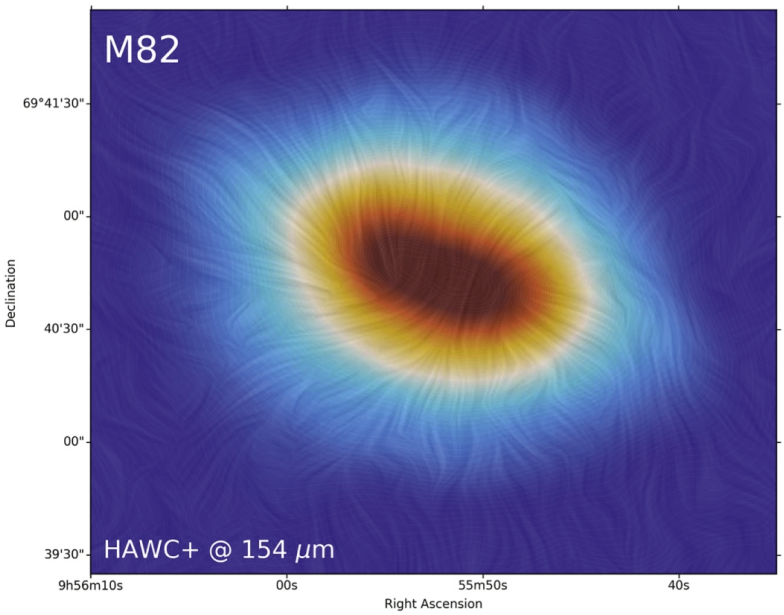

Figure 120. A LIC map of M82, taken with the 154 μm channel of the HAWC+ instrument on SOFIA (Jones et al., 2019).

6.2 BLAST-TNG External Galaxy Targets

The five external galaxies which have been selected for observation by BLAST-TNG were chosen based on several criteria. The primary constraint comes from the telescope's visibility, which will change throughout the flight as the stratospheric circumpolar vortex winds carry the telescope around the Antarctic continent. During the flight, the telescope's latitude might change by up to 10 degrees. The telescope can move in elevation between 22 and 50 degrees, with limits on azimuth imposed by the position of the sun (it should be kept out of view of the primary mirror). Roughly, observable targets must fall within the range of 5 to 18 hours of RA, and -62 to -23 degrees of declination.

The second selection criterion is that the galaxies should be bright enough for BLAST-TNG to be able to measure an SNR on the polarization-fraction of $\gtrsim$3 (see Section 6.2.2). Galaxies which satisfy the first two criteria are deemed more desirable if previous observations of their B-fields exist, particularly those based on measurements of polarized synchrotron emission or Faraday RMs.

6.2.1 Target List

The five galaxy targets, along with their sky coordinates, distance, resolvable spatial scale at 250 μm and type are listed in Table 27.

NGC 4945: NGC 4945 is an edge-on, Seyfert 2 (Sy2) type spiral galaxy which is otherwise thought to be similar to the Milky Way (Spoon et al., 2000). To date, no radio halo has been detected in NGC 4945 (Elmouttie et al., 1997).

Object	RA (J2000)	DEC (J2000)	Distance (kpc)	Scale_{250} (kpc)	Type
NGC 4945	13:05:27.279	-49:28:04.44	3,900	0.57	SB(s)cd
ESO 97-G13	14:13:09.906	-65:20:20.47	4,200	0.61	SA(s)b
NGC 3621	11:18:16.5	-32:48:51	6,700	0.97	SA(s)d
NGC M83	13:37:00.919	-29:51:56.74	8,900	1.29	SA(s)ab
NGC 1808	05:07:42.343	-37:30:46.98	10,900	1.59	(R′)SAB(s)b

Table 27. A summary of the five nearby galaxy targets selected for observation by BLAST-TNG, listed in order of their distance.

ESO 97-G13: ESO 97-G13 (AKA the Circinus Galaxy) is a near edge-on Sy2 galaxy located close to the Galactic plane. Its bipolar radio lobes have been mapped at centimeter wavelengths using ATCA (Elmouttie et al., 1998). The radio lobes are associated with ionized plumes being ejected from the AGN. The B-field in the radio lobes is observed to be aligned with the plume axis.

NGC 3621: NGC 3621 is a near edge-on spiral field galaxy. There is no evidence for a bulge in the galaxy's center, but it is thought to be a Sy2 type (Barth et al., 2008). The HI continuum was recently mapped at 6″ resolution using the Very Large Array (VLA), as part of the THINGS survey of nearby galaxies (Walter et al., 2008).

M83: M83 (NGC 5236, AKA the Southern Pinwheel Galaxy) is a nearby, nearly face-on grand-design spiral galaxy of the starburst variety. It is located near the dwarf galaxy, NGC 5253, which also hosts a high level of active star formation. M83 has been observed in a wide range of wavebands, and the large-scale structure of its B-field has been mapped in polarized synchrotron emission (Sukumar and Allen (1989); Frick et al. (2016)). The B-field is of the inter-arm spiral type. The authors of Frick et al. (2016) compared the radio observations of M83 to sub-mm total intensity maps,

and found that the orientation of one of the main B-field arms is misaligned with its corresponding gaseous arms.

NGC 1808: NGC 1808 is a Sy2 type galaxy which was mapped by BLAST 2006 (Wiebe et al., 2009). Its large-scale B-field has been mapped in optical polarization, which revealed a coherent toroidal magnetic spiral, with no visible poloidal component (Scarrott et al., 1993). Siebenmorgen et al. (2001) detected polarized emission at 170 μm along four sightlines through the galaxy, and found that the polarization angles differ significantly from those measured from nonthermal polarized radio emission (Siebenmorgen et al., 2001).

6.2.2 Required Observing Times

In this section we present estimates of the minimum integration times which are required to make 3-σ measurements of p for each of the galaxy targets described in Section 6.2.1[29].

The minimum beam-averaged intensity (MJy/sr) which is required to achieve 1-σ error bars on p of σ_p, over a map of area $A_{\rm map}$ (deg^2), in t hours, is:

$$I_{\rm req} = I_{\rm ref} \left(\frac{A_{\rm map}}{1\ {\rm deg}^2} \right)^{1/2} \left(\frac{5\ {\rm hr}}{t} \right)^{1/2} \left(\frac{0.005}{\sigma_p} \right) \qquad [{\rm MJy/sr}] \tag{6.2}$$

where $I_{\rm ref}$ (MJy/sr) is the intensity which is required to measure $\sigma_p = 0.5\%$, after a 5 hr integration over a 1 deg^2 map (see Table 28). The required observing time to achieve 1-σ error bars on p of σ_p over a map of area $A_{\rm map}$ (deg^2), for a source with intensity $I_{\rm req}$, is:

[29] A detailed description of the methods presented here can be found at: https://sites.northwestern.edu/blast/

$$t_{\rm req} = t_{\rm ref} \left(\frac{A_{\rm map}}{1 \text{ deg}^2} \right) \left(\frac{0.005}{\sigma_p} \right)^2 \left(\frac{100 \text{ MJy/sr}}{I_{\rm req}} \right)^2 \qquad [\text{hr}] \tag{6.3}$$

where $t_{\rm ref}$ is the time required to obtain $\sigma_p = 0.5\%$ for a source with $I = 100$ MJy/sr (see Table 28).

	250 μm	350 μm	500 μm
$I_{\rm ref}$ (MJy/sr)	188.7	113.7	42.4
$t_{\rm ref}$ (hr)	17.8	6.4	0.9
Beam FWHM (″)	30	41	59

Table 28. Reference values for calculating the required target intensity $I_{\rm req}$ and observation time $t_{\rm req}$ to achieve a desired maximum uncertainty on the polarization-fraction σ_p. The reference intensity $I_{\rm ref}$ (MJy/sr) is the beam-averaged source intensity which is required to measure $\sigma_p = 0.5\%$, after a 5 hr integration over a 1 deg^2 map. The reference integration time $t_{\rm ref}$ (hr) is the time required to obtain $\sigma_p = 0.5\%$ for a source with an intensity of 100 MJy/sr. Table values are taken from https://sites.northwestern.edu/blast/. The beam FWHM for each observation band (in arcsec) is listed in the third row.

The smallest maps that BLAST-TNG will make will be ∼0.5 deg^2. Because the galaxy targets fit into this area, their map size will be uniformly ∼0.5 deg^2. Table 29 lists the required integration time $t_{\rm req}$, in hours, needed to achieve $\sigma_p = 0.5\%$ over the region of each galaxy which is visible in the the Herschel 250 μm total intensity maps. Three of the targets require at least 7 hr of integration time. The brightest targets, NGC 4945 and ESO 97-G13, require at least 3 hr.

Figures 121 to 125 show the Herschel 250 μm maps for each galaxy target. The colorscales are (left frame) $t_{\rm req}$ in hours, and (right frame) intensity in MJy/sr.

Object	A_{map} (deg^2)	σ_p (%)	t_{req} (hr)	I_{250} center (MJy/sr)
NGC 4945	0.5	0.5	3	10,000
ESO 97-G13	0.5	0.5	3	3,600
NGC 3621	0.5	0.5	7	150
M83	0.5	0.5	7	1,000
NGC 1808	0.5	0.5	7	2,500

Table 29. Map area A_{map}, 1-σ polarization-fraction σ_p, integration time t_{req}, and central beam-averaged intensity from Herschel 250 μm maps for each BLAST-TNG galaxy target

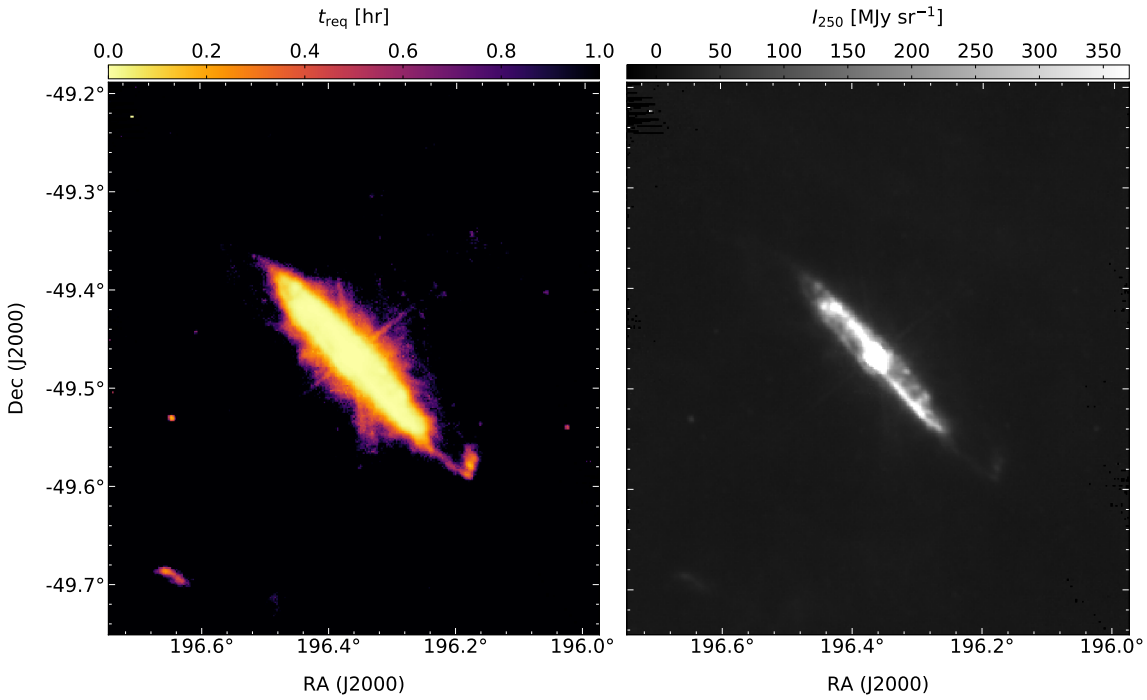

Figure 121. NGC 4945: Required mapping time (left) and 250 μm intensity, from Herschel SPIRE.

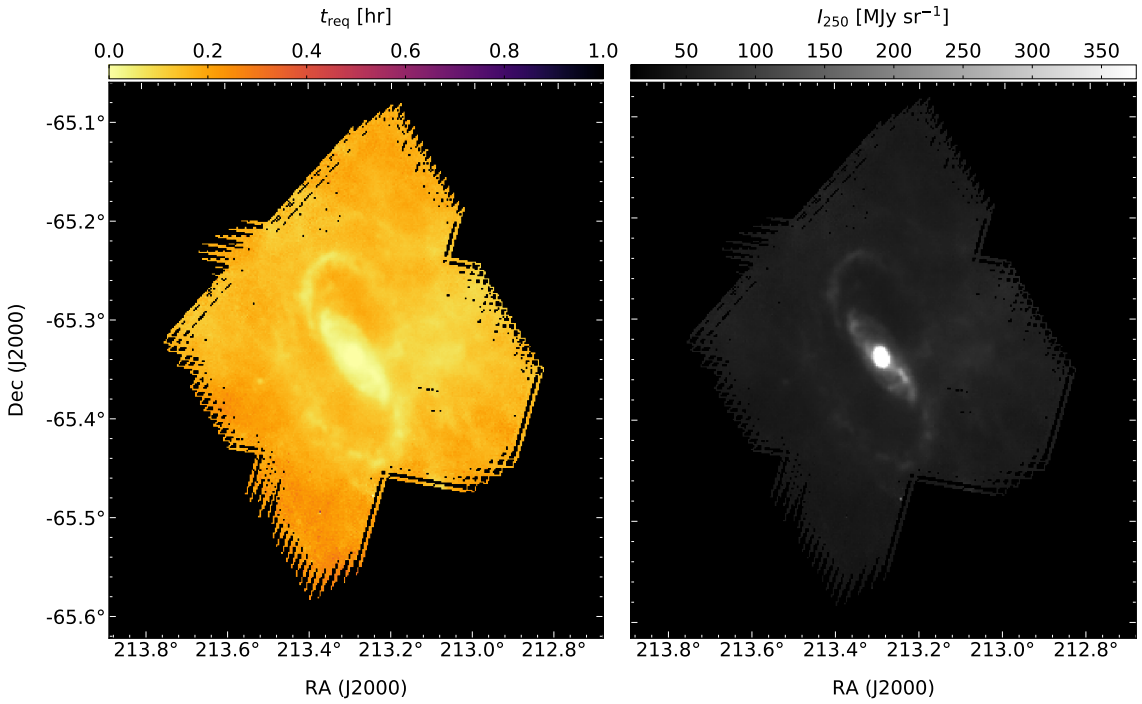

Figure 122. ESO 97-G13: Required mapping time (left) and 250 μm intensity, from Herschel SPIRE.

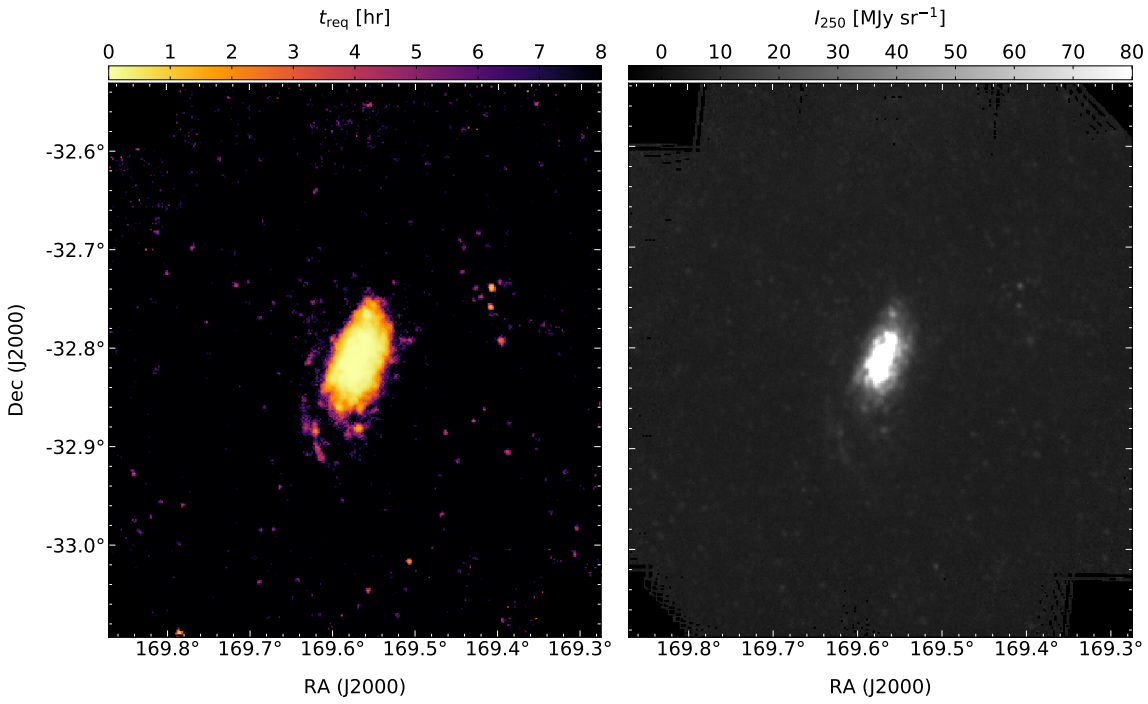

Figure 123. NGC 3621: Required mapping time (left) and 250 μm intensity, from Herschel SPIRE.

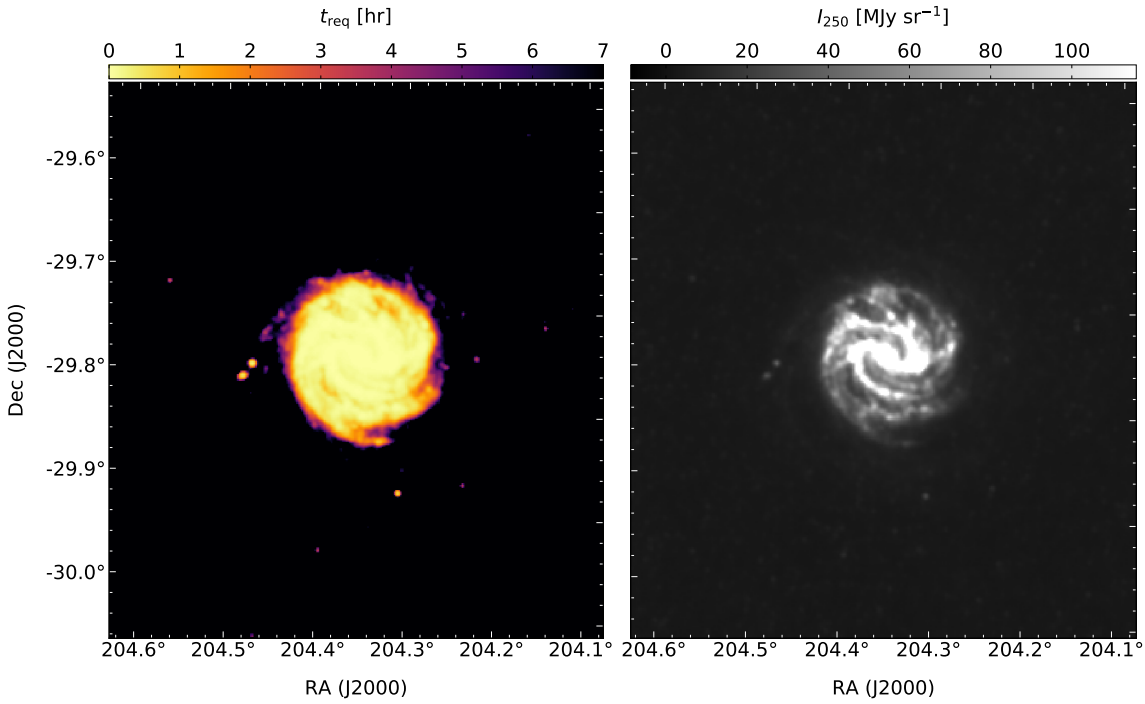

Figure 124. M83: Required mapping time (left) and 250 μm intensity, from Herschel SPIRE.

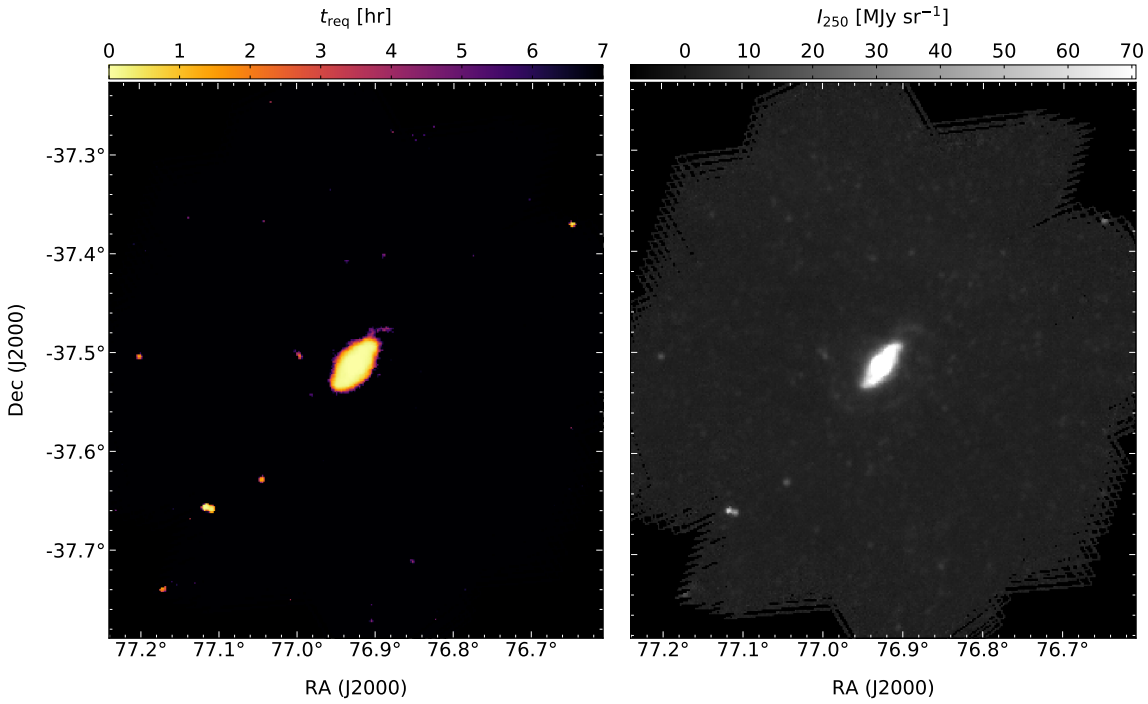

Figure 125. NGC 1808: Required mapping time (left) and 250 μm intensity, from Herschel SPIRE.

Chapter 7

CONCLUSION

7.1 Science-Driven Technological Development

The ambitious science goals of current and next generation sub-mm/FIR/mm-wave observatories drives the development of new technologies at a rapid pace. Upcoming ground-based observatories which will probe the inflationary epoch through measurements of B-mode polarization in the CMB (e.g., Simons Observatory (Ade et al., 2019), AliCPT (Li et al., 2018), CMB-S4 (Abitbol et al., 2017)) require detector multiplexing factors of $\mathcal{O}(10^5 - 10^6)$. These multiplexing factors exceed the capabilities of preexisting DSP boards, such as the CASPER ROACH2, by an order of magnitude.

From a SWaP-C standpoint, the need for higher multiplexing factors cannot be met by simply scaling up current systems. The current per-pixel cost of MKID readout systems with multiplexing factors of $\sim 10^3$ (e.g., BLAST-TNG, TolTEC and DARKNESS) is $\sim$1 USD/pixel. Additionally, the increase in size, weight and power dissipation of the scaled-up systems would impose thermal and mass budgets which prohibit their use on planned space-based sub-mm/FIR surveyors (e.g., LiteBIRD (Matsumura et al., 2014), Origins Space Telescope (Battersby et al., 2018)).

The solutions to the limitations inherent in current systems will incorporate novel electronics and DSP algorithms developed by scientific collaborations, as well as new technologies which are emerging from the commercial sector. The scaling down of hardware and electronics includes both the digital and analog (IF/RF) sets of

electronics. This process has already begun for the systems which will be used by upcoming sub-mm balloon experiments (e.g., EXCLAIM and TIM).

Figure 126 shows one of the TolTEC IF electronics slices, which was assembled at ASU. It is based on the BLAST-TNG IF electronics, combining COTS components, such as the modulator and demodulator, with custom components designed at ASU, such as the anti-aliasing filters and programmable attenuators. Figure 127 shows an equivalent system (also designed at ASU) which has been scaled down to a single multi-layer PCB which is many times smaller than the baseline IF system. Single-board systems like this offer significant SWaP-C advantages.

The SWaP-C of the digital electronics is also undergoing rapid improvements. One example of this is the Xilinx RF System on a Chip[30] (RFSoC) family of integrated RF-FPGA systems. The RFSoC, which is based on the Zynq UltraScale+ FPGA, incorporates up to 16 14-bit DACs and 16 12-bit ADCs on a single chip, and has an instantaneous RF bandwidth of several gigahertz. The expanded signal bandwidth is accompanied by SWaP-C improvements of several orders of magnitude, which directly translate into lower cost-per-pixel. Efforts to port existing MKID and TES FPGA firmware designs, including the BLAST-TNG firmware, are already underway.

The larger FPGA fabric provided by these new commercial boards allows for the implementation of improved detector readout algorithms for both MKIDs and TESs which up until now have only been possible with the use of custom electronics. One example of this is the 'tone-tracking' algorithm which is used by the SLAC Microresonator Radio Frequency (SMuRF) readout (Henderson et al., 2018). SMuRF is the readout system for Simons Observatory, which uses a detector technology called microwave frequency domain multiplexing (μMUX, Irwin and Lehnert (2004)).

[30]https://www.xilinx.com

μMUX systems use TES sensors which are inductively coupled to RF SQUIDs that form part of a resonator circuit. In this way, they can be frequency domain multiplexed and readout using FDM algorithms similar to those which are used for KID readout. Tone-tracking refers to the process of making high speed changes to the amplitude and frequency of each probe tone in a multitone readout comb in order to compensate for changes in the responsivity of each pixel. The level of difficulty which is involved in implementing tone-tracking depends on the method which is used to calculate the required adjustments, as well as the method which is used to synthesize the probe tones.

In the ASU LEKID readout system, the tone comb is synthesized in software, and then stored in RAM on the ROACH2 board. One shortcoming of this method of tone synthesis is that in order to change either the frequency, amplitude or phase of a single tone in the comb, the entire comb must be resynthesized. Rewriting the tone comb to RAM takes $\sim$3 seconds, making resonator tuning a very slow process.

During the 2018 OLIMPO flight, it took $\sim$1 hr to tune 120 pixels (Masi et al., 2019). Using the CORDIC algorithm (COordinate Rotation DIgital Computer) for tone synthesis, the parameters of each probe tone can be tweaked independently of all of the others, at a very high rate. This opens up the possibility of trying many different tone-tracking algorithms, which are not possible to implement on a ROACH2 system due to the limited size of the FPGA fabric.

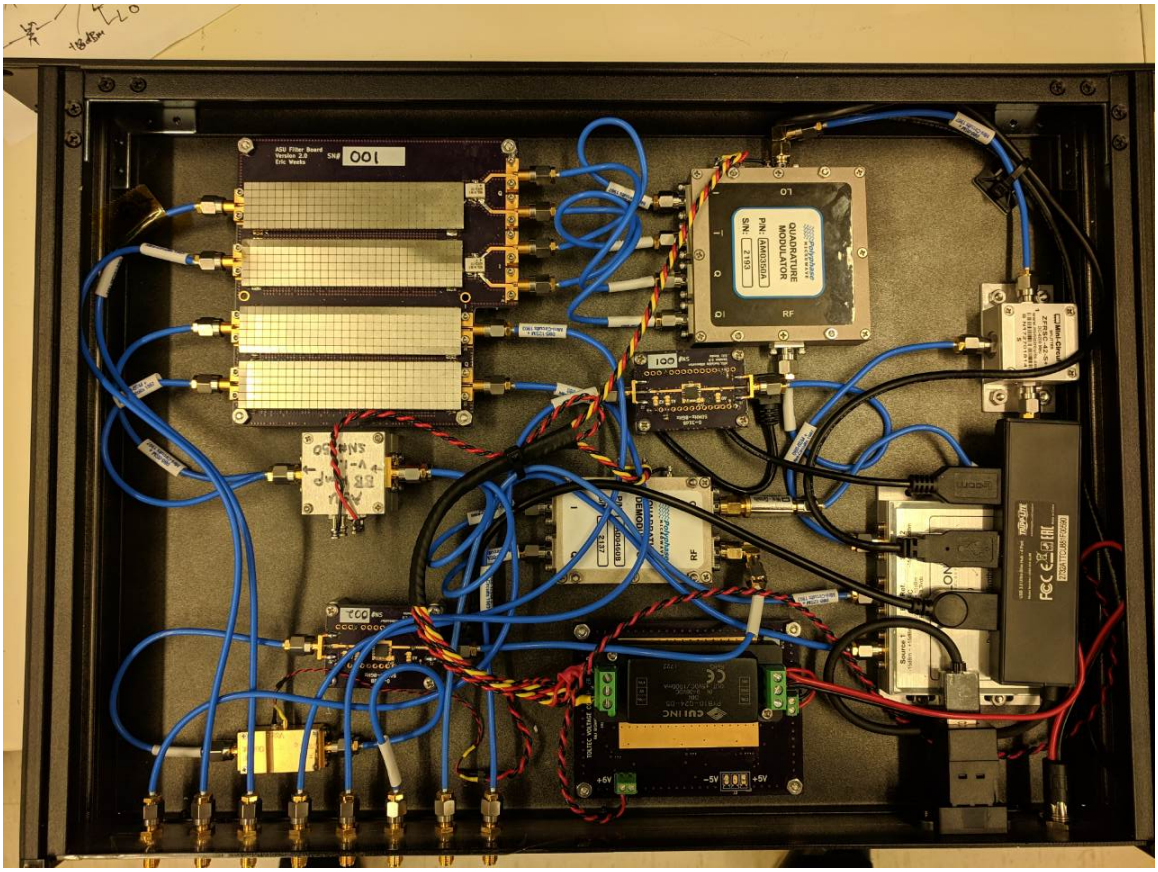

Figure 126. An IF electronics slice for TolTEC. Image from Eric Weeks.

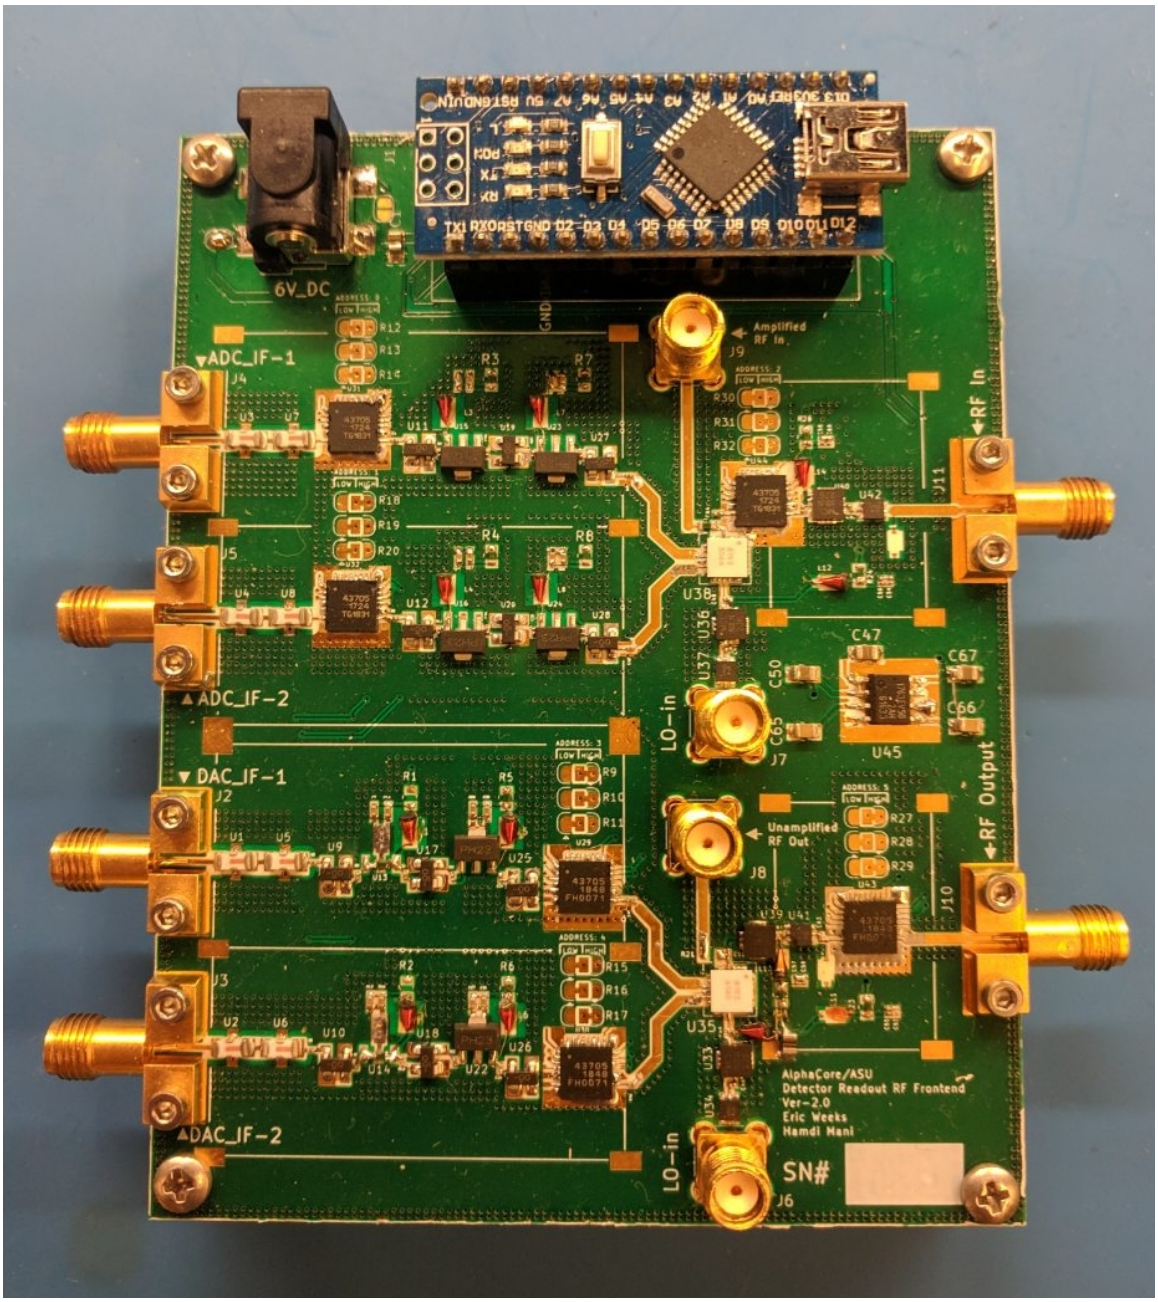

Figure 127. A single-board board implementation of the TolTEC IF electronics. Image from Eric Weeks.

REFERENCES

Abergel, A., Ade, P. A., Aghanim, N., Alves, M., Aniano, G., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., et al. (2014). Planck 2013 results. xi. all-sky model of thermal dust emission. *Astronomy & Astrophysics*, 571:A11.

Abitbol, M. H., Ahmed, Z., Barron, D., Thakur, R. B., Bender, A. N., Benson, B. A., Bischoff, C. A., Bryan, S. A., Carlstrom, J. E., Chang, C. L., et al. (2017). Cmb-s4 technology book. *arXiv preprint arXiv:1706.02464*.

Adam, R., Ade, P. A., Aghanim, N., Arnaud, M., Aumont, J., Baccigalupi, C., Banday, A. J., Barreiro, R., Bartlett, J., Bartolo, N., et al. (2016). Planck intermediate results-xxx. the angular power spectrum of polarized dust emission at intermediate and high galactic latitudes. *Astronomy & Astrophysics*, 586:A133.

Ade, P., Aghanim, N., Alves, M., Arnaud, M., Arzoumanian, D., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A., Barreiro, R., et al. (2016). Planck intermediate results-xxxv. probing the role of the magnetic field in the formation of structure in molecular clouds. *Astronomy & Astrophysics*, 586:A138.

Ade, P., Aguirre, J., Ahmed, Z., Aiola, S., Ali, A., Alonso, D., Alvarez, M. A., Arnold, K., Ashton, P., Austermann, J., et al. (2019). The simons observatory: science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, 2019(02):056.

Ade, P. A., Aghanim, N., Alina, D., Alves, M., Armitage-Caplan, C., Arnaud, M., Arzoumanian, D., Ashdown, M., Atrio-Barandela, F., Aumont, J., et al. (2015). Planck intermediate results. xix. an overview of the polarized thermal emission from galactic dust. *Astronomy & Astrophysics*, 576:A104.

Aghanim, N., Alves, M., Arnaud, M., Arzoumanian, D., Aumont, J., Baccigalupi, C., Banday, A., Barreiro, R., Bartolo, N., Battaner, E., et al. (2016). Planck intermediate results-xxxiv. the magnetic field structure in the rosette nebula. *Astronomy & Astrophysics*, 586:A137.

Aguirre, J., Collaboration, S., et al. (2018). Starfire: The spectroscopic terahertz airborne receiver for far-infrared exploration. In *American Astronomical Society Meeting Abstracts# 231*, volume 231.

Allen, D. and Hillier, D. (1993). The shape of the homunculus nebula around η carinae. *Publications of the Astronomical Society of Australia*, 10(4):338–341.

Anderson, L., Zavagno, A., Deharveng, L., Abergel, A., Motte, F., André, P., Bernard, J.-P., Bontemps, S., Hennemann, M., Hill, T., et al. (2012). The dust properties of bubble h ii regions as seen by herschel. *Astronomy & Astrophysics*, 542:A10.

Andersson, B., Lazarian, A., and Vaillancourt, J. E. (2015). Interstellar dust grain alignment. *Annual Review of Astronomy and Astrophysics*, 53:501–539.

Anlage, S., Snortland, H., and Beasley, M. (1989). A current controlled variable delay superconducting transmission line. *IEEE Transactions on Magnetics*, 25(2):1388–1391.

Annunziata, A. J., Santavicca, D. F., Frunzio, L., Catelani, G., Rooks, M. J., Frydman, A., and Prober, D. E. (2010). Tunable superconducting nanoinductors. *Nanotechnology*, 21(44):445202.

Ashton, P. C., Ade, P. A., Angilè, F. E., Benton, S. J., Devlin, M. J., Dober, B., Fissel, L. M., Fukui, Y., Galitzki, N., Gandilo, N. N., et al. (2018). First observation of the submillimeter polarization spectrum in a translucent molecular cloud. *The Astrophysical Journal*, 857(1):10.

Austermann, J., Beall, J., Bryan, S., Dober, B., Gao, J., Hilton, G., Hubmayr, J., Mauskopf, P., McKenney, C., Simon, S., et al. (2018). Millimeter-wave polarimeters using kinetic inductance detectors for toltec and beyond. *Journal of Low Temperature Physics*, 193(3-4):120–127.

Bardeen, J., Cooper, L. N., and Schrieffer, J. R. (1957). Theory of superconductivity. *Physical Review*, 108(5):1175.

Barry, P. (2014). *On the development of SuperSpec: a fully integrated on-chip spectrometer for far-infrared astronomy*. PhD thesis, Cardiff University.

Barth, A. J., Strigari, L. E., Bentz, M. C., Greene, J. E., and Ho, L. C. (2008). Dynamical constraints on the masses of the nuclear star cluster and black hole in the late-type spiral galaxy ngc 3621. *The Astrophysical Journal*, 690(1):1031.

Battersby, C., Armus, L., Bergin, E., Kataria, T., Meixner, M., Pope, A., Stevenson, K. B., Cooray, A., Leisawitz, D., Scott, D., et al. (2018). The origins space telescope. *Nature Astronomy*, 2(8):596.

Beck, R. (2006). The origin of magnetic fields in galaxies: Observational tests with the square kilometre array. *Astronomische Nachrichten: Astronomical Notes*, 327(5-6):512–516.

Beck, R. (2008). Measuring interstellar magnetic fields by radio synchrotron emission. *Proceedings of the International Astronomical Union*, 4(S259):3–14.

Beck, R. (2016). Magnetic fields in spiral galaxies. *The Astronomy and Astrophysics Review*, 24(1):4.

Beck, R., Fletcher, A., Shukurov, A., Snodin, A., Sokoloff, D., Ehle, M., Moss, D., and Shoutenkov, V. (2005). Magnetic fields in barred galaxies-iv. ngc 1097 and ngc 1365. *Astronomy & Astrophysics*, 444(3):739–765.

Benton, S. J. (2015). *Mapping submillimetre polarization with BLASTPol*. PhD thesis.

Bohlin, R. (1978). Savage. bd, & drake, jf 1978. *ApJ*, 224:132.

Bourrion, O., Benoit, A., Bouly, J., Bouvier, J., Bosson, G., Calvo, M., Catalano, A., Goupy, J., Li, C., Macias-Perez, J., et al. (2016). Nikel_amc: Readout electronics for the nika2 experiment. *Journal of Instrumentation*, 11(11):P11001.

Bourrion, O., Vescovi, C., Catalano, A., Calvo, M., D'Addabbo, A., Goupy, J., Boudou, N., Macias-Perez, J., and Monfardini, A. (2013). High speed readout electronics development for frequency-multiplexed kinetic inductance detector design optimization. *Journal of Instrumentation*, 8(12):C12006.

Boyd, S. (1986). Multitone signals with low crest factor. *IEEE transactions on circuits and systems*, 33(10):1018–1022.

Brien, T. L., Ade, P. A., Barry, P. S., Castillo-Domìnguez, E., Ferrusca, D., Gascard, T., Gómez, V., Hargrave, P. C., Hornsby, A. L., Hughes, D., et al. (2018). Muscat: the mexico-uk sub-millimetre camera for astronomy. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, page 107080M. International Society for Optics and Photonics.

Cabral, B. and Leedom, L. C. (1993). Imaging vector fields using line integral convolution. Technical report, Lawrence Livermore National Lab., CA (United States).

Catalano, A., Calvo, M., Ponthieu, N., Adam, R., Adane, A., Ade, P., André, P., Beelen, A., Belier, B., Benoit, A., et al. (2014). Performance and calibration of the nika camera at the iram 30 m telescope. *Astronomy & Astrophysics*, 569:A9.

Chandrasekhar, S. and Fermi, E. (1953). Magnetic fields in spiral arms. *The Astrophysical Journal*, 118:113.

Cho, J. and Yoo, H. (2016). A technique for constraining the driving scale of turbulence and a modified chandrasekhar–fermi method. *The Astrophysical Journal*, 821(1):21.

Churchwell, E., Povich, M., Allen, D., Taylor, M., Meade, M., Babler, B., Indebetouw, R., Watson, C., Whitney, B., Wolfire, M., et al. (2006). The bubbling galactic disk. *The Astrophysical Journal*, 649(2):759.

Crutcher, R. M., Nutter, D., Ward-Thompson, D., and Kirk, J. (2004). Scuba polarization measurements of the magnetic field strengths in the l183, l1544, and l43 prestellar cores. *The Astrophysical Journal*, 600(1):279.

da Cunha, E., Eminian, C., Charlot, S., and Blaizot, J. (2010). New insight into the relation between star formation activity and dust content in galaxies. *Monthly Notices of the Royal Astronomical Society*, 403(4):1894–1908.

D'Addabbo, A., Adam, R., Adane, A., Ade, P., Andre, P., Beelen, A., Belier, B., Benoît, A., Bideaud, A., Billot, N., et al. (2013). The nika instrument: results and perspectives towards a permanent kid based camera for the pico veleta observatory. *arXiv preprint arXiv:1312.4801*.

Davidson, K., Helmel, G., and Humphreys, R. M. (2018). Gaia, trumpler 16, and eta carinae. *arXiv preprint arXiv:1808.02073*.

Davis Jr, L. (1951). The strength of interstellar magnetic fields. *Physical Review*, 81(5):890.

Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A., and Zmuidzinas, J. (2003). A broadband superconducting detector suitable for use in large arrays. *Nature*, 425(6960):817.

de Jong, T., Klein, U., Wielebinski, R., and Wunderlich, E. (1985). Astronomy and astrophysics, 147.

De Visser, P. (2014). *Quasiparticle dynamics in aluminium superconducting microwave resonators*. PhD thesis, TU Delft, Delft University of Technology.

Devlin, M. J., Ade, P. A., Aretxaga, I., Bock, J. J., Chapin, E. L., Griffin, M., Gundersen, J. O., Halpern, M., Hargrave, P. C., Hughes, D. H., et al. (2009a). Over half of the far-infrared background light comes from galaxies at z≥ 1.2. *Nature*, 458(7239):737.

Devlin, M. J. et al. (2009b). Over half of the far-infrared background light comes from galaxies at z ≥ 1.2. *Nature*, 458(7239):737–739.

Dicker, S., Ade, P., Aguirre, J., Brevik, J., Cho, H., Datta, R., Devlin, M., Dober, B., Egan, D., Ford, J., et al. (2014). Mustang2: a large focal plan array for the 100 meter green bank telescope. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, volume 9153, page 91530J. International Society for Optics and Photonics.

Dober, B., Austermann, J., Beall, J., Becker, D., Che, G., Cho, H., Devlin, M., Duff, S., Galitzki, N., Gao, J., et al. (2016). Optical demonstration of thz, dual-polarization sensitive microwave kinetic inductance detectors. *Journal of Low Temperature Physics*, 184(1-2):173–179.

Dober, B. et al. (2014). The next-generation BLASTPol experiment. In *SPIE Astronomical Telescopes+ Instrumentation*, pages 91530H–91530H. International Society for Optics and Photonics.

Dodkins, R., Mahashabde, S., O'Brien, K., Thatte, N., Fruitwala, N., Walter, A., Meeker, S., Szypryt, P., and Mazin, B. (2018). Mkid digital readout tuning with deep learning. *Astronomy and computing*, 23:60–71.

Dole, H., Lagache, G., Puget, J.-L., Caputi, K. I., Fernandez-Conde, N., Le Floc'h, E., Papovich, C., Pérez-González, P. G., Rieke, G. H., and Blaylock, M. (2006). The cosmic infrared background resolved by spitzer-contributions of mid-infrared galaxies to the far-infrared background. *Astronomy & Astrophysics*, 451(2):417–429.

Doyle, S., Mauskopf, P., Naylon, J., Porch, A., and Duncombe, C. (2008). Lumped element kinetic inductance detectors. *Journal of Low Temperature Physics*, 151(1-2):530–536.

Draine, B. T. (2003). Interstellar dust grains. *Annual Review of Astronomy and Astrophysics*, 41(1):241–289.

Duan, R., McHugh, S., Serfass, B., Mazin, B. A., Merrill, A., Golwala, S. R., Downes, T. P., Czakon, N. G., Day, P. K., Gao, J., et al. (2010). An open-source readout for mkids. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741, page 77411V. International Society for Optics and Photonics.

Duffing, G. (1918). *Erzwungene Schwingungen bei veränderlicher Eigenfrequenz und ihre technische Bedeutung*. Number 41-42. F. Vieweg & sohn.

Elmouttie, M., Haynes, R., Jones, K., Ehle, M., Beck, R., Harnett, J., and Wielebinski, R. (1997). The radio continuum structure of the edge-on spiral galaxy ngc 4945. *Monthly Notices of the Royal Astronomical Society*, 284(4):830–838.

Elmouttie, M., Haynes, R., Jones, K., Sadler, E., and Ehle, M. (1998). Radio continuum evidence for nuclear outflow in the circinus galaxy. *Monthly Notices of the Royal Astronomical Society*, 297(4):1202–1218.

Endo, A., Van der Werf, P., Janssen, R., De Visser, P., Klapwijk, T., Baselmans, J., Ferrari, L., Baryshev, A., and Yates, S. (2012). Design of an integrated filterbank for deshima: On-chip submillimeter imaging spectrograph based on superconducting resonators. *Journal of Low Temperature Physics*, 167(3-4):341–346.

Fendt, C., Beck, R., and Neininger, N. (1998). Spiral pattern in the optical polarization of ngc 6946. *Astronomy and Astrophysics*, 335:123–133.

Ferriere, K. and Terral, P. (2014). Analytical models of x-shape magnetic fields in galactic halos. *Astronomy & Astrophysics*, 561:A100.

Fissel, L. M., Ade, P. A., Angilè, F. E., Ashton, P., Benton, S. J., Chen, C.-Y., Cunningham, M., Devlin, M. J., Dober, B., Friesen, R., et al. (2018). Relative alignment between the magnetic field and molecular gas structure in the vela c giant molecular cloud using low and high density tracers. *arXiv preprint arXiv:1804.08979*.

Fissel, L. M. et al. (2016). Balloon-borne submillimeter polarimetry of the Vela C molecular cloud: Systematic dependence of polarization fraction on column density and local polarization-angle dispersion. *The Astrophysical Journal*, 824(2):134.

Flanigan, D., McCarrick, H., Jones, G., Johnson, B. R., Abitbol, M. H., Ade, P., Araujo, D., Bradford, K., Cantor, R., Che, G., et al. (2016). Photon noise from chaotic and coherent millimeter-wave sources measured with horn-coupled, aluminum lumped-element kinetic inductance detectors. *Applied Physics Letters*, 108(8):083504.

Fletcher, A., Beck, R., Shukurov, A., Berkhuijsen, E., and Horellou, C. (2011). Magnetic fields and spiral arms in the galaxy m51. *Monthly Notices of the Royal Astronomical Society*, 412(4):2396–2416.

Fox, M. (2006). *Quantum optics: an introduction*, volume 15. OUP Oxford.

Franco, G. and Alves, F. (2015). Tracing the magnetic field morphology of the lupus i molecular cloud. *The Astrophysical Journal*, 807(1):5.

Frick, P., Stepanov, R., Beck, R., Sokoloff, D., Shukurov, A., Ehle, M., and Lundgren, A. (2016). Magnetic and gaseous spiral arms in m83. *Astronomy & Astrophysics*, 585:A21.

Gaczkowski, B., Preibisch, T., Ratzka, T., Roccatagliata, V., Ohlendorf, H., and Zinnecker, H. (2013). Herschel far-infrared observations of the carina nebula complex-ii. the embedded young stellar and protostellar population. *Astronomy & Astrophysics*, 549:A67.

Galitzki, N. et al. (2014a). The balloon-borne large aperture submillimeter telescope for polarimetry-blastpol: Performance and results from the 2012 antarctic flight. In *Ground-based and Airborne Telescopes V*, volume 9145, page 91450R. International Society for Optics and Photonics.

Galitzki, N. et al. (2014b). The next generation BLAST experiment. *Journal of Astronomical Instrumentation*, 3(02):1440001.

Gandilo, N. N. et al. (2016). Submillimeter polarization spectrum in the vela c molecular cloud. *The Astrophysical Journal*, 824(2):84.

Gao, J. (2008). *The physics of superconducting microwave resonators*. PhD thesis, California Institute of Technology.

Gao, J., Daal, M., Martinis, J. M., Vayonakis, A., Zmuidzinas, J., Sadoulet, B., Mazin, B. A., Day, P. K., and Leduc, H. G. (2008). A semiempirical model for two-level system noise in superconducting microresonators. *Applied Physics Letters*, 92(21):212504.

Gendelev, L. and Krumholz, M. R. (2012). Evolution of blister-type h ii regions in a magnetized medium. *The Astrophysical Journal*, 745(2):158.

Glover, R. and Tinkham, M. (1957). Conductivity of superconducting films for photon energies between 0.3 and 4 0 k t c. *Physical Review*, 108(2):243.

Golwala, S. R., Bockstiegel, C., Brugger, S., Czakon, N. G., Day, P. K., Downes, T. P., Duan, R., Gao, J., Gill, A. K., Glenn, J., et al. (2012). Status of music, the multiwavelength sub/millimeter inductance camera. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, volume 8452, page 845205. International Society for Optics and Photonics.

Gordon, S. et al. (2016). An open source, FPGA-based LeKID readout for BLAST-TNG: Pre-flight results. *Journal of Astronomical Instrumentation*, 5(4).

Greaves, J., Holland, W., Jenness, T., and Hawarden, T. (2000). Magnetic field surrounding the starburst nucleus of the galaxy m82 from polarized dust emission. *Nature*, 404(6779):732.

Griffin, M. J., Swinyard, B. M., and Vigroux, L. G. (2003). Spire-herschel's submillimetre camera and spectrometer. In *IR Space Telescopes and Instruments*, volume 4850, pages 686–698. International Society for Optics and Photonics.

Hailey-Dunsheath, S., Shirokoff, E., Barry, P., Bradford, C., Chapman, S., Che, G., Glenn, J., Hollister, M., Kovács, A., LeDuc, H., et al. (2016). Low noise titanium nitride kids for superspec: A millimeter-wave on-chip spectrometer. *Journal of Low Temperature Physics*, 184(1-2):180–187.

Hall, J. S. (1949). Observations of the polarized light from stars. *Science*, 109(2825):166–167.

Hargrave, P., Waskett, T., Lim, T., and Swinyard, B. (2006). Performance of flight-model on-board calibration sources on herschel-spire. In *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III*, volume 6275, page 627514. International Society for Optics and Photonics.

Hecht, E. (2002). *Optics*. Pearson education. Addison-Wesley.

Heitsch, F., Zweibel, E. G., Mac Low, M.-M., Li, P., and Norman, M. L. (2001). Magnetic field diagnostics based on far-infrared polarimetry: tests using numerical simulations. *The Astrophysical Journal*, 561(2):800.

Henderson, S., Allison, R., Austermann, J., Baildon, T., Battaglia, N., Beall, J., Becker, D., De Bernardis, F., Bond, J., Calabrese, E., et al. (2016). Advanced actpol cryogenic detector arrays and readout. *Journal of Low Temperature Physics*, 184(3-4):772–779.

Henderson, S. W., Ahmed, Z., Austermann, J., Becker, D., Bennett, D. A., Brown, D., Chaudhuri, S., Cho, H.-M. S., D'Ewart, J. M., Dober, B., et al. (2018). Highly-multiplexed microwave squid readout using the slac microresonator radio frequency (smurf) electronics for future cmb and sub-millimeter surveys. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, page 1070819. International Society for Optics and Photonics.

Hickish, J., Abdurashidova, Z., Ali, Z., Buch, K. D., Chaudhari, S. C., Chen, H., Dexter, M., Domagalski, R. S., Ford, J., Foster, G., et al. (2016). A decade of developing radio-astronomy instrumentation using casper open-source technology. *Journal of Astronomical Instrumentation*, 5(04):1641001.

Hildebrand, R. H., Kirby, L., Dotson, J. L., Houde, M., and Vaillancourt, J. E. (2009). Dispersion of magnetic fields in molecular clouds. i. *The Astrophysical Journal*, 696(1):567.

Hiltner, W. (1949). Polarization of light from distant stars by interstellar medium. *Science*, 109(2825):165–165.

Houde, M. (2004). Evaluating the magnetic field strength in molecular clouds. *The Astrophysical Journal Letters*, 616(2):L111.

Houde, M., Vaillancourt, J. E., Hildebrand, R. H., Chitsazzadeh, S., and Kirby, L. (2009). Dispersion of magnetic fields in molecular clouds. ii. *The Astrophysical Journal*, 706(2):1504.

Hubmayr, J., Beall, J., Becker, D., Brevik, J., Cho, H., Che, G., Devlin, M., Dober, B., Gao, J., Galitzki, N., et al. (2014). Dual-polarization-sensitive kinetic inductance detectors for balloon-borne sub-millimeter polarimetry. *Journal of Low Temperature Physics*, 176(3-4):490–496.

Hubmayr, J., Beall, J., Becker, D., Cho, H.-M., Devlin, M., Dober, B., Groppi, C., Hilton, G. C., Irwin, K. D., Li, D., et al. (2015). Photon-noise limited sensitivity in titanium nitride kinetic inductance detectors. *Applied Physics Letters*, 106(7):073505.

Irwin, K. D. and Lehnert, K. W. (2004). Microwave squid multiplexer. *Applied physics letters*, 85(11):2107–2109.

Johnson, J. B. (1928). Thermal agitation of electricity in conductors. *Phys. Rev.*, 32:97–109.

Jones, T. J., Dowell, C. D., Rodriguez, E. L., Zweibel, E. G., Berthoud, M., Chuss, D. T., Goldsmith, P. F., Hamilton, R. T., Hanany, S., Harper, D. A., et al. (2019). Sofia far-infrared imaging polarimetry of m82 and ngc 253: Exploring the supergalactic wind. *The Astrophysical Journal Letters*, 870(1):L9.

Kandus, A., Kunze, K. E., and Tsagas, C. G. (2011). Primordial magnetogenesis. *Physics Reports*, 505(1):1–58.

Kittel, C. and Kroemer, H. (1998). Thermal physics.

Kramer, C., Cubick, M., Röllig, M., Sun, K., Yonekura, Y., Aravena, M., Bensch, F., Bertoldi, F., Bronfman, L., Fujishita, M., et al. (2008). Clumpy photon-dominated regions in carina-i.[c i] and mid-j co lines in two 4″ x 4″. *Astronomy & Astrophysics*, 477(2):547–555.

Kudoh, T. and Basu, S. (2003). Nonlinear hydromagnetic wave support of a stratified molecular cloud. *The Astrophysical Journal*, 595(2):842.

Lasker, B. M., Lattanzi, M. G., McLean, B. J., Bucciarelli, B., Drimmel, R., Garcia, J., Greene, G., Guglielmetti, F., Hanley, C., Hawkins, G., et al. (2008). The second-generation guide star catalog: description and properties. *The Astronomical Journal*, 136(2):735.

Lazarian, A. and Hoang, T. (2007). Radiative torques: analytical model and basic properties. *Monthly Notices of the Royal Astronomical Society*, 378(3):910–946.

Lesson, D. (1966). A simple model of feedback oscillator noise spectrum. *proc. IEEE*, 54(2):329–330.

Li, H., Li, S.-Y., Liu, Y., Li, Y.-P., Cai, Y., Li, M., Zhao, G.-B., Liu, C.-Z., Li, Z.-W., Xu, H., et al. (2018). Probing primordial gravitational waves: Ali cmb polarization telescope. *National Science Review*, 6(1):145–154.

Li, H.-b., Griffin, G., Krejny, M., Novak, G., Loewenstein, R., Newcomb, M., Calisse, P., and Chuss, D. (2006). Results of sparo 2003: mapping magnetic fields in giant molecular clouds. *The Astrophysical Journal*, 648(1):340.

Lourie, N. P., Angilé, F. E., Ashton, P. C., Catanzaro, B., Devlin, M. J., Dicker, S., Didier, J., Dober, B., Fissel, L. M., Galitzki, N., et al. (2018). Design and characterization of a balloon-borne diffraction-limited submillimeter telescope platform for blast-tng. In *Ground-based and Airborne Telescopes VII*, volume 10700, page 1070022. International Society for Optics and Photonics.

Lucy, L. B. (1974). An iterative technique for the rectification of observed distributions. *The astronomical journal*, 79:745.

Lyons, R. G. (2004). *Understanding digital signal processing, 3/E*. Pearson Education India.

Ma, K.-L. (1996). Texture synthesis with line integral convolution. *ACM SIGGRAPH97 Course Notes*, 8.

Marsden, G., Ade, P. A., Bock, J. J., Chapin, E. L., Devlin, M. J., Dicker, S. R., Griffin, M., Gundersen, J. O., Halpern, M., Hargrave, P. C., et al. (2009). Blast: Resolving the cosmic submillimeter background. *The Astrophysical Journal*, 707(2):1729.

Martinis, J. M., Cooper, K. B., McDermott, R., Steffen, M., Ansmann, M., Osborn, K., Cicak, K., Oh, S., Pappas, D. P., Simmonds, R. W., et al. (2005). Decoherence in josephson qubits from dielectric loss. *Physical review letters*, 95(21):210503.

Masi, S., Battistelli, E., Brienza, D., Conversi, L., Cruciani, A., de Bernardis, P., De Petris, M., Fiadino, P., Iacoangeli, A., Lamagna, L., et al. (2008). Olimpo. *Mem. S. It.*, 79:887.

Masi, S., de Bernardis, P., Paiella, A., Piacentini, F., Lamagna, L., Coppolecchia, A., Ade, P., Battistelli, E., Castellano, M., Colantoni, I., et al. (2019). Kinetic inductance detectors for the olimpo experiment: in–flight operation and performance. *arXiv preprint arXiv:1902.08993*.

Mates, J., Becker, D. T., Bennett, D. A., Dober, B., Gard, J., Hays-Wehle, J., Fowler, J., Hilton, G., Reintsema, C., Schmidt, D., et al. (2017). Simultaneous readout of 128 x-ray and gamma-ray transition-edge microcalorimeters using microwave squid multiplexing. *Applied Physics Letters*, 111(6):062601.

Mathis, J. S. (1990). Interstellar dust and extinction. *Annual Review of Astronomy and Astrophysics*, 28(1):37–70.

Matsumura, T., Akiba, Y., Borrill, J., Chinone, Y., Dobbs, M., Fuke, H., Ghribi, A., Hasegawa, M., Hattori, K., Hattori, M., et al. (2014). Mission design of litebird. *Journal of Low Temperature Physics*, 176(5-6):733–740.

Matthews, T. G., Ade, P. A., Angilè, F. E., Benton, S. J., Chapin, E. L., Chapman, N. L., Devlin, M. J., Fissel, L. M., Fukui, Y., Gandilo, N. N., et al. (2014). Lupus i observations from the 2010 flight of the balloon-borne large aperture submillimeter telescope for polarimetry. *The Astrophysical Journal*, 784(2):116.

Mattis, D. and Bardeen, J. (1958). Theory of the anomalous skin effect in normal and superconducting metals. *Physical Review*, 111(2):412.

Mauskopf, P. (2018). Transition edge sensors and kinetic inductance detectors in astronomical instruments. *Publications of the Astronomical Society of the Pacific*, 130(990):082001.

Mauskopf, P., Doyle, S., Barry, P., Rowe, S., Bidead, A., Ade, P., Tucker, C., Castillo, E., Monfardini, A., Goupy, J., et al. (2014). Photon-noise limited performance in aluminum lekids. *Journal of Low Temperature Physics*, 176(3-4):545–552.

Mazin, B., Meeker, S. R., Strader, M., Szypryt, P., Marsden, D., van Eyken, J., Duggan, G., Walter, A., Ulbricht, G., Johnson, M., et al. (2013). Arcons: A 2024 pixel optical through near-ir cryogenic imaging spectrophotometer. *Publications of the Astronomical Society of the Pacific*, 125(933):1348.

Mazin, B. A. (2005). Microwave kinetic inductance detectors. Technical report.

McGeehan, R., Barry, P., Shirokoff, E., Bradford, C., Che, G., Glenn, J., Gordon, S., Hailey-Dunsheath, S., Hollister, M., Kovács, A., et al. (2018). Low-temperature noise performance of superspec and other developments on the path to deployment. *Journal of Low Temperature Physics*, 193(5-6):1024–1032.

McHugh, S., Mazin, B. A., Serfass, B., Meeker, S., O'Brien, K., Duan, R., Raffanti, R., and Werthimer, D. (2012). A readout for large arrays of microwave kinetic inductance detectors. *Review of Scientific Instruments*, 83(4):044702.

Meeker, S. R., Mazin, B. A., Walter, A. B., Strader, P., Fruitwala, N., Bockstiegel, C., Szypryt, P., Ulbricht, G., Coiffard, G., Bumble, B., et al. (2018). Darkness: A microwave kinetic inductance detector integral field spectrograph for high-contrast astronomy. *Publications of the Astronomical Society of the Pacific*, 130(988):065001.

Monfardini, A., Adam, R., Adane, A., Ade, P., André, P., Beelen, A., Belier, B., Benoit, A., Bideaud, A., Billot, N., et al. (2014). Latest nika results and the nika-2 project. *Journal of Low Temperature Physics*, 176(5-6):787–795.

Noguchi, T., Dominjon, A., and Sekimoto, Y. (2018). Analysis of characteristics of al mkid resonators. *IEEE Transactions on Applied Superconductivity*, 28(4):1–6.

Nyquist, H. (1928). Thermal agitation of electric charge in conductors. *Physical review*, 32(1):110.

Oberst, T., Parshley, S., Stacey, G., Nikola, T., Löhr, A., Harnett, J., Tothill, N., Lane, A., Stark, A., and Tucker, C. (2006). Detection of the 205 μm [n ii] line from the carina nebula. *The Astrophysical Journal Letters*, 652(2):L125.

Ostriker, E. C., Stone, J. M., and Gammie, C. F. (2001). Density, velocity, and magnetic field structure in turbulent molecular cloud models. *The Astrophysical Journal*, 546(2):980.

Owen, C. S. and Scalapino, D. J. (1972). Superconducting state under the influence of external dynamic pair breaking. *Phys. Rev. Lett.*, 28:1559–1561.

Padoan, P., Goodman, A., Draine, B., Juvela, M., Nordlund, Å., and Rögnvaldsson, Ö. E. (2001). Theoretical models of polarized dust emission from protostellar cores. *The Astrophysical Journal*, 559(2):1005.

Paiella, A., Battistelli, E., Castellano, M., Colantoni, I., Columbro, F., Coppolecchia, A., D'Alessandro, G., de Bernardis, P., Gordon, S., Lamagna, L., et al. (2019). Kinetic inductance detectors and readout electronics for the olimpo experiment. In *Journal of Physics: Conference Series*, volume 1182, page 012005. IOP Publishing.

Pascale, E., Ade, P. A., Bock, J. J., Chapin, E. L., Devlin, M. J., Dye, S., Eales, S. A., Griffin, M., Gundersen, J. O., Halpern, M., et al. (2009). Blast: A far-infrared measurement of the history of star formation. *The Astrophysical Journal*, 707(2):1740.

Pavel, M. D. and Clemens, D. P. (2012). H ii region driven galactic bubbles and their relationship to the galactic magnetic field. *The Astrophysical Journal*, 760(2):150.

Pellegrini, E., Baldwin, J., Brogan, C., Hanson, M., Abel, N., Ferland, G. J., Nemala, H., Shaw, G., and Troland, T. H. (2007). A magnetically supported photodissociation region in m17. *The Astrophysical Journal*, 658(2):1119.

Preibisch, T. et al. (2012). Herschel far-infrared observations of the carina nebula complex-i. introduction and global cloud structure. *Astronomy & Astrophysics*, 541:A132.

Preibisch, T., Ratzka, T., Kuderna, B., Ohlendorf, H., King, R. R., Hodgkin, S., Irwin, M., Lewis, J. R., McCaughrean, M. J., and Zinnecker, H. (2011a). Deep wide-field near-infrared survey of the carina nebula. *Astronomy & Astrophysics*, 530:A34.

Preibisch, T., Schuller, F., Ohlendorf, H., Pekruhl, S., Menten, K. M., and Zinnecker, H. (2011b). A deep wide-field sub-mm survey of the carina nebula complex. *Astronomy & Astrophysics*, 525:A92.

Price, D. C. (2016). Spectrometers and polyphase filterbanks in radio astronomy. *arXiv preprint arXiv:1607.03579*.

Rahner, D., Pellegrini, E. W., Glover, S. C., and Klessen, R. S. (2017). Winds and radiation in unison: a new semi-analytic feedback model for cloud dissolution. *Monthly Notices of the Royal Astronomical Society*, 470(4):4453–4472.

Rebolledo, D., Green, A. J., Burton, M., Brooks, K., Breen, S. L., Gaensler, B., Contreras, Y., Braiding, C., and Purcell, C. (2017). The carina nebula and gum 31 molecular complex–ii. the distribution of the atomic gas revealed in unprecedented detail. *Monthly Notices of the Royal Astronomical Society*, 472(2):1685–1704.

Richardson, W. H. (1972). Bayesian-based iterative method of image restoration. *JOSA*, 62(1):55–59.

Roccatagliata, V., Preibisch, T., Ratzka, T., and Gaczkowski, B. (2013). Herschel far-infrared observations of the carina nebula complex-iii. detailed cloud structure and feedback effects. *Astronomy & Astrophysics*, 554:A6.

Rowe, S. (2015). *Passive terahertz imaging with lumped element kinetic inductance detectors*. PhD thesis, Cardiff University.

Scarrott, S., Draper, P., Stockdale, D., and Wolstencroft, R. (1993). Imaging polarimetry of the starburst galaxy ngc 1808: another m82? *Monthly Notices of the Royal Astronomical Society*, 264(1):L7–L12.

Seiffert, M., Borys, C., Scott, D., and Halpern, M. (2006). An upper limit to polarized submillimetre emission in arp 220. *Monthly Notices of the Royal Astronomical Society*, 374(2):409–414.

Shariff, J. A. (2015). *Polarimetry from the stratosphere with SPIDER and BLASTPol*. University of Toronto (Canada).

Shariff, J. A., Ade, P. A., Angilè, F. E., Ashton, P., Benton, S. J., Devlin, M. J., Dober, B., Fissel, L. M., Fukui, Y., Galitzki, N., et al. (2019). Submillimeter polarization spectrum of the carina nebula. *The Astrophysical Journal*, 872(2):197.

Shirokoff, E. et al. (2012). Mkid development for superspec: an on-chip, mm-wave, filter-bank spectrometer. In *SPIE Astronomical Telescopes+ Instrumentation*, pages 84520R–84520R. International Society for Optics and Photonics.

Siebenmorgen, R., Krugel, E., and Laureijs, R. (2001). The infrared continuum radiation of ngc 1808-a pah and polarisation study. *Astronomy & Astrophysics*, 377(3):735–744.

Smith, N. (2006). The structure of the homunculus. i. shape and latitude dependence from h2 and [fe ii] velocity maps of η carinae. *The Astrophysical Journal*, 644(2):1151.

Smith, N. and Brooks, K. J. (2008). The carina nebula: A laboratory for feedback and triggered star formation. *arXiv preprint arXiv:0809.5081*.

Smith, N., Egan, M. P., Carey, S., Price, S. D., Morse, J. A., and Price, P. A. (2000). Large-scale structure of the carina nebula. *The Astrophysical Journal Letters*, 532(2):L145.

Smith, N., Povich, M. S., Whitney, B. A., Churchwell, E., Babler, B. L., Meade, M. R., Bally, J., Gehrz, R. D., Robitaille, T. P., and Stassun, K. G. (2010). Spitzer space telescope observations of the carina nebula: the steady march of feedback-driven star formation. *Monthly Notices of the Royal Astronomical Society*, 406(2):952–974.

Sofue, Y., Klein, U., Beck, R., and Wielebinski, R. (1985). Large-scale configuration of magnetic fields in spiral galaxies. *Astronomy and Astrophysics*, 144:257–260.

Soler, J., Ade, P., Angilè, F., Ashton, P., Benton, S., Devlin, M., Dober, B., Fissel, L., Fukui, Y., Galitzki, N., et al. (2017). The relation between the column density structures and the magnetic field orientation in the vela c molecular complex. *Astronomy & Astrophysics*, 603:A64.

Soler, J., Bracco, A., and Pon, A. (2018). The magnetic environment of the Orion-Eridanus superbubble as revealed by planck. *Astronomy & Astrophysics*, 609:L3.

Soler, J. D., Hennebelle, P., Martin, P., Miville-Deschenes, M.-A., Netterfield, C. B., and Fissel, L. (2013). An imprint of molecular cloud magnetization in the morphology of the dust polarized emission. *The Astrophysical Journal*, 774(2):128.

Spoon, H., Koornneef, J., Moorwood, A., Lutz, D., and Tielens, A. (2000). Mid-infrared iso spectroscopy of ngc 4945. *ASTRONOMY & ASTROPHYSICS*, 357:898.

Stanchfield, S., Ade, P., Aguirre, J., Brevik, J., Cho, H., Datta, R., Devlin, M., Dicker, S., Dober, B., Egan, D., et al. (2016). Development of a microwave squid-multiplexed tes array for mustang-2. *Journal of Low Temperature Physics*, 184(1-2):460–465.

Strader, M. J. (2016). *Digitial readout for microwave kinetic inductance detectors and applications in high time resolution astronomy*. University of California, Santa Barbara.

Sukumar, S. and Allen, R. (1989). Large-scale magnetic-field structure in the spiral galaxy m83. *Nature*, 340(6234):537.

Swenson, L. et al. (2013). Operation of a titanium nitride superconducting microresonator detector in the nonlinear regime. *Journal of Applied Physics*, 113(10):104501.

Swenson, L. J., Day, P. K., Dowell, C. D., Eom, B. H., Hollister, M. I., Jarnot, R., Kovács, A., Leduc, H. G., McKenney, C. M., Monroe, R., et al. (2012). Mako: a pathfinder instrument for on-sky demonstration of low-cost 350 micron imaging arrays. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, volume 8452, page 84520P. International Society for Optics and Photonics.

Switzer, E. (2017). Measuring the cosmological evolution of gas and galaxies with the experiment for cryogenic large-aperture intensity mapping (exclaim). *NASA APRA Proposal*.

Tinkham, M. (2004). *Introduction to superconductivity*. Courier Corporation.

Townsley, L. K., Broos, P. S., Corcoran, M. F., Feigelson, E. D., Gagné, M., Montmerle, T., Oey, M., Smith, N., Garmire, G. P., Getman, K. V., et al. (2011). An introduction to the chandra carina complex project. *The Astrophysical Journal Supplement Series*, 194(1):1.

Tremblin, P., Anderson, L., Didelon, P., Raga, A. C., Minier, V., Ntormousi, E., Pettitt, A., Pinto, C., Samal, M., Schneider, N., et al. (2014). Age, size, and position of h ii regions in the galaxy-expansion of ionized gas in turbulent molecular clouds. *Astronomy & Astrophysics*, 568:A4.

Vaillancourt, J. E. and Matthews, B. C. (2012). Submillimeter polarization of galactic clouds: a comparison of 350 μm and 850 μm data. *The Astrophysical Journal Supplement Series*, 201(2):13.

Van Eck, C., Brown, J., Shukurov, A., and Fletcher, A. (2015). Magnetic fields in a sample of nearby spiral galaxies. *The Astrophysical Journal*, 799(1):35.

van Rantwijk, J., Grim, M., van Loon, D., Yates, S., Baryshev, A., and Baselmans, J. (2016). Multiplexed readout for 1000-pixel arrays of microwave kinetic inductance detectors. *IEEE Transactions on Microwave Theory and Techniques*, 64(6):1876–1883.

Velazquez, M. et al. (2016). Design of a 2-mm wavelength kids prototype camera for the large millimeter telescope. *Journal of Low Temperature Physics*, 184(3-4):799–804.

Walter, F., Brinks, E., De Blok, W., Bigiel, F., Kennicutt Jr, R. C., Thornley, M. D., and Leroy, A. (2008). Things: The hi nearby galaxy survey. *The Astronomical Journal*, 136(6):2563.

Weber, S., Murch, K., Slichter, D., Vijay, R., and Siddiqi, I. (2011). Single crystal silicon capacitors with low microwave loss in the single photon regime. *Applied Physics Letters*, 98(17):172510.

Welch, P. (1967). The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics*, 15(2):70–73.

Werthimer, D. (2011). The casper collaboration for high-performance open source digital radio astronomy instrumentation. In *2011 XXXth URSI general assembly and scientific symposium*, pages 1–4. IEEE.

Wheeler, J., Hailey-Dunsheath, S., Shirokoff, E., Barry, P., Bradford, C., Chapman, S., Che, G., Doyle, S., Glenn, J., Gordon, S., et al. (2018). Superspec, the on-chip spectrometer: improved nep and antenna performance. *Journal of Low Temperature Physics*, 193(3-4):408–414.

Whittet, D. C. (2002). *Dust in the galactic environment*. CRC press.

Wiebe, D. V. et al. (2009). BLAST observations of resolved galaxies: temperature profiles and the effect of active galactic nuclei on fir to submillimeter emission. *The Astrophysical Journal*, 707(2):1809.

Zmuidzinas, J. (2012). Superconducting microresonators: Physics and applications. *Annu. Rev. Condens. Matter Phys.*, 3(1):169–214.

APPENDIX A

BLAST-TNG LEKID READOUT OPERATOR'S MANUAL

The following sections constitute a manual for the BLAST-TNG readout flight software, which is embedded in the Master Control Program (MCP). This software was completed during the run-up to the 2018/2019 Antarctic flight attempts, and will be used during the 2019/2020 flight.

The BLAST-TNG readout consists of five identical electronics slices. With firmware, software and hardware developed by members of the collaboration between 2014–2018, these slices perform all of the digital and analog signal processing which is required to readout large arrays (large, at the time of writing, being $\mathcal{O}(10^3)$) of microwave kinetic inductance detectors (MKIDs) per slice.

The information contained in this operator's manual is intended to make the in-flight (or ground-based) operation of the readout system a relatively painless experience for first-time, as well as for more experienced users.

It is organized as follows:

- Section A.1 is a glossary of terms.
- Section A.2 provides essential system paths and hardware addresses on the flight computers and groundstation.
- Section A.3 describes the normal modes of readout operation during the flight.
- Section A.4 describes sequences of actions that can be triggered during both manual and automatic operation of the readout.
- Section A.5 describes the OWL graphical user interface (created by Javier Romualdez), which will be used to interface with the software during ground-based testing as well as during the flight.
- Section A.6 describes the available set of readout commands, and how and when they should be used (as of January, 2019).

A.1 Glossary of Terms

1GbE: One-gigabit-per-second Ethernet interface. Maximum bit rate is less than 1 Gbps.
ADC: Also: A/D. Analog-to-digital-converter.
Attenuator server: A TCP server running in the background on each Raspberry Pi 3. The TCP server is created by a Python script, `rudat_server_boot.py`, which runs in the background on each Pi. MCP is the client for the server (see: `roach.c`, `atten_client()`). See also: RUDAT.
Baseband: Also: Complex baseband. A frequency spectrum centered at 0 Hz (DC). This is also the band which is synthesized/digitized by the digital electronics (MUSIC board and FPGA). It spans -256 to +256 MHz (512 MHz total). Each baseband frequency has a corresponding RF carrier frequency (see: probe tone).
BORPH: The Berkeley Operating System for Reprogrammable Hardware. A Linux-

based operating system running on the PPCs. See: BORPH
CASPER: The Collaboration for Astronomical Signal Processing and Electronics Research. See: CASPER
Channel: A unique detector timestream, containing the 32-b *I* and 32-b *Q* values for a single resonator on the array, sampled at 488.28125 Hz. Channel values are saved to disk. Each channel number maps to a single baseband/RF frequency pair. Channel numbers are zero-indexed, and increment in baseband frequency order, i.e., channel 10 corresponds to the 10th baseband frequency in the list of baseband frequencies (e.g., `bb_target_freqs.dat`).
DAC: Also: D/A. Digital-to-analog-converter.
Demodulator: Also: Downconverter. Polyphase AD0540B. A quadrature modulator used to downconvert RF/IF to baseband *I* and *Q*. There is one in each ROACH slice.
Inputs:

- RF signal from the cryostat, containing the multiplexed waveforms from each pixel. LO signal generated by the Valon 5009 Synthesizer. The default LO frequency is 750 MHz (i.e., on system restart).

Outputs:

- Baseband *I* and Q.

Power: +5 and -5 VDC, fed by the Vicor DC/DC converter inside the ROACH enclosure. Datasheet: Polyphase AD0540B
df: 'delta-f'. The amount in Hz by which a channel's resonant (not carrier) frequency has drifted. In the ideal, linear operational regime, *df* is inversely proportional to the amount of absorbed power in a channel.
Downconverter: See Demodulator.
fpg: A CASPER firmware bitstream file type. Unlike a typical bit file, an fpg file contains ASCII-formatted metadata describing its contents. The metadata is located at the end of the binary data. The ROACH FPGA firmware is located on FC1 and FC2 at: /data/etc/blast/roachFirmware
FPGA: Field Programmable Gate Array. A reprogrammable integrated-circuit found on each ROACH-2 board, which runs the ASU open-source KID readout firmware. The FPGAs perform the real-time digital signal processing (DSP) which is necessary to readout the BLAST detectors. These particular FPGAs are Xilinx Virtex-6.
I: The in-phase (real, cosine) component of a quadrature waveform.
IF: Intermediate-frequency. The frequency spectrum produced by mixing baseband frequencies with the LO. In BLAST-TNG, IF coincides with RF, which is roughly 284–1056 MHz.
KATCP: The Karoo Array Telescope Protocol. A KATCP server running on each ROACH PPC provides I/O access to the FPGA firmware registers, as well as to a basic set of commands which can be used to query and control the PPC.
See: KATCP

LEKID: Also, KID. Lumped element kinetic inductance detector.
LO: A variable tone produced by the Valon 5009 Synthesizer, and used to up/downconvert baseband to IF/RF in the mod/demodulators. On system start, each of the LOs are set to 750 MHz. During normal operation, the LOs remain at the center frequencies of each detector array. During sweeps, they are moved. The LO frequencies can be set and read by command (see: LO commands).
LO center frequencies for each array:

- 500 (ROACH1): 540 MHz
- 250U (ROACH2): 827 MHz
- 250W (ROACH5): 828 MHz

LUT: Look-up-table.
MKID: Also, KID. Microwave kinetic inductance detector.
MODULATOR: See Upconverter. Polyphase Microwave AM0350A. A quadrature mixer used to upconvert baseband *I* and *Q* to IF/RF *I* and Q. There is one in each ROACH slice.
Inputs:

- Baseband *I* and Q, synthesized by the MUSIC board *I* and *Q* DACs.
- LO signal generated by the Valon 5009 Synthesizer. The default LO frequency is 750 MHz (i.e., on system restart).

Outputs:

- IF/RF signal which is sent to the cryostat.

Power: +5 and -5 VDC, fed by Vicor DC/DC converter in Roach enclosure. See: Datasheet: Polyphase AM0350A
MUSIC Board: A DAC/ADC board originally developed for the MUSIC camera (Multi-color Sub-millimeter Kinetic Inductance Camera). It is a daughter board of the ROACH-2 in each ROACH slice. Each MUSIC board has:

- 2X ADCs (*I* and *Q*): Texas Instruments ADS54RF63
- 2X DACs (*I* and *Q*): Texas Instruments DAC5681

Datasheet: MUSIC Board
OctoClock: An 8-channel clock distribution module located on the inner-frame which provides the 10 MHz reference and PPS to each ROACH slice. The PPS is also fed to the flight computers. Datasheet: OctoClock
PLL: Phase-locked-loop.
PPC: PowerPC 440EPx. This is the microprocessor onboard each ROACH-2. It runs a Linux OS called BORPH. The PPCs are accessible by SSH, from the ground stations and flight computers. From a flight operations standpoint, the only reason for SSHing into them would be to shut them down. To bring a ROACH back up requires re-cycling the power to the entire enclosure.

```
> ssh root@roach<which roach>
```

Password: `root`
To shutdown:

```
> halt
```

Datasheet: PPC440EPx
Probe tone: A carrier tone generated by the ROACH at each resonator frequency, which is sent through the array and modulated in phase and amplitude. The modulated tone is looped back into the ROACH slice and demodulated by the FPGA firmware and electronics.
Raspberry PI 3: A low-cost, single-board computer found in each ROACH slice. Each Pi serves as a control interface for one Valon 5009 Synthesizer and two RUDAT programmable attenuators.
Inputs/Outputs:

- USB-2 × 3: Control signals for one Valon and two RUDATs.
- Ethernet: Direct into the inner frame switches. See Hardware Addresses.

Power: USB provided by ROACH-2 board. If needed, the Pis are accessible by SSH from either the ground stations or FC1/FC2. To SSH:

```
> ssh pi@pi<which pi>
```

If prompted for a password, it is: `raspberry`
Reference parameters: Several quantities which are calculated for each channel from a target sweep, and used in the *df* calculation. They are: dI/df, dQ/df, I on res, Q on res.
ROACH slice: A unit of detector readout hardware. The BLAST ROACH-2 enclosure contains five identical ROACH slices. Each slice reads out one detector array (up to 1,000 frequency channels), and comprises:

- 1 × Roach-2 FPGA board
- 1 × MUSIC DAC/ADC board
- 1 × Raspberry Pi 3

A set of IF electronics (mod/demodulator pair, RUDAT attenuator pair, Valon synthesizer, anti-aliasing filters, 2 × baseband amplifiers, RF amplifier, 2 × 1 GHz low-pass filters), several meters of coax. . .
RF: The frequency spectrum occupied by the LeKIDs (roughly 300–1100 MHz). In this special case (homodyne system), it's interchangeable with IF.
ROACH-2: Reconfigurable Open Architecture Computing Hardware. The ROACH-2 is the motherboard for each ROACH slice. It can be thought of as a single-board computer having an FPGA as a PCI device.
Inputs:

- 1PPS signal, driven by OctoClock.
- Signals from MUSIC DAC/ADC board, routed through ZDOK bus.
- 1GbE link to the onboard PPC.

Outputs:

- Signals to MUSIC DAC/ADC board, routed through ZDOK bus.
- 1GbE. UDP packets sent to inner-frame ethernet switches and saved to disk.

- 1GbE link to the onboard PPC.

Power: 28 VDC from gondola inner frame. See: CASPER ROACH-2

RUDAT: MiniCircuits RUDAT-6000-30. A programmable step-attenuator with 30 dB of dynamic range that can be stepped in 0.5 dB increments. There are two RUDATs in each ROACH slice:

- Output atten: The RUDAT located at each ROACH slice's output (connected to the cryostat input ports).
- Input atten: The RUDAT located at each ROACH slice's input (connected to the cryostat output ports).

Power: 5 VDC via Raspberry Pi 3 USB. See: RUDAT-6000-30

Q: The quadrature (90-deg out-of-phase, sine) component of a quadrature waveform.

QDR: Quad-data-rate RAM. This is solid-state RAM on the ROACH board which contains the DAC tone LUTs.

S21: The forward scattering parameter term found in a two-port scattering matrix. Here it's used to extend the VNA analogy to ROACH readout. $S_{21}V_{\text{in}}$ (in raw units) is the quantity which is saved to disk, at 488.28125 Hz, for each detector channel.

- $S_{21} = I + jQ$
- $|S_{21}| = \sqrt{I^2 + Q^2}$
- $\phi(S_{21}) = \arctan 2(Q, I)$

Target comb: A frequency comb composed of MKID carrier tones. The length of the comb will range from 50 to 750 tones.

Target sweep: A narrow (150–200 MHz) frequency sweep performed with the target comb, around each resonant frequency. Used to establish a reference sweep and reference parameters for either calculating *df* or refitting the target tones. **Upconverter**: See Modulator.

Valon 5009: A dual-channel synthesizer (23–6000 MHz) found in each ROACH slice.

Outputs:

- Channel 1: A fixed tone at 512 MHz serving as the clock signal for the ROACH-2 FPGA and MUSIC boards. The RF power of the clock signal is fixed at +6 dBm.
- Channel 2: A variable tone ranging between 284–1056 MHz serving as the LO for the mod/demodulators. The RF signal power for the LO is fixed at +5 dBm.

Inputs:

- External Reference: A fixed tone at 10 MHz generated by the OctoClock. This tone is the phase reference for the PLL used by the VCO. The RF signal power is +10 dBm.

Power: 6 VDC @ 560 mA when both outputs are enabled (default state). Power for all five Valons is supplied by a single Vicor DC/DC converter attached to the back plate of the ROACH box. Control signals to/from each Valon are routed through a USB link to the Pis. Datasheet: Valon 5009

Valon server: A TCP server running in the background on each Pi. The TCP server is created by a Python script (`valon_server_boot.py`) running in the background on each Pi. MCP is the client for the server (see: `roach.c`, `valon_client()`).

VCO: Voltage controlled oscillator.
VNA: Vector-Network Analyzer.
VNA comb: Also, search comb. A frequency comb comprised of 1,000 evenly spaced tones. This comb is used to search for resonances.
VNA sweep: A frequency sweep of the entire RF bandwidth covered by the ROACH frequency comb (512 MHz). Used as the preliminary step in 'finding the KIDS'.
UDP: User Datagram Protocol.

A.2 Data Paths and Hardware Addresses

On FC1/FC2:

ROACH root directory: `/home/fc1user/roach_flight`

On BLAST Groundstation 1:

Python analysis scripts are here: `/home/blast/detectors`

Downlinked tarballs go here: `/data/etc/downloaded_files`

Hardware addresses: See: `/etc/hosts` and `/etc/ethers` on FC1 or FC2
PowerPC addresses:

```
192.168.40.51 roach1
192.168.40.52 roach2
192.168.40.53 roach3
192.168.40.54 roach4
192.168.40.55 roach5
```

UDP packet source addresses:

```
192.168.40.71 roach1-udp
192.168.40.72 roach2-udp
192.168.40.73 roach3-udp
192.168.40.74 roach4-udp
192.168.40.75 roach5-udp
```

UDP packet destination addresses:

```
192.168.40.3 roach-udp-dest (on FC1)
192.168.40.4 roach-udp-dest (on FC2)
```

Raspberry Pi 3 addresses:

```
192.168.40.61 pi1
192.168.40.62 pi2
192.168.40.63 pi3
192.168.40.64 pi4
192.168.40.65 pi5
```

PowerPC MAC addresses:

```
02:44:01:02:0B:03 roach1
02:44:01:02:0D:17 roach2
02:44:01:02:0D:16 roach3
02:44:01:02:11:0C roach4
02:44:01:02:11:0B roach5
```

ROACH 1GbE MAC addresses:

```
02:44:02:02:0B:03 roach1-udp
02:44:02:02:0D:17 roach2-udp
02:44:02:02:0D:16 roach3-udp
02:44:02:02:11:0C roach4-udp
02:44:02:02:11:0B roach5-udp
```

Raspberry Pi 3 MAC addresses:

```
b8:27:eb:f3:28:b6 pi1
b8:27:eb:01:f1:31 pi2
b8:27:eb:2a:c9:f0 pi3
b8:27:eb:4c:2f:7d pi4
b8:27:eb:f7:16:a5 pi5
```

A.3 Normal Operation

The detector readout has two primary functions: readout and housekeeping. In this context, readout means to probe each LEKID resonance with a unique carrier tone, demodulate the phase and amplitude modulation which each pixel imprints onto it, and record a continuous timestream of that demodulated signal. Housekeeping means to monitor, as autonomously as possible, the envelope of the changes in resonator responsivity which are due to a combination of optical loading and temperature drift in the cryostat, and to either automatically correct for these changes, or indicate to the user that significant changes have occurred. The readout and housekeeping modes, as well as the user responsibilities pertaining to each, are described below.

A.3.1 Readout Mode

The readout function of the system is performed by the FPGA firmware and ROACH electronics. Therefore, it's mostly autonomous. When MCP starts, its goal is to get the ROACHs into their idle state as quickly as possible. The idle state is called streaming. To get to the streaming state, the firmware and electronics must progress through a series of initialization states (see: Owl Fields-STATE). In streaming state, the ROACHs continuously probe, demodulate and save data to disk.

On start-up, the sequence of events for each ROACH slice is:

1. Connect to ROACH KATCP server.
2. Upload FPGA firmware.
3. Initialize firmware.
4. Write VNA comb (packet streaming starts).
5. Do full-loop (see: Mult-command loops).
6. Acquire data.

If all environmental conditions, e.g., optical loading, telescope elevation and fridge temperature, were static, there would be no need to do anything further with the detector readout. Because environmental conditions are not static during flight, the resonant frequencies and quality factors of the LEKIDs drift over time, resulting in changes in detector responsivity. On the other hand, the readout carrier frequencies are static. Consequently, regular responsivity checks and corresponding adjustments to the carrier frequencies are required.

A.3.2 Housekeeping Mode

Housekeeping functions fall under two categories: Automatic and manual.

Automatic functions: Once the readout reaches its streaming state, various sequences of events are triggered during an observation. Automatic triggering is enabled by default, but can be disabled by the user.

Elevation turnarounds: At elevation turnarounds, the pointing code triggers the turnaround loop. See Multi-command loops.

Azimuth turnarounds: At azimuth turnarounds, the pointing code triggers a three-point LO chop. The frequency chop step sizes are 10 kHz for the 500 μm array, and 2.5 kHz for all other arrays. The purpose of the LO chop is to imprint a regular calibration metric for *df* into each of the channel timestreams. If the probe

tones are on or very near a channel's resonant frequency, the steps which occur in the *df* timestreams should be symmetric around the mean *df* value for that channel, with an amplitude equaling the frequency chop step size (i.e., 2.5 or 10 kHz). If the channel's resonant frequency has drifted, the chop appearing in the *df* timestreams will be asymmetric around the mean *df* value.

Manual Functions: There are several sequences of commands which can be used to manually perform housekeeping functions. During normal operation, the multi-command loops (some of which are triggered automatically) should suffice.

A.4 Multi-command Loops

Full loop: See: `full_loop`, `full_loop_all`, `full_loop_default`. The full loop is the most complete of all of the loops. It starts with writing the VNA comb, and ends with a set of target tones, in the final streaming state. The full-loop is executed automatically on system-start. It is also done when the telescope moves to a new observation target (according to the schedule files) or when the fridge has completed a new cycle. The sequence of events in the full-loop are:

1. If target comb is loaded, write VNA comb.
2. Set the input and output attenuators to the last written power-per-tone (dBm) setting (automatically adjusted for the number of tones. If no power/tone has been written by the operator, the default value is used. See: `POWER_PER_TONE`, `set_atten_calc`).
3. Do a VNA sweep.
4. Do KID-finding, write found frequencies.
5. Set the input and output attenuators to correspond to new number of tones.
6. Pulse the cal lamp.
7. Do a target sweep.
8. Calculate the offset between the probe tone frequencies and the estimated resonant frequencies of each channel.
9. Rewrite target comb.
10. Do another target sweep.
11. Calculate reference parameters for *df* calculation.
12. Pulse the cal lamp.
13. If the full-loop fails at any point, retry from the beginning (once).

Refit freqs: See: `refit_freqs`, `refit_freqs_all`. The purpose of the refit sequence is to get the target tone frequencies closer to the resonant frequency of each channel, in the event that the resonances have shifted by less than a line-width from their previous rest positions. If the resonances have shifted much more than a line-

width, a full loop will be required to re-optimize the target tone frequencies. The sequence of refit events is:

1. Pulse cal lamp (see: `check_lamp_retune`).
2. Do target sweep.
3. Calculate frequency offset between target tones and resonant frequencies, which are taken to be $f_{\mathrm{probe}} - \min(|S_{21}|)$.
4. Rewrite target tones.
5. Pulse cal lamp.

Check_lamp_retune: See: `check_lamp_retune`, `check_lamp_retune_all`. When the cal lamp is pulsed by the ROACH thread, a user-specified length of I/Q data is saved (0.5–1 s is typical) with the lamp off, and then with the lamp on. The following values are calculated and saved to the ROACH root directory on whichever flight computer is in charge. $I_{\mathrm{on}} - I_{\mathrm{off}}$, $Q_{\mathrm{on}} - Q_{\mathrm{off}}$, $df_{\mathrm{on}} - df_{\mathrm{off}}$ (see: Downlinking). These difference quantities are proportional to the detector responsivities in I, Q and df. The average of df over the channels for each ROACH is saved in the frame, and displayed in the `LAST_DF_RESPONSE` `Owl` field.
Turnaround loop: See: `turnaround_loop`, `turnaround_loop_all`. The turnaround loop is triggered at elevation turnarounds (see: `roach_allow_scan_check`). It does the following:

1. `check_lamp_retune`
2. `refit_freqs`
3. `check_lamp_retune`

Find kids loop: See: `find_kids_loop`, `find_kids_loop_all`. This loop is intended for use when the operator needs to troubleshoot the VNA sweep and KID-finding algorithm. Under normal circumstances, there should be no need to run this loop. The sequence of events is:

1. Write VNA comb.
2. Do VNA sweep.
3. Do KID-finding, without writing newly found target frequencies.

If the operator would like to write the target frequencies, see: `load_freqs`, `load_freqs_all`.

A.5 Owl Frontend

During regular operation of the readout, real-time system status and health information is displayed in the fields of the roach.owl file. The information in these fields should provide the operator with all they need to successfully navigate through the various normal modes of operation, as well as to perform basic debugging and

sanity checks. The content of the fields is sampled at multiple rates, as indicated in the following list:

IS_STREAMING: Indicates that UDP data is being received by MCP. The data streaming status is checked automatically every 10 seconds by each ROACH commanding thread. If the data stream stops, `DATA_STREAM_ERROR` should go high.
STATE: One of the following: `BOOT`, `CONNECTED`, `PROGRAMMED`, `CONFIGURED`, `STREAMING`. `STREAMING` is the final state, where the ROACHs should spend 99% of their time. For most ROACHs, STREAMING should be reached within a minute or so after a cold start. ROACH1's boot sequence lags behind the others by 30 s to 1 min, so expect it to spend more time in the earlier states during start-up.

- `BOOT`: The start-up state. In `BOOT` state, the command thread attempts to establish a connection to the KATCP server running on each PPC. After several failed connection attempts, `KATCP_CONNECT_ERROR` is set high to alert the operator. The ROACH command thread will then sleep for a few seconds before retrying the connection.
  Troubleshooting:
  If a ROACH seems to be stuck in the BOOT state (i.e., a few minutes have passed with no joy), it's possible that the ROACH is powered down. If able, try the following:
  1. Check to see if the ROACH PPC responds to ping from either the flight computers or groundstation.
  2. If ping succeeds, it could mean that the ROACH is up, but the ROACH command thread is in a bad state. Try to:
     a) Send the `reset_roach` command.
     b) Restart MCP.

     If ping fails, it's likely that the ROACH is down. Try to power cycle the ROACHs with commands: `roach_cycle`, `roaches_off`, `roach_on`.
- `CONNECTED`: After successful connection to the KATCP server, the command thread attempts to upload the FPGA firmware (fpg file). If this fails, `FIRMWARE_UPLOAD_FAIL` will go high, and MCP will continue to attempt to upload the fpg file. If it succeeds, `HAS_FIRMARE` goes high.
  Troubleshooting:
  It's uncommon for a firmware upload attempt to fail. If this occurs, we suggest following the troubleshooting protocol for BOOT state. If the flight computers are accessible by SSH, SG and AS have a method of manually uploading the fpg using a Python script located in: `/data/etc/blast/roachPython`.
- `PROGRAMMED`: After successful firmware upload, MCP programs the firmware registers and prepares to write the frequency comb. A critical step in this process is the 'calibration' of the QDR RAM. Calibration is performed by a Python script, called by MCP. Successful calibration results in:

* HAS_QDR_CAL = 1
* FPGA_CLOCK_FREQUENCY reads between 255–259 MHz (256 MHz is desired). Failure to calibrate the QDR results in:
* QDR_CAL_FAIL flag = 1.
* The FPGA_CLOCK_FREQUENCY will be out of range.

Troubleshooting:
QDR calibration failures almost always stem from problems with the 512 MHz clock signal provided by the Valon. Try to:

1. Power cycle the Pi. See: reboot_pi
2. Power cycle the ROACHs.
3. If power cycling doesn't help, it's possible that the 10 MHz external reference signal has dropped out. See: roach_set_extref.

If none of the above works, contact SG or AS.

– CONFIGURED: In this state, MCP creates the VNA comb, uploads it to the FPGA, turns on the DACs, and then checks that the data socket is receiving UDP packets. While tones are being written, IS_WRITING goes high. Note: Data streaming is automatically triggered within the firmware after tone writing succeeds. If all goes well, you should see that:

1. CURRENT_NTONES = 1000
2. HAS_VNA_TONES = 1
3. STATE = STREAMING

Troubleshooting: If MCP is not receiving data, DATA_STREAM_ERROR will go high (see: IS_STREAMING), and MCP will reset the state to BOOT in order to restart the initialization process. If tone writing has succeeded (see: CURRENT_NTONES), but no data is detected, it is likely that the problem lies with MCP and/or the flight computers. If possible, try to:

1. SSH into FC1/FC2. To check if data is making it to the ethernet card, do:

```
> sudo tcpdump -i eth0:0 | grep udp
```

   If UDP packets are arriving at roach-udp-dest from the ROACH in question, the reason for the error likely has to do with MCP.
2. Restart MCP.

– STREAMING: This is the final state, in which all normal flight-mode readout operations take place. In this state, IS_STREAMING will remain high, and MCP will continue to monitor the streaming status on 10 second intervals.

DATA_STREAM_ERROR: Indicates that UDP data is not being received by the MCP data socket. The data stream is checked every 10 seconds.
ALLOW_AUTO_TURNAROUND: See: roach_allow_scan_check_all. Indicates that automatic triggering of the elevation turnaround loop is enabled. State is high by default.

ALLOW_CHOP_LO: See: `enable_chop_lo_all`. Enables automatic azimuth turnaround LO chops. State is high by default.
CURRENT_NTONES: Current number of tones running in the system. When `HAS_VNA_TONES` = 1, `CURRENT_NTONES` should be 1000. Once channels have been assigned, this number will range from 50–750, and should match `N_CHANNELS`.
N_CHANNELS: Current number of channels, which should match `CURRENT_NTONES`.
PREV_N_CHANNELS: Previous number of channels. This is shown to provide a reference for the current number of tones, or channels. There may be slight variations in `N_CHANNELS` between consecutive VNA sweeps. If the difference between `N_CHANNELS` and `PREV_N_CHANNELS` is significant (e.g., > 50), possible reasons include but are not limited to: The optical loading has changed, or there is spurious noise in the system that is interfering with the channel finding algorithm.
HAS_VNA_TONES: Indicates that the VNA/search comb is loaded. When the system is running the search comb (as at the beginning of the full or find kids loop), `CURRENT_NTONES` should be 1000. Once channels have been assigned, `HAS_VNA_TONES` should go low. To load the VNA comb, see: `reload_vna_freqs`
HAS_TARG_TONES: Indicates that the target comb is loaded. After channels have been assigned, this field should equal 1, and `HAS_VNA_TONES` should equal 0.
POWER_PER_TONE: Indicates the estimated power level per tone at the output of the ROACH, in units of dBm (0 dBm = 1 mW). The desired power level per tone is chosen as an option when sending certain commands (see: `full_loop`, `turnaround_loop`). Also see: `set_attens_calc`.
SET_ATTEN_OUT: dB of output attenuation (ROACH output) requested by command (see, e.g., `set_attens`).
READ_ATTEN_OUT: dB of output attenuation (ROACH output) currently running (see `read_attens`).
SET_ATTEN_IN: dB of input attenuation (ROACH input) requested by command (see, e.g., `set_attens`).
READ_ATTEN_IN: dB of input attenuation (ROACH input) currently running (see `read_attens`).
ADC_RMS_I: An estimate of the root-mean-square voltage level at the I-ADC input, in millivolts. This level will vary as a function of the attenuator settings and number of tones. Under typical operating conditions it should not exceed 300 mV (over 500 mV will clip and/or saturate the ADC). To refresh this value, see: `show_adc_rms`.
ADC_RMS_Q: Same as above, for the Q-ADC.
LO_CENTER_FREQ: The rest frequency of the LO, in megahertz. During normal operation the LO should only deviate from this frequency during sweeps and LO chops. Center frequencies are hard-coded in roach.c.
LO_FREQ_READ: The current LO frequency, in megahertz. See: `read_lo`.

DOING_FULL_LOOP: Indicates that MCP is doing a full-loop. See: Multi-command loops. Also: `full_loop`, `full_loop_all`.
DOING_FIND_KIDS_LOOP: Indicates that MCP is doing a find kids loop. See: Multi-command loops. Also: `find_kids_loop`, `find_kids_loop_all`.
DOING_TURNAROUND_LOOP: Indicates that MCP is doing a turnaround loop. See: Multi-command loops. Also: `turnaround_loop`, `turnaround_loop_all`.
IS_WRITING: Indicates that a new frequency comb is being written.
IS_SWEEPING: Indicates that a sweep is progress.
IS_FINDING_KIDS: Indicates that MCP is running the channel assignment/KID finding algorithm. This process takes 15–20 seconds to complete.
IS_COMPRESSING: Indicates that data is being compressed in the background. Data is compressed to tar.gz for downlink (see: Downlinking data. Also: `compress_roach_data`).
FPGA_CLOCK_FREQ_MHZ: The estimated FPGA clock frequency reported after QDR RAM calibration, when the ROACH is in the `PROGRAMMED` state. The FPGA derives its clock from the 512 MHz Valon signal, which it halves. Normal clock frequencies range from 255–257 Hz. Anything below or above this range indicates that there is either a problem with the Valon or the external reference, and the ROACH state will not proceed past `PROGRAMMED`. For help troubleshooting this scenario, see `PROGRAMMED` above.
HAS_VNA_SWEEP: Indicates that a VNA sweep has been completed.
HAS_TARG_SWEEP: Indicates that target sweep has been completed.
HAS_REF_PARAMS: See: *df* calculation, reference parameters. Indicates that the reference parameters required to calculate *df* have been calculated from the last target sweep, and saved to disk.
TONE_FINDING_ERROR: Indicates that an error has occurred during channel assignment. The most likely sources of the problem are either too many or too few channels having been found. See: `set_minkids`, `set_max_nkids`. Return codes:

1. Number of channels found < currently allowed min
2. Number of channels found > currently allowed max
3. Failure to find any channels

SWEEP_FAIL: Indicates that a sweep in-progress has failed. Sweep fails are usually caused by Pi read/write errors. Check `PI_ERROR_COUNT`. Depending on the command which was used to call the sweep, there may or may not be an automatic retry.
FIRMWARE_UPLOAD_FAIL: Indicates that the firmware has failed to upload. See: `STATE/BOOT`.
QDR_CAL_FAIL: Indicates that the QDR calibration has failed. See: `STATE/PROGRAMMED`.
KATCP_CONNECT_ERROR:

HAS_QDR_CAL: Indicates that the QDR RAM has been calibrated.
PI_ERROR_COUNT: Indicates that either the Valon or attenuator clients have either failed to contact the Pi server, or received bad data from the server. If the error count reaches 10, the Pi will automatically be rebooted, which reloads the Valon and attenuator server. See: `reboot_pi`.
PI_REBOOT: Indicates that a Pi reboot command has been issued, either by MCP or by the operator.
LAST_DF_RESPONSE: See: `check_lamp_retune`, `check_lamp_retune_all`. The average *df* lamp response, in Hz, over all channels.
DF_RETUNE_THRESH: The *df* retune threshold, in Hz. If `N_OUTOFRANGE` channels have a lamp response lower than this threshold, retuning is recommended, but neither triggered automatically nor required. See: `set_df_diff_threshold`.
N_OUT_OF_RANGE_THRESH: The number of channels which need to have a *df* response lower than `DF_RETUNE_THRESHOLD` to merit a retune. See: `set_noutofrange_thresh`.
N_OUTOFRANGE: The number of channels with *df* lamp responses which are out of range.
IS_CHOPPING_LO: Indicates that the LO is being chopped. If `AUTO_CHOP_LO` = 1, this field will go high during azimuth turnarounds. If the operator wants to chop the LO by command (see: `chop_lo`), they must first disable the automatic elevation turnaround retuning. See: `roach_disallow_scan_check_all`.
HAS_LAMP_CONTROL: Indicates that the ROACH has exclusive control over the cal lamp. To set control, see: `roach_has_lamp`.
WAITING_FOR_LAMP: When high, indicates that a ROACH is ready and waiting to save lamp response data. The lamp will be pulsed once the ROACH with lamp control is also ready, and all ROACHs will save data simultaneously.
FULL_LOOP_FAIL: When high, indicates that an error has occurred during the full-loop.
TURNAROUND_LOOP_FAIL: When high, indicates that an error has occurred during the full loop.
EXT_REF: Indicates which reference is being used by the Valon synthesizers: 1 (external, from the OctoClock), or 0 (internal, from the Valons). See: Valon, OctoClock.
PI_TEMP: The temperature of a Pi CPU, in C. See `read_pi_temps`.

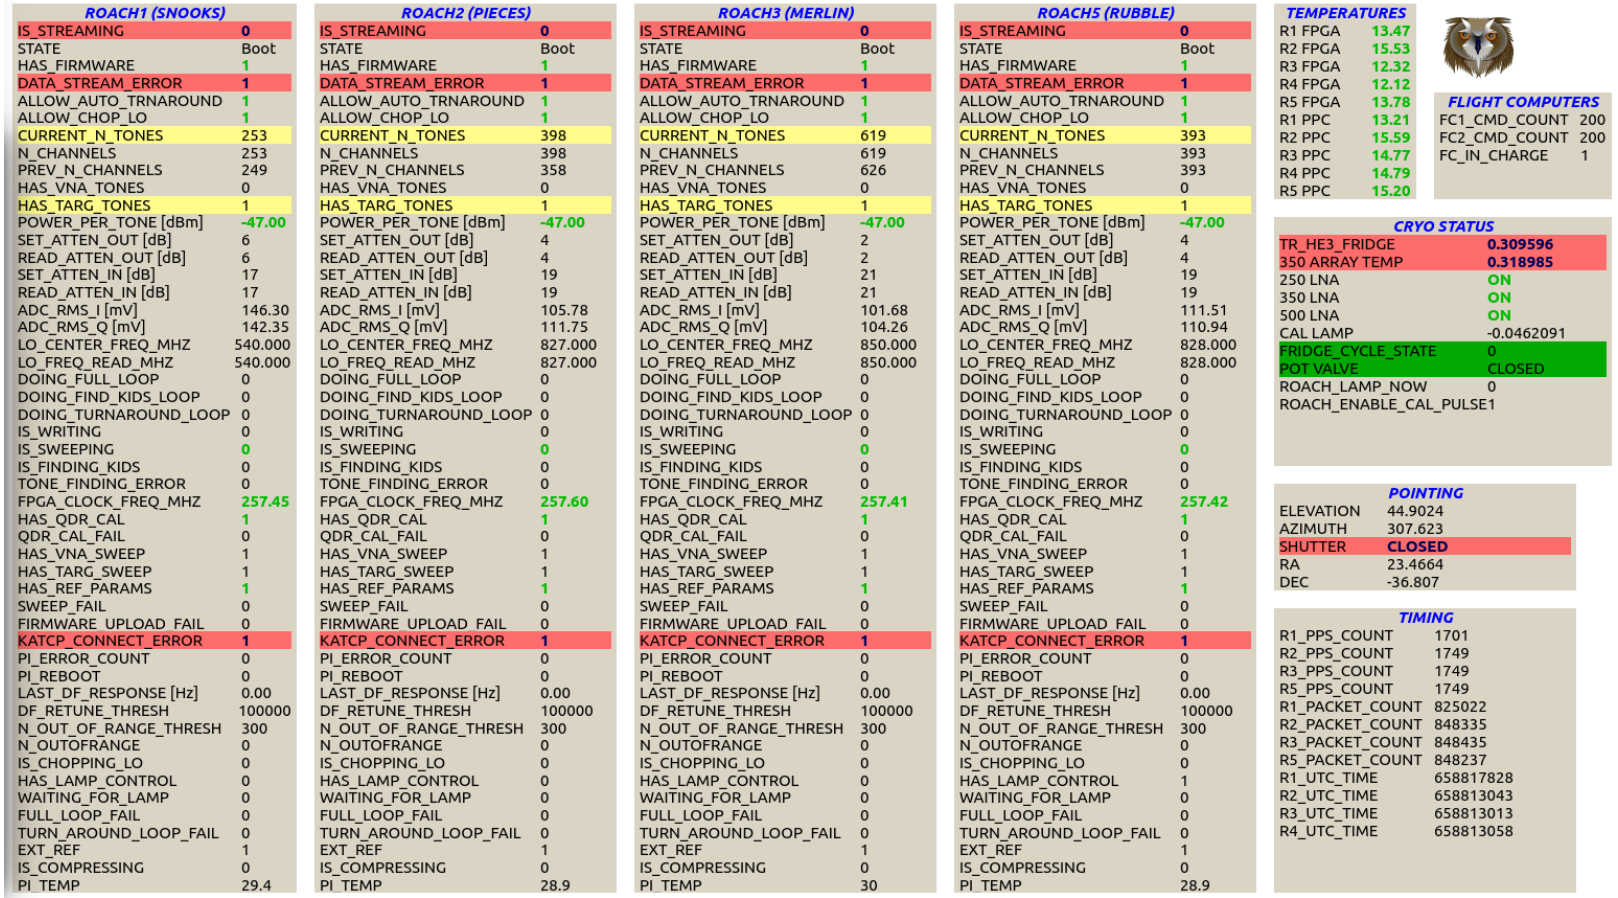

ROACH1 (SNOOKS)	
IS_STREAMING	0
STATE	Boot
HAS_FIRMWARE	1
DATA_STREAM_ERROR	1
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	253
N_CHANNELS	253
PREV_N_CHANNELS	249
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	6
READ_ATTEN_OUT [dB]	6
SET_ATTEN_IN [dB]	17
READ_ATTEN_IN [dB]	17
ADC_RMS_I [mV]	146.30
ADC_RMS_Q [mV]	142.35
LO_CENTER_FREQ_MHZ	540.000
LO_FREQ_READ_MHZ	540.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.45
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	1
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	29.4

ROACH2 (PIECES)	
IS_STREAMING	0
STATE	Boot
HAS_FIRMWARE	1
DATA_STREAM_ERROR	1
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	398
N_CHANNELS	398
PREV_N_CHANNELS	358
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	4
READ_ATTEN_OUT [dB]	4
SET_ATTEN_IN [dB]	19
READ_ATTEN_IN [dB]	19
ADC_RMS_I [mV]	105.78
ADC_RMS_Q [mV]	111.75
LO_CENTER_FREQ_MHZ	827.000
LO_FREQ_READ_MHZ	827.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.60
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	1
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	28.9

ROACH3 (MERLIN)	
IS_STREAMING	0
STATE	Boot
HAS_FIRMWARE	1
DATA_STREAM_ERROR	1
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	619
N_CHANNELS	619
PREV_N_CHANNELS	626
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	2
READ_ATTEN_OUT [dB]	2
SET_ATTEN_IN [dB]	21
READ_ATTEN_IN [dB]	21
ADC_RMS_I [mV]	101.68
ADC_RMS_Q [mV]	104.26
LO_CENTER_FREQ_MHZ	850.000
LO_FREQ_READ_MHZ	850.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.41
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	1
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	30

ROACH5 (RUBBLE)	
IS_STREAMING	0
STATE	Boot
HAS_FIRMWARE	1
DATA_STREAM_ERROR	1
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	393
N_CHANNELS	393
PREV_N_CHANNELS	393
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	4
READ_ATTEN_OUT [dB]	4
SET_ATTEN_IN [dB]	19
READ_ATTEN_IN [dB]	19
ADC_RMS_I [mV]	111.51
ADC_RMS_Q [mV]	110.94
LO_CENTER_FREQ_MHZ	828.000
LO_FREQ_READ_MHZ	828.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.42
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	1
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	1
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	28.9

TEMPERATURES	
R1 FPGA	13.47
R2 FPGA	15.53
R3 FPGA	12.32
R4 FPGA	12.12
R5 FPGA	13.78
R1 PPC	13.21
R2 PPC	15.59
R3 PPC	14.77
R4 PPC	14.79
R5 PPC	15.20

FLIGHT COMPUTERS	
FC1_CMD_COUNT	200
FC2_CMD_COUNT	200
FC_IN_CHARGE	1

CRYO STATUS	
TR_HE3_FRIDGE	0.309596
350 ARRAY TEMP	0.318985
250 LNA	ON
350 LNA	ON
500 LNA	ON
CAL LAMP	-0.0462091
FRIDGE_CYCLE_STATE	0
POT VALVE	CLOSED
ROACH_LAMP_NOW	0
ROACH_ENABLE_CAL_PULSE1	

POINTING	
ELEVATION	44.9024
AZIMUTH	307.623
SHUTTER	CLOSED
RA	23.4664
DEC	-36.807

TIMING	
R1_PPS_COUNT	1701
R2_PPS_COUNT	1749
R3_PPS_COUNT	1749
R5_PPS_COUNT	1749
R1_PACKET_COUNT	825022
R2_PACKET_COUNT	848335
R3_PACKET_COUNT	848435
R5_PACKET_COUNT	848237
R1_UTC_TIME	658817828
R2_UTC_TIME	658813043
R3_UTC_TIME	658813013
R4_UTC_TIME	658813058

Figure 128. Example `Owl` screen when the ROACHs are powered down.

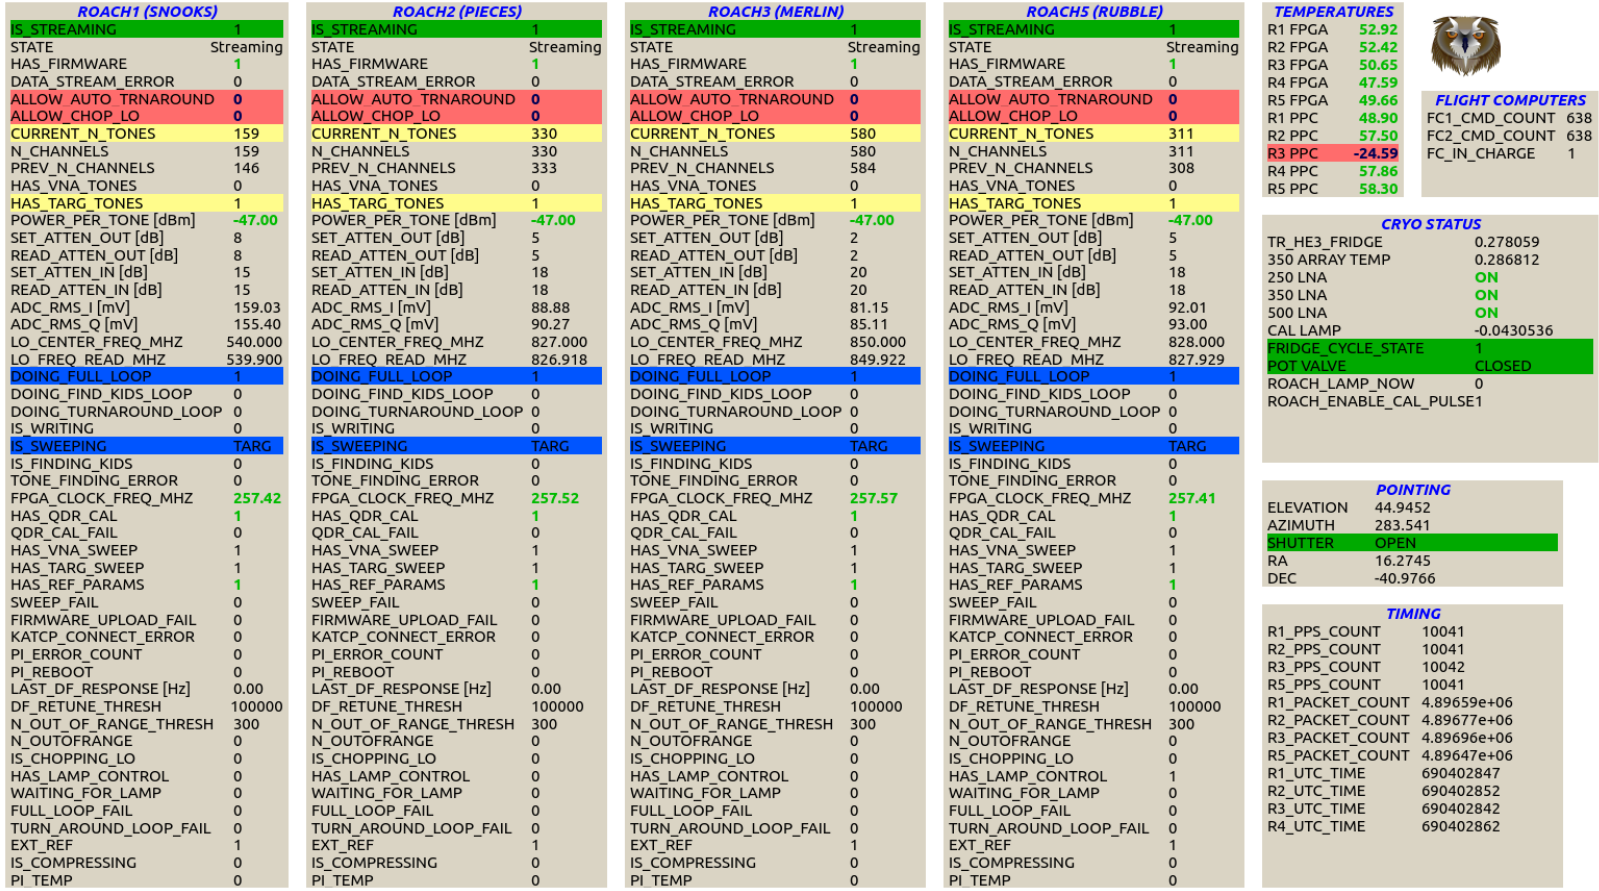

ROACH1 (SNOOKS)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	0
ALLOW_CHOP_LO	0
CURRENT_N_TONES	159
N_CHANNELS	159
PREV_N_CHANNELS	146
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	8
READ_ATTEN_OUT [dB]	8
SET_ATTEN_IN [dB]	15
READ_ATTEN_IN [dB]	15
ADC_RMS_I [mV]	159.03
ADC_RMS_Q [mV]	155.40
LO_CENTER_FREQ_MHZ	540.000
LO_FREQ_READ_MHZ	539.900
DOING_FULL_LOOP	1
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	TARG
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.42
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	0

ROACH2 (PIECES)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	0
ALLOW_CHOP_LO	0
CURRENT_N_TONES	330
N_CHANNELS	330
PREV_N_CHANNELS	333
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	5
READ_ATTEN_OUT [dB]	5
SET_ATTEN_IN [dB]	18
READ_ATTEN_IN [dB]	18
ADC_RMS_I [mV]	88.88
ADC_RMS_Q [mV]	90.27
LO_CENTER_FREQ_MHZ	827.000
LO_FREQ_READ_MHZ	826.918
DOING_FULL_LOOP	1
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	TARG
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.52
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	0

ROACH3 (MERLIN)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	0
ALLOW_CHOP_LO	0
CURRENT_N_TONES	580
N_CHANNELS	580
PREV_N_CHANNELS	584
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	2
READ_ATTEN_OUT [dB]	2
SET_ATTEN_IN [dB]	20
READ_ATTEN_IN [dB]	20
ADC_RMS_I [mV]	81.15
ADC_RMS_Q [mV]	85.11
LO_CENTER_FREQ_MHZ	850.000
LO_FREQ_READ_MHZ	849.922
DOING_FULL_LOOP	1
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	TARG
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.57
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	0

ROACH5 (RUBBLE)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	0
ALLOW_CHOP_LO	0
CURRENT_N_TONES	311
N_CHANNELS	311
PREV_N_CHANNELS	308
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	5
READ_ATTEN_OUT [dB]	5
SET_ATTEN_IN [dB]	18
READ_ATTEN_IN [dB]	18
ADC_RMS_I [mV]	92.01
ADC_RMS_Q [mV]	93.00
LO_CENTER_FREQ_MHZ	828.000
LO_FREQ_READ_MHZ	827.929
DOING_FULL_LOOP	1
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	TARG
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.41
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	1
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	0

TEMPERATURES	
R1 FPGA	52.92
R2 FPGA	52.42
R3 FPGA	50.65
R4 FPGA	47.59
R5 FPGA	49.66
R1 PPC	48.90
R2 PPC	57.50
R3 PPC	-24.59
R4 PPC	57.86
R5 PPC	58.30

FLIGHT COMPUTERS	
FC1_CMD_COUNT	638
FC2_CMD_COUNT	638
FC_IN_CHARGE	1

CRYO STATUS	
TR_HE3_FRIDGE	0.278059
350 ARRAY TEMP	0.286812
250 LNA	ON
350 LNA	ON
500 LNA	ON
CAL LAMP	-0.0430536
FRIDGE_CYCLE_STATE	1
POT VALVE	CLOSED
ROACH_LAMP_NOW	0
ROACH_ENABLE_CAL_PULSE	1

POINTING	
ELEVATION	44.9452
AZIMUTH	283.541
SHUTTER	OPEN
RA	16.2745
DEC	-40.9766

TIMING	
R1_PPS_COUNT	10041
R2_PPS_COUNT	10041
R3_PPS_COUNT	10042
R5_PPS_COUNT	10041
R1_PACKET_COUNT	4.89659e+06
R2_PACKET_COUNT	4.89677e+06
R3_PACKET_COUNT	4.89696e+06
R5_PACKET_COUNT	4.89647e+06
R1_UTC_TIME	690402847
R2_UTC_TIME	690402852
R3_UTC_TIME	690402842
R4_UTC_TIME	690402862

Figure 129. Example `Owl` screen when the ROACHs are performing a full loop.

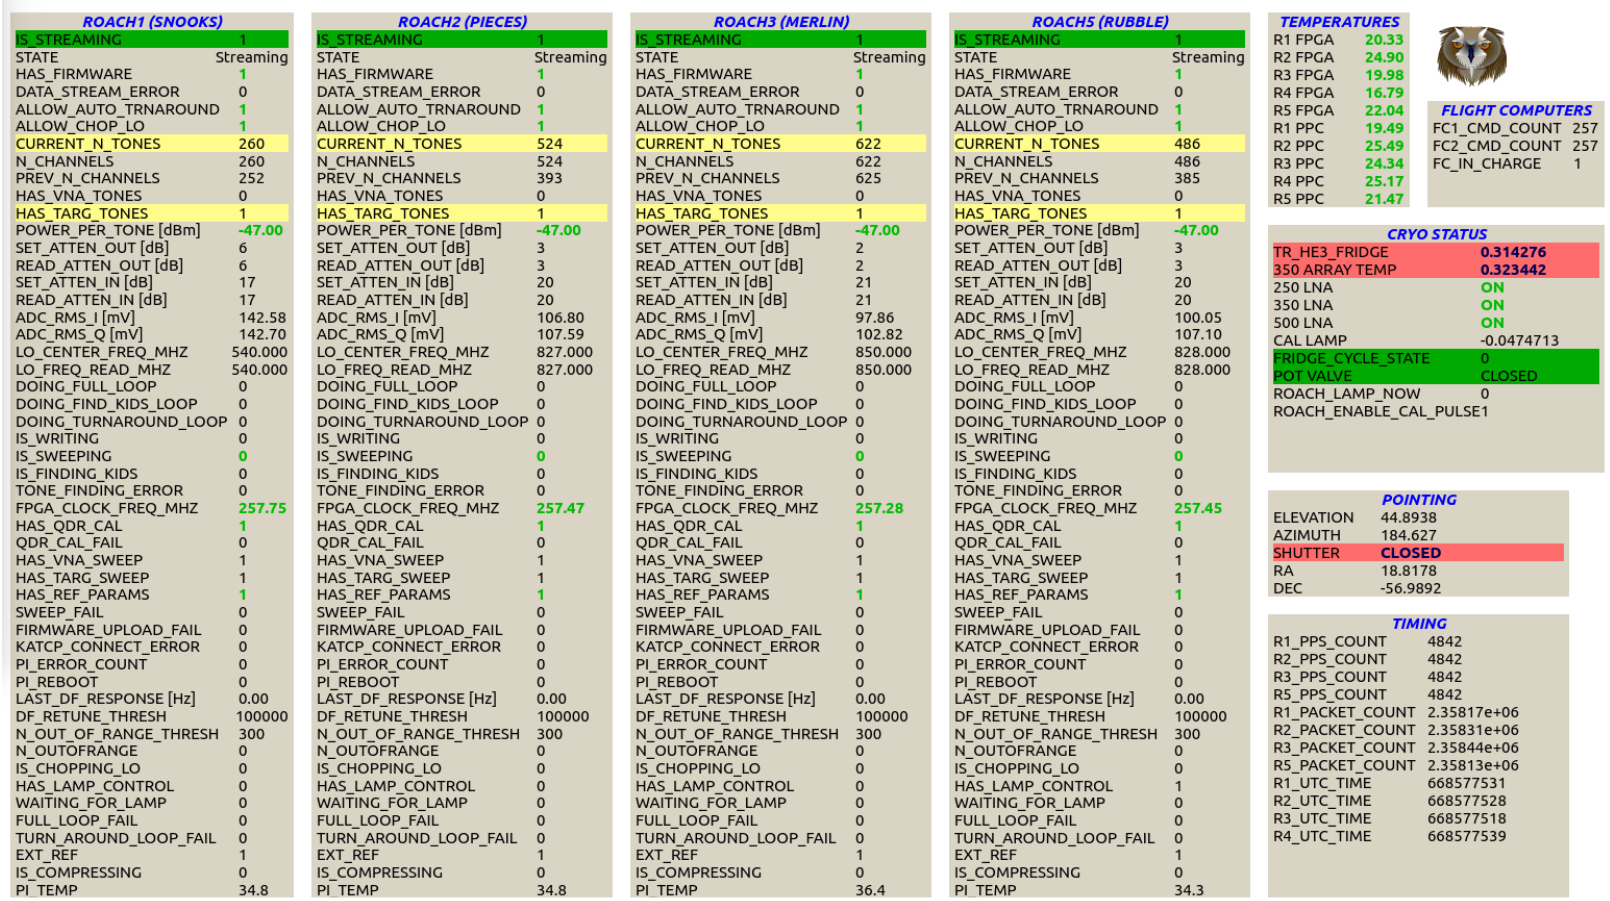

ROACH1 (SNOOKS)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	260
N_CHANNELS	260
PREV_N_CHANNELS	252
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	6
READ_ATTEN_OUT [dB]	6
SET_ATTEN_IN [dB]	17
READ_ATTEN_IN [dB]	17
ADC_RMS_I [mV]	142.58
ADC_RMS_Q [mV]	142.70
LO_CENTER_FREQ_MHZ	540.000
LO_FREQ_READ_MHZ	540.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.75
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	34.8

ROACH2 (PIECES)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	524
N_CHANNELS	524
PREV_N_CHANNELS	393
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	3
READ_ATTEN_OUT [dB]	3
SET_ATTEN_IN [dB]	20
READ_ATTEN_IN [dB]	20
ADC_RMS_I [mV]	106.80
ADC_RMS_Q [mV]	107.59
LO_CENTER_FREQ_MHZ	827.000
LO_FREQ_READ_MHZ	827.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.47
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	34.8

ROACH3 (MERLIN)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	622
N_CHANNELS	622
PREV_N_CHANNELS	625
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	2
READ_ATTEN_OUT [dB]	2
SET_ATTEN_IN [dB]	21
READ_ATTEN_IN [dB]	21
ADC_RMS_I [mV]	97.86
ADC_RMS_Q [mV]	102.82
LO_CENTER_FREQ_MHZ	850.000
LO_FREQ_READ_MHZ	850.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.28
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	0
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	36.4

ROACH5 (RUBBLE)	
IS_STREAMING	1
STATE	Streaming
HAS_FIRMWARE	1
DATA_STREAM_ERROR	0
ALLOW_AUTO_TRNAROUND	1
ALLOW_CHOP_LO	1
CURRENT_N_TONES	486
N_CHANNELS	486
PREV_N_CHANNELS	385
HAS_VNA_TONES	0
HAS_TARG_TONES	1
POWER_PER_TONE [dBm]	-47.00
SET_ATTEN_OUT [dB]	3
READ_ATTEN_OUT [dB]	3
SET_ATTEN_IN [dB]	20
READ_ATTEN_IN [dB]	20
ADC_RMS_I [mV]	100.05
ADC_RMS_Q [mV]	107.10
LO_CENTER_FREQ_MHZ	828.000
LO_FREQ_READ_MHZ	828.000
DOING_FULL_LOOP	0
DOING_FIND_KIDS_LOOP	0
DOING_TURNAROUND_LOOP	0
IS_WRITING	0
IS_SWEEPING	0
IS_FINDING_KIDS	0
TONE_FINDING_ERROR	0
FPGA_CLOCK_FREQ_MHZ	257.45
HAS_QDR_CAL	1
QDR_CAL_FAIL	0
HAS_VNA_SWEEP	1
HAS_TARG_SWEEP	1
HAS_REF_PARAMS	1
SWEEP_FAIL	0
FIRMWARE_UPLOAD_FAIL	0
KATCP_CONNECT_ERROR	0
PI_ERROR_COUNT	0
PI_REBOOT	0
LAST_DF_RESPONSE [Hz]	0.00
DF_RETUNE_THRESH	100000
N_OUT_OF_RANGE_THRESH	300
N_OUTOFRANGE	0
IS_CHOPPING_LO	0
HAS_LAMP_CONTROL	1
WAITING_FOR_LAMP	0
FULL_LOOP_FAIL	0
TURN_AROUND_LOOP_FAIL	0
EXT_REF	1
IS_COMPRESSING	0
PI_TEMP	34.3

TEMPERATURES	
R1 FPGA	20.33
R2 FPGA	24.90
R3 FPGA	19.98
R4 FPGA	16.79
R5 FPGA	22.04
R1 PPC	19.49
R2 PPC	25.49
R3 PPC	24.34
R4 PPC	25.17
R5 PPC	21.47

FLIGHT COMPUTERS	
FC1_CMD_COUNT	257
FC2_CMD_COUNT	257
FC_IN_CHARGE	1

CRYO STATUS	
TR_HE3_FRIDGE	0.314276
350 ARRAY TEMP	0.323442
250 LNA	ON
350 LNA	ON
500 LNA	ON
CAL LAMP	-0.0474713
FRIDGE_CYCLE_STATE	0
POT VALVE	CLOSED
ROACH_LAMP_NOW	0
ROACH_ENABLE_CAL_PULSE	1

POINTING	
ELEVATION	44.8938
AZIMUTH	184.627
SHUTTER	CLOSED
RA	18.8178
DEC	-56.9892

TIMING	
R1_PPS_COUNT	4842
R2_PPS_COUNT	4842
R3_PPS_COUNT	4842
R5_PPS_COUNT	4842
R1_PACKET_COUNT	2.35817e+06
R2_PACKET_COUNT	2.35831e+06
R3_PACKET_COUNT	2.35844e+06
R5_PACKET_COUNT	2.35813e+06
R1_UTC_TIME	668577531
R2_UTC_TIME	668577528
R3_UTC_TIME	668577518
R4_UTC_TIME	668577539

Figure 130. Example `Owl` screen when the ROACHs are streaming, with target tones.

A.6 Commands

The following is a description of the core set of commands which can be used to control the readout during normal lab, or flight operations. Syntax is given for manually inputting the commands into `blastcmd`, although during flight they will most likely be executed using `Cow`.

*CAUTION— don't execute these commands unless you have a plan in mind.
**EXTRA CAUTION- improper use of these commands will lead to a decrease in the number of future BLAST-TNG science publications. If you think you might need to use one, first try to contact SG or AS.

Note: <which roach> is an integer in (1,5).

A.6.1 Commands Which Execute Sequences Of Events

full_loop: Triggers a full-loop. See: Multi-command loops. Arguments:

- Which ROACH (1–5)
- KID-finding mode (see: KID-finding) 1 = default 2 = use configurable parameters
- Target power/tone, in dBm (-47 dBm typical for full-loop).

Syntax: `full_loop <which roach> <KID-finding mode> <power/tone in dBm>`
full_loop_all: Triggers a full-loop for all ROACHs. See: Multi-command loops. Arguments: See `full_loop`
Syntax: `full_loop_all <KID-finding mode <power/tone in dBm>`
full_loop_default: Triggers a full-loop for all ROACHs, using previously chosen KID-finding mode and power/tone. See: Multi-command loops, `full_loop`, `full_loop_all`.
Syntax: `full_loop_default`, `full_loop_default_all`
turnaround_loop: Triggers an elevation turnaround loop for a single ROACH. See: Normal operation. Arguments:

- Which ROACH (1–5)
- KID-finding mode (see: KID-finding) 1 = Default 2 = Use configurable parameters
- Target power/tone, in dBm (-47 dBm typical).

Syntax: `turnaround_loop <which roach> <KID-finding mode> <power/tone in dBm>`

turnaround_loop_all: Triggers an elevation turnaround loop for all ROACHs. See: Multi-command loops. Arguments: See `turnaround_loop`
Syntax: `turnaround_loop_all <KID-finding mode> <power/tone in dBm>`
find_kids_loop: Triggers the find-kids-loop for a single ROACH. Arguments:

- Which ROACH (1–5)
- KID-finding mode (see: KID-finding) 1 = Default 2 = Use configurable parameters
- Target power/tone, in dBm (-47 dBm typical).

Syntax: `find_kids_loop <which roach>`
find_kids_loop_all: Triggers the find-kids-loop for all ROACHs. Arguments: See `find_kids_loop`
Syntax: `find_kids_loop_all <KID-finding mode> <power/tone in dBm>`
refit_freqs: Does a target sweep, calculates the *df* between carrier and resonant frequencies, adjusts the tone frequencies, and re-sweeps to store a new reference sweep and set of reference parameters.
Syntax: `refit_freqs <which roach> 1`
refit_freqs_all: As above, for all ROACHs.
Syntax: `refit_freqs_all`

A.6.2 Commands Which Are Useful For Data Visualization

zero_df_all: Calculates (for all ROACHs) the mean *df* value from the last 50 values, and subtracts it from all future *df* values. This centers the *df* timestreams around DC, which is useful for measuring frequency drift as function of time in the KST detector timestreams (see: `Mole`).
Syntax: `zero_dfs`

A.6.3 Commands That Enable Triggering Of Periodic Events

enable_chop_lo_all: If set high, allows the LO to perform a three-point frequency chop at azimuth turnarounds.
Syntax: `enable_chop_lo_all <0--1>`
roach_allow_scan_check: Enables the automatic elevation turnaround-loop for a single ROACH.
roach_allow_scan_check_all: Enables the automatic elevation turnaround-loop for all ROACHs.
Syntax: `roach_allow_scan_check_all`

roach_disallow_scan_check: Disables the automatic elevation turnaround-loop for a single ROACH.
Syntax: `roach_disallow_scan_check <which roach>`
roach_disallow_scan_check_all: Disables the automatic elevation turnaround-loop for all ROACHs.
Syntax: `roach_disallow_scan_check_all`

A.6.4 Tone Writing Commands

*__reload_vna_all__: Loads/reloads the VNA comb to all ROACHs. Use only when debugging.
Syntax: `reload_vna_all`
*__load_freqs__: Loads the last set of target frequencies which was found during channel assignment. Frequencies are read from the file `bb_target_freqs.dat`, stored in the root directory of each ROACH.
Syntax: `load_freqs <which roach>`
*__load_freqs_all__: Same as above, for all ROACHs.
Syntax: `load_freqs_all`

A.6.5 LO Commands

center_lo: Sets the LO frequency for a single Valon to its nominal center frequency (hard-coded in `roach.c`).
Syntax: `center_lo <which roach>`
center_lo_all: Same as `center_lo`, but for all ROACHs.
Syntax: `center_lo_all`
chop_lo: See: Housekeeping/automatic. Does a three-point LO frequency chop about the LO center frequency.
Syntax: `chop_lo <which roach>` For the command to execute, the following conditions must be true:

1. `ALLOW_AUTO_CHOP` = 1 (see `enable_chop_lo`)
2. `ALLOW_AUTO_TURNAROUND` = 0 (see `roach_allow_scan_check`)
3. `IS_SWEEPING` = 0

*__offset_lo__: Offsets the LO frequency by specified amount in Hz, relative to center frequency.
Syntax: `offset_lo <which roach> <some number of Hz (can be <0)>`
Note: If you offset Roach1's LO by 1000 Hz, e.g.,

```
> offset_lo 1 1000
```

and then decide that you actually wanted to offset it by 2000 Hz, there is no need to recenter the LO. Just do:

```
> offset_lo 1 2000
```

The LO frequency will now be +2000 Hz above its center frequency. However, before resuming normal operation, you must manually recenter the LO, by using the `center_lo` command.
*__offset_lo_all__: Same as `offset_lo`, but for all ROACHs.
Syntax: `offset_lo_all`
__recenter_lo__: Returns the LO frequency to the default center frequency (hard-coded for each ROACH in `roach.c`).
Syntax: `recenter_lo <which roach>`
__recenter_lo_all__: Same as `recenter_lo`, but for all ROACHs.
Syntax: `recenter_lo_all`
__read_lo__: Read the current LO frequency. Frequency is returned in MHz.
Syntax: `read_lo <which roach>`
__set_lo_MHz__: Set the current LO frequency, in MHz.
Syntax: `set_lo_MHz <which roach>`

A.6.6 Tone Power/Attenuator Commands

__read_attens__: Reads the current attenuator settings. The result, in dB, will appear in the Owl fields: `READ_ATTEN_OUT`, `READ_ATTEN_IN`.
Syntax: `read_attens <which roach>`
__show_adc_rms__: Reads the current root-mean-square voltage level of the signals at the input to the ADCs (I and Q), in millivolts. The values are displayed in the `ADC_RMS_I` and `ADC_RMS_Q` Owl fields.
Syntax: `show_adc_rms <which roach>`
__set_attens__: Set the RUDATs for a single ROACH, in dB. The requested values will appear in the Owl fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
Syntax: `set_attens <which roach> <float output atten, dB> <float input atten, dB>`
__set_attens_default__: Sets the RUDATs to their default values, which are defined in roach.c as 4 dB (output), 19 dB (input). The requested values will appear in the Owl fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
Syntax: `set_attens_default <which roach>`
__set_attens_default_all__: Same as above, or all ROACHs. The set values will appear in the Owl fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
__set_attens_calc__: Automatically set the output attenuator to target a power/tone, in dBm. The calculated power level takes into account the current number of tones.

The input attenuator value is chosen to enforce a total attenuation of 23 dB between the input and output. The default is -47 dBm/tone. The set values will appear in the Owl fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
Syntax: `set_attens_calc <which roach> <float power/tone, in dBm>`
set_attens_last_all: Sets the input and output attenuators to their last set values. These values are read from the file: `last_attens.dat`, located on FC1/FC2 in: `/roach_flight/roach<which roach>` The set values will appear in the Owl fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
Syntax: `set_attens_last_all <float output atten, dB> <float input atten, dB>`
set_attens_min_output: Sets the output attenuators of all ROACHs to 30 dB. This mode is used to measure the warm electronics noise contribution to the total noise budget. It's flagged because it would be bad to set this and then forget to revert back to normal settings. The set values will appear in the `Ow` fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
Syntax: `set_attens_min_output_all`
set_attens_conserved: Sets the output attenuator to specified number of dB, while choosing the input attenuator level to conserve a total of 23 dB between the pair. The set values will appear in the `Owl` fields: `SET_ATTEN_OUT`, `SET_ATTEN_IN`.
Syntax: `set_attens_conserved <which roach> <float output atten, dB>`

A.6.7 Pi Commands

***reboot_pi**: Reboots the specified Pi, by sending `sudo reboot` to the netcat server listening on ports 12345 or 12346.
Syntax: `reboot_pi <which roach> <PORT#>` Note: MCP will automatically reboot a Pi after 10 errors are registered by the Valon and/or attenuator clients. This command should only be used in the event that automatic rebooting appears to have failed. If the Pi is already powered down, the command will have no effect. If needed, verify that the Pi is powered up by pinging it at the IP address specified in the Hardware Addresses section. After rebooting, the Valon and attenuator servers start up automatically (they are in `/etc/rc.local` on the Pis). There is no need to manually configure anything after reboot.
****roach_set_extref**: Configure the Valons to use either the external (1, default for flight-mode) or internal (0) PLL reference.
Syntax: `roach_set_extref <which roach> < 0 | 1 >` Note: By default, the Valon uses the external 10 MHz reference signal provided by the OctoClock. If the external reference signal drops out or changes frequency, the LO and FPGA clock signals generated by the Valon will become unstable and/or drop out entirely. This

command is intended for use only in this scenario. Signs that the external reference has dropped out may include:

- Accompanied loss of PPS signal (check PPS counter fields in `Owl`).
- Nonsensical channel data lasting for more than a few seconds.
- No channel timestreams at all (flat line, possibly with a DC offset).

The clearest indication that the external reference has dropped out will be a QDR RAM calibration fail after resetting the ROACHs, with an out-of-range FPGA clock frequency reported by Owl (see `FPGA_CLOCK_FREQ` field).
If you have no other option but to use this command, follow this procedure:

1. Send the command to switch one or more (see *all* version) to internal reference:

   ```
   > roach_set_extref <which roach> 0
   ```

2. Power cycle the readout (see: Power commands)

read_pi_temps: Read the temperatures of all Pi CPUs. The temperatures will appear in the `PI_TEMP` field in `Owl` for each ROACH.

A.6.8 Power Status Commands

roach_cycle: Power cycles all five ROACH slices.
Syntax: `roach_cycle`
roaches_off: Powers down all five ROACH slices.
Syntax: `roaches_off`
roaches_on: Powers up all five ROACH slices.
Syntax: `roaches_on`
kill_roach: Halts a ROACH's PPC. The ROACH may be brought back online by power cycling the entire readout stack.
Syntax: `kill_roach <which roach>`

A.6.9 Commands That Set Global Parameters

roach_has_lamp_control: Gives exclusive control of the cal lamp to the ROACH. No other ROACHs will be able to trigger the pulse.
Syntax: `roach_has_lamp_control <which roach>`
set_find_kids_params: Set the KID-finding parameters. On system start-up, these parameters are automatically set to default values. See: KID-finding Arguments:

- Which ROACH (1–5)
- Smoothing scale (kHz, 1000.0–100000.0)

- Dip-depth threshold (dB, 0.1, 100)
- Spacing threshold (kHz, 80.0, 10000.0)

Syntax: `set_find_kids_params <which roach> <smoothing scale> <dip-depth threshold> <spacing threshold>`

*__set_max_nkids__: Sets the maximum number of channels that can be found from a VNA sweep without triggering an error. The maximum number of channels should be 500–800. See: KID-finding, `TONE_FINDING_ERROR`.
Syntax: `set_max_nkids <which roach> <max nkids>`

*__set_max_nkids_all__: As above, for all ROACHs.
Syntax: `set_max_nkids <max nkids>`

*__set_min_nkids__: Sets the minimum number of channels that can be found from a VNA sweep without triggering an error. The minimum number of channels should be 1–50. See: KID-finding, `TONE_FINDING_ERROR`.
Syntax: `set_min_nkids <which roach> <min nkids>`

*__set_min_nkids_all__: As above, for all ROACHs.
Syntax: `set_min_nkids_all <min nkids>`

__set_df_diff_retune_threshold__: See `DF_RETUNE_THRESH`. Sets the threshold for the cal lamp *df* response ($df_{on} - df_{off}$), in Hz.
Syntax: `set_df_diff_retune_threshold <which roach> <Hz>`

__set_df_diff_retune_threshold_all__: As above, for all ROACHs.
Syntax: `set_df_diff_retune_threshold_all <Hz>`

__set_n_outofrange_thresh__: See `N_OUTOFRANGE_THRESH`. Sets the threshold for number of channels which can be out of range (df lamp response).
Syntax: `set_n_outofrange_thresh <which roach> <number of chan>`

__set_n_outofrange_thresh_all__: As above, for all ROACHs.
Syntax: `set_n_outofrange_thresh <number of chan>`

__set_default_tone_power__: Sets the default power/tone (dBm) used by `full_loop_default`. Typical power/tone is -47 dBm.
Syntax: `set_default_tone_power <which roach> <power/tone, in dBm>`

__set_default_tone_power_all__: As above, for all ROACHs.
Syntax: `set_default_tone_power_all <power/tone, in dBm>`

A.6.10 Data Handling Commands

__compress_roach_data__: Compresses the most recent set of the specified data type for each ROACH slice into a single tar.gz. Each tarball is assigned a local environment variables based on its data type, which should be specified as the file path when downlinking data.
Syntax: `compress_roach_data <0--6>`
where:

0 = VNA sweeps
1 = Target sweeps
2 = IQ timestreams
3 = DF timestreams
4 = Lamp responses (*df*, *I*, *Q*)
5 = Noise comparison data
6 = Baseband target frequency lists

Environment variables:
0 = $ALL_VNA_SWEEPS
1 = $ALL_TARG_SWEEPS
2 = $ALL_IQ_DATA
3 = $ALL_DF_DATA
4 = $ALL_LAMP_DATA
5 = $ALL_NOISE_COMP
6 = $ALL_BB_FREQS

A.7 Data Downlinking

Several commands result in data being saved on FC1/FC2. Some of the data is overwritten each time it's created (e.g., the lamp responses in *df*, *I*, and *Q* (`lamp_response.dat`)). Other data (e.g., the VNA and target sweep data) is time-stamped and stored on the hard drives as a record of the flight. When data which would be useful to view during the flight is created, MCP automatically creates a tarball of the data for each ROACH, which is overwritten with each consecutive data creation. These tarballs are given unique environment variables corresponding to their data type. These environment variables can be input into `Cow` to downlink the data. The variables are:

```
$R<1--5>_LAST_VNA_SWEEP
$R<1--5>_LAST_TARG_SWEEP
$R<1--5>_LAST_DF_DATA
$R<1--5>_LAST_IQ_DATA
$R<1--5>_LAST_LAMP_DATA
$R<1--5>_LAST_BB_TARG_FREQS
$R<1--5>_LAST_TARG_FREQS_MAGS
$R<1--5>_LAST_NOISE_COMP
```

Example I: Getting data for all ROACHs To downlink a tarball containing all the most recent VNA sweeps for each ROACH, from `Cow`:

1. In ROACH commands: `compress_roach_data 0`
2. Go to Telemetry screen
3. Choose: `Request_stream_file` Choose options: Downlink = Fragment # = 0 File block number = 1 Absolute file path = `$ALL_VNA_SWEEPS`
4. On GROUNDSTATION-1, go to: `/data/etc/downloaded_files`
5. The downlinked file will appear as: `file_block_1`
6. A bash script (`untarloop.sh`) running in the background will untar the file when downlink is complete, and move the data to the appropriate directories.

Example II: Getting data for a single ROACH To downlink a VNA sweep tarball containing data for a single ROACH, `Cow`:

1. Go to Telemetry screen
2. Choose: `Request_stream_file` Choose options: Downlink = Fragment # = 0 File block number = 1 Absolute file path = `$R1_LAST_VNA_SWEEP`
3. On GROUNDSTATION-1, go to: `/data/etc/downloaded_files`
4. The downlinked file will appear as: `file_block_1`
5. A bash script (`untarloop.sh`) running in the background will untar the file when downlink is complete, and move the data to the appropriate directories.